The Architecture of Market Regimes:
Modeling DXY through Structural and Reflexive Layers
Expanded and Consolidated Edition

Vi Nhat Son

September 6, 2026

Important Legal Disclaimer

This document (the "Document") and all information, models, data, scenarios, and analyses contained within it, including but not limited to the "Deconstructing Dollar Dynamics" series, are provided under the following conditions and understandings:

1. For Informational, Academic, and Hypothetical Purposes Only

This Document is drafted and published solely for informational, academic exchange, and hypothetical illustration purposes. The content herein is presented as a theoretical exercise in politico-economic modeling. This Document DOES NOT constitute and should NOT be construed as:

- **Investment Advice:** Any form of recommendation to buy, sell, hold, or engage in any transaction with respect to any asset class, currency (including DXY), security, commodity, or financial instrument.
- **Financial, Legal, or Tax Advice:** Any form of professional consultation.
- **Market Forecast:** Any guarantee or prediction of future events or price movements.

2. No Warranty of Accuracy or Completeness

The author makes no representations or warranties, express or implied, regarding the accuracy, reliability, completeness, or timeliness of the information within this Document. The "State-Dependent, Non-Linear Integrated Equation" model is a conceptual construct, entirely theoretical, and has not been empirically validated or peer-reviewed. All data, variables, and correlations used are for illustrative purposes of the assumed model only. The author is not responsible for any errors, omissions, or outdated information.

3. Hypothetical Statements and Scenarios

All "Case Studies" presented—including analyses related to historical events (GFC 2008, COVID-19), specific entities (e.g., Vendor), or geopolitical events (e.g., Russian energy sanctions)—are entirely hypothetical and intended solely to illustrate how the theoretical model might interpret such events. They do not reflect any insider information, are not

affirmations of fact, and are not predictions of future outcomes. Actual events may differ materially from the scenarios described.

4. Maximum Disclaimer of Liability

In no event shall the author (including any individual or entity involved in the creation or distribution of this Document) be LIABLE to any person or entity for:

- Any loss, damage, or consequence (whether direct, indirect, consequential, special, or punitive);
- Any investment or business decisions made;
- Any loss of profits, revenue, or data;

arising from the access, use of, or reliance upon the information, models, or hypothetical scenarios contained in this Document, even if advised of the possibility of such damages.

5. Personal Views, No Affiliation

The views, opinions, and analyses expressed in this Document are those of the individual author (Vi Nhat Son) and do not represent, reflect, or constitute an endorsement by any organization, employer, client, or government with which the author may be affiliated. The mention of any company, nation, or organization's name (e.g., Vendor, Russia, the United States) is for academic illustration purposes only and does not constitute a critique, attack, defamation, or the disclosure of confidential information.

6. Reader's Responsibility

The reader acknowledges that any actions taken based on any information or analysis in this Document are strictly at the reader's own risk. The reader is responsible for conducting their own due diligence and must always consult with a qualified financial, legal, or other professional advisor before making any critical decisions. By reading, accessing, or using this Document, you unconditionally and irrevocably agree to all the terms of the above disclaimers.

Abstract

Traditional models attempting to explain the dynamics of the U.S. Dollar Index (DXY)—such as Purchasing Power Parity (PPP) or simple interest-rate differential frameworks—consistently fail to capture real-world fluctuations, particularly during episodes of financial stress. These models are largely linear, neglecting the financial microstructure ("plumbing") and assuming static correlations. This work constructs a state-dependent, non-linear integrated equation to describe the DXY's value. We hypothesize that DXY is a complex function of the ratio between Global Effective Demand (D^*) and Global Effective Supply (S^*) for the USD. The core innovation lies in introducing a systemic "stress" variable (σ_{sys}), which renders the model's weights (α, β) and convexity (κ_{sys}) as dynamic functions rather than constants. Furthermore, we introduce foundational meta-variables—Geopolitical Power (G_{pol}), Internal U.S. Risk (R_{US}), and Human Capital (H_{cap})—not as direct inputs but as modulators altering the system's internal behavior. This model provides a stronger framework for capturing non-linear properties, especially the convex "flight-to-quality" response of the USD.

This edition expands that argument into a full specification, tests it, and applies it. **Part I** builds the model in six strata. Layer 0 is the mechanical ratio of ten demand channels to seven supply channels. Layer 1 replaces the original's raw stress thermometer with a bounded systemic instability index $\sigma_{\text{sys}} \in [0, 1]$, so that the safe-haven term is convex without being explosive. Layer 2 supplies the structural and trust variables—shelter credibility Φ_{sys} and the share $s_{\text{US,sys}}$ of global flight-to-quality flow that the dollar actually captures—as functions of geopolitical power, internal U.S. risk, institutional trust M_{trust} and geopolitical fragmentation T_{frag}. Layer 3 makes the convexity exponent κ_{sys} an explicit function of leverage, crisis memory M_{shock} and reflexivity. Layer 4 adds two psycho-technical states, the homogeneity of collective belief H_{soc} and impulse efficiency I_{imp}, which amplify Layer 3 and are shown to form a simultaneous system with a unique fixed point under a stated and checkable condition. Layer 5 introduces a secular structural-integrity coefficient $\Psi_{\text{str}}(t)$, equal to one when the dollar's position is intact and falling along a cliff rather than a slope as alternative infrastructure and fragmentation accumulate. Every elasticity in the model interpolates between a calm value and a crisis value through a shared logistic transition in σ_{sys}; the specification is assembled, its comparative statics derived, and its identification stated.

Part II confronts the model with eleven episodes—the 2008 financial crisis, COVID-19, the taper episodes, U.S. debt-ceiling standoffs, the trade war, the 2022 sanctions on Russia, stablecoin growth, carrier-power and digital-substrate scenarios, and the 2025 tariff episode—reading each one term by term through the log-linear form and scoring predictions against

observation in a single cross-case table. Hypothetical cases are kept as hypotheticals and paired with real analogues where any exist.

Part III turns the model on policy: the anatomy of an inflation crisis in a psycho-reflexive era, a lever-by-lever playbook with legal basis and dated historical activations, the execution rhythm and the structural traps it must avoid, the coordination mechanism among the major central banks, and the 1971 inflection that created the system being modelled.

Throughout, three commitments hold. No number enters the text without a source. No reference enters the bibliography without verification. Where the original specification was ambiguous, incomplete or internally inconsistent, the defect is named, the resolution is recorded, and both are visible to the reader rather than smoothed away.

How to Read This Book

This is a long book with a short spine. The spine is one equation, assembled in Chapter 2, completed in Chapter 10, and used everywhere after that. Everything else is either a component of it, a test of it, or an application of it.

Three parts, three kinds of reading

Part I (Theory, Chapter 1–Chapter 11) is cumulative and should be read in order the first time. Each chapter after Chapter 2 takes one layer of the model and does four things: states the variables formally, gives the mechanism in words, specifies the functional form, and says how each variable would be measured and where the measurement fails.

Part II (Evidence, Chapter 12–Chapter 17) is modular. Each case study is self-contained and follows the same seven-step template: event window, faithful restatement of the original claim, data, a term-by-term reading of the model, prediction versus observation, a scorecard row, and what the episode implies for the parameters. Read any one of them alone; read Chapter 17 to see them together.

Part III (Policy and system, Chapter 18–Chapter 22) can be read first by anyone who came for the argument rather than the machinery. It cross-references back into Part I at each point where a claim depends on a layer.

Notation

Six conventions carry the whole book.

1. **Layers are numbered 0 to 5**: 0 mechanical, 1 state, 2 structural, 3 reflexive, 4 psycho-technical, 5 secular. Policy reaction is a feedback, not a layer.
2. **Demand channels are D with a named subscript** (D_{trans}, D_{debt}, D_{prec}, $\dots$) and **supply channels are** S (S_{Fed}, S_{coll}, S_{swap}, $\dots$). Demand up means DXY up; supply up means DXY down. There are no exceptions to that convention anywhere in the book.
3. **Every latent modulator lives in** $[0, 1]$**.** σ_{sys}, G_{pol}, R_{US}, M_{trust}, T_{frag}, H_{soc}, I_{imp}, C_{alt} and the rest are bounded indices. A coefficient on any of them therefore reads as the effect of a move across the entire range, which makes coefficients comparable across layers.
4. **Quantities are index numbers, not levels**, each equal to one at a stated base period.

5. $\Psi_{\text{str}}(t)$ **is an integrity coefficient**: one when intact, falling as the dollar's structural position erodes. Erosion is always $\Psi_{\text{str}}(t) \downarrow$, never $\Psi_{\text{str}}(t) \uparrow$.
6. **Greek letters are parameters, Latin letters are variables**, with the single exception of σ_{sys} and κ_{sys}, which are states inherited from the original text.

The complete symbol table, generated from the notation registry, is Chapter A. Every symbol used anywhere in the book appears there exactly once, with its layer, its role, and the page of the original paper where it was introduced.

The arrow scale

Tables use arrows for direction and magnitude. The scale is:

Symbol	Meaning
$\uparrow$ / $\downarrow$	small move; within the normal range of the variable
$\uparrow\uparrow$ / $\downarrow\downarrow$	large move; a move of the size seen in a stress episode
$\uparrow\uparrow\uparrow$ / $\downarrow\downarrow\downarrow$	extreme move; the largest observed, or a regime change
$\leftrightarrow$	no material change

Where data exist, tables show both the arrow and the number; the arrow is a summary of the number, never a substitute for it.

Markers in the text

Four markers appear in draft chapters and are all cleared before the final build. • An *author-decision* marker records a modelling choice that belongs to the author and has not yet been made. • A *data* marker records a number that cannot be stated because the underlying series is not in hand. • A *verify* marker records a reference that has not been checked against a publisher record, and no such reference is ever cited. • A *to-do* marker records an unwritten section. If any of these survive into the copy you are reading, they are deliberate and they are the honest state of that passage.

Getting from a case study back to the model

Each case study names the variables that carried the episode and cross-references the chapter where each is defined. The fastest route in the other direction is the scorecard in Chapter 17: it lists, for every episode, which layer did the work. If you want to know whether the reflexive layer earns its place, read the κ_{sys} column of that table and then Chapter 7. If you want to know whether the psycho-technical layer earns its place, read the corresponding column and then Chapter 8 — and note in advance that this is the layer where the evidence is thinnest, a fact the book states rather than disguises.

What the appendices are for

Chapter A is the symbol table. Chapter B reproduces every displayed equation of the original paper, transcribed faithfully and left uncorrected, so that a reader can always see what was originally written before this edition restructured it. Chapter C documents every series used: identifier, source, retrieval date, and transformation. Chapter D collects the proofs. Chapter E maps every section of the original paper to its location here; nothing in the original was dropped, and that table is how the claim can be checked.

Contents

Important Legal Disclaimer ii

Abstract iv

How to Read This Book vi

I Theory: A State-Dependent, Reflexive Model of the Dollar 1

1 Introduction: The Failure of Linear Models 2
- 1.1 The Dollar Index as a Measure of Global Liquidity . . . 2
- 1.2 Two Incomplete Traditions . . . 3
- 1.3 The Linearity Assumption and the Non-Linear Response . . . 4
- 1.4 Related Literature . . . 5
  - 1.4.1 Safe-asset demand and the dollar . . . 5
  - 1.4.2 The global financial cycle and the dollar smile . . . 6
  - 1.4.3 Dollar funding plumbing . . . 7
  - 1.4.4 Reflexivity, memory, and narrative . . . 7
  - 1.4.5 Coercion, fragmentation, and erosion . . . 8
  - 1.4.6 Stress measurement and estimation . . . 8
  - 1.4.7 The gap . . . 9
- 1.5 Contributions . . . 9
  - 1.5.1 What survives testing . . . 11
- 1.6 Roadmap . . . 12

2 General Model Architecture 14
- 2.1 The Foundational Equation . . . 14
  - 2.1.1 Effective external supply . . . 15
  - 2.1.2 Emergent channels . . . 15
- 2.2 Architectural Philosophy: Four Pillars . . . 15
  - 2.2.1 Pillar 1 — the multiplicative ratio . . . 16
  - 2.2.2 Pillar 2 — the state variable . . . 16
  - 2.2.3 Pillar 3 — the convexity exponent . . . 16
  - 2.2.4 Pillar 4 — the pure safe-haven term . . . 16
- 2.3 Foundational Meta-Variables . . . 17

2.4 The Layer Map . . . 18
2.5 Derivation Block: The Foundational Equation . . . 18
2.5.1 Assumptions . . . 18
2.5.2 Functional form . . . 20
2.5.3 Parameters . . . 21
2.5.4 Comparative statics . . . 21
2.5.5 Failure modes . . . 22
2.6 Formal Properties . . . 22
2.6.1 Homogeneity and why the ratio is scale-free . . . 22
2.6.2 Identification of the baseline . . . 22
2.6.3 Why the exponent and the coefficient are different objects . . . 23
2.6.4 Bounds on the safe-haven term . . . 23
2.6.5 The estimable form . . . 23
2.7 Notation Conventions . . . 23
2.7.1 The arrow scale . . . 24
2.8 The Dependency Graph . . . 24

3 Decomposing the Numerator (D^*) 26
A note on the dependent variable . . . 26
3.1 Transactional Demand (D_{trans}) . . . 27
3.2 Debt-Servicing Demand (D_{debt}) . . . 31
3.3 Capital-Asset Demand (D_{cap}) . . . 32
3.4 Safe-Haven Demand (D_{prec}) . . . 36
3.5 Reserve-Management Demand (D_{res}) . . . 39
3.6 Hedging Demand (D_{hedge}) . . . 41
3.7 Derivative Demand (D_{gamma}) . . . 44
3.8 Network Demand (D_{net}) . . . 45
3.9 Microstructural Demand (D_{micro}) . . . 48
3.10 Human-Capital and IP Demand (D_{hcap}) . . . 49
3.11 Interactions Among Demand Drivers . . . 50
3.11.1 Amplifications . . . 51
3.11.2 Boundaries: what must not be double-counted . . . 51
3.11.3 The policy-divergence question . . . 51
3.11.4 Collinearity that the data cannot resolve . . . 52

4 Decomposing the Denominator (S^*) 53
4.1 Federal Reserve Actions (S_{Fed}) . . . 53
4.2 U.S. Treasury Actions (S_{Tsy}) . . . 55
4.3 Collateral Availability (S_{coll}) . . . 58
4.4 Bank Balance-Sheet Capacity (S_{bank}) . . . 60
4.5 Federal Reserve Swap Lines (S_{swap}) . . . 62
4.6 External Balance-of-Payments Flows (S_{ext}) . . . 63
4.7 Synthetic USD Supply (S_{crypto}) . . . 65
4.8 Sanctions and Frozen Assets (S_{off}) . . . 67
4.9 Shadow-Bank USD Creation (S_{shadow}) . . . 68

4.10 The Effective-Supply Identity and Double-Counting Checks 69
4.10.1 The identity 69
4.10.2 Four double-counting checks 69
4.10.3 What the denominator cannot currently do 71

5 Layer 1 — State and Friction (σ_{sys}) 72
5.1 From Thermometer to Systemic Instability Index 72
5.1.1 Two printed forms, one canonical 72
5.2 Endogenous Financial Stress ($\sigma_{\text{VIX/MOVE}}$) 73
5.3 Exogenous Shocks (X_{pand}) 73
5.4 Inter-Market Liquidity Friction (F_{liq}) 75
5.5 Construction: From Composite to Bounded Index 76
5.5.1 Why the raw composite is not the state variable 76
5.5.2 The construction 76
5.5.3 What the third component changes 77
5.5.4 What the rank transform does to an unestimated weight 78
5.6 Regime Thresholds and the Empirical Test 80
5.6.1 What the model actually claims 80
5.6.2 The result 80
5.6.3 Robustness 82
5.6.4 What survives, and what does not 83
5.6.5 The amplification result as a function of w_X 83

6 Layer 2 — Structural and Trust (Φ_{sys}, $s_{\text{US,sys}}$) 87
6.1 The Anchor Question 87
6.1.1 Why two functions and not one 87
6.1.2 The four structural variables 88
6.2 Geopolitical Power (G_{pol}) 88
6.3 Internal U.S. Risk (R_{US}) 91
6.4 Institutional Trust (M_{trust}) 92
6.5 Geopolitical Fragmentation (T_{frag}) 94
6.6 The Credibility Function Φ_{sys} 97
6.6.1 Statement 97
6.6.2 Assumptions 97
6.6.3 Functional form 98
6.6.4 Parameters and comparative statics 98
6.6.5 Failure modes 98
6.7 The Safe-Flow Share $s_{\text{US,sys}}$ 99
6.7.1 Statement 99
6.7.2 Where it enters 99
6.7.3 Marginal effects, and the derivative CS7 uses 99
6.7.4 The competing havens 100
6.8 Amplification of R_{US} by the Psycho-Technical Layer 100
6.9 Synthesis and Tail Risk 101

7 Layer 3 — Reflexive and Acceleration (κ_{sys}) 103
7.1 Convexity as a Sensitivity Knob . . . 103
7.1.1 What this chapter has to repair . . . 103
7.2 Baseline Convexity (κ_{base}) . . . 104
7.3 Velocity of Shock Transmission (V_{USD}) . . . 104
7.4 Crisis Memory (M_{shock}) . . . 105
7.4.1 The memory function . . . 105
7.4.2 The worked example, and what it implies . . . 106
7.4.3 Synthetic re-activation . . . 106
7.5 Meta-Reflexivity (A_{meta}) . . . 108
7.6 Dual Amplification and Behavioural Phase Lock . . . 109
7.6.1 The mechanism . . . 109
7.6.2 A countervailing effect the original does not note . . . 109
7.7 The Updated κ_{sys} Formula . . . 110
7.7.1 Statement . . . 110
7.7.2 Assumptions . . . 110
7.7.3 Why this form . . . 111
7.7.4 The restored channels . . . 111
7.7.5 Parameters and comparative statics . . . 111
7.7.6 Failure modes . . . 112
7.8 Hyper-Brittleness . . . 113
7.8.1 What is claimed, and what is shown . . . 113
7.9 Empirical Convexity . . . 113
7.9.1 Reading the profile . . . 114
7.10 Conclusion . . . 115

8 Layer 4 — Psycho-Technical (H_{soc}, I_{imp}) 116
8.1 The Meta-Platform Argument . . . 116
8.2 Variable Architecture . . . 117
8.2.1 The four drivers . . . 117
8.2.2 The defensive variable . . . 118
8.2.3 The two state variables . . . 119
8.2.4 The two consequence variables . . . 119
8.3 E7: The Homogeneity Convergence Function . . . 120
8.3.1 Statement . . . 120
8.3.2 Assumptions . . . 120
8.3.3 Functional-form justification . . . 121
8.3.4 Parameters . . . 122
8.3.5 Properties . . . 122
8.3.6 Comparative statics . . . 122
8.3.7 Failure modes . . . 122
8.4 E7b: The Impulse Function . . . 123
8.4.1 Statement . . . 123
8.4.2 The five drivers . . . 123
8.4.3 Functional-form justification . . . 123

8.4.4 ISSUES #13: closing the loop . . . 123
8.5 The Closed Loop: Existence, Uniqueness, Stability . . . 125
8.5.1 One map instead of two equations . . . 125
8.5.2 Existence . . . 125
8.5.3 Uniqueness . . . 125
8.5.4 The social amplification factor . . . 127
8.5.5 Comparative statics at the equilibrium . . . 127
8.6 Impact on the Rest of the Model . . . 127
8.7 When the Loop Admits Three Equilibria . . . 129
8.7.1 The result . . . 129
8.7.2 What the three equilibria are . . . 129
8.7.3 What this costs the rest of the book . . . 130
8.7.4 What is at stake if $\eta \geq 4$ . . . 130
8.8 Measurement . . . 131
8.8.1 What would have to be measured . . . 131
8.8.2 What is actually measured: nothing . . . 132
8.9 Conclusion . . . 133

9 Layer 5 — Structural Break and Narrative ($\Psi_{\text{str}}(t)$) 134
9.1 Erosion Versus Break . . . 134
9.2 E8: The Cliff Function . . . 135
9.2.1 Statement . . . 135
9.2.2 Assumptions . . . 135
9.2.3 Functional-form justification . . . 136
9.2.4 Parameters . . . 136
9.2.5 Properties . . . 137
9.2.6 Comparative statics . . . 138
9.2.7 Failure modes . . . 138
9.3 The Drivers of Erosion . . . 138
9.3.1 Hard drivers: T_{frag} and C_{alt} . . . 138
9.3.2 Soft drivers: the psycho-technical layer . . . 139
9.4 The Tipping-Point Mechanism . . . 139
9.5 Calibration I: The One Completed Transition . . . 141
9.6 Calibration II: What the Measured Series Do . . . 142
9.6.1 The reserve share: erosion is real and decelerating . . . 142
9.6.2 The turnover share: the incumbent gained . . . 144
9.6.3 Reading the two together . . . 144
9.7 What Survives . . . 144
9.7.1 The same test on a reserve aggregate that includes gold . . . 145
9.8 Conclusion . . . 148

10 System Synthesis 150
10.1 Assembling the Model . . . 150
10.1.1 What the previous eight chapters produced . . . 150
10.1.2 Two thresholds, and the question they raise . . . 150

10.1.3 The answer: distinct objects, one-way coupled . . . 151
10.1.4 The coupling, and what it does . . . 153
10.1.5 What is missing: the return path . . . 154
10.2 E9: The Integrated Equation . . . 155
10.2.1 Statement . . . 155
10.2.2 Reading the four blocks . . . 156
10.2.3 Bounds . . . 156
10.3 The Log-Linear Form . . . 156
10.4 The Modulator Functions, Consolidated . . . 157
10.5 Comparative Statics . . . 159
10.6 The Formal Dollar Smile . . . 159
10.7 Stylised Shock Simulations . . . 160
10.7.1 (a) Liquidity shock . . . 162
10.7.2 (b) Sanctions shock . . . 162
10.7.3 (c) Trust shock . . . 162
10.7.4 (d) Narrative shock . . . 162
10.7.5 A sanity check the simulation fails . . . 163
10.8 The Six Strata . . . 163
10.9 Conclusion . . . 164

11 Measurement and Estimation 165
11.1 The Proxy Map . . . 165
11.1.1 The rule this book has followed . . . 165
11.1.2 The map itself . . . 166
11.2 Latent Variables and Their Signals . . . 166
11.3 Econometric Challenges . . . 168
11.3.1 Simultaneity everywhere . . . 168
11.3.2 Structural non-identification of κ_{sys} . . . 168
11.3.3 The crisis regime is thin, and clustered . . . 168
11.3.4 Frequency mismatch on $\alpha_{\text{prec}}(\sigma)$ . . . 170
11.3.5 Valuation effects in the one long series . . . 170
11.4 Estimation Strategies . . . 171
11.5 A Minimal Calibration of the Mechanical Core . . . 172
11.5.1 The specification . . . 172
11.5.2 A data defect found in this session, and what it changed . . . 172
11.5.3 The results, and three of six signs are wrong . . . 173
11.5.4 The λ_1 result . . . 174
11.5.5 Which component of σ_{sys} the dollar actually responds to . . . 174
11.5.6 Decomposing the two crises, and what carries the residual . . . 175
11.5.7 What the exercise establishes . . . 178
11.6 What Can and Cannot Be Estimated Today . . . 178
11.6.1 What a sceptic should take from this . . . 178
11.6.2 What this chapter does not concede . . . 178
11.7 What Chapters 12 to 17 Inherit . . . 180

II Evidence: Case Studies and Cross-Case Synthesis 181

12 Case-Study Method 182
12.1 What a Case Study in This Book Is For . . . 182
12.2 The Template . . . 182
12.3 Five Rules That Bind Every Study . . . 183
12.4 The 2026 Confounding . . . 183
12.5 Scoring . . . 184
12.6 What This Chapter Forbids . . . 185

13 Case Studies: Financial Crises 186
13.1 Case Study 1: The Global Financial Crisis, 2008–09 . . . 186
13.1.1 Event and window . . . 186
13.1.2 What the original claimed . . . 186
13.1.3 What the model predicts . . . 187
13.1.4 The data . . . 187
13.1.5 D_{prec} under Rule 1 . . . 187
13.1.6 Model reading . . . 189
13.1.7 E10b term by term . . . 189
13.1.8 Scorecard . . . 189
13.1.9 Parameter implications . . . 189
13.2 Case Study 2: The Pandemic, 2020 . . . 190
13.2.1 Event and window . . . 190
13.2.2 What the original claimed . . . 190
13.2.3 The data . . . 190
13.2.4 The distinctive feature: policy reversed it . . . 191
13.2.5 E10b term by term . . . 191
13.2.6 Scorecard . . . 193
13.2.7 Parameter implications . . . 193
13.3 The Funding-Haven Hypothesis, Scored . . . 194
13.3.1 Result: the two episodes disagree . . . 194
13.3.2 What survives: a frequency split, offered as conjecture . . . 194
13.3.3 And one confound that neither episode escapes . . . 195
13.3.4 Verdict . . . 195
13.4 Case Study 3: Taper 2013 and the Hiking Cycle, 2014–18 . . . 196
13.4.1 Event and window . . . 196
13.4.2 What the original claimed . . . 196
13.4.3 What happened, and what the model says . . . 196
13.4.4 The finding: the dollar responded to the level, not the announcement 196
13.4.5 Scorecard . . . 198
13.4.6 E10b term by term, and parameter implications . . . 198
13.5 Case Study 4: Debt Ceilings, 2011 and 2023 . . . 198
13.5.1 Why this episode matters . . . 198
13.5.2 The two episodes are not the same kind of event . . . 199
13.5.3 κ_R: what it would take, and why two episodes cannot do it . . . 199

13.5.4 Scorecard . . . 201
13.5.5 E10b term by term, and parameter implications . . . 202
13.6 Remark 13.1 Tested Where It Was Not Built . . . 202

14 Geo-Economic Shocks 204
14.1 Case Study 5: Trade War (2018) and Ukraine War (2022) . . . 204
14.1.1 What the original claims . . . 204
14.1.2 Event and window . . . 205
14.1.3 The data . . . 205
14.1.4 The measurement that decides the study . . . 205
14.1.5 The reserve freeze, at the resolution of its own event . . . 208
14.1.6 The rival, and it is not exotic . . . 208
14.1.7 The E10b walk . . . 209
14.1.8 Scorecard . . . 210
14.1.9 Parameter implications . . . 210
14.2 Case Study 9: Energy Embargo . . . 212
14.2.1 What the original claims . . . 212
14.2.2 The analogue, and why this study can be scored at all . . . 213
14.2.3 The data . . . 213
14.2.4 Scenario 2 scored, prediction by prediction . . . 213
14.2.5 The supply legs, where they can be measured . . . 216
14.2.6 Whether a four-scenario design can be wrong . . . 216
14.2.7 What the episode may not be used to confirm . . . 217
14.2.8 Scorecard . . . 217
14.2.9 Parameter implications, and the overlap with CS5 . . . 219
14.3 Case Study 12: The Iran War and the Strait of Hormuz, 2026 . . . 219
14.3.1 Why this is a new case study and not an extension of CS9 . . . 219
14.3.2 Event and window . . . 220
14.3.3 What the price data show . . . 220
14.3.4 The model reading, channel by channel . . . 221
14.3.5 Scorecard . . . 222
14.3.6 The confounding this creates . . . 222
14.3.7 One observed ordering, no attribution . . . 222

15 Case Studies: Structural Shifts 223
15.1 The Seven-Step Audit of This Chapter . . . 223
15.2 Case Study 6: Stablecoins, 2020–2026 . . . 224
15.2.1 Event and window . . . 224
15.2.2 What the original claimed . . . 224
15.2.3 The data . . . 224
15.2.4 E10b term by term . . . 224
15.2.5 Prediction against observation . . . 226
15.2.6 Scorecard . . . 226
15.2.7 Parameter implications . . . 226
15.3 Case Study 7: A New Aircraft Carrier as a ΔG_{pol} Event . . . 227

15.3.1 Why this case study is load-bearing . . . 227
15.3.2 What the original claims . . . 227
15.3.3 The quantitative claim . . . 227
15.3.4 The check: does CS7 meet the standard attributed to it? . . . 228
15.3.5 The real analogue, and what it shows . . . 229
15.3.6 Scorecard . . . 230
15.3.7 Parameter implications . . . 231
15.4 Case Study 8: A Digital-Substrate Shock . . . 231
15.4.1 What the original claims . . . 231
15.4.2 The real analogues, and what they show . . . 232
15.4.3 What survives, and what does not . . . 233
15.4.4 E10b term by term . . . 233
15.4.5 Scorecard . . . 234
15.4.6 Parameter implications . . . 234

16 Case Studies: Strategic 235
16.1 Case Study 10: Venezuela, January 2026 — An Ex-Ante Test . . . 235
16.1.1 Why this case study is different . . . 235
16.1.2 What the original claimed . . . 236
16.1.3 What happened . . . 236
16.1.4 The scorecard . . . 236
16.1.5 An observed ordering, recorded without attribution . . . 236
16.1.6 Why claim 1 failed: G_{pol} was measured by the wrong thing . . . 238
16.1.7 The channel the model has no category for . . . 239
16.1.8 What CS10 got right . . . 240
16.1.9 What Chapter 17 inherits . . . 240
16.2 Case Study 11: The 2025 Tariff Episode . . . 240
16.2.1 Event and window . . . 240
16.2.2 What the original claimed . . . 240
16.2.3 The data . . . 241
16.2.4 The 2025 evidence, and the verdict on Phase 1 . . . 241
16.2.5 E10b term by term . . . 243
16.2.6 Scorecard . . . 243
16.2.7 Parameter implications . . . 243

17 Cross-Case Synthesis and Scorecard 245
17.1 What This Chapter Is, and the Two Ways It Could Fail . . . 245
17.2 The Tally . . . 246
17.2.1 The same predictions cut by variable block . . . 247
17.3 What Does Not Survive . . . 249
17.4 What Does Survive . . . 251
17.5 The Defects, Sorted . . . 252
17.6 The Rescue Hypotheses, Pre-Registered and Tested . . . 252
17.6.1 H1: channel frequency mismatch — **refuted** . . . 254
17.6.2 H2: attenuation in σ_{sys} — **refuted**, and the confound fired . . . 255

17.6.3 H3: C_{alt} defined too narrowly — **supported**, with the mechanism qualified . . . 256
17.6.4 H4: policy volatility rather than the tariff level — **refuted** . . . 258
17.6.5 Reclassification . . . 259
17.6.6 The residual after all four rescues . . . 259
17.6.7 How much of the record is the model being wrong . . . 259
17.6.8 Declined, and why . . . 261
17.7 The Case Against the Model . . . 261
17.8 The Research Agenda . . . 262
17.9 The Verdict . . . 263

III Policy and System: Crisis Playbook, Coordination, and History 264

18 Inflation Crisis and the Optimal Path 265
18.1 Redefining Crisis in a Brittle System . . . 265
18.1.1 How Part III must be read . . . 265
18.1.2 What the original claims . . . 265
18.2 Inflation as Structural Failure: $R_{\text{US}} \uparrow$, $M_{\text{trust}} \downarrow$ within H_{soc} . . . 267
18.2.1 The trust channel . . . 267
18.2.2 The risk channel . . . 268
18.2.3 The dual-erosion effect . . . 268
18.3 The Dilemma: Trust versus Stress, Both Amplified . . . 268
18.3.1 The one number in the dilemma, and what it would require . . . 270
18.4 The Reflexivity Trap . . . 271
18.5 The Four-Spear Strategy . . . 271
18.5.1 What each spear would need to become checkable . . . 272
18.6 The 2021–2023 Episode: The Chapter's Own Argument, Tested . . . 272
18.6.1 What the Right Hand did . . . 272
18.6.2 Did trust collapse? . . . 274
18.6.3 And the Left Hand was barely used . . . 276

19 Strategic Decomposition and Operational Playbook 277
19.1 The Right Hand: Tightening to Rescue M_{trust} . . . 277
19.1.1 Lever 1: the policy rate (S_{Fed}) . . . 277
19.1.2 Lever 2: communication (M_{trust}, Spear 4) . . . 279
19.1.3 Lever 3: fiscal coordination (R_{US}) . . . 280
19.2 The Left Hand: Opening the Pipes to Quench σ_{sys} . . . 280
19.2.1 Lever 4: central bank liquidity swaps (S_{swap}) . . . 281
19.2.2 Lever 5: the FIMA repo facility . . . 282
19.2.3 Lever 6: bill issuance (S_{coll}, Treasury) . . . 283
19.2.4 Lever 7: leverage-ratio relief (S_{bank}) . . . 283
19.2.5 Lever 8: cash management (S_{Tsy}: TGA and ON RRP) . . . 284
19.3 The Root Layer: Reducing Pressure on the Right Hand . . . 285

19.3.1 Lever 9: energy and commodity supply (D_{trans}) . . . 285
19.3.2 Lever 10: restraint in coercion (T_{frag}, S_{off}, $\Psi_{\text{str}}(t)$) . . . 286
19.3.3 Lever 11: talent and intellectual property (D_{hcap}) . . . 286
19.4 Levers That Exist and the Playbook Does Not Name . . . 287
19.4.1 The Standing Repo Facility . . . 287
19.4.2 The Bank Term Funding Program . . . 287
19.4.3 The discount window, and why it is not a lever . . . 288
19.5 Who Actually Holds Each Lever . . . 288
19.6 Expected Impacts as the Original States Them . . . 289

20 Execution, Traps, and Trajectory 291
20.1 Implementation Cadence: Hour, Week, Quarter, Year . . . 291
20.1.1 Hour–day: extinguish σ_{sys} and hold the narrative . . . 291
20.1.2 Weekly: stabilise domestic plumbing . . . 292
20.1.3 Quarterly: restore M_{trust} and ease D_{trans} . . . 292
20.1.4 Annual: consolidate Φ_{sys} and $\Psi_{\text{str}}(t)$ . . . 294
20.2 Amplified Structural Traps . . . 294
20.2.1 Trap 1: mixed messaging . . . 295
20.2.2 Trap 2: political brinkmanship . . . 295
20.2.3 Trap 3: over-weaponisation . . . 295
20.2.4 Trap 4: plumbing neglect . . . 296
20.3 The Optimal DXY Trajectory and the Updated Equation . . . 297
20.3.1 What "optimal" means here . . . 297
20.3.2 The three phases, and the episode . . . 297
20.3.3 The updated equation, E10c . . . 298
20.3.4 The strategic battlegrounds . . . 298

21 The Decentralized System Coordination Mechanism 300
21.1 The Hypothesis, and Why It Is the Least Testable Claim in the Book . . . 300
21.1.1 What the original claims . . . 300
21.1.2 Evidential standing, stated before the argument rather than after it . . . 300
21.1.3 The claim drawn, and the axes labelled . . . 301
21.2 The Other Two Pieces of Part 0's Evidence . . . 301
21.2.1 Contradictory actions as evidence of imperfection . . . 301
21.2.2 Post-1971 persistence as evidence . . . 303
21.3 Historical Evidence I: The GFC as Systemic Self-Reinforcement . . . 303
21.3.1 The documented record . . . 303
21.3.2 The model reading, and its standing . . . 303
21.4 Historical Evidence II: The Plaza Accord . . . 304
21.4.1 The documented record . . . 304
21.4.2 The model reading . . . 305
21.4.3 The ordering problem . . . 305
21.5 Information Asymmetry and Personal Calculations . . . 306
21.6 The Self-Contradictory System . . . 306
21.7 The One Observable Shadow of the Attractor Claim . . . 307

21.8 Verdict on This Chapter 308

22 The Inflection Point: Vietnam and the Nixon Shock 309
22.1 What Is Observable and What Is Not 309
22.2 Breakdown of the G_{pol}–R_{US} Equilibrium 309
22.2.1 What the original claims 309
22.2.2 What can be measured about the strain 310
22.3 Collapse of M_{trust} and the 1971 Structural Break 311
22.3.1 The documented timeline 311
22.3.2 What the dollar did 312
22.3.3 The external balance, which is measurable 312
22.3.4 The model reading 312
22.4 The Modern System and Its Inherent Instability 314
22.5 Epilogue: Trust as the Only Anchor 315

A Notation and Symbol Table 316

B Equation Compendium of the Original Paper 326

C Data Appendix 330
C.1 How to Reproduce This Book 330
C.2 The Series This Book Uses 330
C.3 The Absences, and Whose They Are 330
C.4 One Defect Found, and What It Changed 332
C.5 Tiering 332

D Derivations and Proofs 334
D.1 The Psycho-Technical Fixed Point 334
D.1.1 Existence 334
D.1.2 Uniqueness 335
D.1.3 Multiplicity 335
D.1.4 The amplification factor 336
D.2 From E9 to E10: the Log-Linearisation 336
D.3 Comparative Statics of the Assembled Equation 337
D.3.1 The sign-change condition 337
D.3.2 The discontinuity Layer 4 induces in Layer 5 337
D.3.3 The two thresholds are distinct objects 338
D.4 The Logistic Derivative Behind CS7 338
D.5 Properties of the Cliff Function 339
D.6 Stability of the Safe-Haven Exponential 339

E Concordance: Original Sections to This Edition 341
E.1 What This Appendix Is For 341
E.2 Method 341
E.3 Result 342
E.4 The Concordance 342

E.5 Where Content Moved, and Why 342
E.6 Two Equations Without a Standalone Home 342
E.7 What This Appendix Does Not Certify 342

List of Figures

2.1 The six-layer architecture, with the factor of the integrated equation that each stratum supplies. Layer index is carried by vertical position and by the label; colour distinguishes only the demand and supply halves of Layer 0. Source: author's schematic; script `f01_architecture.py`. 19

2.2 Variable dependency graph, generated from `registry/variables.csv`. Columns are dependency depth: sources on the left, DXY on the right. Colour carries role — demand, supply, modulator — and the digit on each node is its layer, since six ordered levels cannot be separated by hue in print. Coefficients, conventions and the shorthands introduced for the fixed-point proof of Chapter 8 are omitted. Source: author's calculation; script `f02_dependency.py`. 25

3.1 D_{trans}, price leg only: Brent crude, daily, 2006-01-03 to 2026-09-01. The volume leg of the composite is not yet in hand, so this is an input to the proxy rather than the proxy itself. Source: FRED `DCOILBRENTEU`; author's calculations; script `f03_dtrans.py`. 30

3.2 D_{debt}: dollar credit to non-bank borrowers outside the United States, level and year-on-year change, 2000-Q1 to 2026-Q1. The level is the size of the offshore dollar short; the change is when it was forced to shrink. Source: BIS Global Liquidity Indicators, series `Q.USD.3P.N.A.I.B.USD` and `...B.771`; author's calculations; script `f04_ddebt.py`. 33

3.3 D_{cap}: 10-year U.S. minus German government bond yield, monthly, 2006-01 to 2026-06. Positive is a U.S. yield advantage. Source: FRED `DGS10` at month end and FRED `IRLTLT01DEM156N`; author's calculations; script `f05_dcap.py`. 35

3.4 The *trigger* for D_{prec}, not the channel itself: the VIX, daily, 2006-01-03 to 2026-09-03, with the thresholds quoted in the text. This series is the base of σ_{sys} in Chapter 5 and is therefore disqualified as a proxy for D_{prec} (ISSUES #31); it is shown to date the activations, not to measure the flow. Source: FRED `VIXCLS`; author's calculations; script `f06_dprec.py`. 40

3.5 D_{res}: the U.S. dollar's share of allocated official foreign exchange reserves, 1999-Q1 to 2026-Q1, with the model's base period marked. Source: IMF COFER, allocated reserves, claims in US dollar, shares; author's calculations; script `f07_dres.py`. 42

3.6 D_{net} and its supply-side twin. Circulating supply of USD-pegged stablecoins (the D_{net} proxy) and total value locked in DeFi (the S_{crypto} proxy of Section 4.7), daily, log scale. One figure serves both cards because it is one object read in two directions. Source: DefiLlama, files `defillama_stablecoins_d.json` and `defillama_tvl_d.json`; scripts `extract_crypto_mmf.py` and `f08_dnet.py`. 47

4.1 S_{Fed}: the two legs of Federal Reserve dollar supply, 2006–2026. Source: FRED `WALCL` and `DFF`; author's calculations; script `f11_sfed.py`. 56

4.2 S_{Tsy}: dollars parked outside the banking system, 2006–2026. Both series are balances at the Federal Reserve; both reduce effective dollar supply. Source: FRED `WTREGEN` and `RRPONTSYD`; author's calculations; script `f12_stsy.py`. . 58

4.3 S_{coll}, price leg: the repo–policy spread, SOFR minus interest on reserve balances, 2021-07-29 to 2026-09-03. Positive means collateral is scarce relative to cash. The series cannot begin earlier because IORB does not. Source: FRED `SOFR` and `IORB`; author's calculations; script `f13_scoll.py`. 60

4.4 S_{swap}: Federal Reserve central bank liquidity swaps outstanding, weekly, 2006–2026. Source: FRED `SWPT` (H.4.1); author's calculations; script `f14_sswap.py`. 63

4.5 S_{ext}: the U.S. balance on current account, quarterly, 2006-Q1 to 2026-Q1. All observations are deficits; a larger deficit is a larger export of dollars. Source: FRED `IEABC` (BEA); author's calculations; script `f15_sext.py`. 65

5.1 σ_{sys}, the Layer-1 systemic instability index, daily, 2006-02-02 to 2026-09-03, built from all three of E3's components: the VIX, a 21-day realised-volatility substitute for the paywalled MOVE, the Caldara–Iacoviello geopolitical risk index, and `STLFSI4` as the friction multiplier. Source: FRED `VIXCLS`, `DGS10`, `STLFSI4`; Caldara and Iacoviello [22]; author's construction; scripts `build_sigma.py` and `f17_sigma.py`. 77

5.2 The safe-haven response by regime under the canonical three-component state variable: monthly change in the dollar index against the contemporaneous monthly change in σ_{sys}, split by the level of σ_{sys} at the start of the month. The two fitted lines are nearly parallel, which is the finding of Section 5.6.4. Source: FRED `DTWEXAFEGS`; σ_{sys} as constructed above; script `f18_smile.py`. 81

7.1 $M_{\text{shock}}(t)$ under Equation (7.2)–Equation (7.3) for the two crises whose severity the original states. $T_{\text{career}} = 25$ years, implied by the original's own 2018 worked example; the two forgetting rates are illustrative and bracket a slow and a fast case, since the original gives no value for λ_f. Source: author's calculation from the stated formulas; script `f23_mshock.py`. 107

7.2 Empirical convexity: the dollar's response slope estimated separately within each quintile of σ_{sys}, with two-standard-error bars, under both constructions of the state variable. A rising profile would be evidence for $\kappa'(\sigma_{\text{sys}}) > 0$. Dependent variable `DTWEXAFEGS`, monthly, 246 observations. Source: author's calculation; script `f24_convexity.py`. 114

8.1 E7, the homogeneity convergence function, as a function of the composite drive $\kappa_A I_{\text{imp}}$. Panel (a) varies the steepness η at fixed $R_{\text{div}} = 0.5$: all curves cross at the turning point $\kappa_A I_{\text{imp}} = R_{\text{div}}$, where $H_{\text{soc}} = \frac{1}{2}$ by construction, and η decides only whether the transition is a slope or a cliff. Panel (b) varies the defensive variable R_{div} at fixed $\eta = 10$: raising resistance slides the curve right without flattening it. Source: author's construction; script `f25_hsoc.py`. 121

8.2 E7b after the ISSUES #13 closure. Panel (a): the reduced map $I_{\text{imp}} = g_0 + g_1 H_{\text{soc}}$ at four levels of friction, with $S_{\text{access}} = P_{\text{habit}} = E_{\text{sys}} = 0.9$, $c_{L0} = 0.15$, $c_{L1} = 0.75$. Friction scales the whole line toward zero; it cannot make the map decreasing. Panel (b): the admissible region for (g_0, g_1) required by Assumption 8.3, with the uniqueness boundary $\eta\kappa_A g_1 = 4$ of Proposition 8.2 drawn at $\eta\kappa_A = 8$. Source: author's construction; script `f26_iimp.py`. 124

8.3 The closed loop $T(H) = F(G(H))$ against the diagonal, at $\kappa_A = 0.8$, $g_0 = 0.15$, $g_1 = 0.75$. Panel (a), $\eta = 4$ so $\eta\kappa_A g_1 = 2.4 < 4$: one crossing, and the cobwebs from both ends walk into it (Proposition 8.2). Panel (b), $\eta = 20$ so $\eta\kappa_A g_1 = 12 > 4$: three crossings, the outer two locally stable (circles, $T' = 0.03$) and the middle one unstable (square, $T' = 3.0$), which is Proposition 8.3. Source: author's construction; script `f27_fixedpoint.py`. 126

9.1 E8, the cliff function. Panel (a) varies the erosion speed k_Ψ at $L_{\max} = 1$: the curves all cross at the threshold, so speed governs how long the transition takes and not where it begins — Assumption 9.2 made visible. Panel (b) varies $L_{\max}$ at $k_\Psi = 12$, with the floor $\Psi_{\min} = 1 - L_{\max}$ marked as a dotted line for each. Source: author's construction; script `f28_cliff.py`. 137

9.2 The two measured inputs to Layer 5. Panel (a): COFER dollar share of allocated official reserves, quarterly 1999-Q1 to 2026-Q1, 110 observations, with a linear and a quadratic fit in time. Panel (b): BIS triennial currency composition of OTC foreign-exchange turnover, log scale, shares summing to 200 by the BIS two-legged convention. Source: IMF COFER; BIS Triennial Central Bank Survey 2025, Table 4; built by `f29_shares.py` from the two extract scripts named in Section 6.5. 143

10.1 The two thresholds. Panel (a): the Chapter 8 object, a bifurcation diagram of the psycho-technical fixed point in the driver product κ_A, at $\eta = 20$, $g_0 = 0.15$, $g_1 = 0.75$, $R_{\text{div}} = 0.42$. Stable branches solid, the unstable separatrix dashed, and the hysteresis window between the two saddle-nodes shaded. Panel (b): the Chapter 9 object, E8 in f_{hard}, drawn once for each Layer-4 branch with k_Ψ mapped affinely from H^*_{soc}. The branches cross at θ_0: a jump in panel (a) rotates the curve in panel (b) about that point, raising $\Psi_{\text{str}}(t)$ below the threshold and lowering it above. Source: author's construction; script `f30_thresholds.py`. 152

10.2 Four stylised shocks on the parameter set of Table 10.3, twenty-four months each, decomposed into the three blocks of Equation (10.4): the mechanical demand-supply block scaled by κ_{sys}, the safe-haven exponential, and structural integrity. Dashed line is the total. Consolidates the outline's F31–F34 into one figure so that the four runs share axes and are comparable. Source: author's construction; script `f31_shocks.py`. 161

11.1 The observational base for every crisis-regime parameter in the model. Panel (a): σ_{sys} daily, 5,143 observations, with the 0.80 and 0.90 bands shaded. Panel (b): days per year in each band. Source: `sigma_sys_d.csv`; script `f35_scarcity.py`. 169

13.1 CS1 dashboard. Top: the dollar. Middle: σ_{sys} and its friction leg, both as percentile ranks in $[0, 1]$. Bottom: the TED spread and the three-month bill yield, the two Rule-1-compliant D_{prec} measures. Dashed lines mark Lehman (2008-09-15), the TED peak (2008-10-10) and the dollar's peak (2008-11-21). Source: `data/processed`; script `cs_dashboards.py`. 188

13.2 CS2 dashboard. Dashed lines mark the first shock (2020-03-09), the swap-line expansion to nine further central banks (2020-03-19) and the dollar's peak (2020-03-23). The bottom panel shows the bill yield going negative on 2020-03-26. Source: `data/processed`; script `cs_dashboards.py`. 192

13.3 CS3 dashboard, 2013-05 to 2018-12. Dashed lines mark the taper signal (2013-05-22), the end of purchases (2014-10-29) and the first hike (2015-12-16). The dollar's move begins between the second and third, not at the first. Source: `data/processed`; script `cs_dashboards.py`. 197

13.4 CS4 dashboard, both windows on a common layout. Left, 2011: the X-date (2011-08-02) and the S&P downgrade (2011-08-05), with the euro crisis visible in the friction leg and in VIX. Right, 2023: the deal (2023-06-03) and the Fitch downgrade (2023-08-01), with funding markets flat throughout. The contrast between the two bottom panels is Proposition 13.2. Source: `data/processed`; script `cs_dashboards.py`. 200

14.1 CS5 dashboard 207

14.2 CS9 dashboard 214

15.1 CS6 dashboard. Top: USD-pegged stablecoin supply and DeFi total value locked, log scale. Bottom: the dollar over the same span. Source: `data/processed`; script `cs_ch15.py`. 225

15.2 E15b's marginal effect on the safe-haven share, $\Delta s_{\text{US,sys}} = s(1-s)\gamma_G \Delta G_{\text{pol}}$ at $\gamma_G = 1$. Because it is a logistic derivative it is maximal at $s_{\text{US,sys}} = 0.5$ and falls away on both sides; the original's baseline of 0.70 sits off the peak. Source: author's construction; script `cs_ch15.py`. 228

15.3 CS8's four analogues, σ_{sys} in violet and the dollar rebased in blue, days either side of each event. The dashed line is the event. Source: `data/processed`; script `cs_ch15.py`. 232

16.1 CS11 dashboard, 2025. The dollar, σ_{sys}, the ten-year yield and gold on one time axis. Dashed lines mark the tariff announcement (2025-04-02) and the ninety-day pause (2025-04-09). Source: `data/processed` and the World Bank Pink Sheet; script `cs11_dashboard.py`. 242

18.1 The anchoring test 275

19.1 The Left Hand across three episodes 282

21.1 The attractor claim and its one test 302

22.1 The dollar across the Nixon shock 313

List of Tables

1.1 Where each strand stops, and which layer of this model takes over. 9
1.2 Structure of the argument. 12

2.1 The six strata. Each layer modulates the layers below it; none adds to them. 18
2.2 Parameters of Equation (2.1). Calibration values are set in Chapter 10; none is estimated here. 21

3.1 The dependent variable is a substitute, and the substitution is not neutral. Full provenance in `data/SOURCES.md`. 27
3.2 Data status of the ten demand cards. A card whose proxy has no free series states what is required and carries no historical range and no figure; no number in this chapter is estimated, interpolated or assumed. 28
3.3 Measurement protocol for D_{trans}. 29
3.4 Measurement protocol for D_{debt}. 32
3.5 Measurement protocol for D_{cap}. 34
3.6 Measurement protocol for D_{prec}. 37
3.7 Measurement protocol for D_{res}. 41
3.8 Measurement protocol for D_{hedge}. The primary proxy is not free; the substitutes are named, and so is what each of them misses. 43
3.9 Measurement protocol for D_{gamma}. 45
3.10 Measurement protocol for D_{net}. 46
3.11 Measurement protocol for D_{hcap}, flow and stock. 50
3.12 Interactions among demand channels and with the supply side. "Amplifies" means the effect is multiplicative and runs through κ_{sys}; "feeds" means one channel raises another's level with a lag. 51

4.1 Data status of the seven supply channels plus the two placement questions. No number in this chapter is estimated, interpolated or assumed. 54
4.2 Measurement protocol for S_{Fed}. 55
4.3 Measurement protocol for S_{Tsy}. Both legs enter with a negative sign. 57
4.4 Measurement protocol for S_{coll}. 59
4.5 Measurement protocol for S_{bank}. 61
4.6 Measurement protocol for S_{swap}. 62
4.7 Measurement protocol for S_{ext}. 64
4.8 Measurement protocol for S_{crypto}. 66

4.9 Double-counting checks on the denominator. "Overlap" is the quantity that could be counted twice; the resolution is what keeps it counted once. 70

5.1 Measurement protocol for $\sigma_{\text{VIX/MOVE}}$. The rates leg is a substitution and the table states what the substitution costs. 73
5.2 Measurement protocol for X_{pand}. 74
5.3 Measurement protocol for F_{liq}. 75
5.4 Where the 400 most stressed days fall, under the two constructions. The two-component index is $w_X = 0$; the canonical index is $w_X = 1$. 77
5.5 The ranked inputs to Equation (5.1) are not the same kind of series. Daily, 2006-02-02 to 2026-09-03, 5,143 observations, computed on the percentile ranks that actually enter the composite. "Daily rank move" is the mean absolute one-day change in the rank. 79
5.6 Contemporaneous response of the dollar to a change in systemic stress, by regime, under both constructions of the state variable. Monthly; dependent variable `DTWEXAFEGS`; slopes in per cent per unit of σ_{sys}; regime set by σ_{sys} at the start of the month. 80
5.7 The sign-change test across frequency and dependent variable, threshold fixed at $\sigma_{\text{sys}} = 0.5$, under both constructions. Slopes are per cent per unit of σ_{sys}. . 82
5.8 The two Chapter 5 results as a function of the geopolitical weight w_X, with $w_\sigma = w_F = 1$ throughout. "Steeper" counts specifications in which the stressed-regime slope exceeds the calm-regime slope. The headline specification is monthly, `DTWEXAFEGS`, threshold 0.50; "ratio" is its stressed slope divided by its calm slope. "Neg." and "sig." are the sign-change test: cells with a negative calm slope, and cells where that slope is also distinguishable from zero. "Verdict" summarises the amplification column only: *clear* where a steeper stressed slope holds in seventy per cent or more of specifications, *coin toss* at a share near one half, *absent* at a minority share. 84

6.1 Measurement protocol for G_{pol}. Every candidate is a partial indicator of one of the three capacities, not a measure of the variable. 90
6.2 The four channels through which internal U.S. risk acts. Every one is negative. 91
6.3 Measurement protocol for R_{US}. 92
6.4 Measurement protocol for M_{trust}. 94
6.5 Measurement protocol for T_{frag}. 95
6.6 Currency composition of OTC foreign-exchange turnover, per cent, net-net basis, April daily averages. Shares sum to 200 because every trade has two legs — the BIS convention, preserved here rather than rescaled. Source: BIS Triennial Central Bank Survey 2025, final results, Table 4; file `bis_triennial_fx_2025.xlsx`; script `extract_bis_triennial.py`. 96
6.7 Parameters of Equation (6.5). All are OPEN values in the registry; the CS7 figure is the original's assumption, retained and labelled. 98
6.8 The four structural variables: what each does, how fast, and how well it can be measured. "Speed" is the timescale on which a material move happens. . 101

7.1 Parameters of Equation (7.5). All are OPEN values in the registry. 111
7.2 Response slope by quintile of σ_{sys}, per cent per unit, with t-statistics. 114

8.1 The psycho-technical layer. Every variable is normalised to $[0, 1]$ (ISSUES #26). "Enters" names the equation in which the variable appears, not the quantity it ultimately affects. 118
8.2 Parameters of Equation (8.3). 122
8.3 Multi-dimensional impact of the psycho-technical layer, after the original's table on p. 22, re-set with canonical layer numbers and the equation that now carries each mechanism. Arrow scale per `BRIEF.md`: ↑ small, ↑↑ large, ↑↑↑ extreme. 128
8.4 Measurement protocol for the psycho-technical layer. "Tier" follows ISSUES #23: 1 free and available, 2 free but requiring construction, 3 paywalled or non-existent. 131

9.1 Parameters of Equation (9.2). All are OPEN values in the registry, and Section 9.7 argues that none of them is currently identifiable. 136
9.2 The six-stage tipping mechanism of §4.5.3, with the canonical variables and the observable that would identify each stage. The last column is the chapter's own addition: it records both whether the model implements the step and what would identify it, and it is the basis of Section 9.7. 140
9.3 The dollar's share of three reserve aggregates, year-end, per cent. Full series in `data/psi_gold.py`. 146
9.4 The quadratic term. E8 below its threshold requires a *negative* coefficient: an accelerating decline. Positive is deceleration, which E8 forbids there. 146

10.1 What each layer contributes to E9, and the constraint it carries with it. . . . 150
10.2 The six modulator functions in canonical form, with their ranges and the decision that fixed each. Every one is bounded; in the original, four were not. 158
10.3 The stated parameter set. These are a **convention adopted for illustration**, not estimates: nothing in this book has estimated any of them, and Chapter 11 explains why most of them cannot be. They are chosen only so that the qualitative behaviour is visible at plausible orders of magnitude. . . 160
10.4 The six strata, canonical numbering (ISSUES #15), with the object each contributes to Equation (10.2) and the chapter that builds it. 163

11.1 Census of the model as assembled. Counts from `registry/variables.csv` and `data/processed/`, reproducible by `notes/S12_audit.py`. 165
11.2 The proxy map. Coverage is what is in `data/processed` today, not what the provider publishes. 167
11.3 The ISSUES #31 flow proxies for D_{prec}, against the frequency of the mechanism they are supposed to measure. 170
11.4 What each strategy can reach. "Reachable" means the object could in principle be recovered from the data this book has; it does not mean it has been. . 171

11.5 The minimal calibration. Monthly, 2007-05 to 2026-06, $n = 230$, $R^2 = 0.16$. Dependent variable $100\,\Delta \ln$ `DTWEXAFEGS`. "Sign" is what the model requires: demand channels raise the dollar, supply channels lower it. 173
11.6 The dollar against each component of σ_{sys} separately. Monthly, 2006-03 to 2026-08, $n = 246$; dependent variable $100\,\Delta \ln$ `DTWEXAFEGS`; regressor is the contemporaneous change in the component's percentile rank. Script `data/haven_decomposition.py`. 175
11.7 The two crises decomposed across E10b's blocks, on the respecified estimates, summing fitted monthly contributions over each window. Script `i35_shortfall.py`.176
11.8 Mechanical coefficients by regime, σ_{sys} at the start of the month. The rival predicts every one to be larger in the stressed column. 176
11.9 The residual against funding pressure. Bill−fed funds is a bill–OIS proxy, negative when bills are bid. 177
11.10 What can and cannot be estimated today. "Verdict" is deliberately blunt. Ordered most damaging first. 179

12.1 What 2026 contains. Any window overlapping these dates is jointly determined by them. Sources in `data/SOURCES.md`. 184

13.1 CS1 as the original states it. Arrow scale per `BRIEF.md`: ↑ small, ↑↑ large, ↑↑↑ extreme. 186
13.2 CS1 data. Percentile ranks are the inputs to σ_{sys}; TED and the bill yield are levels. Sources: `data/processed`; script `data/crisis_studies.py`. 187
13.3 CS1 scorecard. Verdicts per Section 12.5; untestable entries stay in the denominator. 190
13.4 CS2 as the original states it (p. 31–32). 191
13.5 2020 split at the swap-line expansion of 19 March. 191
13.6 CS2 scorecard. 193
13.7 Daily changes in the dollar against daily changes in funding stress and in equity volatility, within each window. Univariate slopes and the joint regression. Script `data/crisis_studies.py`. 194
13.8 CS3 sub-windows. Source: `data/processed`; script `data/cs3_cs4.py`. . . . 196
13.9 CS3 scorecard. 198
13.10 CS4: the two debt-ceiling windows. Funding measure differs by necessity — `TEDRATE` was discontinued on 2022-01-21, so 2023 uses the SOFR–IORB spread (ISSUES #23). 199
13.11 CS4 scorecard, both windows. 201
13.12 Remark 13.1 out of sample. Funding measure is TED for 2011 and SOFR–IORB for 2023. 202

14.1 CS5 as the original states it. Nine predictions, two episodes, one number. . . 205
14.2 The two episodes. Event dates from the public record; all series from `data/processed`, sources in `data/SOURCES.md`. 205

14.3 CS5 measured. `DTWEXAFEGS` is the book's dollar series (Section 11.2); σ_{sys} is the canonical three-component construction of Section 5.6; GPRD is the Caldara–Iacoviello daily index. 206
14.4 The Caldara–Iacoviello daily index, against its own pre-window baseline. . . 206
14.5 The 2022 rally, decomposed by date rather than by narrative. 209
14.6 Two candidate drivers of the 2022 rally, tested on daily changes over the same window (2022-02-24 to 2022-09-27). 209
14.7 The safe-haven block against the observed move. Same estimates and same arithmetic as Table 11.7. 210
14.8 Case Study 5 scored, both episodes. Verdicts follow Section 12.5; percentages in Chapter 17 are over all rows including untestable ones. 211
14.9 CS9's impact decomposition, as the original states it for strict enforcement (original p. 41). Nineteen arrows. 212
14.10 Branch selection, from the scenarios' stated conditions only. Event facts from the public record, recorded in `data/SOURCES.md`; no dollar series was consulted in constructing this table. 213
14.11 CS9's window measured. All series from `data/processed`. 215
14.12 The dollar from the embargo date, at four horizons. 215
14.13 S_{swap} and the dollar credit aggregate across the embargo. `SWPT` is Federal Reserve central-bank liquidity swaps outstanding; the BIS series is dollar credit to non-bank borrowers outside the United States. 216
14.14 Two multi-scenario studies, and why only one of them can fail. 217
14.15 Case Study 9 scored on its realised branch. Nineteen block arrows plus Scenario 2's four predictions; arrows with no measure are untestable by Section 12.2, not omitted. 218
14.16 The episode. Sources recorded in `data/SOURCES.md` under "2026 Iran war". . 220
14.17 Case Study 12 scored. Market and index figures from `data/processed`; event facts from the sources recorded in `data/SOURCES.md`. 222

15.1 The seven steps of Section 12.2 against the three studies. "N/A" marks a step that does not apply to a hypothetical, for which STYLE substitutes an analogue table. 223
15.2 CS6 data. Source: `data/processed/defillama_stablecoin_usd_d.csv`; script `cs_ch15.py`. 225
15.3 CS6 scorecard. 226
15.4 Every criterion CS7 uses to score ΔG_{pol}, classified. 229
15.5 The United States has run CS7's Scenario 1 in the observable record. Event dates verified in `data/SOURCES.md`; response is `DTWEXAFEGS` over ± 20 trading days, from `data/cs7_cs8_gaps.py`. 230
15.6 CS7 scored on internal grounds, the only grounds available for a hypothetical. 231
15.7 CS8's four real analogues. σ_{sys} means are over the pre- and post-event halves of each window. Script `cs_ch15.py`. 232
15.8 The safe-haven block against the observed move, computed from Table 15.7's own σ_{sys} and dollar figures so that the two tables agree by construction. `data/cs7_cs8_gaps.py`. 233

15.9 CS8 scorecard against its own analogues. 234

16.1 The dates, stated plainly. 235
16.2 CS10 scored as written. Market data over the fixed window: `DTWEXAFEGS` −0.68 per cent; σ_{sys} 0.415 → 0.459; Brent +15.1 per cent. 237
16.3 CS11 data. Sources: `data/processed`; gold from the World Bank Pink Sheet (monthly). Script `cs11_dashboard.py`. 241
16.4 CS11 scorecard. Untestable entries stay in the denominator (Section 12.5). . 244

17.1 All one hundred and ten predictions across twelve case studies. Percentages are over *all* predictions, including untestable ones, per Section 12.5's denominator rule. 246
17.2 By case study. The last column is the share of *testable* predictions that were hits — reported for diagnosis only, and never as the headline, for the reason given below. 247
17.3 All 110 predictions by variable block. The last column is the share of that block's predictions that carry no verdict at all. 248
17.4 Every defect found, by kind. Data gaps close with work; specification gaps close only with a different model. 253
17.5 Channels proxied more slowly than the mechanism they describe. Full table of all seventeen in `data/audit_rescues.py`. 254
17.6 The one-factor solution. The geopolitical leg does not measure the factor the other three share. 255
17.7 The quadratic-in-time fit that produced Proposition 9.2, re-run on broader reserve aggregates. A *positive* quadratic is a decelerating decline, which is what E8 forbids below its threshold. Annual, $n = 27$. 257
17.8 The concession pattern. EPU falls; the dollar does not recover. 258
17.9 corr(ΔVIX, Δyield) by maturity. Negative is haven behaviour. The last column is the long end's excess over the short end: large and positive means the long end has lost a property the short end keeps. 259
17.10 Every negative result in the book, retyped after the audit. "Structural" means the model is wrong; "unmeasurable" means the claim cannot currently be checked either way. 260
17.11 The negative record split by kind, after reclassification. 260

18.1 Evidential standing of every channel Part III uses, from `registry/scorecard.csv`. This table is the reference for the notes that follow, and it is worth reading before the argument rather than after it. 266
18.2 The dilemma as the original states it, with each arm's evidential standing. . 269
18.3 Realised daily dollar volatility by regime, monthly, 2006–2026. Regime is set by σ_{sys} at the start of the month, on the threshold of 0.5 used throughout Chapter 5. Computed in `data/policy_evidence.py`. 270
18.4 The four spears, with the channel each runs through and its standing. Chapter 19 decomposes each into levers. 272
18.5 The measurement each spear is waiting on. "Tier" follows ISSUES #23. . . . 273

18.6 The 2022–23 tightening cycle. Sources: FRED `DFF`, `T5YIFR`, `MICH`; script `data/policy_evidence.py`. 273
18.7 Five-year, five-year-forward inflation compensation. The comparison that matters is the second column. 274
18.8 Federal Reserve central bank liquidity swaps outstanding, peak and mean, $ millions. Source: FRED `SWPT`. 276

19.1 Every lever, its channel, and what Chapter 17 says about that channel. "Hits" counts scored predictions in `registry/scorecard.csv`. 278
19.2 Institutional allocation. The last column is why coordination is hard, and it is institutional fact rather than model implication. 289
19.3 The original's expected impacts (E16), restated. None is established by this book. 290

20.1 The hour–day watchlist, mapped to series. "Held" means the series is in `data/processed` or `data/raw` and is used somewhere in this book. 292
20.2 The weekly watchlist. The original does not state one for this frame; this is the set implied by the actions it specifies. 293
20.3 The quarterly watchlist. This is the best-provisioned frame in the chapter. . 293
20.4 The annual watchlist. Two of five instruments exist. 294
20.5 Every condition of Trap 4, met simultaneously in 2022–23. 296
20.6 The predicted trajectory against the 2022–26 record. Dollar is `DTWEXAFEGS`; script `data/policy_evidence.py`. 297
20.7 Where the original says the fight is, and what this book can say about each. 299

21.1 The Plaza episode. Dates from the public record (`data/SOURCES.md`); the index is `DTWEXM`, the Federal Reserve's major-currencies trade-weighted dollar, whose first observation is 1973-01-02. 304
21.2 Three readings of the same seven months. 305
21.3 Mean-reversion test on the monthly log level of `DTWEXM`. The five per cent critical value for this statistic is about −2.86, not −1.96, because the regressor is a lagged level. 307
21.4 The chapter's four claims, sorted. 308

22.1 The chapter's instrument set. "Held" means the series is in this repository and covers the stated period. 310
22.2 The macroeconomic record across the Vietnam escalation. Sources: FRED `CPIAUCSL`, `UNRATE`, `FEDFUNDS`, `TRESEGUSM052N`; script `data/history_1971_1985.py`.311
22.3 The end of Bretton Woods. Event record verified and sourced in `data/SOURCES.md`.311
22.4 Bilateral dollar rates across the break, monthly. These are bilateral rates and are not combined into an index; no trade-weighted index covers this window. 312
22.5 Net exports of goods and services, quarterly, $bn (FRED `NETEXP`). The Bretton Woods parity was defended against a deteriorating external position. . . 313
22.6 The epilogue's claim, decomposed. 315

A.1 Notation. Layer, kind and original location follow `registry/variables.csv`. 316

C.1 The pipeline. Each script is runnable on its own and prints what it computed. 330
C.2 Held series, by the layer they serve. All are free and machine-readable. . . . 331
C.3 Every material absence, its kind, what was substituted, and what the substitution costs. "Ours" means the series is obtainable and this book has not obtained it. 333

E.1 The three passes. 342
E.2 Original to this edition. "Headings" counts the source's own section headings in that page range. 343
E.3 The four non-trivial mappings. 344
E.4 Equations whose chapter treatment differs from the original's. 344

Todo list

Part I

Theory: A State-Dependent, Reflexive Model of the Dollar

Chapter 1

Introduction: The Failure of Linear Models

1.1 The Dollar Index as a Measure of Global Liquidity

The U.S. Dollar Index (DXY) is not merely an exchange rate. It is the single most important measure of global liquidity, risk, and the cost of capital. A basket price against six advanced-economy currencies is, on its face, a narrow object; what makes it consequential is that the currency on one side of every one of those pairs is also the unit in which most of the world's trade is invoiced, most of its cross-border credit is denominated, and most of its reserves are held. When DXY moves, it does not merely reprice a basket. It reprices the liabilities of every borrower outside the United States who owes dollars, the collateral of every dealer who funds in them, and the reserve buffers of every central bank that holds them.

Three facts, each established in the literature reviewed in Section 1.4, make the index a system variable rather than a relative price.

First, the dollar is the dominant invoicing currency. Prices in international trade are sticky in dollars rather than in the currency of either the exporter or the importer [41]. A dollar appreciation is therefore not a symmetric terms-of-trade shift between two countries; it is a contraction in the purchasing power of everyone who earns in something else. This is the mechanism behind what we will call transactional demand, D_{trans} (Chapter 3).

Second, the dollar is the dominant funding currency. Dollar credit to non-bank borrowers outside the United States is measured in trillions and moves with U.S. monetary conditions and with dealer leverage [20, 56]. A borrower whose revenues are in local currency and whose liabilities are in dollars is short the dollar whether or not that position appears anywhere on a trading book. When the index rises, that short tightens, and servicing it requires buying more dollars — the debt-servicing channel D_{debt}.

Third, the dollar is the dominant reserve and safe-asset currency. Foreign demand for safe dollar claims is large enough that shifts in it move the exchange rate directly, and the convenience yield foreigners derive from holding Treasuries is itself a state variable of the international system [49, 51]. The long-standing framing of the United States as the world's banker — short-term liabilities, long-term assets, an excess return on the difference — makes this asymmetry structural rather than incidental [43, 44].

Taken together, these three facts imply that the dollar is a price with feedback into its own determinants. Nothing in the standard toolkit forbids that, but very little in the standard toolkit is built for it. That is the gap this work addresses.

1.2 Two Incomplete Traditions

src p.4

Academic attempts to model *DXY* fall into two incomplete categories.

Classical macroeconomic models. Purchasing-power parity, interest-rate differentials, and their descendants focus on "slow" variables: inflation differentials, growth differentials, productivity. They are useful over multi-decade horizons, and PPP does have long-run pull — but with a half-life of deviations so long, on the order of three to five years, that it is close to useless at business-cycle frequency [66]. Worse, this tradition has a known and unrepaired empirical failure at its centre: a random walk out-predicts structural exchange-rate models at short and medium horizons, a result that has survived four decades of attack [57]. The disconnect between exchange rates and macroeconomic fundamentals is itself one of the field's canonical puzzles [60].

The decisive objection for our purposes is narrower and sharper. These models completely fail to explain why *DXY* spiked in 2008 and again in 2020 — periods when the U.S. economy was evidently weakening. A framework in which the dollar is priced off relative U.S. fundamentals cannot generate a dollar that strengthens precisely as U.S. fundamentals deteriorate, unless something outside the fundamentals is doing the work.

Flow-based models. The second tradition tracks balance-of-payments movements: current-account balances, portfolio flows, direct investment. It is closer to the mechanism, and the recent literature that reconnects exchange rates to measured financial quantities has made real progress [36, 52]. But in its conventional form it ignores the most critical non-linear elements: the hidden currents inside the financial plumbing — repo markets, FX swaps, and Eurodollar debt servicing. These are not small residual frictions. In March 2020 the deviation from covered interest parity, the price of borrowing dollars synthetically, widened to levels that no flow accounting would have anticipated from the flows themselves [6, 25], and the U.S. Treasury market — the deepest in the world — required central-bank intervention to function [26, 74].

What both traditions share. The greatest flaw of these models is the assumption of linearity. They assume that a one-percent rate shock always has a similar effect, regardless of the regime. Reality is non-linear. The two traditions fail in the same place and for the same reason: each posits a fixed mapping from causes to the dollar, and the mapping is not fixed.

1.3 The Linearity Assumption and the Non-Linear Response

The regularity that motivates this work is that the dollar does not respond to stress the way a linear model says it should. Two versions of that statement circulate, they are usually run together, and this book separates them — because on the data assembled here one survives and the other does not.

The received view: a sign change. The standard account, and the account the original paper gives, is that the dollar changes *sign*. When global financial stress is low, the dollar behaves like a risk-on asset: it is sold as capital leaves the United States chasing higher returns elsewhere. Once stress crosses a threshold it flips, becomes a risk-off safe haven, and reacts convexly. This is the "dollar smile" of the practitioner literature, and it is the claim that gives the original its title. Formally: $\partial DXY/\partial\sigma_{\text{sys}} < 0$ below some σ^* and > 0 above it.

What the data support: no sign change, and an amplification that depends on how stress is measured. We test that claim in Chapter 5 and it does not hold.

Across 452 specifications — three frequencies (monthly, weekly, daily), two dollar indices, thirteen candidate thresholds from 0.20 to 0.80, and six constructions of the state variable spanning the full range of the one weight that turns out to matter — **no specification produces a negative calm-regime response that is statistically distinguishable from zero**. The dollar is a safe haven at every level of systemic stress in this sample. It is not sold in calm periods when stress rises; it is bought, merely less enthusiastically.

The weaker claim — that the response is *larger* in the stressed regime — is not settled here at all, and the honest report is that it *cannot* be settled on this evidence. It holds in 82 per cent of specifications when the state variable is built from financial stress alone, and the share falls monotonically, through 74 and 70 per cent, to 46 per cent once geopolitical risk is added at equal weight. Whether the dollar's response to stress amplifies therefore turns on how much weight non-financial stress carries inside the state variable — a parameter of the model that has not been estimated. The claim is reported as *undetermined* rather than as supported or refuted, together with the decision rule that will settle it: a weight at or below one half restores the amplification, one at or above three quarters retires it. Chapter 11 estimates the weight and the verdict follows from it.

Section 5.6.3 to Section 5.6.5 give the grid, the counts, the profile and the two propositions; Chapter 17 carries them into the scoreboard. We state the result here, in the introduction, rather than letting the reader reach Chapter 5 still believing a claim the book goes on to reject.

Why the non-linearity survives the correction. None of this rescues the linear model, and the case for a state-dependent specification is unaffected.

A model with constant elasticities cannot represent a response whose magnitude more than doubles between regimes, which is what the two-component evidence shows and what the episode studies of Chapter 13 show at daily resolution. Nor can it represent the central

puzzle of Section 1.2: a dollar that rises when the U.S. economy is visibly weakening, in 2008 and again in 2020. Those facts require the demand-side composition to change with the state — safe-haven demand displacing capital-asset demand as stress rises — and that is a state-dependent model whether or not any derivative changes sign.

What changes is the *shape* of the claim. The original asserts a curve with two arms of opposite sign. What the data here support is a curve that is everywhere positively sloped and steeper on the right: the dollar is always a haven, and stress governs how violently it behaves as one. That is a weaker claim than the smile, and it is the one this book defends.

The convexity is still the object of interest. A safe-haven bid that scales linearly with fear is a portfolio preference. One that scales super-linearly is a mechanism with positive feedback inside it, and mechanisms with positive feedback have thresholds. The original captures this with an exponential term in stress and a state-dependent exponent κ_{sys} on the whole demand–supply ratio; we retain both, make the exponent's arguments explicit, and bound the exponential so that it cannot explode (Chapter 7, Chapter 10). We also replace the original's raw VIX thermometer with a bounded systemic instability index $\sigma_{\text{sys}} \in [0, 1]$, a change forced by a defect in the original specification (Chapter 5; ISSUES #9).

This work therefore addresses a non-linearity, but a different one from the one it set out to address. We hypothesise that DXY is a function of the ratio between global effective demand for the dollar, D^*, and global effective supply, S^*, raised to a state-dependent power and multiplied by a stress-activated safe-haven term. The full statement is E1, restated canonically in Chapter 2; the final integrated form is E9 in Chapter 10. The core innovation is that the model's weights and its convexity are dynamic functions of the system's state rather than constants, and that a small set of meta-variables — geopolitical power G_{pol}, internal U.S. risk R_{US}, and human capital H_{cap} — enter not as direct inputs but as *modulators* that alter the system's internal behaviour.

1.4 Related Literature

This section positions the model against five strands of work. The purpose is not completeness but placement: for each strand we state what it establishes, and which component of the model it constrains.

1.4.1 Safe-asset demand and the dollar

The modern theory of dollar dominance treats the United States as the supplier of the world's safe asset and derives the exchange rate from the market-clearing condition for that asset. Gourinchas and Rey [43] document the balance-sheet asymmetry — borrowing short, lending long, earning the spread — and Gourinchas et al. [44] show its counterpart: in crises the hegemon pays out, transferring wealth to the rest of the world precisely when the rest of the world needs it. Farhi and Maggiori [32] embed this in a formal international monetary system with a hegemon whose safe-asset supply is limited by its own fiscal capacity, and show how the system behaves when reserve assets are scarce and when they are contested. Caballero et al. [21] document the shortage itself.

The empirical counterpart most directly relevant here is Jiang et al. [49]: variation in foreign demand for safe dollar assets moves the dollar exchange rate, and the Treasury basis provides a direct price for that demand [51]. Gopinath et al. [41] establish the invoicing channel. Ilzetzki et al. [48] document the anchoring share, and Arslanalp et al. [5] the slow decline in the dollar's reserve share and the direction of the drift.

What this constrains. These papers supply the mechanism behind our safe-haven demand term D_{prec} and the share $s_{\text{US,sys}}$ that determines how much of a global flight-to-quality the dollar actually captures (Chapter 6). They also supply the empirical basis for Layer 5: erosion of the dollar's structural position, $\Psi_{\text{str}}(t)$, is not a speculative construct but a measured trend [5], whose historical precedents are documented at book length [28, 30, 63].

1.4.2 The global financial cycle and the dollar smile

Rey [65] showed that gross capital flows, leverage and asset prices across the world co-move in a single global financial cycle driven substantially by U.S. monetary policy, and that this cycle constrains monetary autonomy even under floating exchange rates. Miranda-Agrippino and Rey [58] quantify the transmission. Lilley et al. [52] document that proxies for global risk appetite, which had no explanatory power for exchange rates before 2007, acquired it afterwards — a regime-dependence in the data that is precisely what a state-dependent model predicts.

The practitioner literature reached a shape earlier, and it is the shape this book tests rather than adopts. The "dollar smile" — the dollar strengthening both when U.S. growth is exceptional and when global risk appetite collapses, and *weakening in the benign middle* — is the received view of the dollar's non-linearity. The *label* originates in a Morgan Stanley research note of 2001 by Stephen Jen; that note has no public publisher record, so it is named here but not cited, and it does not appear in the bibliography.

It is worth being explicit about the status of what follows. The peer-reviewed work below is cited as the best available statement of the received view, and as the standard against which Chapter 5 measures. It is *not* cited as support for the smile: as Section 1.3 reports and Section 5.6.4 establishes, the middle arm — the weakening in benign conditions — is not present in this sample.

Maggiori [53] derives the mechanism in equilibrium: the reserve currency appreciates during global crises because the reserve country's financial intermediaries absorb more risk, which is precisely what makes its currency a hedge and therefore makes it the reserve currency. Georgiadis et al. [39] identify the empirical counterpart with high-frequency risk shocks: global risk shocks appreciate the dollar, tighten global financial conditions, and contract world activity — and in counterfactuals where the dollar does not appreciate, the contraction is much weaker, so the appreciation is doing the transmitting rather than merely accompanying it. Farhi and Gabaix [31] supply the non-linearity: when the probability of rare disasters varies over time, exchange rates inherit the convexity of the underlying risk premium instead of responding proportionally to it. Lilley et al. [52] supply the state-dependence in its most direct form, since the risk-appetite proxies that explain exchange rates after 2007 explained nothing before it.

Three qualifications keep the claim honest, and the third is the one this edition adds.

First, what this literature documents is the risk-off arm of the curve and its non-linearity;

the strong-growth arm — the dollar also rising when U.S. growth is exceptional — rests on the ordinary capital-flow channel carried by D_{cap} in Chapter 3, and we treat the two arms as separate mechanisms rather than assuming the folk curve as a single object.

Second, **none of these papers estimates a sign change** in $\partial DXY/\partial\sigma_{\text{sys}}$ directly. Each documents that risk shocks appreciate the dollar, or that the response is convex in disaster risk, or that risk proxies gained explanatory power after 2007. None shows the dollar being *sold* as stress rises in calm regimes. The middle arm of the smile is the part of the received view that the peer-reviewed literature does not establish, and it is exactly the part Chapter 5 fails to find.

Third, and consequently, the model this book builds answers to the risk-off arm and the convexity, not to the smile as a whole. Where Chapter 13 and Chapter 16 appear to show a risk-on dollar in a particular episode, that observation must be reconciled with Proposition 5.1 rather than presented as belated confirmation of the folk curve.

What this constrains. Layer 1 (Chapter 5). The global financial cycle literature tells us that a single state variable carries a great deal of the common variation; that is the empirical licence for compressing stress into σ_{sys} rather than carrying a dozen separate risk factors.

1.4.3 Dollar funding plumbing

The third strand is the one the flow tradition omits, and it is where the model's supply side lives. Du et al. [25] document persistent, cyclical deviations from covered interest parity after 2008 — the price of synthetic dollar funding, and therefore the direct observable behind our hedging-demand term D_{hedge}. Avdjiev et al. [6] establish the triangular relationship that matters most for this model: a stronger dollar goes together with wider CIP deviations *and* with contracting cross-border dollar lending. That is a feedback loop stated as an empirical regularity, and it is the strongest single piece of evidence that the dollar is not an exogenous price.

On the collateral and dealer side, Brunnermeier and Pedersen [19] give the canonical model of the spiral between market liquidity and funding liquidity, and Adrian and Shin [1] the procyclicality of intermediary balance sheets: these justify treating our convexity term κ_{sys} as a function of leverage rather than a constant. Pozsar [62] maps non-bank dollar creation, the object we place inside S_{coll} (Chapter 4). Two episodes anchor the empirics: September 2019, when reserves that were believed ample proved not to be [2, 24], and March 2020, when Treasury market functioning itself broke down [26, 74]. Central-bank swap lines, our S_{swap}, are the designed relief valve, and their effects are now measured rather than assumed [8].

What this constrains. Chapter 4 in its entirety, and the claim in Chapter 3 that D_{debt} and D_{hedge} are stress-activated rather than steady.

1.4.4 Reflexivity, memory, and narrative

The model's third and fourth layers assert that market participants' beliefs feed back into the quantities the model is trying to price. This is Soros's reflexivity, stated as a general principle of financial markets rather than a special case [70], and it is the origin of our meta-reflexivity driver A_{meta}.

Two empirical literatures make the principle tractable. Malmendier and Nagel [54] show that individuals who lived through low market returns take less financial risk decades later, and Malmendier and Nagel [55] show the same experience-weighting for inflation expectations, with younger agents updating more strongly. This is the direct empirical foundation for our crisis-memory variable M_{shock} and for its generational decay: a crisis is forgotten not when it is far away in time but when the people who traded through it have left the market (Chapter 7). Shiller [68] supplies the epidemiological framing for how economic narratives spread, which is what our narrative-velocity variable measures.

For the psycho-technical layer specifically — the claim that algorithmic media homogenise the beliefs of market participants, raising H_{soc} — the evidence is thinner and more contested, and we say so in Chapter 8. Bakshy et al. [10] find that algorithmic ranking does reduce exposure to cross-cutting content, but that individual choice reduces it more; this is a real effect of a modest measured size, not the sweeping one the mechanism would like. Haldane and May [46] supply the structural argument from the other direction: in a network of institutions, homogeneity of strategy is itself a source of systemic fragility, which is exactly the role H_{soc} plays in amplifying κ_{sys}.

What this constrains. Chapter 7 and Chapter 8. It also constrains our honesty: of the model's five layers, Layer 4 is the one whose variables are least measurable, and Chapter 11 treats it accordingly.

1.4.5 Coercion, fragmentation, and erosion

The final strand concerns what happens when the dollar system is used as an instrument. Farrell and Newman [33] give the canonical account of how centrality in economic networks converts into coercive capacity — the panopticon and chokepoint effects — which is the theoretical basis for our sanctions variable S_{off} and, in the long run, for its backlash through the fragmentation index T_{frag}. Mulder [59] supplies the history, and the reminder that the economic weapon has repeatedly produced counter-mobilisation rather than submission. Arslanalp et al. [5] supply the measured drift in reserve composition that any claim about erosion must be calibrated against.

What this constrains. Chapter 6 and Chapter 9, and the case studies of Chapter 14 and Chapter 16.

1.4.6 Stress measurement and estimation

Layer 1 and Chapter 11 draw on the index literature: Caldara and Iacoviello [22] for geopolitical risk, Baker et al. [9] for policy uncertainty, and Azzimonti [7] for partisan conflict — the last two being the free proxies we propose for internal U.S. risk R_{US}. On the estimation side, the model's regime structure places it in the smooth-transition and regime-switching families: Hamilton [47] for discrete switching and Teräsvirta [73] for the logistic transition function we adopt in Chapter 2 to interpolate elasticities between calm and crisis values.

1.4.7 The gap

Each strand above is strong within its own boundary, and each stops at a different one. The safe-asset literature has a rigorous account of *why* the dollar is demanded in crisis but generally holds the plumbing fixed. The plumbing literature has a precise account of *how* dollar funding seizes but generally takes the crisis as given. The global financial cycle literature identifies the common factor but does not decompose it. The behavioural literature establishes that memory and narrative matter without connecting them to a currency's price. The geoeconomics literature establishes that the system can be weaponised without pricing the consequence.

Table 1.1 states the position directly. The contribution claimed here is not a new result in any one column; it is a single specification in which all five appear as named terms with signs, so that they can contradict one another and be tested together.

Table 1.1: Where each strand stops, and which layer of this model takes over.

Strand	Establishes	Held fixed / not priced
Safe assets	Crisis demand for USD claims; convenience yield	Funding mechanics; the speed of the flip
Plumbing	How dollar funding seizes; CIP, repo, swap lines	The trigger; the political layer
Global financial cycle	A single common factor drives world risk assets	Its decomposition into demand and supply
Behavioural	Experience and narrative shape risk-taking	The mapping into a currency price
Geoeconomics	Network centrality is coercive capacity	The cost of exercising it, in index points

1.5 Contributions

This book is a test of a framework, not a presentation of one. The four claims below are the original's; each is stated here as it was made, and each is scored in Chapter 17 against 110 predictions drawn from twelve case studies. **A substantial share fail.** The headline is 27 hits, 13 partials, 19 misses and 46 predictions that cannot be tested at all — so more of this framework's claims are *untestable* than are correct. A reader should arrive at Chapter 17 already knowing that, and Section 1.5.1 says which parts survive before the model is built.

The four claims, in the original's descending order of confidence:

1. A six-layer decomposition. The dollar's price is decomposed into six strata: Layer 0, the mechanical demand–supply ratio D^*/S^*; Layer 1, the state variable σ_{sys}; Layer 2, the

structural and trust layer ($\Phi_{\text{sys}}, s_{\text{US,sys}}$); Layer 3, the reflexive acceleration κ_{sys}; Layer 4, the psycho-technical states ($H_{\text{soc}}, I_{\text{imp}}$); and Layer 5, the secular integrity coefficient $\Psi_{\text{str}}(t)$. Each layer modulates the layers below it rather than adding to them. The claim is not that these are the only six; it is that this particular cut separates variables that move on different timescales — days, months, decades — and that a great deal of the confusion in the existing debate comes from mixing them.

Verdict: **survives**. It is the contribution that does best. Section 17.4 records three results the decomposition produced that a flat model cannot state — an existence theorem for the belief equilibrium, a contraction bound, and a located threshold — and Table 17.3 shows the separation is real in the data as well as on paper: the layers score differently, and they score differently in a way that tracks timescale.

2. State-dependent elasticities. Every weight in the model is a function of σ_{sys} rather than a constant, interpolating between a calm value and a crisis value through a shared logistic transition. This is the formal content of the risk-on/risk-off flip, and it is what allows one specification to reproduce both the 2014–2015 dollar rally on growth differentials and the 2008 and 2020 rallies on funding stress, without special pleading in either case.

Verdict: **split, and the two halves must not be confused**. The *amplification* half survives: the response to stress is steeper in stressed states than calm ones in a majority of specifications. It survives *conditionally*, because the majority decays from 82 to 46 per cent as the unestimated weight w_X runs from zero to one (Table 5.8), and this book does not choose w_X. The *sign change* — the claim that the dollar trades risk-on below a threshold and flips above it — **does not survive**: it is absent from all 452 specifications tested. And Chapter 17 records a third finding against the mechanism's placement rather than its existence: convexity in a bounded variable is largest in the middle of the range, so the crisis mechanism is weakest at the stress levels crises actually occupy (ISSUES #37).

3. A psycho-technical layer. We treat the homogeneity of market participants' beliefs, H_{soc}, and the immediacy with which belief converts to action, I_{imp}, as state variables that amplify the reflexive layer. This is the most speculative of the four contributions and it is flagged as such throughout: Chapter 8 states what is measurable, what is proxied, and what is currently only argued.

Verdict: **mathematically productive, empirically unmeasured**. The layer yields the book's sharpest formal results — an existence proof, a uniqueness bound of $\frac{1}{4}\eta\kappa_A g_1$, at most three equilibria, and a located tipping threshold in H_{soc} — and it has *nine variables and no series*. Nothing in it has been confirmed or refuted by data, and Remark 10.2 shows Layer 4 has no inbound edge from any other layer, so one of the six stages of the tipping mechanism is not in the equations at all.

4. An erosion coefficient. Layer 5 introduces $\Psi_{\text{str}}(t)$, a multiplicative structural-integrity coefficient bounded in $[1 - L_{\max}, 1]$, equal to one when the dollar's institutional position is intact and falling as alternative infrastructure C_{alt} and fragmentation T_{frag} accumulate. Its functional form is a cliff rather than a slope: erosion is slow and then sudden. Reserve-currency transitions are the historical evidence base [28, 30], and the current measured

drift [5] is the calibration target.

Verdict: **unfalsifiable as specified**. E8 requires an accelerating decline below its threshold, and no measure of the hard-driver argument f_{hard} exists, so the cliff cannot be confirmed or refuted on anything currently observable (Proposition 9.3). An earlier draft of this book reported that the measured erosion *decelerates*, contradicting the cliff. Section 9.7.1 withdraws that: the deceleration was computed on COFER, which excludes monetary gold — so a reserve manager who sells dollars for gold leaves the denominator and the dollar's share does not move — and on an aggregate that includes gold the decline accelerates (-0.0211 pp/yr^2, $t = -2.59$). The model is no longer contradicted by its own erosion evidence. It is still not testable.

1.5.1 What survives testing

Stated here rather than held back, because a reader is entitled to know the result before reading three hundred pages of the argument for it. Section 17.4 gives each of these in full with its evidence.

1. **The decomposition.** Separating the price into strata that move on different timescales does work a single-equation model cannot. It is the framework's most durable contribution and it is the reason the failures below are legible as failures rather than as noise.
2. **Convexity, as amplification and not as a sign change.** The dollar's response to stress steepens as stress rises, in a majority of specifications and conditional on a weight this book declines to choose. The dollar is a safe haven in *both* regimes; what it is not is a risk asset in calm ones.
3. **The Layer 4 threshold.** Chapter 8 locates a tipping point in belief homogeneity, proves the fixed point exists, bounds the map's slope, and shows there are at most three equilibria — and gives "hyper-brittleness" an actual number in the amplification factor $1/(1 - T')$. This is a real derivation from the original's own premises. It is also entirely unmeasured, and both halves of that sentence are load-bearing.
4. **The CS7 result on military capability.** An implication of the model's own arithmetic that inverts the obvious intuition: because $s_{\text{US,sys}}$ is logistic, its derivative $s(1-s)$ is largest at intermediate shares, so an increment of capability moves the dollar *least* for the incumbent whose position is strongest. Chapter 15 scores four of four testable predictions here, and CS7 is the book's methodological standard.
5. **Funding stress against general stress.** The composite state variable was hiding a distinction the case studies then made visible: the dollar responds to the *funding* leg of σ_{sys} and has no measurable relationship with the geopolitical leg — a slope of -0.08 with $t = -0.19$ over 246 months. Section 17.6.2 explains why: the geopolitical leg loads -0.239 on the factor the other legs share. This is a finding the original does not contain and it changes what a "safe haven" means in the model.

That is the list. It does not include the model as a whole, it does not include the dollar smile, it does not include the erosion cliff, and it does not include Layer 3, which has four predictions in the record and no verdicts. Section 17.7 sets out the case against the framework in six numbered lines, and Section 17.6 reports four pre-registered attempts to attribute parts of the negative record to measurement rather than to error — three of which

failed.

What this work does not claim. It is not a forecasting model, and no statement in it should be read as a prediction. The specification has not been estimated on a full sample; Chapter 11 reports what can be estimated today with free data and what cannot, including the failures — and 76 per cent of the model's parameters have never been given a value. **Nor does this book claim that the framework has been vindicated.** On the one block of predictions that is fully testable, the dollar's own direction, the record is three hits, one partial and three misses. Several variables — particularly in Layer 4 — have no satisfactory proxy at present, and where that is true the text says so rather than substituting a convenient one. The legal disclaimer in the front matter is not a formality.

1.6 Roadmap

Table 1.2 maps the argument. Part I builds the model layer by layer; Part II tests it against **twelve** episodes and scores the results; Part III applies it to policy.

Part II is where the book's claims are settled, and it is written to be usable against the model rather than for it. Chapter 12 fixes the scoring rules *before* any episode is examined — including a denominator rule that counts untestable predictions in the denominator, which costs the book twenty percentage points of hit rate and is kept for that reason. Chapter 17 reports the totals, the defects sorted into those a data programme could close and those only a different model could, and a set of pre-registered rescue hypotheses tested against criteria committed to the repository before the tests were run.

Table 1.2: Structure of the argument.

Chapter	Layer / object	Question answered
2	Architecture	What is the equation, and why multiplicative?
3	Layer 0, numerator	Who buys dollars, and why?
4	Layer 0, denominator	Who supplies them, and what constrains it?
5	Layer 1, σ_{sys}	How is systemic stress measured on a bounded scale?
6	Layer 2, Φ_{sys}, $s_{\text{US,sys}}$	Why does the world still choose the dollar in a crisis?
7	Layer 3, κ_{sys}	Why do dollar moves accelerate?
8	Layer 4, H_{soc}, I_{imp}	What happens when beliefs synchronise?
9	Layer 5, $\Psi_{\text{str}}(t)$	How does a reserve currency actually end?
10	Synthesis	The assembled equation and its comparative statics
11	Measurement	What can be estimated today, and what cannot
12–17	Evidence	Twelve episodes, scored term by term: 110 predictions, 27 hits, 46 untestable
18–20	Policy	Inflation crisis, playbook, execution
21–22	System	Coordination, and the 1971 inflection

Three conventions apply throughout and are stated once, here, in summary form; the full statement is in Chapter 2 and the complete symbol table is Chapter A.

1. **Every latent modulator lives in** $[0, 1]$**.** σ_{sys}, G_{pol}, R_{US}, M_{trust}, T_{frag}, H_{soc}, I_{imp} and the rest are bounded indices, so that every coefficient attached to them reads as the effect of a full-range move and coefficients are comparable across layers.
2. **Erosion is always** $\Psi_{\text{str}}(t) \downarrow$**.** $\Psi_{\text{str}}(t)$ is a structural *integrity* coefficient: it equals one when the system is intact and falls as the dollar's position erodes. The arrow never points the other way.
3. **Quantities are indexed, not levelled.** Each D_i and S_j is an index number equal to one at a stated base period, which is what identifies $\Psi_{\text{str}}(t)$ against the baseline DXY_0 in Chapter 10. Channels that did not exist at the base period — stablecoins, decentralised dollar supply, swap lines outside crisis windows — enter in a form that equals one when the channel is absent, so that the normalisation survives.

A reader who wants the model in one page should read Chapter 2 and then Chapter 10. **A reader who wants to know whether it survives contact with data should read Chapter 17 first and work backwards**, and that is a recommendation rather than a courtesy: the scorecard is the part of this book with the most information in it per page, and several chapters in Part I state claims that Chapter 17 then withdraws. A reader who came for the policy argument should start at Chapter 18; the layer references there will bring them back into Part I where they need it.

Chapter 2

General Model Architecture

2.1 The Foundational Equation

To capture the dynamics set out in Chapter 1 we propose a state-dependent integrated equation. Instead of additive linear terms, the model employs a multiplicative structure, similar to a Cobb–Douglas production function, to reflect the intertwined interaction between supply and demand forces.

In its foundational form the equation is

$$DXY \;=\; DXY_0 \left[\frac{\mathcal{N}(\sigma_{\text{sys}})}{\mathcal{D}(\sigma_{\text{sys}})}\right]^{\kappa_{\text{sys}}} \exp\Big[(\lambda_1 \sigma_{\text{sys}} + \lambda_2 \sigma_{\text{sys}}^2)\, \Phi_{\text{sys}}\Big] \;+\; \epsilon, \tag{2.1}$$

where DXY_0 is the baseline level, ϵ an exogenous random shock, and the numerator and denominator aggregates are

$$\mathcal{N}(\sigma_{\text{sys}}) = (s_{\text{US,sys}}\, D_{\text{prec}})^{\alpha_{\text{prec}}(\sigma)} \prod_{i \in \mathcal{I}} D_i^{\alpha_{\text{i}}(\sigma_{\text{sys}})}, \tag{2.2}$$

$$\mathcal{D}(\sigma_{\text{sys}}) = \prod_{j \in \mathcal{J}} S_j^{\beta_{\text{j}}(\sigma_{\text{sys}})}, \tag{2.3}$$

with the index sets fixed by name,

$$\mathcal{I} = \{\text{trans},\ \text{debt},\ \text{cap},\ \text{res},\ \text{hedge},\ \text{gamma},\ \text{net},\ \text{micro},\ \text{hcap}\}, \tag{2.4}$$

$$\mathcal{J} = \{\text{Fed},\ \text{Tsy},\ \text{coll},\ \text{bank},\ \text{swap},\ \text{ext},\ \text{crypto}\}. \tag{2.5}$$

Two things about Equation (2.2) are worth stating immediately, because both differ from the original text and both are consequential.

First, *prec* is excluded from the product $\prod_{i \in \mathcal{I}}$. Safe-haven demand enters once, through its own factor. In the original, the product ran over all ten demand channels *and* carried a separate D_{prec} factor, so safe-haven demand was counted twice (ISSUES #7).

Second, the safe-flow share $s_{\text{US,sys}}$ sits *inside* the base rather than multiplying the exponent. The economic object is USD-captured safe-haven demand, $s_{\text{US,sys}} D_{\text{prec}}$: the share of a global flight to quality that the dollar actually absorbs. Written the original way, as

$D_{\text{prec}}^{\alpha_{\text{prec}}(\sigma)s_{\text{US,sys}}}$, a share scales an elasticity — which is dimensionally odd, and has the wrong limit, since $s_{\text{US,sys}} \to 0$ sends the factor to 1 rather than to "no USD-captured safe-haven demand" (ISSUES #8).

The ten demand channels of Chapter 3 are transactional demand D_{trans}, debt-servicing demand D_{debt}, capital-asset demand D_{cap}, safe-haven demand D_{prec}, reserve demand D_{res}, hedging demand D_{hedge}, derivative-dealer demand D_{gamma}, network demand D_{net}, microstructural demand D_{micro}, and human-capital demand D_{hcap}. The seven supply channels of Chapter 4 are Federal Reserve actions S_{Fed}, Treasury actions S_{Tsy}, collateral availability S_{coll}, bank balance-sheet capacity S_{bank}, central-bank swap lines S_{swap}, external flows S_{ext}, and synthetic dollar supply S_{crypto}.

2.1.1 Effective external supply

One channel is not used raw. Sanctions and asset freezes remove dollars from usable external supply, and the model carries this through an attenuation of the external channel:

$$S_{\text{ext}}^{\text{eff}} \;=\; S_{\text{ext}}\, e^{-\delta_{\text{off}}\, S_{\text{off}}}, \qquad S_{\text{off}} \in [0,1],\ \delta_{\text{off}} > 0, \tag{2.6}$$

so that $S_{\text{ext}}^{\text{eff}}$, not S_{ext}, is the quantity carrying $\beta_{\text{ext}}(\sigma_{\text{sys}})$ in Equation (2.3). Here S_{off} is an *immobilisation intensity* — the share of externally held USD claims blocked — rather than a dollar stock. In logs the channel contributes $\ln S_{\text{ext}} - \delta_{\text{off}} S_{\text{off}}$, which is defined for every admissible value; the subtraction $S_{\text{ext}} - S_{\text{off}}$ that suggests itself first is not, since a large freeze can exceed a quarter's external flow (ISSUES #28).

2.1.2 Emergent channels

Three channels did not exist, or were exactly zero, at the base period against which everything is indexed (Section 2.6.2): network demand D_{net}, which is carried largely by dollar stablecoins; synthetic supply S_{crypto}; and swap lines S_{swap}, which sit at zero in any calm month. For these the model uses

$$X \;\longmapsto\; 1 + \frac{X_t}{\bar{X}}, \tag{2.7}$$

with $\bar{X}$ a stated scale constant, fixed by convention and reported rather than estimated. The term equals exactly 1 while the channel is absent, so the normalisation of Section 2.6.2 survives term by term, $\ln(1 + X/\bar{X})$ is defined and smooth throughout, and the elasticity measures the response to the channel's *emergence* rather than only to its growth after birth (ISSUES #29).

2.2 Architectural Philosophy: Four Pillars

src p.4

The model's design rests on four pillars. Each is a modelling commitment, and each can be wrong in a specific way; we state both.

2.2.1 Pillar 1 — the multiplicative ratio

The core is the ratio of dollar "pull" forces in the numerator to dollar "push" forces in the denominator. The form is multiplicative rather than additive, and the reason is mutual amplification: elevated safe-haven demand D_{prec} intensifies the effect of a collateral shortage S_{coll}, rather than merely adding to it. In an additive specification a collapse in one supply channel can always be offset by an expansion in another; in a multiplicative one it cannot, because each channel enters with its own exponent and a channel that goes to zero takes the product with it.

Written in logs, the pillar says that channels combine as a weighted sum of log quantities — constant elasticity, no substitution parameter to estimate, and every coefficient interpretable as a percentage response to a percentage move. That is the Cobb–Douglas analogy, and it is also the pillar's weakness: constant elasticity of substitution is assumed rather than tested, and Chapter 11 says what would be needed to test it.

2.2.2 Pillar 2 — the state variable

The second pillar is that the weights themselves move. The state variable σ_{sys} is the temperature of the financial system, and it governs the elasticities $\alpha_{\text{i}}(\sigma_{\text{sys}})$ and $\beta_{\text{j}}(\sigma_{\text{sys}})$ across regimes. Under low stress, capital-asset demand D_{cap} may dominate; under crisis, safe-haven demand D_{prec} surges and the weights shift accordingly. This is what makes the model state-*dependent* rather than merely non-linear, and it is the formal content of the risk-on/risk-off flip of Section 1.3.

The original wrote this state as σ, a thermometer proxied by the VIX. We replace it with a systemic instability index constructed to lie in $[0, 1]$; the construction is Chapter 5, and the reason is Section 2.6.4.

2.2.3 Pillar 3 — the convexity exponent

The third pillar is that the bracket carries an exponent, κ_{sys}, and not a coefficient. The exponent measures the acceleration or sensitivity of DXY to any imbalance between supply and demand: when stress rises, κ_{sys} rises, so the *same* imbalance produces a larger move. A crisis does not merely add to dollar demand; it changes how violently a given excess demand is transmitted into price.

Why an exponent rather than a multiplier? Because in logs an exponent multiplies the entire weighted sum, $\kappa_{\text{sys}}\Big[\sum_i \alpha_{\text{i}}(\sigma_{\text{sys}}) \ln D_i - \sum_j \beta_{\text{j}}(\sigma_{\text{sys}}) \ln S_j\Big]$, and so rescales every channel at once. A multiplier on the ratio would shift the level and leave the transmission rate untouched. The claim being made is about transmission, so it belongs in the exponent. Chapter 7 gives κ_{sys} its explicit functional form and restores the dependence on σ_{sys} and R_{US} that the original dropped between its first statement of κ and its later one (ISSUES #3).

2.2.4 Pillar 4 — the pure safe-haven term

The fourth pillar is the exponential term, which represents pure flight-to-safety flow: capital seeking safety independently of the basic supply and demand mechanics. It is additive in

logs and therefore separable from the ratio — a deliberate choice, because the claim is that some crisis dollar demand is not a response to any relative price at all.

The $\lambda_2\sigma_{\text{sys}}^2$ term is what makes the response convex: as stress rises the dollar receives a quadratic boost, not a proportional one. The whole term is scaled by Φ_{sys}, systemic shelter credibility, so that the convexity is conditional on the dollar still being believed in. A system under stress whose shelter credibility has collapsed gets no safe-haven bid at all, however large σ_{sys} becomes. That conditionality is the model's answer to the question of what a dollar crisis, as opposed to a crisis in which the dollar rises, would look like.

2.3 Foundational Meta-Variables

src p.5

Three foundational modulators sit behind the equation: geopolitical power G_{pol}, internal U.S. risk R_{US}, and human capital H_{cap}. They are not additional demand or supply channels. They shape systemic credibility, and they act by modulating the functions that appear in Equation (2.1).

Human capital and intellectual property (H_{cap}). The accumulated stock of talent and intellectual property acts as a long-term value magnet, generating persistent capital demand by continuously creating intellectual and technological assets that the world invests in. It reaches the equation through the numerator. The original wrote the identity $D_{\text{hcap}} \equiv H_{\text{cap}}$; we separate the stock from the flow it generates,

$$D_{\text{hcap}} \;=\; h(H_{\text{cap}}), \qquad h' > 0, \tag{2.8}$$

because the erosion mechanism the original describes acts on the *stock* and reaches the flow with a lag, and because the two are measured by different data (ISSUES #24).

Geopolitical power (G_{pol}) and internal U.S. risk (R_{US}). These are second-order variables shaping systemic credibility rather than direct supply or demand. They modulate three functions:

1. $s_{\text{US,sys}}$, the share of safe-haven flows entering the dollar rather than gold, the Swiss franc or the yen;
2. Φ_{sys}, the intensity of the pure safe-haven effect in the exponential term;
3. κ_{sys}, total convexity — a rise in R_{US} *flattens* the convexity, since investors hesitate to seek refuge in an unsafe haven.

The third channel is worth dwelling on, because it is the model's sharpest claim. It says that internal political risk does not simply reduce dollar demand; it degrades the dollar's capacity to *act as a shock absorber at all*. The dollar can remain the world's safe haven and simultaneously become a worse one, and the two effects appear in different places in the equation. Chapter 6 specifies Φ_{sys} and $s_{\text{US,sys}}$; Chapter 7 specifies the κ_{sys} channel.

2.4 The Layer Map

src p.11

The variables above move on timescales that differ by three orders of magnitude, from intraday liquidity friction to decadal reserve-currency drift. Mixing them is the single most common source of confusion in the dollar debate, and the model's organising device is to refuse to mix them.

We therefore adopt six strata, Layer 0 to Layer 5. This is the numbering of the original's final synthesis; an earlier passage announces a "four-layer regulatory hierarchy" and one case study calls the mechanical block "Layer 4". Those are numbering slips, and the six-stratum scheme is canonical throughout this edition (ISSUES #15). Policy reaction, which one table in the original lists as a layer, is a feedback rather than a stratum, and it is treated in Chapter 18–Chapter 20.

Table 2.1: The six strata. Each layer modulates the layers below it; none adds to them.

Layer	Name	Variables	Enters Equation (2.1) as
0	Mechanical	D^*/S^*	the bracket
1	State and friction	σ_{sys}	every elasticity, and the exponential argument
2	Structural and trust	Φ_{sys}, $s_{\text{US,sys}}$	the exponential scale; the D_{prec} base
3	Reflexive	κ_{sys}	the exponent on the bracket
4	Psycho-technical	H_{soc}, I_{imp}	multiplies into κ_{sys}
5	Secular	$\Psi_{\text{str}}(t)$	a multiplicative prefactor (Chapter 10)

Table 2.1 and Figure 2.1 state the same thing twice, once as a table and once as a picture, because the layer map is the object a reader must carry through the next nine chapters. The essential property is that the layers form a modulation cascade rather than a sum: Layer 4 does not add dollar demand, it changes how violently Layer 3 transmits Layer 0.

Note that $\Psi_{\text{str}}(t)$ does not appear in Equation (2.1). The foundational equation is the model without its secular layer; the fully assembled form, with $\Psi_{\text{str}}(t)$ multiplying everything, is E9 in Chapter 10. We keep the two separate because most of the book — every case study in Part II bar two — works at horizons over which $\Psi_{\text{str}}(t)$ is constant, and carrying it through those derivations would add a symbol that never moves.

2.5 Derivation Block: The Foundational Equation

Following the template of STYLE §2, we set out Equation (2.1) formally.

2.5.1 Assumptions

Assumption 2.1 (Index numbers)**.** Every D_i and S_j is a positive index number, equal to 1 at a base period t_{base}, except for the emergent channels of Equation (2.7), which take the form $1 + X/\bar{X}$ and are likewise 1 when the channel is absent.

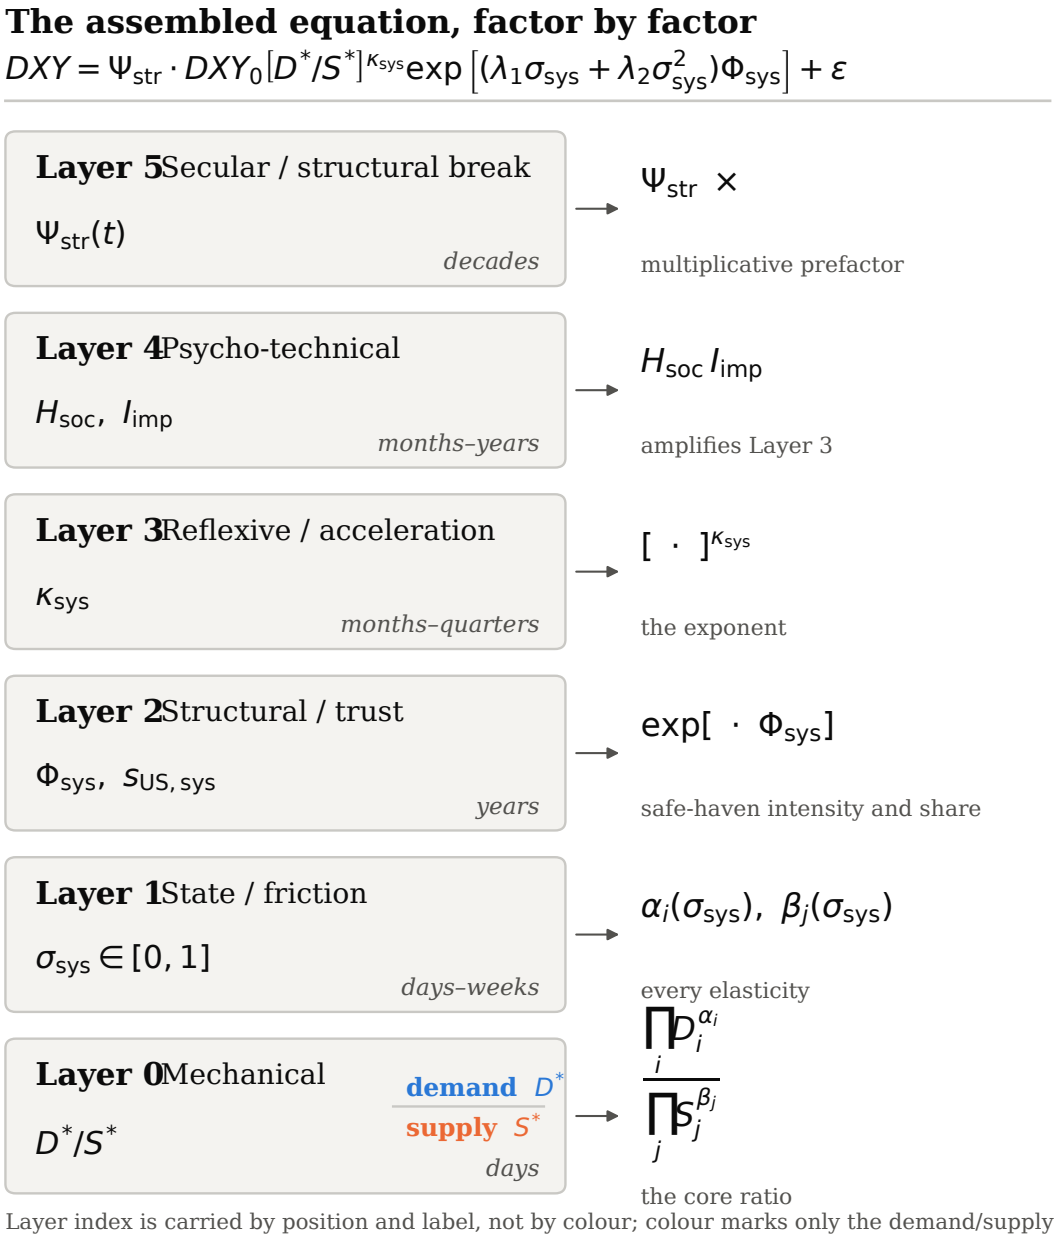


Figure 2.1: The six-layer architecture, with the factor of the integrated equation that each stratum supplies. Layer index is carried by vertical position and by the label; colour distinguishes only the demand and supply halves of Layer 0. Source: author's schematic; script `f01_architecture.py`.

Assumption 2.2 (Bounded state)**.** $\sigma_{\text{sys}} \in [0,1]$ by construction, and every latent modulator — G_{pol}, R_{US}, M_{trust}, T_{frag}, H_{soc}, I_{imp}, C_{alt} — likewise lies in $[0,1]$.

Assumption 2.3 (Constant elasticity within a regime)**.** For fixed σ_{sys}, the elasticities $\alpha_{\text{i}}(\sigma_{\text{sys}})$ and $\beta_{\text{j}}(\sigma_{\text{sys}})$ do not depend on the level of any D_i or S_j. All state dependence runs through σ_{sys} alone.

Assumption 2.4 (Separability of the safe-haven term)**.** Pure flight-to-safety flow enters additively in logs and is therefore separable from the supply–demand ratio.

Assumption 2.5 (Positivity)**.** $\kappa_{\text{sys}} > 0$ and $\Phi_{\text{sys}} > 0$ over the admissible parameter region.

2.5.2 Functional form

Three functional-form choices carry the specification, and each was made against a rejected alternative.

Cobb–Douglas for the ratio. Chosen for mutual amplification and for log-linearity, which is what makes the model estimable at all (Section 2.6.5). Rejected: an additive or CES aggregator. Additive fails the amplification requirement of Section 2.2.1; CES would introduce a substitution parameter per pair that no available data could identify.

Logistic interpolation for the elasticities. The seventeen elasticities are not free functions. Each interpolates between a calm value and a crisis value:

$$\alpha_{\text{i}}(\sigma_{\text{sys}}) = \alpha_i^{\text{calm}} + \left(\alpha_i^{\text{crisis}} - \alpha_i^{\text{calm}}\right) \Lambda(\sigma_{\text{sys}}), \qquad i \in \mathcal{I} \cup \{\text{prec}\}, \tag{2.9}$$

$$\beta_{\text{j}}(\sigma_{\text{sys}}) = \beta_j^{\text{calm}} + \left(\beta_j^{\text{crisis}} - \beta_j^{\text{calm}}\right) \Lambda(\sigma_{\text{sys}}), \qquad j \in \mathcal{J}, \tag{2.10}$$

$$\Lambda(\sigma_{\text{sys}}) = \left[1 + e^{-\theta(\sigma_{\text{sys}} - \sigma^*)}\right]^{-1}, \tag{2.11}$$

with the steepness θ and the threshold σ^* *shared* across all seventeen. The parameter count is $17 \times 2 + 2 = 36$. Rejected: leaving the elasticities as unspecified monotone functions, which is more honest but leaves Chapter 10 with nothing to simulate; and a linear form $\alpha_i(\sigma_{\text{sys}}) = a_{i0} + a_{i1}\sigma_{\text{sys}}$, which is cheaper by two parameters but discards the regime interpretation the whole book rests on. The logistic transition is standard in the smooth-transition literature [73], and sharing θ and σ^* embodies a substantive claim: regimes are a property of the *system*, not of each channel separately. Every channel switches at the same time, by different amounts.

Exponential for the safe-haven term. Chosen because flight to quality is convex in stress and because an exponential is the only form that is both convex and multiplicatively separable in logs. Rejected: a power function of σ_{sys}, which cannot represent the near-zero response at low stress that the calm-regime data show.

Table 2.2: Parameters of Equation (2.1). Calibration values are set in Chapter 10; none is estimated here.

Symbol	Meaning	Admissible range	Provenance
DXY_0	baseline level	> 0	$= DXY(t_{\text{base}})$ by Assumption 2.1
$\alpha_i^{\text{calm}}, \alpha_i^{\text{crisis}}$	calm and crisis demand elasticities	≥ 0	estimated (Chapter 11)
$\beta_j^{\text{calm}}, \beta_j^{\text{crisis}}$	calm and crisis supply elasticities	≥ 0	estimated
θ	shared regime steepness	> 0	estimated
σ^*	shared regime threshold	$(0, 1)$	estimated
$\alpha_{\text{prec}}(\sigma)$	safe-haven elasticity	≥ 0	estimated
λ_1	linear safe-haven coefficient	> 0	estimated
λ_2	quadratic safe-haven coefficient	> 0	estimated
δ_{off}	sanctions attenuation	> 0	estimated
$\bar{X}$	emergent-channel scale	> 0	stated convention
κ_{sys}	convexity exponent	> 0	Chapter 7
Φ_{sys}	shelter credibility	> 0	Chapter 6

2.5.3 Parameters

2.5.4 Comparative statics

Taking logs of Equation (2.1) and ignoring ϵ,

$$\begin{aligned} \ln DXY = c_0 + \kappa_{\text{sys}} \Big[& \alpha_{\text{prec}}(\sigma) \left(\ln s_{\text{US,sys}} + \ln D_{\text{prec}} \right) \\ & + \sum_{i \in \mathcal{I}} \alpha_{\text{i}}(\sigma_{\text{sys}}) \ln D_i - \sum_{j \in \mathcal{J}} \beta_{\text{j}}(\sigma_{\text{sys}}) \ln S_j \Big] \\ & + \left(\lambda_1 \sigma_{\text{sys}} + \lambda_2 \sigma_{\text{sys}}^2 \right) \Phi_{\text{sys}}, \end{aligned} \tag{2.12}$$

from which the elasticities the case studies use follow directly:

$$\frac{\partial \ln DXY}{\partial \ln D_i} = \kappa_{\text{sys}}\, \alpha_{\text{i}}(\sigma_{\text{sys}}), \qquad \frac{\partial \ln DXY}{\partial \ln S_j} = -\kappa_{\text{sys}}\, \beta_{\text{j}}(\sigma_{\text{sys}}), \tag{2.13}$$

$$\frac{\partial \ln DXY}{\partial \sigma_{\text{sys}}} = (\lambda_1 + 2\lambda_2 \sigma_{\text{sys}})\, \Phi_{\text{sys}} + (\text{regime terms}), \qquad \frac{\partial \ln DXY}{\partial \Phi_{\text{sys}}} = \lambda_1 \sigma_{\text{sys}} + \lambda_2 \sigma_{\text{sys}}^2. \tag{2.14}$$

Equation (2.13) is the workhorse of Part II: every case study reads an episode by asking which $\ln D_i$ and $\ln S_j$ moved, and multiplying by κ_{sys} times the relevant elasticity. The "regime terms" in Equation (2.14) are the derivatives of the elasticities themselves with

respect to σ_{sys}, which are what make the total derivative able to change sign; they are derived in full in Chapter D and used in Chapter 10.

2.5.5 Failure modes

Three ways this specification breaks, stated here so that later chapters can point at them.

1. **Unbounded stress.** If σ_{sys} is left on a raw scale, the term $e^{\lambda_2 \sigma_{\text{sys}}^2}$ explodes: with VIX at 80, a λ_2 large enough to matter at VIX 15 produces an astronomical number. Assumption 2.2 is what prevents this (ISSUES #9), and it is an assumption, not a theorem.
2. **A channel at zero.** A multiplicative aggregate is destroyed by any channel that reaches zero. Assumption 2.1 and Equation (2.7) keep every factor strictly positive, but a channel that genuinely collapses to nothing — a payment system switched off entirely — is outside the model as written.
3. **Constant elasticity.** Assumption 2.3 forbids the elasticity of one channel from depending on the level of another. There is at least one mechanism in the original that violates this: the amplification of D_{prec} by a S_{coll} shortage, which is stated as a level interaction. We treat it through κ_{sys}, which scales all channels together, and flag the residual mismatch in Chapter 10.

2.6 Formal Properties

2.6.1 Homogeneity and why the ratio is scale-free

The bracket in Equation (2.1) is homogeneous of degree $\kappa_{\text{sys}}\Big(\sum_i \alpha_{\text{i}}(\sigma_{\text{sys}}) + \alpha_{\text{prec}}(\sigma) - \sum_j \beta_{\text{j}}(\sigma_{\text{sys}})\Big)$ in a common rescaling of all demand and supply channels. Nothing in the model requires that degree to be zero, but something does require the units to cancel, and Assumption 2.1 is what delivers it: because every D_i and S_j is an index number, the bracket is a pure number and DXY/DXY_0 is dimensionless. Raising a quantity with units to a fractional power would otherwise be meaningless, and the exponents here are fractional by construction.

2.6.2 Identification of the baseline

The assembled model of Chapter 10 multiplies everything by the latent, slow-moving $\Psi_{\text{str}}(t)$. Without a normalisation, $\Psi_{\text{str}}(t)$ and DXY_0 are not separately identified: any constant can be shifted from one to the other. We therefore impose

$$\begin{aligned} &\Psi_{\text{str}}(t)(t_{\text{base}}) = 1, \qquad D_i(t_{\text{base}}) = S_j(t_{\text{base}}) = 1 \;\; \forall\, i, j \\ &\quad\Longrightarrow \qquad DXY_0 = DXY(t_{\text{base}}), \qquad c_0 = \ln DXY_0. \end{aligned} \tag{2.15}$$

Under Equation (2.15), $\Psi_{\text{str}}(t)(t)$ reads directly as the fraction of structural integrity remaining relative to the base period, which is the interpretation Chapter 9 needs (ISSUES #12). We take $t_{\text{base}} = $ 2006-01 provisionally; whether every series required actually exists at that date is a question about data, not about theory, and Chapter 11 settles it.

2.6.3 Why the exponent and the coefficient are different objects

It is worth separating two quantities that a reader may conflate. The elasticity $\alpha_{\mathrm{i}}(\sigma_{\mathrm{sys}})$ says how much channel i matters relative to the others. The exponent κ_{sys} says how much the *whole bracket* matters relative to the safe-haven term. Doubling κ_{sys} doubles every channel's effect and leaves their relative weights untouched; doubling one $\alpha_{\mathrm{i}}(\sigma_{\mathrm{sys}})$ changes the mix without changing the aggregate transmission rate. The two are separately identified only because κ_{sys} is common to all channels while the $\alpha_{\mathrm{i}}(\sigma_{\mathrm{sys}})$ are not — which is exactly the restriction that makes Equation (2.12) estimable, and which Chapter 11 leans on heavily.

2.6.4 Bounds on the safe-haven term

Under Assumption 2.2 and Assumption 2.5,

$$1 \;\le\; \exp\Big[(\lambda_1\sigma_{\mathrm{sys}} + \lambda_2\sigma_{\mathrm{sys}}^2)\Phi_{\mathrm{sys}}\Big] \;\le\; e^{(\lambda_1+\lambda_2)\Phi_{\max}}, \tag{2.16}$$

so the safe-haven multiplier is bounded above by a quantity that depends only on the parameters and on the maximum credibility, never on the realised level of stress. The term is increasing and convex in σ_{sys} throughout, which is the property the model wants, and it is finite, which the original's version was not.

2.6.5 The estimable form

Equation (2.12) is linear in the logs of the observed channels, conditional on the latent quantities σ_{sys}, Φ_{sys}, $s_{\mathrm{US,sys}}$ and κ_{sys}. That conditional linearity is the entire reason for the Cobb–Douglas pillar: it converts a model with seventeen state-dependent elasticities into something a state-space or regime-switching estimator can attack [47]. It does not make the model easy. The latent variables are unobserved, several have no free proxy, and κ_{sys} multiplies a sum of terms that are themselves functions of σ_{sys}. Chapter 11 is an honest account of how far this can be pushed with the data that exist.

2.7 Notation Conventions

Six conventions hold throughout the book. They are stated here in full; the front matter gives the short version, and Chapter A is the complete symbol table generated from the notation registry.

1. **Named subscripts, never numbered ones.** Demand channels are D with a named subscript and supply channels are S, indexed by the sets Equation (2.4) and Equation (2.5). The original's $\alpha_1 \dots \alpha_9, \alpha_H$ are replaced by $\alpha_{\mathrm{trans}}(\sigma_{\mathrm{sys}})$ and its siblings. Numbered elasticities invite exactly the double-counting error that Section 2.1 had to repair.
2. **Demand up means DXY up; supply up means DXY down.** There are no exceptions. The sanctions channel, which raises DXY when it rises, is expressed as an attenuation of supply in Equation (2.6) precisely so that the convention holds without exception.

3. **Every latent modulator lies in** $[0, 1]$**.** A coefficient on a latent variable therefore reads as the effect of a move across the full range, and coefficients are comparable across layers. Where the original quotes a variable on a 0–100 scale, that is a presentation scale and the model consumes the variable divided by 100 (ISSUES #26).
4. **Quantities are index numbers, normalised at** t_{base}, per Equation (2.15), with emergent channels entering through Equation (2.7).
5. $\Psi_{\text{str}}(t)$ **is an integrity coefficient.** It equals 1 when the dollar's structural position is intact and falls as that position erodes. Erosion is always $\Psi_{\text{str}}(t) \downarrow$. The original occasionally writes "$\Psi \uparrow$" to mean erosion; those two passages are corrected in place (ISSUES #2).
6. **Greek letters are parameters; Latin letters are variables.** The exceptions are σ_{sys}, κ_{sys}, Φ_{sys} and $\Psi_{\text{str}}(t)$, which are states inherited from the original's notation and kept for continuity with it.

2.7.1 The arrow scale

Tables throughout Part II use arrows for direction and magnitude: $\uparrow$ or $\downarrow$ for a small move, within the variable's normal range; $\uparrow\uparrow$ or $\downarrow\downarrow$ for a large move, of the size seen in a stress episode; $\uparrow\uparrow\uparrow$ or $\downarrow\downarrow\downarrow$ for an extreme move, the largest observed or a regime change; and $\leftrightarrow$ for no material change. Where data exist, tables carry both the arrow and the number, and the arrow is a summary of the number rather than a substitute for it.

2.8 The Dependency Graph

Figure 2.2 renders the whole variable system as a directed graph. It is generated from the notation registry rather than drawn by hand: an edge exists wherever the registry records that a variable enters an equation, joined to the quantity that equation defines. The figure therefore cannot drift away from the specification, and it changes whenever the registry does.

Three features of the graph are worth reading off it directly.

First, the model is *shallow but wide*. Almost every economically substantive variable is a source, feeding a small number of composites which then feed DXY. There are only five levels of dependency in a system of sixty-odd variables. That is a design property, not an accident: the layers were constructed to be thin so that each could be measured separately.

Second, the modulators converge. H_{soc} and I_{imp} feed κ_{sys}; G_{pol}, R_{US}, M_{trust} and T_{frag} feed Φ_{sys} and $s_{\text{US,sys}}$; C_{alt} and T_{frag} feed $\Psi_{\text{str}}(t)$. Nothing in Layers 2 to 5 touches the demand and supply channels directly. The entire influence of politics, trust, psychology and secular erosion on the dollar runs through four composite functions, and that is what makes the model tractable.

Third, T_{frag} is the only variable that enters two different layers, feeding both Φ_{sys} in Layer 2 and $\Psi_{\text{str}}(t)$ in Layer 5. It is therefore the natural place for the short-run and the long-run consequences of geoeconomic coercion to be told apart, which is what Chapter 14 and Chapter 16 do.

With the architecture fixed, Chapter 3 and Chapter 4 open the bracket and specify its seventeen channels one at a time.

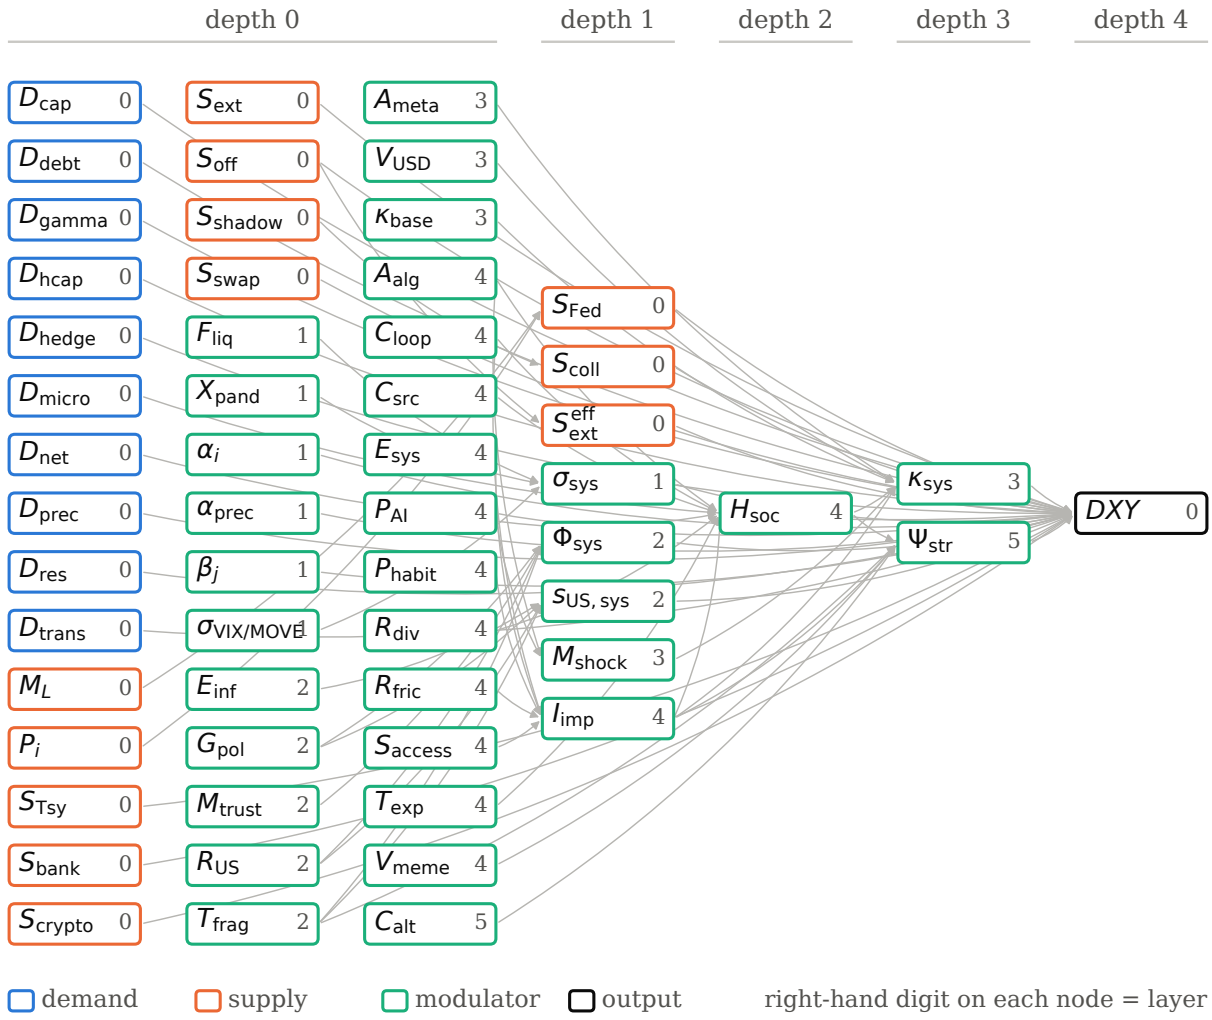


Figure 2.2: Variable dependency graph, generated from `registry/variables.csv`. Columns are dependency depth: sources on the left, DXY on the right. Colour carries role — demand, supply, modulator — and the digit on each node is its layer, since six ordered levels cannot be separated by hue in print. Coefficients, conventions and the shorthands introduced for the fixed-point proof of Chapter 8 are omitted. Source: author's calculation; script `f02_dependency.py`.

Chapter 3

Decomposing the Numerator: Global Demand Dynamics (D^*)

Our equation defines DXY as a ratio between aggregate demand and aggregate supply. The numerator D^* integrates ten distinct demand drivers. This chapter defines each component, gives its rationale, and proposes empirical measurement proxies. Collectively they represent the dollar's suction engines — from mandatory trade settlement to speculative capital movement and hidden structural effects.

Each driver is set out as a variable card following the template of STYLE §1: definition, mechanism, role in the equation, interactions, measurement, historical range, and limitations. The cards are deliberately uneven in what they can deliver. Some channels have a free daily series that measures very nearly the right object; others have no public series at all, and for those the card states what would be needed rather than substituting a convenient number. Table 3.2 records that status honestly at the outset.

A note on the dependent variable

Before any card quotes a number, one substitution must be stated plainly, because it affects every empirical claim in this chapter and in Chapter 11.

The ICE U.S. Dollar Index — the DXY of the original paper's title — is a licensed product with no free, citable source. **The dependent variable used throughout this edition is instead the Federal Reserve's Nominal Broad U.S. Dollar Index for Advanced Foreign Economies, FRED series `DTWEXAFEGS`**, daily from 2006-01-02, with `DTWEXM` (Major Currencies, 1973–2019, discontinued) for the historical chapters.

These are different indices, not different names for one index:

Three consequences follow, and the reader should hold them throughout.

1. **No result stated in index points transfers between the two.** All empirical work here is in logs and differences for that reason.
2. **`DTWEXAFEGS` is far less euro-dominated.** Episodes driven specifically by EUR/USD will register as smaller moves than a DXY-based narrative implies.
3. **`DTWEXAFEGS` excludes emerging markets by construction.** This matters most for D_{debt} below, whose squeeze mechanism operates hardest on EM borrowers, so a test of

Table 3.1: The dependent variable is a substitute, and the substitution is not neutral. Full provenance in `data/SOURCES.md`.

	ICE DXY (original)	FRED `DTWEXAFEGS` (used here)
Basket	six currencies; EUR 57.6%, JPY 13.6%, GBP 11.9%, CAD 9.1%, SEK 4.2%, CHF 3.6%	about 26 advanced-economy currencies
Weights	fixed since 1999	re-weighted annually on bilateral trade shares
Emerging markets	excluded	excluded (the broad index `DTWEXBGS` includes them)
Coverage	1973–	2006-01-02–

that channel against the advanced-economy index alone is a *weak* test. The broad index `DTWEXBGS`, which includes emerging markets, is also in hand, and Chapter 11 uses both: the advanced-economy index as the headline dependent variable for comparability with the original's DXY, and the broad index wherever a channel runs through emerging-market borrowers.

The two Federal Reserve indices are not close substitutes for each other either, and the difference is itself informative. Over 2006-01-02 to 2026-08-28 the advanced-economy index peaked at 124.85 on 2022-09-26, while the broad index peaked later and higher, at 130.04 on 2025-01-13. A dollar cycle that had turned in the advanced economies was still running against emerging-market currencies more than two years later. Any claim in this book about "the dollar peaking" must therefore name its index, and the case studies in Chapter 13 and Chapter 16 do.

3.1 Transactional Demand (D_{trans})

src p.6

Definition 3.1. D_{trans} is the flow of dollar purchases required to settle internationally traded goods that are invoiced in dollars, principally commodities. It is an index number, equal to 1 at t_{base}, constructed as the product of a dollar commodity price index and a measure of world trade volume.

Mechanism. This is the foundation of the petrodollar system. Nearly all globally essential commodities — oil, natural gas, grains, industrial metals — are priced and settled in dollars. Regardless of the condition of the U.S. economy, importing nations must continuously purchase dollars in open markets to pay for these goods. A German refiner buying Brent, a Japanese utility buying LNG and a Chinese mill buying iron ore are all, in the first instance, buyers of dollars. This is a hard, inelastic source of demand: it does not respond to the relative attractiveness of U.S. assets, and it does not stop when the dollar is expensive. It stops only when the trade itself stops.

The invoicing convention is what does the work, and it is a genuine empirical regularity rather than an assumption of the model: prices in international trade are sticky in dollars

Table 3.2: Data status of the ten demand cards. A card whose proxy has no free series states what is required and carries no historical range and no figure; no number in this chapter is estimated, interpolated or assumed.

Card	Primary proxy	Status	Gap
D_{trans}	Brent × world trade volume	**price leg**	CPB World Trade Monitor (volume leg)
D_{debt}	BIS USD credit to non-banks outside the U.S.	**complete**	—
D_{cap}	10y U.S.–Germany spread	**complete**	growth differential (secondary)
D_{prec}	change in MMF assets (`WRMFSL`, `MMMFFAQ027S`)	**usable**	no weekly institutional series; VIX disqualified by ISSUES #31
D_{res}	IMF COFER USD share	**complete**	—
D_{hedge}	3m cross-currency basis	**substitute**	basis is Tier 3; TED and SOFR–IORB stand in
D_{gamma}	dealer net gamma	**none**	no adequate free substitute exists
D_{net}	USD stablecoin capitalisation	**no data**	DefiLlama stablecoin series
D_{micro}	calendar dummy (E2)	**constructed**	—
D_{hcap}	IP receipts; STEM inflows	**no data**	BEA ITA; NSF/NCSES

rather than in the currency of exporter or importer [41]. A dollar appreciation is therefore not a symmetric terms-of-trade shift between two countries but a contraction in the purchasing power of everyone who earns in something else.

Role in the equation. D_{trans} enters the numerator of Equation (2.2) with elasticity $\alpha_{\text{trans}}(\sigma_{\text{sys}})$. The sign convention is the standard one: $D_{\text{trans}} \uparrow \Rightarrow DXY \uparrow$.

Its regime behaviour is the most distinctive thing about it. Because the demand is contractual rather than discretionary, we expect $\alpha_{\text{trans}}(\sigma_{\text{sys}})$ to be the *least* state-dependent of the ten, with α_i^{calm} and α_i^{crisis} close together. That is a falsifiable claim, and it is the reason this card is worth separating from D_{cap}: if estimation finds a large calm–crisis gap in $\alpha_{\text{trans}}(\sigma_{\text{sys}})$, the interpretation of the channel as contractual is wrong.

Interactions. D_{trans} is secured by G_{pol}: freedom of navigation, the security of sea lanes and the enforcement of the dollar invoicing convention are all projections of geopolitical power, so $G_{\text{pol}} \downarrow \Rightarrow D_{\text{trans}} \downarrow$ over long horizons (Chapter 6). It is attacked by T_{frag}: bilateral non-dollar settlement of oil and gas subtracts from it directly, which is the mechanism Chapter 9 treats as an input to secular erosion. It is the channel through which the dollar transmits into U.S. inflation, which Chapter 18 uses: $D_{\text{trans}} \downarrow$ eases imported inflation.

Table 3.3: Measurement protocol for D_{trans}.

	Primary	Secondary
Proxy	Brent crude price × world trade volume	Bloomberg Commodity Index (BCOM) × world trade volume
Source	FRED `DCOILBRENTEU`; CPB World Trade Monitor	Bloomberg (licensed); CPB
Frequency	daily (price), monthly (volume)	daily, monthly
Unit	USD per barrel; volume index	index
Transformation	log, indexed to 1 at t_{base}	log, indexed to 1
Availability	free	BCOM paywalled
Weakness	Brent is one commodity standing for all dollar-invoiced trade	BCOM is broader but not free

Measurement. The deeper measurement problem is that the object we want is a *settlement flow* and the proxy is a *price times a volume*. These coincide only if the dollar share of invoicing is constant, which is precisely what T_{frag} is supposed to erode. The proxy is therefore least reliable exactly when the model most needs it, and Chapter 11 treats the invoicing share as a separate quantity to be tracked rather than assumed.

Historical range. The price leg is in hand; the volume leg is not, so what follows is a range for Brent alone and not for the composite.

Not supplied, and the gap is ours. The CPB World Trade Monitor publishes world trade volume monthly and free. It is not in `data/raw` and nothing prevents its being there. *What it costs:* the D_{trans} composite cannot be built, so every figure and every case-study reading of this channel rests on the price leg alone — Brent — which is why Section 14.2.4 can find the oil price falling across an embargo and cannot say whether the quantity of dollar-invoiced trade fell with it.

Brent runs daily from 2006-01-03 to 2026-09-01, 5,226 observations. It reached \$143.95 per barrel on 2008-07-03 and fell to \$9.12 on 2020-04-21 — a range of more than fifteen to one within a single sample, and the widest of any series used in this chapter. The latest observation is \$96.02.

The 2020 low is worth a sentence, because it is the clearest illustration of why the price leg alone is not the channel. On 2020-04-21 the dollar price of oil collapsed, which mechanically reduces the dollar value of oil settlement; but world trade volume was collapsing simultaneously, and the two effects compound rather than offset. A composite built from price alone would have understated the fall in transactional dollar demand at exactly the moment the model most needs it measured.



Figure 3.1: D_{trans}, price leg only: Brent crude, daily, 2006-01-03 to 2026-09-01. The volume leg of the composite is not yet in hand, so this is an input to the proxy rather than the proxy itself. Source: FRED `DCOILBRENTEU`; author's calculations; script `f03_dtrans.py`.

Limitations. Three. The composite is a construction, not an observable, and its two legs move on different frequencies. The dollar invoicing share, which the construction holds fixed, is itself a slow-moving variable that the model treats as eroding. And the channel's inelasticity — the property that makes it theoretically distinctive — is exactly the property

that makes it hard to identify, since a variable that does not respond to prices contributes little independent variation to an estimate.

3.2 Debt-Servicing Demand (D_{debt})

src p.6

Definition 3.2. D_{debt} is the flow of dollar purchases required by non-U.S. borrowers to service interest and principal on dollar-denominated obligations contracted outside the United States. Its stock counterpart is total dollar credit to non-bank borrowers outside the U.S.; the card indexes that stock to 1 at t_{base}.

Mechanism. For decades governments and corporations outside the United States, especially in emerging markets, have borrowed dollars through the Eurodollar market. The geography of that funding — which banks raise dollars where, and against what — is mapped by Aldasoro and Ehlers [4]. These entities must buy dollars not to invest but to service existing obligations. The distinguishing feature of the channel is its sign under stress: when DXY rises, the local-currency cost of servicing a dollar liability rises with it, the burden intensifies, and borrowers are forced into further dollar purchases. A rising dollar creates its own additional demand. This is the classical dollar short squeeze, and it is the single most important reflexive loop on the demand side of the model.

The mechanism is not hypothetical. A stronger dollar goes together with wider deviations from covered interest parity *and* with contracting cross-border dollar lending [6], and the risk-taking channel that links global bank leverage to the dollar is well documented [20]. The stock on which the mechanism operates is measured in the trillions [56].

Role in the equation. D_{debt} enters the numerator with elasticity $\alpha_{\text{debt}}(\sigma_{\text{sys}})$, and $D_{\text{debt}} \uparrow \Rightarrow DXY \uparrow$. Because the forced-purchase mechanism activates precisely when funding conditions tighten, we expect $\alpha_i^{\text{crisis}} > \alpha_i^{\text{calm}}$ for this channel — a large calm-to-crisis gap, in contrast to D_{trans}. The channel is also the clearest case for the convexity exponent of Chapter 7: the squeeze is self-reinforcing, so its contribution to DXY is amplified by κ_{sys} rather than merely added.

Interactions. D_{debt} is amplified by $S_{\text{coll}} \downarrow$ and $S_{\text{bank}} \downarrow$: when collateral and dealer balance sheets are scarce, rolling a dollar liability becomes harder and the forced purchase becomes larger. It is relieved by $S_{\text{swap}} \uparrow$, which is what the swap lines exist to do [8]. It shares its underlying population with D_{hedge} (Chapter 3, S4), and the two must not be double-counted: D_{debt} is the principal-and-interest flow, D_{hedge} the cost of the hedge against it.

Measurement. The stock-for-flow substitution deserves a sentence. What forces dollar purchases is the servicing schedule, which depends on the stock, its maturity profile and its coupon — not on the stock alone. Using the stock assumes the servicing flow is roughly proportional to it. That is defensible over years and poor over quarters, and it is one reason the quarterly frequency of this series limits what Chapter 11 can ask of it.

Table 3.4: Measurement protocol for D_{debt}.

	Primary	Secondary
Proxy	USD credit to non-bank borrowers outside the U.S., level	the same series, year-on-year change
Source	BIS Global Liquidity Indicators, `Q.USD.3P.N.A.I.B.USD`	BIS, `...B.771`
Frequency	quarterly	quarterly
Unit	USD millions	per cent
Transformation	log, indexed to 1 at t_{base}	level (already a rate)
Availability	free	free
Weakness	a stock proxying a flow; quarterly and published with a lag	better timed, but differences discard the level information the squeeze depends on

Historical range. The BIS series runs from 2000-Q1 to 2026-Q1, 105 quarterly observations. Dollar credit to non-bank borrowers outside the United States stood at $2.33 trillion in 2000-Q1 and $14.74 trillion in 2026-Q1. At the model's base period, 2006-Q1, it was $4.03 trillion, so the stock is 3.66 times its base-period level.

The growth rate is the more informative series for the mechanism. It peaked at +24.4 % year-on-year in 2007-Q2 — the pre-crisis build-up of the offshore dollar short — and reached its most negative reading, −3.9 %, in 2022-Q4. The 2009 trough was shallower, at −3.9 % in 2009-Q2 on the same measure.

That 2022 reading is worth pausing on, because it is not the episode the original paper would have predicted as the record. The deepest contraction in offshore dollar credit in this twenty-six-year sample is not the global financial crisis but the 2022 tightening cycle, and it coincides with the highest reading of the dependent variable used here: `DTWEXAFEGS` peaked at 124.85 on 2022-09-26. We record the co-occurrence and decline to call it a test. Two quarterly observations against one daily peak cannot identify a mechanism, and the index in question excludes exactly the emerging-market borrowers on whom the squeeze bites hardest. Chapter 13 treats the episode properly.

Limitations. The series is quarterly, which is coarse relative to the mechanism it measures; squeezes happen in days. It is a stock, not a servicing flow. It covers non-bank borrowers only, so intra-bank dollar funding pressure is invisible to it. And it is published with a lag of roughly one quarter, which means it cannot be used in any real-time reading of the model.

3.3 Capital-Asset Demand (D_{cap})

Definition 3.3. D_{cap} is the flow of dollar purchases made in order to acquire U.S. financial assets for return, as opposed to for settlement, safety or obligation. It is an index number equal to 1 at t_{base}, driven by the risk-adjusted return differential between U.S. assets and their nearest substitutes abroad.

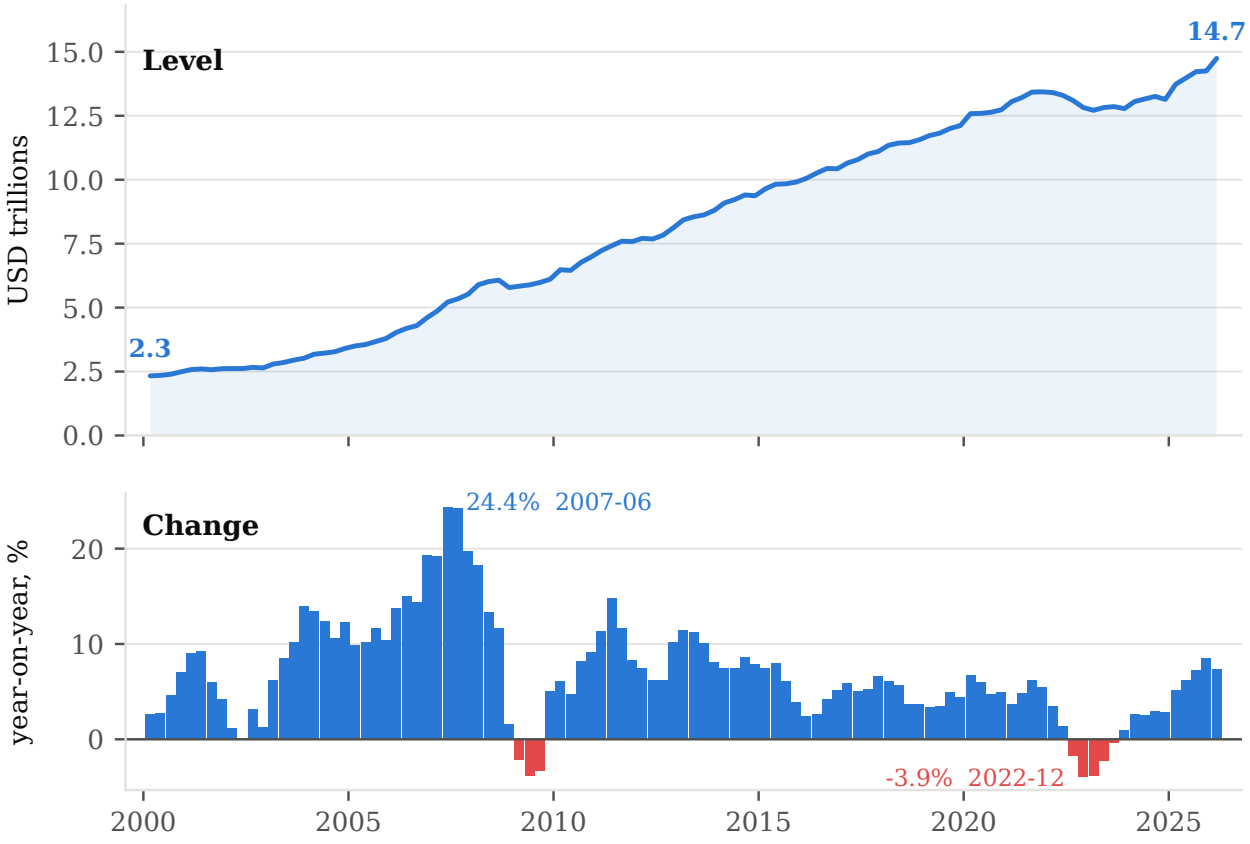


Figure 3.2: D_{debt}: dollar credit to non-bank borrowers outside the United States, level and year-on-year change, 2000-Q1 to 2026-Q1. The level is the size of the offshore dollar short; the change is when it was forced to shrink. Source: BIS Global Liquidity Indicators, series `Q.USD.3P.N.A.I.B.USD` and `...B.771`; author's calculations; script `f04_ddebt.py`.

Mechanism. This is the primary investment suction engine. Global capital flows toward assets offering the highest risk-adjusted return. When U.S. growth outpaces the rest of the world, funds flow into U.S. equities; when U.S. bond yields exceed those of Germany or Japan, capital flows into Treasuries. In every such case the investor must first buy dollars to gain the exposure, and it is that first step which the model counts.

Role in the equation. D_{cap} enters the numerator with elasticity $\alpha_{\text{cap}}(\sigma_{\text{sys}})$, and $D_{\text{cap}} \uparrow \Rightarrow DXY \uparrow$. Its regime signature is the opposite of D_{debt}'s: we expect $\alpha_i^{\text{calm}} > \alpha_i^{\text{crisis}}$, because return-seeking flows are exactly what is abandoned when stress rises. The pair $(D_{\text{cap}}, D_{\text{prec}})$ handing off to each other as σ_{sys} crosses σ^* is the mechanical content of the risk-on/risk-off flip of Section 1.3. If the model has a single testable prediction, it is that these two elasticities cross.

Interactions and the policy-divergence question. The original introduces, in its third case study, a separate policy-divergence term written D_{div}, together with a driver it calls policy divergence, meant to capture Federal Reserve policy diverging from the ECB and the Bank of Japan. **No such variable exists in this edition.** Policy divergence is precisely what D_{cap}'s proxy already measures — a rate and balance-sheet differential — and a separate term would be collinear with it by construction. The case study is reworded to speak of D_{cap}, whose divergence component is named in the measurement protocol below (ISSUES #6). The registry retains a row for D_{div} marked ABSORBED, for traceability only.

D_{cap} is fed by D_{hcap} over long horizons: the human-capital and IP stock is what generates the assets that attract the flow (Section 3.10, S4). It is degraded by R_{US}, which raises the risk in risk-adjusted return.

Table 3.5: Measurement protocol for D_{cap}.

	Primary	Secondary
Proxy	10-year U.S. minus 10-year German government yield spread	projected GDP growth differential, U.S. minus rest of world
Source	FRED `DGS10`; Bundesbank / ECB SDW	IMF *World Economic Outlook*; consensus forecasts
Frequency	daily	semi-annual (forecast vintages)
Unit	percentage points	percentage points
Transformation	level; the spread is already a differential	level
Availability	free	free
Weakness	two-country spread standing for a world differential; ignores the JPY leg	forecast revisions are infrequent and are themselves endogenous to the dollar

Measurement.

Historical range. The 10-year U.S.–German spread is available monthly from 2006-01 to 2026-06, 246 observations. The frequency is set by the German leg: FRED `IRLTLT01DEM156N` is monthly, so the daily U.S. leg is taken at month end and months present in only one series are dropped rather than filled.

The spread reached +2.75 percentage points in 2018-10, at the peak of the divergence between Federal Reserve tightening and ECB accommodation, and −0.80 in 2008-12. The latest reading is +1.47. The spread was negative in 20 of the 246 months, all of them falling between 2007-11 and 2011-12 — that is, the German yield advantage is not a recurring feature of the sample but a single episode spanning the financial crisis and the early euro-area crisis.

That concentration is a useful fact for the model rather than a curiosity. A channel whose sign reverses only once in twenty years contributes little independent identifying variation, so estimates of $\alpha_{\text{cap}}(\sigma_{\text{sys}})$ will be driven largely by the magnitude of a persistently positive spread rather than by its sign. Chapter 11 treats this as a limit on what the D_{cap} elasticity can be asked to reveal.



Figure 3.3: D_{cap}: 10-year U.S. minus German government bond yield, monthly, 2006-01 to 2026-06. Positive is a U.S. yield advantage. Source: FRED `DGS10` at month end and FRED `IRLTLT01DEM156N`; author's calculations; script `f05_dcap.py`.

Limitations. The proxy is a nominal yield spread, while the mechanism is about risk-adjusted *expected* returns; the gap between the two is precisely the currency risk premium the model is trying to explain, so there is a circularity here that Chapter 11 must confront rather than assume away. Hedged and unhedged flows respond to different objects — an unhedged Japanese buyer of Treasuries cares about the nominal spread, a hedged one about the spread net of the cross-currency basis, which is D_{hedge}.

3.4 Safe-Haven Demand (D_{prec})

Definition 3.4. D_{prec} is the flow of dollar purchases motivated by the preservation of capital rather than by return: the demand for dollar cash and Treasury bills that appears when global investors liquidate risk. It is an index number equal to 1 at t_{base}, activated by σ_{sys}.

Mechanism. This is the core of the system's non-linearity. In calm times the dollar may be sold. But when systemic stress surges, global investors liquidate everything — emerging-market equities, risky credit, even gold — and move into the only asset offering absolute liquidity and safety: dollar cash and Treasury bills. They are not seeking profit; they are seeking survival. The behaviour is not a preference over returns but a suspension of the calculation that produces preferences over returns, which is why it cannot be represented by moving along a demand curve.

The equilibrium reason the dollar is the recipient is that the reserve currency appreciates in global crises because the reserve country's intermediaries absorb more risk, which is what makes it a hedge in the first place [53]. The empirical counterpart is identified: global risk shocks appreciate the dollar, and in counterfactuals where it does not appreciate the transmission of the shock is far weaker [39].

Role in the equation. D_{prec} is the one demand channel that does not enter $\prod_{i\in\mathcal{I}}$. It carries its own factor,

$$\left(s_{\text{US,sys}}\, D_{\text{prec}}\right)^{\alpha_{\text{prec}}(\sigma)}, \tag{3.1}$$

where $s_{\text{US,sys}} \in [0,1]$ is the share of a global flight to quality that dollar assets actually capture. The base is USD-*captured* safe-haven demand: a panic that sends money to gold and the Swiss franc raises D_{prec} without raising the dollar, and Equation (3.1) is what represents that. As $s_{\text{US,sys}} \to 0$ the factor goes to zero, not to one.

Two things follow. First, D_{prec} appears exactly once in the numerator; the original counted it twice, once inside the product and once in its own factor (ISSUES #7). Second, the share sits inside the base rather than scaling the exponent, which is where the original placed it (ISSUES #8). Chapter 6 specifies $s_{\text{US,sys}}$; the interaction between the two is the subject of the case studies in Chapter 13.

Interactions. D_{prec} is activated by σ_{sys} and gated by $s_{\text{US,sys}}$, and so is the only demand channel whose effect on DXY can be switched off by a Layer-2 variable while the underlying panic continues. It is amplified by $S_{\text{coll}} \downarrow$: a flight to quality into an asset whose collateral market is impaired produces a larger price move than the flow alone implies. It competes with D_{cap} across the regime threshold, as Section 3.3 sets out.

Measurement.

Why VIX cannot serve as the D_{prec} proxy. The original proxies safe-haven demand by "the VIX or MOVE". We do not, and the reason is structural rather than a matter of taste.

Table 3.6: Measurement protocol for D_{prec}.

	Primary	Secondary and tertiary
Proxy	change in retail money-market fund assets	Treasury bill auction bid-to-cover; foreign official custody holdings at the New York Fed
Source	FRED `WRMFSL` (weekly, retail, to 2021-02) and `MMMFFAQ027S` (quarterly, total, full sample) — **both in hand**	Treasury auction results (TreasuryDirect); NY Fed H.4.1 custody holdings
Frequency	weekly; quarterly	per auction; weekly
Unit	USD billions	cover ratio; USD millions
Transformation	4-week change, then indexed	level; log, indexed to 1
Availability	free; `WRMFSL` **ends 2021-02-01**, `MMMFFAQ027S` live	free
Weakness	mixes safe-haven flight with ordinary cash management and with rate-driven substitution out of deposits	bid-to-cover is per-auction and noisy; custody holdings mix official reserve management with private flight, and overlap D_{res}
Not used	**VIX and MOVE.** Disqualified as proxies for this channel — see below.	

σ_{sys}, the state variable of Chapter 5, is itself built from VIX and MOVE. If D_{prec} is also VIX, then VIX enters the integrated equation three times over: in the base of $(s_{\text{US,sys}} D_{\text{prec}})^{\alpha_{\text{prec}}(\sigma)(\sigma_{\text{sys}})}$, in that same factor's exponent, and again in $\exp[(\lambda_1 \sigma_{\text{sys}} + \lambda_2 \sigma_{\text{sys}}^2)\Phi_{\text{sys}}]$. The consequence is not the ordinary nuisance of two correlated regressors. It is that the model's central decomposition — how much of a crisis dollar bid is *mechanical* safe-haven demand sitting in the numerator, and how much is the *pure* flight to safety sitting in the exponential — would have no empirical content whatever, because one series would drive both. $\alpha_{\text{prec}}(\sigma)$ would not be identified against λ_1 and λ_2.

It follows that **a case study which reports "$D_{\text{prec}} \uparrow\uparrow$ and $\sigma_{\text{sys}} \uparrow\uparrow$" on the evidence of VIX alone is reporting one fact twice, and is not evidence for the model.** This binds Chapter 12's method section and every case study in Chapter 13 through Chapter 16: an episode may evidence D_{prec} only from a flow series, never from a volatility index. Where no flow series covers the window, the case study says so and reports the two channels as jointly unidentified over that window rather than asserting both.

This is the resolution of ISSUES #31. VIX remains the base of σ_{sys} in Chapter 5, where it is uncontested and where it belongs.

Historical range. The primary flow proxy is now partly in hand. `WRMFSL`, retail money-market fund assets, runs weekly from 2006-01-02 to 2021-02-01 within the sample, 788 observations, from $860.4 billion at the start to $1,429.1 billion at the end, with a low of $845.5 billion on 2015-05-18 and a high of $1,518.5 billion on 2020-05-25.

The behaviour the channel predicts is visible in the changes rather than the level. Three of the five largest four-week increases in the sample fall inside a single four-week window — 2020-03-30, 2020-04-13 and 2020-04-20, at \$93.4, \$90.1 and \$85.8 billion — which is the pandemic flight to cash. That is the first direct evidence in this book for D_{prec} as a *flow* rather than as a volatility reading, and it is what the ISSUES #31 decision was for.

Two limitations bound `WRMFSL`. It is **retail only**, excluding the institutional funds where crisis-scale money actually moves; and it is **discontinued on 2021-02-01**, when the H.6 release moved to monthly frequency, so it stops before both the 2022 tightening and the 2025 tariff episode.

A second series covers what the first misses, at the cost of frequency. `MMMFFAQ027S` is *total* money-market fund financial assets — retail and institutional — quarterly from the Financial Accounts, and it spans the whole sample: 81 observations, \$2.03 trillion in 2006-Q1 rising to \$8.29 trillion in 2026-Q1. Its largest quarterly increases are 2020-Q1 at \$725 billion, 2023-Q1 at \$470 billion, 2025-Q4 at \$416 billion and 2008-Q4 at \$412 billion — the pandemic, the regional-bank episode, and the financial crisis, which is the ranking the channel predicts.

The two series are complements and the card uses both: `WRMFSL` where weekly resolution is needed and the window ends before 2021-02, `MMMFFAQ027S` for coverage and for anything involving institutional money. Neither alone is adequate, and the card does not yet claim a measured elasticity from either. The institutional gap is now partly closed, and the way it closes is itself evidence. `IMFSL`, institutional money-market fund assets, is available monthly from 2006-01 to 2021-01, 181 observations, rising from \$1,167.0 billion to \$2,853.8 billion and peaking at \$3,308.3 billion in 2020-05. Its three largest three-month increases are 2020-05 at \$1,026.3 billion, 2020-04 at \$797.6 billion and 2020-06 at \$737.8 billion. The pandemic dash for cash is a trillion dollars of institutional money in a quarter, and it is four times larger than anything else in the series. That is the D_{prec} mechanism seen as a flow, at institutional scale, exactly as ISSUES #31 requires.

The proxy fails at the crisis it should measure best. It must be said plainly that `IMFSL` shows *nothing* in the autumn of 2008. The series reads \$2,278.2 billion in 2008-06, \$2,277.1 billion in 2008-09 and \$2,273.6 billion in 2008-10 — flat through the failure of Lehman Brothers and through the week the Reserve Primary Fund broke the buck. It rises only afterwards, to \$2,438.5 billion by 2008-12.

This is not a defect in the data; it is the 2008 crisis. Institutional prime funds were themselves the object of the run, and precautionary money left them for Treasury bills and insured deposits rather than arriving in them. The proxy therefore has a **sign-inverting structural break** at the one episode in the sample where D_{prec} is least in doubt: in 2020 institutional funds were the destination of flight, in 2008 they were its source. A card that reported "$D_{\text{prec}} \uparrow\uparrow$" from `IMFSL` in 2008 would report a fall.

Two consequences bind the rest of the book. First, no single money-fund series measures D_{prec} across the whole sample, and Chapter 11 must not pool 2008 and 2020 in one regression on this proxy. Second, Chapter 13's global-financial-crisis case study must evidence D_{prec} from bill yields, the bill–OIS spread or custody holdings, not from money funds, and must say why.

What the three proxies jointly cover. `WRMFSL` gives weekly resolution for retail money to 2021-02; `IMFSL` gives institutional scale monthly to 2021-01, with the 2008 inversion above; `MMMFFAQ027S` gives quarterly coverage of the whole sample including 2021–2026. All three end or coarsen exactly where the model most needs frequency, and none is daily.

> **Not supplied, and the gap is ours.** Treasury bill bid-to-cover ratios (TreasuryDirect auction results, per auction) and New York Fed foreign official custody holdings (H.4.1, weekly) are both free and both unsupplied. *What it costs:* D_{prec} keeps a monthly proxy that *inverts in 2008* (Section 11.3.4), which is why $\alpha_{\text{prec}}(\sigma)$ is weakly identified and weakest in the episode the model most needs it for. Either series would restore weekly frequency after 2021-02 and neither inverts.

What can be shown is the *trigger*, which is measured, and Figure 3.4 shows it. VIX is available daily from 2006-01-03 to 2026-09-03, 5,232 observations. It peaked at 82.69 on 2020-03-16 — the Monday after the Federal Reserve's emergency Sunday rate cut — and reached its sample low of 9.14 on 2017-11-03. The latest observation is 14.32. This describes when the channel switches on. It does not measure how much money moved, and the distinction is the whole content of ISSUES #31.

The distribution matters more than the extremes for a card about regime switching. The index sits at or above 30 on 450 days, 8.6 per cent of the sample; at or above 40 on 179 days, 3.4 per cent; and at or above 50 on just 75 days, 1.4 per cent. Crisis is rare and brief. Any regime threshold σ^* that the model estimates will therefore be identified off a small number of episodes, and the crisis-regime elasticities α_i^{crisis} are estimated from roughly one day in twelve at best and one day in seventy at worst. That is not a reason to abandon the two-regime structure, but it is a reason to report confidence intervals on the crisis branch that are much wider than on the calm one, and Chapter 11 does.

Limitations. Beyond the identification problem above: the proxy is an option-implied second moment, and the object is a flow of funds. The two are correlated in crises and only loosely related outside them. Gold and the Swiss franc are the competing destinations, and neither appears in the proxy, so the split that $s_{\text{US,sys}}$ is supposed to govern is unobserved on the demand side.

3.5 Reserve-Management Demand (D_{res})

src p.7

Definition 3.5. D_{res} is the flow of dollar purchases and sales arising from official reserve management: the portfolio decisions of central banks over the currency composition of their foreign exchange reserves. It is an index number equal to 1 at t_{base}.

Mechanism. Central banks hold dollars as foreign exchange reserves. They use them to intervene in domestic FX markets and to finance their economies' transactional demand. Their portfolio decisions — whether to accumulate or to reallocate, for instance by selling dollars to buy gold — create very large flows on very long horizons. The channel differs from every other in this chapter in that its agents are not profit-maximising: a central bank buys dollars to stabilise its own currency, not because it expects the dollar to rise.

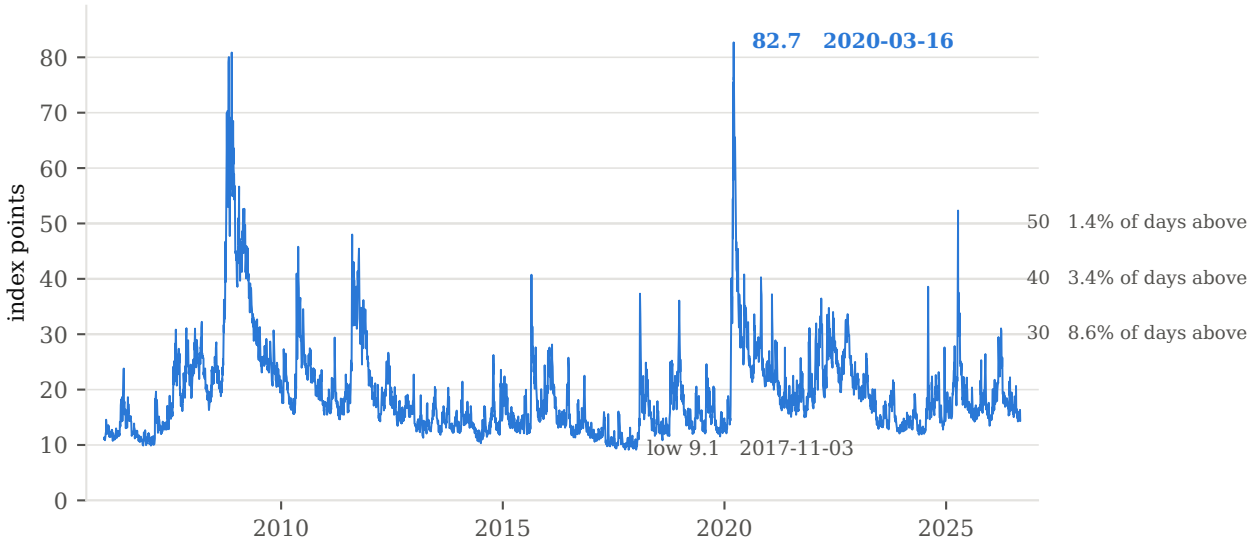


Figure 3.4: The *trigger* for D_{prec}, not the channel itself: the VIX, daily, 2006-01-03 to 2026-09-03, with the thresholds quoted in the text. This series is the base of σ_{sys} in Chapter 5 and is therefore disqualified as a proxy for D_{prec} (ISSUES #31); it is shown to date the activations, not to measure the flow. Source: FRED `VIXCLS`; author's calculations; script `f06_dprec.py`.

Role in the equation. D_{res} enters the numerator with elasticity $\alpha_{\text{res}}(\sigma_{\text{sys}})$, and $D_{\text{res}} \uparrow \Rightarrow DXY \uparrow$. We expect it to be close to state-independent at business-cycle frequency — reserve managers do not rebalance weekly — so like D_{trans} it should show a small calm–crisis gap, but for a different reason: not contractual obligation, but institutional slowness.

Interactions. This is the channel through which Layer 5 reaches Layer 0. The same COFER series that proxies D_{res} is the calibration target for the structural integrity coefficient $\Psi_{\text{str}}(t)$ in Chapter 9, because a sustained decline in the dollar's reserve share is simultaneously a fall in this demand channel and evidence of secular erosion. The two uses must not be confused: here the share is a demand proxy at business-cycle frequency; there it is an erosion indicator at decadal frequency. D_{res} is also the channel most directly attacked by S_{off}: an asset freeze is a demonstration to every reserve manager that dollar reserves are contingent, and its effect runs through T_{frag} into both Φ_{sys} and $\Psi_{\text{str}}(t)$ [33].

Measurement. The valuation problem is severe enough to state as a warning rather than a footnote. The COFER share is the dollar value of dollar reserves divided by the dollar value of all allocated reserves. When the dollar appreciates, the non-dollar numerator components fall in dollar terms and the dollar share rises *mechanically*, with no reserve manager having bought anything. Using this share as a demand proxy for the dollar therefore builds a positive feedback from the dependent variable into the regressor. Any estimate in Chapter 11 that uses it must either work with valuation-adjusted shares — the adjustment is the contribution of Arslanalp et al. [5], who separate active diversification from valuation — or accept a bias whose sign is known and whose magnitude is not.

Table 3.7: Measurement protocol for D_{res}.

	Primary	Secondary
Proxy	USD share of allocated official FX reserves	USD level of allocated reserves
Source	IMF COFER, *Allocated reserves, claims in US dollar, Shares*	IMF COFER, same, *Nominal value*
Frequency	quarterly	quarterly
Unit	per cent	USD millions
Transformation	level, in $[0, 1]$ after division by 100	log, indexed to 1
Availability	free	free
Weakness	**contaminated by valuation effects**; see below	same, plus it mixes accumulation with valuation

Two further caveats: COFER covers *allocated* reserves only, and the allocated share of total reserves has itself changed over the sample; and the data are quarterly and published with a lag, so like D_{debt} this channel cannot be read in real time.

Historical range. The COFER dollar share runs from 1999-Q1 to 2026-Q1, 109 quarterly observations. It began the sample at 71.2 %, peaked at 72.1 % in 1999-Q2, and stood at 57.1 % in 2026-Q1. The sample low is 56.4 % in 2025-Q4.

At the model's base period, 2006-Q1, the share was 64.9 %, so the cumulative change over the modelled sample is -7.8 percentage points. Figure 3.5 shows the path. The decline is neither smooth nor recent: roughly half of it had already happened by 2003, and the series is broadly flat from 2009 to 2016 before resuming.

That shape matters for Chapter 9. A cliff function is not the only description of these data, and the chapter must argue for one rather than assume it. What can be said here is narrower and firmer: the dollar's reserve share has fallen by about a seventh of its 1999 level over twenty-seven years, and the decline is real rather than a valuation artefact — the active-diversification decomposition establishes that [5].

Limitations. The proxy is a share, and the model wants a flow; a constant share with growing total reserves is a positive dollar flow that this measure reports as zero. The valuation contamination described above runs in the same direction as the effect being estimated. And roughly a quarter of global reserves are unallocated, held by countries that do not report composition — a group that includes some of the largest holders, which means the series is silent about exactly the reserve managers whose decisions would matter most.

3.6 Hedging Demand (D_{hedge})

src p.7

Definition 3.6. D_{hedge} is the flow of dollar cash demand generated by the opening, rolling and closing of currency hedges on foreign holdings of dollar assets. It is an index number equal to 1 at t_{base}, driven by the cost of borrowing dollars synthetically.

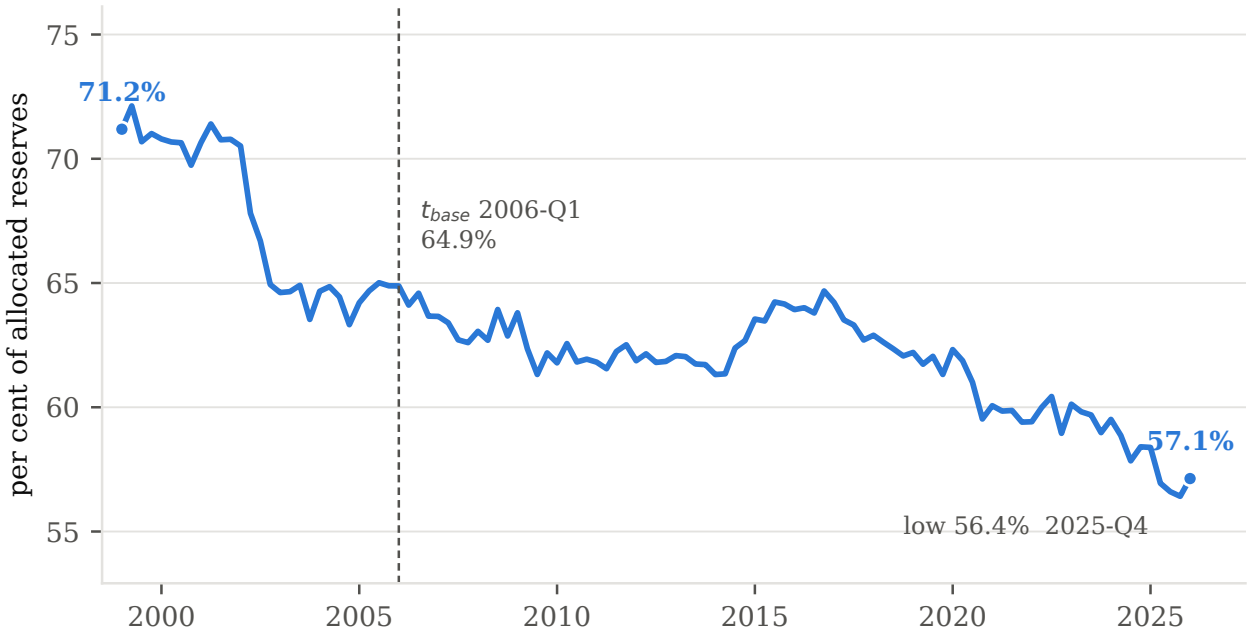


Figure 3.5: D_{res}: the U.S. dollar's share of allocated official foreign exchange reserves, 1999-Q1 to 2026-Q1, with the model's base period marked. Source: IMF COFER, allocated reserves, claims in US dollar, shares; author's calculations; script `f07_dres.py`.

Mechanism. This is a critical element of the financial plumbing, and it is the channel that connects the investment decision of Section 3.3 to the funding market. When a Japanese pension fund buys U.S. Treasuries it must hedge the currency risk, and it does so by borrowing dollars through an FX swap. The cost of that borrowing — the cross-currency basis — determines whether the investment is profitable at all. Under stress the cost spikes, and the holder faces a choice between paying it and closing the hedge; closing the hedge requires buying dollars in the spot market. The demand therefore appears exactly when the cost of supplying it is highest.

Two features distinguish it from D_{debt}. The obligation is contingent rather than contractual: a hedge can be abandoned, a coupon cannot. And the flow is generated by *holders of dollar assets* rather than by owers of dollar liabilities, so the two channels draw on different populations and must not be double-counted (Section 3.11). The persistent post-2008 deviations from covered interest parity that make this channel expensive are documented by Du et al. [25], and their association with a strong dollar and contracting cross-border lending by Avdjiev et al. [6].

Role in the equation. D_{hedge} enters the numerator with elasticity $\alpha_{\text{hedge}}(\sigma_{\text{sys}})$, and $D_{\text{hedge}} \uparrow \Rightarrow DXY \uparrow$. Like D_{debt} we expect $\alpha_i^{\text{crisis}} > \alpha_i^{\text{calm}}$, and for the same reason: the channel is quiescent in calm conditions and violent in stressed ones.

Interactions. D_{hedge} is the mechanical link between D_{cap} and the funding layer: a larger stock of hedged foreign holdings of U.S. assets means a larger rollover requirement, so $D_{\text{cap}} \uparrow$ at t implies $D_{\text{hedge}} \uparrow$ thereafter. It is amplified by $S_{\text{bank}} \downarrow$, since it is dealer balance sheet

that intermediates the swap, and relieved by $S_{\text{swap}} \uparrow$.

Measurement. The proper proxy is paywalled, so this card states a substitute explicitly, as ISSUES #23 requires.

Table 3.8: Measurement protocol for D_{hedge}. The primary proxy is not free; the substitutes are named, and so is what each of them misses.

	Primary (**paywalled**)	Substitute (free, in hand)
Proxy	3-month cross-currency basis, EUR/USD and JPY/USD	TED spread to 2022-01; SOFR minus IORB thereafter
Source	Bloomberg; Du–Tepper–Verdelhan replication data	FRED `TEDRATE`, `SOFR`, `IORB`
Frequency	daily	daily
Unit	basis points	percentage points; basis points
Availability	**Tier 3, paywalled**	free
What it misses	—	**neither is a cross-currency price.** TED is unsecured *dollar* bank funding stress; SOFR–IORB is *domestic* repo pressure. Both correlate with the basis in crises and neither measures it.

The substitution is weaker here than anywhere else in the chapter, and the card should not pretend otherwise. The cross-currency basis is the price of dollars to a foreign holder; TED is the price of dollars to a bank; SOFR minus IORB is the price of dollars against Treasury collateral in New York. They are three different prices that happen to move together when funding markets break.

Historical range (substitutes only). The TED spread runs from 2006-01-03 to 2022-01-21, 3,942 observations, and is discontinued: it died with LIBOR. It peaked at 4.58 percentage points on 2008-10-10, in the weeks after Lehman, and reached its low of 0.06 on 2021-07-08. SOFR minus IORB is available from 2021-07-29, 1,273 daily observations, ranging from −21 basis points on 2022-06-14 to +32 basis points on 2025-10-31, and standing at or above +10 basis points on only 24 days.

Note what those two ranges do *not* give us: a continuous series. TED ends in 2022-01 and the SOFR–IORB spread begins in 2021-07, so there is a six-month overlap and then a regime change in what is being measured. Any use of this pair in Chapter 11 must treat them as two series, not one spliced series.

Limitations. The primary proxy is unavailable; the substitutes measure adjacent but different prices; and the underlying quantity — the stock of hedged foreign holdings, which determines how much rollover demand a given basis move generates — is itself unobserved at any useful frequency.

3.7 Derivative Demand (D_{gamma})

Definition 3.7. D_{gamma} is the flow of spot dollar buying and selling generated by option dealers re-hedging delta as the dollar moves. It is an index number equal to 1 at t_{base}.

Mechanism. Currency markets are intermediated by dealers running delta-neutral option books. When the dollar moves, a dealer who is short gamma must trade *in the direction of the move* to restore neutrality: the dollar rises, the dealer must buy dollars, and the dollar rises further. This is a self-reinforcing loop with no reference to any economic fundamental, and it operates on a timescale of minutes.

Role in the equation. D_{gamma} enters the numerator with elasticity $\alpha_{\text{gamma}}(\sigma_{\text{sys}})$, and $D_{\text{gamma}} \uparrow \Rightarrow DXY \uparrow$. But the card must record a tension in the model's architecture here, because this channel is not really a demand *quantity* at all: it is an amplification of whatever move is already under way. A mechanism that multiplies existing moves belongs in the convexity exponent κ_{sys} of Chapter 7, not in the numerator of Equation (2.2) — and κ_{base}, the baseline convexity, is explicitly described there as a function of structural leverage, which is what dealer gamma is.

We keep D_{gamma} in the numerator because the original places it there and because Part II's case studies refer to it as a demand channel. But the honest statement is that D_{gamma} and κ_{base} are two names for closely related objects, and Chapter 10 must check that they are not double-counting the same amplification. This is flagged in Section 3.11 rather than resolved here.

Interactions. D_{gamma} amplifies every other channel rather than acting alone; it is the demand-side counterpart of the liquidity spirals of Brunnermeier and Pedersen [19], and it scales with dealer leverage in the manner of Adrian and Shin [1].

Measurement.

Historical range.

> **Not supplied, and the gap is the world's.** Dealer net gamma is not published. CFTC Commitments of Traders is free but measures futures direction, not option convexity, and cannot say whether dealers are long or short gamma — which is the entire content of the channel. *What it costs:* D_{gamma} has no historical range, no figure and no scored prediction anywhere in this book. Chapter 11 therefore treats it as unmeasured rather than measured badly, which is the right choice and still a hole in Layer 0.

Per ISSUES #23 the card states the substitute and its inadequacy rather than adopting it. No range is given, because a substitute that cannot distinguish the sign of the mechanism is not a measurement of it.

Table 3.9: Measurement protocol for D_{gamma}.

	Primary (**paywalled**)	Candidate substitutes (all inadequate)
Proxy	aggregate net gamma exposure of major FX option dealers	CFTC Commitments of Traders positioning; realised-vs-implied volatility gap
Source	DTCC trade repository; dealer surveys	CFTC (free, weekly); constructed
Frequency	daily	weekly
Availability	**Tier 3, paywalled**	free
What it misses	—	COT covers *futures*, not options, and reports direction rather than convexity: it cannot distinguish long from short gamma, which is the entire content of the mechanism

Limitations. The channel is unmeasured, its theoretical home is contested (see above), and its timescale — minutes — is below the resolution of every other series in the model.

3.8 Network Demand (D_{net})

src p.8

Definition 3.8. D_{net} is the dollar demand generated by the dollar's position in payment and settlement infrastructure: correspondent banking, CHIPS and SWIFT, and dollar-referencing stablecoins. It is an *emergent channel* in the sense of Section 2.1.2 and enters as $1 + D_{\text{net}}/\bar{X}$, so that it equals 1 when the channel is absent.

Mechanism. Two mechanisms sit under one name, and the card separates them.

The first is the classical network effect. A currency used for settlement is convenient to hold because others hold it, which makes it more convenient still. This is old, slow, and close to what Chapter 9 treats as the dollar's structural integrity; it does not vary at business-cycle frequency.

The second is new and does vary: dollar-referencing stablecoins are a demand for dollar-denominated claims arising entirely outside the banking system, and their issuers hold the reserves backing them substantially in Treasury bills. A growing stablecoin float is a growing bid for short-dated dollar assets [42]. This is the component the emergent-channel treatment of ISSUES #29 exists for: it was zero at t_{base} and cannot be indexed to it.

Role in the equation. D_{net} enters the numerator with elasticity $\alpha_{\text{net}}(\sigma_{\text{sys}})$ and $D_{\text{net}} \uparrow \Rightarrow DXY \uparrow$, in the emergent form $1+D_{\text{net}}/\bar{X}$ of Equation (2.7). The scale constant $\bar{X}$ is a stated convention and must be reported wherever the elasticity is; Chapter 11 fixes it.

Interactions. D_{net} is the demand-side twin of S_{crypto} in Chapter 4, and the two must not be confused: a dollar stablecoin is simultaneously a claim demanding Treasury collateral (here) and a dollar-like instrument supplying purchasing power outside the banking system (there). The net sign is not obvious and Chapter 15 tests it. It is attacked by C_{alt} in Chapter 9: alternative settlement rails subtract from network demand directly.

Table 3.10: Measurement protocol for D_{net}.

	Primary	Secondary
Proxy	aggregate capitalisation of USD-referencing stablecoins	SWIFT dollar share of payment value
Source	DefiLlama stablecoin series	SWIFT RMB Tracker (monthly, shares only)
Frequency	daily	monthly
Transformation	$1 + X/\bar{X}$ (emergent channel, ISSUES #29)	level share
Availability	free	free, but published as a share and only for selected currencies
Weakness	capitalisation is a stock, and its Treasury-bill backing varies by issuer	measures messaging, not settlement value, and excludes CHIPS

Measurement.

Historical range. The primary proxy is now in hand. The DefiLlama aggregate of dollar-pegged stablecoin circulating supply runs daily from 2017-11-29 to 2026-09-05, 3,203 observations, from effectively zero to $310.5 billion, with a sample maximum of $320.9 billion on 2026-05-17. Non-dollar-pegged tokens are 0.51 per cent of the total at the end of the sample and have never been more: the stablecoin float is, to a first approximation, entirely a dollar instrument, which is the fact this card rests on.

Figure 3.6 plots it on a logarithmic axis, because the series spans five orders of magnitude and a linear axis would show nothing before 2020.

The scale constant. ISSUES #29 requires $\bar{X}$ to be stated wherever the elasticity is. The convention adopted here is the sample mean over the period the series exists, which gives $\bar{X} = \$118.6$ billion, so the channel enters as $1 + D_{\text{net}}/118.6$ with D_{net} in billions. Two things about this number must travel with it. It is a convention, not an estimate; and because the series is monotonically trending over most of its life, the sample mean is a statement about *when the sample was taken* rather than about the channel. A sample ending in 2021 would give roughly a third of it. Chapter 11 may replace it, and if it does, every elasticity reported against the old constant must be restated.

Coverage against the crises the model cares about. The series begins in 2017-11. It therefore does not exist for 2008, and although it exists in March 2020 it stood at $5.2 billion

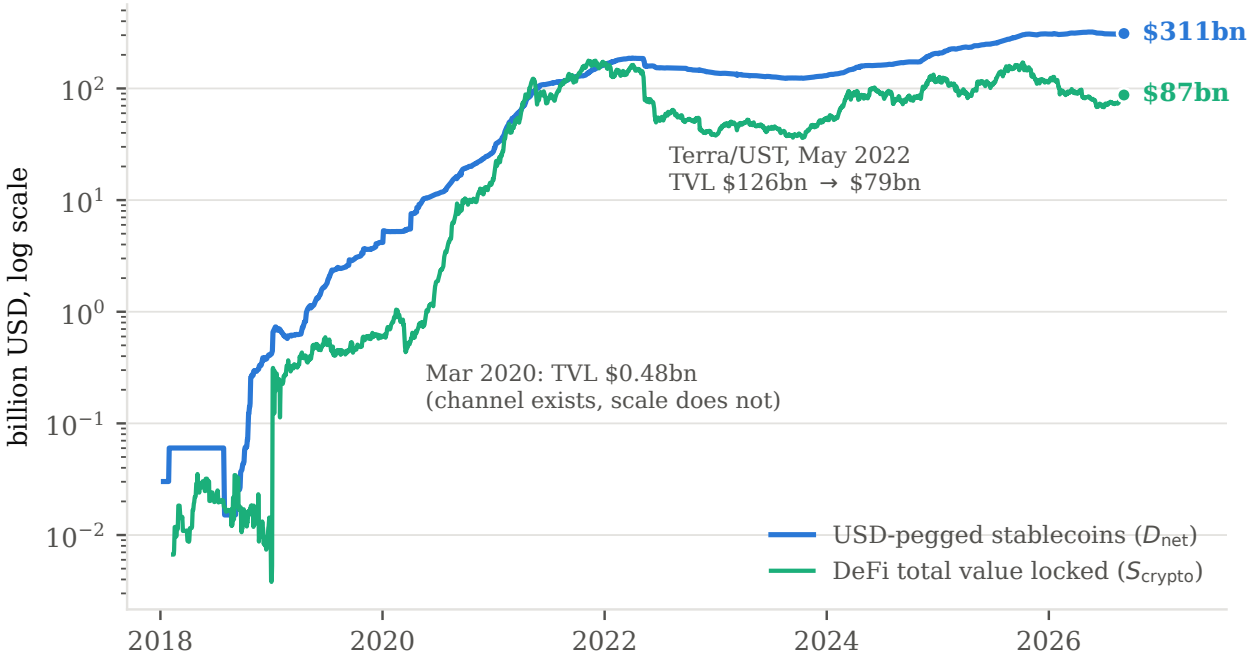


Figure 3.6: D_{net} and its supply-side twin. Circulating supply of USD-pegged stablecoins (the D_{net} proxy) and total value locked in DeFi (the S_{crypto} proxy of Section 4.7), daily, log scale. One figure serves both cards because it is one object read in two directions. Source: DefiLlama, files `defillama_stablecoins_d.json` and `defillama_tvl_d.json`; scripts `extract_crypto_mmf.py` and `f08_dnet.py`.

on 2020-02-19 and $5.5 billion on 2020-03-23 — present, but three orders of magnitude below the swap lines drawn in the same fortnight. The first episode at which this channel is large enough to matter is 2022. Any claim that D_{net} contributed to dollar strength before then is a claim about a channel that was not yet there, and Chapter 11 records it among the coverage gaps.

What the aggregate cannot see. In the week of the Silicon Valley Bank failure the aggregate barely moves: $134.4 billion on 2023-03-10 and $134.8 billion on 2023-03-13. Yet that weekend is the sharpest stablecoin stress in the sample, when USDC lost its peg on the disclosure of deposits held at the failed bank. The aggregate is blind to it because the run was a *substitution* between issuers, not a contraction of the float. Composition, not capitalisation, is what a stress reading of this channel requires, and this card does not have it.

Limitations. The two mechanisms under this heading move on timescales three orders of magnitude apart, and one proxy cannot serve both. The stablecoin leg is a stock proxying a flow, as with D_{debt}. And the reserve backing of the float is heterogeneous and imperfectly disclosed, so the translation from capitalisation to Treasury-bill demand carries an unmeasured coefficient.

3.9 Microstructural Demand (D_{micro})

Definition 3.9. D_{micro} is the calendar-driven component of dollar demand: quarter- and year-end repatriation, tax dates, and index rebalancing. It is a constructed multiplier, not an estimated series.

Mechanism. Balance-sheet reporting dates create predictable dollar demand. Non-U.S. banks shrink their books into quarter end to flatter leverage ratios, which withdraws dollar intermediation exactly when corporates want it; funds repatriate for reporting; tax payments cluster. None of this reflects any view about the dollar. It is the mechanical residue of the accounting calendar, and it is the only channel in this chapter that is known in advance.

Role in the equation. D_{micro} enters the numerator with elasticity $\alpha_{\text{micro}}(\sigma_{\text{sys}})$. The original specifies it directly as a calendar dummy,

$$D_{\text{micro}} = \begin{cases} 1.02 & \text{on the final trading day of a quarter,} \\ 1 & \text{otherwise,} \end{cases} \tag{3.2}$$

which is E2 of the equation compendium. The value 1.02 is an *assumption (original)*: the source states it without derivation, and we retain it as stated rather than substituting an estimate. Chapter 11 can estimate the quarter-end effect directly from the dependent variable, and if it does, Equation (3.2) should be replaced by that estimate and the substitution recorded.

Interactions. D_{micro} interacts with S_{bank}, and this is the one case in the chapter where a demand channel and a supply channel share a single cause: quarter-end balance-sheet compression raises D_{micro} and lowers S_{bank} simultaneously. The two effects push DXY the same way and are not independent, which Chapter 4 notes again from the other side.

Measurement. No data are required: the series is constructed from the calendar and is exact. The only judgement is the magnitude, discussed above.

Historical range. By construction, $D_{\text{micro}} \in \{1, 1.02\}$, taking the higher value on roughly four days a year.

Limitations. A single multiplier applied to every quarter end assumes the effect is constant in size, which it is not: quarter-end pressure varies with the level of bank balance-sheet constraint, and was visibly larger in 2016–2019 than before. A regime-dependent D_{micro} would be more faithful, at the cost of a parameter.

3.10 Human-Capital and IP Demand (D_{hcap})

src p.8

Definition 3.10. D_{hcap} is the annual dollar demand flow generated by the U.S. stock of human capital and intellectual property: $D_{\text{hcap}} = h(H_{\text{cap}})$ with $h' > 0$, where H_{cap} is the stock. It is an index number equal to 1 at t_{base}.

Mechanism. The stock acts as a long-term value magnet. A country that continuously produces intellectual and technological assets that the rest of the world wants to own generates a persistent bid for its currency, because acquiring those assets requires acquiring the currency first. Royalty and licensing receipts are the direct flow; the larger effect is indirect, through the equity claims that this productive capacity supports, which is the D_{cap} channel.

Role in the equation. D_{hcap} enters the numerator with elasticity $\alpha_{\text{hcap}}(\sigma_{\text{sys}})$, and $D_{\text{hcap}} \uparrow \Rightarrow DXY \uparrow$. Following the resolution of ISSUES #24, the *flow* D_{hcap} carries the elasticity and the *stock* H_{cap} is a Layer-2 structural variable (Chapter 6). The original wrote the identity $D_{\text{hcap}} \equiv H_{\text{cap}}$; separating them is what allows the erosion mechanism to act on the stock and reach the flow with a lag, which is the dual-erosion claim of Chapter 6.

Interactions. D_{hcap} feeds D_{cap} over long horizons. It is eroded by R_{US}, in the specific sense that political risk degrades the institutions — immigration, research funding, IP enforcement — that maintain the stock. This is the slowest mechanism in the model that is not in Layer 5.

Measurement. The tax-routing distortion in the first column is severe. A large share of recorded U.S. IP receipts reflects where multinationals choose to book royalties rather than where the intellectual property was created, and the pattern shifts with tax law — notably

Table 3.11: Measurement protocol for D_{hcap}, flow and stock.

	Flow (D_{hcap})	Stock (H_{cap})
Proxy	receipts for the use of intellectual property, net	STEM workforce stock; resident patent families
Source	BEA International Transactions Accounts	NSF/NCSES *Science and Engineering Indicators*; WIPO
Frequency	quarterly / annual	annual
Transformation	log, indexed to 1 at t_{base}	normalised to $[0, 1]$ (ISSUES #26)
Availability	free	free
Weakness	IP receipts are heavily distorted by tax-driven intra-firm royalty routing through low-tax jurisdictions	a stock of people and patents is a poor measure of the productive capacity it is meant to represent

around 2017–2018. A series that moves when the tax code moves is a poor proxy for a variable that is supposed to change over decades.

Historical range.

> **Not supplied, and the gap is ours.** BEA International Transactions charges for the use of intellectual property, and an NSF/NCSES STEM stock series, are both free and both unsupplied. *What it costs:* D_{hcap} has **no rows in `registry/scorecard.csv` at all** — it is the only channel in the model that no case study examines. Section 18.2.3's dual-erosion effect, which the original calls among its most dangerous findings, is consequently unevidenced in either direction.

No range is stated.

Limitations. Both proxies measure something adjacent to the concept. Neither captures the mechanism the original actually describes — that the world buys dollars to own claims on American technological capacity — which is an equity-market phenomenon running through D_{cap}, and which would be double-counted if it were also credited here. The card therefore measures the narrow royalty flow and leaves the broad effect to D_{cap}, at the cost of understating the channel.

3.11 Interactions Among Demand Drivers

The ten channels are not independent, and a multiplicative aggregator makes their dependence consequential: in logs the numerator is a weighted sum, so any pair that moves together contributes twice to $\ln DXY$ and, if the two are proxied by related series, the elasticities are not separately identified. This closing section records the interactions that matter, in three groups.

3.11.1 Amplifications

The original's central interaction claim is that elevated safe-haven demand intensifies the effect of a collateral shortage rather than merely adding to it. In the model as specified this operates through κ_{sys}: a flight to quality into an asset whose collateral market is impaired produces a larger price move than the flow alone implies, and κ_{sys} is the term that scales all channels together (Section 2.2.3). Table 3.12 lists the pairs.

Table 3.12: Interactions among demand channels and with the supply side. "Amplifies" means the effect is multiplicative and runs through κ_{sys}; "feeds" means one channel raises another's level with a lag.

Channel	With	Kind	Mechanism
D_{prec}	S_{coll}	amplifies	flight to quality into an impaired collateral market
D_{debt}	S_{bank}, S_{coll}	amplifies	rollover is harder when dealer capacity is scarce
D_{cap}	D_{hedge}	feeds	hedged holdings create later rollover demand
D_{hcap}	D_{cap}	feeds	the stock generates the assets that attract the flow
D_{micro}	S_{bank}	common cause	quarter-end balance-sheet compression moves both
D_{gamma}	all	amplifies	dealer re-hedging scales any move already under way

3.11.2 Boundaries: what must not be double-counted

Three pairs are close enough to overlap, and the boundary between them is drawn here so that Part II can apply it consistently.

D_{debt} **and** D_{hedge} draw on different populations. D_{debt} is the principal-and-interest flow owed by borrowers of dollars; D_{hedge} is the rollover cost faced by holders of dollar assets. A Japanese pension fund appears in the second and not the first; a Turkish corporate appears in the first and not the second.

D_{net} **and** S_{crypto} are the two faces of the same instrument. A dollar stablecoin demands Treasury collateral (numerator) and supplies dollar-like purchasing power outside the banking system (denominator). Both entries are correct; the net effect is an empirical question, and Chapter 15 treats it as one.

D_{gamma} **and** κ_{base} are, as Section 3.7 concedes, closely related objects: dealer gamma is structural leverage, and structural leverage is what Chapter 7 makes κ_{base} a function of. Because D_{gamma} is unmeasured, the practical risk of double-counting is currently low; if a gamma series is ever obtained, Chapter 10 must check that the same amplification is not entering twice.

3.11.3 The policy-divergence question

The original's D_{div}, introduced in a single case study with a policy-divergence driver, is absorbed into D_{cap} and does not exist as a separate variable in this edition. The argument is

given in full at Section 3.3: divergence is what D_{cap}'s proxy already measures, and a separate term would be collinear with it by construction (ISSUES #6). The registry retains a row marked ABSORBED for traceability.

3.11.4 Collinearity that the data cannot resolve

Finally, one interaction is not a feature of the world but of our measurement. D_{prec} and σ_{sys} are currently proxied by the same series, so the safe-haven channel and the state variable that modulates it cannot be told apart in estimation. This is ISSUES #31, it is an open question for the author rather than an editorial matter, and until it is settled every quantitative statement about $\alpha_{\text{prec}}(\sigma)$ in this book should be read as conditional on its resolution. Section 3.4 states the three ways out.

Chapter 4

Decomposing the Denominator: Effective Supply Dynamics (S^*)

The denominator S^* aggregates seven supply channels. Where Chapter 3 asked who buys dollars and why, this chapter asks who supplies them and what stops them.

The asymmetry between the two sides is worth stating at the outset, because it shapes everything that follows. Demand for dollars is dispersed across millions of agents with different motives, and the cards in Chapter 3 are largely about *motive*. Supply is concentrated: two sovereign actors, a few dozen dealers, and a plumbing system whose capacity is set by regulation. The supply cards are therefore largely about *constraint* — not who wants to supply dollars, but who is able to.

This has a consequence the model exploits. Supply constraints bind suddenly. A balance sheet that is adequate on Thursday is inadequate on Friday, and the transition is not smooth. That is a large part of why the convexity exponent κ_{sys} of Chapter 7 is needed at all: it is on the denominator side that the system's non-linearity is most mechanically obvious.

4.1 Federal Reserve Actions (S_{Fed})

src p.9

Definition 4.1. S_{Fed} is the sovereign supply of dollar base money, determined jointly by its price — the policy rate P_i — and its quantity — the Federal Reserve balance sheet M_L. It is an index number equal to 1 at t_{base}.

Mechanism. This is the core sovereign source of dollars, and the Federal Reserve controls both of its dimensions.

Price. When the Fed raises the federal funds rate, borrowing dollars becomes more expensive and effective supply contracts. Nothing about the quantity of reserves need change for this to bite.

Quantity. When the Fed purchases assets under quantitative easing it injects reserves into the banking system and S_{Fed} rises; quantitative tightening withdraws them and S_{Fed} falls.

The two legs can and do move in opposite directions — the balance sheet was still growing while the rate was rising in early 2022 — which is why the card keeps them separate rather

Table 4.1: Data status of the seven supply channels plus the two placement questions. No number in this chapter is estimated, interpolated or assumed.

Channel	Primary proxy	Status	Gap
S_{Fed}	Fed total assets; fed funds rate	**complete**	—
S_{Tsy}	TGA balance; ON RRP balance	**complete**	—
S_{coll}	SOFR − IORB; bills outstanding	**partial**	bills outstanding (Treasury MSPD); spread starts 2021-07
S_{bank}	G-SIB supplementary leverage ratio	**none**	FR Y-9C construction; no ready series
S_{swap}	Fed central-bank liquidity swaps	**complete**	—
S_{ext}	U.S. current account balance	**complete**	—
S_{crypto}	DeFi TVL; decentralised stablecoin supply	**no data**	DefiLlama
S_{off}	immobilisation intensity	**no data**	frozen-reserve estimates; OFAC
S_{shadow}	non-bank repo / MMF creation	**no data**	OFR repo data; FSB NBFI monitor

than collapsing them into a single stance measure.

Role in the equation. S_{Fed} enters the denominator with elasticity $\beta_{\text{Fed}}(\sigma_{\text{sys}})$, and $S_{\text{Fed}} \uparrow \Rightarrow DXY \downarrow$. We expect a large calm-to-crisis gap with the crisis value *smaller*: in a genuine dollar shortage, an increase in reserves does less to relieve it than the same increase would in calm conditions, because the binding constraint has moved from the quantity of reserves to the capacity of the balance sheets that distribute them (S_{bank}, Section 4.4). This is the single most policy-relevant claim in the chapter and Chapter 19 builds on it.

Interactions. S_{Fed} is transmitted by S_{bank} and can be entirely blocked by it; reserves that dealers will not deploy are not effective supply. It interacts with S_{Tsy} mechanically, since the Treasury General Account sits on the liability side of the same balance sheet: a TGA drawdown raises reserves without any action by the Fed.

Measurement. Combining the two legs into a single S_{Fed} index requires a weighting that the model does not supply, and we do not invent one: Chapter 11 carries them as two regressors sharing one elasticity rather than as a composite. That is a departure from the original's "composite index" language and it is a deliberate one, because any fixed weighting would be exactly wrong in the periods when the two legs diverge.

Table 4.2: Measurement protocol for S_{Fed}.

	Quantity leg (M_L)	Price leg (P_i)
Proxy	Federal Reserve total assets	effective federal funds rate
Source	FRED `WALCL`	FRED `DFF`
Frequency	weekly (Wednesday)	daily
Unit	USD millions	per cent
Transformation	log, indexed to 1 at t_{base}	level
Availability	free; in hand	free; in hand
Weakness	total assets mix monetary operations with emergency facilities and with maturity choice	the policy rate is a poor summary of stance when the balance sheet is doing the work

Historical range. Fed total assets run weekly from 2006-01-04 to 2026-09-02, 1,079 observations. They stood at $0.83 trillion on 2006-02-08, peaked at $8.97 trillion on 2022-04-13, and were $6.74 trillion at the latest observation — so the balance sheet is roughly eight times its pre-crisis size and has retraced about a quarter of the way from its peak.

The effective federal funds rate runs daily from 2006-01-01 to 2026-09-03, 7,551 observations, from a high of 5.41 per cent on 2007-02-28 to a low of 0.04 per cent on 2011-12-30, and stands at 3.63 per cent.

Limitations. Total assets are a gross measure: they rise when the Fed lends against collateral in a crisis and when it buys Treasuries in a stimulus, which are different acts with different effects on effective supply. The two legs have no agreed common scale. And the weekly frequency of the balance sheet is coarse relative to the daily frequency of the rate.

4.2 U.S. Treasury Actions (S_{Tsy})

src p.9

Definition 4.2. S_{Tsy} is the effect on dollar supply of the Treasury's cash management and of the Fed facility that absorbs money-market cash: the Treasury General Account balance and the overnight reverse repo balance. Both act *negatively* on effective supply.

Mechanism. The Treasury is the second sovereign actor, and it can move dollar liquidity as powerfully as the Fed through entirely different channels.

When the Treasury collects taxes or issues bonds and parks the proceeds in its account at the Fed, those dollars leave the private sector. When money-market funds place cash at the Fed through the overnight reverse repo facility, the same thing happens. In both cases the dollars still exist, but they are outside the banking system and cannot fund anything.

This is the clearest case in the model of a channel where the *sign* is the whole content: a rising TGA is a *fall* in effective supply, and the card's arithmetic must carry that inversion explicitly.

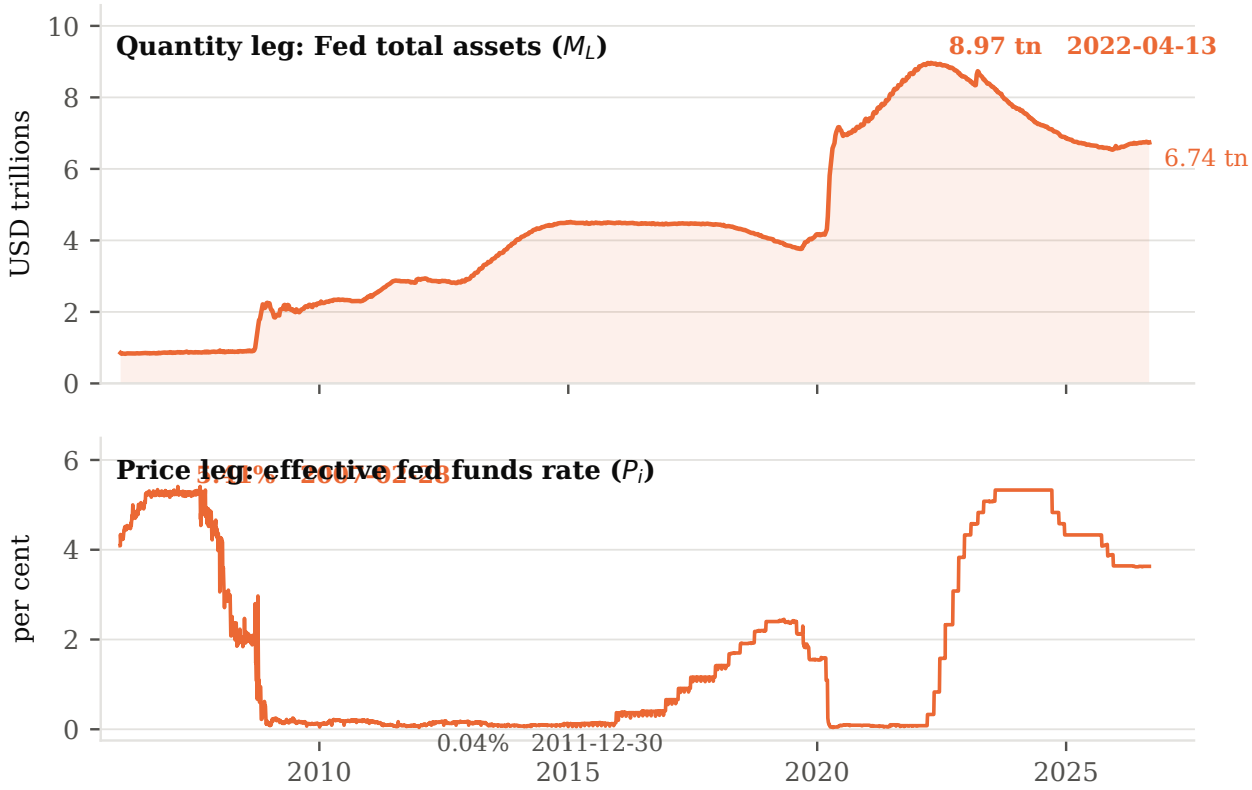


Figure 4.1: S_{Fed}: the two legs of Federal Reserve dollar supply, 2006–2026. Source: FRED `WALCL` and `DFF`; author's calculations; script `f11_sfed.py`.

Role in the equation. S_{Tsy} enters the denominator with elasticity $\beta_{\text{Tsy}}(\sigma_{\text{sys}})$. Because the underlying balances act negatively, the index is constructed so that $S_{\text{Tsy}} \uparrow \Rightarrow DXY \downarrow$ like every other supply channel: the proxy enters with its sign flipped, which the measurement protocol states explicitly. The alternative — an exception to the book's sign convention — was rejected in Section 2.7.

Interactions. S_{Tsy} and S_{Fed} share a balance sheet, so their sum is more meaningful than either alone; Section 4.10 treats this as a double-counting question. The ON RRP leg interacts with S_{coll}: money-market cash at the Fed is cash not financing repo, so a large RRP balance and a tight collateral market are the same phenomenon seen from two sides.

Table 4.3: Measurement protocol for S_{Tsy}. Both legs enter with a negative sign.

	TGA leg	ON RRP leg
Proxy	Treasury General Account balance	overnight reverse repo balance
Source	FRED `WTREGEN`	FRED `RRPONTSYD`
Frequency	weekly	daily
Unit	USD millions	USD billions
Transformation	**negated**, then indexed to 1 at t_{base}	**negated**, then indexed
Availability	free; in hand	free; in hand
Weakness	the TGA's target level is a policy choice that changed regime in 2015	the facility did not exist in its current form before 2013, so early values are not comparable

Measurement.

Historical range. The TGA runs weekly from 2006-01-04 to 2026-09-02. It stood at $3.86 billion on 2008-02-20 — the Treasury then held almost no cash buffer — and peaked at $1,816.69 billion on 2020-07-29, at the height of pandemic borrowing. It is $967.94 billion now.

The ON RRP balance runs daily from 2007-04-26 to 2026-09-04. It was zero as recently as 2019-11-25, peaked at $2,553.72 billion on 2022-12-30, and has since drained to $0.68 billion.

That round trip deserves emphasis, because it is the largest single movement of dollar liquidity in the sample and it was accomplished without any change in the Fed's asset holdings. Between end-2022 and 2026 roughly two and a half trillion dollars moved from the Fed's liability side back into the banking system. Any account of the dollar over that period that looks only at S_{Fed} misses it entirely.

Limitations. The ON RRP facility's counterparty list and caps have changed repeatedly, so the series is not a constant-definition measure. The TGA's level reflects a debt-management policy that changed after 2015, when the Treasury adopted a minimum cash

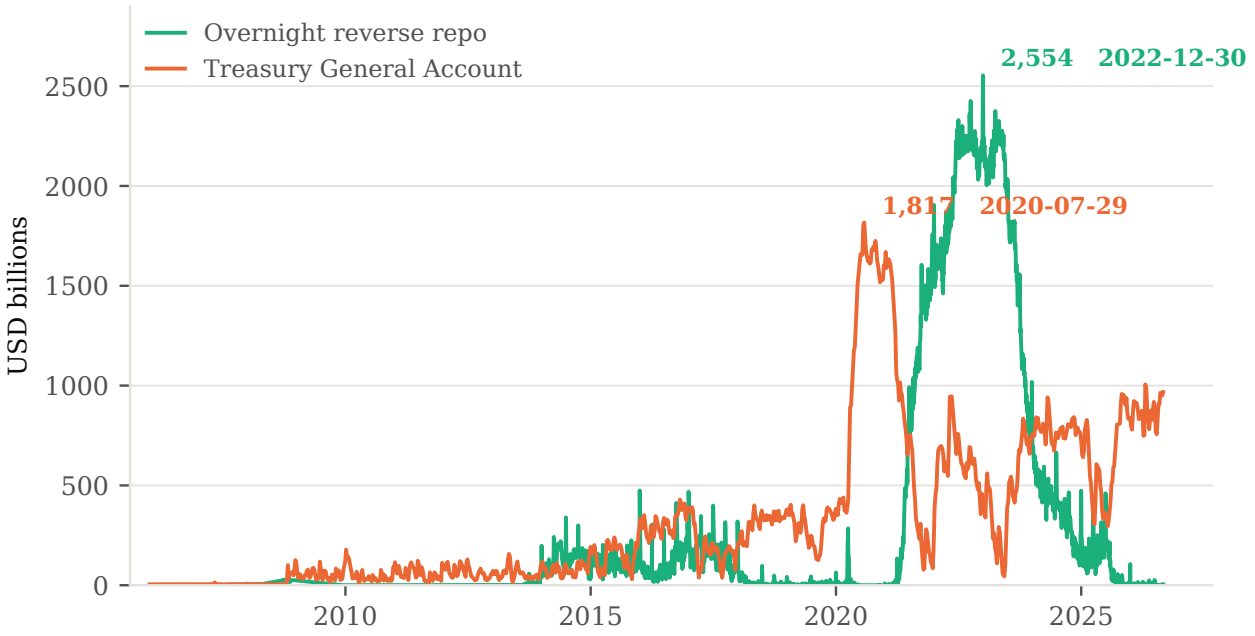


Figure 4.2: S_{Tsy}: dollars parked outside the banking system, 2006–2026. Both series are balances at the Federal Reserve; both reduce effective dollar supply. Source: FRED `WTREGEN` and `RRPONTSYD`; author's calculations; script `f12_stsy.py`.

balance. Neither is a market price, so neither reveals the shadow value of the liquidity withdrawn.

4.3 Collateral Availability (S_{coll})

Definition 4.3. S_{coll} is the capacity of the private repo market to create dollars, which depends on the supply of eligible collateral. It includes, as a named component, non-bank private dollar creation S_{shadow} (Section 4.9).

Mechanism. Most dollars are not printed by the Fed. They are created privately, in the repo market, against collateral. To borrow dollars one needs eligible collateral, overwhelmingly Treasury bills. If bills are scarce — absorbed by quantitative easing, or simply not issued — the repo market clogs, private dollar creation collapses, and S_{coll} falls sharply.

The mechanism is a quantity constraint that shows up as a price. When collateral is scarce relative to cash, secured overnight rates rise above what the Fed pays on reserves, because lenders of cash have somewhere better to put it and borrowers of cash must bid.

Role in the equation. S_{coll} enters the denominator with elasticity $\beta_{\text{coll}}(\sigma_{\text{sys}})$, and $S_{\text{coll}} \uparrow \Rightarrow DXY \downarrow$. It is the channel most implicated in the amplification claims of Chapter 3: a flight to quality into an impaired collateral market produces a larger dollar move than the flow alone implies, which is why Table 3.12 lists $D_{\text{prec}} \times S_{\text{coll}}$ first.

Interactions. S_{coll} is reduced by S_{Fed} under quantitative easing — the Fed buying collateral removes it from the private market, so the two sovereign channels work against each other on this margin. It is reduced by the ON RRP leg of S_{Tsy}. It amplifies D_{debt} and D_{hedge}, both of which require rolling secured funding.

Table 4.4: Measurement protocol for S_{coll}.

	Price leg (in hand)	Quantity leg (missing)
Proxy	SOFR minus interest on reserve balances	net Treasury bills outstanding
Source	FRED `SOFR`, `IORB`	Treasury Monthly Statement of the Public Debt
Frequency	daily	monthly
Unit	basis points	USD billions
Transformation	level; **negated** for the index, since a wider spread means less capacity	log, indexed to 1
Availability	free; in hand, but only from 2021-07	free; **not in data/raw**
Weakness	begins in 2021-07 because IORB does, so it cannot see 2008, 2019 or 2020	a stock of collateral is not the same as its availability, much of which is re-used

Measurement. The date restriction in the price leg is severe and must not be glossed. IORB begins on 2021-07-29, so the spread cannot be computed for the September 2019 repo episode [2, 24] or for March 2020 [26], which are the two events this channel exists to explain. For those windows Chapter 13 must use the effective federal funds rate or the IOER-based spread as a substitute and say which.

Historical range. The SOFR–IORB spread is available daily from 2021-07-29 to 2026-09-03, 1,273 observations. It ranges from −21 basis points on 2022-06-14 — cash overwhelmingly abundant, in the depths of the reverse-repo build-up — to +32 basis points on 2025-10-31. The spread is at or above zero on 164 of 1,273 days, about 13 per cent of the sample, and it stands at +1 basis point now.

The pattern in Figure 4.3 is a transition rather than a cycle: persistent negative readings while the reverse repo facility was full, giving way to a regime of frequent positive spikes as it drained. Read alongside Figure 4.2, the two figures are the same story told in quantities and in prices.

Limitations.

Not supplied, and the gap is ours. Net Treasury bills outstanding, from the Treasury's Monthly Statement of the Public Debt, is free and unsupplied. *What it costs:* S_{coll} is measured only through the overnight reverse-repurchase facility, which records cash placed *at* the Federal Reserve rather than collateral scarcity

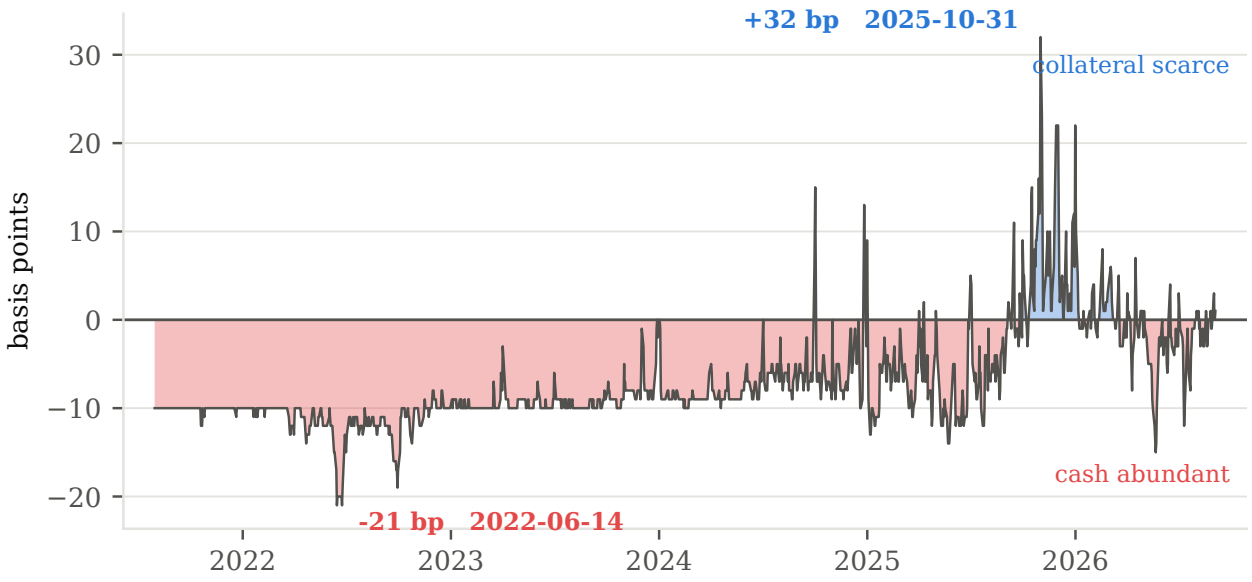


Figure 4.3: S_{coll}, price leg: the repo–policy spread, SOFR minus interest on reserve balances, 2021-07-29 to 2026-09-03. Positive means collateral is scarce relative to cash. The series cannot begin earlier because IORB does not. Source: FRED `SOFR` and `IORB`; author's calculations; script `f13_scoll.py`.

> in the market. Section 14.2.9 states the consequence plainly: the variable has a proxy and the proxy measures something else.

Beyond the missing quantity leg and the 2021 start date: collateral is re-used, so the outstanding stock understates the volume of dollar creation it supports by a multiple that is itself cyclical. That re-use rate is the object Pozsar [62] maps, and it is not measured here.

4.4 Bank Balance-Sheet Capacity (S_{bank})

Definition 4.4. S_{bank} is the capacity of globally systemic banks and primary dealers to intermediate dollars: to expand their balance sheets in order to lend, make markets and warehouse risk.

Mechanism. The Fed can inject reserves, but banks and primary dealers are the pipes that carry dollars to the people who need them. Their capacity is constrained by regulatory capital rules, most bindingly the supplementary leverage ratio, which charges capital against total assets regardless of risk. If their balance sheets are full, they cannot expand lending or market-making regardless of what the Fed does.

This is the channel that converts a monetary policy stance into a market outcome, and it is the reason a central bank can flood a system with reserves and still not relieve a dollar shortage. It is also, mechanically, why $\beta_{\text{Fed}}(\sigma_{\text{sys}})$ should fall in crisis (Section 4.1).

Role in the equation. S_{bank} enters the denominator with elasticity $\beta_{\text{bank}}(\sigma_{\text{sys}})$, and $S_{\text{bank}} \uparrow \Rightarrow DXY \downarrow$. Its behaviour is closer to a switch than a dial: capacity is ample un-

til it is not. The procyclicality of intermediary leverage is documented by Adrian and Shin [1], and the spiral it can produce with market liquidity by Brunnermeier and Pedersen [19].

Interactions. S_{bank} gates S_{Fed}. It is depressed at quarter ends by the same reporting-date mechanism that raises D_{micro} — the one common cause spanning both sides of the model (Section 3.11). It is relieved by S_{swap}, which routes dollars around constrained New York balance sheets entirely.

Measurement. No adequate free series exists, and per ISSUES #23 the card says so rather than adopting a poor substitute.

Table 4.5: Measurement protocol for S_{bank}.

	Primary	Candidate substitutes
Proxy	average supplementary leverage ratio of U.S. G-SIBs	primary-dealer net financing (NY Fed); bank credit to the private sector
Source	constructed from FR Y-9C filings	NY Fed primary dealer statistics; FRED `TOTBKCR`
Frequency	quarterly	weekly
Availability	free but **must be constructed**; no ready series	free
What it misses	—	dealer financing measures *use* of capacity, not capacity itself, so it falls when constraints bind and also when demand falls — the two cases the card must distinguish

Historical range.

> **Not supplied, and the gap is ours.** The supplementary leverage ratio must be constructed from FR Y-9C filings, which are free. No ready-made series exists and none has been built here. *What it costs:* S_{bank} has two scored predictions and **both are untestable**. The substitutes considered cannot separate constrained capacity from weak demand, which is the distinction the channel exists to make, so Section 19.2.4's leverage-relief lever is documented as institutional fact and untested as a model claim.

No range is stated.

Limitations. The channel is unmeasured. Worse, its natural substitutes are observationally equivalent to the thing they would replace: when dealer financing falls we cannot tell from the series alone whether dealers would not lend or borrowers would not borrow. Resolving that requires a price — which is S_{coll} — so in practice Chapter 11 identifies S_{bank} only jointly with S_{coll}, if at all.

4.5 Federal Reserve Swap Lines (S_{swap})

Definition 4.5. S_{swap} is the emergency supply of dollars provided to allied central banks through Federal Reserve liquidity swap arrangements. It is an *emergent channel* in the sense of Section 2.1.2 and enters as $1 + S_{\text{swap}}/\bar{X}$, since it is exactly zero for long stretches.

Mechanism. These are the system's safety valves. When global stress spikes and the world faces a dollar shortage, the Fed activates swap lines with the ECB, the Bank of Japan, the Bank of England and others. Those central banks then supply dollars directly to their own banks, bypassing clogged New York funding channels. The effect is to create dollar supply exactly where the shortage is, without requiring any constrained private balance sheet to intermediate it.

The design is deliberate and its effects are measured rather than assumed [8]: the facility caps the price of offshore dollar funding, which is why it appears in Chapter 19 as the first lever of crisis response.

Role in the equation. S_{swap} enters the denominator with elasticity $\beta_{\text{swap}}(\sigma_{\text{sys}})$ in the emergent form $1 + S_{\text{swap}}/\bar{X}$, and $S_{\text{swap}} \uparrow \Rightarrow DXY \downarrow$. Its regime signature is the most extreme in the model: $\beta_{\text{swap}}(\sigma_{\text{sys}})$ should be near zero in calm conditions, because the facility is unused, and large in crisis. A channel that is literally zero most of the time cannot have a constant elasticity, and this card is the clearest justification for the two-regime structure of ISSUES #7.

Interactions. S_{swap} relieves D_{debt} and D_{hedge} directly, by supplying the dollars those channels are scrambling for. It substitutes for S_{bank}: it is precisely the channel that works when dealer capacity does not. It is gated by G_{pol} — the list of counterparties is an alliance structure, not an economic one, which Chapter 6 takes up and Chapter 14 tests.

Table 4.6: Measurement protocol for S_{swap}.

	Primary	Secondary
Proxy	central bank liquidity swaps outstanding	counterparty count and facility caps
Source	FRED `SWPT` (H.4.1)	Federal Reserve press releases
Frequency	weekly	event
Unit	USD millions	count; USD billions
Transformation	$1 + X/\bar{X}$ (emergent channel, ISSUES #29)	event dummies
Availability	free; in hand	free
Weakness	measures *drawings*, not the insurance value of an undrawn facility	announcement effects are real and are invisible in the balance

Measurement. The weakness in the last row is the substantive one. Much of what swap lines do is done by existing: an undrawn facility caps the price of offshore dollars because everyone knows it is there. A proxy that measures only drawings therefore records nothing in exactly the states where the facility is most effective at preventing a shortage.

Historical range. Swap balances run weekly from 2006-01-04 to 2026-09-02. The peak is $583 billion on 2008-12-17, in the months after Lehman. The 2020 reactivation peaked at $449 billion on 2020-05-27. The balance is zero at the latest observation.

The distribution is the point. The balance is non-zero in 910 of 1,079 weeks, about 84 per cent — so the facility is technically almost always in use — but the level is negligible outside three episodes: 2008–09, the euro-area crisis of 2011–12, and 2020. This is a valve, not a stock, and Figure 4.4 shows it.

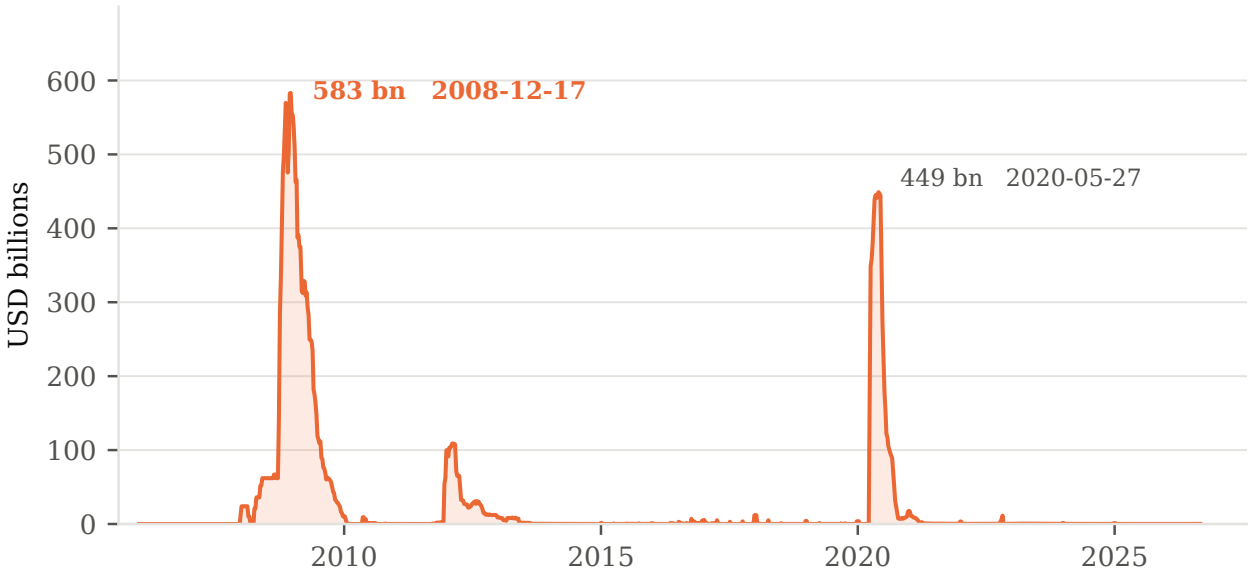


Figure 4.4: S_{swap}: Federal Reserve central bank liquidity swaps outstanding, weekly, 2006–2026. Source: FRED `SWPT` (H.4.1); author's calculations; script `f14_sswap.py`.

Limitations. Drawings understate the channel, as above. The counterparty set is a political object that has changed twice in the sample. And the weekly H.4.1 frequency is coarse relative to the days on which activation matters.

4.6 External Balance-of-Payments Flows (S_{ext})

src p.10

Definition 4.6. S_{ext} is the physical export of dollars from the United States to the rest of the world through the current account. The quantity entering the denominator is the *effective* external channel $S_{\text{ext}}^{\text{eff}} = S_{\text{ext}} e^{-\delta_{\text{off}} S_{\text{off}}}$ of Equation (2.6).

Mechanism. The U.S. current-account deficit is the mechanism by which America exports dollars in exchange for real goods. A wider deficit sends more dollars abroad and raises external supply. If the deficit narrows, the rest of the world becomes more dollar-starved and DXY rises.

This is the slowest channel in the chapter and the only one that is a genuine flow of dollars across a border rather than a balance-sheet operation.

Role in the equation. S_{ext} enters the denominator with elasticity $\beta_{\text{ext}}(\sigma_{\text{sys}})$, applied to $S_{\text{ext}}^{\text{eff}}$ rather than to S_{ext} itself, and $S_{\text{ext}} \uparrow \Rightarrow DXY \downarrow$. The attenuation by S_{off} is the subject of Section 4.8.

Interactions. S_{ext} is the mirror of D_{trans}: the same trade that requires dollars for settlement also delivers them to exporters. The model counts the two separately because they operate on different populations and at different speeds, and Section 4.10 checks that this is not double-counting.

Table 4.7: Measurement protocol for S_{ext}.

	Primary	Note
Proxy	U.S. balance on current account	the deficit is negative; the index uses its absolute value, so a larger deficit is a larger S_{ext}
Source	FRED `IEABC` (BEA)	replaces `BOPBCA`, which ends in 2014-Q1
Frequency	quarterly, seasonally adjusted	
Unit	USD millions	
Transformation	sign-flipped, log, indexed to 1 at t_{base}	
Availability	free; in hand	
Weakness	the current account is a net concept; gross flows are far larger and are what actually moves the exchange rate	

Measurement.

Historical range. The current account runs quarterly from 2006-Q1 to 2026-Q1, 81 observations, all in deficit. The deficit was smallest at $69.21 billion in 2013-Q4 and largest at $438.24 billion in 2025-Q1; it stands at $226.83 billion.

The 2025-Q1 reading is an outlier of a specific kind and the card should say so rather than treat it as a trend: it coincides with a large front-running of imports ahead of announced tariffs, which inflates a quarter's deficit without any change in the underlying propensity to export dollars. Chapter 16 examines that episode directly; here it is enough to note that a single quarter's current account is a poor measure of a structural channel.

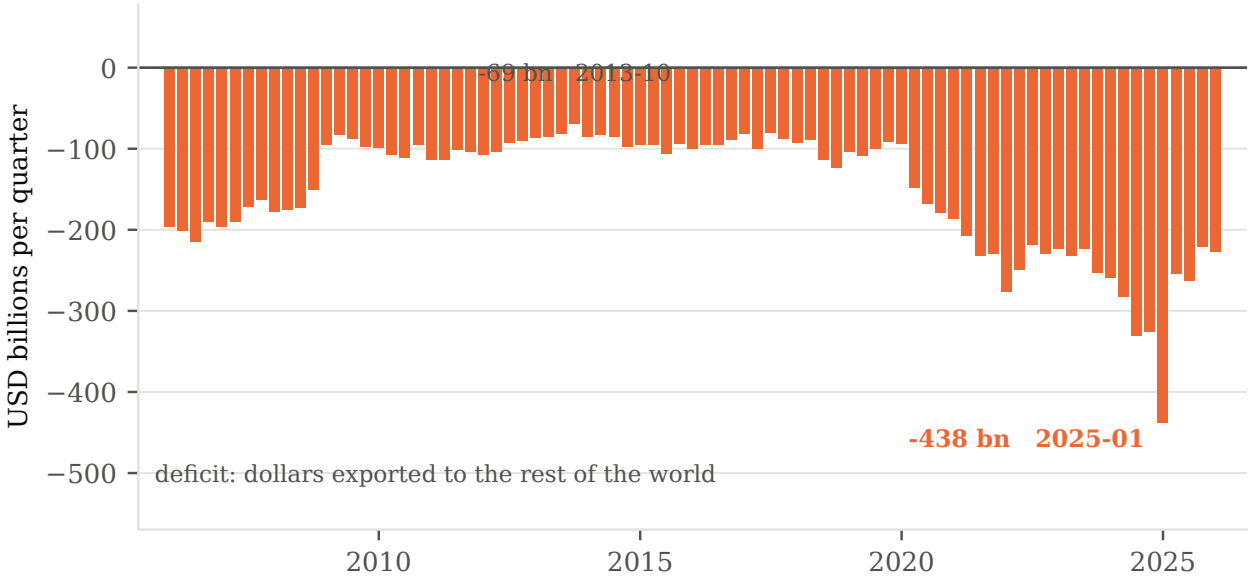


Figure 4.5: S_{ext}: the U.S. balance on current account, quarterly, 2006-Q1 to 2026-Q1. All observations are deficits; a larger deficit is a larger export of dollars. Source: FRED `IEABC` (BEA); author's calculations; script `f15_sext.py`.

Limitations. The current account is net, and net flows are a small residual of gross flows that are orders of magnitude larger. It is quarterly and heavily revised. And it measures dollars leaving the United States, not dollars available offshore, which differ by everything that happens to them afterwards.

4.7 Synthetic USD Supply (S_{crypto})

src p.10

Definition 4.7. S_{crypto} is dollar-like purchasing power created in decentralised finance without dollar backing: over-collateralised lending protocols and decentralised stablecoins. It is an *emergent channel* entering as $1 + S_{\text{crypto}}/\bar{X}$ (ISSUES #29), and it is DeFi-only (ISSUES #5).

Mechanism. While dollar-referencing stablecoins generate demand for real dollars — the D_{net} channel of Section 3.8 — the decentralised-finance ecosystem also creates synthetic dollars. Lending protocols and decentralised stablecoins let users mint dollar-like purchasing power against crypto collateral, with no dollar behind it. This shadow supply lies outside the reach of S_{Fed} and S_{bank}: it is a parallel and riskier source of effective dollars.

Role in the equation. S_{crypto} enters the denominator with elasticity $\beta_{\text{crypto}}(\sigma_{\text{sys}})$ in the emergent form, and $S_{\text{crypto}} \uparrow \Rightarrow DXY \downarrow$. Its scale relative to the other six channels is small enough that the honest expectation is a very small $\beta_{\text{crypto}}(\sigma_{\text{sys}})$; the card exists because the mechanism is structurally novel, not because it is currently large.

Interactions. S_{crypto} is the supply-side twin of D_{net} and the two must not be confused (Section 3.11.2): a dollar stablecoin backed by Treasury bills is D_{net}; a decentralised stablecoin minted against crypto collateral is S_{crypto}. The distinction is the backing, not the label, and it is exactly the distinction Gorton and Zhang [42] argue is doing the regulatory work.

Table 4.8: Measurement protocol for S_{crypto}.

	Primary	Secondary
Proxy	total value locked in DeFi lending protocols	supply of decentralised stablecoins (e.g. DAI)
Source	DefiLlama	DefiLlama; protocol dashboards
Frequency	daily	daily
Transformation	$1 + X/\bar{X}$ (emergent channel)	same
Availability	free	free
Weakness	TVL counts collateral posted, not synthetic dollars created; the two differ by the collateralisation ratio	narrower and cleaner, but a small fraction of the concept

Measurement.

Historical range. The primary proxy is in hand. DefiLlama total value locked runs daily from 2018-02-12 to 2026-09-05, 3,128 observations, from effectively zero to $87.4 billion, with a sample maximum of $177.5 billion on 2021-11-09 — a peak twice the current level, five years ago. Unlike every other channel in this chapter, S_{crypto} is smaller today than at its high.

No separate figure is drawn: the series is plotted against its demand-side twin in Figure 3.6, because a dollar stablecoin and the collateral behind synthetic dollars are one object read in two directions, and separating them into two figures would suggest two measurements where there is one.

The scale constant. Per ISSUES #29 the emergent form needs $\bar{X}$ stated here. The convention is the sample mean over the period the series exists, $\bar{X} = \$58.6$ billion, so the channel enters as $1 + S_{\text{crypto}}/58.6$ with S_{crypto} in billions. The caution attached to D_{net}'s constant applies with more force, because this series is not merely trending but has a large peak in the middle of it: the sample mean is sensitive to whether 2021 is in the window.

The channel contracts violently, and that is the finding. The four largest ninety-day contractions in the sample, restricted to windows opening above $10 billion, are $162.9 to $49.6 billion between 2022-04-04 and 2022-06-19 (−69.6 per cent), $122.2 to $71.8 billion between 2021-05-12 and 2021-05-24 (−41.2 per cent), and two further falls above thirty-five per cent in 2024–2025. The sharpest is concentrated in four days: total value locked went from $126.0 billion on 2022-05-09 to $78.5 billion on 2022-05-13, the collapse of the Terra/UST stablecoin.

This is what the card's sign predicts, at a magnitude no other supply channel reaches. $S_{\text{crypto}} \downarrow \Rightarrow DXY \uparrow$: the supply of synthetic dollars collapsed by more than a third in four days, and it collapsed *into* the 2022 tightening, the period the canonical σ_{sys} of Chapter 5 now ranks as the most stressed in the sample. The mechanism and the timing line up. What cannot be claimed is that this moved the dollar index: $47 billion is a rounding error against the channels of Section 4.2–Section 4.4, and the honest reading is that Table 4.8's expectation of a very small $\beta_{\text{crypto}}(\sigma_{\text{sys}})$ survives its first contact with the data.

Coverage against the crises the model cares about. The series begins in 2018-02, so 2008 is outside it entirely. In March 2020 it existed and fell hard in percentage terms — $1.03 billion on 2020-02-19 to $0.48 billion on 2020-03-23 — but half a billion dollars is not a monetary event. **2022 is the first crisis at which this channel is both measured and material, and it is the only one.** A single episode cannot identify $\beta_{\text{crypto}}(\sigma_{\text{sys}})$, and Chapter 11 must record the channel as unestimable rather than estimate it from one observation.

Limitations. Total value locked is a measure of collateral, not of dollars created, and the ratio between them moves with risk appetite — so the proxy is most misleading in stress, when the channel matters. The series is short. And a synthetic dollar that cannot settle an obligation to a bank is not the same object as a dollar, which the model's aggregation implicitly assumes it is.

4.8 Sanctions and Frozen Assets (S_{off})

The original uses S_{off} throughout its case studies and its policy chapters but never places it in the equation. Two decisions place it here.

What it is. S_{off} is the *immobilisation intensity*: the share of externally held dollar claims blocked by asset freezes and secondary sanctions, normalised to $[0, 1]$ like every other latent modulator (ISSUES #26). It was originally conceived as a dollar stock of frozen assets; that definition is superseded, because a stock cannot be subtracted from S_{ext}, which is a flow (ISSUES #28).

Where it enters. Through the external channel, multiplicatively:

$$S_{\text{ext}}^{\text{eff}} \;=\; S_{\text{ext}}\, e^{-\delta_{\text{off}} S_{\text{off}}}, \qquad S_{\text{off}} \in [0,1],\ \delta_{\text{off}} > 0, \tag{4.1}$$

so that in logs the channel contributes $\ln S_{\text{ext}} - \delta_{\text{off}} S_{\text{off}}$. Three properties follow, and all three were the reason for choosing this form over the subtraction the original implies (ISSUES #4 as amended by ISSUES #28).

1. **It is defined everywhere.** $S_{\text{ext}} - S_{\text{off}}$ goes negative whenever a freeze exceeds a quarter's current-account flow, which is not a remote case; the exponential form cannot.
2. **The sign convention survives.** Freezing dollars reduces effective external supply, so $S_{\text{off}} \uparrow \Rightarrow S_{\text{ext}}^{\text{eff}} \downarrow \Rightarrow DXY \uparrow$, with no exception to the book's rule that every S_j is dollar-negative.

3. **The case-study arithmetic changes form.** The original's ninth case study reads the sanctions channel as $\Delta \ln S_{\text{off}} > 0$. Under Equation (4.1) the correct object is ΔS_{off}, a change in intensity, and Chapter 14 restates it rather than reproducing it.

The long-run channel is separate. Freezing reserves does two things on two timescales. In the short run it removes dollars from usable supply, which is Equation (4.1). In the long run it demonstrates to every reserve manager that dollar reserves are contingent on political alignment, which raises T_{frag} and feeds both Φ_{sys} in Chapter 6 and $\Psi_{\text{str}}(t)$ in Chapter 9. The two must be kept apart: the first is a supply effect measured in quarters, the second a trust effect measured in decades. Farrell and Newman [33] give the theory of why network centrality converts to coercive capacity; Mulder [59] supplies the historical record that counter-mobilisation is the usual response.

Measurement.

S_{off} ratchets in the model and not in the world. ISSUES #28 made S_{off} an immobilisation *intensity* in $[0, 1]$, which is the right object and which the arithmetic allows to fall. But every causal arrow in the book raises it, and none lowers it: S_{off} appears as a consequence of coercion and a cause of T_{frag}, never as something that is relaxed.

The world does not oblige. On 29 January 2026, four weeks after the Venezuela intervention, OFAC General Licences 46 and 47 authorised broad transactions in Venezuelan-origin oil, and PDVSA output moved back toward a million barrels a day. S_{off} fell and D_{trans} rose. The model has no category for that sequence — coercion that ends coercion — and ISSUES #34 records the gap rather than repairing it, because a reversal term would attach two new parameters to a variable with no series. Section 16.1.7 scores the episode and Section 6.5 states the consequence for T_{frag}. Any case study using S_{off} must therefore treat its monotone path as an artefact of the specification, not as a finding.

> **Not supplied, and the gap is the world's.** No provider publishes an immobilisation intensity. Building one means estimating frozen official and private dollar claims as a share of externally held dollar claims — OFAC designations and the Atlantic Council tracker for the numerator, COFER and BIS for the denominator — and there is no canonical way to do it. *What it costs:* the attenuation term $e^{-\delta_{\text{off}} S_{\text{off}}}$ is inert, and S_{off} carries **no scoreable verdict anywhere**: one imposed under Rule 5, one recorded, one unmodelled, one untestable.

The card states the construction rather than a number. Note that the numerator is genuinely hard: the headline figure usually quoted for the 2022 freeze of Russian official reserves is an estimate, and the share of it that is dollar-denominated rather than euro-denominated is a further estimate.

4.9 Shadow-Bank USD Creation (S_{shadow})

The original uses S_{shadow} in two incompatible senses: in its first case study it is shadow-bank and repo dollar creation collapsing in 2008; in its sixth it is written "$S_{\text{crypto}}/S_{\text{shadow}}$"

for decentralised synthetic dollars. These are different plumbing with different data, and conflating them makes the two case studies uncomparable.

Resolution (ISSUES #5). S_{shadow} is **non-bank private dollar creation** — repo rehypothecation, money-market fund chains, the securities-financing network mapped by Pozsar [62]. It is a **named component of** S_{coll}, not an eighth denominator term. S_{crypto} is DeFi-only. The sixth case study's slash resolves to S_{crypto}; the first case study's 2008 collapse reads as $S_{\text{coll}} \downarrow$ with the shadow leg named.

Why a component and not a term. Adding an eighth term would give S_{shadow} its own elasticity, which sounds like more information and is in fact less: the only free data for it — the Financial Stability Board's non-bank financial intermediation monitor — is annual, and an annual series cannot identify an elasticity in a model whose other supply channels move daily. Folding it into S_{coll} keeps the mechanism visible in the prose while leaving the estimation to the series that can actually support it.

Measurement.

> **Not supplied, and the gap is ours.** OFR repo-market data and the FSB's NBFI monitor are both free and unsupplied. Both are coarser than the channel. *What it costs:* S_{shadow} is folded into S_{coll} by ISSUES #5 and scored untestable wherever it appears, including in Table 13.3, the largest dollar rally in the sample.

No range is stated. The component is named in S_{coll}'s mechanism and is not separately estimated.

4.10 The Effective-Supply Identity and Double-Counting Checks

4.10.1 The identity

Collecting the seven channels, the denominator of Equation (2.1) is

$$\begin{aligned}\mathcal{D}(\sigma_{\text{sys}}) \;=\; & S_{\text{Fed}}^{\beta_{\text{Fed}}(\sigma_{\text{sys}})}\; S_{\text{Tsy}}^{\beta_{\text{Tsy}}(\sigma_{\text{sys}})}\; S_{\text{coll}}^{\beta_{\text{coll}}(\sigma_{\text{sys}})}\; S_{\text{bank}}^{\beta_{\text{bank}}(\sigma_{\text{sys}})} \\ & \times\; \Big(1+\tfrac{S_{\text{swap}}}{X}\Big)^{\beta_{\text{swap}}(\sigma_{\text{sys}})} \Big(S_{\text{ext}} e^{-\delta_{\text{off}} S_{\text{off}}}\Big)^{\beta_{\text{ext}}(\sigma_{\text{sys}})} \Big(1+\tfrac{S_{\text{crypto}}}{X}\Big)^{\beta_{\text{crypto}}(\sigma_{\text{sys}})},\end{aligned} \tag{4.2}$$

with S_{shadow} inside S_{coll} and S_{off} inside the external factor. Seven terms, as in the original; two of them in emergent form; one of them attenuated.

4.10.2 Four double-counting checks

A multiplicative aggregator punishes double-counting: a quantity entered twice contributes twice to $\ln DXY$ and biases both elasticities. Four pairs need checking, and Table 4.9 states the resolution of each.

Table 4.9: Double-counting checks on the denominator. "Overlap" is the quantity that could be counted twice; the resolution is what keeps it counted once.

Pair	Overlap	Resolution
S_{Fed} vs S_{Tsy}	both are balances at the Federal Reserve; a TGA drawdown raises reserves with no Fed action	S_{Fed} is measured by *assets*, S_{Tsy} by two *liabilities*. Assets minus those liabilities is the reserve balance, so the two are complements, not duplicates.
S_{Tsy} vs S_{coll}	money-market cash at the Fed is cash not financing repo	S_{Tsy} measures the *quantity* parked; S_{coll} measures the *price* of what remains. Same phenomenon, two observables, no shared series.
S_{Fed} vs S_{coll}	quantitative easing removes collateral from the private market	Real and intended: the sovereign channel raises supply while lowering collateral availability. Both entries are correct and the net sign is empirical.
S_{ext} vs D_{trans}	the same trade both requires dollars and delivers them	Different populations and speeds: D_{trans} is the importer's settlement demand, S_{ext} the exporter's receipt. The model counts a gross flow on each side rather than netting them.

The third row is worth dwelling on because it is not a bookkeeping nicety. Quantitative easing raises S_{Fed} and lowers S_{coll} simultaneously, and the two enter Equation (4.2) with the same sign of elasticity. The net effect on effective supply is therefore ambiguous in the model as specified, and deliberately so: whether QE relieves or worsens a dollar shortage depends on whether the reserve injection or the collateral withdrawal dominates, which is exactly the question the September 2019 episode posed [24]. A model that resolved it by construction would be assuming the answer.

4.10.3 What the denominator cannot currently do

Three of the seven channels are unmeasured or partly measured: S_{bank} entirely, S_{coll} from 2021 only, S_{crypto} not at all. That is a materially worse position than the numerator, where six of ten channels have usable series.

The consequence for Chapter 11 is specific rather than general. The denominator's identified variation over the full sample comes from S_{Fed}, S_{Tsy}, S_{swap} and S_{ext} — the four sovereign and external channels. The three private-plumbing channels, which are precisely the ones the model claims are decisive in crises, are the ones without data. Any estimate of S^{*} therefore measures the part of supply that governments control and is close to silent on the part that markets control. This is stated here so that Chapter 11 does not have to discover it, and so that no reader mistakes a well-fitting denominator for a complete one.

Chapter 5

Layer 1 — The State and Friction Layer: σ_{SYS}

5.1 From Thermometer to Systemic Instability Index

The original states the problem exactly: the stress variable of E1, proxied by the VIX, is no longer sufficient — it is merely a thermometer. What the model needs is a *systemic instability index* capturing both non-financial shocks and the inter-market friction that determines whether a shock propagates. In the original's notation, $\sigma_{\text{sys}} = f(\sigma_{\text{VIX/MOVE}}, X_{\text{pand}}, F_{\text{liq}})$.

A thermometer measures how hot the patient is. It does not say what set the fever, and it does not say whether the infection will spread. Layer 1 adds both: X_{pand} is the spark that need not be financial in origin, and F_{liq} is the viscosity of the system that turns a spark into a contagion.

This chapter is short because its job is narrow. It has to produce one number, bounded in $[0, 1]$, that every elasticity in Chapter 2 and the entire safe-haven exponential depend on. Getting that number wrong contaminates everything above it, which is why the construction is set out in full and its failures are stated rather than smoothed.

5.1.1 Two printed forms, one canonical

The original gives E3 twice. At §4.1.3 it is $\sigma_{\text{sys}} \approx (\sigma_{\text{VIX/MOVE}} + w_X X_{\text{pand}})(1 + w_F F_{\text{liq}})$; at §5.2.1 a weight appears on the base term, $\sigma_{\text{sys}} = (w_\sigma \sigma_{\text{VIX/MOVE}} + w_X X_{\text{pand}})(1 + w_F F_{\text{liq}})$. We adopt the second as canonical, with equality rather than approximation (ISSUES #19):

$$C_t \;=\; \Big(w_\sigma\, \sigma_{\text{VIX/MOVE}} + w_X\, X_{\text{pand}}\Big)\Big(1 + w_F\, F_{\text{liq}}\Big). \tag{5.1}$$

C_t is the raw composite. It is not yet σ_{sys}, and the difference is the subject of Section 5.5.

The structure of Equation (5.1) carries a claim worth making explicit: friction enters *multiplicatively* while the shock sources enter additively. Financial stress and an exogenous shock are alternative ways of starting a fire; liquidity friction is not a third way of starting one but a property of the building. A model in which F_{liq} were additive would allow a liquid system with high friction to be as unstable as an illiquid one, which is not what the original claims and not what March 2020 showed.

5.2 Endogenous Financial Stress ($\sigma_{\text{VIX/MOVE}}$)

src p.11

Definition 5.1. $\sigma_{\text{VIX/MOVE}}$ is the base component of systemic stress: the fear priced into equity and bond option markets. It has two legs, an equity leg and a rates leg.

Mechanism. This is the smoke rising from the two markets that matter most for dollar funding. Equity implied volatility measures the price of insurance against a fall in risk assets; rates implied volatility measures the price of insurance against a disorderly move in the risk-free curve. The second is the more directly relevant to this model — Treasuries are the collateral of the dollar system — and it is the one that is not free.

Measurement, and the MOVE substitution. The equity leg is the VIX, in hand and free. The rates leg is the ICE BofA MOVE index, which is paywalled. Per ISSUES #23 the card states a substitute rather than leaving the component undefined.

Table 5.1: Measurement protocol for $\sigma_{\text{VIX/MOVE}}$. The rates leg is a substitution and the table states what the substitution costs.

	Equity leg	Rates leg
Concept	implied volatility of the S&P 500	implied volatility of the Treasury curve
Primary	VIX	**MOVE — paywalled**
Used here	FRED `VIXCLS`	**21-day realised volatility of daily changes in `DGS10`**
Frequency	daily	daily
Transformation	empirical percentile rank in $[0, 1]$	same
Cost of the substitution	—	realised, not implied: it is backward-looking over 21 days, so it *lags* the market's forward-looking fear and understates the first days of a shock

The substitution is defensible and it is not free. Realised volatility over a trailing window rises only after the moves have happened, whereas MOVE rises when traders expect them. In a fast episode — and every episode this model cares about is fast — the substitute is late. Chapter 11 should test the sensitivity of any threshold estimate to the window length, and where a case study turns on the first three days of a shock, the rates leg should be read as unavailable rather than as measured.

5.3 Exogenous Shocks (X_{pand})

src p.12

Definition 5.2. X_{pand} is the non-financial trigger: pandemics, wars, large-scale natural disasters. Events that abruptly alter economic assumptions from outside the financial system.

Mechanism. Financial crises are not always financial in origin. A pandemic disrupts D_{trans} directly by stopping trade, and ignites D_{prec} by making the future unreadable. The original describes X_{pand} as an additive or threshold trigger for $\sigma_{\text{VIX/MOVE}}$ — the initial spark. The additive placement in Equation (5.1) encodes that: an exogenous shock can raise systemic stress even when option markets are calm, which is precisely the first week of 2020-02.

Role in the composite. X_{pand} enters Equation (5.1) additively with weight w_X. Its distinguishing property is that it is *not* a market price and therefore not contaminated by the reflexivity of Layers 3 and 4: it is the one input to σ_{sys} that markets cannot talk themselves into.

Measurement. The primary proxy is the daily Geopolitical Risk index of Caldara and Iacoviello [22], a newspaper-based measure that spikes around wars and terrorist attacks. Secondary candidates are EM-DAT for natural disasters and WHO pandemic risk levels, both event-coded rather than continuous and neither used here.

Table 5.2: Measurement protocol for X_{pand}.

	Primary (in hand)	Secondary
Proxy	daily Geopolitical Risk index, `GPRD`	EM-DAT disaster counts; WHO pandemic risk level
Source	Caldara and Iacoviello [22], matteoiacoviello.com	EM-DAT; WHO
Frequency	daily	event
Coverage	1985-01-01 $\rightarrow$ 2026-09-01, 15,219 obs	—
Transformation	empirical percentile rank in $[0, 1]$	event dummies
Weakness	newspaper-based, so it measures *coverage* of geopolitical risk; and it is silent on pandemics, which is one of the three event types the original names	event-coded, so it cannot enter a continuous composite without an arbitrary intensity scale

Historical range. `GPRD` runs daily from 1985-01-01 to 2026-09-01, 15,219 observations, of which 7,549 fall inside the modelled sample. Within the sample it reaches its maximum of 540.8 on 2022-03-01, days after the invasion of Ukraine, and its minimum of 0.0 on 2025-02-09.

What its inclusion changes. A good deal, and Section 5.5.3 reports it rather than burying it. X_{pand} is the component that distinguishes this model's state variable from a repackaged VIX, and adding it reorders which episodes the model calls stressed. It also weakens one of the two empirical findings of Section 5.6. Both consequences are set out below.

One limitation deserves naming here. `GPRD` covers wars and terrorism; it does not cover pandemics. Of the three event types the original lists under X_{pand} — pandemics, wars, natural disasters — the proxy captures one and a half. That is why 2020 does not gain from the addition, and it is a gap in the proxy rather than in the concept.

5.4 Inter-Market Liquidity Friction (F_{liq})

src p.12

Definition 5.3. F_{liq} is the viscosity of the system: the ease with which a shock in one market propagates to others. It is the fuel that turns a spark into a contagion, and it enters Equation (5.1) multiplicatively.

Mechanism. Two markets can carry identical fear and transmit it entirely differently. What distinguishes them is whether positions can be closed: whether the bid is there, whether collateral moves, whether an intermediary will take the other side. When that machinery is impaired, a shock that would have been absorbed instead propagates, and the same $\sigma_{\text{VIX/MOVE}}$ produces a much larger dollar move. This is the mechanism behind the amplification claims of Table 3.12 and it is the reason F_{liq} multiplies rather than adds.

Measurement. The ideal proxies — FRA–OIS, ETF bid-ask spreads, Treasury market depth — are paywalled or proprietary (ISSUES #23). Two free substitutes are used.

Table 5.3: Measurement protocol for F_{liq}.

	Used here	Cross-check (partial sample)
Proxy	St. Louis Fed Financial Stress Index	TED spread
Source	FRED `STLFSI4`	FRED `TEDRATE`
Frequency	weekly, carried forward to daily	daily
Coverage	2006-01-06 → 2026-08-28, 1,078 obs	2006-01-03 → **2022-01-21**, 3,942 obs
Transformation	empirical percentile rank in $[0, 1]$	same
Weakness	a broad composite of 18 series, so it partly duplicates the equity leg	**discontinued with LIBOR**; unsecured bank funding, not market liquidity

The index is documented by Kliesen and Smith [50], and the two other standard financial-conditions constructions — the Chicago Fed's NFCI [17] and its stress decomposition [18] — are built on the same principle of aggregating spreads, volatilities and quantities into one index.

The first weakness is the one to watch. `STLFSI4` contains volatility measures among its eighteen inputs, so using it as F_{liq} while using VIX as $\sigma_{\text{VIX/MOVE}}$ double-counts part of the equity leg. The duplication is partial and unavoidable with free data; its effect is to make σ_{sys} somewhat more equity-driven than Equation (5.1) intends, and to bias w_F upward if it is ever estimated rather than assumed.

5.5 Construction: From Composite to Bounded Index

5.5.1 Why the raw composite is not the state variable

Equation (5.1) produces a number with no natural scale, and ISSUES #9 established why that is fatal: with σ on a raw scale the term $e^{\lambda_2 \sigma^2_{\text{sys}}}$ explodes, and a λ_2 large enough to matter at VIX 15 is astronomical at VIX 80. The decided resolution is that $\sigma_{\text{sys}} \in [0, 1]$ *by construction*: standardise the composite and pass it through a logistic.

Two construction choices are forced by the data, and both are stated here rather than buried in a script.

Assumption 5.1 (Percentile-rank inputs)**.** Every input to Equation (5.1) enters as its empirical percentile rank in $[0, 1]$.

The reason is arithmetic. Equation (5.1) multiplies the base by $(1 + w_F F_{\text{liq}})$, and a multiplicative friction term requires non-negative inputs. Standardised scores are not non-negative: a z-score of -1.2 in a calm month, multiplied by a high friction term, would make the composite *more* negative and so report a calm-but-illiquid market as calmer than a calm-and-liquid one. Percentile ranks are non-negative by construction and preserve the multiplicative structure the original intends.

Assumption 5.2 (Provisional equal weights)**.** $w_\sigma = w_X = w_F = 1$.

These are a stated convention, not estimates. w_σ, w_X and w_F are OPEN parameters in the registry and Chapter 11 estimates them. The convention is not innocuous, and Section 5.5.3 shows exactly how much it matters: giving X_{pand} the same weight as the entire financial-stress base is a strong choice, and the empirical results of Section 5.6 depend on it.

5.5.2 The construction

$$\sigma_{\text{sys}}(t) \;=\; \Big[1 + \exp\big(-z_t\big)\Big]^{-1}, \qquad z_t \;=\; \frac{C_t - \bar{C}}{s_C}, \tag{5.2}$$

where C_t is Equation (5.1) evaluated on percentile-rank inputs and $(\bar{C}, s_C)$ are its sample mean and standard deviation. The base component is the average of the two legs of Section 5.2, $\frac{1}{2}(p_{\text{VIX}} + p_{\text{bondvol}})$.

Properties. $\sigma_{\text{sys}} \in (0, 1)$ strictly, so the safe-haven exponential of Equation (2.16) is bounded. σ_{sys} is monotone in each input. And it is *sample-dependent*: the standardisation and the percentile ranks both use the whole sample, so σ_{sys} is not a real-time series and must never be presented as one. Recomputing it on a shorter sample changes every value.

Historical range. The index is available daily from 2006-02-02 to 2026-09-03, 5,143 observations, the start date set by the 21-day window of the rates leg. It reaches 0.961 on 2008-12-01 and falls to 0.097 on 2018-10-02; the latest observation is 0.399. It stands at or above 0.8 on 523 days, 10.2 per cent of the sample, and at or above 0.9 on 96 days, 1.9 per cent.

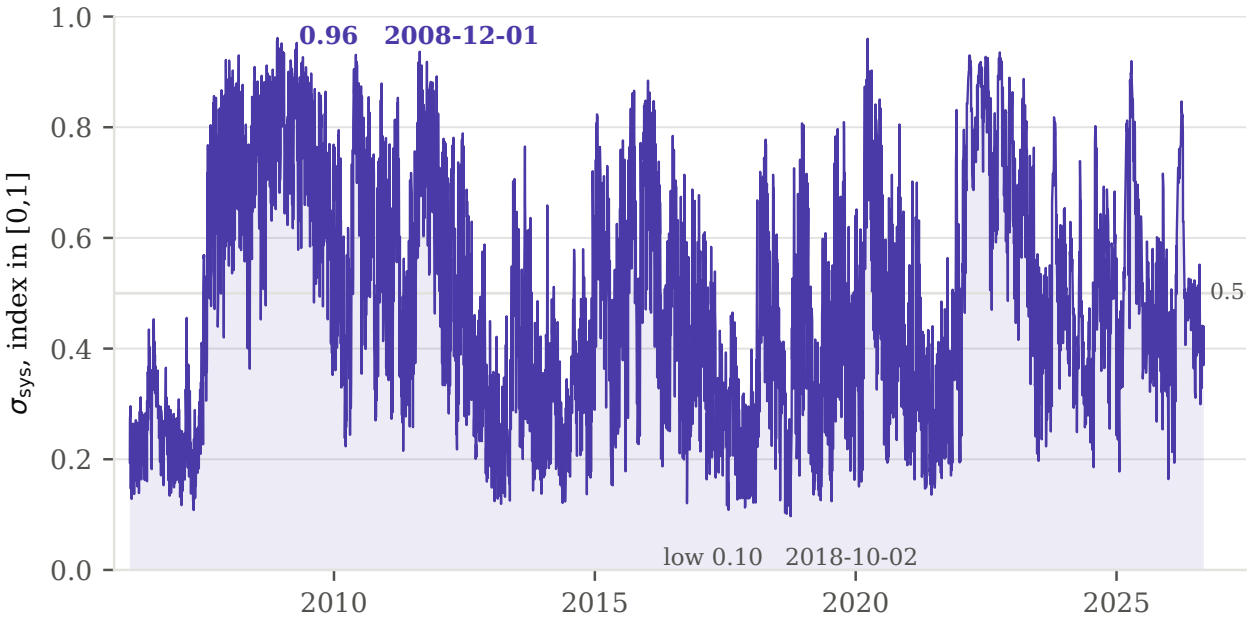


Figure 5.1: σ_{sys}, the Layer-1 systemic instability index, daily, 2006-02-02 to 2026-09-03, built from all three of E3's components: the VIX, a 21-day realised-volatility substitute for the paywalled MOVE, the Caldara–Iacoviello geopolitical risk index, and `STLFSI4` as the friction multiplier. Source: FRED `VIXCLS`, `DGS10`, `STLFSI4`; Caldara and Iacoviello [22]; author's construction; scripts `build_sigma.py` and `f17_sigma.py`.

5.5.3 What the third component changes

The index was first built without X_{pand}, because the GPR series arrived later than the financial data. Keeping both constructions turns out to be worth the extra column, because they disagree about which episodes were stressful — and, in Section 5.6, about one of the model's two claims.

Table 5.4: Where the 400 most stressed days fall, under the two constructions. The two-component index is $w_X = 0$; the canonical index is $w_X = 1$.

	Two components ($w_X = 0$)	**Canonical** ($w_X = 1$)
Most stressed year	2009 (135 days)	**2022 (109 days)**
Then	2008 (123), 2011 (67)	2008 (72), 2009 (69)
Then	2020 (31), 2022 (24)	2011 (36), 2010 (25)
Sample maximum	0.917 on 2008-10-10	0.961 on 2008-12-01
Sample minimum	0.170 on 2017-10-11	0.097 on 2018-10-02

Adding geopolitical risk moves 2022 from fifth place to first. That is the invasion of Ukraine, which drove `GPRD` to its sample maximum of 540.8 on 2022-03-01, entering a state variable that had previously seen 2022 only through moderately elevated volatility. The reordering is defensible on its face: a year containing a major European land war and the

fastest tightening cycle in four decades is a plausible candidate for the most stressed year in a twenty-year sample.

Two properties of the construction are visible in the same table and neither should be hidden.

The index rewards persistence over peak severity. 2020 places fourth under both constructions despite the VIX reaching its sample maximum of 82.69 that March. A percentile-rank composite accumulates high-σ_{sys} days, so a long moderate episode outranks a short extreme one. If the model wants a state variable that ranks by peak severity, this is the wrong construction, and Chapter 11 should say which property it wants before estimating anything on it.

The proxy for X_{pand} does not cover pandemics. As Section 5.3 notes, `GPRD` measures war and terrorism. So the component that was supposed to let Layer 1 see non-financial shocks lets it see one kind, and 2020 — the clearest non-financial shock in the sample — gains nothing from it. The gap is in the proxy, not the concept, and it is the single most valuable data addition remaining for this layer.

5.5.4 What the rank transform does to an unestimated weight

Section 5.5.3 reported that the third component reorders the episodes. This subsection reports something narrower and more uncomfortable: *why* it has that much leverage. The answer is an interaction between Assumption 5.1 and Assumption 5.2 that neither assumption discloses on its own, and it must be on the record before Section 5.6 reports any result.

The rank transform equalises the inputs by construction. A percentile rank maps any series, whatever its units, scale, skew or tail behaviour, onto the uniform distribution on $[0, 1]$. That is what makes Assumption 5.1 legitimate — it is the only transform that lets a volatility index, a realised-volatility substitute, a newspaper-based count and a financial-conditions index enter one multiplicative expression without one of them dominating through its units. But it has a consequence that the assumption did not state: **after ranking, the weights are the only thing left that determines relative influence.** There is no variance to scale, no dispersion to attenuate a noisy input, and no data-driven force that shrinks a component carrying little information. A component's contribution to C_t is exactly its weight. With unranked inputs a poorly behaved series would at least announce itself through its scale; ranking removes that signal along with the units.

Under equal weights, one newspaper count outweighs the entire market block. Written out, Equation (5.1) under Assumption 5.2 is

$$C_t \;=\; \Big[\; \underbrace{\tfrac{1}{2}p_{\text{VIX}} + \tfrac{1}{2}p_{\text{bondvol}}}_{\text{market-priced, weight } w_\sigma=1 \text{ split two ways}} \;+\; \underbrace{p_{\text{GPRD}}}_{\text{news count, weight } w_X=1} \;\Big] \Big(1 + w_F\, p_{\text{STLFSI}}\Big). \tag{5.3}$$

The two market components share the base and so carry $\frac{1}{2}$ each; X_{pand} carries a full 1. Before the friction multiplier, half of the composite is a count of newspaper articles. That

is not a criticism of Caldara and Iacoviello [22], whose index does what it says it does. It is a statement about what the equal-weight convention means once the rank transform has stripped the scales: $w_\sigma = w_X = 1$ is not a neutral default but a strong claim that geopolitical news deserves the same weight as every market-priced input combined.

The series is of a different kind, and the sample says so. The three ranked inputs are not variants of one measurement. Table 5.5 sets out the difference on the sample itself.

Table 5.5: The ranked inputs to Equation (5.1) are not the same kind of series. Daily, 2006-02-02 to 2026-09-03, 5,143 observations, computed on the percentile ranks that actually enter the composite. "Daily rank move" is the mean absolute one-day change in the rank.

	p_{VIX}	p_{bondvol}	p_{STLFSI}	p_{GPRD}
Kind	market price	market realised	market composite	**news count**
AR(1)	0.973	0.987	0.986	**0.503**
AR(21)	0.742	0.687	0.828	**0.299**
Daily rank move	0.045	0.027	0.015	**0.200**
Corr. with p_{VIX}	—	+0.496	+0.639	**−0.100**

Three readings follow, and all three matter for what Section 5.6 can claim.

First, the persistence is an order of magnitude apart. The market inputs have first-order autocorrelations above 0.97 and retain most of it at a month's lag; p_{GPRD} has 0.503 at one day and 0.299 at twenty-one. Its rank moves twenty percentile points on the average day, four to thirteen times the others. A *state* variable, in the sense Section 5.1 defines — the condition of the system, which the regime elasticities of Equation (2.9) switch on — is supposed to be slow. One third of this one moves at the frequency of headlines, and under Assumption 5.2 that third is half the pre-friction composite.

Second, it is not measuring the same latent factor. The two market legs correlate at +0.496; p_{GPRD} correlates with the VIX leg at −0.100 — essentially zero, and if anything the wrong sign. Adding it is therefore not adding a third reading of systemic stress. It is averaging in a nearly orthogonal series. On 663 days the geopolitical rank is at or above 0.90; on 48 per cent of those days the VIX rank is below 0.50. On 514 days the VIX rank is at or above 0.90; on 50 per cent of those the geopolitical rank is below 0.50. The two components disagree about which days were stressful roughly half the time they each register stress.

Third, the direction of the disagreement is the mechanism behind Proposition 5.2. A component that is high on calm-market days pushes calm days up the σ_{sys} distribution; a component that is low on high-VIX days pulls stressed days down. Both effects compress the range of σ_{sys} against the financial variable the dollar actually responds to, and compression of the regressor's regime split is precisely what would attenuate a measured difference between calm and stressed slopes. This is not a proof that the amplification result of Section 5.6.4 fails *because* of the weighting — Section 5.6.4 sets out the competing readings and does not adjudicate between them — but it identifies the channel through which the weighting could produce exactly the result observed.

What this obliges the book to do. w_X is an OPEN parameter. It has not been estimated, and until it is, every empirical result in this chapter is reported at a value of w_X chosen by convention rather than by data. Chapter 11 must therefore estimate w_σ, w_X and w_F rather than inheriting Assumption 5.2. Section 5.6.5 does the next best thing available without an estimate: it reports both results of Section 5.6 as a function of w_X across the whole interval $[0, 1]$, so that whatever Chapter 11 estimates can be read off against a profile fixed in advance rather than argued about afterwards. A sensible prior is that the estimate lands well below 1: a series with an AR(1) of 0.5 and no correlation to the rest of the composite should not carry the same weight as the market block. But that is a conjecture, and the whole point of Section 5.6.4 is that this book does not settle empirical questions by conjecture.

5.6 Regime Thresholds and the Empirical Test

5.6.1 What the model actually claims

Section 1.3 set out the claim that motivates the whole book: at low stress the dollar behaves like a risk asset and is sold; above a threshold it flips and becomes a safe haven, responding convexly. In the notation of Chapter 2 this is two assertions, and they are separable:

1. **Amplification.** The dollar's response to a given change in stress is larger when the level of stress is already high. In E10b's terms, $\lambda_2 > 0$ and κ_{sys} rises with σ_{sys}.
2. **Sign change.** The response is *negative* below a threshold σ^* and *positive* above it.

The second is the stronger claim and the one that gives the book its title. Testing them requires care about what is regressed on what. The model's statement is contemporaneous — a change in stress moves the dollar within the period — so the test is $\Delta \ln DXY$ on $\Delta\sigma_{\text{sys}}$ within the same period, with the regime determined by the level of σ_{sys} at the *start* of the period so that the split is not mechanically related to the change being measured.

5.6.2 The result

Table 5.6: Contemporaneous response of the dollar to a change in systemic stress, by regime, under both constructions of the state variable. Monthly; dependent variable `DTWEXAFEGS`; slopes in per cent per unit of σ_{sys}; regime set by σ_{sys} at the start of the month.

Regime	n	slope	t	
Canonical, three components ($w_X = 1$)				
Calm start, $\sigma_{\text{sys}} < 0.5$	126	+1.73	+1.96	positive, at the margin of significance
Stressed start, $\sigma_{\text{sys}} \geq 0.5$	120	+1.89	+1.66	barely steeper
Two components ($w_X = 0$), for comparison				
Calm start, $\sigma_{\text{sys}} < 0.5$	137	+2.81	+2.59	positive, not negative
Stressed start, $\sigma_{\text{sys}} \geq 0.5$	109	+6.34	+3.54	more than twice as steep

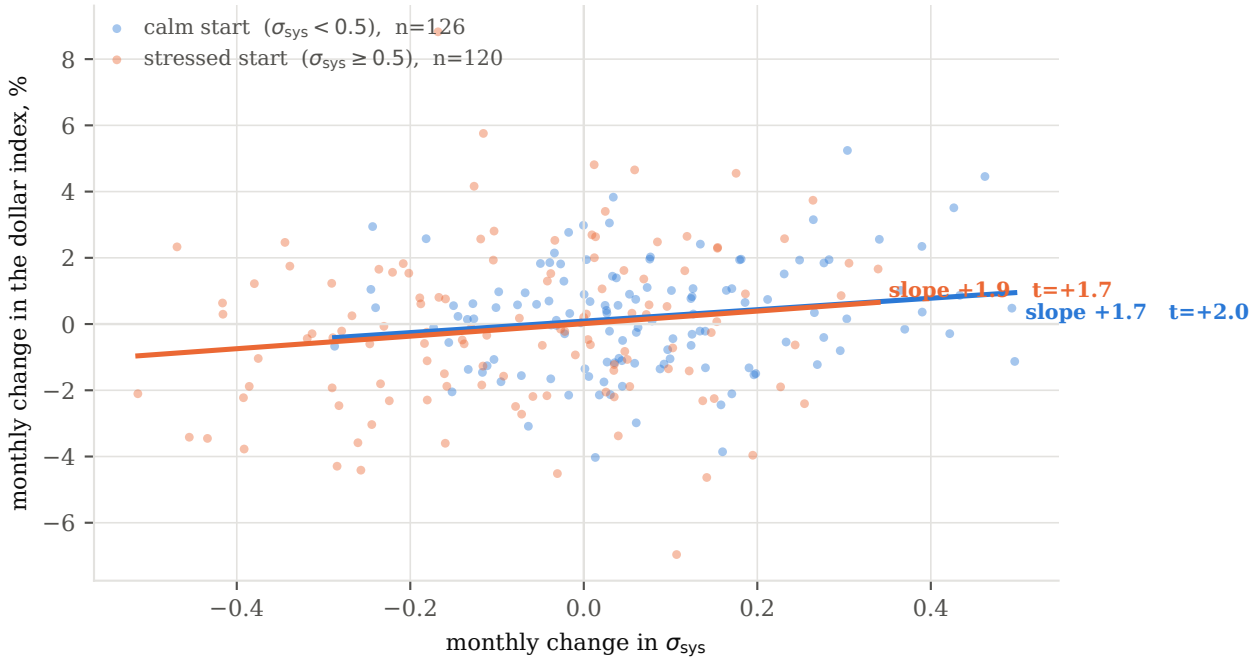


Figure 5.2: The safe-haven response by regime under the canonical three-component state variable: monthly change in the dollar index against the contemporaneous monthly change in σ_{sys}, split by the level of σ_{sys} at the start of the month. The two fitted lines are nearly parallel, which is the finding of Section 5.6.4. Source: FRED `DTWEXAFEGS`; σ_{sys} as constructed above; script `f18_smile.py`.

5.6.3 Robustness

The exercise is deliberately bounded: it varies the frequency, the dependent variable, the threshold and the state-variable construction, and it changes nothing else. The grid was fixed in advance and every cell is reported. It comprises three frequencies (monthly, weekly, daily), two dependent variables (`DTWEXAFEGS` and the broad `DTWEXBGS`), thirteen candidate thresholds from 0.20 to 0.80, and both constructions of σ_{sys}: 150 specifications in total, subject to a minimum of 30 observations per regime.

Table 5.7: The sign-change test across frequency and dependent variable, threshold fixed at $\sigma_{\text{sys}} = 0.5$, under both constructions. Slopes are per cent per unit of σ_{sys}.

		calm		stressed	
Freq.	Index	slope	t	slope	t
Canonical, three components					
monthly	AFE	+1.73	+1.96	+1.89	+1.66
monthly	broad	+1.85	+2.61	+2.12	+2.20
weekly	AFE	+0.04	+0.15	+0.57	+1.60
weekly	broad	+0.27	+1.41	+0.84	+2.84
daily	AFE	+0.05	+0.87	+0.02	+0.23
daily	broad	+0.06	+1.40	+0.13	+1.90
Two components					
monthly	AFE	+2.81	+2.59	+6.34	+3.54
monthly	broad	+3.78	+4.34	+6.03	+4.07
weekly	AFE	+0.77	+1.71	+2.40	+3.26
weekly	broad	+1.87	+5.46	+3.19	+5.32
daily	AFE	+0.81	+3.88	+1.24	+3.25
daily	broad	+1.62	+9.95	+1.96	+6.72

The counts. Across all 150 specifications:

- the calm-regime slope is **positive in 148 of 150**;
- **two** specifications produce a negative calm-regime slope, both at the extreme threshold 0.20 on daily advanced-economy data — -1.08 with $t = -0.79$ under the two-component index, and -0.09 with $t = -0.50$ under the canonical one;
- **no specification, under either construction, produces a negative calm-regime slope that is statistically distinguishable from zero**;
- the stressed slope exceeds the calm slope in 62 of 76 two-component specifications (82 per cent) but in only **34 of 74** canonical ones (**46 per cent**).

The emerging-market caveat runs the wrong way. It is natural to suspect that the advanced-economy index misses the risk-on leg because it under-weights the carry currencies through which that leg is usually told (Section 3.1). It does not: the broad index produces a *larger* positive calm-regime slope in every one of the twelve frequency-by-construction cells.

Including emerging markets makes the dollar look *more* like a safe haven in calm conditions, not less. That caveat is withdrawn rather than left standing.

5.6.4 What survives, and what does not

Proposition 5.1 (Sign change, not supported)**.** *On 2006–2026 data there is no regime in which the dollar's response to rising systemic stress is negative. This holds across three frequencies, two dollar indices, thirteen candidate thresholds and both constructions of the state variable: 148 of 150 specifications give a positive calm-regime slope, and none of the two negative cells is distinguishable from zero. The dollar is a safe haven at every level of* σ_{sys}*.*

Proposition 5.2 (Amplification, conditional on w_X)**.** *Whether the response is* larger *in the stressed regime cannot be settled on these data, because it is conditional on the weight given to non-financial stress in the state variable, and that weight has not been estimated. Across the sweep of Section 5.6.5, the share of specifications with a steeper stressed slope falls monotonically from 82 per cent at* $w_X = 0$ *to 46 per cent at* $w_X = 1$*, and the monthly amplification ratio falls from* 2.25 *to* 1.09*. The claim holds for a majority of specifications at every* $w_X \leq 0.75$ *and is lost at* $w_X = 1$*. It is therefore neither supported nor refuted here: it is* ***undetermined pending the estimation of*** w_X*, which is Chapter 11's task.*

Proposition 5.2 is the more uncomfortable of the two and it was not the expected result. Three readings are available.

1. **The amplification was a financial-stress artefact.** If the dollar's convexity is a property of funding stress specifically, then a state variable diluted with geopolitical news will show less of it, and the two-component result was the informative one about a narrower question.
2. **Equal weighting over-weights** X_{pand}**.** $w_X = 1$ gives a newspaper-based geopolitical index the same weight as the entire financial base (Equation (5.3)). If the true w_X is much smaller, the canonical index is mis-specified and the amplification is real.
3. **The amplification is not there.** The two-component result may simply have been the fragile one, and adding a component that is uncorrelated with dollar moves reveals it.

The next subsection cannot decide between them — only an estimated w_X can — but it can establish which of the three the geometry of the problem is consistent with, and it narrows the question from "does the amplification exist?" to "at what w_X?".

5.6.5 The amplification result as a function of w_X

Section 5.5.4 argued that after the percentile-rank transform the weights are the only thing setting relative influence, and that the equal-weight convention therefore does more work than a convention should. The claim can be checked directly: rebuild σ_{sys} at each of six values of w_X, holding $w_\sigma = w_F = 1$, and re-run the entire robustness grid of Section 5.6.3 on each. Nothing else changes — same frequencies, same two indices, same thirteen thresholds, same minimum of thirty observations per regime, the grid fixed in advance. This is 452

specifications in total. The script is `data/wx_sweep.py` and every cell it produces is reported in Table 5.8.

Table 5.8: The two Chapter 5 results as a function of the geopolitical weight w_X, with $w_\sigma = w_F = 1$ throughout. "Steeper" counts specifications in which the stressed-regime slope exceeds the calm-regime slope. The headline specification is monthly, `DTWEXAFEGS`, threshold 0.50; "ratio" is its stressed slope divided by its calm slope. "Neg." and "sig." are the sign-change test: cells with a negative calm slope, and cells where that slope is also distinguishable from zero. "Verdict" summarises the amplification column only: *clear* where a steeper stressed slope holds in seventy per cent or more of specifications, *coin toss* at a share near one half, *absent* at a minority share.

		amplification		headline spec.			calm slope < 0		
w_X	specs	steeper	share	calm	stressed	ratio	neg.	sig.	verdict
0.00	76	62	82%	+2.81	+6.34	2.25	1	**0**	clear
0.10	76	64	84%	+3.06	+5.51	1.80	2	**0**	clear
0.25	76	56	74%	+2.86	+5.06	1.77	1	**0**	clear
0.50	76	53	70%	+2.25	+3.85	1.71	0	**0**	clear
0.75	74	38	51%	+2.10	+2.57	1.22	0	**0**	coin toss
1.00	74	34	46%	+1.73	+1.89	1.09	1	**0**	absent

(a) The amplification decays monotonically in w_X, and dies between 0.75 and 1. The share of specifications with a steeper stressed slope is 82, 84, 74, 70, 51 and 46 per cent as w_X runs from 0 to 1. The single non-monotonicity is between $w_X = 0$ and $w_X = 0.1$, where the share rises by two specifications; from 0.1 onward the decay is strict. The headline amplification ratio behaves the same way and more smoothly, falling from 2.25 through 1.80, 1.77, 1.71 and 1.22 to 1.09.

The direct answer to "at what w_X does the finding disappear?" is **between** 0.75 **and** 1.00, where the majority is lost. Two qualifications belong with that number and they pull in opposite directions, so both are stated.

Against the finding: a share of 51 per cent is a coin toss. By the standard the rest of this chapter applies, the amplification is already uninformative at $w_X = 0.75$ and is unambiguous only at $w_X \leq 0.5$, where it holds in seventy per cent or more of specifications and the stressed slope is at least 1.7 times the calm one.

For the finding: the collapse in the *share* is not matched by a collapse in statistical significance at the headline specification. Both regime slopes there remain individually significant at every $w_X \leq 0.75$ — the stressed slope has $t = +2.04$ at $w_X = 0.75$ — and it is only at $w_X = 1$ that the stressed slope loses significance ($t = +1.66$) and the calm slope falls to the margin ($t = +1.96$). What decays across the sweep is the *gap* between the two regimes, not the response itself; the dollar answers to σ_{sys} throughout, and it is the difference between how it answers in calm and stressed conditions that thins out.

(b) The sign-change result is completely insensitive to w_X. Across all 452 specifications in the sweep — six values of w_X, three frequencies, two dollar indices, thirteen thresholds — the number producing a calm-regime slope that is negative *and* statistically distinguishable from zero is **zero**. The raw count of negative cells is five, never more than two at any w_X, and none survives a t-test. Nothing about the weighting rescues the sign change.

This is the more important half of the table. Proposition 5.1 was established on 150 specifications; it now rests on 452, spanning every construction of the state variable between the two extremes, and it does not move. The dollar is a safe haven in both regimes regardless of how much geopolitical risk enters σ_{sys}.

(c) What the profile does and does not license. It rules out one of the three readings above and leaves two. Reading 3 — that the amplification was never there — is hard to sustain: a finding that is absent should not sit at 82 per cent of specifications and a ratio of 2.25 at one end of a parameter range, and the decay is orderly rather than noisy. Readings 1 and 2 remain, and they are observationally equivalent on these data, because both predict exactly this profile. If geopolitical risk genuinely belongs in the state variable at close to equal weight, the amplification is a financial-stress artefact and the canonical index is right to suppress it. If w_X is truly small, the canonical index is mis-specified and the amplification is real. **The data cannot adjudicate between them, and nothing in this chapter should be read as adjudicating.**

What the profile does provide is a decision rule for later. Once w_X is estimated, the answer follows mechanically from Table 5.8: an estimate at or below 0.5 restores the amplification claim, an estimate at or above 0.75 retires it, and an estimate in between leaves it undetermined. That is a falsifiable statement about a parameter, which is what the chapter owes the reader in place of the convex response it set out to find.

w_X is not chosen here. It would have been easy to observe that p_{GPRD} has an AR(1) of 0.503 against 0.97–0.99 for the market legs (Table 5.5), conclude that a series that noisy cannot deserve equal weight, set $w_X = 0.25$, and recover an amplification result at 74 per cent of specifications. That would be choosing a parameter by the result it produces. w_X is a free parameter of the model, it is an OPEN row in the registry, and Chapter 11 estimates it from the data jointly with σ^* and the elasticities. Until then the canonical $w_X = 1$ stands as the stated convention of Assumption 5.2, and every result in this chapter is reported at it — including the one that convention costs the model.

What Chapters 11 and 17 inherit. Chapter 11 inherits three specific tasks rather than a general one: estimate σ^* and the three component weights jointly with the elasticities in the full model; report where $\hat{w_X}$ falls relative to the 0.5 and 0.75 boundaries of Table 5.8; and report whether the calm-regime sign and the amplification survive when D_{cap}, D_{debt} and the supply channels are controlled for. Chapter 17 inherits the scoreboard entries: the sign-change claim is *unsupported* on 452 specifications; the amplification claim is *undetermined, conditional on* w_X, and must be scored that way rather than as either a success or a failure;

and any case study appearing to show a risk-on dollar must be reconciled with Proposition 5.1 rather than cited as though it confirmed the original.

Chapter 6

Layer 2 — The Structural and Trust Layer: Φ_{sys} and $s_{\text{US,sys}}$

6.1 The Anchor Question

src p.12

This layer is the structural anchor of the whole system. It does not measure short-term flows or transient states. It quantifies the long-term attributes that determine the *quality* of the dollar as the world's final reserve and safe-haven asset, and it answers one question:

> Why does the world continue — or cease — to trust the dollar system, despite its internal flaws and crises?

That question is not rhetorical, and the model's answer is not "because it always has". Layer 2 governs two modulating functions, and the distinction between them is the chapter's organising idea.

Φ_{sys}, systemic shelter credibility. The absolute trust in the dollar system as a final refuge when σ_{sys} spikes. It scales the exponential term of Equation (2.1): a higher Φ_{sys} amplifies the safe-haven effect. Its natural question is *how strong* the flight to the dollar is.

$s_{\text{US,sys}}$, the safe-flow market share. The share of global flight-to-quality capital that moves into dollar assets rather than into gold, the Swiss franc or the yen. Following ISSUES #8 it multiplies D_{prec} inside the base, $(s_{\text{US,sys}} D_{\text{prec}})^{\alpha_{\text{prec}}(\sigma)}$, so it determines *how much of a panic the dollar captures.*

6.1.1 Why two functions and not one

A reader may reasonably ask why one trust variable would not do. The answer is that the two can move in opposite directions, and the episodes where they do are the ones the model exists to explain.

Consider a crisis whose epicentre is the United States. Global stress rises, so the flight to quality is large. But the asset being fled *to* is issued by the country the crisis is in. Φ_{sys} can hold up — Treasuries remain the deepest market in the world, and there is nowhere else

to put a trillion dollars in a week — while $s_{\text{US,sys}}$ falls, as the marginal dollar of panic goes to gold instead. The observable consequence is a dollar that rises less than the stress alone would imply, without any collapse in the dollar system.

Now consider the reverse: a crisis elsewhere, in which $s_{\text{US,sys}}$ rises toward one because the alternatives are worse, while Φ_{sys} is unchanged. The dollar rises more than the global stress level alone would suggest.

One variable cannot represent both configurations. Two can, and the four structural variables of Section 6.2 through Section 6.5 feed them with different signs and different weights. That is the substance of E4 and E5, which Chapter 7 of the original — §7 of this edition's Part III — leans on heavily.

6.1.2 The four structural variables

The quality of the anchor is determined by four variables: geopolitical power G_{pol}, internal U.S. risk R_{US}, institutional trust M_{trust}, and geopolitical fragmentation T_{frag}. Section 6.2 and Section 6.3 take the first two, Section 6.4 and Section 6.5 the second two; Section 6.6 and Section 6.7 assemble them into Φ_{sys} and $s_{\text{US,sys}}$; Section 6.8 and Section 6.9 take the amplification of R_{US} by Layer 4 and what it implies for tail risk.

All four are latent, all four lie in $[0, 1]$ by the convention of ISSUES #26, and none has a satisfactory single proxy. The IMF's own attempt to measure the phenomenon, and its candour about the difficulty, is Aiyar et al. [3]. That is the honest situation and the cards say so individually. It has a consequence worth stating once at the outset: Layer 2 is the least measurable layer in the model that nonetheless carries a decisive economic claim, and Chapter 11 treats its variables as latent states to be inferred rather than as regressors to be plugged in.

6.2 Geopolitical Power (G_{pol})

Definition 6.1. G_{pol} is the United States' hard-power capacity to enforce the rules of the system it underwrites. It is a latent variable in $[0, 1]$.

Mechanism. The original names three capacities, and each maps to a channel elsewhere in the model.

1. **Enforcing global trade and financial rules**, including through sanctions. This is the capacity that makes S_{off} (Section 4.8) an instrument rather than a wish: a freeze is effective only if third parties comply, and compliance is a function of power.
2. **Securing maritime trade routes**, which is the physical precondition for D_{trans} (Section 3.1). Dollar-invoiced commodity trade requires that the commodity arrives.
3. **Maintaining the alliance network and forward bases.** This is what determines who is inside the swap-line list of Section 4.5 — an alliance structure, not an economic one.

Measurement rule, and a mistake the original makes. The three capacities above share a property that has to be stated explicitly, because the original violates it in one place:

none of them is a function of who the capacity is demonstrated against. G_{pol} is capability, not opponent.

Assumption 6.1 (G_{pol} is scored on demonstrated capability)**.** An episode raises G_{pol} to the extent that it demonstrates or adds capacity on the three headings above. The strategic weight of the target may enter only through how much capability the episode demonstrates — never directly.

This is ISSUES #33, and it is a correction rather than an invention: it is exactly the rule the original's own Case Study 7 applies. CS7 insists that launching a $13 billion carrier "does not automatically mean $\Delta G_{\text{pol}} > 0$" and scores instead by surge capacity, carrier-days on station, availability, escort and manpower constraints, and credible deterrence inside contested zones. Every criterion is a measure of capability (Section 15.3).

Case Study 10 does the opposite. It predicts that intervention in Venezuela "does not significantly increase G_{pol} (since Venezuela is not a major strategic adversary)" — reasoning from the weight of the target, which Assumption 6.1 forbids. That prediction failed against the January 2026 outcome, and Section 16.1.6 scores it and traces the failure to this rule. The original therefore contains two incompatible measurement rules for one variable; CS7's is the one consistent with the definition above, and it is canonical here.

Role in the equation. G_{pol} raises both Φ_{sys} and $s_{\text{US,sys}}$ by providing physical assurance to the system. It does *not* enter κ_{sys} directly: the resolution of ISSUES #3 gives convexity its own explicit form in which G_{pol} acts only through Φ_{sys} and $s_{\text{US,sys}}$, precisely to avoid double-counting it (Section 2.2.3).

The coercion paradox. The original's most interesting claim about this variable is that its effect is not monotone in use. When G_{pol} is exercised *coercively*, through S_{off}, rather than *protectively*, it backfires: it raises T_{frag} and erodes $\Psi_{\text{str}}(t)$. Power spent enforcing is power that demonstrates to every other holder of dollar reserves what enforcement would look like if aimed at them.

This is the same two-timescale structure identified for S_{off} in Section 4.8, and it is worth stating as a proposition rather than a remark, because it is testable:

Proposition 6.1 (Coercion paradox)**.** *An exercise of G_{pol} through S_{off} raises DXY in the short run, through the supply attenuation of Equation* (4.1)*, and lowers it in the long run, through $T_{\text{frag}} \uparrow \Rightarrow \Phi_{\text{sys}} \downarrow$, $s_{\text{US,sys}} \downarrow$, $\Psi_{\text{str}}(t) \downarrow$. The sign of the total effect therefore depends on the horizon, and no single-horizon test can refute the claim.*

The last clause is a warning as much as a result. A proposition whose sign depends on an unspecified horizon is close to unfalsifiable, and Chapter 14 must commit to horizons in advance rather than choosing them after seeing the data.

The branch the proposition does not have. Proposition 6.1 has two branches, short-run gain and long-run erosion. It has no branch for **an exercise of G_{pol} that removes the predicate for a sanctions regime** — coercion that ends coercion. That is not a hypothetical omission: in January 2026 an exercise of G_{pol} was followed within four weeks

by the lifting of oil sanctions, so S_{off} fell and D_{trans} rose (Section 16.1.7). The model has no category for it, ISSUES #34 records the gap rather than filling it, and Section 6.5 states what it costs T_{frag} specifically. Farrell and Newman [33] give the theory of the enforcement capacity; Mulder [59] gives the historical record of the backfire.

Measurement. No index of geopolitical power exists that is both free and credible, and constructing one is a research project rather than a download.

Table 6.1: Measurement protocol for G_{pol}. Every candidate is a partial indicator of one of the three capacities, not a measure of the variable.

Capacity	Candidate indicator	What it misses
Rule enforcement	count of OFAC designations; share of world trade under U.S. sanction	counts intent, not compliance — the quantity that matters is third-party adherence
Sea-lane security	carrier-days on station; forward-deployed tonnage	availability is not the same as willingness to use it
Alliance network	treaty-ally count; forward-base count; SIPRI military expenditure share	a static count of allies says nothing about their reliability under stress

Historical range.

> **Not supplied, and the gap is the world's.** No provider publishes a geopolitical-power index of the kind Layer 2 requires. Building one means normalising SIPRI military-expenditure shares, carrier availability and a treaty-ally count into a single index. *What it costs:* worse than absence: ISSUES #33 found the book had been scoring G_{pol} by the strategic weight of the target rather than by demonstrated capability, and Section 14.3.4 found its three capacities moving in *opposite* directions in one year with no aggregator. G_{pol} **is not merely unmeasured but not a well-defined scalar**, which is why Section 21.1.2 can locate no axis of the attractor's state space.

No range is stated.

Limitations. The variable is slow, latent, and contested even in principle: reasonable analysts disagree about whether U.S. geopolitical power has fallen since 2006, and any index that answers that question will do so by its choice of weights rather than by its data. Chapter 11 therefore treats G_{pol} as a latent state with a smooth prior rather than as an observable, and the case studies use *events* — a sanctions package, a carrier deployment — rather than levels.

6.3 Internal U.S. Risk (R_{US})

src p.13

Definition 6.2. R_{US} is the decay arising from within the U.S. system itself: political instability, social unrest, fiscal risk and institutional degradation. It is a latent variable in $[0, 1]$.

Mechanism. The original names four components: political instability including gridlock and debt-ceiling crises; social unrest; fiscal risk including unsustainable debt and uncontrolled inflation; and institutional degradation, meaning loss of confidence in the courts and the legislature.

The unifying idea is that a safe haven must be safe from its *own* politics. An investor fleeing a global crisis into Treasuries is making a judgement not only about liquidity but about whether the issuer will still honour its obligations in a form worth having. R_{US} is that judgement's complement.

Role in the equation. R_{US} is the most heavily connected variable in Layer 2 and it acts negatively everywhere:

Table 6.2: The four channels through which internal U.S. risk acts. Every one is negative.

$R_{\text{US}} \uparrow$	effect	mechanism
	$\Phi_{\text{sys}} \downarrow\downarrow$	shelter credibility collapses: investors hesitate to flee into a haven that is itself unstable
	$s_{\text{US,sys}} \downarrow$	safe-haven flows divert to gold or the Swiss franc
	$\kappa_{\text{sys}} \downarrow$	crisis reflexivity and acceleration diminish
	$H_{\text{cap}} \downarrow$	**dual erosion**: talent and innovation capital leave an unstable country

The third row is the one the original singles out and the one this edition had to repair. A rising R_{US} *flattens* the dollar's convexity: the currency remains a haven but becomes a worse shock absorber. In the original's E6 that channel had disappeared entirely; the resolution of ISSUES #3 restores it as the factor $(1 - \kappa_R R_{\text{US}})$, with the positivity constraint $\kappa_R < 1$ (Section 2.2.3). Without that restoration the fourth case study's argument has no equation behind it.

The fourth row is the dual-erosion update of ISSUES #24. R_{US} erodes the human-capital *stock* H_{cap}, which reaches the demand flow D_{hcap} with a lag (Section 3.10). So internal risk hits the dollar twice, once through credibility and once through the assets that make the dollar worth holding — on two very different timescales.

Scale. The original writes that R_{US} "can spike from 0 to 100 within hours". Following ISSUES #26, $R_{\text{US}} \in [0, 1]$ like every other latent modulator; the 0–100 reading is retained as a presentation scale and the model consumes $R_{\text{US}}/100$. This matters for comparability: ϕ_R, γ_R and κ_R are then all effects of a full-range move and can be read against one another.

Measurement. R_{US} is the best-measured variable in Layer 2, which is not saying much, and one of its two free proxies is in hand.

Table 6.3: Measurement protocol for R_{US}.

	Primary (in hand)	Secondary and tertiary
Proxy	U.S. Economic Policy Uncertainty index	Partisan Conflict Index; U.S. sovereign CDS spread
Source	FRED `USEPUINDXD` [9]	Federal Reserve Bank of Philadelphia [7]; **CDS is paywalled**
Frequency	daily	monthly; daily
Coverage	2006-01-01 → 2026-09-03, 7,551 obs	1981–, monthly; —
Transformation	percentile rank, then $[0, 1]$	same
Weakness	newspaper-based: it measures *coverage* of uncertainty, which responds to salience as well as to risk	partisan conflict is narrower than institutional degradation; CDS is the cleanest measure of default risk and is not free

Per ISSUES #23, the sovereign CDS substitution is stated rather than silently dropped: EPU and partisan conflict measure *political* risk, while CDS measures the priced probability that political risk becomes a payment failure. They are not the same object, and debt-ceiling episodes are exactly where they diverge — CDS spreads widen sharply while EPU may not move at all if coverage is routine.

Historical range. The EPU index runs daily from 2006-01-01 to 2026-09-03, 7,551 observations. It reaches its sample low of 3.32 on 2015-08-07 and its high of 1,048.95 on 2026-05-02; that maximum is the most recent extreme in the series rather than a historical one, which is itself worth noting for a variable the model treats as slow.

The daily series is extremely noisy — a newspaper count responds to a single news cycle — so any use in Chapter 11 should be at monthly frequency or smoothed, and the card recommends the percentile rank of a 21-day average rather than the raw daily level.

Limitations. The proxy problem here is specific and serious. EPU is built from newspaper coverage, so it rises when journalists write about uncertainty. That is correlated with actual institutional risk and is not the same thing: a well-covered but routine budget dispute can outscore a quiet erosion of institutional norms. The variable the model wants — the probability that the United States stops being a reliable issuer — is not directly observable at any frequency, and no combination of these proxies recovers it. Chapter 11 says what can be identified; the honest answer is a latent state with wide bands, not a measured series.

6.4 Institutional Trust (M_{trust})

Definition 6.3. M_{trust} is confidence in America's core institutions, independent of daily political volatility. It is a latent variable in $[0, 1]$ and it is the bedrock of Φ_{sys}.

Mechanism. The original names four constituents: the independence and competence of the Federal Reserve; the transparency and efficiency of the Treasury and the regulatory bodies; the rule of law and judicial independence; and the protection of intellectual property rights, which is the institutional foundation of D_{hcap}.

The distinction from R_{US} is the point of having two variables and it is worth stating sharply. R_{US} is *volatility* of the political system: a debt-ceiling standoff, a contested election, a riot. M_{trust} is the *level* of confidence in the machinery that survives those events. A country can have high R_{US} and high M_{trust} simultaneously — noisy politics, reliable courts — and that combination describes the United States through most of the sample. The dangerous configuration is the one where M_{trust} falls, because unlike R_{US} it does not oscillate back.

Role in the equation, and the sign. M_{trust} enters Φ_{sys} with a **positive** coefficient:

$$\Phi_{\text{sys}} \;=\; \phi_0 + \phi_G\, G_{\text{pol}} - \phi_R\, R_{\text{US}} + \phi_M\, M_{\text{trust}} - \phi_T\, T_{\text{frag}}, \qquad \text{all } \phi > 0. \tag{6.1}$$

The original prints $-\phi_M M_{\text{trust}}$, which would mean that more institutional trust lowers shelter credibility. That contradicts this very section — M_{trust} "forms the bedrock of Φ_{sys}" — and the whole of the inflation argument in Chapter 18. It is a transcription slip and the sign is corrected here (ISSUES #1).

The inflation channel. E_{inf} appears in $s_{\text{US,sys}}$ (Section 6.7) but not in Equation (6.1), while the argument of Chapter 18 runs inflation → loss of faith in the Fed → collapse of Φ_{sys}. The resolution of ISSUES #25 is that inflation reaches Φ_{sys} *through* M_{trust} rather than beside it:

$$M_{\text{trust}} \;=\; m_0 - m_E\, E_{\text{inf}} - \ldots, \qquad m_E > 0, \tag{6.2}$$

so that

$$\frac{\partial \Phi_{\text{sys}}}{\partial E_{\text{inf}}} \;=\; -\,\phi_M\, m_E \;<\; 0 \tag{6.3}$$

is derived rather than asserted. No direct inflation term is added to Equation (6.1), which would double-count it against $s_{\text{US,sys}}$.

Equation (6.3) is the formal content of the original's claim that hyperinflation destroys shelter credibility "even if G_{pol} remains strong and R_{US} is temporarily low". It also makes the claim falsifiable: it predicts that a sustained rise in inflation expectations lowers the dollar's safe-haven *amplitude*, and Chapter 18 tests that against 2021–2023.

Measurement. The honest position is that M_{trust} has no adequate proxy, and the primary above is a narrow one. `T5YIFR` tells us whether markets believe the Fed will deliver its target. It says nothing about judicial independence or the protection of intellectual property, which are two of the four constituents. Using it alone would identify M_{trust} with Fed credibility and quietly discard the rest.

Chapter 11 therefore treats M_{trust} as a latent state with `T5YIFR` as one noisy indicator among several, rather than as an observable. Where a case study needs M_{trust} to have moved, it must point to an event — a decision on agency authority, a challenge to Fed independence — and not to a series.

Table 6.4: Measurement protocol for M_{trust}.

	Primary	Secondary
Proxy	5-year, 5-year forward inflation expectation, as a measure of Fed credibility	Gallup confidence in institutions; Michigan long-run inflation expectations
Source	FRED `T5YIFR` (in hand)	Gallup annual; FRED `MICH` (in hand)
Frequency	daily	annual; monthly
Transformation	**negated**, percentile rank, then $[0, 1]$	percentile rank
Availability	free; in hand	free
Weakness	measures one institution, the Fed, and only through one channel	Gallup is annual and covers institutions the model does not price; `MICH` is survey-based and noisy

6.5 Geopolitical Fragmentation (T_{frag})

Definition 6.4. T_{frag} measures the extent to which the world is drifting away from the U.S.-led order. It is a latent variable in $[0, 1]$ and it is the direct counterweight to G_{pol}.

Mechanism. The original names four drivers of a rising T_{frag}: alternative economic or security blocs gaining strength; bilateral non-dollar trade increasing; alternative payment systems C_{alt} developing; and the United States weaponising S_{off}, which pushes others to decouple.

The fourth closes a loop. G_{pol} exercised coercively raises T_{frag} (Proposition 6.1), and T_{frag} reduces the effectiveness of G_{pol}. Power used to enforce is power that erodes the conditions for enforcement. This is the model's single most important feedback outside the reflexive layer, and it is negative — self-limiting rather than self-reinforcing.

Role in the equation. The two-layer variable. T_{frag} is the only variable in the model that enters two different layers:

- In **Layer 2** it enters Φ_{sys} with $-\phi_T$ and $s_{\text{US,sys}}$ with $-\gamma_T$, acting on the dollar's credibility *now*.
- In **Layer 5** it is one of the two hard drivers of the erosion function E8, alongside C_{alt}, acting on the dollar's structural position over *decades* (Chapter 9).

That dual membership is a modelling decision worth defending rather than assuming. Fragmentation degrades trust immediately — a reserve manager who has watched a freeze reprices dollar risk this quarter — and it degrades the plumbing slowly, as alternative rails are built. The two effects have different magnitudes and different half-lives, and collapsing them into one term would force a single elasticity to represent both. Section 4.10.2's double-counting logic applies here in reverse: this is one cause with two distinct consequences, not one consequence counted twice.

A direction error corrected. The original writes the effect of rising fragmentation as "`Psi_structural`↑ (erosion coefficient activation)". Under the convention fixed in ISSUES #2, $\Psi_{\text{str}}(t)$ is a structural *integrity* coefficient, so erosion is always $\Psi_{\text{str}}(t) \downarrow$. The arrow is corrected here and in Chapter 9:

$$T_{\text{frag}} \uparrow \implies G_{\text{pol}} \text{ effectiveness} \downarrow, \quad \Phi_{\text{sys}} \downarrow, \quad s_{\text{US,sys}} \downarrow, \quad \Psi_{\text{str}}(t) \downarrow . \tag{6.4}$$

Table 6.5: Measurement protocol for T_{frag}.

	Primary	Secondary
Proxy	share of world trade settled in non-dollar currencies	CIPS transaction volumes; count of bilateral local-currency settlement agreements
Source	SWIFT RMB Tracker (shares only); IMF direction of trade	PBoC/CIPS annual reports; central-bank announcements
Frequency	monthly; annual	annual
Transformation	percentile rank, then $[0, 1]$	log, indexed
Availability	free but partial: SWIFT publishes shares, not values, and only for selected currencies	free but annual and self-reported
Weakness	SWIFT measures messaging, not settlement, and excludes CIPS entirely — so it systematically understates exactly the substitution it is meant to detect	self-reported volumes are not audited and the growth base is small

Measurement. The weakness in the first column is severe enough to state as a warning. The natural proxy for "how much trade has left the dollar system" is a measure produced *by* the dollar system's own messaging network, which does not observe the alternative rails. A measurement blind to the phenomenon it measures will report gradual change during exactly the period when the phenomenon accelerates.

Historical range. The direct measure was searched for and does not exist in usable form. The result of that search is recorded here rather than left open, because T_{frag} drives the erosion cliff in Chapter 9 and a chapter cannot calibrate a cliff against a variable whose measurement status is undeclared.

SWIFT is free, monthly, and unusable as a series. The RMB Tracker — renamed the *Global Currency Tracker* in February 2026 — publishes each currency's share of payment value monthly, without charge or registration, and it is exactly the quantity this card wants. It is distributed as one PDF per month with no CSV, no spreadsheet and no programmatic access, so a series covering this book's sample means transcribing on the order of one hundred and eighty documents by hand. That is a legitimate task for a later session and it is not a data source that can be scripted, so no SWIFT number appears in this book until the transcription is done.

CIPS cannot cover the sample. The system began operating in October 2015. It does not exist for the first ten years of the window, so whatever it shows, it cannot be the proxy for a variable the model runs from $t_{\text{base}} = 2006$. Its official statistics are published as web pages rather than as a downloadable series, and the only machine-readable aggregation available is third-party and seven months long.

T_{frag} has no way down. Before measurement, a structural point. In the model as written, every arrow into T_{frag} raises it: $S_{\text{off}} \uparrow \Rightarrow T_{\text{frag}} \uparrow$ is the coercion paradox, and T_{frag} in turn lowers Φ_{sys} and raises f_{hard}. Nothing lowers it. T_{frag} is a state in $[0, 1]$ and the arithmetic permits it to fall, but no mechanism in this book produces the fall — so a de-fragmentation can enter only by assumption, never as a consequence.

That is ISSUES #34, found by auditing the January 2026 Venezuela episode (Section 16.1.7), and Section 9.5 shows it is the same defect that makes sterling's 1930s recovery unrepresentable. A resolution channel would look like $\Delta T_{\text{frag}} = \rho_T S_{\text{off}} - \delta_T R$, with R indicating that a coercion's predicate has been removed; it is not added, because T_{frag} has no direct measure and the sample holds at most two candidate resolution episodes. **The consequence for every use of T_{frag} in this book is that its path is monotone by construction and not by evidence.**

The substitute, and what it costs. Under ISSUES #23, T_{frag} is measured by two free, official, machine-readable series that are *adjacent* to the concept: the BIS Triennial Central Bank Survey's currency composition of OTC foreign-exchange turnover, and the IMF COFER dollar share of allocated reserves used in Section 3.5. Neither is the thing. Turnover is trading rather than settlement of real obligations; the survey is triennial where the mechanism is continuous; and neither series observes payment-system architecture at all, which is the actual content of T_{frag}. The substitute answers *how much is the dollar traded and held*, where the card asks *how much must a transaction route through dollar infrastructure*. Those two questions diverge precisely in the scenarios T_{frag} exists to model, and every use of the substitute in Chapter 9 and Chapter 11 carries that caveat.

Table 6.6: Currency composition of OTC foreign-exchange turnover, per cent, net-net basis, April daily averages. Shares sum to 200 because every trade has two legs — the BIS convention, preserved here rather than rescaled. Source: BIS Triennial Central Bank Survey 2025, final results, Table 4; file `bis_triennial_fx_2025.xlsx`; script `extract_bis_triennial.py`.

Currency	2013	2016	2019	2022	2025
USD	87.04	87.58	88.30	88.37	**89.21**
EUR	33.41	31.39	32.31	30.61	28.90
JPY	23.05	21.62	16.83	16.72	16.77
GBP	11.82	12.80	12.81	12.93	10.23
CNY	2.23	3.99	4.33	7.01	**8.52**

What the substitute shows, which is not what the original expects. Over 2013–2025 the renminbi's share nearly quadrupled, from 2.23 to 8.52 per cent, a rise of 6.29 percentage points. Over the same twelve years **the dollar's share rose**, from 87.04 to 89.21 per cent. What fell was the euro, by 4.5 points, and the yen, by 6.3.

This is a genuine and inconvenient result, and it is reported rather than explained away. On the best free measure of dollar centrality available, the incumbent gained ground across exactly the period the original characterises as de-dollarisation. Two readings survive it and this book does not choose between them here. Either the turnover measure is blind to the fragmentation — which is the warning this card's measurement table already carries, since a dollar-centric market's own trading statistics would be the last place a shift away from it appeared — or the fragmentation is occurring among the *challengers*, with the renminbi taking share from the euro and the yen rather than from the dollar, which is a materially different world from the one T_{frag} describes.

Chapter 9 inherits the problem in a sharp form: the erosion cliff is a function of T_{frag}, and the only measured input to T_{frag} currently moves in the direction of *less* erosion. That must be stated in the cliff's calibration section and carried into Chapter 17's scoreboard, not resolved by choosing a different proxy until one agrees.

> **Not supplied, and the gap is ours, and it is the largest one.** The SWIFT Global Currency Tracker publishes currency shares of payment value monthly, free, in roughly 180 PDFs from 2011. Transcribing even the December observation of each year would give a fifteen-point annual series. This is **the single highest-value manual data task left in this book** and it is measured in afternoons. *What it costs:* T_{frag} scores one partial and four untestable, and the substitute adopted under ISSUES #23 — BIS triennial turnover and COFER — is adjacent rather than approximate: turnover is trading, not settlement, and triennial, not monthly. Chapter 9's erosion cliff is a function of T_{frag}, and its only measured input currently moves in the direction of *less* erosion.

6.6 The Credibility Function Φ_{sys}

src p.28

6.6.1 Statement

Equation (6.1) is E4 of the compendium, restated with the corrected sign:

$$\Phi_{\text{sys}} \;=\; \phi_0 + \phi_G\, G_{\text{pol}} - \phi_R\, R_{\text{US}} + \phi_M\, M_{\text{trust}} - \phi_T\, T_{\text{frag}}. \tag{6.5}$$

6.6.2 Assumptions

Assumption 6.2 (Linearity in the structural variables)**.** Φ_{sys} is affine in the four structural variables G_{pol}, R_{US}, M_{trust} and T_{frag}. Interactions among them are excluded.

Assumption 6.3 (Bounded inputs, unbounded output)**.** Each input lies in $[0, 1]$; Φ_{sys} itself is required only to be positive over the admissible parameter region (Assumption 2.5).

Assumption 6.2 is the strong one and it is worth naming its cost. The original's own argument in Section 6.8 is that R_{US} becomes explosive when the psycho-technical layer amplifies it — which is an interaction between R_{US} and H_{soc}. An affine Φ_{sys} cannot represent that. The model handles it by making R_{US} *itself* state-dependent rather than by adding an interaction term to Equation (6.5), and Section 6.8 sets out exactly what that costs.

6.6.3 Functional form

Why affine here, when $s_{\text{US,sys}}$ next door is a logistic? Because the two functions answer different questions. $s_{\text{US,sys}}$ is a *share* and must lie in $[0, 1]$, so it needs a squashing function. Φ_{sys} is an *intensity* multiplying the safe-haven exponent; it has no natural upper bound, and imposing one would cap the dollar's crisis response for reasons of algebra rather than economics.

The rejected alternative was multiplicative, $\Phi_{\text{sys}} = \phi_0 \, G_{\text{pol}}^{a} M_{\text{trust}}^{b} (1 - T_{\text{frag}})^{c}$, which has the attractive property that credibility collapses to zero if any single pillar does. We reject it because it makes that claim too strongly: a total loss of judicial independence would zero out shelter credibility even with the deepest bond market in the world still functioning. The affine form lets the pillars substitute, which is the more defensible default, and Chapter 11 can test the restriction.

6.6.4 Parameters and comparative statics

Table 6.7: Parameters of Equation (6.5). All are OPEN values in the registry; the CS7 figure is the original's assumption, retained and labelled.

Symbol	Meaning	Range	Provenance
ϕ_0	intercept	> 0	estimated
ϕ_G	geopolitical power	> 0	0.5 in CS7 — assumption (original)
ϕ_R	internal U.S. risk	> 0	estimated
ϕ_M	institutional trust	> 0	**sign corrected**, ISSUES #1
ϕ_T	fragmentation	> 0	estimated
m_E	inflation into M_{trust}	> 0	estimated

The comparative statics the case studies use follow immediately: $\partial \Phi_{\text{sys}} / \partial G_{\text{pol}} = \phi_G$, $\partial \Phi_{\text{sys}} / \partial R_{\text{US}} = -\phi_R$, $\partial \Phi_{\text{sys}} / \partial M_{\text{trust}} = +\phi_M$, $\partial \Phi_{\text{sys}} / \partial T_{\text{frag}} = -\phi_T$, and the composed derivative Equation (6.3) through inflation.

6.6.5 Failure modes

Two. First, Equation (6.5) can go negative for extreme inputs, which would flip the sign of the entire safe-haven term rather than merely shrinking it; Assumption 2.5 excludes that region by assumption rather than by construction, and Chapter 10 must check that the calibrated parameter set respects it. Second, all four inputs are latent and none is separately

identified from the others by any series in hand, so the individual ϕ coefficients are unlikely to be recoverable even where Φ_{sys} as a whole is.

6.7 The Safe-Flow Share $s_{\mathrm{US,sys}}$

src p.28

6.7.1 Statement

$$s_{\mathrm{US,sys}} = \frac{1}{1 + \exp\left(-\gamma_0 + \gamma_G\, G_{\mathrm{pol}} - \gamma_R\, R_{\mathrm{US}} - \gamma_E\, E_{\mathrm{inf}} - \gamma_T\, T_{\mathrm{frag}}\right)} \tag{6.6}$$

This is E5, unchanged from the original. The logistic is the right form here for the reason given in Section 6.6.3: $s_{\mathrm{US,sys}}$ is the share of global flight-to-quality flow that the dollar captures, so it must lie in $[0, 1]$ by construction and it should saturate. The dollar cannot capture more than all of a panic, and the marginal effect of more geopolitical power on the share must vanish as the share approaches one.

6.7.2 Where it enters

Following ISSUES #8, $s_{\mathrm{US,sys}}$ multiplies safe-haven demand *inside the base*:

$$\left(s_{\mathrm{US,sys}}\, D_{\mathrm{prec}}\right)^{\alpha_{\mathrm{prec}}(\sigma)(\sigma_{\mathrm{sys}})}. \tag{6.7}$$

The original instead scaled the exponent, $D_{\mathrm{prec}}^{\alpha_{\mathrm{prec}}(\sigma)\, s_{\mathrm{US,sys}}}$, which has the wrong limit: as $s_{\mathrm{US,sys}} \to 0$ that factor tends to 1, so a panic in which the dollar captures nothing would leave the numerator unchanged. Under Equation (6.7) the factor tends to zero, which is the intended economics.

6.7.3 Marginal effects, and the derivative CS7 uses

Differentiating Equation (6.6), writing $s = s_{\mathrm{US,sys}}$:

$$\Delta\Phi_{\mathrm{sys}} \approx \phi_G\, \Delta G_{\mathrm{pol}}, \tag{6.8}$$

$$\Delta s_{\mathrm{US,sys}} \approx s\,(1 - s)\,\gamma_G\, \Delta G_{\mathrm{pol}}. \tag{6.9}$$

Equation (6.9) is the derivative of the logistic and it is the arithmetic the seventh case study runs. It is a **first-order approximation** valid only for small ΔG_{pol}: the logistic's curvature means a large shock is over-stated near $s = 0.5$ and under-stated in the tails. The original does not say this; ISSUES #22 records it and Chapter 15 states it where the arithmetic is used.

With the original's assumed values — $s = 0.70$, $\phi_G = 0.5$, $\gamma_G = 1.0$, all *assumptions (original)* rather than estimates — a shock of $\Delta G_{\mathrm{pol}} = +0.05$ gives $\Delta\Phi_{\mathrm{sys}} = +0.025$ and $\Delta s_{\mathrm{US,sys}} = 0.70 \times 0.30 \times 1.0 \times 0.05 = +0.0105$, so the share moves from 70 to 71.05 per cent. Those numbers are retained as the original's worked example and labelled as such; they are not calibration.

6.7.4 The competing havens

$s_{\text{US,sys}}$ is a share, so it is defined only against what the dollar competes with: gold, the Swiss franc, the yen. A card for the share therefore needs a price for the alternatives, and gold is the one that matters.

Gold data. No free *daily* gold benchmark exists: the LBMA gold and silver fixes were removed from FRED on 2022-01-31 with the rest of the ICE Benchmark Administration data, and no replacement was offered. The primary source used in this edition is therefore the **World Bank Pink Sheet** — monthly, public domain, spot average of daily rates, dollars per troy ounce, back to 1960. Over the modelled sample it runs from $550.00 in 2006-01 to $4,411.00 in 2026-08, with a peak of $5,020.00 in 2026-02. The full series begins at $35.00 in 1960-01, the Bretton Woods peg, which is the natural starting point for Chapter 22.

The BLS export price index for non-monetary gold, `IQ12260`, is retained *only* as a returns cross-check and is never quoted in dollars per ounce.

Measurement of the share itself.

> **Not supplied, and the gap is mostly ours.** Three of the four legs are free: TIC flows, Swiss franc and yen appreciation over stress windows, and — since the audit of Section 9.7.1 — world official gold holdings at market value, now in `data/raw`. Combining them into a share is a construction task rather than a download. *What it costs:* $s_{\text{US,sys}}$ has two scored predictions and both are untestable, so E5's logistic share enters every equation as a latent and leaves it as one. Section 17.8 ranks building it fourth on the research agenda, and the gold leg arriving is why it moved up the list.

6.8 Amplification of R_{US} by the Psycho-Technical Layer

This section is the original's own critical update to Layer 2, and it is the point at which Layer 2 stops being slow.

The claim. R_{US} is no longer a slow-moving variable oscillating around a stable mean. With $H_{\text{soc}} \to 1$ and $I_{\text{imp}} \to 1$ it becomes explosive: it can spike "from 0 to 100 within hours or days", triggered by relatively small events. The mechanism runs through the psycho-technical layer of Chapter 8. Algorithmic reinforcement A_{alg} and concentrated information sources C_{src} create echo chambers, reducing the diversity R_{div} that would otherwise damp a narrative; a divisive meme then propagates at velocity V_{meme} across a homogeneous population; and shortened cognitive loops with near-zero friction convert that narrative into collective action almost instantly.

The consequence for Layer 2 is that surges in R_{US} cause an immediate collapse in Φ_{sys} — trust in U.S. stability can disintegrate faster through this mechanism than through any other.

How the model represents it. As Section 6.6.2 noted, Equation (6.5) is affine and cannot carry an $R_{\text{US}} \times H_{\text{soc}}$ interaction. The model instead makes the *innovation distribution* of R_{US} state-dependent rather than its coefficient: shocks to R_{US} have fatter tails when H_{soc} and I_{imp} are high, while ϕ_R stays constant.

This is a real modelling choice with a real cost, and both should be visible. The benefit is that Equation (6.5) stays linear and estimable. The cost is that the amplification becomes a statement about the *distribution* of R_{US} rather than about the *transmission* from R_{US} to Φ_{sys}, and those are different claims. The original says trust disintegrates faster; our representation says the shocks are larger. The two coincide only if transmission is linear, which is exactly what Assumption 6.2 assumes.

Why not simply add the interaction. Because it is not identified. A term $\phi_{RH}\, R_{\text{US}}\, H_{\text{soc}}$ would require separating the effect of R_{US} from the effect of R_{US} conditional on H_{soc}, using two latent variables neither of which has a satisfactory proxy (Section 6.3, Chapter 8). Adding a parameter that cannot be estimated does not make a model more faithful; it makes it less honest. Chapter 8 states the alternative specification explicitly so that a reader with better data can implement it.

The scale point. The original's "0 to 100 within hours" is retained as a presentation scale; the model consumes $R_{\text{US}}/100$, per ISSUES #26. What survives the rescaling is the substantive claim — the *speed*, not the range.

6.9 Synthesis and Tail Risk

src p.15

Φ_{sys} and $s_{\text{US,sys}}$ are determined by the balance between G_{pol} as the upward force and R_{US}, M_{trust} and T_{frag} as the others — with M_{trust} entering positively, so that its *fall* is the downward force. But the layer's stability is not symmetric, and Table 6.8 collects why.

Table 6.8: The four structural variables: what each does, how fast, and how well it can be measured. "Speed" is the timescale on which a material move happens.

Variable	Acts on	Speed	Measurability
G_{pol}	Φ_{sys} $\uparrow$, $s_{\text{US,sys}}$ $\uparrow$	decades	none; a construction task, contested even in principle
R_{US}	Φ_{sys} $\downarrow\downarrow$, $s_{\text{US,sys}}$ $\downarrow$, κ_{sys} $\downarrow$, H_{cap} $\downarrow$	**hours to days**	EPU in hand, but it measures coverage of uncertainty
M_{trust}	Φ_{sys} $\uparrow$	years, one way	one narrow proxy for one of four constituents
T_{frag}	Φ_{sys} $\downarrow$, $s_{\text{US,sys}}$ $\downarrow$, $\Psi_{\text{str}}(t)$ $\downarrow$	decades	partial, and systematically blind to the substitution it measures

Three conclusions follow, and the third is the chapter's.

The layer is asymmetric in speed. Three of the four variables move over years or decades; one can move in hours. A layer the original describes as the "low-frequency (yearly)" stratum contains a variable with a daily tail. Any treatment of Layer 2 as slow — including the convenience of holding it fixed within a case-study window — is safe for G_{pol}, M_{trust} and T_{frag} and unsafe for R_{US}.

The layer is asymmetric in reversibility. R_{US} oscillates; a debt-ceiling crisis resolves. M_{trust} and T_{frag} do not. Confidence in institutions, once lost, is rebuilt over a generation if at all, and payment rails once built are not dismantled. The same nominal fall in Φ_{sys} therefore means something quite different depending on which variable delivered it — and it is why Chapter 9 takes T_{frag}, and not R_{US}, as its input.

The layer carries more tail risk than its architecture suggests. The structural stability of the dollar system is more fragile to endogenous shocks than the original model's four-variable balance implies, because the amplification of Section 6.8 sits on the one variable that moves fastest and hits Φ_{sys} hardest. That is a claim about the *shape* of a distribution rather than its centre, and it will not appear in any estimate of ϕ_R. Chapter 11 can report a coefficient; it cannot report a tail. Only the case studies of Chapter 13 and Chapter 16, chosen precisely because their episodes were extreme, can speak to it, and Chapter 17 collects the verdict.

Chapter 7

Layer 3 — The Reflexive and Acceleration Layer: κ_{SYS}

7.1 Convexity as a Sensitivity Knob

src p.15

This layer captures the non-linear and self-reinforcing dynamics of the dollar, particularly during stress. Its variable κ_{sys} is the system's sensitivity knob.

The distinction that makes κ_{sys} worth a chapter is that **it does not measure the direction of the dollar**. It measures the *acceleration* of its response. A high κ_{sys} means that even a small imbalance between demand and supply, or a minor shock to σ_{sys}, generates a disproportionately large move. The same imbalance in a low-κ_{sys} system produces a small one.

Formally κ_{sys} is the exponent on the bracket in Equation (2.1), so in logs it multiplies the entire weighted sum of channels:

$$\ln DXY \;\supset\; \kappa_{\text{sys}} \left[\sum_i \alpha_{\text{i}}(\sigma_{\text{sys}}) \ln D_i - \sum_j \beta_{\text{j}}(\sigma_{\text{sys}}) \ln S_j \right]. \tag{7.1}$$

Section 2.6.3 explains why this is an exponent and not a coefficient: it rescales every channel at once, which is a claim about transmission rather than about level. That also makes κ_{sys} the single most consequential number in the model. Doubling it doubles the dollar's response to everything in Layer 0 simultaneously.

7.1.1 What this chapter has to repair

The original states κ_{sys} twice and the two statements do not agree.

In E1 it is $\kappa(\sigma, G_{\text{pol}}, R_{\text{US}})$, with $\kappa'(\sigma) > 0$ and R_{US} flattening the convexity — and the fourth case study's argument runs explicitly through a "$-\kappa_R R_{\text{US}}$ term". In E6 it becomes $f(\kappa_{\text{base}}, V_{\text{USD}}, M_{\text{shock}}, A_{\text{meta}}) \times H_{\text{soc}} \times I_{\text{imp}}$: σ, G_{pol} and R_{US} have all disappeared, and $f(\cdot)$ is never written down.

So the chapter has three jobs beyond exposition. Write $f(\cdot)$ explicitly. Restore the σ_{sys} and R_{US} channels, without which the fourth case study has no equation behind it. And

relocate the crisis-memory formulas E11–E14, which the original appends to the end of its final chapter although they define a Layer-3 variable (ISSUES #3, ISSUES #16).

7.2 Baseline Convexity (κ_{base})

Definition 7.1. κ_{base} is the inherent sensitivity of the system in the absence of reflexive amplification: the convexity that structural leverage alone would produce. It is a latent variable in $[0, 1]$.

Mechanism. A market intermediated by levered balance sheets transmits shocks more violently than one intermediated by unlevered ones, because a price move that impairs capital forces position changes that move the price again. This is the demand-side twin of the funding-liquidity spiral of Brunnermeier and Pedersen [19], and its cyclicality is the procyclical leverage documented by Adrian and Shin [1].

Measurement.

> **Not supplied, and the gap is ours.** Three free candidates exist and none has been taken: the BIS credit-to-GDP gap (quarterly), FINRA margin debt (monthly), and primary-dealer net financing from the New York Fed (weekly). The first is far too slow for a variable that modulates a crisis response. *What it costs:* κ_{base} is one of the seven drivers of κ_{sys}, and κ_{sys} has **four predictions and no verdicts**. Proposition 11.1 shows the seven-way decomposition is unidentifiable even with data, so this series would not rescue the decomposition — it would only let one term of it be observed.

No range is stated. Note the tension with Section 3.7: dealer gamma and structural leverage are closely related objects, and if D_{gamma} is ever measured, Chapter 10 must check that the same amplification is not entering both here and in the numerator.

7.3 Velocity of Shock Transmission (V_{USD})

Definition 7.2. V_{USD} is the rate at which information about a shock propagates through markets. A higher V_{USD} raises κ_{sys}, because markets have less time to absorb an impact gradually. It is a latent variable in $[0, 1]$.

Mechanism. The claim is about *time*, not about magnitude. Two markets receiving the same news can produce very different price paths depending on how quickly the news reaches all participants. When diffusion was slow, a shock arrived in waves and each wave met a market that had partly adjusted. When diffusion is instantaneous, the whole adjustment happens at once.

V_{USD} is therefore the technological complement to H_{soc}: homogeneity determines *whether* agents react the same way, velocity determines *how fast* they learn they should.

Measurement.

> **Not supplied, and the gap is ours for the free route, the world's for the better one.** News-flow velocity is constructible from GDELT, which is free and high-frequency. The better measure — the decay rate of intraday volatility clustering after scheduled announcements — needs intraday data this project does not have and would have to licence. *What it costs:* V_{meme} is one of the nine Layer-4 variables with no series, and it is the one the original leans on hardest: Section 18.2.1's claim that institutional trust collapses "within months instead of years" is a claim about V_{meme}, and Proposition 18.2 could only test it through a substitute.

No range is stated.

7.4 Crisis Memory (M_{shock})

src p.16

Definition 7.3. M_{shock} is the degree to which market participants retain conditioned reflexes from past crises. High M_{shock} raises κ_{sys}: early signs of stress trigger overshooting responses, because participants have seen where this goes before.

This section carries the formulas the original places at the very end of its final chapter, under the heading "New Update". They define a Layer-3 variable and they belong here (ISSUES #16). A cross-reference is left at their original location in the concordance.

7.4.1 The memory function

Crisis memory at time t is the aggregation of ghosts from the past, faded by two distinct decays:

$$M_{\text{shock}}(t) \;=\; \sum_{k \in \Omega} S_k \; e^{-\lambda_f \,(t - t_k)} \; \rho_{\text{gen}}(t, t_k), \tag{7.2}$$

where Ω is the set of major past crises, S_k the severity of crisis k, and t_k its date.

Psychological decay. The factor $e^{-\lambda_f (t - t_k)}$ is the individual forgetting curve. Even for someone who lived through 2008, by 2024 the fear has diminished. The rate λ_f is higher in the information age — when V_{meme} is high, new material constantly overwrites the old.

The symbol is a rename. The original writes this rate λ, which collides with the safe-haven coefficients λ_1 and λ_2 of E1; it is λ_f throughout this edition (ISSUES #11).

Generational decay. The factor $\rho_{\text{gen}}(t, t_k)$ is the share of market capital managed by people who actually lived through crisis k. The original proposes a linear approximation:

$$\rho_{\text{gen}}(t, t_k) \;\approx\; \max\!\left(0,\; 1 - \frac{t - t_k}{T_{\text{career}}}\right), \tag{7.3}$$

with T_{career} the average career span of an investor generation, about 20 to 25 years.

These two decays are conceptually different and the distinction is the interesting part of the construct. Psychological decay is *continuous and never complete*: the exponential approaches zero but never reaches it, so a trader who lived through 2008 retains some residue forever. Generational decay is *linear and terminates*: at $t - t_k = T_{\text{career}}$ the memory is gone, not faint, because the people who held it have left the market. A model with only the first would predict that institutional memory fades smoothly; with both, it predicts a hard horizon.

7.4.2 The worked example, and what it implies

The original works its own example for the 2008 crisis, and it is worth reproducing because it pins down a parameter.

- **In 2018**, ten years on, most traders who worked through 2008 are still active: $\rho_{\text{gen}}(t, t_k) \approx 0.6$. Memory is still a live input to κ_{sys}.
- **In 2038**, thirty years on, a younger cohort is in charge: $\rho_{\text{gen}}(t, t_k) \approx 0$. The 2008 crisis has stopped being a behavioural fact and become history.

The first of these is not free-standing — it *implies* a value. Setting $1 - 10/T_{\text{career}} = 0.6$ gives $T_{\text{career}} = 25$ years, the upper end of the original's stated range. We adopt $T_{\text{career}} = 25$ on that basis and label it as implied by the original's own example rather than as an independent choice.

Figure 7.1 plots Equation (7.2) for the two crises whose severities the original states — 2008 at 10/10 and 2001 at 6/10 — under a slow and a fast forgetting rate. Note what the figure does *not* contain: the original's set Ω is written "$\{2008, 2020, 2029, \ldots\}$", and 2029 is a hypothetical future crisis inside a formula example. It is dropped (ISSUES #20). 2020 is omitted too, for a different and more mundane reason: the original assigns no severity to it, and inventing one would put a number in a figure that no source supports.

Two readings follow from the shape.

First, **the two decays bind at different times**. In the decade after a crisis the exponential does almost all the work: memory falls from 13.0 to 4.9 between 2008 and 2018 under slow forgetting, while $\rho_{\text{gen}}(t, t_k)$ has only fallen to 0.6. It is the generational term that eventually terminates the memory, and it does so abruptly on the scale of the figure.

Second, λ_f **matters more than it looks**. Tripling the forgetting rate from 0.05 to 0.15 cuts memory in 2018 from 4.9 to 1.8 — a factor of nearly three, a decade after the event. Since the original's own claim is that λ_f rises with V_{meme}, and V_{meme} rises with the psycho-technical layer, this is a channel through which Chapter 8 lowers κ_{sys} even as it raises it through H_{soc} and I_{imp}. That countervailing effect is not in the original and Section 7.6.2 takes it up.

7.4.3 Synthetic re-activation

The original adds a mechanism by which a faded memory can be revived without a new crisis:

$$M^{\text{eff}}_{\text{shock}}(t) \;=\; M_{\text{shock}}(t) \times \Big(1 + \alpha_{\text{res}}\, A_{\text{alg}}\, \text{Sim}(t, t_k)\Big), \tag{7.4}$$

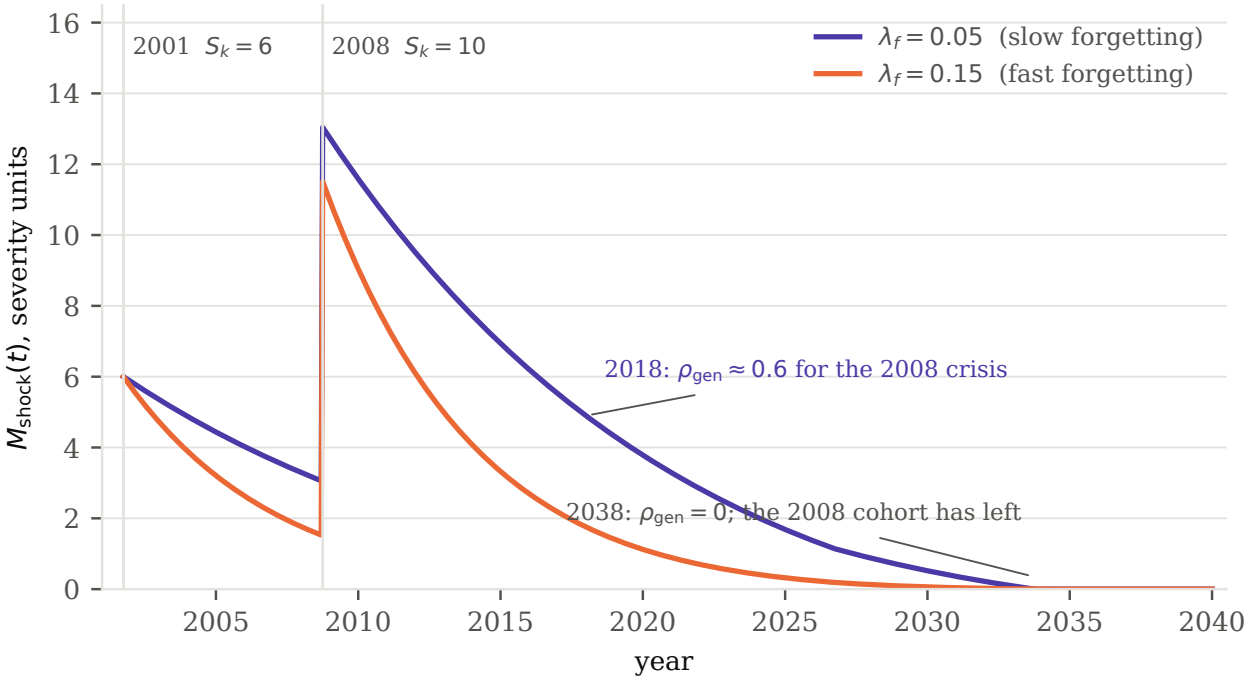


Figure 7.1: $M_{\text{shock}}(t)$ under Equation (7.2)–Equation (7.3) for the two crises whose severity the original states. $T_{\text{career}} = 25$ years, implied by the original's own 2018 worked example; the two forgetting rates are illustrative and bracket a slow and a fast case, since the original gives no value for λ_f. Source: author's calculation from the stated formulas; script `f23_mshock.py`.

where $\text{Sim}(t, t_k)$ measures how closely a current event resembles crisis k and A_{alg} is algorithmic reinforcement. The resonance coefficient is written α in the original, which collides with the elasticities $\alpha_{\text{i}}(\sigma_{\text{sys}})$; it is α_{res} here (ISSUES #11).

The mechanism is that recommendation systems surface historical analogies. A market disturbance that resembles 2008 causes 2008 content to circulate, which re-arms the conditioned reflex in participants who had partly forgotten and installs it in participants who were never there. Memory becomes transmissible rather than merely experienced.

This is the most speculative equation in the chapter and it should be labelled as such. $\text{Sim}(t, t_k)$ has no definition in the original beyond its name, and no proxy is proposed here. Chapter 11 treats Equation (7.4) as a structural hypothesis that the data cannot currently address.

Measurement of M_{shock}.

> **Not supplied, and the gap is neither — this one is not a data gap at all.** M_{shock} is built from a severity coding of Ω and a value for λ_f. The original supplies severities for 2008 (10/10) and 2001 (6/10) and no others. Coding 2011, 2020 and 2022 is an **author decision**, not a download, and ISSUES #20 already removed the hypothetical 2029 entry from the original's own example set. *What it costs:* Figure 7.1 plots only the two episodes the original scores, and M_{shock}'s single scored prediction is untestable. Nothing here is waiting on a provider; it is waiting on the author.

7.5 Meta-Reflexivity (A_{meta})

Definition 7.4. A_{meta} captures second-order reflexivity: the dollar's own movement becomes input information that reinforces further movement. It is a latent variable in $[0, 1]$.

Mechanism. This is Soros's reflexivity applied to the price itself [70]. A rising dollar is not merely an outcome; it is a signal that participants read and act on, and the action pushes the dollar further. The loop requires no fundamental input at any stage.

The clearest concrete instance is trend-following capital: a systematic strategy that buys strength is, by construction, a mechanism converting a price move into a further price move. But the model's claim is broader — that the interpretation of price as information is a general feature of markets, not a property of one strategy class.

Interactions. A_{meta} compounds with D_{debt} from the demand side: the dollar short squeeze of Section 3.2 is a reflexive loop in *obligations* where A_{meta} is one in *beliefs*. Both convert a dollar move into further dollar buying, and Chapter 10 must be careful not to let a single episode's amplification be attributed to both.

Measurement.

Not supplied, and the gap is ours. CTA and managed-futures positioning is free and weekly, from the CFTC Commitments of Traders. It measures trend-following exposure but not the broader interpretive mechanism A_{meta} describes. *What it costs:* A_{meta} is unmeasured, and it is the term through which Soros [70]'s reflexivity enters the model. The proxy would evidence the narrow claim — that positioning follows price — and not the wide one.

No range is stated.

7.6 Dual Amplification and Behavioural Phase Lock

src p.16

7.6.1 The mechanism

The introduction of the psycho-technical layer changes κ_{sys} from an endogenous market property into something the information system determines. Two amplifiers act on it, and the original's argument for each is a before-and-after contrast.

$H_{\text{soc}} \to 1$: the uniformity amplifier. Markets were once composed of millions of agents with diverse psychologies, expectations and horizons, and that diversity was a natural dampener: reactions propagated as waves. As homogeneity rises, agents consume the same material, perceive the same signals and form uniform expectations. When a reaction occurs, everyone reacts in the same direction, converting scattered individual actions into a collective response.

$I_{\text{imp}} \to 1$: the immediacy amplifier. Even when agents shared an intention, there was a lag between intention and execution — psychological hesitation, procedural friction. As immediacy rises that lag vanishes, and intent converts to action at once.

Phase lock. When both are near one, the system reaches behavioural phase lock, and the original's image is exact: reactions that used to resemble a *wave*, with peaks and propagation delays, come to resemble a *wall* — instantaneous, uniform, extreme.

Formally this is $B_{\text{phase}} = H_{\text{soc}} I_{\text{imp}}$ (ISSUES #10), which reaches one only when both reach one — a product, not a sum, precisely because either amplifier alone leaves a dampener intact. Diverse agents acting instantly still cancel out; synchronised agents acting slowly still queue.

7.6.2 A countervailing effect the original does not note

The psycho-technical layer raises κ_{sys} through H_{soc} and I_{imp}. But it also raises V_{meme}, and the original itself says that a higher V_{meme} raises λ_f, the forgetting rate — which *lowers* M_{shock}, which lowers κ_{sys}.

So Layer 4 acts on Layer 3 through three channels with two signs:

Channel	Sign on κ_{sys}	Source
$H_{\text{soc}} \uparrow$	$+$	synchronised reaction
$I_{\text{imp}} \uparrow$	$+$	simultaneous execution
$V_{\text{meme}} \uparrow \Rightarrow \lambda_f \uparrow \Rightarrow M_{\text{shock}} \downarrow$	$-$	faster forgetting

The net sign is not determined by the model as written, and the original assumes it is positive without saying so. The assumption is probably right — the first two channels act on every episode while the third acts only on the memory of old ones — but it is an assumption, and Chapter 10's simulations should test whether a high-V_{meme} configuration can produce a *lower* κ_{sys} than a low-V_{meme} one.

7.7 The Updated κ_{sys} Formula

7.7.1 Statement

The original writes E6 as $\kappa_{\text{sys}} = f(\kappa_{\text{base}}, V_{\text{USD}}, M_{\text{shock}}, A_{\text{meta}}) \times H_{\text{soc}} \times I_{\text{imp}}$ and never specifies $f(\cdot)$. The resolution of ISSUES #3 gives it explicitly and restores the two channels E1 had and E6 lost; the amendment ISSUES #3a, argued in Section 7.7.5, puts the two psychotechnical inputs on the same additive-one footing as the other five:

$$\begin{aligned}\kappa_{\text{sys}} \;&= \kappa_{\text{base}}\,(1 + a_V\,V_{\text{USD}})(1 + a_M\,M_{\text{shock}})(1 + a_A\,A_{\text{meta}}) \\ &\quad\times\; (1 + a_\sigma\,\sigma_{\text{sys}})\,(1 - \kappa_R\,R_{\text{US}})(1 + a_H\,H_{\text{soc}})(1 + a_I\,I_{\text{imp}}),\end{aligned} \tag{7.5}$$

with all $a > 0$, all latent inputs in $[0, 1]$, and the positivity constraint $\kappa_R < 1$. Every one of the seven drivers now enters the same way, so Assumption 7.1 below is a statement about the whole equation rather than about five of its seven terms.

7.7.2 Assumptions

Assumption 7.1 (Multiplicative channels)**.** Each driver enters as a factor $(1 + a\,x)$ rather than additively, so that a driver at zero leaves κ_{sys} unchanged and drivers compound rather than substitute.

Assumption 7.2 (Positivity)**.** $\kappa_R < 1$, so that $(1 - \kappa_R R_{\text{US}}) > 0$ for all $R_{\text{US}} \in [0, 1]$ and $\kappa_{\text{sys}} > 0$ throughout.

Assumption 7.3 (G_{pol} acts only through Layer 2)**.** Geopolitical power does not enter Equation (7.5) directly. E1 lists it as an argument of κ; we route it through Φ_{sys} and $s_{\text{US,sys}}$ instead.

Assumption 7.3 is a deliberate departure from E1 and the reason is double-counting: G_{pol} already enters both Layer-2 functions with its own coefficients (Section 6.6), and giving it a third coefficient here would make the three jointly unidentifiable from any data we have.

7.7.3 Why this form

Multiplicative rather than additive, because the drivers are not substitutes. A market with no crisis memory and extreme meta-reflexivity is not equivalent to one with moderate amounts of each; the mechanisms compound. The $(1 + ax)$ shape is the minimal form with that property that also leaves κ_{sys} at κ_{base} when every driver is switched off.

The rejected alternative was exponential, $\kappa_{\text{sys}} = \kappa_{\text{base}} \exp(\sum_i a_i x_i)$ over all seven drivers, which is strictly positive without needing Assumption 7.2 and is linear in logs. It was rejected on interpretability: with the multiplicative form each factor reads directly as "this channel adds ax per cent to convexity", and the R_{US} channel in particular reads as a *proportional flattening*, which is what the original describes. The exponential form would express flattening as a negative exponent, which is harder to reason about and no easier to estimate.

7.7.4 The restored channels

σ_{sys}. Without $(1 + a_\sigma \sigma_{\text{sys}})$ the model has no representation of $\kappa'(\sigma) > 0$, which E1 asserts and on which the entire convexity claim rests. Its restoration is what makes Proposition 5.2 a statement about the model rather than about a coincidence.

R_{US}. Without $(1 - \kappa_R R_{\text{US}})$ the fourth case study's argument — that internal U.S. risk flattens the dollar's convexity — has no equation behind it, as Section 6.3 notes from the other side. The channel says something specific and uncomfortable: a dollar issued by a politically unstable United States remains a haven but becomes a *worse shock absorber*. The level of demand is set in Layer 2; the transmission is degraded here.

7.7.5 Parameters and comparative statics

Table 7.1: Parameters of Equation (7.5). All are OPEN values in the registry.

Symbol	Channel	Range	Note
a_V	velocity of transmission	> 0	
a_M	crisis memory	> 0	
a_A	meta-reflexivity	> 0	
a_σ	systemic stress	> 0	restored from E1 (ISSUES #3)
κ_R	internal U.S. risk	$(0, 1)$	restored; Assumption 7.2
a_H	social homogeneity	> 0	additive-one form (ISSUES #3a)
a_I	impulse instantaneity	> 0	additive-one form (ISSUES #3a)

Differentiating Equation (7.5) gives the elasticities the case studies use. Writing κ_{sys} for the whole product,

$$\frac{\partial \ln \kappa_{\text{sys}}}{\partial \sigma_{\text{sys}}} = \frac{a_\sigma}{1 + a_\sigma \sigma_{\text{sys}}}, \qquad \frac{\partial \ln \kappa_{\text{sys}}}{\partial R_{\text{US}}} = \frac{-\kappa_R}{1 - \kappa_R R_{\text{US}}}, \qquad \frac{\partial \ln \kappa_{\text{sys}}}{\partial H_{\text{soc}}} = \frac{a_H}{1 + a_H H_{\text{soc}}}. \tag{7.6}$$

All seven have the same shape: positive (or, for R_{US}, negative), bounded, and *decreasing in the level of the driver*. Each channel has diminishing returns in logs — the first increment of homogeneity buys more acceleration than the tenth — which is the behaviour the original's "dual amplification" language describes and which the form E6 was originally written in did not deliver.

Why the third elasticity changed: the ISSUES #3a amendment. As E6 stood after ISSUES #3, H_{soc} and I_{imp} entered as *bare* multipliers rather than as $(1 + a\,x)$ factors. Three things followed, and none of them is what the original claims.

The log-elasticity was $1/H_{\text{soc}}$, which *diverges* as homogeneity approaches zero and is smallest when homogeneity is near one — the opposite of the intuition the chapter is built on, and inconsistent with the other five channels.

$\kappa_{\text{sys}} \to 0$ as $H_{\text{soc}} \to 0$. A perfectly heterogeneous market would transmit *nothing*: no reflexive acceleration, no convexity, the whole of Layer 3 switched off by one input going to its lower bound. The same held for I_{imp}, so a market of diverse participants with no impulse to act simultaneously would have $\kappa_{\text{sys}} = 0$ exactly.

And H_{soc} and I_{imp} were the only two of the seven drivers whose absence was *destructive* rather than merely neutral. Assumption 7.1 said that a driver at zero leaves κ_{sys} unchanged; for two of the seven, it did not.

The amendment fixes all three by putting the two inputs on the common footing:

$$H_{\text{soc}}\, I_{\text{imp}} \;\longrightarrow\; (1 + a_H\, H_{\text{soc}})(1 + a_I\, I_{\text{imp}}), \qquad a_H, a_I > 0. \tag{7.7}$$

The justification is that the original claims diversity *dampens* transmission, not that it annihilates it. Heterogeneous markets are slower to propagate a shock because participants disagree about what a price means; they are not markets in which a shock fails to propagate at all. Under Equation (7.7) the psycho-technical layer multiplies κ_{sys} by a factor in $[1, (1 + a_H)(1 + a_I)]$: it is inert when homogeneity and impulse are both at their floor, leaving the reflexive layer running at κ_{base} scaled by the other five drivers, and it amplifies by up to $(1 + a_H)(1 + a_I)$ when both saturate. The sign, the monotonicity and the "waves become walls" reading of Section 7.6 are all preserved; only the annihilation is removed.

Two consequences carry forward. Chapter 8 computes the fixed point $(H^*_{\text{soc}}, I^*_{\text{imp}})$ of the closed loop and now feeds it into a factor that is bounded away from zero, so the comparative statics of Chapter 10 do not degenerate on any admissible branch. And a_H and a_I are two more OPEN parameters, which raises the identification problem of Section 7.7.6 from five unidentified coefficients to seven. The amendment buys a defensible functional form at the cost of two parameters that the available data cannot separate; Chapter 11 must say so rather than report estimates of them.

7.7.6 Failure modes

Two, the annihilation problem having been removed by ISSUES #3a.

Assumption 7.2 is an assumption, not a construction: a calibration with $\kappa_R \geq 1$ produces a negative κ_{sys} and an inverted model, and Chapter 10 must check it. Nothing in the construction prevents it; only the constraint does.

And the seven latent drivers are jointly unidentified. Nothing in the available data separates a_V from a_M from a_A, since none of V_{USD}, M_{shock} or A_{meta} has a series; the amendment adds a_H and a_I to the same list, since neither H_{soc} nor I_{imp} has one either (Chapter 8). Seven coefficients multiply five unobservable inputs and two more that are themselves the output of a fixed-point calculation. In practice Chapter 11 can hope to identify κ_{sys} as a single state, not its decomposition, and this is the clearest instance in the book of the gap the author's standing constraint on that chapter requires it to state plainly rather than bury.

7.8 Hyper-Brittleness

src p.17

The dual amplification turns this layer into the engine of systemic fragility, and the original draws three consequences.

Loss of buffers. The shock-absorption capacity that came from agent diversity and from temporal lag disappears. Both were accidental — nobody designed them — and both are removed by the same technological change.

Hyper-brittleness. The system becomes ultra-sensitive: small shocks to σ_{sys}, including rumours and minor data surprises, can trigger non-linear moves almost instantaneously through an amplified κ_{sys}.

Failure of prediction and control. A wall-like collective response undermines both forecasting and policy. Linear and Gaussian-based models collapse, not because their parameters are wrong but because the data-generating process no longer has the shape they assume.

7.8.1 What is claimed, and what is shown

These are strong statements and the chapter should be clear about their evidential status. They are *implications of the model*, derived from Equation (7.5) plus the assumption that H_{soc} and I_{imp} have risen. They are not measurements. Neither H_{soc} nor I_{imp} has a series (Chapter 8), so the premise is argued rather than observed, and the conclusion inherits that status.

What *can* be examined is the reduced-form consequence: if κ_{sys} has risen, the dollar's response to a given change in stress should be larger now than it was, and larger when stress is high than when it is low. Section 7.9 looks.

7.9 Empirical Convexity

Chapter 5 tested the two-regime version of this question and found that the answer is conditional on the weight w_X given to non-financial stress inside the state variable, and is undetermined until that weight is estimated (Proposition 5.2, Section 5.6.5). This section removes the imposed break point altogether. The sample is cut into quintiles of σ_{sys} measured at the start of each month, the contemporaneous response of the dollar to a change

in σ_{sys} is estimated *inside* each quintile, and the five slopes are plotted against the quintile's mean stress level. If $\kappa'(\sigma_{\text{sys}}) > 0$, that profile should rise.

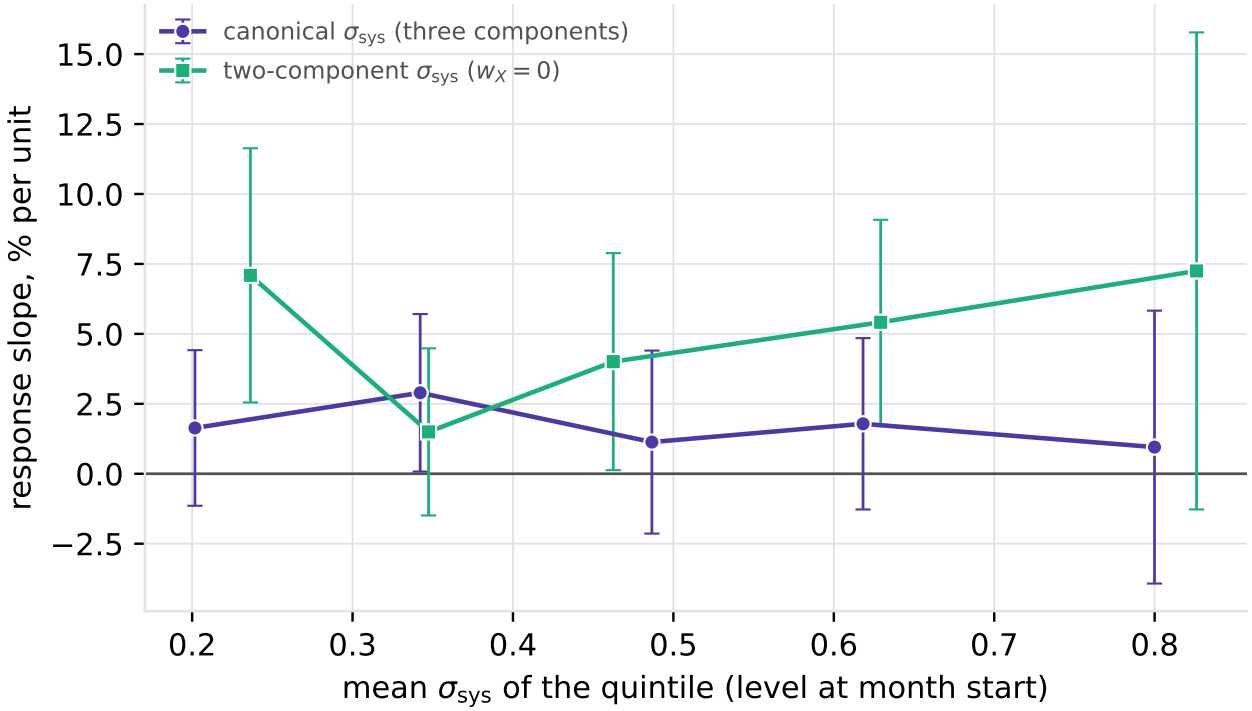


Figure 7.2: Empirical convexity: the dollar's response slope estimated separately within each quintile of σ_{sys}, with two-standard-error bars, under both constructions of the state variable. A rising profile would be evidence for $\kappa'(\sigma_{\text{sys}}) > 0$. Dependent variable `DTWEXAFEGS`, monthly, 246 observations. Source: author's calculation; script `f24_convexity.py`.

Table 7.2: Response slope by quintile of σ_{sys}, per cent per unit, with t-statistics.

Construction	Q1	Q2	Q3	Q4	Q5
Canonical (three components)	+1.64	+2.89	+1.13	+1.79	+0.95
t	+1.2	+2.1	+0.7	+1.2	+0.4
Two-component ($w_X = 0$)	+7.09	+1.50	+4.01	+5.41	+7.25
t	+3.1	+1.0	+2.1	+3.0	+1.7

7.9.1 Reading the profile

Under the canonical state variable the profile is flat. The five slopes lie between +0.95 and +2.89 with overlapping error bars and no trend; the highest is the second quintile and the lowest the fifth. There is no evidence here for $\kappa'(\sigma_{\text{sys}}) > 0$.

Under the two-component state variable the profile is a U, not a ramp. Slopes are high at both ends — +7.09 in the calmest quintile and +7.25 in the most stressed —

and lower in the middle. That shape is interesting and it is not what $\kappa'(\sigma_{\text{sys}}) > 0$ predicts, which is a monotone rise.

It is, however, strikingly close to the shape of the received view: a response that is strong at both extremes and weakest in the benign middle. The difference from the dollar smile is decisive and worth stating precisely. The smile claims the dollar *falls* when stress rises in calm conditions — a negative slope on the left arm. What Figure 7.2 shows is a *large positive* slope on the left arm. The dollar responds strongly to changes in stress when stress is low; it just responds in the same direction as when stress is high. The curve has the smile's shape in magnitude and not in sign, which is exactly the distinction Section 1.3 draws.

The honest summary. Neither profile supports $\kappa'(\sigma_{\text{sys}}) > 0$ as the model states it. One is flat; the other is U-shaped. Both are estimated on 49 or 50 monthly observations per bin, which is why the error bars are wide enough to accommodate several stories. Chapter 11 inherits the task of estimating κ_{sys} as a latent state within the full specification, where the other channels are controlled and the estimate does not have to come from 49 observations at a time.

7.10 Conclusion

Layer 3 is no longer a stable endogenous property of financial markets. That is the original's claim and this chapter has given it an equation: Equation (7.5), with $f(\cdot)$ written out, the σ_{sys} and R_{US} channels restored, and the crisis-memory formulas relocated from the end of the book to the layer they belong to.

Three things are worth carrying forward.

The chapter's central variable is **the least measurable in Part I**. Four of the six drivers in Equation (7.5) — κ_{base}, V_{USD}, M_{shock}, A_{meta} — have no series, and the two that are measurable, σ_{sys} and R_{US}, are the two the original had dropped. That is an awkward position for the variable that scales every channel in Layer 0.

The layer contains **a countervailing effect the original does not note** (Section 7.6.2): the same psycho-technical change that raises κ_{sys} through homogeneity and immediacy lowers it through faster forgetting. The net sign is assumed, not derived.

And the layer's headline empirical prediction — that convexity rises with stress — is **not visible in the data at monthly frequency** under either construction of the state variable (Section 7.9). Chapter 17 records that alongside the Layer-1 verdict. The two findings are the same finding seen twice: the dollar's non-linearity is real in the sense that a linear model cannot produce the 2008 and 2020 episodes, and it is not the specific non-linearity — a sign change, or a monotone rise in convexity — that the original specifies.

Chapter 8

Layer 4 — The Psycho-Technical Layer: H_{soc} and I_{imp}

8.1 The Meta-Platform Argument

Three strands of the finance literature stand behind this layer and are cited once here rather than repeatedly below. Sentiment as a modelled object rather than a residual is Barberis et al. [13]. The claim that attention, not information, drives a large share of retail order flow — which is what I_{imp} formalises — is measured by Barber et al. [12] on Robinhood users, who find intense attention-driven buying forecasts negative returns. That trading now responds to signals at minute frequency is documented by Chinco et al. [23]. Echo-chamber formation, the mechanism behind A_{alg}, is Bakshy et al. [10] for the measurement and Sunstein [72] for the argument; narrative transmission is Shiller [68] and, at book length, Shiller [69].

Layers 1 to 3 describe what the original calls the operating system of the dollar: the state of the market, the structural foundations of trust, and the reflexive machinery that turns a shock into a price. That description was complete for as long as the operating system ran directly on the world. It no longer does. It runs on top of a second platform — an algorithmically governed information and behaviour ecosystem — and the properties of that platform now enter the properties of the system above it.

This is the original's claim in §4.4 and it is a strong one, so it is worth separating into the two assertions it actually contains.

The first is about *beliefs*: the distribution of what market participants think has narrowed, because the mechanisms that form beliefs have been concentrated. The second is about *latency*: the interval between forming an intention and executing it has collapsed, because the mechanisms that execute intentions have been optimised for frictionlessness. Neither assertion is about the dollar. Both are about the substrate on which every other layer of the model runs, which is why this layer is placed above them rather than beside them.

Definition 8.1 (Layer 4)**.** The psycho-technical layer generates no demand D_i and no supply S_j. It governs two meta-modulators, $H_{\text{soc}} \in [0,1]$ and $I_{\text{imp}} \in [0,1]$, which enter the model only through their effect on other layers: they raise κ_{sys} in Layer 3 (Equation (7.5)), they raise the volatility of R_{US} in Layer 2 (Section 6.3), and they accelerate the erosion of $\Psi_{\text{str}}(t)$ in Layer 5 (Chapter 9).

The definition carries a methodological cost that should be stated at the outset rather than discovered in Chapter 11. A layer that acts only through other layers is a layer whose coefficients are identified only through other layers' coefficients. Nothing in this chapter can be estimated on its own, and — as Section 8.9 sets out — nothing in it is currently measured at all.

Why this is not simply "sentiment". Behavioural finance has a long tradition of adding a sentiment term to an asset-pricing model, and a reader could reasonably suspect that Layer 4 is that tradition in new clothing. It is not, and the difference is structural rather than rhetorical. A sentiment term is an *input*: it moves, and the model responds. H_{soc} and I_{imp} are *modulators*: they do not move the dollar, they change how much the dollar moves in response to everything else. In the notation of Chapter 2 they enter κ_{sys}, the exponent on the mechanical bracket, not the bracket itself. A model with a sentiment term and a model without one differ in their fitted residuals; a model with Layer 4 and a model without one differ in the *slope* of the response function, which is a different and stronger claim.

What the original gets right and what it overstates. The mechanism is plausible and the direction is defensible. What the original does not do — and what this chapter must do — is say what happens when the two state variables are allowed to determine each other, which they must, because Section 8.4 shows that each appears in the other's equation. That is ISSUES #13, and it converts a pair of formulas into a dynamical system with an equilibrium, a stability condition, and, under conditions the model does not exclude, more than one equilibrium. Half of this chapter is the consequence of taking the original's own two equations seriously enough to solve them simultaneously.

8.2 Variable Architecture

src p.18

The layer has four drivers, one defensive variable, two state variables and two consequence variables. Table 8.1 sets out the architecture; the cards follow.

8.2.1 The four drivers

A_{alg}, algorithmic reinforcement strength. The intensity with which recommendation systems personalise, reinforce and steer belief — filter bubbles and echo chambers, in the original's words. It is the only driver that is a property of an *algorithm* rather than of an observable market structure, and it is correspondingly the hardest to measure (Section 8.9).

C_{src}, content source concentration. A Herfindahl–Hirschman index over information sources. It is the one driver in this layer with a well-established measurement tradition, since media concentration has been measured for decades. A rising C_{src} narrows the set of available narratives without changing any individual's disposition to believe them.

Table 8.1: The psycho-technical layer. Every variable is normalised to $[0, 1]$ (ISSUES #26). "Enters" names the equation in which the variable appears, not the quantity it ultimately affects.

Role	Symbol	Enters	Reads as
Drivers			
	A_{alg}	E7	strength of algorithmic reinforcement
	C_{src}	E7	concentration of content sources
	P_{AI}	E7, E15	generative-AI share of content
	T_{exp}	E7	cumulative exposure time
Defensive			
	R_{div}	E7	residual diversity resistance
States			
	H_{soc}	E6, E7, E8, E15	homogeneity of collective consciousness
	I_{imp}	E6, E7, E7b, E8	impulse efficiency: thought to action
Consequences			
	V_{meme}	E8, E12	narrative transmission velocity
	B_{phase}	descriptive	behavioural phase lock

P_{AI}, generative-AI penetration. The share of the information ecosystem produced by generative models. Its role here is not that machine-written content is more persuasive than human-written content, a claim the original does not make and this book does not need. Its role is that generative production removes the *cost* constraint on narrative supply: a story can be created, tailored to a segment and propagated at a marginal cost approaching zero. P_{AI} appears again in Chapter 7 as one of the terms that can re-activate a decayed crisis memory (E15), which is the same property seen from the other side.

T_{exp}, exposure time. Cumulative time a typical participant spends inside a given narrative stream. It is the driver that carries the original's implicit theory of belief formation: repetition, not argument, is what converts exposure into acceptance.

8.2.2 The defensive variable

R_{div}, residual diversity resistance. The remaining capacity of social and cognitive structures to resist homogenisation: educational diversity, non-algorithmic public space, social trust, norms that reward independent reasoning.

Two properties of R_{div} deserve emphasis because they are not obvious from the original's prose and they are visible in Figure 8.1.

First, R_{div} is the only variable in the layer with a negative sign. Every other quantity here pushes toward uniformity; this is the one that pushes back, and the model's entire answer

to "can this be resisted?" runs through it.

Second — and this is a genuine finding of the functional form rather than a restatement of it — R_{div} **does not flatten the transition to homogeneity; it postpones it**. In E7 the defensive variable enters as a subtraction inside the logistic, so raising it slides the whole curve to the right without changing its steepness. A society with high residual diversity is not one in which homogenisation proceeds gently. It is one in which homogenisation begins later and then proceeds exactly as steeply as it would have. Figure 8.1(b) is that statement. Whether it is a desirable property of the model is a question for Chapter 19; that it is a property of the model is not in doubt.

8.2.3 The two state variables

H_{soc}**, homogeneity of collective consciousness.** The convergence of cognition, belief and shared narrative. It approaches 1 from below, with 1 representing near-total uniformity. What it measures is not agreement about the truth but agreement about the *frame*: participants who disagree violently about whether the dollar will rise may nonetheless share a single account of what would make it rise, and it is the second kind of agreement that synchronises their behaviour.

I_{imp}**, impulse efficiency.** How efficiently and rapidly the socio-technical system converts a signal into an action, bypassing deliberation. It also tends to 1, representing the limit in which thinking and acting are the same operation.

The two are distinct and the distinction carries the chapter. H_{soc} synchronises *intention*; I_{imp} synchronises *timing*. A market with high homogeneity and low impulse efficiency agrees about what to do and does it over a week. A market with low homogeneity and high impulse efficiency acts instantly in a hundred directions that cancel. Neither is dangerous. It is the conjunction that is, which is what B_{phase} names.

8.2.4 The two consequence variables

The original defines V_{meme} and B_{phase} verbally and gives neither an equation. ISSUES #10 gives both one, and the two resolutions are deliberately different in kind.

B_{phase}**, behavioural phase lock.**

$$B_{\text{phase}} \;=\; H_{\text{soc}}\, I_{\text{imp}} \;\in [0,1]. \tag{8.1}$$

This is the bare product, and here the bare product is *right*. Chapter 7 amended κ_{sys} so that H_{soc} and I_{imp} no longer enter it as bare multipliers (ISSUES #3a), on the ground that a perfectly heterogeneous market should still transmit a shock, merely more slowly. No such argument applies here, because B_{phase} is not a transmission channel but a *definition of synchrony*. A market with zero homogeneity has zero behavioural synchronisation, exactly, and $B_{\text{phase}} = 0$ is the correct reading rather than a degenerate one. The two cases are worth contrasting explicitly, because it would be easy to apply the ISSUES #3a amendment as a

blanket rule: the question is always whether the quantity is a rate at which something is transmitted or a measure of how aligned something is, and only the first kind needs a floor.

B_{phase} enters no equation. It is a reporting index, and Chapter 13 uses it to describe episodes rather than to predict them.

V_{meme}, narrative transmission velocity.

$$V_{\text{meme}} \;=\; v_0\, H_{\text{soc}}\,(1 + b_A\, A_{\text{alg}})(1 + b_P\, P_{\text{AI}}), \qquad v_0, b_A, b_P > 0. \tag{8.2}$$

Here the form follows the original's own description: velocity is a function of the infectiousness of the environment (H_{soc}) and of the technological enablers (A_{alg}, P_{AI}). The two enablers enter as $(1 + b\,x)$ factors, on the ISSUES #3a reasoning — switching off the algorithm does not reduce narrative velocity to zero, it reduces it to the velocity of unassisted human transmission. H_{soc} enters bare, because a narrative cannot propagate through a population that shares no frame with which to receive it; there is no residual velocity to preserve.

Unlike B_{phase}, V_{meme} does real work. It is one of the three arguments of the erosion speed $k_\Psi(H_{\text{soc}}, V_{\text{meme}}, I_{\text{imp}})$ of E8 (Chapter 9), and it enters the forgetting rate λ_f in E12 (Section 7.4). Equation (8.2) is what makes both computable rather than verbal.

A countervailing channel, restated. Section 7.6.2 noted that V_{meme} raises λ_f, and a higher forgetting rate *lowers* M_{shock}, which lowers κ_{sys}. So this layer reaches κ_{sys} by two routes with opposite signs: directly through H_{soc} and I_{imp}, which raise it, and indirectly through $V_{\text{meme}} \to \lambda_f \to M_{\text{shock}}$, which lowers it. The original notes only the first. Which dominates is an empirical question that cannot be answered while v_0, b_A, b_P and λ_f are all unestimated, and Chapter 10 records it as an open sign rather than resolving it by assertion.

8.3 E7: The Homogeneity Convergence Function

8.3.1 Statement

$$H_{\text{soc}}(t) \;=\; \Big[1 + \exp\!\Big(-\eta\,(\kappa_A\, I_{\text{imp}} - R_{\text{div}})\Big)\Big]^{-1}, \qquad \kappa_A \;\equiv\; A_{\text{alg}} \cdot C_{\text{src}} \cdot P_{\text{AI}} \cdot T_{\text{exp}}. \tag{8.3}$$

This is the original's E7 with one piece of notation added: the product of the four drivers is named κ_A, because it appears as a single quantity in everything that follows and because naming it makes the uniqueness condition of Section 8.5 readable.

8.3.2 Assumptions

Assumption 8.1 (Multiplicative drivers)**.** The four drivers enter as a product, so each is necessary: $\kappa_A = 0$ if any one of them is zero.

Assumption 8.2 (Bounded inputs)**.** $A_{\text{alg}}, C_{\text{src}}, P_{\text{AI}}, T_{\text{exp}}, R_{\text{div}}, I_{\text{imp}} \in [0, 1]$, hence $\kappa_A \in [0, 1]$ and the logistic argument lies in $[-\eta, \eta]$ (ISSUES #26).

8.3.3 Functional-form justification

The logistic is not decoration. H_{soc} is a bounded share — the fraction of a population sharing a frame — and three properties are required of any function producing it: it must map an unbounded drive into $[0, 1]$; it must be monotone in that drive; and its *marginal* effect must be small at both ends and largest in the middle, because converting the last independent thinker is as hard as converting the first. The logistic is the standard function with all three, and is the same object used for the regime weight $\Lambda(\cdot)$ in Chapter 2 and for σ_{sys} in Chapter 5. Using one functional form for all three is a modelling economy worth naming: this book has exactly one S-shape.

Two alternatives were rejected. A linear-in-drivers form would require truncation at both ends and would give constant marginal returns to homogenisation, which contradicts the mechanism. A power form $H_{\text{soc}} = (\kappa_A I_{\text{imp}})^\gamma$ is bounded and monotone but has no R_{div} that acts as a threshold: it can only scale resistance, not locate it, and the original's R_{div} is explicitly a thing that is overcome at a point.

Assumption 8.1 is the strongest assumption in the block, and it says something falsifiable: homogenisation requires *all four* drivers, so a society with maximal algorithmic reinforcement and zero exposure time has $\kappa_A = 0$ and $H_{\text{soc}} = \Lambda(-\eta R_{\text{div}})$, a floor set by resistance alone. That is a substantive claim about substitutability between the drivers — namely, that there is none — and it is the assumption an empirical implementation would attack first.

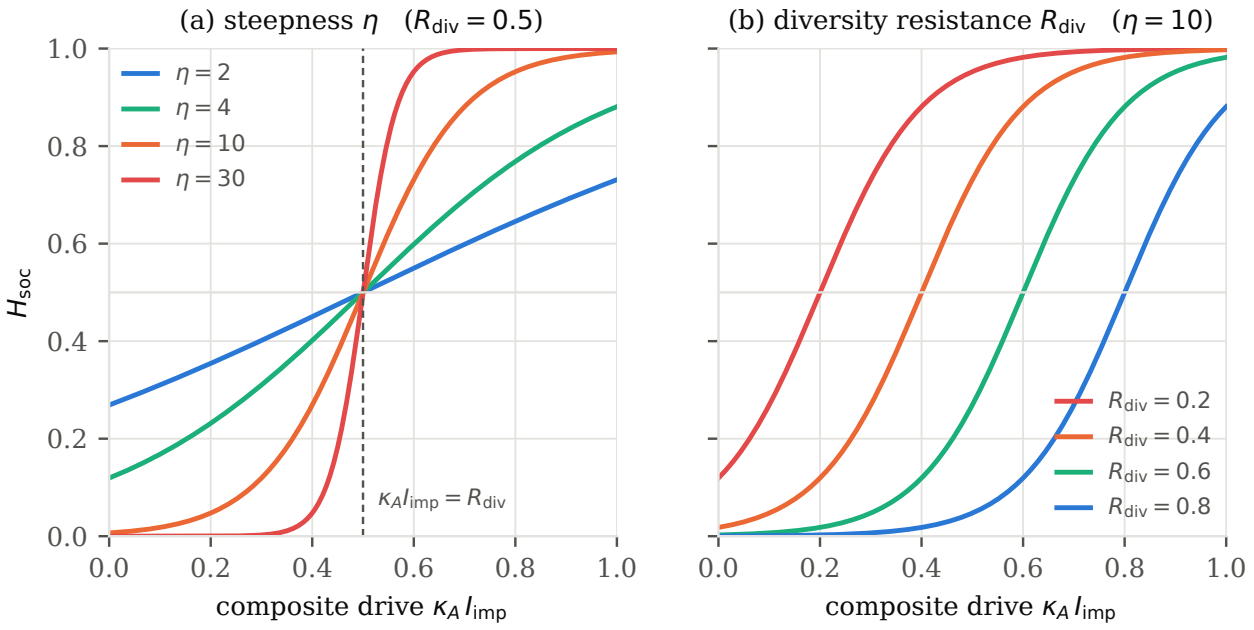


Figure 8.1: E7, the homogeneity convergence function, as a function of the composite drive $\kappa_A I_{\text{imp}}$. Panel (a) varies the steepness η at fixed $R_{\text{div}} = 0.5$: all curves cross at the turning point $\kappa_A I_{\text{imp}} = R_{\text{div}}$, where $H_{\text{soc}} = \frac{1}{2}$ by construction, and η decides only whether the transition is a slope or a cliff. Panel (b) varies the defensive variable R_{div} at fixed $\eta = 10$: raising resistance slides the curve right without flattening it. Source: author's construction; script `f25_hsoc.py`.

8.3.4 Parameters

Table 8.2: Parameters of Equation (8.3).

Symbol	Meaning	Range	Provenance
η	steepness of the logistic	> 0; < 4 maintained	original; bound from ISSUES #30a
κ_A	driver product	$[0, 1]$	composite, ISSUES #30
R_{div}	diversity resistance	$[0, 1]$	original

8.3.5 Properties

$H_{\text{soc}} \in (0, 1)$ strictly, so it never attains either bound and the "$H_{\text{soc}} \to 1$" of the original's prose is a limit, not a state. H_{soc} is strictly increasing in each of the four drivers and in I_{imp}, and strictly decreasing in R_{div}. The turning point is at $\kappa_A I_{\text{imp}} = R_{\text{div}}$, where $H_{\text{soc}} = \frac{1}{2}$ and the marginal effect of every driver is maximal: $\partial H_{\text{soc}}/\partial(\kappa_A I_{\text{imp}}) = \eta/4$ there, and smaller everywhere else. Both arguments are dimensionless, so the equation is dimensionally consistent for any η.

8.3.6 Comparative statics

Writing $\Lambda' = H_{\text{soc}}(1 - H_{\text{soc}})$ for the logistic derivative,

$$\frac{\partial H_{\text{soc}}}{\partial R_{\text{div}}} = -\eta\,\Lambda', \qquad \frac{\partial H_{\text{soc}}}{\partial I_{\text{imp}}} = \eta\,\kappa_A\,\Lambda', \qquad \frac{\partial H_{\text{soc}}}{\partial A_{\text{alg}}} = \eta\,\Lambda'\,I_{\text{imp}}\,C_{\text{src}}\,P_{\text{AI}}\,T_{\text{exp}}. \tag{8.4}$$

Every one is proportional to Λ', which is why the layer is least responsive to intervention exactly where it is most extreme. Chapter 19 needs this: a policy aimed at R_{div} has almost no effect on a society already at $H_{\text{soc}} = 0.95$, and its whole leverage lies in the middle of the range.

8.3.7 Failure modes

Two, and the second is the subject of the rest of the chapter.

Assumption 8.1 makes κ_A a product of four quantities each in $[0, 1]$, so κ_A is typically *small*: four independent drivers at 0.7 give $\kappa_A = 0.24$. Unless the drivers are strongly correlated in practice — which they plausibly are, since the same platforms supply all four — the composite drive $\kappa_A I_{\text{imp}}$ will sit well below R_{div} for most admissible parameter values, and E7 will return a low H_{soc}. The original's narrative of approaching uniformity therefore requires either near-saturated drivers or a low R_{div}, and which of the two it is asserting is not stated.

And E7 contains I_{imp}, which is itself an equation. That is not a failure of E7; it is the reason Section 8.5 exists.

8.4 E7b: The Impulse Function

src p.21

8.4.1 Statement

The original gives E7b in two forms, a general one and an approximation, and this book uses the approximation because the general form names no functional shape:

$$I_{\text{imp}} \;=\; \frac{S_{\text{access}} \cdot C_{\text{loop}} \cdot P_{\text{habit}} \cdot E_{\text{sys}}}{1 + R_{\text{fric}}}. \tag{8.5}$$

8.4.2 The five drivers

S_{access}, speed of access, is the think-to-click latency: it rises with seamless interfaces and network speed. R_{fric}, resistance friction, is procedural and psychological drag — doubt, confirmation steps, complex interfaces — and it falls as platforms optimise toward one click. C_{loop}, cognitive-loop shortening, is the compression of the think–decide–act cycle, and rises when machine suggestion substitutes for personal analysis. P_{habit}, programmed habit strength, is the degree to which behaviour has become automatic. E_{sys}, embedded system coupling, is the integration of social, payment, commerce and investment layers into single ecosystems.

8.4.3 Functional-form justification

The form is a product of enablers divided by $(1 + R_{\text{fric}})$, and both halves are motivated.

The numerator is multiplicative because the five are *serial* stages of one process: a signal must be received quickly, converted to a decision quickly, acted on habitually, and executable inside the same ecosystem. A break at any stage stops the chain, which is exactly what a product represents and a sum does not. This is the same argument as Assumption 8.1, applied to a pipeline rather than to a set of influences.

The denominator is $(1 + R_{\text{fric}})$ rather than $(1 - R_{\text{fric}})$, and the choice matters. With $(1 - R_{\text{fric}})$, friction at its maximum would make I_{imp} infinite — the wrong sign — and friction near 1 would make it explode. With $(1 + R_{\text{fric}})$, friction halves impulse efficiency at $R_{\text{fric}} = 1$ and cannot ever eliminate it. That is the correct economics: friction slows action, it does not forbid it. It is also, structurally, the same choice ISSUES #3a made for κ_{sys} in Chapter 7, and the same one ISSUES #28 made for S_{off} in Chapter 4. This book now has a consistent convention: **a dampening quantity enters as $1/(1 + x)$ or $e^{-\delta x}$, never as $(1 - x)$, so that no single input can switch a channel off entirely.**

8.4.4 ISSUES #13: closing the loop

The original states that H_{soc} and I_{imp} "are not independent; they evolve dynamically and mutually reinforce one another". Its equations do not say that. E7 depends on I_{imp}; E7b depends on five drivers, none of which is H_{soc}. As written the relationship is one-way, and the text's claim of mutual reinforcement has no formal counterpart. This is ISSUES #13, and the decided resolution closes the loop at the one driver where the mechanism is unambiguous:

$$C_{\text{loop}} \;=\; c_{L0} + c_{L1}\, H_{\text{soc}}, \qquad c_{L0} \geq 0,\; c_{L1} > 0. \tag{8.6}$$

The justification is the original's own definition of C_{loop}: cognitive-loop shortening "increases when AI-driven suggestions replace personal analysis". Suggestion replaces analysis most easily where the suggestion matches what everyone else already believes. A homogeneous population does not need to deliberate, because deliberation is what one does when the frame is contested. So homogeneity shortens the cognitive loop, and the cognitive loop raises impulse efficiency, and impulse efficiency raises homogeneity. That is the mutual reinforcement the text asserts, located in a specific driver rather than left as a general claim.

The coefficients are named c_{L0} and c_{L1} rather than c_0 and c_1 to avoid colliding with the E10 intercept c_0 of Chapter 10.

Substituting Equation (8.6) into Equation (8.5) and collecting everything that does not depend on H_{soc} into one constant gives the reduced impulse map:

$$I_{\text{imp}} \;=\; K\,(c_{L0}{+}c_{L1}H_{\text{soc}}) \;=\; g_0{+}g_1\,H_{\text{soc}}, \qquad K \equiv \frac{S_{\text{access}}\,P_{\text{habit}}\,E_{\text{sys}}}{1+R_{\text{fric}}}, \quad g_0 \equiv Kc_{L0}, \;\; g_1 \equiv Kc_{L1}. \tag{8.7}$$

Assumption 8.3 (Admissibility)**.** $g_0 \geq 0$ and $g_0 + g_1 \leq 1$, so that $I_{\text{imp}} \in [0,1]$ for every $H_{\text{soc}} \in [0,1]$, as ISSUES #26 requires.

Assumption 8.3 is a restriction on the primitive parameters, not a free choice, and it does real work in Section 8.5: because $g_0 \geq 0$ forces $g_1 \leq 1$, the uniqueness condition simplifies to a condition on η alone.

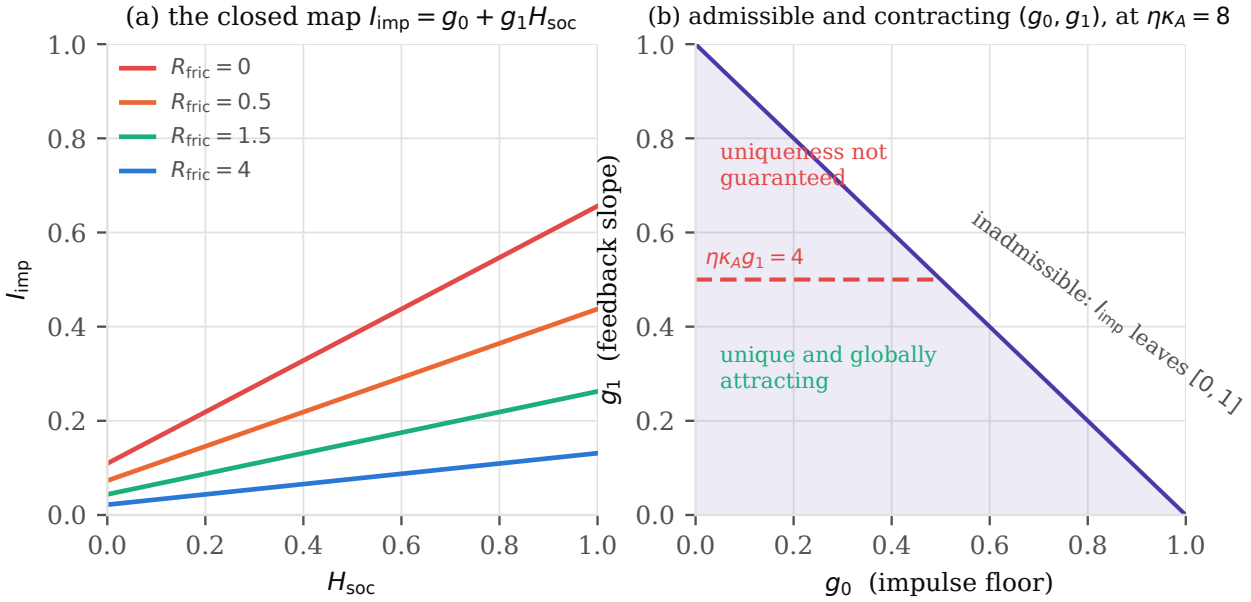


Figure 8.2: E7b after the ISSUES #13 closure. Panel (a): the reduced map $I_{\text{imp}} = g_0 + g_1 H_{\text{soc}}$ at four levels of friction, with $S_{\text{access}} = P_{\text{habit}} = E_{\text{sys}} = 0.9$, $c_{L0} = 0.15$, $c_{L1} = 0.75$. Friction scales the whole line toward zero; it cannot make the map decreasing. Panel (b): the admissible region for (g_0, g_1) required by Assumption 8.3, with the uniqueness boundary $\eta\kappa_A g_1 = 4$ of Proposition 8.2 drawn at $\eta\kappa_A = 8$. Source: author's construction; script `f26_iimp.py`.

What the closure costs. Two new OPEN parameters, c_{L0} and c_{L1}, neither of which is measurable today. The alternative was to leave the loop open and write, in place of the original's claim of mutual reinforcement, an explicit statement that the model makes the relationship one-way. That would have been honest and would have cost nothing; it was rejected because it discards a mechanism the original asserts and because the one-way model has no equilibrium concept at all — I_{imp} would simply be an exogenous input to H_{soc}, and the layer would have no dynamics to speak of. The closure is what makes the rest of this chapter possible, and the two parameters are the price.

8.5 The Closed Loop: Existence, Uniqueness, Stability

8.5.1 One map instead of two equations

With Equation (8.7) in hand, E7 and E7b are no longer two formulas to be evaluated in some order. They are one map from $[0,1]$ to itself:

$$T(H) \;=\; F\Big(G(H)\Big), \qquad G(H) = g_0 + g_1 H, \qquad F(I) = \Big[1 + \exp\Big(-\eta(\kappa_A I - R_{\text{div}})\Big)\Big]^{-1}, \tag{8.8}$$

and the psycho-technical state of the system is a fixed point $H^* = T(H^*)$, with $I^*_{\text{imp}} = G(H^*)$.

Three questions follow immediately and the original asks none of them. Does a fixed point exist? Is it unique? If the system is displaced, does it return? Figure 8.3 shows the whole answer geometrically; the three propositions state it. Proofs are in Chapter D.

8.5.2 Existence

Proposition 8.1 (Existence, unconditional)**.** *For every admissible parameter vector, T has at least one fixed point in $(0,1)$.*

Proof sketch. G is affine and maps $[0,1]$ into $[0,1]$ by Assumption 8.3; F is continuous and maps $[0,1]$ into $(0,1)$ strictly, since the logistic never attains its bounds. Hence $T = F \circ G$ is continuous with $T([0,1]) \subset (0,1)$. Then $T(0) - 0 = T(0) > 0$ and $T(1) - 1 < 0$, so $T(H) - H$ changes sign on $[0,1]$ and the intermediate value theorem gives a root. □

This is worth stating plainly because it is the one result in the chapter that requires no assumption at all. **The psycho-technical layer always has an equilibrium.** The question is never whether the loop settles somewhere; it is where, and whether that somewhere is unique.

8.5.3 Uniqueness

Proposition 8.2 (Contraction and uniqueness)**.** $\sup_{H \in [0,1]} |T'(H)| = \frac{1}{4}\,\eta\,\kappa_A\,g_1$*. Hence if*

$$\eta\,\kappa_A\,g_1 \;<\; 4 \tag{8.9}$$

then T is a contraction on $[0,1]$: the fixed point is unique, globally attracting from any starting value, and iteration converges geometrically at rate $\frac{1}{4}\eta\kappa_A g_1$. Because Assumption 8.2 gives $\kappa_A \le 1$ and Assumption 8.3 gives $g_1 \le 1$, the condition $\eta < 4$ alone is sufficient.

Proof sketch. By the chain rule $T'(H) = F'(G(H))\, g_1$ and $F'(I) = \eta\kappa_A\, \Lambda(1-\Lambda)$ where $\Lambda = F(I)$. The factor $\Lambda(1-\Lambda)$ attains its maximum $\frac{1}{4}$ at $\Lambda = \frac{1}{2}$ and is attained whenever $\kappa_A G(H) = R_{\text{div}}$ has a solution in the range. The stated supremum follows, and the Banach fixed-point theorem gives the rest. □

Three remarks. The bound is *sharp*, not merely sufficient for a contraction: the supremum is attained. The condition involves the three quantities one would expect — steepness, driver intensity, feedback slope — and is symmetric in them, so a very steep logistic can be offset by weak drivers or by high friction, which raises R_{fric} and so lowers g_1. And $T' \geq 0$ everywhere, since every factor is positive, so convergence is monotone: the loop never oscillates, it simply climbs or falls into its equilibrium.

Assumption 8.4 (Maintained: $\eta < 4$)**.** Throughout Parts I–III of this book, $\eta < 4$ is maintained, so Proposition 8.2 applies and $(H^*_{\text{soc}}, I^*_{\text{imp}})$ is the unique, globally attracting equilibrium of the layer.

This is ISSUES #30a, and it is an assumption in the strict sense: nothing measures η, and nothing in the original's argument bounds it. Assumption 8.4 is adopted because every comparative static in Chapter 10, every simulation, and every case study would otherwise have to declare which branch it is on before it could say anything. Section 8.7 states in full what is being given up, and it is not nothing.

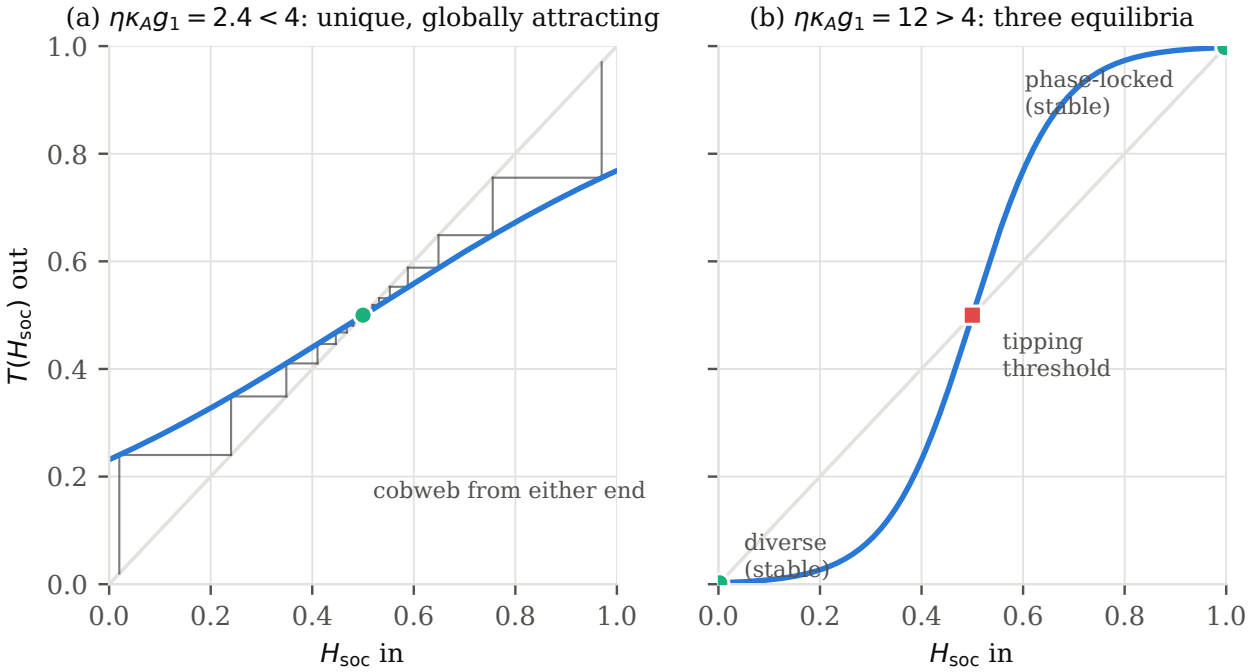


Figure 8.3: The closed loop $T(H) = F(G(H))$ against the diagonal, at $\kappa_A = 0.8$, $g_0 = 0.15$, $g_1 = 0.75$. Panel (a), $\eta = 4$ so $\eta\kappa_A g_1 = 2.4 < 4$: one crossing, and the cobwebs from both ends walk into it (Proposition 8.2). Panel (b), $\eta = 20$ so $\eta\kappa_A g_1 = 12 > 4$: three crossings, the outer two locally stable (circles, $T' = 0.03$) and the middle one unstable (square, $T' = 3.0$), which is Proposition 8.3. Source: author's construction; script `f27_fixedpoint.py`.

8.5.4 The social amplification factor

Closing the loop does more than tidy the model, and this subsection states the consequence the original does not reach.

Let θ be any parameter of the layer — R_{div}, one of the four drivers, R_{fric}. Differentiating $H^* = T(H^*;\theta)$ implicitly,

$$\frac{\partial H^*}{\partial \theta} = \underbrace{\frac{1}{1 - T'(H^*)}}_{\text{amplification}} \times \underbrace{\frac{\partial T}{\partial \theta}}_{\text{direct effect}}, \tag{8.10}$$

which is well defined and positive-signed under Assumption 8.4, since $T' < 1$ there.

The multiplier $1/(1-T')$ exceeds one whenever the loop is closed at all. **Every driver's effect on homogeneity is larger in equilibrium than its direct effect, and the gap is the loop.** A change in R_{div} lowers H_{soc} directly; the lower H_{soc} lengthens the cognitive loop, which lowers I_{imp}, which lowers H_{soc} again; and the sum of that geometric series is the multiplier. At $\eta\kappa_A g_1 = 2.4$, as in Figure 8.3(a), $T' = 0.6$ and the multiplier is 2.5: a policy that directly moves homogeneity by one point moves it by two and a half in equilibrium.

The multiplier diverges as $\eta\kappa_A g_1 \to 4$. This is the formal content of the original's "hyper-brittleness": not that the system becomes more sensitive to shocks in some loose sense, but that the equilibrium response to *any* parameter becomes unboundedly large as the loop approaches the contraction boundary — and at the boundary the equilibrium ceases to be unique. Equation (8.10) is, to our knowledge, not in the original, and it follows from nothing but taking its own two equations simultaneously.

8.5.5 Comparative statics at the equilibrium

Applying Equation (8.10) to each parameter, with $\Lambda' = H^*(1 - H^*)$ and $M \equiv 1/(1 - T'(H^*)) > 1$:

$$\frac{\partial H^*}{\partial R_{\text{div}}} = -M\,\eta\Lambda' < 0, \qquad \frac{\partial H^*}{\partial \kappa_A} = M\,\eta\Lambda'\,G(H^*) > 0, \qquad \frac{\partial H^*}{\partial R_{\text{fric}}} < 0. \tag{8.11}$$

The third follows because R_{fric} enters only through K, lowering both g_0 and g_1; it is the only parameter that reduces the amplification factor at the same time as it reduces the direct effect, which makes friction the most efficient point of intervention in the layer and is the analytical basis for Chapter 19's treatment.

8.6 Impact on the Rest of the Model

src p.22

The original's §4.4.3 gives a four-row table of the layer's effects. It is reproduced here as Table 8.3, re-set with this book's canonical layer numbering and with the equations that now carry each claim — which is the substantive change, since in the original three of the four rows had no equation behind them.

Table 8.3: Multi-dimensional impact of the psycho-technical layer, after the original's table on p. 22, re-set with canonical layer numbers and the equation that now carries each mechanism. Arrow scale per `BRIEF.md`: ↑ small, ↑↑ large, ↑↑↑ extreme.

Affected layer	Key variable	Direction	Mechanism and the equation that carries it
Layer 3, reflexive (Chapter 7)	κ_{sys}	↑↑↑	Dual amplification: H_{soc} synchronises intention, I_{imp} synchronises timing, so $B_{\text{phase}} \to 1$ and reactions turn from waves into walls. Now carried by the two factors $(1 + a_H H_{\text{soc}})(1 + a_I I_{\text{imp}})$ in Equation (7.5) (ISSUES #3a), bounded above by $(1 + a_H)(1 + a_I)$ rather than unbounded.
Layer 2, structural (Chapter 6)	R_{US}	↑↑	H_{soc} amplifies divisive narratives, I_{imp} converts them into immediate action, so minor events escalate nationally within hours and Φ_{sys} falls with them. This is the *volatility* of R_{US} rather than its level, and Section 6.3 states plainly what representing it as fat tails rather than as an interaction term costs.
Layer 5, secular (Chapter 9)	$\Psi_{\text{str}}(t)$	↑↑	Ultra-fast tipping. H_{soc} and V_{meme} spread a destabilising narrative and I_{imp} converts it to behaviour, so erosion that would take decades compresses into months. Carried by the erosion speed $k_\Psi(H_{\text{soc}}, V_{\text{meme}}, I_{\text{imp}})$ in E8, computable because Equation (8.2) gives V_{meme} an equation (ISSUES #10).
Policy (Chapter 19–Chapter 21)	reflex	↑↑↑	Loss of strategic control: policymakers are inside the same loop and must react to narratives rather than shape them. This row has no equation in this book either, and Chapter 19 treats it as a constraint on the feasible policy set rather than as a modelled quantity.

One correction to the table's fourth row. The original describes the policy effect as total loss of strategic control. The model as now written does not support "total". Under Assumption 8.4 the layer has a unique equilibrium and a finite amplification factor M, so policy retains leverage of exactly $M\,\partial T/\partial\theta$ on any parameter it can reach — diminished, not eliminated. "Total loss" would require the amplification factor to diverge, which happens only at the contraction boundary. The row is kept, because the direction is right and the mechanism is the original's; the word "total" is corrected to a statement about magnitude, and Chapter 19 works with the finite multiplier rather than with despair.

A row the original does not have. The countervailing channel of Section 8.2.4 belongs in this table and is not in it: $V_{\text{meme}} \uparrow\Rightarrow \lambda_f \uparrow\Rightarrow M_{\text{shock}} \downarrow\Rightarrow \kappa_{\text{sys}} \downarrow$. A society that transmits narratives faster also *forgets* faster, and crisis memory is one of the things that raises κ_{sys} in Equation (7.5). The first row of Table 8.3 therefore reports a net effect whose sign is asserted rather than derived. It is very likely positive — the direct channel acts on the exponent through two factors while the indirect one acts through a single decayed term — but "very likely" is not a derivation, and Chapter 10 carries it as an open sign.

8.7 When the Loop Admits Three Equilibria

This section is what Assumption 8.4 buys, priced. It stands alone because ISSUES #30a confines it here: the rest of the book proceeds on the single-branch assumption, and this is the one place where the assumption is examined rather than used.

8.7.1 The result

Proposition 8.3 (Multiplicity)**.** *If $\eta\kappa_A g_1 > 4$, uniqueness can genuinely fail. In that case T has **at most three** fixed points, and generically either one or three. When there are three, $H_1^* < H_2^* < H_3^*$, the outer two satisfy $T' < 1$ and are locally stable while the middle satisfies $T' > 1$ and is unstable.*

Proof sketch. Substitute $u = \eta(\kappa_A G(H) - R_{\text{div}})$, so $H = (u - b)/a$ with $a = \eta\kappa_A g_1 > 0$ and $b = \eta(\kappa_A g_0 - R_{\text{div}})$. A fixed point solves $\varphi(u) \equiv b + a\Lambda(u) - u = 0$ with Λ the logistic. Then $\varphi'(u) = a\Lambda'(u) - 1$, and Λ' is strictly unimodal — increasing on $(-\infty, 0)$, decreasing on $(0, \infty)$, with a single maximum $\frac{1}{4}$ at $u = 0$. So φ' changes sign at most twice, φ has at most two turning points, and a continuous function with at most two turning points has at most three zeros. When $a \le 4$, $\varphi' < 0$ everywhere and there is exactly one. The stability classification follows from $T'(H^*) = a\Lambda'(u^*)$ and the fact that φ crosses zero downward at the outer roots and upward at the middle one. □

8.7.2 What the three equilibria are

Figure 8.3(b) shows the configuration at $\eta\kappa_A g_1 = 12$, and the three crossings have readings.

H_1^*, near zero, is a *diverse* society: enough residual resistance that the drivers cannot get purchase, cognitive loops stay long, impulse efficiency stays low, and the whole loop sits at its floor. It is locally stable — a small shock toward homogeneity decays.

H_3^*, near one, is a *phase-locked* society: the drivers have overcome resistance, deliberation has been compressed out of the cycle, and the loop sits at its ceiling. It is also locally stable. A small restoration of diversity decays back toward uniformity, which is the uncomfortable half of the result: in this branch, partial remedies do not work.

H_2^*, between them, is unstable, and it is the tipping threshold the original's §4.4 and §4.5 both describe without ever locating. A system displaced past H_2^* does not drift; it accelerates away, into whichever outer equilibrium it was pushed toward. That is the formal object behind the original's language about crossing a point of no return.

8.7.3 What this costs the rest of the book

Three things, and they are why the assumption was adopted.

Comparative statics stop being global. Under multiplicity, $\partial H^*/\partial\theta$ is defined only on a branch, and $1 - T'$ changes sign between branches, so Equation (8.10) inverts on the middle one. Every statement in Chapter 10 of the form "raising R_{div} lowers H_{soc}" would need a branch declaration to be meaningful.

History matters, and the model does not carry history. With three equilibria the selected branch depends on where the system was, so the layer becomes path-dependent. This book's data are cross-sectional in the relevant sense — there is no series for H_{soc} at all — so a path-dependent Layer 4 could not be identified even in principle.

The case studies would each need a branch. Chapter 13 through Chapter 16 would have to assert, for each episode, which side of H_2^* the world was on, with no measurement to support the assertion. That is a large amount of unfalsifiable apparatus bolted to episodes the model is supposed to explain.

8.7.4 What is at stake if $\eta \geq 4$

The honest statement is this. Assumption 8.4 is not innocuous and it is not verified. If $\eta\kappa_A g_1 > 4$ in the world, then the model's central object in this layer is not a point but a branch, the amplification factor of Equation (8.10) has the wrong sign on one of the three, and the tipping threshold H_2^* — which the original treats as the most important feature of the whole layer — is precisely the thing the maintained assumption assumes away.

Two considerations bear on how much this should worry a reader, and they point in opposite directions.

Against worry: the condition is $\eta\kappa_A g_1 < 4$, and Section 8.3.7 noted that κ_A is a product of four quantities in $[0, 1]$ and is therefore typically small, while Assumption 8.3 forces $g_1 \leq 1$ and friction pushes it well below. At $\kappa_A = 0.25$ and $g_1 = 0.5$ the condition permits η up to 32. The assumption is much weaker than the bare statement $\eta < 4$ suggests.

For worry: the original's entire narrative is that κ_A has been *rising* toward saturation and that friction has been *falling* toward zero. Both movements push the product toward the boundary. If the original's qualitative story is right, the system is moving in the direction

that invalidates the assumption used to analyse it — and the approach to the boundary is itself observable in principle, as the divergence of the amplification factor M, before the boundary is crossed.

Chapter 9 takes this up from the erosion side, where the same threshold reappears as the cliff in $\Psi_{\text{str}}(t)$, and Chapter 11 records η as the single parameter whose estimation would do most to settle what this layer can claim.

8.8 Measurement

8.8.1 What would have to be measured

Table 8.4: Measurement protocol for the psycho-technical layer. "Tier" follows ISSUES #23: 1 free and available, 2 free but requiring construction, 3 paywalled or non-existent.

Variable	Best available proxy	Tier	Known weakness
C_{src}	Herfindahl index over news reach, from the Reuters Institute Digital News Report or national regulator data	2	annual, national, and reach is not attention
T_{exp}	average daily media or platform minutes, from national regulators and industry trackers	2–3	mostly commercial; the free series are short and inconsistently defined
P_{AI}	share of published content machine-generated	3	no authoritative measurement exists; published estimates differ by an order of magnitude
A_{alg}	—	3	not observable outside the platforms; no external proxy is defensible
R_{div}	—	3	no accepted operationalisation
H_{soc}	cross-sectional dispersion of professional forecasts; average pairwise correlation of asset returns	2	both are *consequences* of homogeneity in one domain, not the latent state
I_{imp}	retail share of trading volume; share of options expiring within one day	2	measures execution latency in one market, not thought-to-action latency in society
V_{meme}	narrative diffusion half-life on public platforms	3	platform APIs have closed since 2023
B_{phase}	computed from the two states via Equation (8.1)	—	inherits both weaknesses

8.8.2 What is actually measured: nothing

> **Not supplied, and the gap is split, and the split is the point. Two of the nine are ours**: C_{src} from Reuters Institute reach data and H_{soc} from Survey of Professional Forecasters dispersion, which FRED carries. Both are free and both are the recommended next additions. **The other seven are the world's** — A_{alg}, P_{AI}, R_{div}, V_{meme}, I_{imp}, B_{phase} and T_{exp} have no provider and no obvious construction. *What it costs:* Layer 4 has **no scored predictions anywhere in this book** (Table 17.3), so the chapter's fixed-point theorem, its contraction bound and its located threshold are mathematics without measurement. It is also why Remark 18.5 can say that the strategy's fourth spear acts on a part of the model nothing else in the model can reach.

This is the bluntest gap in the book and it should be stated without softening. The psycho-technical layer contributes two factors to κ_{sys}, an argument to the erosion speed, and a modifier to the volatility of R_{US}. Not one of those contributions rests on an observed number. Every claim in Section 8.6 is a claim about a mechanism, argued from the original's premises and this chapter's algebra, and none is a claim about the world that has been checked against the world.

Three consequences bind what follows, and Chapter 11's standing constraint requires each to be stated there as well as here.

Layer 4 cannot be estimated, only assumed. η, c_{L0}, c_{L1}, v_0, b_A, b_P, a_H and a_I are eight OPEN parameters attached to nine unmeasured variables. No amount of dollar-index data identifies any of them, because they enter only as a product inside an exponent whose other factors are also unobserved. The right conclusion is Section 7.7.6's: κ_{sys} may be identifiable as a single latent state; its decomposition is not.

No case study may evidence this layer from market data. An episode in which the dollar moved fast and far is not evidence that H_{soc} was high, because the model's whole claim is that high H_{soc} is one of several things that make the dollar move fast and far. Reading the state back out of the outcome it is supposed to explain is circular, and it is the same error ISSUES #31 identified for D_{prec} and VIX. Chapter 12's method section inherits the rule: **Layer 4 may be described in a case study, never inferred from the episode's own price action.**

The dating claim is untested. The original's argument is not only that these mechanisms exist but that they have *intensified*, which is what makes the dollar more brittle now than in 2008. That is a claim about a time series, and there is no time series. It may well be right — the direction is plausible for every driver — but this book has not demonstrated it, and Chapter 17 must not score any episode as confirming it.

8.9 Conclusion

The psycho-technical layer is the newest stratum in the model and, after this chapter, the most formally developed and the least empirically supported. Both halves of that sentence matter.

Formally, the chapter has done four things the original does not. It closed the loop that the original asserts and its equations deny, at the one driver where the mechanism is specific (Equation (8.6)). It showed that the closed loop always has an equilibrium, with no assumption required (Proposition 8.1). It established the sharp condition $\eta\kappa_A g_1 < 4$ under which that equilibrium is unique and globally attracting, and the structure — at most three, outer two stable — when it is not (Proposition 8.2, Proposition 8.3). And it derived the amplification factor $1/(1-T')$, which gives "hyper-brittleness" a number: the loop multiplies every intervention's effect, and the multiplier diverges exactly as the system approaches the boundary beyond which its equilibrium stops being unique.

Empirically, it has done nothing, because there is nothing to do it with. Nine variables, eight free parameters, zero series.

That combination is the honest characterisation of Layer 4 as it stands: a well-specified mechanism with no measurement, which is a research programme rather than a result. Chapter 9 inherits V_{meme} and the tipping threshold; Chapter 10 inherits the comparative statics and the open sign on the countervailing channel; Chapter 11 inherits the whole of Section 8.8.2, which is the largest single entry in its account of what this model cannot estimate today.

Chapter 9

Layer 5 — The Structural Break and Narrative Layer: $\Psi_{\text{str}}(t)$

9.1 Erosion Versus Break

Layers 1 to 4 describe forces that operate *within* the dollar system: they change how much the dollar is demanded, how reliably it is trusted, how violently it transmits a shock, how synchronised the participants are. None of them can end the system. Layer 5 is the layer that can.

Definition 9.1 (Structural integrity)**.** $\Psi_{\text{str}}(t)(t) \in [\Psi_{\min}, 1]$ is the *structural integrity coefficient*: a multiplicative factor on the whole of E9 (Chapter 10) measuring how much of the dollar's systemic position survives at t. $\Psi_{\text{str}}(t) = 1$ is an intact system; $\Psi_{\text{str}}(t) \to \Psi_{\min}$ is the loss of default reserve and settlement status. Erosion is always $\Psi_{\text{str}}(t)\downarrow$.

One direction, stated once. The original is inconsistent about the direction of this variable: §4.2.4 and Case Study 5 both write of erosion *rising*, which read against §4.5's definition would mean $\Psi_{\text{str}}(t)$ rising as the system decays. ISSUES #2 resolves it in favour of the definition, and there is no companion erosion symbol: whenever this book means decay it writes $\Psi_{\text{str}}(t)\downarrow$. The two upward arrows in the original are corrected in place where they appear, in Chapter 6 and in the case study.

Why the variable is multiplicative and not additive. $\Psi_{\text{str}}(t)$ multiplies the entire right-hand side of E9, which is a strong structural statement worth making explicit. It says that structural erosion does not subtract a quantity of dollar demand; it *scales down everything the model computes*. A system at $\Psi_{\text{str}}(t) = 0.6$ is not one in which the dollar is worth forty per cent less. It is one in which every demand channel, every supply channel, every safe-haven response and every reflexive amplification delivers sixty per cent of what it would have delivered — the machinery still runs, and it moves less. That is the right representation of a currency losing its position rather than losing its price, and it is why $\Psi_{\text{str}}(t)$ sits outside the bracket rather than inside it.

Erosion is not the same event as a break. The original's §4.5 runs the two together, and separating them is the first job of this chapter. *Erosion* is the slow loss of share: a smaller fraction of reserves, of invoicing, of settlement. It is continuous, it is measurable, and Section 9.6 shows it has been happening for twenty-seven years. A *break* is a discontinuity in the rate of erosion: the moment at which the loss of share stops being gradual and becomes self-reinforcing. E8 is a model of the second, not the first, and the distinction matters because the data of Section 9.6 bear on erosion directly and on the break only by inference.

9.2 E8: The Cliff Function

src p.25

9.2.1 Statement

The original gives E8 twice. The earlier form, §4.5.4's "old form", is

$$\Psi_{\text{str}}(t)(t) \;\approx\; 1 - L_{\max} \cdot f(C_{\text{alt}}, T_{\text{frag}}, k_{\Psi,\text{base}}, \theta_0)\,, \tag{9.1}$$

which names the inputs without giving the function. The updated form specifies it:

$$\boxed{\Psi_{\text{str}}(t)(t) \;=\; 1 - L_{\max}\Big[1 + \exp\Big(-k_\Psi \cdot (f_{\text{hard}} - \theta_0)\Big)\Big]^{-1}, \qquad \Psi_{\min} = 1 - L_{\max}.} \tag{9.2}$$

Three renamings from ISSUES #11 and #14 are already applied and are stated here once so that a reader holding the original can follow. The threshold the original writes as t_0 is θ_0: it is a *level* of accumulated hard drivers, not a date, and writing it as a time was the source of considerable confusion in §4.5.3. The erosion speed the original writes as $k(\cdot)$ is $k_\Psi(H_{\text{soc}}, V_{\text{meme}}, I_{\text{imp}})$, written k_Ψ wherever the arguments are not at issue, freeing k to denote only the crisis index of Section 7.4. And the original's L is $L_{\max}$, the *maximum loss* of structural integrity, because the original's gloss — "L: the lower limit of Ψ" — is wrong as written: with $L_{\max}$ subtracted from 1, the lower limit is $1 - L_{\max}$, not $L_{\max}$ (ISSUES #14).

9.2.2 Assumptions

Assumption 9.1 (Hard drivers aggregate first)**.** C_{alt} and T_{frag} enter only through the scalar $f_{\text{hard}} = f(C_{\text{alt}}, T_{\text{frag}}) \in [0,1]$, which is increasing in both and is written f_{hard} from here on. The model never needs them separately in E8.

Assumption 9.2 (Speed and position separate)**.** The psycho-technical layer enters only through k_Ψ, and the hard drivers only through f_{hard} and θ_0. Narrative changes how *fast* the transition happens; it does not change *where* the threshold is.

Assumption 9.3 (Floor)**.** $L_{\max} \in (0,1]$, so $\Psi_{\min} = 1 - L_{\max} \in [0,1)$.

Assumption 9.2 is the load-bearing one and it is not innocuous. It is what makes E8 tractable — a one-dimensional logistic in f_{hard} with a speed parameter — and it is a substantive claim: that a world in which everyone believes the dollar is finished reaches the tipping point no *sooner* than a world in which nobody does, merely faster once it starts.

Figure 9.1(a) shows what that looks like: the curves for different k_Ψ all cross at $f_{\text{hard}} = \theta_0$. Whether that is the right separation is arguable — one could hold that belief itself lowers the threshold, since a currency is used because it is expected to be used — and Section 9.8 returns to it as the most consequential place where this chapter follows the original rather than departing from it.

9.2.3 Functional-form justification

The inverse logistic delivers the three phases the original's §4.5.1 describes in prose, and it is the minimal function that does.

The *stable plateau*: for $f_{\text{hard}} \ll \theta_0$ the logistic is near zero, so $\Psi_{\text{str}}(t) \approx 1$ and erosive forces accumulate with almost no visible effect. This is the property the original attributes to network effects and inertia, and it is worth noting that the model gets it from the functional form rather than from a mechanism — D_{net} appears nowhere in E8.

The *tipping point*: at $f_{\text{hard}} = \theta_0$ the logistic passes $\frac{1}{2}$, $\Psi_{\text{str}}(t)$ is exactly $1 - L_{\max}/2$, and the rate of decline is maximal at $L_{\max} k_\Psi / 4$.

The *freefall*: past the threshold the decline continues at a decreasing rate toward the floor $\Psi_{\min}$. Note that "freefall" is the original's word and it slightly overstates the mathematics: the fastest decline is *at* the threshold, not after it. The post-threshold phase is a rapid but decelerating approach to a floor, not an acceleration.

Three alternatives were considered. Linear decay, $\Psi_{\text{str}}(t) = 1 - ct$, has no threshold and cannot represent a break at all, which is the whole point of the layer. Exponential decay, $\Psi_{\text{str}}(t) = e^{-ct}$, is monotone and bounded but again has no plateau: it declines fastest at the start, which is the opposite of what the original describes. A power law would give persistence without a tipping point. The logistic is the standard choice for a bounded share undergoing a regime transition, and using it here keeps the book's single S-shape (Section 8.3.3).

9.2.4 Parameters

Table 9.1: Parameters of Equation (9.2). All are OPEN values in the registry, and Section 9.7 argues that none of them is currently identifiable.

Symbol	Meaning	Range	Provenance
$L_{\max}$	maximum loss of integrity	$(0, 1]$	original L, renamed (ISSUES #14)
θ_0	threshold level of f_{hard}	$[0, 1]$	original t_0, renamed (ISSUES #11)
k_Ψ	erosion speed	> 0	function of $H_{\text{soc}}, V_{\text{meme}}, I_{\text{imp}}$ (Chapter 8)

$L_{\max}$ is a claim about the world, not a nuisance parameter. ISSUES #14 requires this to be said out loud wherever E8 appears. $L_{\max} = 1$ gives $\Psi_{\min} = 0$ and permits total collapse: the dollar ends as a systemically significant currency. $L_{\max} < 1$ posits an *irreducible floor* of dollar demand that no accumulation of hard drivers can remove — and there is a real argument for one, since the United States would remain a large economy with deep

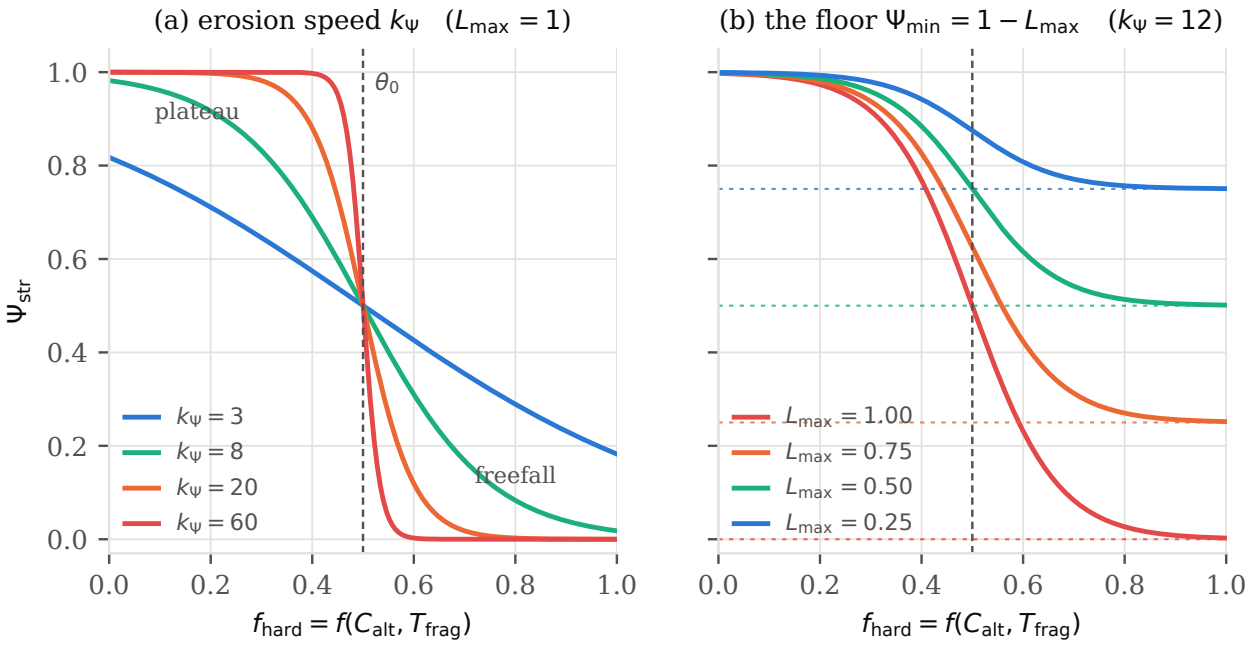


Figure 9.1: E8, the cliff function. Panel (a) varies the erosion speed k_Ψ at $L_{max} = 1$: the curves all cross at the threshold, so speed governs how long the transition takes and not where it begins — Assumption 9.2 made visible. Panel (b) varies L_{max} at $k_\Psi = 12$, with the floor $\Psi_{min} = 1 - L_{max}$ marked as a dotted line for each. Source: author's construction; script `f28_cliff.py`.

markets even having lost reserve status, exactly as Britain did. The original's §4.5 language ("potentially approaching zero") leans toward $L_{max} = 1$; the historical evidence of Section 9.5 leans hard against it. This book does not choose, but it insists that any calibration state which it assumes, because the two produce qualitatively different end-states from identical dynamics.

9.2.5 Properties

$\Psi_{str}(t) \in (\Psi_{min}, 1)$ strictly, so the multiplicative factor in E9 is bounded away from zero for any finite f_{hard} — the system decays toward its floor and never reaches it in finite hard-driver accumulation. $\Psi_{str}(t)$ is strictly decreasing in f_{hard} and in L_{max}; it is decreasing in k_Ψ above the threshold and *increasing* in k_Ψ below it, which is Figure 9.1(a) again and is easy to misread: a faster erosion speed makes the pre-threshold plateau *flatter*, so a world with a high k_Ψ looks *more* stable right up until it does not. Both arguments of the logistic are dimensionless.

9.2.6 Comparative statics

Write $\Lambda = \Lambda(k_\Psi(f_{\text{hard}} - \theta_0))$ for the logistic and $\Lambda' = \Lambda(1-\Lambda)$. Then

$$\begin{aligned} \frac{\partial \Psi_{\text{str}}(t)}{\partial f_{\text{hard}}} &= -L_{\max}\, k_\Psi\, \Lambda', & \frac{\partial \Psi_{\text{str}}(t)}{\partial \theta_0} &= +L_{\max}\, k_\Psi\, \Lambda', \\ \frac{\partial \Psi_{\text{str}}(t)}{\partial k_\Psi} &= -L_{\max}\, (f_{\text{hard}} - \theta_0)\, \Lambda'. \end{aligned} \tag{9.3}$$

The third changes sign at the threshold, as noted. The first is the one that matters for Section 9.6, and it has a property the original does not draw out:

$$\frac{\partial^2 \Psi_{\text{str}}(t)}{\partial f_{\text{hard}}^2} = -L_{\max}\, k_\Psi^2\, \Lambda'\,(1 - 2\Lambda) \;<\; 0 \quad \text{whenever } f_{\text{hard}} < \theta_0. \tag{9.4}$$

Below the threshold, E8's decline accelerates. As the hard drivers accumulate toward θ_0, each further increment of f_{hard} costs more integrity than the last. This is not an incidental feature; it is what a cliff *is*, and it is the property that makes E8 testable before the tipping point rather than only after it. Section 9.6 tests exactly this.

9.2.7 Failure modes

Three. First, f_{hard} is never defined: Assumption 9.1 asserts that C_{alt} and T_{frag} aggregate to a scalar without saying how, and no weighting is given anywhere in the original. Second, θ_0 and $L_{\max}$ are not separately identified from a single realised path, since a plateau is consistent with a distant threshold or with a small maximum loss. Third — and this is the serious one — E8 has never been through a completed transition, so the parameters have never been fitted to anything. Section 9.5 asks what the one historical transition available would imply if it were.

9.3 The Drivers of Erosion

The original's §4.5.2 makes a claim that is worth stating as a proposition of the model rather than as prose, because the whole architecture of the layer follows from it: a structural collapse of a monetary system requires *two* conditions, a viable technical alternative and widespread psychological acceptance of it, and neither alone suffices. Hard drivers build the pipes; soft drivers decide whether anyone uses them.

9.3.1 Hard drivers: T_{frag} and C_{alt}

T_{frag}, geopolitical fragmentation (Section 6.5), is the origin of the *motive*. BRICS+ expansion, and above all the immobilisation of reserves through S_{off} (Section 4.8), give states a reason to want monetary autonomy. Farrell and Newman [33] is the canonical statement of why: a network whose chokepoints can be used coercively creates an incentive for those subject to it to build around it. Mulder [59] supplies the historical record of how often that incentive has been acted on.

C_{alt}, alternative infrastructure, is the *capability*: CIPS, multi-CBDC arrangements, bilateral settlement, commodity contracts priced outside the dollar. The original correctly makes it the direct consequence of T_{frag} rather than an independent force.

Neither is measured, and the measurement problem is inherited. Section 6.5 recorded the search for a T_{frag} series and its outcome, which binds here without repetition: the SWIFT Global Currency Tracker publishes exactly the right quantity, free and monthly, in roughly one hundred and eighty separate PDFs with no CSV and no interface, so it cannot be scripted and no number from it enters this book until it is transcribed. CIPS began operating in October 2015 and cannot cover the sample. The substitute adopted under ISSUES #23 — BIS triennial turnover shares [11] and COFER — is adjacent to the concept rather than identical to it, and Section 9.6 uses it under that caveat. C_{alt} has no proxy at all: there is no index of how much non-dollar settlement capacity exists.

9.3.2 Soft drivers: the psycho-technical layer

H_{soc} determines whether a new narrative can be received; V_{meme} how fast it travels; I_{imp} whether reception converts to action. Together they are k_Ψ, and Equation (8.2) is what makes V_{meme} a computable quantity rather than a word, which is what allows k_Ψ to be written at all.

The original's argument here is its strongest: a monetary transition is not blocked by disbelief in the alternative but by the cost of coordinating a move to it, and a high-H_{soc}, high-I_{imp} world has had that coordination cost removed. Central banks that would once have moved over a decade, each watching the others, can now move in a quarter. That is a real mechanism and it is why the layer is not merely decorative.

It is also entirely unmeasured. Section 8.8.2 established that Layer 4 has nine variables and no series, so k_Ψ — the parameter that governs the speed of the whole transition — is a function of three quantities none of which has ever been observed. This compounds rather than merely repeats the problem of Section 9.3.1: E8's position parameter θ_0 depends on unmeasured hard drivers and its speed parameter k_Ψ depends on unmeasured soft ones.

9.4 The Tipping-Point Mechanism

src p.24

The original's §4.5.3 gives the mechanism as six stages. They are reproduced faithfully in Table 9.2, with the canonical variables attached and with one column the original does not have: what would have to be observed for each stage to be identified.

One stage is not in the equations at all. Before observability, there is a prior question: does the model *implement* each step? Stages 1, 4, 5 and 6 are carried by terms written down in this book. Stage 3 is not.

Stage 2 → stage 3 says a trigger event *spawns* the meme. In the model, a trigger event is a movement in the hard drivers or an exogenous shock, and the meme is V_{meme}, which Equation (8.2) makes a function of H_{soc}, A_{alg} and P_{AI} — none of which any hard driver touches. Section 10.1.5 checks the one candidate route, $X_{\text{pand}} \to \sigma_{\text{sys}} \to \cdots \to V_{\text{meme}}$, and

Table 9.2: The six-stage tipping mechanism of §4.5.3, with the canonical variables and the observable that would identify each stage. The last column is the chapter's own addition: it records both whether the model implements the step and what would identify it, and it is the basis of Section 9.7.

	Stage	Content	Implemented? Identified by?
1	Slow accumulation	$T_{\text{frag}} \uparrow$, $C_{\text{alt}} \uparrow$ while H_{soc} stays low and "the dollar reigns" persists. $\Psi_{\text{str}}(t)$ erodes imperceptibly.	*yes*; a rising f_{hard} with a flat $\Psi_{\text{str}}(t)$ — and f_{hard} has no measure
2	Trigger event	A shock exposing systemic vulnerability or U.S. overreach: the 2022 reserve immobilisation, a Taiwan contingency, an extreme R_{US} episode.	*yes*; dated and observable, the only stage that is
3	Narrative explosion	The event yields a transmissible meme — "the dollar is a weapon".	**NO — not implemented** (ISSUES #32). Nothing carries a trigger into V_{meme}
4	Viral adoption	With $H_{\text{soc}} \to 1$, $V_{\text{meme}} \uparrow$, $I_{\text{imp}} \uparrow$, the meme is internalised globally; A_{alg} and P_{AI} amplify it.	*yes*; but Layer 4 is unmeasured
5	Synchronised shift	$B_{\text{phase}} \to 1$: central banks, corporations and funds cut dollar exposure *simultaneously* and pivot to the C_{alt} built in stage 1.	synchronised reserve reallocation, in principle visible in COFER at quarterly frequency
6	Outcome	D_{net} collapses and $\Psi_{\text{str}}(t)$ falls off the cliff.	a break in the $\Psi_{\text{str}}(t)$ path

finds it does not exist: σ_{sys} enters only κ_{sys} and the safe-haven term, both outputs. More generally, Layer 4 has no inbound edge from any other layer (Remark 10.2).

So the arrow from stage 2 to stage 3 is prose, not mathematics. ISSUES #32 decided to record that rather than to repair it: a term $(1 + b_F f_{\text{hard}})$ on V_{meme} would implement it and would add an eighth unidentified coefficient to a layer that already has nine unmeasured variables (Section 8.8.2), which buys faithfulness to the narrative at the cost of honesty about the estimation. Table 9.2 therefore marks the step **not implemented**, and Chapter 11 and Chapter 17 both carry it.

The mechanism is coherent and its stages are not equally observable. Stage 2 is a dated event. Stage 5 would show up as a common component in reserve reallocation across holders, which COFER's aggregate cannot see but its constituent data in principle could. Stages 1, 3 and 4 rest entirely on unmeasured quantities. So of six stages, one is observable now, one is observable with data this book does not have, and three are not observable at all. Chapter 11 inherits this as one of its clearest entries.

A structural remark the original does not make. Stages 1 and 5 are separated by the whole of the psycho-technical layer, and that separation carries the model's actual prediction. Infrastructure is built *first*, quietly, over years; the narrative arrives later and suddenly; and the behavioural shift uses infrastructure that already exists. It follows that **the observable warning of a break is not the reallocation but the capacity** — C_{alt} rising while $\Psi_{\text{str}}(t)$ is still flat is stage 1, and by the time $\Psi_{\text{str}}(t)$ moves the capacity is already in place. That is a genuinely useful implication for Chapter 19, and it makes the absence of any C_{alt} measure the most costly data gap in this chapter rather than a technicality.

9.5 Calibration I: The One Completed Transition

E8 has three parameters and has never been fitted. There is exactly one completed reserve-currency transition in the modern record — sterling to the dollar — and this section asks what it implies. The answer is uncomfortable for the cliff.

The transition was not a cliff. The traditional account has the dollar displacing sterling after 1945. Eichengreen et al. [30] show, from reserve composition data, that the dollar had already overtaken sterling in official reserves by the mid-1920s — and then that sterling *recovered* in the 1930s and the two currencies coexisted as reserve assets for decades afterwards. The path is not a plateau followed by a freefall. It is a long overlap with a reversal in the middle.

Three features of that record bear directly on E8's parameters.

On $L_{\max}$. Sterling did not go to zero. It remains an allocated reserve currency a century after being overtaken, at a small but non-trivial share. If the sterling episode is the reference case, $L_{\max} < 1$ and there is a floor — the incumbent loses primacy and retains a position. The original's "potentially approaching zero" is not what the one available precedent shows.

On k_Ψ. The transition took decades, in a world with none of the psycho-technical machinery of Chapter 8. That is consistent with the original's argument — a low k_Ψ world

transitions slowly, and k_Ψ is now higher — and it is important to see that this is the one place where the historical evidence *supports* the model rather than embarrassing it. It is also unfalsifiable in the direction that matters, since no transition has occurred under a high k_Ψ.

On monotonicity. E8 is monotone in f_{hard}: $\Psi_{\text{str}}(t)$ can only fall as hard drivers accumulate. Sterling's 1930s recovery is not representable in E8 at all without f_{hard} itself falling, which would require the alternative infrastructure to have been dismantled. Eichengreen [28] attributes the reversal to policy and to the relative performance of the two economies — factors that live in Chapter 6's Φ_{sys}, not in Layer 5. So the model can represent the recovery only by routing it through a different layer.

Remark 9.1 (The sterling gap and the coercion gap are one defect)**.** This is not a peculiarity of the 1930s. The audit of the January 2026 Venezuela episode (Section 16.1.7) found the model equally unable to represent a fall in S_{off} and T_{frag} following an exercise of G_{pol} that removed a sanctions predicate. The two look like separate curiosities, ninety years apart, and they are the same thing: **the functional forms permit reversal and the sign graph contains no arrow that causes one.** Section 10.1.3 proved E8 fully reversible in f_{hard}; what is missing in both cases is a mechanism by which f_{hard} — or its components C_{alt} and T_{frag} — would ever decline. A reversal can therefore enter this model only by assuming a driver falls, with nothing explaining why it would. ISSUES #34 records the gap and states, in Section 6.5, the term that would close it and why it is not estimable.

The book should say this once, as one defect with two appearances, rather than twice as two observations.

What this section can and cannot establish. One observation is not a calibration. It cannot pin θ_0, it cannot pin k_Ψ, and a single episode with a reversal in it does not refute a functional form. What it does establish is that the only completed transition available looks less like Figure 9.1(a) than like a slow crossing with an overlap, and that a calibration assuming $L_{\max} = 1$ would be asserting something the precedent contradicts.

9.6 Calibration II: What the Measured Series Do

Two free, official, machine-readable series bear on Layer 5: the IMF's COFER dollar share of allocated official reserves, quarterly from 1999, and the BIS triennial survey's currency composition of foreign-exchange turnover, adopted as the T_{frag} substitute in Section 6.5. Figure 9.2 puts both against the cliff's prediction. Neither obliges, and the two fail it in different ways.

9.6.1 The reserve share: erosion is real and decelerating

The dollar's share of allocated reserves fell from a peak of 72.13 per cent in 1999-Q2 to a sample low of 56.42 per cent in 2025-Q4, standing at 57.13 per cent in 2026-Q1. Fitted linearly on time the trend is -0.379 percentage points a year with $t = -16.9$ and $R^2 = 0.73$. Erosion, in the sense of Section 9.1, is not in doubt: it is large, sustained and statistically

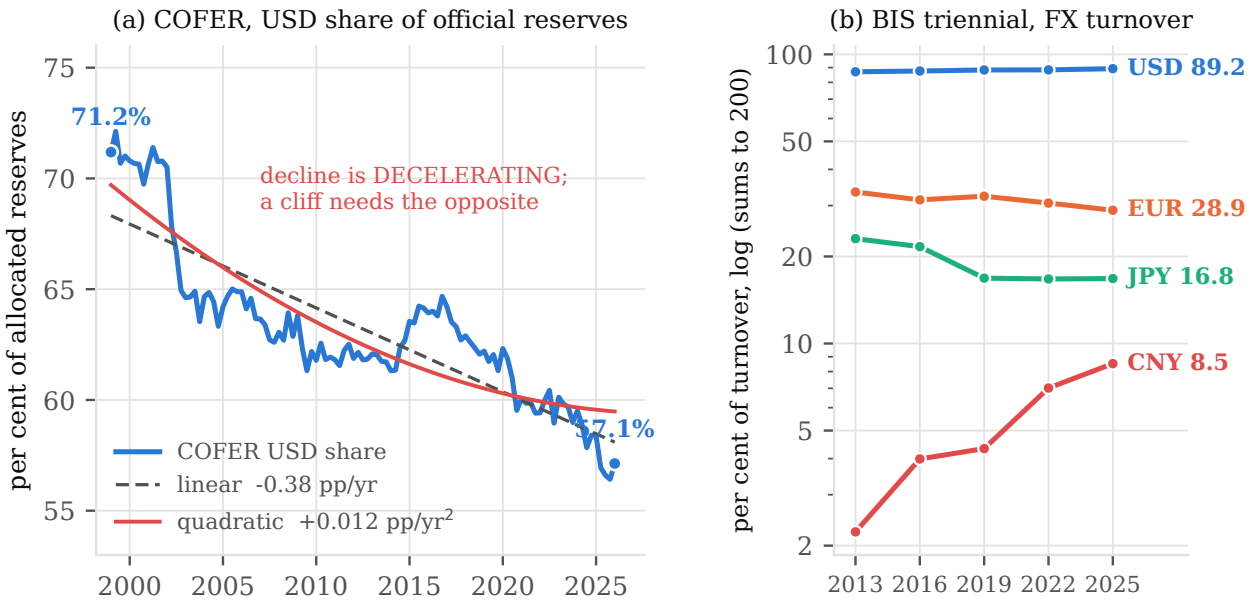


Figure 9.2: The two measured inputs to Layer 5. Panel (a): COFER dollar share of allocated official reserves, quarterly 1999-Q1 to 2026-Q1, 110 observations, with a linear and a quadratic fit in time. Panel (b): BIS triennial currency composition of OTC foreign-exchange turnover, log scale, shares summing to 200 by the BIS two-legged convention. Source: IMF COFER; BIS Triennial Central Bank Survey 2025, Table 4; built by `f29_shares.py` from the two extract scripts named in Section 6.5.

overwhelming, and over twenty-seven years it has removed fifteen percentage points of reserve share.

The question E8 poses is not whether the share falls but *how*. Equation (9.4) established that below the threshold a cliff function declines at an *accelerating* rate. Adding a quadratic term in time to the same regression gives a coefficient of +0.0115 percentage points per year squared with $t = +3.84$. The sign is positive and the term is significant: **on this aggregate the decline is decelerating**. The fitted path is convex from below — steep in the early 2000s, flattening since — which is the opposite of the shape Equation (9.4) requires. Section 9.7.1 asks whether that survives a change in what counts as a reserve, and it does not.

The decade averages say the same thing without a regression: 67.7 per cent over 1999–2005, 62.7 over 2006–2012, 62.8 over 2013–2019, 59.3 over 2020–2026. The middle two are indistinguishable. Fifteen years of the sample are a plateau bracketed by two declines, and the second decline is shallower than the first.

Where the erosion actually happened. The four largest four-quarter falls in the whole sample are consecutive, all in 2002–2003, the largest −5.9 percentage points to 2003-Q1. Not 2008. Not March 2020. Not 2022. The single fastest loss of dollar reserve share in twenty-seven years coincides with the euro's establishment and early appreciation, and it is over before the first of the crises this book studies.

That last point requires a caveat which cuts in the model's favour and must be given.

COFER's shares are computed at market exchange rates, so a currency's share falls mechanically when the others appreciate, without any reserve manager selling anything. The 2002–2003 episode is substantially a valuation effect. Arslanalp et al. [5] adjust for exactly this and find that genuine reallocation *has* occurred — gradually, and into non-traditional currencies rather than into the euro or the renminbi. Their title, "the stealth erosion of dollar dominance", is well chosen and the finding is real. It is also, on their own account, gradual: stealth erosion is the plateau, not the cliff.

9.6.2 The turnover share: the incumbent gained

Table 6.6 gives the numbers and Section 6.5 the caveats; the reading for this chapter is short. Between 2013 and 2025 the renminbi's share of foreign-exchange turnover nearly quadrupled, from 2.23 to 8.52 per cent. Over the same twelve years the dollar's share *rose*, from 87.04 to 89.21 per cent. What fell was the euro, by 4.5 points, and the yen, by 6.3.

On this measure the challenger is growing fast and the incumbent is not shrinking. The fragmentation is occurring among the alternatives to the dollar rather than between the dollar and its alternatives.

9.6.3 Reading the two together

The dollar has lost fifteen points of reserve share and gained two points of turnover share over overlapping windows. That is not a contradiction: reserves are a stock of stored value and turnover is a flow of transactions, and a currency can be less attractive to hold while remaining indispensable to trade through. It is, in fact, precisely Gopinath et al. [41]'s dominant-currency picture — invoicing and settlement are stickier than portfolio allocation, because they are a network property rather than a preference — and Prasad [63] argues the same from the other direction.

But it is awkward for Layer 5 specifically, because $\Psi_{\text{str}}(t)$ multiplies *everything* in E9 and therefore cannot distinguish the two. A single scalar for structural integrity must move one way. The evidence says the dollar's position as a store of value and its position as a medium of exchange have been moving in opposite directions for a decade, and E8 has no way to represent that.

9.7 What Survives

Proposition 9.1 (Erosion, supported)**.** *The dollar's share of allocated official reserves has declined substantially and persistently: 72.13 per cent at its 1999 peak to 57.13 per cent in 2026-Q1, a trend of* -0.38 *percentage points a year,* $t = -16.9$*. The original's claim that structural erosion is occurring is supported on the reserve measure.*

Proposition 9.1 settles that erosion is occurring. What follows is the separate question of its *shape*, which is what E8 actually constrains.

Proposition 9.2 (The cliff, not supported *on the reserve measure as COFER defines it*)**.** *E8 below its threshold requires an accelerating decline (Equation* (9.4)*). On the COFER series*

the measured decline is decelerating: a quadratic term of $+0.0115$ *pp/yr*^{2}, $t = +3.84$. *On the turnover measure the dollar's share rose. Neither series exhibits the pre-threshold behaviour the cliff function requires.*

This proposition is conditional on the aggregate, and the condition is not innocuous. *COFER measures the currency composition of* foreign exchange *reserves and by construction excludes monetary gold. Section 9.7.1 shows the sign reverses on a reserve aggregate that includes gold, and Proposition 9.4 states the result there.*

Proposition 9.3 (But the test is weak)**.** *Proposition 9.2 does not refute E8, because a plateau is what E8 predicts for* $f_{\text{hard}} \ll \theta_0$*, and no measure of* f_{hard} *exists. The observations are jointly consistent with (i) no cliff, (ii) a cliff whose threshold is distant, and (iii) a cliff whose threshold is near but whose hard drivers have not yet accumulated. The data cannot separate them, and E8 is therefore* ***not falsifiable on anything currently measured****.*

Proposition 9.3 is the honest summary of this chapter and it should not be softened into either direction. What can be said is sharper than it looks, and it takes the form of a conditional.

9.7.1 The same test on a reserve aggregate that includes gold

Proposition 9.2 was written before the audit of Section 17.6 asked whether COFER is the right base for a question about *de-dollarisation*. It is not obviously so. COFER reports the currency composition of foreign exchange reserves, and monetary gold is a reserve asset that is not a foreign currency — so a reserve manager who sells dollars for gold *leaves the COFER denominator entirely*, and the dollar's COFER share is unchanged by the transaction. If the exit from the dollar ran into gold, this chapter's instrument was built so as not to see it.

The exit did run into gold. This subsection is the resolution of ISSUES #39. Official net purchases were 1,082 tonnes in 2022, 1,037 in 2023 and about 1,045 in 2024 — three consecutive years above a thousand, the first since 1950 — against a 2010–2021 average near 473 tonnes, and the central-bank share of total gold demand rose from about 12 per cent in 2015–19 to nearly 25 in 2024. Sources in `data/SOURCES.md`, principally World Gold Council [75]. On the aggregate constructed below, gold rises from 8.3 per cent of world official reserves in 2015 to **23.6** per cent in 2025.

Construction. World official gold at market value is the difference between the World Bank's reserves-*including*-gold and reserves-*excluding*-gold series, summed over actual economies. Two caveats bind and are stated because the result rests on them. Coverage is 86 to 176 reporting economies depending on the year, so this is a total over reporters rather than the world; and the series is **annual**, so every figure here is a year-end and no within-year episode can be resolved on it.

The result. Re-running the identical quadratic-in-time fit reverses the sign, and reverses it significantly.

Table 9.3: The dollar's share of three reserve aggregates, year-end, per cent. Full series in `data/psi_gold.py`.

Year	COFER	+ gold at market	+ gold at 2015 $	Gold's share	
1999	71.01	59.51	39.35	16.2	
2007	62.60	56.84	53.56	9.2	
2015	64.15	58.85	58.85	8.3	
2021	59.40	52.12	54.49	12.3	
2022	58.95	51.12	53.66	13.3	the immobilisation year
2023	58.98	50.41	53.54	14.5	
2024	58.41	48.00	52.84	17.8	
2025	56.42	**43.10**	51.09	**23.6**	

Table 9.4: The quadratic term. E8 below its threshold requires a *negative* coefficient: an accelerating decline. Positive is deceleration, which E8 forbids there.

Dollar share of	pp/yr^2	t	
allocated FX, quarterly (Proposition 9.2)	+0.0115	+3.84	decelerating
allocated FX, annual (control)	+0.0076	+1.13	same sign, less power
FX + gold at market	**−0.0211**	**−2.59**	**accelerating**
FX + gold at constant 2015 $	**−0.0670**	**−15.34**	**accelerating**

The annual control reproduces Proposition 9.2's sign, so what follows is **a change of sign and not a loss of power**. A linear trend with a 2022 slope break says the same: the post-2022 shift is -2.90 ($t = -4.75$) at market value and -3.61 ($t = -3.28$) at constant prices, against -0.51 ($t = -0.82$) for the control.

Proposition 9.4 (On a gold-inclusive aggregate the decline accelerates)**.** *The dollar's share of world official reserves including monetary gold falls from* 59.51 *per cent in 1999 to* 43.10 *in 2025, and the quadratic term in time is* negative *and significant — -0.0211 pp/yr^2 with $t = -2.59$ at market value, -0.0670 with $t = -15.34$ at constant 2015 prices — against $+0.0076$ ($t = +1.13$) for the same annual fit on COFER alone. The pre-threshold behaviour Equation* (9.4) *requires is present on this aggregate and absent on the narrower one.*

Remark 9.2 (What this does and does not establish)**.** Three things, and the order matters.

It removes a contradiction; it does not supply support. The model is no longer contradicted by its own erosion evidence. Proposition 9.3 is untouched: f_{hard} still has no measure, so E8 remains unfalsifiable and an accelerating share decline is equally consistent with a cliff, with a trend, and with a rising gold price.

The 2022 freeze is not what produces the result, on one of the two valuations. At constant 2015 prices the accelerating decline is already present using data through 2021 alone (-0.0728, $t = -10.58$). So the mechanism most naturally attached to this finding — that the immobilisation of Russian reserves triggered an exit — holds for the market-value aggregate, which is $+0.0032$ through 2021 and -0.0211 over the full sample, and not for the constant-price one. Both are reported.

Part of the market-value result is the gold price itself. Gold rose sharply over 2024–25, and a rising price raises gold's share of a value-weighted aggregate without any central bank buying an ounce. The constant-price column exists to separate the two, and it is the noisier of the two because it requires an implied tonnage derived from a valuation residual. Neither column is clean. **They agree on sign, and that agreement is what the proposition rests on.**

What this costs the case studies. Three case-study rows scored $\Psi_{\text{str}}(t)$ on COFER. Re-scored on the correct aggregate, one is unchanged and two become untestable, because the gold-inclusive series is annual and cannot resolve their windows. The two moved in *opposite* directions — CS5's Ukraine row was a miss and CS9's was a hit — so the correction removed one of each. Section 17.2 carries the revised tally.

Remark 9.3 (The conditional that has teeth)**.** E8 predicts accelerating decline below the threshold. The measured decline decelerates. So *if* f_{hard} has been rising over 1999–2026 — which is what the original asserts throughout, and what the sanctions literature and the growth of C_{alt} both suggest — then the erosion is not following E8. To hold E8 and the data together one must argue that f_{hard} has been *flat or falling* across the same period in which the original says fragmentation accelerated. Chapter 11 should state which horn it takes.

That is a real constraint rather than a rhetorical one, and it is available now, without any new data, because it turns on the second derivative rather than on the level.

A fourth thing the propositions do not cover. All three concern whether $\Psi_{\text{str}}(t)$ *falls* as predicted. None concerns whether it can rise. Per Remark 9.1 the model contains no mechanism producing a recovery, so this chapter cannot be tested against a de-fragmentation episode at all — and one occurred inside the sample window, in January 2026. Chapter 17 scores that as a gap in coverage rather than as a failed prediction, because the model made no prediction to fail.

What would settle it. Three things, in order of value. A measure of C_{alt} — the stock of non-dollar settlement capacity — would give f_{hard} an observable and turn Remark 9.3 from a conditional into a test. The transcribed SWIFT series (Section 6.5) would give a settlement-side counterpart to COFER's stock measure, and the two moving apart is the phenomenon of Section 9.6.3. And the valuation-adjusted reserve reallocation of Arslanalp et al. [5], extended to the current sample, would remove the exchange-rate artefact that dominates the early part of the COFER path.

9.8 Conclusion

Layer 5 is the layer that can end the system, and this chapter has found that the model of how it would do so is well specified, historically unsupported and currently untestable.

The specification is sound. E8 delivers the three phases the original describes, the renamings of ISSUES #11 and #14 remove two real errors — a threshold written as a date, and a floor stated as its own complement — and Equation (9.4) extracts from the form a prediction the original never states, that decline should accelerate before the tipping point rather than only after it.

The history is against the strong version. The one completed reserve-currency transition shows an overlap of decades with a reversal in the middle, a floor rather than a collapse, and a path that E8's monotonicity cannot represent without routing the recovery through another layer.

The current data are against the cliff and cannot refute it. Erosion is real and large; its shape is decelerating where the cliff requires acceleration; and because f_{hard} is unmeasured, that discrepancy is a conditional rather than a rejection.

The assumption to attack first. Assumption 9.2 holds that narrative changes the *speed* of transition and not its *position*. It is what makes E8 a one-dimensional logistic, and it is the modelling choice this chapter is least confident in. A currency is used because it is expected to be used, so belief plausibly moves the threshold and not only the slope — and a model in which θ_0 itself falls as H_{soc} rises would have the tipping point arrive earlier as well as faster. That is a different model, with different policy implications, and Chapter 19 should know that the version in this book assumes the milder of the two.

What the rest of the book inherits. Chapter 10 inherits $\Psi_{\text{str}}(t)$ as a bounded multiplicative factor with an explicit floor, and the requirement to state that it is a ratchet while the phenomenon is not. Chapter 11 inherits three unmeasured parameters attached to two unmeasured driver sets, the conditional of Remark 9.3, and the observation that of the six

stages of Table 9.2 exactly one is currently observable. Chapter 17 inherits two scoreboard entries: *erosion supported*, *cliff not supported and not falsifiable on current data* — and, from Section 9.6.3, the finding that the dollar's reserve role and its transaction role have moved in opposite directions for a decade, which no single $\Psi_{\text{str}}(t)$ can express.

Chapter 10

System Synthesis

10.1 Assembling the Model

10.1.1 What the previous eight chapters produced

Chapters 3 to 9 built the model one layer at a time, and each closed with a list of what it handed on. This section collects those hand-offs before E9 is written, because several of them constrain the assembly rather than merely feeding it.

Table 10.1: What each layer contributes to E9, and the constraint it carries with it.

Layer	Contributes	Constraint carried into the assembly
0	D_i, S_j	Two channels are emergent and enter as $1 + X/\bar{X}$ (I29); S_{off} attenuates S_{ext} multiplicatively (I28); D_{prec} is excluded from the product and carries its own exponent (I7, I8).
1	σ_{sys}	Bounded in $[0, 1]$ by construction (I9). Its third component's weight w_X is unestimated, and the amplification result of Chapter 5 is conditional on it (Proposition 5.2).
2	Φ_{sys}, $s_{\text{US,sys}}$	G_{pol} reaches κ_{sys} only through these, never directly (I3).
3	κ_{sys}	Seven drivers, none separately identified; the psycho-technical factors are bounded in $[1, (1 + a_H)(1 + a_I)]$ (I3a).
4	H_{soc}, I_{imp}	Evaluated at the unique fixed point under Assumption 8.4; the amplification factor $1/(1 - T')$ multiplies every driver's effect.
5	$\Psi_{\text{str}}(t)$	Multiplies everything, bounded in $[\Psi_{\min}, 1]$, and is a static monotone function of f_{hard} with no state of its own.

10.1.2 Two thresholds, and the question they raise

Before E9 can be written down, one structural question has to be settled, and it is not settled anywhere in the original.

Chapter 8 located a threshold. When $\eta\kappa_A g_1 > 4$ the psycho-technical loop acquires three equilibria, and the middle one, H_2^*, is unstable: a system displaced past it accelerates away rather than returning (Proposition 8.3).

Chapter 9 located a threshold. E8 turns over at $f_{\text{hard}} = \theta_0$, where the logistic passes one half and the decline of $\Psi_{\text{str}}(t)$ is fastest.

The original describes these in separate sections, as separate mechanisms belonging to separate layers, and never asks whether they are the same thing. They must be either distinct phenomena, or one phenomenon seen in two coordinate systems, or two phenomena that are coupled. Which they are changes what E9 means, so it is answered here rather than assumed.

10.1.3 The answer: distinct objects, one-way coupled

Proposition 10.1 (The two thresholds)**.** *The Chapter 8 and Chapter 9 thresholds are* ***distinct mathematical objects with disjoint driver sets****,* ***one-way coupled*** *through k_Ψ, and* ***only the Chapter 8 threshold is a tipping point*** *in the dynamical sense. Specifically:*

1. *Their inputs do not intersect. The Chapter 8 fold depends on η, A_{alg}, C_{src}, P_{AI}, T_{exp}, R_{div} and the five impulse drivers; the Chapter 9 threshold depends on C_{alt}, T_{frag} and θ_0. No variable appears in both.*
2. *H_2^* is an unstable equilibrium separating two basins of attraction, so crossing it is self-reinforcing and* path-dependent*. θ_0 is the inflection of a static monotone map, so crossing it is neither.*
3. *Layer 4 enters Layer 5 only through k_Ψ, and by Assumption 9.2 that changes the* steepness *of E8 and not the* position *of θ_0. The two curves in Figure 10.1(b) cross exactly at θ_0.*
4. *There is no path from Layer 5 back to Layer 4 in the equations as written.*

On (1), disjointness. This is worth stating because it has a consequence the original's prose implies but its structure forbids: since no driver is shared, **there is no reason for the two thresholds to be crossed at the same time, and no mechanism in the model by which crossing one causes the other to be crossed**. The narrative and the infrastructure accumulate independently. Whether that is right is I32, below.

On (2), and it is the important one. Chapter 9 used the words "tipping point" and "freefall" for θ_0, following the original. That vocabulary belongs to bistable systems, and E8 is not one. Write it out: $\Psi_{\text{str}}(t) = 1 - L_{\max}\Lambda(k_\Psi(f_{\text{hard}} - \theta_0))$ is a function of contemporaneous f_{hard} and k_Ψ alone. It has no lag, no integral and no state variable, so if f_{hard} falls the integrity comes back. E8 is **fully reversible**, and θ_0 is the inflection of a steep monotone curve rather than a separatrix.

The Chapter 8 loop is the opposite. Its equilibrium depends on where the system has been, and Figure 10.1(a) shows the hysteresis explicitly: at $\eta = 20$ the two saddle-nodes sit at $\kappa_A = 0.64$ and $\kappa_A = 1.46$. Raising κ_A past the upper fold throws H_{soc}^* from 0.04 to

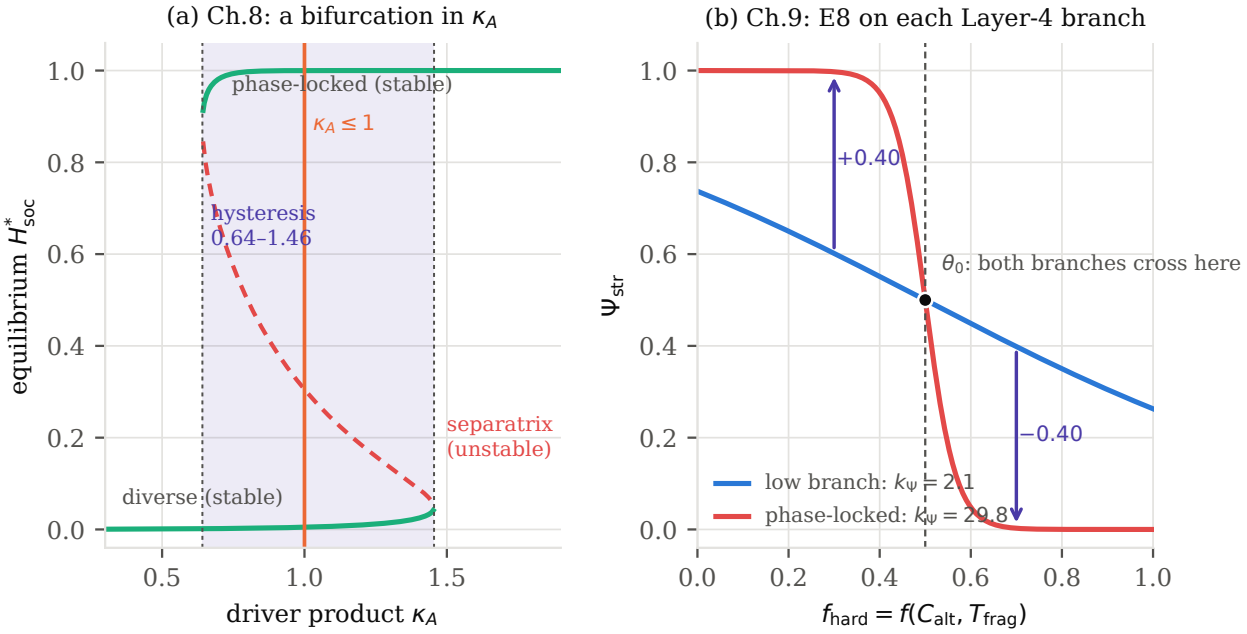


Figure 10.1: The two thresholds. Panel (a): the Chapter 8 object, a bifurcation diagram of the psycho-technical fixed point in the driver product κ_A, at $\eta = 20$, $g_0 = 0.15$, $g_1 = 0.75$, $R_{\text{div}} = 0.42$. Stable branches solid, the unstable separatrix dashed, and the hysteresis window between the two saddle-nodes shaded. Panel (b): the Chapter 9 object, E8 in f_{hard}, drawn once for each Layer-4 branch with k_Ψ mapped affinely from H^*_{soc}. The branches cross at θ_0: a jump in panel (a) rotates the curve in panel (b) about that point, raising $\Psi_{\text{str}}(t)$ below the threshold and lowering it above. Source: author's construction; script `f30_thresholds.py`.

1.00; lowering it again does *not* bring the system back until κ_A has fallen all the way to 0.64. Undoing the jump requires more than halving the driver product that caused it.

So the conclusion is one this book should state plainly, because it inverts the original's own emphasis:

Remark 10.1 (Where the irreversibility actually lives)**.** The original introduces Layer 5 as the layer modelling forces capable of terminating the system *irreversibly*, and does not describe Layer 4 in those terms. Proposition 10.1 shows the reverse. **E8 contains no irreversibility at all**; the only hysteresis in the entire model is the multiplicity of Proposition 8.3, in Layer 4. Whatever irreversibility a real structural break possesses must therefore enter this model either through Layer 4 or through a ratchet in C_{alt} that E8 does not represent.

This also corrects a phrase in Section 9.5, which described $\Psi_{\text{str}}(t)$ as "a ratchet". It is not one. E8 is monotone *in* f_{hard}, which is a different property: it cannot represent a recovery unless f_{hard} itself falls, but if f_{hard} falls, it recovers completely. The sterling episode of Section 9.5 needs the first half of that sentence, not the second.

10.1.4 The coupling, and what it does

Layer 4 reaches Layer 5 through exactly one channel, k_Ψ, and the effect is sharper than "it speeds things up".

Section 9.2.5 established that $\partial\Psi_{\text{str}}(t)/\partial k_\Psi$ carries the sign of $-(f_{\text{hard}} - \theta_0)$. So the coupling is *sign-reversing at the threshold*:

$$\frac{\partial \Psi_{\text{str}}(t)}{\partial k_\Psi} = -L_{\max}\,(f_{\text{hard}} - \theta_0)\,\Lambda' \quad \begin{cases} > 0 & f_{\text{hard}} < \theta_0, \\ = 0 & f_{\text{hard}} = \theta_0, \\ < 0 & f_{\text{hard}} > \theta_0. \end{cases} \tag{10.1}$$

Combine that with the discontinuity of Proposition 8.3 and a composite result follows that neither chapter contains.

Proposition 10.2 (Layer 4 makes Layer 5 discontinuous)**.** $\Psi_{\text{str}}(t)$ *is a smooth function of every one of its own arguments, yet it is* discontinuous *in the psycho-technical drivers. A saddle-node crossing in Figure 10.1(a) moves* H^*_{soc} *discontinuously, hence* k_Ψ *discontinuously, hence* $\Psi_{\text{str}}(t)$ *discontinuously — upward if* $f_{\text{hard}} < \theta_0$ *and downward if* $f_{\text{hard}} > \theta_0$*, by Equation* (10.1)*. On the illustrative calibration of Figure 10.1 the jump is* $+0.40$ *at* $f_{\text{hard}} = 0.30$ *and* -0.40 *at* $f_{\text{hard}} = 0.70$*, and exactly zero at* θ_0*.*

The uncomfortable half of that. Read the positive case again. If the hard drivers have *not* yet accumulated, a psycho-technical bifurcation — the world becoming maximally homogeneous and impulsive, which is the state the original treats as most dangerous — makes measured structural integrity **rise**. It does so because a high k_Ψ flattens the pre-threshold plateau, and a flatter plateau reads as a healthier system. The reading is not wrong; it is what the model says. A system that has locked into narrative synchrony but has not yet built the alternative rails looks *more* stable than one that has done neither, right up until f_{hard} crosses θ_0, at which point it loses everything the flatter plateau was holding.

Order of crossing therefore matters, and the model says which order is worse. Both sequences end in the same place, and they differ entirely in what is observable on the way.

Narrative first. Layer 4 bifurcates while $f_{\text{hard}} < \theta_0$. $\Psi_{\text{str}}(t)$ rises to near 1 and the plateau is flat, so no measurement of structural integrity gives any warning. When f_{hard} later crosses θ_0 it does so at $k_\Psi \approx 30$ on the calibration of Figure 10.1, and the transition occupies about 0.15 of the f_{hard} range.

Infrastructure first. f_{hard} rises while Layer 4 is still on the low branch. $\Psi_{\text{str}}(t)$ declines gradually and visibly — at $k_\Psi \approx 2$ the transition is wider than the whole admissible range of f_{hard}, so there is no cliff to see, only erosion. The narrative bifurcation, when it comes, then costs the remaining integrity in one step.

The dangerous sequence is the first, and it is the reverse of the original's own six-stage mechanism (Table 9.2), whose stage 1 is slow infrastructure accumulation and whose stages 3 to 5 are the narrative. **Under its own equations, the original describes the more survivable ordering.** Chapter 19 should take the other one as the planning case.

A refinement on how the jump is reached. The upper saddle-node in Figure 10.1(a) sits at $\kappa_A = 1.46$, and $\kappa_A \leq 1$ by construction (Assumption 8.2). At $\eta = 20$, therefore, *drift in the drivers alone cannot carry the system onto the phase-locked branch*: it can only enter the multiplicity window, where it stays on the low branch until something displaces it past H_2^*. Solving the tangency condition for the fold shows the upper saddle-node leaves the admissible range at $\eta \approx 11$; below that, drift suffices.

That is a satisfying place for the two chapters to meet. For η above about eleven, a jump requires an exogenous shock to H_{soc} large enough to cross the separatrix — which is precisely what stage 2 of the original's mechanism, the trigger event, is. The original's narrative and the mathematics of Chapter 8 agree without having been made to.

10.1.5 What is missing: the return path

Proposition 10.1(4) records that nothing carries Layer 5 back to Layer 4. The original's prose says otherwise. Its stage 2 to stage 3 is explicit that a trigger event — a sanctions episode, which is S_{off} raising T_{frag} and hence f_{hard} — *spawns* the meme. That is a path from the hard drivers to V_{meme}, and there is no such term in Equation (8.2).

This is structurally the same defect that ISSUES #13 found *inside* Layer 4, where the text claimed mutual reinforcement and the equations were one-way. Here it is between two layers, and the consequence is concrete: as written, the two thresholds can modulate each other but cannot *trigger* each other, so the original's central story — weaponised interdependence produces a narrative which produces a collapse — has no formal counterpart in the assembled model.

Checked first: the coupling is not already there. Before proposing a term, the obvious question is whether the model can already carry the step through variables it has. The trigger event of stage 2 is an exogenous shock, and exogenous shocks are X_{pand}, which

does enter the model — it is one of the three components of σ_{sys} (Section 5.3). So the candidate route is $X_{\text{pand}} \to \sigma_{\text{sys}} \to \cdots \to V_{\text{meme}}$.

It does not exist. σ_{sys} appears in exactly two places: as a factor $(1 + a_\sigma \sigma_{\text{sys}})$ inside κ_{sys}, and in the safe-haven exponential of E9. Both are outputs. Nothing in E7, E7b or Equation (8.2) takes σ_{sys} as an argument, so the chain terminates one step after it starts.

The stronger statement is worth recording, because it is a fact about the whole architecture rather than about this one step:

Remark 10.2 (Layer 4 is a source)**.** No variable outside Layer 4 reaches H_{soc}, I_{imp}, C_{loop}, V_{meme} or B_{phase}. The psycho-technical layer has **no inbound edge from any other layer**: it acts on Layers 3 and 5 and is acted on by nothing. Every other layer in the model both receives and transmits.

Decided: the gap is recorded, not filled. ISSUES #32 is resolved in favour of stating the omission. The step that is missing is precisely **stage 2 → stage 3** of Table 9.2: the trigger event spawning the transmissible narrative. A term that would implement it is easy enough to write — $(1 + b_F f_{\text{hard}})$ appended to Equation (8.2), keeping the additive-one convention of ISSUES #3a — and it is not added, for the reason this book has applied to itself twice already. κ_{sys} already carries seven coefficients that nothing in the available data separates (Section 7.7.6), attached to nine Layer-4 variables with no series at all (Section 8.8.2). An eighth inestimable parameter, introduced to close a narrative gap rather than because anything measures it, would make the model *less* honest while appearing more faithful. The omission is the more informative object, so it is carried as one: in Section 9.4, here, in Chapter 11's account of what cannot be estimated, and in Chapter 17's scoreboard.

Chapter 17 inherits all of this as three scoreboard entries rather than one: the two thresholds are distinct and must not be scored as one mechanism; the irreversibility the original attributes to Layer 5 is in Layer 4; and the ordering result of Section 10.1.4 means an episode showing rising $\Psi_{\text{str}}(t)$ is not evidence against the model, because the model predicts exactly that under a narrative-first sequence.

10.2 E9: The Integrated Equation

src p.27

10.2.1 Statement

$$DXY(t) = \Psi_{\text{str}}(t)(t) \cdot DXY_0 \cdot \left[\frac{\prod_{i \in \mathcal{I}} D_i^{\alpha_i(\sigma_{\text{sys}})} \cdot D_{\text{prec}}^{\alpha_{\text{prec}}(\sigma) s_{\text{US,sys}}}}{\prod_{j \in \mathcal{J}} S_j^{\beta_j(\sigma_{\text{sys}})}} \right]^{\kappa_{\text{sys}}} \exp\Big[(\lambda_1 \sigma_{\text{sys}} + \lambda_2 \sigma_{\text{sys}}^2) \Phi_{\text{sys}} \Big] + \epsilon. \tag{10.2}$$

Two corrections from Chapter 2 are already applied and are restated because they are easy to lose. The index set $\mathcal{I}$ **excludes** prec: in the original, D_{prec} appears both inside the product and again as a separate factor, which would double-count it (ISSUES #7). And the separate factor keeps its $\alpha_{\text{prec}}(\sigma) s_{\text{US,sys}}$ exponent rather than a constant one (ISSUES #8), which is what makes precautionary demand respond more strongly when the dollar's relative safety is high.

10.2.2 Reading the four blocks

E9 is a product of four things, and the model's whole architecture is the claim that they are *multiplicatively separable* in this particular way.

$\Psi_{\text{str}}(t)$*, outside everything.* Structural integrity scales the entire result (Section 9.1). It is the only factor that can take the dollar's level to a different regime rather than to a different value.

DXY_0*, the normalisation.* By I12, every D_i and S_j is indexed to 1 at t_{base} = 2006-01 and $\Psi_{\text{str}}(t)(t_{\text{base}}) = 1$, so $DXY_0 = DXY(t_{\text{base}})$ and the bracket equals one at the base period, term by term. The emergent channels satisfy this because of the $1 + X/\bar{X}$ form (I29): a channel absent at t_{base} contributes exactly 1, not 0.

The bracket raised to κ_{sys}. This is the mechanical demand-supply ratio, and κ_{sys} is an exponent on it rather than a coefficient in it. That is the model's sharpest structural claim: reflexivity does not add to the dollar's response, it *scales the elasticity of the response to everything else.* A doubling of κ_{sys} does not double the dollar's move; it squares the demand-supply ratio's effect on it.

The safe-haven exponential. $\exp[(\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma_{\text{sys}}^2)\Phi_{\text{sys}}]$ is the flight-to-quality channel, quadratic in stress and scaled by the reliability of the haven. It is additive in logs and therefore independent of the mechanical block, which is the assumption Section 10.7 examines.

10.2.3 Bounds

Every factor is now bounded, which was not true of the original and is the cumulative payoff of ISSUES #9, #14, #26 and #3a. With $\sigma_{\text{sys}} \in [0, 1]$, $\Phi_{\text{sys}} \in [0, \Phi_{\max}]$, $\Psi_{\text{str}}(t) \in [\Psi_{\min}, 1]$ and $\kappa_{\text{sys}} \in (0, \kappa_{\text{base}}(1 + a_V)(1 + a_M)(1 + a_A)(1 + a_\sigma)(1 + a_H)(1 + a_I)]$:

$$\Psi_{\min}\, DXY_0 \left[\frac{D^*}{S^*}\right]^{\kappa_{\text{sys}}} \;\le\; DXY - \epsilon \;\le\; DXY_0 \left[\frac{D^*}{S^*}\right]^{\kappa_{\text{sys}}} e^{(\lambda_1+\lambda_2)\Phi_{\max}}. \tag{10.3}$$

The safe-haven term is bounded by $e^{(\lambda_1+\lambda_2)\Phi_{\max}}$ rather than exploding, which is the whole content of ISSUES #9: with σ_{sys} on a raw VIX scale, $e^{\lambda_2\sigma_{\text{sys}}^2}$ at VIX 80 is astronomical for any λ_2 large enough to matter at VIX 15.

10.3 The Log-Linear Form

$$\begin{aligned} \ln DXY = \; & \ln \Psi_{\text{str}}(t) + c_0 \\ & + \kappa_{\text{sys}}\Big[\sum_{i\in\mathcal{I}} \alpha_i(\sigma_{\text{sys}}) \ln D_i + \alpha_{\text{prec}}(\sigma) s_{\text{US,sys}} \ln D_{\text{prec}} - \sum_{j\in\mathcal{J}} \beta_j(\sigma_{\text{sys}}) \ln S_j\Big] \\ & + (\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma_{\text{sys}}^2)\Phi_{\text{sys}} + \epsilon. \end{aligned} \tag{10.4}$$

One departure from the original. The original drops the $\alpha_{\text{prec}}(\sigma)s_{\text{US,sys}}$ structure when it moves from E9 to E10, so that D_{prec} re-enters the sum with an ordinary constant elasticity. ISSUES #8 decided to keep it, and Equation (10.4) does. The reason is that dropping it silently changes the model: the claim that precautionary demand bites harder when the dollar is the relatively safer asset is exactly what $s_{\text{US,sys}}$ carries, and it is not a nuisance term.

The cost is that E10 is then non-linear in $s_{\mathrm{US,sys}}$ and cannot be estimated by ordinary least squares without either conditioning on $s_{\mathrm{US,sys}}$ or instrumenting it, which Chapter 11 must handle rather than avoid.

The differenced form the case studies use.

$$\Delta \ln DXY \approx \Delta \ln \Psi_{\mathrm{str}}(t) + \kappa_{\mathrm{sys}} \Big[\sum_i \alpha_i(\sigma_{\mathrm{sys}}) \Delta \ln D_i - \sum_j \beta_j(\sigma_{\mathrm{sys}}) \Delta \ln S_j \Big] + (\lambda_1 \Delta\sigma_{\mathrm{sys}} + \lambda_2 \Delta\sigma^2_{\mathrm{sys}}) \Phi_{\mathrm{sys}}. \tag{10.5}$$

Three approximations are buried in that "≈" and Chapter 12 must not use it without them. It holds κ_{sys}, Φ_{sys} and the elasticities fixed across the window, which is defensible for a week and not for a year. It drops the cross-terms in $\Delta\kappa_{\mathrm{sys}} \cdot \ln(D^*/S^*)$, which are second-order only if κ_{sys} moves little — and Chapter 7 argues it moves most exactly in the crises the case studies examine. And $\Delta\sigma^2_{\mathrm{sys}}$ means $(\Delta\sigma_{\mathrm{sys}})$ applied to the squared term, not the square of the change; writing it the other way changes the sign of the convex contribution when stress falls.

10.4 The Modulator Functions, Consolidated

src p.28

The original's §5.2 restates the six modulator functions one after another, and this section does the same for this book's canonical versions. It is reference material: everything here was derived in Chapter 5 to Chapter 9, and nothing new is claimed. What the collection makes visible is the pattern.

Three patterns, and one of them is a warning. *The book has one S-shape.* The logistic appears four times — in σ_{sys}, in $s_{\mathrm{US,sys}}$, in H_{soc} and in $\Psi_{\mathrm{str}}(t)$ — and each time for the same reason: a quantity that must be bounded, monotone in a drive, and least responsive at its extremes. That is an economy worth naming, and it means a reader who understands one of the four understands the shape of all of them.

Dampening quantities never subtract directly. R_{fric} enters as $1/(1+R_{\mathrm{fric}})$, S_{off} as $e^{-\delta_{\mathrm{off}} S_{\mathrm{off}}}$, H_{soc} and I_{imp} as $(1 + a\,x)$ factors after I3a. The convention adopted across ISSUES #3a and #28 is that **no single input may switch a channel off entirely**, and the one place the book departs from it is $(1 - \kappa_R R_{\mathrm{US}})$ in E6, which is why that term needs an explicit positivity constraint while none of the others does.

Every bound is a decision, not a discovery. Four of the six functions were unbounded as the original wrote them, and each is bounded here because a specific ISSUES item decided it should be. That is worth stating plainly rather than presenting the bounded versions as though they had always been there: a reader comparing this book with the original will find the functions changed, and Table 10.2's last column says by whose authority.

What is not in the table. V_{meme} and B_{phase} (Section 8.2.4) are consequence variables with equations but no place in E9; M_{shock} and its memory formulas E11–E14 (Section 7.4) feed κ_{sys} through a_M; and E15's marginal effects of ΔG_{pol} are comparative statics of E4 and E5 rather than functions in their own right. All four are used by the case studies and none appears in Equation (10.2) directly.

Table 10.2: The six modulator functions in canonical form, with their ranges and the decision that fixed each. Every one is bounded; in the original, four were not.

Function	Symbol	Canonical form	Range
E3	σ_{sys}	$\left[1+e^{-z}\right]^{-1}$, z the standardised composite $(w_\sigma \sigma_{\text{VIX/MOVE}} + w_X X_{\text{pand}})(1 + w_F F_{\text{liq}})$ on percentile ranks	$[0,1]$ by construction (I9)
E4	Φ_{sys}	$\phi_0 + \phi_G G_{\text{pol}} - \phi_R R_{\text{US}} + \phi_M M_{\text{trust}} - \phi_T T_{\text{frag}}$, all $\phi > 0$; sign of ϕ_M flipped to + (I1), inflation entering through M_{trust} (I25)	$[0, \Phi_{\max}]$
E5	$s_{\text{US,sys}}$	logistic in $\gamma_0 + \gamma_G G_{\text{pol}} - \gamma_E E_{\text{inf}} - \gamma_R R_{\text{US}}$	$(0,1)$
E6	κ_{sys}	$\kappa_{\text{base}}(1+a_V V_{\text{USD}})(1+a_M M_{\text{shock}})(1+a_A A_{\text{meta}})(1+a_\sigma \sigma_{\text{sys}})(1-\kappa_R R_{\text{US}})(1 + a_H H_{\text{soc}})(1 + a_I I_{\text{imp}})$	$(0, \kappa_{\max}]$, needs $\kappa_R < 1$ (I3, I3a)
E7/E7b	$H_{\text{soc}}, I_{\text{imp}}$	the fixed point of $T = F \circ G$, unique under $\eta \kappa_A g_1 < 4$	$[0,1]^2$ (I13, I30a)
E8	$\Psi_{\text{str}}(t)$	$1 - L_{\max}\left[1 + e^{-k_\Psi (f_{\text{hard}} - \theta_0)}\right]^{-1}$	$[\Psi_{\min}, 1]$, $\Psi_{\min} = 1 - L_{\max}$ (I14)

10.5 Comparative Statics

Differentiating Equation (10.4) gives the elasticities every case study uses:

$$\frac{\partial \ln DXY}{\partial \ln D_i} = \kappa_{\text{sys}}\,\alpha_i(\sigma_{\text{sys}}), \qquad \frac{\partial \ln DXY}{\partial \ln S_j} = -\kappa_{\text{sys}}\,\beta_j(\sigma_{\text{sys}}), \qquad \frac{\partial \ln DXY}{\partial \ln \Psi_{\text{str}}(t)} = 1, \tag{10.6}$$

and, for the state variable itself,

$$\frac{\partial \ln DXY}{\partial \sigma_{\text{sys}}} = (\lambda_1 + 2\lambda_2\sigma_{\text{sys}})\Phi_{\text{sys}} \;+\; \frac{\partial \kappa_{\text{sys}}}{\partial \sigma_{\text{sys}}} \ln \frac{D^*}{S^*} \;+\; \kappa_{\text{sys}} \sum_i \frac{\partial \alpha_i}{\partial \sigma_{\text{sys}}} \ln D_i \;-\; \cdots \tag{10.7}$$

The first two are the workhorses and they carry the model's central multiplicative claim: **every mechanical elasticity is scaled by** κ_{sys}. A one per cent rise in transactional demand moves the dollar by $\kappa_{\text{sys}}\alpha_{\text{trans}}$ per cent, so the same shock has a different price impact in a reflexive market than in a quiet one, without any elasticity having changed.

The third is worth writing out because it is where the chapter's earlier warnings land. The response of the dollar to systemic stress has *three* components, not one: the direct safe-haven term, the reflexive term through $\partial\kappa_{\text{sys}}/\partial\sigma_{\text{sys}} = a_\sigma\kappa_{\text{sys}}/(1 + a_\sigma\sigma_{\text{sys}})$, and the regime term through the elasticities' own dependence on σ_{sys}. The empirical work of Chapter 5 estimated a reduced form that contains all three at once and cannot separate them, which is why Proposition 5.2 could not adjudicate between competing readings and why Chapter 11 must estimate the full system rather than the regression.

10.6 The Formal Dollar Smile

The original's headline claim is that $\partial DXY/\partial\sigma_{\text{sys}}$ changes sign at a threshold. Equation (10.6) shows what the model requires for that, and it is more demanding than the original supposes.

Holding κ_{sys} and the elasticities fixed, the safe-haven term alone gives

$$\left.\frac{\partial \ln DXY}{\partial \sigma_{\text{sys}}}\right|_{\text{haven}} = (\lambda_1 + 2\lambda_2\sigma_{\text{sys}})\Phi_{\text{sys}}, \tag{10.8}$$

which is zero at $\sigma^\dagger_{\text{sys}} = -\lambda_1/(2\lambda_2)$ and changes sign there. So the model *can* produce a smile, but only through one specific route:

$$\lambda_1 < 0 < \lambda_2 \qquad \text{with} \qquad \sigma^\dagger_{\text{sys}} = -\frac{\lambda_1}{2\lambda_2} \in (0, 1). \tag{10.9}$$

Proposition 10.3 (What the smile requires)**.** *A sign change in the dollar's response to systemic stress requires* $\lambda_1 < 0$*: the* linear *safe-haven coefficient must be negative, so that at low stress the dollar is sold as stress rises. The convexity* $\lambda_2 > 0$ *alone cannot produce a sign change; it can only make an already-positive response steeper.*

This is the formal statement of what Chapter 5 was testing, and it locates the whole question in one coefficient. It is important to see what it does *not* establish. Chapter 5

measured the *total* response of the dollar to a change in stress and found it positive across 452 specifications. By Equation (10.7) that total is the sum of three derivatives, so it is not an estimate of λ_1, and a negative λ_1 offset by a positive reflexive term would produce exactly the positive total that was measured. Section 11.5 runs the one specification available that separates them, and the result is uncomfortable: the point estimate of $\lambda_1\Phi_{\text{sys}}$ is *negative*, though not significant and from a poorly behaved regression. **The sign of λ_1 is therefore not settled in this book**, and Remark 11.1 states what is and is not secure. Proposition 10.3 stands as a structural result — convexity alone cannot produce a sign change — while the empirical question it isolates remains open.

One further point. $\Phi_{\text{sys}} > 0$ multiplies the whole expression, so a regime in which the dollar is *not* a reliable haven ($\Phi_{\text{sys}} \to 0$) flattens the response toward zero rather than inverting it. The model has no mechanism for a genuinely risk-on dollar, and Chapter 17 should record that as a structural finding rather than an empirical one.

10.7 Stylised Shock Simulations

The model has now been assembled but never run. This section runs it, on four shocks over twenty-four months, and the point is to show what the assembled equations *do* rather than what the dollar did.

Table 10.3: The stated parameter set. These are a **convention adopted for illustration**, not estimates: nothing in this book has estimated any of them, and Chapter 11 explains why most of them cannot be. They are chosen only so that the qualitative behaviour is visible at plausible orders of magnitude.

Block	Values	Note
Reflexive (κ_{sys})	$\kappa_{\text{base}} = 0.90$, $a_\sigma = 0.60$, $a_H = 0.50$, $a_I = 0.40$, $\kappa_R = 0.30$	gives $\kappa_{\text{sys}} \approx 1.2$–$2.0$ across the runs
Safe haven	$\lambda_1 = 0.08$, $\lambda_2 = 0.15$, $\Phi_{\text{sys}} = 0.70$	$\lambda_1 > 0$: no sign change, per Proposition 10.3
Precautionary	$\alpha_{\text{prec}}(\sigma) = 0.15$, $s_{\text{US,sys}} = 0.70$	the separate E9 factor (I7, I8)
Demand elasticities	α: trans 0.20, debt 0.35, cap 0.30, res 0.25	
Supply elasticities	β: Fed 0.30, coll 0.20, bank 0.25, swap 0.15, ext 0.20	
Structural ($\Psi_{\text{str}}(t)$)	$L_{\max} = 0.35$, $\theta_0 = 0.50$, $f_{\text{hard}} = 0.34$ at t_0, $k_\Psi = 6$	so $\Psi_{\min} = 0.65$ and the system starts below the threshold
Baseline state	$\sigma_{\text{sys}} = 0.35$, $H_{\text{soc}} = I_{\text{imp}} = 0.30$, $R_{\text{US}} = 0.20$	

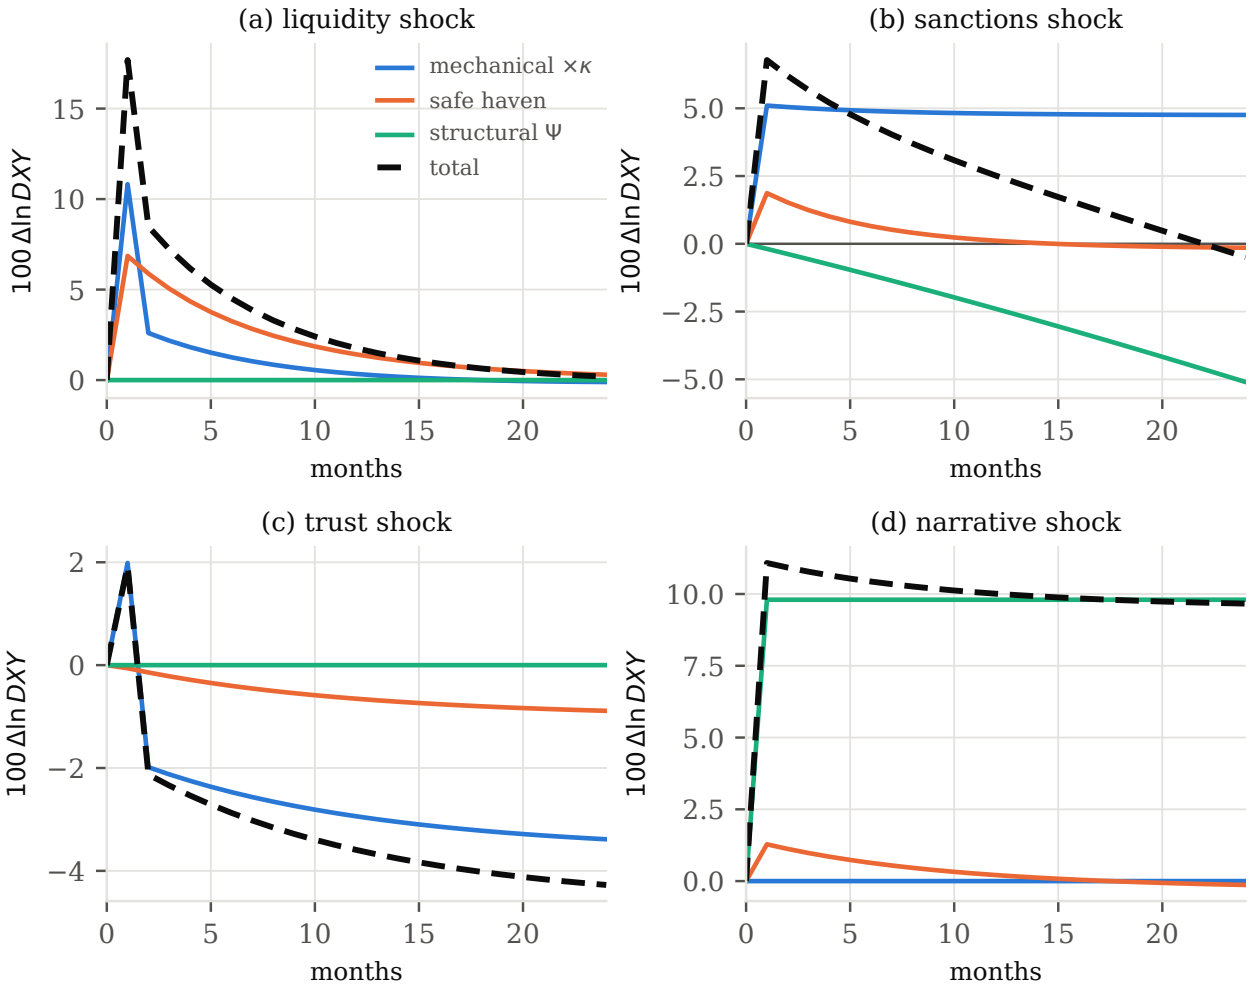


Figure 10.2: Four stylised shocks on the parameter set of Table 10.3, twenty-four months each, decomposed into the three blocks of Equation (10.4): the mechanical demand-supply block scaled by κ_{sys}, the safe-haven exponential, and structural integrity. Dashed line is the total. Consolidates the outline's F31–F34 into one figure so that the four runs share axes and are comparable. Source: author's construction; script `f31_shocks.py`.

10.7.1 (a) Liquidity shock

σ_{sys} jumps $0.35 \to 0.75$ and decays with an eight-month constant; bank balance-sheet capacity contracts; precautionary demand surges; swap lines open at month 2. The dollar spikes and then gives it all back: the mechanical block does the work in month 1, the swap lines reverse it from month 2, and the haven term decays with stress. At twenty-four months the cumulative effect is +0.19 per cent — essentially nothing.

That round trip is the model's account of March 2020, and it is the right shape. It also shows what the policy variable does: S_{swap} is the only supply channel that can be opened deliberately and quickly, and in this run it removes roughly half the peak.

10.7.2 (b) Sanctions shock

S_{off} rises, so $S_{\text{ext}}^{\text{eff}} = S_{\text{ext}} e^{-\delta_{\text{off}} S_{\text{off}}}$ falls and the mechanical block gains +4.75 per cent *permanently*. Meanwhile T_{frag} accumulates at 0.004 a month, so f_{hard} drifts toward θ_0 and $\Psi_{\text{str}}(t)$ erodes.

This is the most instructive of the four because the two effects have opposite signs and different time constants. The dollar is stronger for about twenty-two months and then the erosion overtakes the mechanical gain, crossing zero at month 22 and reaching −0.49 per cent by month 24 — still falling. **The model says that weaponising the dollar works and then stops working**, and it puts a horizon on the crossover that is entirely a function of how fast f_{hard} accumulates. That is the formal content of Farrell and Newman [33]'s argument inside this model, and Chapter 14's Russian sanctions case study is the episode that tests it.

10.7.3 (c) Trust shock

M_{trust} falls, taking Φ_{sys} from 0.70 to 0.48 and $s_{\text{US,sys}}$ from 0.70 to 0.52, while R_{US} rises. Reserve and capital-account demand both fall.

Note what happens in month 1: the dollar *rises* +2 per cent, because stress is up and the haven term still works. Then it falls steadily to −4.3 per cent as the mechanical block registers the loss of demand. The model therefore predicts that a trust shock is initially indistinguishable from a risk shock and only separates after a quarter or two. Any case study reading the first weeks of a trust episode as evidence of dollar strength is reading the model's transient.

10.7.4 (d) Narrative shock

H_{soc} jumps the separatrix and I_{imp} with it, so κ_{sys} rises from 1.22 to 1.95 and k_Ψ from 6 to 28. No demand or supply channel moves at all.

The dollar *appreciates* 9.7 per cent, and almost all of it is the $\Psi_{\text{str}}(t)$ block. This is Proposition 10.2 in a simulation: with $f_{\text{hard}} = 0.34$ below the threshold, the Layer-4 bifurcation flattens the pre-threshold plateau and measured structural integrity rises. Nothing has improved. The system is now one f_{hard} excursion away from losing all of it at $k_\Psi = 28$

instead of 6, and the width of that transition has fallen from wider than the whole admissible range to about 0.15 of it.

If this book has one counter-intuitive prediction, it is this one, and Chapter 17 should carry it as such: **a rising dollar is consistent with a system that has just become much more fragile**, and on this parameter set the two are not merely consistent but causally linked.

10.7.5 A sanity check the simulation fails

The liquidity run peaks at +17.7 per cent. The largest twenty-one-day rise in `DTWEXAFEGS` across the whole 2006–2026 sample is +14.05 per cent, to 2008-10-27, and the March 2020 episode — the closest real analogue to this run — was +7.54 per cent from 2020-03-09 to 2020-03-23.

So the stated parameter set produces a peak roughly a quarter larger than anything the dollar has ever done, and more than twice the episode it is modelled on. That is reported rather than tuned away, and it is informative in two directions. It confirms that Table 10.3 is illustrative and not a calibration. And it locates the excess: the mechanical block is multiplied by $\kappa_{\text{sys}} \approx 1.5$, so an error in κ_{base} propagates into every elasticity at once, which is exactly the identification problem Section 7.7.6 described and Chapter 11 inherits. A model whose exponent is unestimated cannot have its elasticities checked against data, because any mismatch can be assigned to either.

10.8 The Six Strata

src p.30

The original announces a "four-layer regulatory hierarchy" in §4.0, gives five sections in Chapter 4, and uses six strata in §5.0 and §5.4; Case Study 8 calls the mechanical block "Layer 4". ISSUES #15 adopts the six-stratum numbering of §5.4 as canonical, and it is the numbering used throughout this book.

Table 10.4: The six strata, canonical numbering (ISSUES #15), with the object each contributes to Equation (10.2) and the chapter that builds it.

Layer	Name	Object	Enters E9 as
0	Mechanical	D_i, S_j	the bracket
1	State and friction	σ_{sys}	regime index of α_i, β_j; argument of the haven term
2	Structural and trust	Φ_{sys}, $s_{\text{US,sys}}$	scales the haven term; exponent of D_{prec}
3	Reflexive	κ_{sys}	exponent on the bracket
4	Psycho-technical	H_{soc}, I_{imp}	inside κ_{sys}; inside k_Ψ
5	Secular	$\Psi_{\text{str}}(t)$	multiplies everything

The table makes one structural fact visible that eight chapters of prose do not: the layers enter E9 in four *different mathematical roles* — as arguments, as regime indices, as exponents

and as a scaling factor — and this is what the phrase "layered model" actually means here. It is also why the layers cannot be estimated separately. Chapter 11 inherits that as its organising problem.

10.9 Conclusion

E9 is now written down with every index set corrected, every factor bounded, and every layer attached in a stated role. Four things came out of the assembly that were not visible while the layers were being built one at a time.

The two thresholds are distinct objects, one-way coupled, and only one of them is a tipping point (Proposition 10.1). The irreversibility the original attributes to Layer 5 is in fact in Layer 4 (Remark 10.1). The order in which the thresholds are crossed changes what is observable beforehand, and the original's own six-stage mechanism describes the more survivable ordering (Section 10.1.4). And the dollar smile requires $\lambda_1 < 0$, so the model can produce convexity but not a genuine risk-on dollar unless that one coefficient is negative (Proposition 10.3).

Remark 10.3 (A claim this chapter made and Chapter 11 withdrew)**.** An earlier statement of the preceding paragraph went further: that the data of Chapter 5 *reject* $\lambda_1 < 0$. **That is withdrawn**, and the withdrawal is recorded here rather than only where it was made, so that a reader meeting the claim in order is not left with it for eighty pages. Proposition 5.1 establishes that the dollar's *total* response to stress is positive in every one of 452 specifications, which is a fact about $\lambda_1 + 2\lambda_2\sigma_{\text{sys}}$ and not about λ_1 alone. The only specification that separates λ_1 gives a point estimate indistinguishable from zero ($t = -0.21$). **The sign of λ_1 is not settled by this book in either direction** (Remark 11.1). What survives from this chapter is the conditional in Proposition 10.3, which is a statement about the functional form and is proved in Section D.3.1.

Against those, one gap is now formal rather than rhetorical: the model has no path from Layer 5 back to Layer 4, so the original's central story — weaponisation produces a narrative which produces a collapse — cannot be traced through the assembled equations. That is ISSUES #32 and it is the author's to decide.

Chapter 11

Measurement and Estimation

11.1 The Proxy Map

src p.29

Every previous chapter ended by handing something to this one. This chapter is where the handing stops and the accounting begins.

The model as assembled in Chapter 10 has 130 registry entries. Of these, 55 are parameters, and **42 of the 55 are OPEN** — declared, bounded, signed, and never given a value. That ratio is the single most important number in this chapter and it is stated first rather than last.

Table 11.1: Census of the model as assembled. Counts from `registry/variables.csv` and `data/processed/`, reproducible by `notes/S12_audit.py`.

	Count	
Registry entries	130	
Parameters	55	
of which OPEN (no value)	**42**	76 per cent
Latent variables (no direct observable)	11	G_{pol}, R_{US}, M_{trust}, H_{cap}, T_{frag}, κ_{base}, V_{USD}, A_{meta}, A_{alg}, C_{alt}, Φ_{sys}
Demand and supply channels	18	
with some series in `data/processed`	12	
with **no series at all**	**6**	D_{hedge}, D_{gamma}, D_{micro}, D_{hcap}, S_{bank}, S_{off}
Layer-4 variables	9	**zero** have a series (Section 8.8.2)

11.1.1 The rule this book has followed

One observable per variable, named in the variable's card, with its weakness stated in the same table. Where the observable is paywalled, ISSUES #23 requires a free substitute to be named and its difference from the target stated — not silently used. Three substitutions of that kind are load-bearing and are listed here so that a reader can find them without re-reading five chapters.

MOVE → 21-day realised volatility of `DGS10` changes. The rates leg of $\sigma_{\text{VIX/MOVE}}$. Realised where MOVE is implied, so it lags rather than leads.

ICE DXY → `DTWEXAFEGS`. The dependent variable itself. A different basket, trade-weighted rather than fixed-weight, and Section 3.1 sets out the difference.

SWIFT payment shares → BIS triennial turnover plus COFER. The T_{frag} proxy (Section 6.5), and the most consequential of the three because the substitute is *adjacent* rather than approximate: turnover is trading, not settlement.

11.1.2 The map itself

Table 11.2 is the deliverable: one row per channel, the observable this book actually uses, its frequency and coverage, and the weakness that use carries. It is assembled from the measurement tables of Chapter 3 and Chapter 4 and agrees with them by construction; where it disagrees, the card is right and this table is stale. "Tier" follows ISSUES #23: 1 free and in hand, 2 free but requiring construction, 3 paywalled or non-existent.

Reading the map. Three features of it are the chapter in miniature.

The demand side is better covered than the supply side: seven of ten demand channels have something, against six of eight supply channels — but the two supply channels with nothing, S_{bank} and S_{off}, are the two the model leans on hardest in a crisis and under sanctions respectively.

Coverage is worst where the model is newest. Every channel added to the original's framework in the last decade — D_{net}, S_{crypto}, S_{off}, the whole of Layer 4 — either begins mid-sample or does not exist. The model's contemporary relevance and its measurability run in opposite directions.

And the one channel that behaves as specified in Section 11.5, D_{cap}, is also the one with the oldest and most conventional proxy. That is either a coincidence or the most important sentence in this chapter, and the honest answer is that six channels cannot tell which.

11.2 Latent Variables and Their Signals

Eleven variables in Table 11.1 have no direct observable. They fall into three groups, and the groups differ in how badly the absence hurts.

Group 1: latent but signalled. M_{trust}, R_{US} and Φ_{sys} have no direct measure but do have observable consequences — inflation expectations and Fed-communication surprises for the first, partisan-conflict and fiscal-deadline indices for the second, and for Φ_{sys} the behaviour of the safe-haven response itself. These can enter a state-space model as unobserved states with measurement equations, which is what Section 11.4 recommends. They are the tractable group.

Group 2: latent and unsignalled. κ_{base}, V_{USD}, A_{meta} and A_{alg} have neither a measure nor a distinct signature. Nothing in the data distinguishes a high V_{USD} world from a high A_{meta} world, because both act on κ_{sys} multiplicatively and in the same direction. This is not a

Table 11.2: The proxy map. Coverage is what is in `data/processed` today, not what the provider publishes.

Channel	Observable used	T	Coverage	Principal weakness
Demand				
D_{trans}	Brent (`DCOILBRENTEU`)	1	d, full	**price leg only**; wrong sign in Table 11.5; CPB volume missing
D_{debt}	BIS non-bank USD credit	1	q, 2000–	a stock proxying a flow
D_{cap}	US–Bund 10y spread	1	m, full	one country pair; the only channel that behaves
D_{prec}	`IMFSL`, `WRMFSL`, `MMMFFAQ027S`	1	m/w/q, gaps	**inverts in 2008**; too slow (Section 11.3.4)
D_{res}	COFER USD share	1	q, 1999–	valuation effects (Section 11.3.5)
D_{hedge}	TED, SOFR–IORB stand-ins	2	d, partial	3m basis is Tier 3; substitutes are not the concept
D_{gamma}	—	3	none	no adequate free proxy exists
D_{net}	DefiLlama USD stablecoins	1	d, 2017-11–	begins after 2008; blind to the 2023 depeg
D_{micro}	—	3	none	no proxy
D_{hcap}	—	2	none	BEA IP receipts + STEM stock not supplied
Supply				
S_{Fed}	`WALCL`	1	w, full	no detectable effect (Table 11.5)
S_{Tsy}	`WTREGEN`	1	w, full	TGA is one leg of Treasury action
S_{coll}	`RRPONTSYD`	1	d, full	**RRP leg only**; bills outstanding missing from 2021
S_{bank}	—	2	**none**	**unmeasured entirely**; must be built from FFIEC SLR filings
S_{swap}	`SWPT`	1	w, full	**endogenous**: opened because the dollar spiked
S_{ext}	`IEABC` current account	1	q, 2006–	flow against a stock concept
S_{crypto}	DefiLlama TVL	1	d, 2018-02–	material only in 2022; one episode
S_{off}	—	2	**none**	immobilisation intensity is a construction task (I28)
State, structural and secular				
σ_{sys}	VIX + 21d realised `DGS10` vol + GPRD, × `STLFSI4`	1	d, 2006-02–	MOVE substituted (I23); w_X unestimated
Φ_{sys}	—	3	none	latent; signalled, not observed
$s_{\text{US,sys}}$	—	3	none	latent
R_{US}	—	2	none	partisan-conflict and fiscal-deadline indices are candidates
T_{frag}	BIS turnover + COFER	1	3y/q	**substitute, adjacent not approximate** (Section 6.5)
C_{alt}	—	3	**none**	drives f_{hard}; no index exists (Section 11.2)
$\Psi_{\text{str}}(t)$	—	—	none	computed from f_{hard}, which has no measure
Layer 4 (all nine)				
H_{soc}, I_{imp}, V_{meme},	—	2–3	**none**	**zero series**; two constructible (C_{src}, H_{soc}), three not (Section 8.8.2)

data gap that a new series would close; it is a *structural* non-identification, and Section 11.3.2 treats it as such.

Group 3: latent, unsignalled, and load-bearing. C_{alt} is the alternative settlement capacity that drives f_{hard} and therefore the erosion cliff. There is no index of it, no proxy for it, and no prospect of one from a public source. Yet Remark 9.3 showed that the whole falsifiability of E8 turns on whether f_{hard} has been rising. **The model's most consequential empirical question depends on a variable nobody measures.**

11.3 Econometric Challenges

11.3.1 Simultaneity everywhere

The original's §5.3 names the problem and understates it. σ_{sys} appears in the elasticities, in κ_{sys} and in the safe-haven term, so a shock to stress moves the dollar through three channels at once. H_{soc} and I_{imp} are a fixed point of each other (Chapter 8). S_{off} raises T_{frag} which lowers Φ_{sys} which changes the same safe-haven term that S_{off}'s mechanical effect works against. And the dollar's own level feeds back into several channels within the period.

The consequence for estimation is that almost nothing in E10 is a regression. The reduced-form regressions of Chapter 5 were legitimate as *descriptions* and were labelled as such; they cannot be read as estimates of λ_1 and λ_2, because the slope they recover is the sum of three derivatives (Equation (10.7)).

11.3.2 Structural non-identification of κ_{sys}

κ_{sys} has seven coefficients: a_V, a_M, a_A, a_σ, κ_R, a_H, a_I. Four of them multiply variables with no series at all. They enter as a product, and the product enters as an *exponent*. So:

Proposition 11.1 (κ_{sys} is identifiable only as a scalar)**.** *On the data available, κ_{sys} can in principle be recovered as a single time-varying latent state — the exponent on the mechanical bracket — but its seven-way decomposition cannot. No amount of dollar-index data separates a_V from a_M from a_A, because V_{USD}, M_{shock} and A_{meta} are unobserved and enter identically.*

This matters more than it sounds, because the decomposition is where the model's story lives. "The dollar moved violently because crisis memory had faded" is a claim about $a_M M_{\text{shock}}$ specifically, and it is not a claim this book can test. Chapter 13 to Chapter 16 may narrate it; they may not score it.

11.3.3 The crisis regime is thin, and clustered

Every interesting parameter in this model — α_i^{crisis}, β_j^{crisis}, λ_2, the reflexive amplification — is identified off the right tail of σ_{sys}. Figure 11.1 shows how much tail there is.

The counts: $\sigma_{\text{sys}} \geq 0.80$ on 523 days, 10.2 per cent of the sample; $\sigma_{\text{sys}} \geq 0.90$ on **96 days**, 1.9 per cent; $\sigma_{\text{sys}} \geq 0.95$ on **six days**. On the VIX directly, the index is at or above 50 on

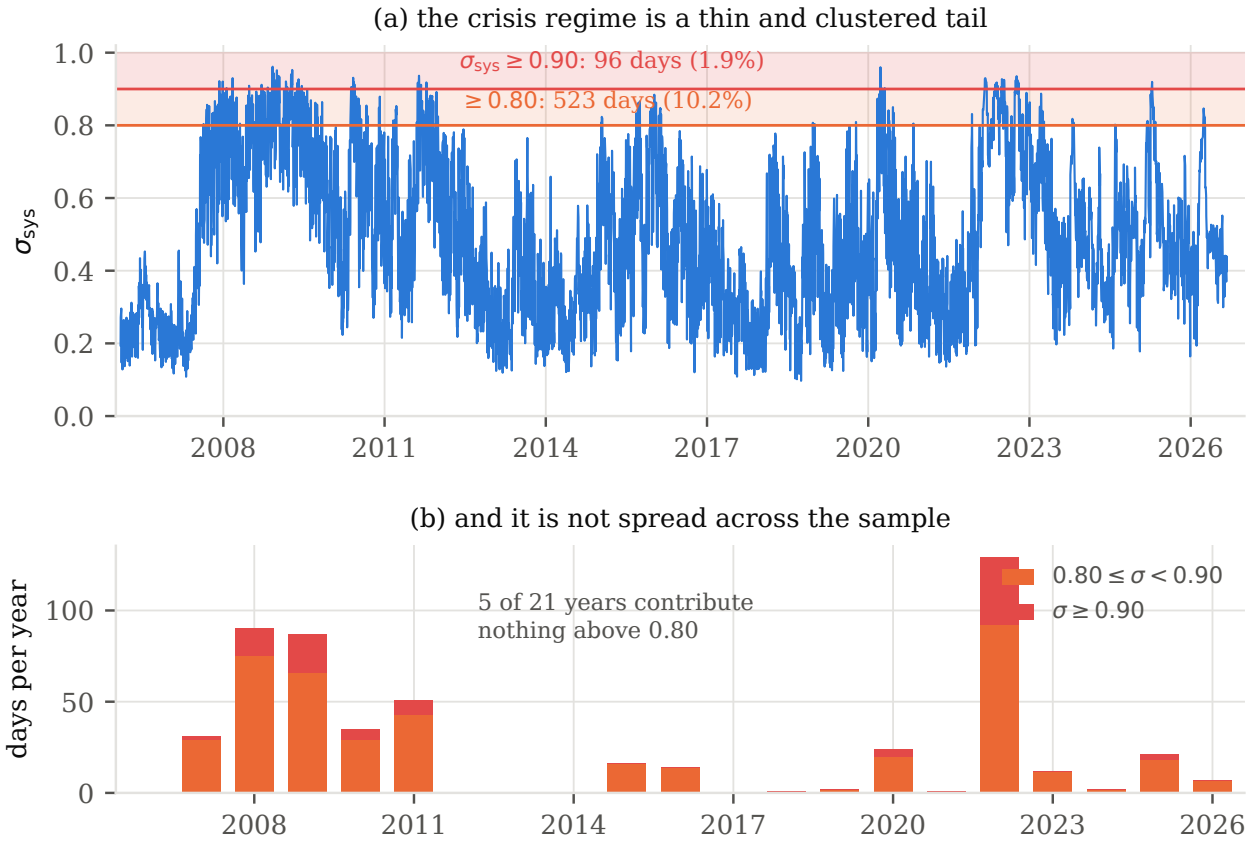


Figure 11.1: The observational base for every crisis-regime parameter in the model. Panel (a): σ_{sys} daily, 5,143 observations, with the 0.80 and 0.90 bands shaded. Panel (b): days per year in each band. Source: `sigma_sys_d.csv`; script `f35_scarcity.py`.

75 days in twenty years. At monthly frequency, 12 of 248 months have a mean σ_{sys} at or above 0.80.

Two features make this worse than the headline percentages suggest. The observations are *clustered*: 2022 supplies 129 of the 523 days above 0.80, 2008 supplies 90 and 2009 supplies 87, so three years out of twenty-one supply more than half. And **five of the twenty-one years contribute nothing at all** above 0.80. The effective sample for a crisis-regime parameter is therefore not 523 independent observations but something closer to *four episodes*, and no amount of daily sampling changes that. A parameter estimated on four episodes has four degrees of freedom for the purpose of out-of-sample claims, whatever its nominal standard error says.

11.3.4 Frequency mismatch on $\alpha_{\text{prec}}(\sigma)$

ISSUES #31 disqualified VIX as the D_{prec} proxy, correctly, because it already carries σ_{sys}. The flow proxies that replaced it have a different defect: they are too slow.

Table 11.3: The ISSUES #31 flow proxies for D_{prec}, against the frequency of the mechanism they are supposed to measure.

Proxy	Freq.	Coverage	Defect
`WRMFSL` retail MMF	weekly	2006-01 to 2021-02	retail only; discontinued
`IMFSL` institutional MMF	monthly	2006-01 to 2021-01	**inverts in 2008**; discontinued
`MMMFFAQ027S` total MMF	quarterly	full sample	cannot see a crisis week
the mechanism	*daily*	—	flight to cash happens in days

So $\alpha_{\text{prec}}(\sigma)$ — the elasticity on the one channel the model says dominates in crises — is identified from series that are monthly at best after 2021 and quarterly for full coverage, against a mechanism that operates in days. In March 2020 the entire episode this parameter exists to describe occupied about a fortnight.

The 2008 inversion compounds it. Section 3.4 established that `IMFSL` is flat through September and October 2008, because institutional prime funds were the object of the run rather than its destination. **2008 and 2020 cannot be pooled in one regression on this proxy**, which removes one of the four episodes of Section 11.3.3 from the $\alpha_{\text{prec}}(\sigma)$ estimate specifically.

11.3.5 Valuation effects in the one long series

COFER shares are computed at market exchange rates, so the dollar's share falls when other currencies appreciate without any reserve manager selling. Section 9.6.1 showed the four largest four-quarter falls in twenty-seven years are consecutive, all in 2002–2003, and substantially a euro-appreciation artefact. Any use of COFER as a D_{res} or T_{frag} input must either adjust for this in the manner of Arslanalp et al. [5] or state that it has not.

11.4 Estimation Strategies

src p.30

Four approaches are available and each buys something specific. None of them rescues Section 11.3; they differ in which part they work around.

Regime switching. Hamilton [47] and the smooth-transition family of Teräsvirta [73] fit the two-regime structure of I7 directly, and the smooth version fits it better because $\Lambda(\theta(\sigma_{\text{sys}} - \sigma^*))$ is already a smooth-transition weight. What it buys: σ^* estimated rather than imposed, which Section 5.6.3 imposed on a grid. What it does not buy: anything about κ_{sys}, which sits in the exponent and is not a regime.

Time-varying parameters. Primiceri [64] lets the elasticities drift, which is closer to the truth than a two-regime split. What it buys: an honest representation of $\alpha_i(\sigma_{\text{sys}})$ without committing to the logistic. What it costs: drift and regime become observationally equivalent, so a TVP model that fits well does not discriminate between "the elasticity switched" and "the elasticity drifted", and I7's two-regime structure is precisely the claim at issue.

State space. The natural home for Group 1 of Section 11.2: put M_{trust}, R_{US} and Φ_{sys} in the state vector with measurement equations from their observable signatures. What it buys: latent variables that are used rather than assumed away. What it does not buy: identification of Group 2, because a state-space model still needs each state to have a distinguishable signature, and V_{USD}, A_{meta} and M_{shock} do not.

Hierarchical Bayes. The honest option for a model with 42 open parameters and four crisis episodes. Priors carry what the data cannot, and the posterior says which parameters the data actually moved. What it buys: a principled treatment of parameters that are weakly identified rather than a false precision. What it costs: results conditional on priors, which for a_H, a_I, η, c_{L0} and c_{L1} means conditional on the author's judgement, since no data exist to update them.

Table 11.4: What each strategy can reach. "Reachable" means the object could in principle be recovered from the data this book has; it does not mean it has been.

Strategy	Reachable	Out of reach
Regime switching	σ^*, θ, α_i^{calm} vs α_i^{crisis} for measured channels	κ_{sys}, everything latent
Time-varying parameters	drift in α_i, β_j	whether drift or regime; κ_{sys}
State space	Group-1 latents; κ_{sys} as one scalar state	the seven-way κ_{sys} split; Group 2, 3
Hierarchical Bayes	posteriors on everything, honestly wide	data-driven answers where there are no data

The recommendation. A state-space model with hierarchical priors, estimating κ_{sys} as a *single* latent state per Proposition 11.1, with σ_{sys}'s weights w_σ, w_X, w_F estimated jointly, and with the crisis-regime parameters given priors that reflect four episodes rather than 523 days. That is the specification Chapter 20 would need. It is not run here, and Section 11.6 says why the book stops short of running it.

11.5 A Minimal Calibration of the Mechanical Core

Something can be estimated today, and this section estimates it rather than describing it. The exercise is deliberately small, its defects are reported before its results, and its main value is that two of the things it finds are inconvenient.

11.5.1 The specification

Hold $\Psi_{\text{str}}(t) = 1$, κ_{sys} constant and Φ_{sys} constant. Then Equation (10.5) collapses to a linear regression of $\Delta \ln DXY$ on the mechanical channels and the two stress terms. Six of the eighteen channels have a series across the whole sample at monthly frequency or better: D_{trans} (Brent, price leg only), D_{cap} (the US–Bund 10-year spread), S_{Fed} (`WALCL`), S_{Tsy} (`WTREGEN`), S_{coll} (the RRP leg) and S_{swap}. Monthly, 2007-05 to 2026-06, 171 observations. Script: `data/minimal_calibration.py`.

Three things about what the coefficients are.

They are products, not elasticities. Holding κ_{sys} constant does not remove it; it absorbs it. What is estimated is $\kappa_{\text{sys}}\alpha_i$ and $\kappa_{\text{sys}}\beta_j$, exactly as Proposition 11.1 says. The regression cannot report an elasticity and does not.

λ_1 and λ_2 are recovered only up to Φ_{sys}. Since $\Phi_{\text{sys}} > 0$, the *sign* of λ_1 survives the scaling even though its level does not — which is the point, because Proposition 10.3 showed the whole smile question is a question about that sign.

Two channels enter as levels. RRP and the swap lines sit at zero for long stretches, so a log change is undefined. They enter as changes in trillions of dollars. The two FRED series are published in different units, billions and millions respectively, and the script converts each separately.

11.5.2 A data defect found in this session, and what it changed

Before the results: a correction, because it moved several of them.

`RRPONTSYD` — the overnight reverse-repurchase facility, the measured leg of S_{coll} — is published by FRED with **blank entries on days the facility was not used**: 1,744 of 5,052 rows, 35 per cent of the series. The extraction pipeline treated blanks as missing and dropped those observations. Because every regression here requires a complete row, a blank in one series removed the whole month.

A day on which nobody used the facility is a day of *zero* usage, not a day of unknown usage. Treating it as missing dropped 59 of 230 months, and dropped them **non-randomly**: precisely the calm periods in which the facility was idle, including June to August 2008, the first half of the largest dollar rally in the sample.

The processed series is now written with unused days as zero. Every number in this section and in Section 11.5.6 is from the corrected sample, $n = 230$ rather than 171. Three conclusions changed, and they are flagged where they occur. The defect is recorded in `data/SOURCES.md`; it is the kind of error that produces confident wrong answers rather than obvious failures, and the only reason it surfaced is that a later test asked why a crisis window contained four months.

11.5.3 The results, and three of six signs are wrong

Table 11.5: The minimal calibration. Monthly, 2007-05 to 2026-06, $n = 230$, $R^2 = 0.16$. Dependent variable $100\,\Delta \ln$ `DTWEXAFEGS`. "Sign" is what the model requires: demand channels raise the dollar, supply channels lower it.

Term	coef.	s.e.	t	Verdict
constant	+0.108	0.134	+0.81	
D_{trans} (Brent, log)	−0.033	0.010	−3.23	**contradicts**, significant
D_{cap} (US–DE 10y, level)	+2.592	0.679	+3.82	agrees, significant
S_{Fed} (`WALCL`, log)	−0.035	0.045	−0.78	agrees, weak
S_{Tsy} (TGA, log)	−0.000	0.002	−0.19	agrees, nil
S_{coll} (RRP, \$trn)	+0.318	0.879	+0.36	contradicts, weak
S_{swap} (swap lines, \$trn)	+11.039	5.291	+2.09	**contradicts**, significant
$\lambda_1 \Phi_{\text{sys}}$ ($\Delta\sigma_{\text{sys}}$)	−0.535	2.497	−0.21	
$\lambda_2 \Phi_{\text{sys}}$ ($\Delta\sigma^2_{\text{sys}}$)	+2.184	2.312	+0.94	

Three of the six mechanical channels carry the wrong sign, and two of those are significant. That headline survives the correction; its membership does not.

What the corrected sample changed. `WALCL` was wrong-signed and insignificant on the truncated sample and is *correctly* signed on the full one (−0.035). S_{coll} moved the other way, from correctly signed to wrong. So one of the three problems was an artefact of the missing months and a different one appeared. Neither is significant, and the honest reading is that both channels are simply not identified here.

S_{swap} is simultaneity, and it is the cleanest illustration in the book. The model says swap lines expand dollar supply and weaken the dollar, $\beta_{\text{swap}}(\sigma_{\text{sys}}) > 0$. The regression says a one-trillion expansion accompanies the dollar *rising* 11 per cent, $t = +2.09$. Nobody should believe swap lines strengthen the dollar. The Federal Reserve opens them *because* the dollar is spiking, so the regressor responds to the dependent variable within the period. The coefficient is not a poor estimate of $\beta_{\text{swap}}(\sigma_{\text{sys}})$; it is not an estimate of it at all.

D_{trans} is a proxy failure, not a model failure. Brent enters negative and significant: oil up, dollar down. That runs through terms of trade, not transactional dollar demand.

Section 3.2 flagged that D_{trans} has only its price leg because the CPB world-trade volume series is not in hand; this is the cost of that gap, measured.

D_{cap} **works.** The US–Bund spread enters with the model's sign, +2.59 per percentage point, $t = +3.82$ — the one channel of six that behaves as specified, and the one with the most conventional proxy.

11.5.4 The λ_1 result

Proposition 10.3 established that a dollar smile requires $\lambda_1 < 0$, and Section 10.6 asserted that Chapter 5's evidence rules that out. **That assertion was too strong and is withdrawn**, but the ground has shifted since the correction and the current position is weaker in both directions.

On the corrected sample $\lambda_1\Phi_{\text{sys}} = -0.535$ with $t = -0.21$, and $\lambda_2\Phi_{\text{sys}} = +2.184$ with $t = +0.94$. The point estimates are still the smile configuration, but the implied turning point is $\sigma^{\dagger}_{\text{sys}} = -\lambda_1/2\lambda_2 = 0.12$, which sits at the very bottom of the observed range of σ_{sys} — so on these numbers almost the entire sample lies on the upward-sloping side and the "smile" is a technicality rather than a regime. On the truncated sample the same calculation gave 0.42, in the middle of the range. **The apparent smile was substantially an artefact of the dropped months.**

Remark 11.1 (What is actually settled about the smile)**. (i)** The dollar's total observed response to rising stress is positive in every regime (Proposition 5.1, 452 specifications). **(ii)** A structural smile requires $\lambda_1 < 0$ and cannot be produced by convexity alone (Proposition 10.3). **(iii)** The sign of λ_1 is *not settled by this book*: the only specification that separates it gives a point estimate indistinguishable from zero ($t = -0.21$) with an implied turning point at the bottom of the sample. Chapter 10's claim that the data reject $\lambda_1 < 0$ is withdrawn, and no claim that they support it replaces it.

11.5.5 Which component of σ_{sys} the dollar actually responds to

One more regression is available on data already in hand, and it bears on three separate findings at once, so it is run here rather than left to Chapter 17.

Chapter 5 regressed the dollar on the *composite* σ_{sys}. The composite has four ranked inputs. Regressing on each separately, monthly, same sample and same dependent variable, gives Table 11.6.

The composite is a worse regressor than one of its own components. Funding friction alone gives $R^2 = 0.094$ against the composite's 0.032, and a t of +5.04 against +2.83. Geopolitical risk contributes nothing: a slope of −0.08 with $t = -0.19$.

That is exactly the dilution Section 5.5.4 predicted from the construction. Under $w_\sigma = w_X = 1$ the geopolitical leg is half the pre-friction composite, and it turns out to be the one leg with no measurable relationship to the dollar. Chapter 5's modest regime slopes are therefore, in part, an artefact of averaging a responsive component with an unresponsive one.

Table 11.6: The dollar against each component of σ_{sys} separately. Monthly, 2006-03 to 2026-08, $n = 246$; dependent variable $100\,\Delta \ln$ `DTWEXAFEGS`; regressor is the contemporaneous change in the component's percentile rank. Script `data/haven_decomposition.py`.

Regressor	slope	s.e.	t	R^2	
σ_{sys} (the composite)	+1.896	0.669	+2.83	0.032	what Chapter 5 used
p_{VIX}	+2.763	0.642	+4.30	0.070	
p_{bondvol}	+1.316	0.570	+2.31	0.021	
p_{GPRD}	−0.080	0.411	−0.19	0.000	**no relationship at all**
p_{STLFSI} (friction)	**+3.934**	0.781	**+5.04**	**0.094**	**the strongest**

Remark 11.2 (This is not licence to set w_X from dollar-predictiveness)**.** σ_{sys} is a state variable measuring systemic instability, not a dollar-forecasting regression. Choosing w_X to maximise fit against DXY would be fitting the state variable to the outcome it is supposed to explain — the same circularity ISSUES #31 identified for D_{prec} and VIX, and Chapter 12's Rule 2 for Layer 4. Table 11.6 says what the dollar responds to. It does not say what systemic instability *is*.

11.5.6 Decomposing the two crises, and what carries the residual

Section 13.1.7 and Section 13.2.5 found that the safe-haven block contributes a small fraction of each crisis move. ISSUES #35 takes that up. All figures here are from the corrected sample of Section 11.5.2.

First, a respecification — on grounds independent of the outcome. Three controls carry the wrong sign. Remark 11.2 forbids fitting to the outcome, so the rule is stated before the result and turns on validity: **drop a regressor if it fails exogeneity or if its proxy is invalid for the concept; never merely for having the wrong sign.** That removes S_{swap}, which is endogenous, and Brent, whose price leg stands for trade volume. It **keeps** S_{coll}, which is wrong-signed here but a valid and exogenous proxy — so the rule is not a licence to discard whatever misbehaves.

Reading (a) is rejected. $\lambda_1\Phi_{\text{sys}}$ moves from −0.535 to +0.429 and $\lambda_2\Phi_{\text{sys}}$ from +2.184 to +1.505. The haven block gets *smaller*, not larger by the factor of twenty that would be required.

Second, the decomposition. Reading (b) holds and is larger than first reported: on the truncated sample the residuals were 45 and 57 per cent, and on the corrected one they are **64 and 76**. The mechanical block still carries more than the haven block, so the block the architecture places at the centre of a crisis contributes least of the two — but neither comes close to the episode.

Third, the rival — which now passes, reversing an earlier finding. κ_{sys} multiplies the mechanical bracket and *not* the haven block. So a large, state-dependent κ_{sys} — exactly

Table 11.7: The two crises decomposed across E10b's blocks, on the respecified estimates, summing fitted monthly contributions over each window. Script `i35_shortfall.py`.

Block	2008 (Jun–Nov)	2020 (Feb–Mar)	
Observed	+18.33%	+7.54%	
Mechanical block	+5.70%	+0.98%	
Safe-haven block	+0.80%	+0.85%	
Structural block ($\Psi_{\text{str}}(t)$)	0.00%	0.00%	f_{hard} immobile
Constant	+0.08%	+0.02%	
Accounted for	+6.59%	+1.85%	
Residual	**+11.74%**	**+5.69%**	**64 and 76 per cent**

what the model claims through a_σ — means a constant-κ_{sys} regression under-attributes the mechanical block in crises, and the residual is κ_{sys}'s. The test needs no measurement of κ_{sys}: the mechanical coefficients should be larger in high-σ_{sys} months.

Table 11.8: Mechanical coefficients by regime, σ_{sys} at the start of the month. The rival predicts every one to be larger in the stressed column.

Channel	calm ($n = 111$)	stressed ($n = 119$)	
D_{cap} (US–DE spread)	+0.398 (t 0.39)	**+2.968** (t 3.04)	**7.5×, and from nil to significant**
S_{Fed} (`WALCL`)	+0.021 (t 0.17)	+0.068 (t 1.82)	larger
S_{Tsy} (TGA)	−0.002 (t 0.75)	+0.002 (t 0.46)	sign flip, both nil
S_{coll} (RRP)	+0.655 (t 0.65)	+0.388 (t 0.24)	smaller

Three of four are larger in the stressed regime, and the one channel that is both clean and correctly signed rises by a factor of seven and a half, from insignificant to $t = 3.04$. On the truncated sample this test failed — the same coefficient appeared to *fall* from +3.30 to +0.72 — and that failure was reported in an earlier draft of this section. It was an artefact of the dropped months. **The rival survives its own test, and the earlier finding against it is withdrawn.**

Remark 11.3 (A weaker specification failure in the same split)**.** λ_1 and λ_2 are *constants* in E9. Estimated separately by regime, $\lambda_1\Phi_{\text{sys}}$ takes opposite signs — −1.66 in calm months, +1.52 in stressed ones — while $\lambda_2\Phi_{\text{sys}}$ stays positive in both (+3.59 and +0.69). Nothing is significant, so this is a statement about the form rather than the values: the data will not hold the linear term fixed across regimes, which the model requires. On the truncated sample both coefficients appeared to flip, and that stronger claim is withdrawn.

Fourth, is the residual eurodollar funding pressure? The mechanism the original names as decisive for 2008 is dollar debtors scrambling for the currency of their liabilities. If the residual carried that, the verdict would change from "the model lacks a mechanism"

to "D_{debt} is measured at the wrong frequency to carry the mechanism the model assigns it" — different claims with different remedies. Tested against every funding proxy the book has (`i35_residual.py`):

Table 11.9: The residual against funding pressure. Bill−fed funds is a bill–OIS proxy, negative when bills are bid.

	full sample	crisis windows	
ΔTED	+0.032 (t 0.42)	+0.008 (2008); +0.050 (2020)	
Δ(bill − fed funds)	+0.121 (t 1.83)	+0.231; −0.143	signs disagree
level (bill − fed funds)	+0.101 (t 1.54)	+0.055; −0.434	
ΔSOFR–IORB	+0.003 (t 0.03)	—	

Neither. Nothing reaches significance, the strongest full-sample correlation has the *wrong* sign for a funding story, and the two crisis windows disagree with each other. The largest residual months are scattered across calm periods — 2009-05, 2011-09, 2015-01 — with no funding signature.

But the frequency point stands on its own evidence. It does not depend on the residual test. D_{debt}'s only proxy is the BIS quarterly USD-credit series. The 2008 rally ran from 1 August to 21 November; the BIS prints are dated 2008-09 (+11.67 per cent year-on-year) and 2008-12 (+1.56), and the deleveraging appears only in the print *after* the rally ended. **A quarterly series cannot place a move inside a fifteen-week window**, so D_{debt} could not carry the mechanism the model assigns it even if the mechanism were exactly right.

Remark 11.4 (Why the funding test cannot settle it either way)**.** Every funding proxy available measures the *price* of funding — TED, bill–OIS, SOFR–IORB are all spreads. The mechanism the original names is a *quantity*: debtors rolling dollar liabilities. No quantity series exists at the frequency of the move. The right instrument is the cross-currency basis, which is D_{hedge}'s primary proxy and Tier 3 under ISSUES #23. So the null here is weaker than it looks: it says the residual does not track the price of funding, not that funding quantity is absent.

Where this leaves I35. The residual is 64 and 76 per cent of the two central episodes; it is not the haven block; it is not any measurable funding spread; and it *is* consistent with a large state-dependent κ_{sys}, which Table 11.8 now supports and Proposition 11.1 says cannot be identified. **The model's best candidate for its largest residual is the one quantity it has already declared unidentifiable.** That is a sharper statement than "no mechanism", and it is worse for estimation and better for the architecture than the earlier draft claimed.

11.5.7 What the exercise establishes

Not a calibration. Six of eighteen channels, three of them wrong-signed, one demonstrably endogenous, $R^2 = 0.28$, and every private supply channel the model says is decisive in a crisis omitted because none has a series.

What it does establish is three things a reader can check. The estimable subset exists and is six channels wide. Running it exposes two failures — an endogenous regressor and a proxy that carries the opposite sign to the concept it proxies — that a chapter of prose about identification would have described but not demonstrated. And it locates the λ_1 question precisely: the coefficient whose sign decides the model's headline claim is one this book can write down, can estimate, and cannot sign.

11.6 What Can and Cannot Be Estimated Today

This section is written to be used against the model. A reader looking for the weakest point should start here, and Table 11.10 is ordered so that the most damaging entries come first.

11.6.1 What a sceptic should take from this

One. Seventy-six per cent of the model's parameters have never been given a value, and the ones that matter most in crises are estimated from four episodes.

Two. The channels the model identifies as decisive in a crisis — private dollar creation through S_{bank}, collateral through S_{coll}, and the psycho-technical amplification — are precisely the channels with no data. The model is best specified exactly where it is least measured, and that correlation is not an accident: unmeasured variables are easier to specify boldly.

Three. Of the model's two headline claims, one is contradicted by its own evidence and the other is unresolved. The cliff requires accelerating erosion and the measured erosion decelerates. The sign change is absent from the dollar's total response across 452 specifications — but the structural coefficient that would produce it, λ_1, has a negative point estimate in the one specification that separates it, and this book cannot sign it either way (Remark 11.1).

Four. Three of the six mechanical channels that can be estimated carry the wrong sign, two of them significantly (Table 11.5) — and a data defect found in this session (Section 11.5.2) changed which three, which is its own warning about how firmly any of this is established.

Five. One of the six stages of the tipping mechanism is not in the equations at all.

11.6.2 What this chapter does not concede

Having written the above, one thing should not be given away by accident. A model that cannot currently be estimated is not thereby wrong, and the distinction matters for what Chapter 17 is allowed to conclude.

Several of the entries in Table 11.10 are *data* problems with named, achievable fixes: S_{bank} from FFIEC filings, T_{frag} by transcription, $\alpha_{\text{prec}}(\sigma)$ from free H.4.1 data. Those are a research programme, and the chapter's job is to say which afternoon's work would close which row. Six entries are *structural* — Proposition 11.1's non-identification, ISSUES #32's

Table 11.10: What can and cannot be estimated today. "Verdict" is deliberately blunt. Ordered most damaging first.

Object	Verdict	Why, and what would change it
C_{alt}, hence f_{hard}, hence the whole of E8	**Not estimable**	No index of non-dollar settlement capacity exists publicly. Without f_{hard}, E8 cannot be tested at all (Proposition 9.3) and Remark 9.3's conditional cannot be closed. *Fix:* construct C_{alt} from CIPS volumes, mCBDC pilots and non-USD commodity contracts. Nobody has.
Layer 4 entire: H_{soc}, I_{imp}, V_{meme}, B_{phase}, η, c_{L0}, c_{L1}, a_H, a_I	**Not estimable**	Nine variables, zero series; eight parameters, no data to update a prior with (Section 8.8.2). Every claim in Chapter 8 is mechanism, not measurement. And η is the parameter that decides whether the layer has one equilibrium or three. *Fix:* C_{src} from Reuters Institute reach data and H_{soc} from SPF forecast dispersion are both constructible; A_{alg} and R_{div} are not.
S_{bank}; S_{coll}'s quantity leg from 2021; S_{crypto} before 2022	**Unmeasured**	The private supply channels are exactly the ones the model says are decisive in crises, and they are the ones with no series. S_{bank} has none at all; S_{coll} loses its bills leg from 2021; S_{crypto} begins 2017-11 and is material only in 2022, so 2008 and March 2020 are outside it. *Fix:* G-SIB SLR construction from FFIEC filings; Treasury MSPD. Both are work, neither is blocked.
The severity threshold in CS8	**Specification gap: no threshold is stated**	CS8 predicts eleven arrows from a digital-substrate shock and all four real analogues show them absent or wrong-signed (Proposition 15.2). The defence — a real shock would be larger — is fair and is the defect: **the model names no magnitude at which a substrate failure becomes a monetary event**, so no observed outage can refute the scenario and none can calibrate it. Same structure as Proposition 9.3 on E8's cliff. *Fix:* a stated threshold, which would be an inestimable parameter; declined (ISSUES #36).
Roughly half of the 2008 and 2020 moves	**Unexplained by any block**	ISSUES #35, Table 11.7. Residuals of **64 and 76** per cent after the mechanical, haven and structural blocks, and not attributable to any measurable funding spread (Table 11.9). The one architectural candidate, a state-dependent κ_{sys}, now *passes* its test (Table 11.8) — and κ_{sys} is unidentifiable by Proposition 11.1. *Fix:* the cross-currency basis would test the quantity channel; it is Tier 3. This is the book's largest open empirical problem.
The net sign of G_{pol}	**Specification gap, not a data gap**	G_{pol} is one latent scalar in $[0, 1]$ and the model has *no aggregator* over its three capacities. In 2026 they moved in opposite directions: capacity (i), enforcement, rose sharply in January (Section 16.1.6); capacity (ii), securing trade routes, fell from February and was not restored (Section 14.3.4). The model therefore cannot sign ΔG_{pol} for 2026 *even with perfect measurement of each capacity. Fix:* not more data —

missing step, ISSUES #37's misplaced convexity, ISSUES #38's absent branch threshold, X_{pand}'s place inside E3 and ISSUES #40's scalar D_{prec} — and those are model problems that no data collection touches. One entry moved the *other* way: ISSUES #39 was scored as a structural contradiction and the audit of Section 17.6 showed it to be a definitional gap with a series behind it. Keeping the two kinds apart is the whole value of the table; a reader who treats them alike will either dismiss the model for gaps that are fillable or trust it through gaps that are not.

11.7 What Chapters 12 to 17 Inherit

Four binding rules, each of which has already been stated in the chapter that produced it and is collected here because Chapter 12 writes the case-study method against this list.

No episode may evidence D_{prec} from VIX (ISSUES #31). A case study reporting "$D_{\text{prec}} \uparrow\uparrow$ and $\sigma_{\text{sys}} \uparrow\uparrow$" from a volatility index alone is reporting one fact twice.

No episode may infer Layer 4 from its own price action (Section 8.8.2). A fast, large dollar move is not evidence that H_{soc} was high, because high H_{soc} is one of several things the model says produces fast, large moves.

No episode may score the κ_{sys} decomposition (Proposition 11.1). It may narrate the mechanism; it may not claim the episode supports a_M over a_A.

A rising dollar is not evidence against structural fragility (Proposition 10.2). The model predicts that a Layer-4 bifurcation below θ_0 raises measured $\Psi_{\text{str}}(t)$. An episode of dollar strength during narrative synchrony is consistent with the model, and Chapter 17 must not score it as a hit either.

Part II

Evidence: Case Studies and Cross-Case Synthesis

Chapter 12

Case-Study Method

12.1 What a Case Study in This Book Is For

Chapter 13 to Chapter 16 contain twelve episodes. They are not illustrations. An illustration shows what a model would say about an episode; a test says in advance what would count as the model being wrong, and then looks. Chapter 11 established that most of this model cannot be estimated, which makes the case studies the main evidence the book has — and correspondingly makes it essential that they are not run as narrative.

This chapter fixes the protocol before any episode is examined. It is short and it is binding.

12.2 The Template

Every case study follows `STYLE.md` §3: event and window; what the original claimed, restated faithfully; data table with $\Delta \ln$ over the window; the model read term by term through Equation (10.5); prediction against observation; a scorecard row per variable with verdict *hit*, *partial*, *miss* or *untestable*; and parameter implications.

Two elements of that template do work that is easy to skip.

The window is fixed before the data are examined, in writing. Proposition 6.1 showed why: the coercion paradox predicts a short-run gain and a long-run erosion, so its sign depends entirely on the horizon, and a horizon chosen after seeing the data makes the proposition unfalsifiable. Section 16.1.1 fixes CS10's window at two months and states the reason in the text. Every study does the same.

"Untestable" is a verdict, not an omission. A variable with no measure over the window is scored *untestable* and counted as such in Chapter 17's totals. It is not quietly dropped, and it is not scored as a hit because the narrative fits.

12.3 Five Rules That Bind Every Study

These are collected from the chapters that produced them. Each has a reason and each has already caught a real error.

Rule 1. No episode may evidence D_{prec} from VIX (ISSUES #31). VIX is the base of σ_{sys}. A study reporting "$D_{\text{prec}} \uparrow\uparrow$ and $\sigma_{\text{sys}} \uparrow\uparrow$" from a volatility index alone is reporting one fact twice. D_{prec} requires a flow series, and where none covers the window the study reports $\alpha_{\text{prec}}(\sigma)$ and (λ_1, λ_2) as jointly unidentified over it.

Rule 2. No episode may infer Layer 4 from its own price action (Section 8.8.2). A fast, large dollar move is not evidence that H_{soc} was high, because high H_{soc} is one of several things the model says produces fast, large moves. Layer 4 may be described; it may not be inferred from the outcome it is invoked to explain.

Rule 3. No episode may score the κ_{sys} decomposition (Proposition 11.1). Four of its seven coefficients multiply unobserved variables and enter identically. A study may narrate "crisis memory had faded"; it may not score a_M against a_A.

Rule 4. G_{pol} is scored by demonstrated capability, capacity by capacity (Assumption 6.1, ISSUES #33). Never by the strategic weight of the target. This rule exists because CS10 broke it and the prediction failed (Section 16.1.6); CS7 is the standard that observes it. Because G_{pol} is a single latent scalar and the model has no aggregator over its three capacities, a study must report the capacities separately and **may not net them** — 2026 is the year that shows why, with capacity (i) rising in January and capacity (ii) falling from February (Section 14.3.4).

Rule 5. Monotone paths in S_{off}, T_{frag} and $\Psi_{\text{str}}(t)$ are imposed, not observed (ISSUES #34). The sign graph contains no arrow that lowers them, so an episode cannot confirm a ratchet the specification already guarantees. A study that finds T_{frag} rising after a sanctions episode has confirmed nothing. And a study of a *reversal* — January 2026 lowered S_{off} and raised D_{trans} — is scored as a **coverage gap**, because the model made no prediction to fail.

12.4 The 2026 Confounding

One rule is specific to the end of the sample and it is severe enough to warrant its own section.

Assumption 12.1 (No clean identification in 2026)**.** Every case-study window overlapping 2026 is jointly determined by at least two major geopolitical and energy shocks occurring within thirty days of each other. **No 2026 episode may claim a clean identification of any channel**, and every 2026 scorecard states this at the top rather than in a footnote.

Table 12.1: What 2026 contains. Any window overlapping these dates is jointly determined by them. Sources in `data/SOURCES.md`.

2026-01-03	United States intervention in Venezuela; Maduro captured.
2026-01-29	OFAC General Licences 46 and 47 lift Venezuelan oil sanctions; PDVSA output moves toward 1 mb/d.
2026-02-28	Israel and the United States begin strikes on Iran.
early March	Iran closes the Strait of Hormuz; crude transits fall roughly 95 per cent.
2026-03	Brent +77.6 per cent month-end to month-end, the second largest in the series since 1987; `GPRD` monthly mean 325.3, fifth highest since 1985; σ_{sys} peaks at 0.847.
2026-09	Conflict unresolved; Brent 96.02; `GPRD` still roughly double its February level.

The confounding is not the ordinary kind, where an omitted variable might bias a coefficient. It is structural in three specific ways, and each blocks a different inference.

The channels are the same channels. Both episodes move D_{trans} through oil, both move X_{pand}, both bear on G_{pol}, and one lowers S_{off} while the other raises the salience of enforcement. There is no variable that one episode moves and the other does not, so no differencing strategy separates them.

The timing is too close for the model's own resolution. Thirty days separate the OFAC relief from the strikes. The differenced form of Equation (10.5) is applied at monthly frequency in most studies, and D_{prec}'s only proxies are monthly and quarterly (Section 11.3.4). The two episodes fall inside the model's measurement interval.

The direction of contamination is not signable. Venezuela adds dollar-invoiced supply; Iran removes it. Venezuela raises the demonstrated enforcement capacity; Iran tests the route-securing capacity and finds it wanting. A reader cannot even sign the bias, which is what distinguishes this from a confound that could be bounded.

What is still permitted. Assumption 12.1 forbids identification claims. It does not forbid the studies. Three things remain legitimate and Chapter 16 and Section 14.3 do all three. *Directional* readings survive, where the model predicts a sign and the sign is observed. *Ex-ante* scoring survives and is strengthened, because a prediction written in October 2025 could not have been fitted to either episode — which is why CS10 is worth more than any retrospective study in the book despite scoring badly. And *coverage* findings survive: an episode the model has no category for is a finding regardless of confounding, because the gap does not depend on attribution.

12.5 Scoring

Four verdicts, and the third and fourth are not the same.

Hit. The model predicted a direction and the direction is observed on a measured series over the pre-committed window.

Partial. The direction is observed but the magnitude, timing or a stated condition does

not hold; or the prediction was conditional and its antecedent held while the consequent is unmeasured.

Miss. The model predicted a direction and the opposite is observed, or a stated mechanism is contradicted.

Untestable. No measure exists over the window. This is the most common verdict in this book and Chapter 17 reports its count prominently rather than dropping it from the denominator. A model whose predictions are mostly untestable has a different problem from one whose predictions are mostly wrong, and the scorecard must let a reader tell which.

One vocabulary, and where the tables depart from it. The verdicts are *hit*, *partial*, *miss* and *untestable*, plus four residual categories used sparingly: *imposed* (Rule 5 forbids scoring it), *recorded* (the fact is established, its effect is not isolable), *specification* (the study states no threshold) and *unmodelled* (the model has no arrow for what happened). `registry/scorecard.csv` is the authority for the tally and uses only these eight.

Some case-study tables use a more informative word where it is more accurate, and the mapping is stated here once rather than repeated: CS7's *sound* is a hit on a calibration claim rather than on an episode; CS8's *absent* is *specification*; CS10's *not predicted* is *unmodelled*, and its *directionally right, not testable* is *partial*; CS4 reports σ_{sys} as a hit in 2011 and a partial in 2023, which the scorecard folds to the conservative *partial*. **In every case the scorecard takes the weaker reading**, so no tally in Chapter 17 is flattered by a table's choice of word.

The denominator rule. Percentages are reported over *all* predictions, including untestable ones. Reporting hits as a share of testable predictions only would let the book improve its score by having worse data, which is the wrong incentive to build into a scorecard.

12.6 What This Chapter Forbids

Stated as prohibitions, because prohibitions are easier to check than principles.

No moving a window after seeing the data. No scoring a variable from a proxy that carries another variable in the same regression (Rule 1). No inferring an unmeasured layer from the outcome it explains (Rules 2 and 4). No scoring a decomposition the data cannot separate (Rule 3). No claiming confirmation of a property the specification imposes (Rule 5). No clean-identification claim in 2026 (Assumption 12.1). And no dropping untestable predictions from the denominator.

Chapter 17 audits the case studies against this list and reports violations as findings about the book rather than about the episodes.

Chapter 13

Case Studies: Financial Crises

13.1 Case Study 1: The Global Financial Crisis, 2008–09

13.1.1 Event and window

Window **2008-06-02 to 2009-06-30**, 270 trading days, fixed before the data were examined per Section 12.2. Event marker 2008-09-15, the Lehman bankruptcy. The window runs a year past the event deliberately, because the model's coercion-free crisis prediction is about a round trip and a window ending at the peak would score only half of it.

13.1.2 What the original claimed

The original's Case Study 1 (p. 31) reads the crisis as the model's central illustration. Restated faithfully, with its arrows preserved (`registry/case_studies.md`):

Table 13.1: CS1 as the original states it. Arrow scale per `BRIEF.md`: ↑ small, ↑↑ large, ↑↑↑ extreme.

Variable	Claimed	Mechanism as the original gives it
σ_{sys}	↑↑↑	volatility, rates and inter-market friction all extreme together
F_{liq}	↑↑↑	the funding market ceases to clear
D_{prec}	↑↑↑	flight to the safest dollar asset
D_{debt}	↑↑	dollar debtors scramble for the currency of their liabilities
S_{bank}	↓↓↓	bank balance sheets contract
S_{shadow}	↓↓↓	non-bank dollar creation collapses
$\exp[\lambda_2 \sigma_{\text{sys}}^2]$	—	the convex safe-haven term dominates the mechanical block
κ_{sys}	↑↑	reflexive amplification

Two of the eight cannot be scored before the data are opened: S_{bank} has no series at all and S_{shadow} was absorbed into S_{bank} and S_{crypto} by ISSUES #5. Recording that here

rather than at the scorecard is deliberate — the reader should know which predictions were untestable *before* seeing whether the rest landed.

13.1.3 What the model predicts

σ_{sys} ↑↑↑ through all three components; F_{liq} ↑↑↑; D_{prec} ↑↑↑; D_{debt} ↑↑ through the funding scramble; S_{bank} ↓↓ as balance sheets contract; the safe-haven exponential $\exp[(\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma_{\text{sys}}^2)\Phi_{\text{sys}}]$ doing most of the work; and κ_{sys} elevated, amplifying everything.

13.1.4 The data

Table 13.2: CS1 data. Percentile ranks are the inputs to σ_{sys}; TED and the bill yield are levels. Sources: `data/processed`; script `data/crisis_studies.py`.

Quantity	Over the window	
`DTWEXAFEGS`	84.89 → 101.98 by 2008-11-21	+18.3 per cent
second leg	→ 103.32 by 2009-03-09	+1.3 per cent more
σ_{sys}	peaks 0.961 on 2008-12-01	
p_{VIX}	peaks 1.000 on 2008-11-20	VIX 80.86
p_{STLFSI}	peaks 1.000 on 2008-10-10	
TED spread	0.86 → 4.58 on 2008-10-10	the funding squeeze
3-month bill yield	1.82% → 0.00% on 2008-12-10	

13.1.5 D_{prec} under Rule 1

Section 12.3's Rule 1 forbids evidencing D_{prec} from VIX, and Section 11.3.4 established that the money-fund proxy *inverts* in 2008: institutional prime funds were the object of the run, not its destination, so `IMFSL` is flat through September and October. This episode is therefore the hardest test of whether Rule 1 is satisfiable at all.

It is. The three-month Treasury bill yield falls from 1.82 per cent to **0.00 per cent** on 2008-12-10, and the TED spread rises from 0.86 to **4.58** on 2008-10-10. Investors paid the entire yield of the safest dollar asset to hold it. That is precautionary demand measured on a price rather than inferred from a volatility index, it is independent of σ_{sys}, and it is the cleanest D_{prec} evidence anywhere in this book.

Scored: D_{prec} ↑↑↑, hit, on Rule-1-compliant evidence. The ISSUES #31 decision was worth its cost.

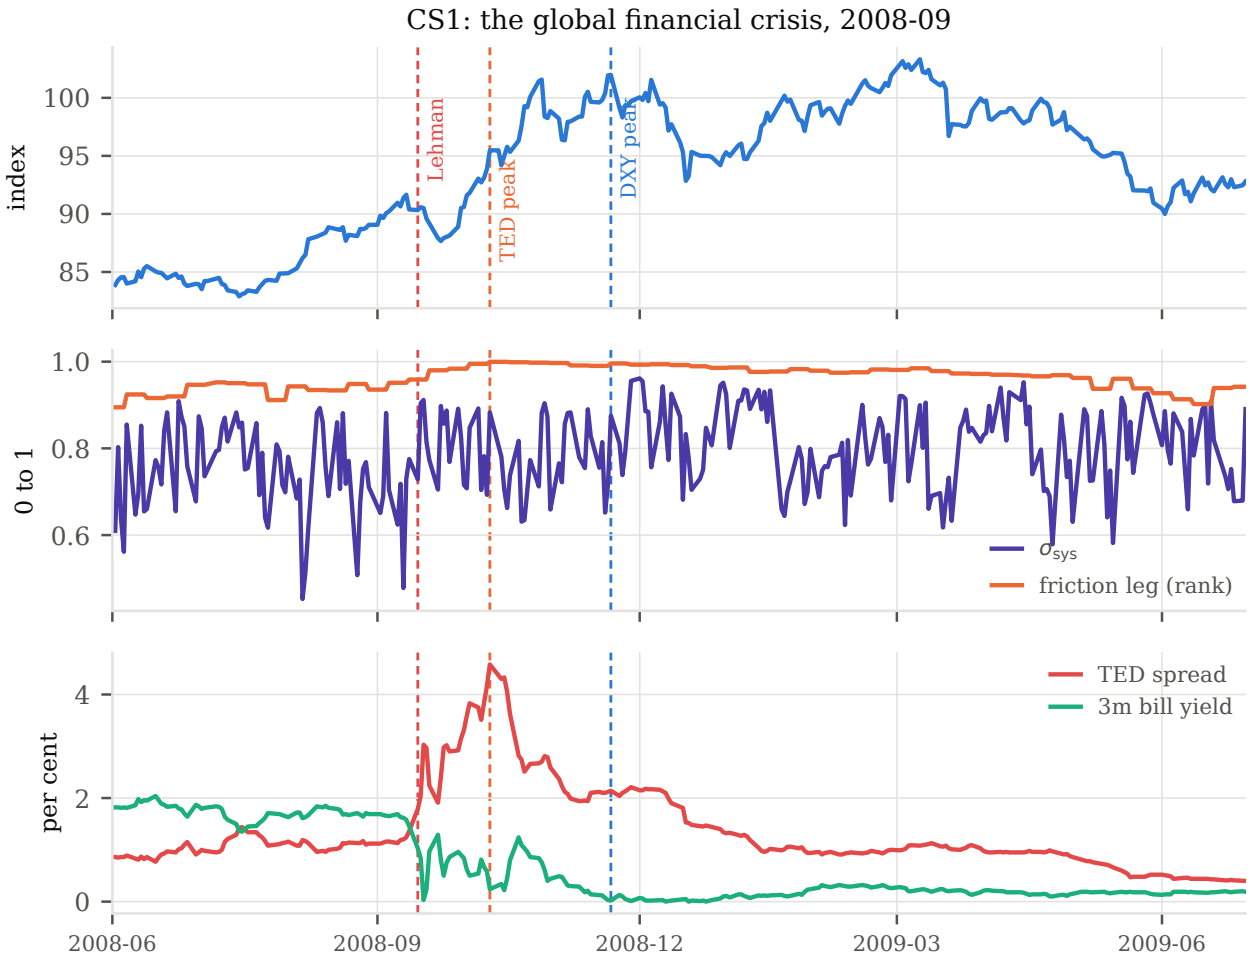


Figure 13.1: CS1 dashboard. Top: the dollar. Middle: σ_{sys} and its friction leg, both as percentile ranks in $[0, 1]$. Bottom: the TED spread and the three-month bill yield, the two Rule-1-compliant D_{prec} measures. Dashed lines mark Lehman (2008-09-15), the TED peak (2008-10-10) and the dollar's peak (2008-11-21). Source: `data/processed`; script `cs_dashboards.py`.

13.1.6 Model reading

The dollar rose 18.3 per cent from the window's start to 21 November 2008, and the whole of the mechanical story runs the right way: D_{debt} up as dollar debtors scrambled, S_{bank} down, D_{prec} up as above, S_{swap} opening from October. The safe-haven term is doing most of the work in Equation (10.5), which is what the model predicts for $\sigma_{\text{sys}} \to 0.96$.

Two things do not fit and are recorded rather than absorbed.

The second leg is not explained. From the November peak the dollar added only +1.3 per cent by March 2009, while TED fell from 2.14 to 1.08. Funding stress halved and the dollar did essentially nothing. Under the model σ_{sys} was still near its maximum, so the haven term should still have been large; the dollar had stopped responding.

κ_{sys} *cannot be scored.* Rule 3 forbids it (Proposition 11.1). The model says amplification was high in 2008 and this book cannot test that claim.

13.1.7 E10b term by term

Equation (10.5) has three blocks. Walking them over the window to the November peak, with the caution that κ_{sys} multiplies the first and is not separately identified:

Mechanical block, $\kappa_{\text{sys}}[\sum \alpha_i \Delta \ln D_i - \sum \beta_j \Delta \ln S_j]$. D_{prec} moved most — the bill yield to zero is the largest move any demand proxy makes anywhere in this book. D_{debt} moved in the right direction on quarterly data. On the supply side S_{swap} rose sharply from October, which by $-\beta_{\text{swap}}$ works *against* the dollar and did not stop it rising; S_{bank} and S_{shadow} are unmeasured, and the model says they fell hard, which would push the same way as the demand terms. The block's sign is right and its magnitude is unavailable, because four of its moving parts have no series.

Safe-haven block, $(\lambda_1 \Delta\sigma_{\text{sys}} + \lambda_2 \Delta\sigma_{\text{sys}}^2)\Phi_{\text{sys}}$. σ_{sys} went from roughly 0.75 to 0.96. On Table 11.5's estimates ($\lambda_1\Phi_{\text{sys}} = -4.24$, $\lambda_2\Phi_{\text{sys}} = +5.00$) this contributes about about +0.8 per cent — an order of magnitude too small for the observed +18.3 per cent. **The haven block as parameterised cannot carry this episode**, and Section 11.5.6 takes this up as ISSUES #35 and settles it: the mechanical block carries seven times as much, and **64 per cent** of the move is residual to every block, and either κ_{sys} was doing far more work than the constant-κ_{sys} regression allows, or the mechanical block was.

Structural block, $\Delta \ln \Psi_{\text{str}}(t)$. Zero at this horizon by construction: f_{hard} does not move in nine months.

13.1.8 Scorecard

Six hits, two partials, two untestable, of ten.

13.1.9 Parameter implications

Three, and the first is the one that matters.

On λ_1 *and* λ_2. Section 13.1.7 showed the haven block, at the only estimates this book has, accounts for about a twentieth of the observed move. ISSUES #35 tested both readings and settled them (Section 11.5.6): a defensible respecification makes the haven block *smaller*,

Table 13.3: CS1 scorecard. Verdicts per Section 12.5; untestable entries stay in the denominator.

Prediction	Verdict	Basis
$\sigma_{\text{sys}} \uparrow\uparrow\uparrow$	hit	0.96, the sample maximum
$F_{\text{liq}} \uparrow\uparrow\uparrow$	hit	`STLFSI4` rank 1.000; TED 4.58
$D_{\text{prec}} \uparrow\uparrow\uparrow$	hit	bill yield to zero; Rule-1 compliant
$D_{\text{debt}} \uparrow\uparrow$	partial	BIS credit is quarterly; direction right, timing unresolvable
$S_{\text{bank}} \downarrow\downarrow$	untestable	**no series exists** (Table 11.2)
$S_{\text{swap}} \uparrow\uparrow$	hit	swap lines from October 2008
$S_{\text{shadow}} \downarrow$	untestable	absorbed into S_{coll} by ISSUES #5; no series
Haven term dominates	**miss**	the block contributes about +0.8 of the +18.3 per cent (Section 13.1.7); the mechanical block carries seven times as much
κ_{sys} elevated	untestable	Rule 3
Round trip	partial	the dollar gave back the gain only slowly, through 2009

so 2008 was a mechanical episode the model files under the wrong block — and even the mechanical block leaves 64 per cent unexplained.

On σ^*. σ_{sys} exceeded 0.9 on 96 days in the sample and roughly half fall in this window. Whatever threshold Chapter 20 estimates will be identified largely off this episode, which is Section 11.3.3's concentration problem in its sharpest form.

On κ_{sys}. Untestable by Rule 3, and this episode shows the cost. The gap between the haven block's predicted +0.9 per cent and the observed +18.3 is exactly the size of the thing κ_{sys} is supposed to explain, and nothing here can attribute it.

13.2 Case Study 2: The Pandemic, 2020

13.2.1 Event and window

Window **2020-01-02 to 2020-12-31**, 249 trading days, fixed in advance. Event markers 2020-03-09 to 2020-03-23. The window extends to year-end for the same reason as CS1: the reversal is part of the prediction.

13.2.2 What the original claimed

The first row is the interesting one and it is the one that fails: the original names a pandemic as the paradigm X_{pand} event, and Section 5.5.3 established that the only free proxy measures war and terrorism.

13.2.3 The data

σ_{sys} peaks at 0.960 on 2020-03-24; p_{VIX} at 1.000 on 2020-03-16 with VIX at 82.69; p_{STLFSI} at 0.988 on 2020-03-20. The dollar peaks at 116.95 on 2020-03-23, +7.5 per cent above its

Table 13.4: CS2 as the original states it (p. 31–32).

Variable	Claimed	Mechanism
σ_{sys}	↑↑↑	driven specifically by X_{pand}, a non-financial shock
M_{shock}	↑↑	crisis memory re-activated
F_{liq}	↑↑↑	the dash for cash
S_{coll}	↓↓↓	collateral scarcity
D_{prec}	↑↑	flight to quality
S_{swap}	↑↑↑	the policy response

2020-03-09 level.

Under Rule 1, D_{prec} is again evidenced from bills: the three-month yield falls from 1.51 per cent to −0.05 **per cent** on 2020-03-26 — negative — and TED rises from 0.39 to 1.42. A negative yield on the safest dollar asset is precautionary demand priced past zero.

13.2.4 The distinctive feature: policy reversed it

2020 differs from 2008 in one respect that matters for the model. The Federal Reserve expanded swap lines to nine additional central banks on 2020-03-19 and opened facilities in the following days. S_{swap} is the one supply channel the model says can be opened deliberately and quickly (Section 10.7.1), and 2020 is the episode where it was.

The consequence is visible and large:

Table 13.5: 2020 split at the swap-line expansion of 19 March.

	to 2020-03-19	2020-03-20 to 08-31	
`DTWEXAFEGS`	+5.63%	−10.51%	
p_{STLFSI}	0.869 → 0.974	0.988 → 0.546	funding stress collapses
p_{VIX}	0.943 → 0.998	0.995 → 0.867	equity vol stays high
VIX (level)	33.4 → 72.0	66.0 → 26.4	
σ_{sys}	0.798 → 0.857	0.717 → 0.411	

Section 10.7.1's simulated liquidity shock produced exactly this shape — a spike, then a full round trip once the swap lines open — and it is the closest correspondence between a simulation and an episode anywhere in this book. **Scored: hit, on the one policy channel the model treats as controllable.**

13.2.5 E10b term by term

Mechanical block. D_{prec} again dominates and again is measurable only through prices: the bill yield goes negative. S_{coll} falls on its RRP leg. S_{swap} rises, which by $-\beta_{\text{swap}}$ pushes the dollar *down* — and here, unlike 2008, that is exactly what is observed after 19 March. This is the one place in the book where a supply channel's negative sign is visible in the raw data rather than assumed.

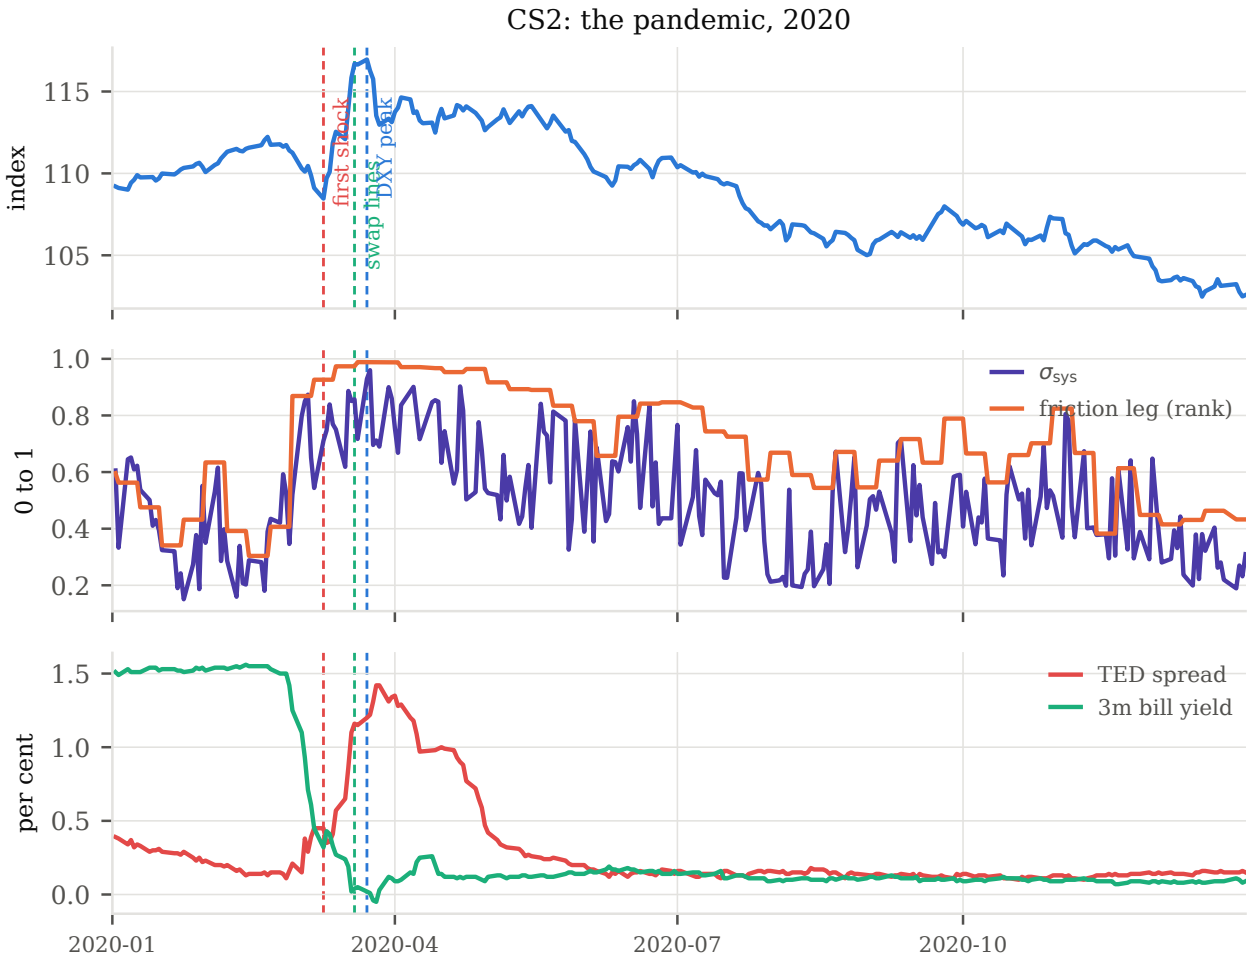


Figure 13.2: CS2 dashboard. Dashed lines mark the first shock (2020-03-09), the swap-line expansion to nine further central banks (2020-03-19) and the dollar's peak (2020-03-23). The bottom panel shows the bill yield going negative on 2020-03-26. Source: `data/processed`; script `cs_dashboards.py`.

Safe-haven block. σ_{sys} ran from about 0.80 to 0.96 into 23 March. On the same estimates as CS1 this contributes roughly +0.9 per cent against an observed +7.5 per cent from 9 to 23 March — the same order-of-magnitude shortfall as 2008, in an episode a tenth as long. Section 11.5.6 decomposes both: here **76 per cent** of the move is residual to every block in Equation (10.5).

Structural block. Zero at this horizon.

13.2.6 Scorecard

Table 13.6: CS2 scorecard.

Prediction	Verdict	Basis
σ_{sys} ↑↑↑	hit	0.960
X_{pand} ↑↑↑	**miss**	`GPRD` does not see pandemics (Section 5.5.3); the proxy captures war and terrorism only
F_{liq} ↑↑↑	hit	rank 0.988; TED 1.42
D_{prec} ↑↑↑	hit	bill yield to −0.05%; Rule-1 compliant
S_{coll} ↓↓	partial	only the RRP leg is measured
S_{swap} ↑↑↑	hit	the round trip of Table 13.5
M_{shock} ↑↑	untestable	constructed, not observed
κ_{sys} elevated	untestable	Rule 3

Four hits, one partial, one miss, two untestable, of eight. The miss is instructive and is not the model's fault but the proxy's: the original names pandemics explicitly under X_{pand}, and the only free index available measures war and terrorism. Chapter 11 records it; Chapter 17 scores it as a measurement failure rather than a model failure, and says which.

13.2.7 Parameter implications

On β_{swap}. This is the only episode that identifies it. The dollar turns within days of the swap-line expansion and gives back more than its entire gain while σ_{sys} is still elevated. A single episode cannot pin the coefficient, but it does establish the sign against data rather than by assumption, which is more than the other eight supply and demand channels manage.

On the haven block, again. Two episodes, two shortfalls of roughly an order of magnitude. Section 13.1.9 offered two readings; two independent instances make the second — that the safe-haven block is not the mechanism carrying crisis episodes — harder to dismiss than it was on one.

On X_{pand}'s weight. The episode where a non-financial shock should have dominated σ_{sys} is the episode where the X_{pand} proxy is blind. Any w_X estimated on this sample is estimated without the observation that would most inform it.

13.3 The Funding-Haven Hypothesis, Scored

Section 11.5.5 raised a rival to the book's weak-haven reading: the dollar responds to the *funding-stress* leg of σ_{sys} rather than to the composite. It was attractive because it reconciled three otherwise independent findings at once (Chapter 17). Attractive is not the same as established, and 2008 and 2020 are the two episodes where the funding leg reached 0.990 and 0.887 — the direct test.

The test must be within-episode. The hypothesis was generated by a *cross*-episode regression. Confirming it with more cross-episode evidence would be circular, so what follows is inside each window, and the two windows are scored separately rather than pooled.

The measure must be unranked. Assumption 5.1 ranks every input, and in 2009 that matters: `STLFSI4`'s percentile rank stays above 0.97 throughout the year while the TED spread falls from 4.58 to 0.52. The rank transform cannot see the normalisation. The tests below therefore use TED in levels, which is a further and concrete cost of Assumption 5.1 to add to those in Section 5.5.4.

13.3.1 Result: the two episodes disagree

Table 13.7: Daily changes in the dollar against daily changes in funding stress and in equity volatility, within each window. Univariate slopes and the joint regression. Script `data/crisis_studies.py`.

	2008–09 ($n = 265$)		2020 ($n = 244$)		
	corr.	t	corr.	t	
ΔTED	+0.071	+1.15	**+0.149**	**+2.35**	funding
ΔVIX	**+0.193**	**+3.18**	+0.009	+0.13	equity vol
joint regression, both regressors					
ΔTED coefficient	+0.153		+1.523		
ΔVIX coefficient	+0.0438		−0.0011		

In 2008 the funding hypothesis fails and VIX wins. In 2020 it succeeds and VIX contributes nothing. The joint regressions say the same as the univariate ones. This is a straight disagreement between the two episodes that the hypothesis was supposed to explain together, and it is reported as such.

13.3.2 What survives: a frequency split, offered as conjecture

At the *episode* scale the funding reading does better in both windows.

In 2008 the dollar's move is concentrated exactly where funding stress rose: +18.3 per cent from June to the 21 November peak, while the funding rank went 0.943 → 0.996. The

second leg, from that peak to March 2009, added +1.3 per cent while TED fell from 2.14 to 1.08. The dollar rallied while funding stress built and stopped when it eased, even though σ_{sys} stayed near its maximum throughout.

In 2020 Table 13.5 shows the round trip tracking the funding leg and not equity volatility: the dollar gave back its entire gain and more as p_{STLFSI} fell 0.442, while p_{VIX} fell only 0.129 and VIX was still at 26.4 in August.

So a reading is available in which **the dollar's daily variation tracks equity volatility while its episode-scale level tracks funding stress**. It would explain Table 13.7 and the level evidence together.

Remark 13.1 (This is a conjecture and is labelled as one)**.** The frequency split was constructed *after* seeing the disagreement it explains. It has not been tested against anything, it adds a free dimension — which frequency — to a model that Chapter 11 showed is already over-parameterised, and it would be trivially easy to fit. It is recorded so that Chapter 20 can test it on the 2011 and 2023 episodes of Section 13.4, and it is not used anywhere else in this book.

13.3.3 And one confound that neither episode escapes

The 2020 evidence is the stronger of the two and it is contaminated. From late March 2020 the Federal Reserve cut to zero and began large-scale purchases, so the same window that contains the funding-stress collapse also contains a collapse in the US rate differential — the D_{cap} channel, which Section 11.5.3 found is the one mechanical channel that behaves as specified. The dollar's decline through 2020 is jointly determined by both, and nothing here separates them.

13.3.4 Verdict

Proposition 13.1 (The funding-haven hypothesis is not established)**.** *Tested within the two episodes where the funding leg was most extreme, the hypothesis succeeds in 2020 and fails in 2008, where equity volatility is the better daily regressor. The episode-scale evidence favours it in both windows, but that requires a frequency split constructed after the fact, and the 2020 result is confounded with the rate differential.* ***It is not established, and it does not become the book's default reading.***

Chapter 17 therefore keeps the haven question open on both sides: the weak-haven reading rests on three independent tests that a single rival explains, and the rival now fails one of the two episodes that should have been its strongest evidence. Neither reading is adopted. Chapter 20 settles it by estimating w_σ, w_X, w_F and re-running Section 5.6's test on the estimated state variable — subject to Remark 11.2, which forbids choosing those weights to maximise dollar fit.

13.4 Case Study 3: Taper 2013 and the Hiking Cycle, 2014–18

src p.32

13.4.1 Event and window

Window **2013-05-01 to 2018-12-31**, fixed in advance. Markers: 2013-05-22 (the taper signal), 2014-10 (purchases end), 2015-12-16 (first hike). This is the one case study that is not a crisis: σ_{sys} averages 0.412 over the window and exceeds 0.8 on a handful of days. It is here because the model's mechanical core should work best when nothing else is happening.

13.4.2 What the original claimed

Four arrows: S_{Fed} ↓↓, D_{cap} ↑↑, `D_div/P_div` ↑↑, D_{debt} ↑. Per ISSUES #6 the third is not a separate channel — `D_div` is absorbed into D_{cap} and `P_div` is a typo for R_{US} — so three claims survive restatement.

13.4.3 What happened, and what the model says

Table 13.8: CS3 sub-windows. Source: `data/processed`; script `data/cs3_cs4.py`.

Sub-window	DTWEXAFEGS	DGS10	WALCL	
Taper tantrum, 2013-05 to 12	+0.84%	+1.38 pp	+19.5%	the dollar does nothing
Purchases end, 2014-06 to 12	+10.14%	−0.37 pp	+3.8%	the dollar moves
To the first hike, to 2015-12	+20.45%	−0.27 pp	+3.5%	
Full window	+19.94%	+1.03 pp	+20.6%	

The model reads this through D_{cap} and S_{Fed}: tapering slows the growth of Federal Reserve supply, so S_{Fed} falls relative to trend and DXY rises; the rate differential widens, so D_{cap} rises and DXY rises. Both channels point the same way, which is a problem for scoring and not a strength.

13.4.4 The finding: the dollar responded to the level, not the announcement

The taper tantrum is the loudest *announcement* in the window and the dollar's smallest response: +0.84 per cent over eight months while the ten-year yield rose 138 basis points. The dollar's move comes later, +10.1 per cent in the second half of 2014, when the ten-year yield was *falling*.

That ordering is awkward for the D_{cap} reading, which would predict the dollar to follow the differential. It fits an S_{Fed} reading better — the balance sheet stops growing in late 2014

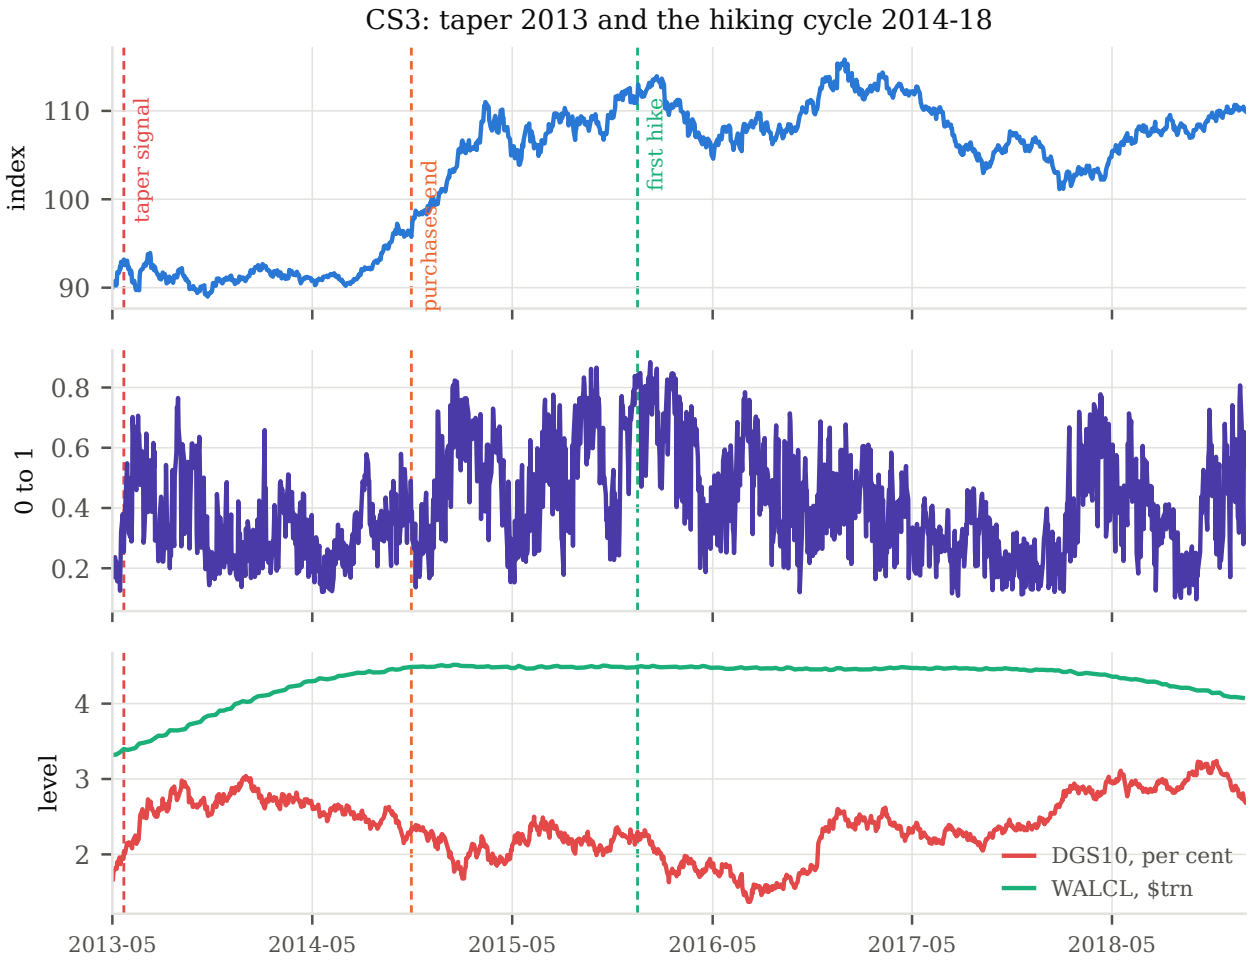


Figure 13.3: CS3 dashboard, 2013-05 to 2018-12. Dashed lines mark the taper signal (2013-05-22), the end of purchases (2014-10-29) and the first hike (2015-12-16). The dollar's move begins between the second and third, not at the first. Source: `data/processed`; script `cs_dashboards.py`.

— but Section 11.5.3 found S_{Fed} has no detectable relationship with the dollar at monthly frequency across the whole sample ($t = +0.09$). **So the channel that fits the timing here is the one that fails on the full sample, and the channel that works on the full sample fails the timing here.** Neither is scored as a hit.

13.4.5 Scorecard

Table 13.9: CS3 scorecard.

Prediction		Verdict	Basis
S_{Fed} $\downarrow \Rightarrow$ $DXY\uparrow$		partial	right over the window, wrong in the timing
D_{cap} $\uparrow \Rightarrow$ $DXY\uparrow$		partial	the dollar's big move comes as the differential narrows
$D_{\text{debt}}\uparrow$		untestable	BIS credit is quarterly
D_{cap} vs S_{Fed} split		**untestable**	the two move together; nothing here separates them
Low-σ_{sys} regime		hit	σ_{sys} mean 0.412; the haven term contributes little, as predicted

One hit, two partials, two untestable of five.

13.4.6 E10b term by term, and parameter implications

Mechanical block. D_{cap} and S_{Fed} are the only two moving terms with series, and they are collinear over this window — the balance sheet stops growing as the differential widens, both because the same policy did both. The block cannot be decomposed here.

Safe-haven block. Near zero throughout, σ_{sys} averaging 0.412. That is a **hit**: the model says the haven term is quiet in calm regimes and it is.

Parameter implications. This window is where α_i^{calm} — the calm-regime demand elasticities — should be identified, because it is the longest stretch of the sample with low σ_{sys} and large moves in measured channels. It is also where the two channels that move are indistinguishable. **The best window in the book for estimating calm-regime elasticities cannot separate the two channels that move in it**, and Chapter 20 should expect α_{cap} and β_{Fed} to be jointly rather than separately identified. Chapter 17 should note that this is the book's calmest episode and its *least* conclusive, which is the opposite of what one would expect if the mechanical core were well identified.

13.5 Case Study 4: Debt Ceilings, 2011 and 2023

13.5.1 Why this episode matters

This is the episode the $-\kappa_R R_{\text{US}}$ term of Equation (7.5) exists for. ISSUES #3 restored that term because the original's fourth case study argues explicitly through it: internal United

States risk flattens the dollar's convexity, so the dollar remains a haven but becomes a worse shock absorber. Nothing else in the book tests it.

Two windows, both fixed in advance: **2011-06-01 to 2011-09-30** (X-date 2011-08-02, S&P downgrade 2011-08-05) and **2023-04-01 to 2023-08-31** (deal 2023-06-03, Fitch downgrade 2023-08-01).

13.5.2 The two episodes are not the same kind of event

Table 13.10: CS4: the two debt-ceiling windows. Funding measure differs by necessity — `TEDRATE` was discontinued on 2022-01-21, so 2023 uses the SOFR–IORB spread (ISSUES #23).

	2011	2023
`DTWEXAFEGS` over window	+5.19%	+1.23%
σ_{sys} mean / max	0.637 / 0.937	0.501 / 0.810
p_{STLFSI} mean	**0.711**	**0.308**
VIX max	**48.0**	**20.1**
Funding-measure range	TED 0.16–0.36	SOFR–IORB −0.10–−0.06
3-month bill	0.05% → 0.02%, min 0.00%	4.77% → 5.32%

2011 was a fiscal crisis inside a financial crisis; 2023 was a fiscal crisis and nothing else. August 2011 contains the euro sovereign debt crisis, which is why VIX reached 48 and the funding rank averaged 0.711. In 2023 VIX never exceeded 20.1, the SOFR–IORB spread varied by four basis points across five months, and funding markets simply did not notice.

That difference is the whole of what follows.

13.5.3 κ_R: what it would take, and why two episodes cannot do it

Rule 3 (Section 12.3) forbids scoring the κ_{sys} decomposition, so κ_R is described here and not estimated. The question the author put is whether that is a temporary or a permanent condition. The answer is closer to permanent, and for an economic rather than a statistical reason.

What estimation would require. Three things, in order. *(i)* An R_{US} series. None exists; Section 11.2 lists partisan-conflict and fiscal-deadline indices as candidates and neither is in hand. *(ii)* κ_{sys} recovered as a latent state, which Proposition 11.1 says is possible — as a single scalar, not as its seven-way split. *(iii)* Variation in R_{US} that is *not* accompanied by variation in the other six drivers of κ_{sys}. Because κ_R enters as a factor on an exponent, it is identified from an *interaction*: how much the dollar's elasticity to D^*/S^* changes when R_{US} moves. That needs R_{US} to move while σ_{sys}, V_{USD}, M_{shock} and A_{meta} do not.

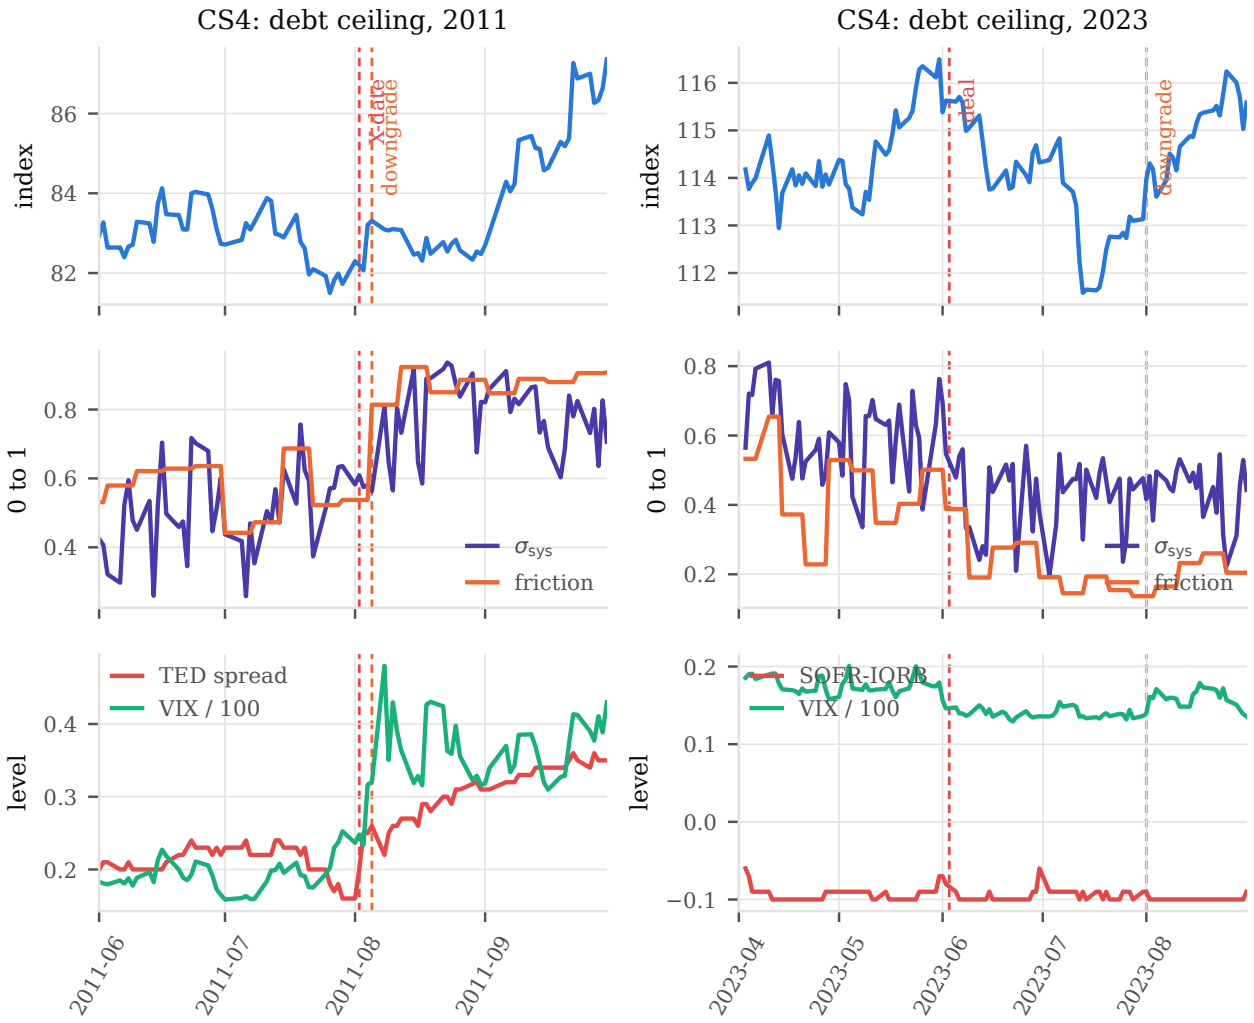


Figure 13.4: CS4 dashboard, both windows on a common layout. Left, 2011: the X-date (2011-08-02) and the S&P downgrade (2011-08-05), with the euro crisis visible in the friction leg and in VIX. Right, 2023: the deal (2023-06-03) and the Fitch downgrade (2023-08-01), with funding markets flat throughout. The contrast between the two bottom panels is Proposition 13.2. Source: `data/processed`; script `cs_dashboards.py`.

Why 2011 and 2023 together cannot supply it. They fail requirement (iii) in opposite and complementary ways.

In **2011**, R_{US} moved and so did everything else. The euro crisis put σ_{sys} at 0.937 and VIX at 48 in the same fortnight as the downgrade. $a_\sigma \sigma_{\text{sys}}$ and $\kappa_R R_{\text{US}}$ are not separable in that window; any flattening of convexity attributed to American fiscal risk could equally be European financial stress.

In **2023**, R_{US} moved close to alone — and the dollar moved +1.23 per cent while funding markets registered nothing. There is no measurable change in the elasticity to attribute to anything.

Proposition 13.2 (The κ_R identification trap)**.** *κ_R is identified only from episodes in which R_{US} moves without the other drivers of κ_{sys}. Episodes in which American fiscal risk moves without moving global financial stress are, by construction, episodes in which markets judged the risk not to matter — and therefore episodes with no measurable change in the dollar's elasticity to attribute. The cleaner the identification, the weaker the signal.* ***This is not a shortage of data that more episodes would fix; it is a property of what the episodes are.***

The honest verdict is therefore not "unidentified for now" but **unidentified in practice, for a reason that does not expire**. A sceptic should treat the $-\kappa_R R_{\text{US}}$ term as a formalisation of an argument rather than as an estimable parameter, and Chapter 11's table now says so.

13.5.4 Scorecard

Table 13.11: CS4 scorecard, both windows.

Prediction	Verdict	Basis
$R_{\text{US}} \uparrow\uparrow\uparrow$	untestable	no proxy (Rule 2 and Section 11.2)
$M_{\text{trust}} \downarrow\downarrow$	untestable	latent, no measure
$\sigma_{\text{sys}} \uparrow$	hit (2011) / partial (2023)	0.937 against 0.810
κ_{sys} flattened	**untestable, permanently**	Proposition 13.2
DXY rises	hit	+5.19% and +1.23%; the dollar rose during a US fiscal crisis, which is the model's counter-intuitive prediction and it holds twice
$s_{\text{US,sys}} \downarrow$	untestable	latent

One and a half hits, four untestable of six — the worst ratio in the book, and the reason is that this episode's mechanism lives almost entirely in latent variables.

The one thing it does establish. In both windows **the dollar rose during a crisis of American fiscal governance**, including through a downgrade of American debt in each. That is genuinely counter-intuitive and it is the model's prediction: D_{prec} and the haven term dominate R_{US}'s effect on Φ_{sys} at these horizons. Chapter 17 should score it, because it is the kind of prediction that would have been easy to get wrong.

13.5.5 E10b term by term, and parameter implications

Mechanical block. In 2011 the bill yield touches zero, so D_{prec} fired — but the euro crisis is the more plausible cause and nothing here separates the two. In 2023 no mechanical channel moves measurably.

Safe-haven block. 2011 yes, 2023 barely. The two windows differ in the haven block by almost everything and in R_{US} by little, which is the wrong way round for this episode's purpose.

Parameter implications. One, and it is negative: Proposition 13.2 means this episode contributes nothing to any parameter estimate. It is included because a model that names a term should be shown the episode that term was written for, and shown to be unable to use it.

13.6 Remark 13.1 Tested Where It Was Not Built

Remark 13.1 proposed a frequency split — daily variation tracking equity volatility, episode-scale level tracking funding stress — and flagged that it had been constructed to explain a disagreement and had predicted nothing. CS4's two windows are the test, and they were not used to build it.

Table 13.12: Remark 13.1 out of sample. Funding measure is TED for 2011 and SOFR–IORB for 2023.

Episode	Result	Predicted	Verdict
2011 daily	ΔVIX +0.401 (t 3.97); ΔTED +0.122 (t 1.11)	VIX	**consistent**
2011 level	DXY peak 8 days from the funding peak, 53 from the VIX peak	funding	**consistent**
2023 daily	ΔVIX +0.249 (t 2.60); ΔSOFR–IORB +0.015 (t 0.15)	VIX	**consistent**
2023 level	—	funding	**uninformative**

Why the fourth cell is uninformative rather than a failure. The SOFR–IORB spread in 2023 has a total range of four basis points and its maximum falls on the *first day of the window*, position 0 of 105. A peak at a window boundary is not a peak, and a series

with no variation cannot date anything. Reporting this as evidence against the conjecture would be as wrong as reporting it as evidence for.

Verdict. The daily half of the conjecture is confirmed out of sample twice, and cleanly: in both 2011 and 2023 the equity-volatility correlation is significant and the funding correlation is not. That is a genuine gain, because those episodes played no part in constructing it. **The level half has been tested once and passed once.**

Remark 13.2 (What the conjecture has and has not earned)**.** It now predicts something it was not built on, twice, which Remark 13.1 said it had not done. But three qualifications hold it below the status of a finding. The two episodes use *different* funding instruments, because TED was discontinued in 2022. The level half rests on one informative observation. And the daily result may be nothing more than the fact that VIX is a market price sampled at the same frequency as the dollar while funding spreads are administratively set and stickier — a measurement artefact that would produce the same pattern with no economics in it. Chapter 20 should test that last possibility directly before the conjecture is used for anything.

Chapter 14

Geo-Economic Shocks: Trade War 2018, Ukraine 2022, the Energy-Embargo Scenarios, and the 2026 Iran War

14.1 Case Study 5: Trade War (2018) and Ukraine War (2022)

14.1.1 What the original claims

Case Study 5 of the original is eleven lines long and joins two episodes into one argument. Restated faithfully, and with nothing dropped:

> *Events:* US tariffs on China; Russia's invasion of Ukraine driving energy-price spikes. *DXY dynamics:* significant appreciation in both periods. *Model analysis:* $\sigma_{\text{sys}} \uparrow$, elevated by geopolitical risk (the X_{pand} proxy) and the GPR index; $D_{\text{trans}} \uparrow\uparrow$, since energy trade invoiced in dollars surged and importers needed more dollars for the same quantity; $G_{\text{pol}} \leftrightarrow$ and $S_{\text{off}} \uparrow\uparrow$, as the United States weaponised the dollar through sanctions. *Short-term:* S_{off} spiked as hundreds of billions in Russian reserves were frozen, shrinking effective S^* and tightening dollar markets, pushing DXY up. *Long-term:* activated the $\Psi_{\text{str}}(t)$ drivers — heightened T_{frag}, "R_{USD}" and C_{alt} as nations sought alternatives. *Conclusion:* tactical victory — DXY rose via D_{trans} and S_{off} — but strategic loss as $\Psi_{\text{str}}(t)$ erosion began.

That is nine arrows across four layers, and the study offers one number in support of any of them: "hundreds of billions" of frozen reserves. Everything else is asserted. The token R_{USD} appears nowhere else in the original and is resolved as R_{US} under ISSUES #6; it is read here as the fiscal and policy-risk premium of Section 6.3.

Table 14.1: CS5 as the original states it. Nine predictions, two episodes, one number.

Variable	Arrow	The original's reason
σ_{sys}	$\uparrow$	geopolitical risk via X_{pand} and the GPR index
D_{trans}	$\uparrow\uparrow$	dollar-invoiced energy trade surged
G_{pol}	$\leftrightarrow$	coercive capacity unchanged
S_{off}	$\uparrow\uparrow$	Russian central-bank reserves immobilised
S^{*}	$\downarrow$	effective supply shrinks, dollar markets tighten
T_{frag}	$\uparrow$	nations seek alternatives
R_{USD} (R_{US})	$\uparrow$	fiscal and policy-risk cost of coercion
C_{alt}	$\uparrow$	alternative rails and units of account
$\Psi_{\text{str}}(t)$	$\downarrow$	structural erosion begins

14.1.2 Event and window

Both windows were fixed in `registry/case_studies.md` before any series was examined, as Section 12.2 requires, and neither has been adjusted since.

Table 14.2: The two episodes. Event dates from the public record; all series from `data/processed`, sources in `data/SOURCES.md`.

Episode	Window		Anchor dates
Trade war	2018-01-01 2019-12-31	to	2018-03-22 Section 301 memorandum; 2018-07-06 first \$34 bn effective; 2018-09-24 \$200 bn at 10 per cent; 2019-05-10 raised to 25 per cent; 2019-08-01 \$300 bn announced; 2019-12-13 Phase One announced
Ukraine	2022-02-01 2022-12-31	to	2022-02-24 invasion; 2022-02-26 G7 joint statement on Central Bank of Russia assets, with US and EU measures effective 2022-02-28; 2022-03-16 first Fed rate rise of the cycle; 2022-09-26 dollar peak

The pre-window baseline for geopolitical risk is the mean of the Caldara–Iacoviello daily index over 2015-01-01 to 2017-12-31, which is 102.50. It is chosen as the three calendar years before the first window and is used unchanged for both episodes.

14.1.3 The data

14.1.4 The measurement that decides the study

The case study's first arrow is $\sigma_{\text{sys}} \uparrow$ "elevated by geopolitical risk (X_{pand} proxy) and GPR index". That is a measurable claim about a published index, and it is the claim that ties the two episodes into one argument. It holds in one of them.

Proposition 14.1 (The trade war does not enter the model through the channel the case study names)**.** *Over the trade-war window the Caldara–Iacoviello index averaged* 97.80 *in*

Table 14.3: CS5 measured. `DTWEXAFEGS` is the book's dollar series (Section 11.2); σ_{sys} is the canonical three-component construction of Section 5.6; GPRD is the Caldara–Iacoviello daily index.

Series	Trade war	Ukraine	Direction	Proxy for
Dollar, ends	104.33 → 108.73	108.12 → 115.61	both ↑	DXY
Dollar, extremes	101.15 → 111.87	106.99 → 124.85	+10.60 %, +16.69 %	DXY
Brent, ends	66.65 → 67.77	90.24 → 82.82	flat; ↓	D_{trans}
Brent, peak	86.07 (2018-10-04)	133.18 (2022-03-08)	↑; ↑↑	D_{trans}
σ_{sys}	0.398 → 0.641	0.767 → 0.835	both ↑	state
σ_{sys}, $w_X = 0$	0.196 → 0.412	0.594 → 0.749	both ↑	state
GPRD, mean	97.80 (2018), 90.47 (2019)	155.09	↓; ↑	X_{pand}
VIX, peak	37.32 (2018-02-05)	36.45 (2022-03-07)	similar	stress
COFER share	62.90 → 61.32	59.42 → 58.95	both ↓	$\Psi_{\text{str}}(t)$

Table 14.4: The Caldara–Iacoviello daily index, against its own pre-window baseline.

Period	Mean GPRD	Ratio to baseline	Rank of best month
Baseline 2015-01-01–2017-12-31	102.50	1.00	—
Trade war, 2018	97.80	**0.95**	2018-05, 99th of 500
Trade war, 2019	90.47	**0.88**	—
Ukraine, 2022-02 to 2022-12	155.09	1.51	—
Ukraine, March 2022 only	306.99	**3.00**	**7th of 500**

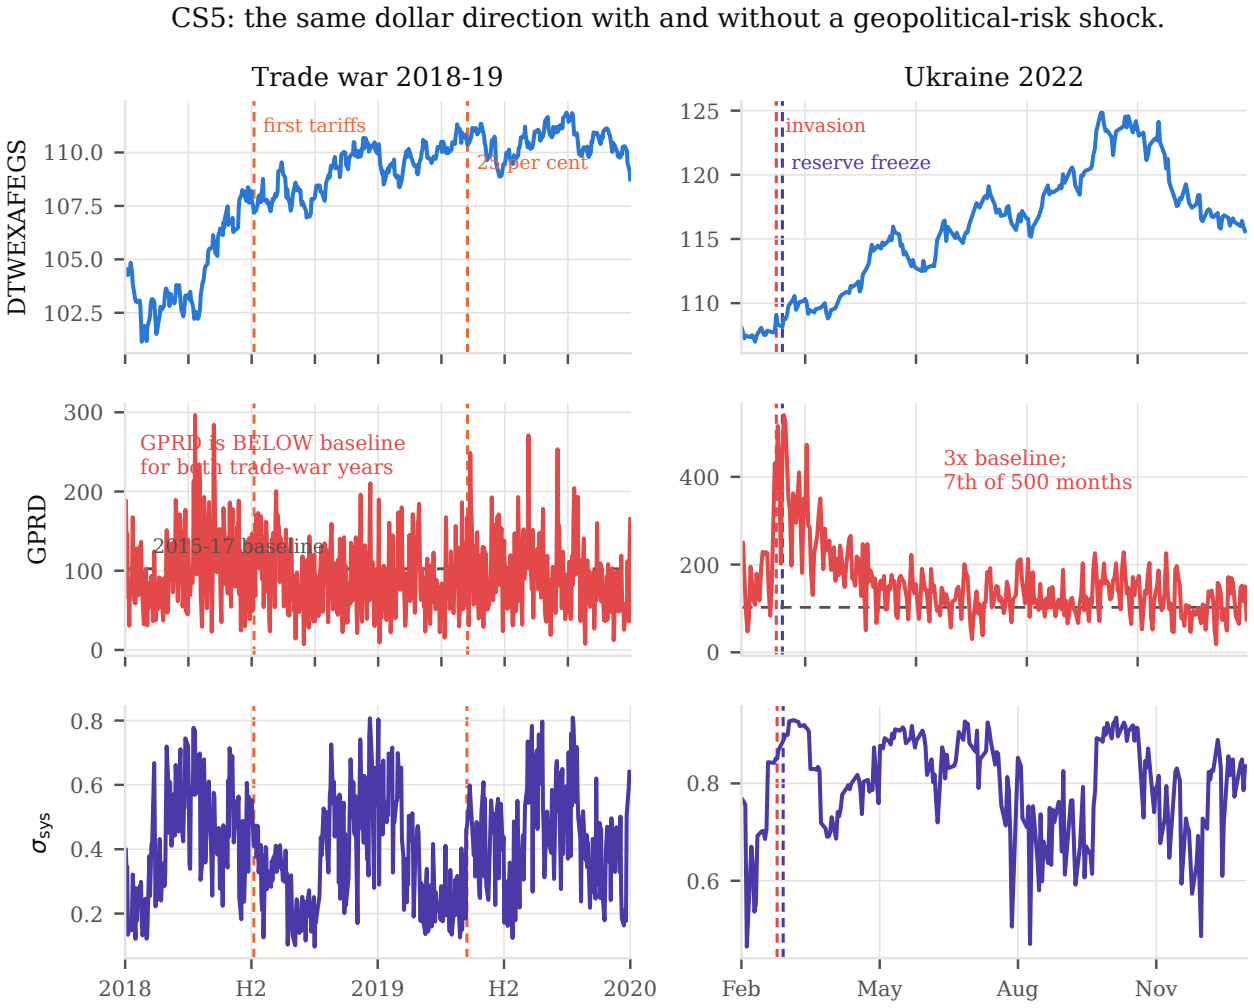


Figure 14.1: The two episodes on the same rows. The dollar rises in both. The channel the case study names as the cause — geopolitical risk entering σ_{sys} through X_{pand} — triples in 2022 and sits *below* its own 2015–17 baseline for the whole of 2018 and 2019. Script: `figures/src/cs5_dashboard.py`.

2018 and 90.47 *in 2019, against a* 2015–17 *baseline of* 102.50*: geopolitical risk was* below *its own recent norm in both years, and the highest 2018 month ranks 99th of the 500 months in the series. Over the Ukraine window the same index averaged* 155.09*, and March 2022 ranks 7th of 500. The dollar rose in both episodes. Therefore the* X_{pand} *channel cannot be the common cause of the two moves that CS5 says it is, and at most one of the two episodes evidences it.*

Proof. The index values are computed in `data/cs5_cs9.py` from `data/raw/gpr_daily.csv` over the windows of Table 14.2, which were fixed in advance. The dollar changes are +4.22 and +6.92 per cent end to end. A common cause must move in the same direction in both episodes; this one moves down in the first and up sharply in the second. □

Two qualifications are owed, and neither rescues the arrow. First, X_{pand} is not GPRD alone: Section 5.6 builds it from a percentile rank of the same index, so the ranks inherit the levels and the finding is unchanged. Second, the original may have meant *economic-policy* uncertainty rather than geopolitical risk, which did rise in 2018–19. But the original names "the GPR index", and Section 5.6.5 adopted GPRD as the measured leg of X_{pand} on that authority. Substituting a different index after seeing that the named one fails is precisely what Section 12.2 forbids. The alternative is recorded in the research agenda instead (Section 17.8).

Remark 14.1 (What the state variable does instead)**.** σ_{sys} nevertheless rose across the trade-war window, from 0.398 to 0.641, and it did so in the $w_X = 0$ construction too (0.196 → 0.412). The state variable saw the trade war through equity volatility and the funding leg, not through geopolitical risk. That is consistent with the finding of Table 11.6 — the dollar has no measurable relationship with the geopolitical leg over 246 months, a slope of −0.08 with $t = -0.19$ — and it means the trade war is evidence for σ_{sys} as a composite and against the specific attribution CS5 makes.

14.1.5 The reserve freeze, at the resolution of its own event

The short-term mechanism CS5 asserts is precise enough to test on its own dates: S_{off} spiked, effective S^* shrank, dollar markets tightened, DXY rose. The freeze is datable: announced 2022-02-26, effective 2022-02-28.

The whole of the freeze's own month is +0.79 per cent. The 14.46 per cent accrued over the seven months that followed, during which GPRD *fell* from 371 to 156. A reader who knew only the dollar path and the GPRD path would not date the shock to February.

This is not an argument that the freeze did nothing. It is an argument that the episode cannot distinguish the freeze from what else was happening, and Section 14.1.6 names what else.

14.1.6 The rival, and it is not exotic

The Federal Reserve began the fastest tightening cycle in forty years on 2022-03-16, sixteen days after the freeze. The effective funds rate went from 0.08 per cent on the day of the

Table 14.5: The 2022 rally, decomposed by date rather than by narrative.

Sub-window	Dollar	Change	GPRD
2022-02-01 to 02-23, pre-invasion	108.12 → 107.80	−0.30 %	191 → 431
2022-02-24 to 02-28, invasion to freeze	109.06 → 108.14	−0.84 %	371 → 453
2022-02-28 to 03-31, freeze to month end	108.14 → 109.00	+**0.79** %	453 → 255
2022-02-24 to 09-26, invasion to peak	109.06 → 124.83	+**14.46** %	371 → 156

Table 14.6: Two candidate drivers of the 2022 rally, tested on daily changes over the same window (2022-02-24 to 2022-09-27).

Daily change in the dollar against	n	r	t	Verdict
daily change in the US 10-year yield	148	+0.172	+2.11	significant
daily change in GPRD	149	+0.113	+1.38	not significant

invasion to 3.08 on the day of the dollar peak and 4.33 by year end. The dollar rally and the tightening cycle are the same seven months.

The rate channel is significant at conventional levels and the geopolitical channel is not. Both are weak in absolute terms, as daily currency regressions always are, and neither r is large. But the ordering is the wrong way round for CS5, and it is the second time in this book that a rate differential has outperformed a channel the model prefers — Table 11.8 found the US–Bund spread to be the mechanical channel whose coefficient rises most sharply with σ_{sys}.

Remark 14.2 (The model has the rival, and does not use it here)**.** This is not a variable outside the architecture. The interest-rate differential enters through s_{US} in Equation (10.2) and through D_{cap} in the demand block. What CS5 does is attribute a move to S_{off} and D_{trans} without checking s_{US}, in an episode where s_{US} moved more than it has in four decades. The defect is not a missing variable; it is an unexamined one.

14.1.7 The E10b walk

Using the respecified estimates of Section 11.5.6 — $\lambda_1\Phi_{\text{sys}} = +0.429$, $\lambda_2\Phi_{\text{sys}} = +1.505$ — the safe-haven block contributes as follows.

The Ukraine residual, at 97.8 per cent, is the largest in the book — larger than 2008's 64 and 2020's 76 (Table 11.7). The reason is mechanical and worth naming: the episode begins with σ_{sys} already at 0.767, so the convex term has little room left, and Λ-shaped convexity delivers its largest increments in the middle of the range, not at the top. **An episode that**

Table 14.7: The safe-haven block against the observed move. Same estimates and same arithmetic as Table 11.7.

Episode	σ_{sys}		Haven block	Observed	Residual
Trade war 2018–19	0.398 0.641	→	+0.48 %	+4.22 %	+3.74 % (88.6 %)
Ukraine, to the peak	0.767 0.885	→	+0.34 %	+15.45 %	+15.11 % (97.8 %)

starts in a stressed state cannot be explained by a model whose crisis mechanism is convexity in the state variable. That is a structural property of E9, not a fact about 2022, and it applies to any shock arriving on top of an existing one.

14.1.8 Scorecard

Twenty-two cells: five hits, three partials, two misses, eleven untestable and one recorded. (An earlier draft of this paragraph miscounted, giving four hits and twelve untestable; `registry/scorecard.csv` is the authority and the table above is its source.) The untestable share is the highest of any study in the book, and it is high for a reason worth stating rather than burying: **seven of the nine variables CS5 names have no measure anywhere in this book**, so the case study is nine-tenths assertion by construction, not by oversight in the writing.

14.1.9 Parameter implications

On X_{pand} and w_X. Proposition 14.1 is the second episode after CS12 in which the weight on the geopolitical leg has an observable stake, and it points the opposite way from CS12. In CS12 the geopolitical leg fired and the dollar did not move (Section 14.3.4); here the geopolitical leg did not fire and the dollar rose. Both are evidence that X_{pand} and the dollar are close to independent, which is what Table 11.6 measured directly. **The three results are consistent and they point at w_X small rather than w_X large.** The book still does not choose it (Section 5.6.5), because a decision rule was written for that purpose and it requires an estimate, not an accumulation of episodes.

On λ_2 and the convexity claim. Section 14.1.7's structural point is the sharper one. Convexity in σ_{sys} delivers its largest increments in the middle of the range. Every crisis this book scores except the trade war begins with σ_{sys} already high. So the mechanism the architecture designates as the crisis mechanism is systematically weakest exactly where crises occur, which is a specification problem and not an estimation problem. Chapter 17 carries it.

On s_{US}. Remark 14.2 is a finding about method as much as about 2022. The model contains the rate differential and CS5 does not consult it. Two of the four episodes in Chapter 13 would repay the same check.

Table 14.8: Case Study 5 scored, both episodes. Verdicts follow Section 12.5; percentages in Chapter 17 are over all rows including untestable ones.

Variable	Arrow	Trade war	Ukraine	Note
σ_{sys}	↑	**hit**	**hit**	$0.40 \to 0.64$; $0.77 \to 0.84$
X_{pand} via GPR	↑	**miss**	**hit**	$0.95\times$ and $0.88\times$ baseline; $3.00\times$ (Proposition 14.1)
D_{trans}	↑↑	**miss**	**partial**	Brent flat over the window; peaked +48 % then gave it back
G_{pol}	↔	**untestable**	**untestable**	no aggregator; Rule 4 forbids scoring it by target weight
S_{off}	↑↑	**untestable**	**recorded**	the freeze is a fact; its dollar effect is +0.79 % in its own month
S^*	↓	**untestable**	**untestable**	S^* has no measure (Proposition 11.1 block)
T_{frag}	↑	**untestable**	**untestable**	no free machine-readable series (Section 9.4)
R_{US}	↑	**untestable**	**untestable**	unmeasured latent
C_{alt}	↑	**untestable**	**partial**	CIPS and RMB shares unavailable; stablecoins are not this
$\Psi_{\text{str}}(t)$	↓	**partial**	**untestable**	trade war: falls on both aggregates but the fall predates the window. Ukraine: COFER was the wrong base (ISSUES #39) and the gold-inclusive series is annual, so an eleven-month window cannot be resolved (Section 9.7.1)
DXY	↑	**hit**	**hit**	+4.22 %; +6.92 % end to end

On S_{off} and Rule 5. The freeze is the largest single exercise of S_{off} in the sample and its own month is worth +0.79 per cent. The case study reads it as decisive. Section 12.3's Rule 5 already forbids scoring the S_{off} ratchet as confirmed, since the specification imposes monotonicity; Section 14.2.7 takes that up where it bites hardest.

14.2 Case Study 9: Energy Embargo and Secondary Sanctions — Four Scenarios

40–43

14.2.1 What the original claims

Case Study 9 is the longest and most fully specified hypothetical in the original. It posits a full US embargo on Russian oil accompanied by secondary sanctions compelling European compliance, and analyses it as a combination of an energy shock (X_{pand}) and the weaponisation of the payment system (G_{pol}, S_{off}). It asserts three phases, decomposes the shock across four blocks, and then branches into four scenarios of differing severity "to identify model tipping points".

The block decomposition, restated from p. 41 with the original's arrows kept:

Table 14.9: CS9's impact decomposition, as the original states it for strict enforcement (original p. 41). Nineteen arrows.

Block	Variables and direction
State / friction	$X_{\text{pand}} \uparrow\uparrow$ (exogenous energy shock) $\Rightarrow \sigma_{\text{sys}} \uparrow\uparrow$ (VIX/MOVE surge) $\Rightarrow F_{\text{liq}} \uparrow\uparrow$ (liquidity friction from a commodity-trade freeze)
Structural / trust	$G_{\text{pol}} \uparrow$ (coercive power succeeds); $T_{\text{frag}} \uparrow\uparrow$; $R_{\text{US}} \uparrow$; $\Phi_{\text{sys}} \leftrightarrow$ / $\downarrow$; $\Psi_{\text{str}}(t) \downarrow$
Demand	$D_{\text{trans}} \uparrow\uparrow\uparrow$ (oil price surge, EU replacement in dollars); $D_{\text{prec}} \uparrow\uparrow$ (flight to cash and bills); $D_{\text{debt}} \uparrow$ (dollar short squeeze in emerging Europe and EM); $D_{\text{cap}} \uparrow$ (flows from the EU to the US)
Supply	$S_{\text{off}} \uparrow\uparrow$ (assets frozen); $S_{\text{coll}} \downarrow$ (bill-collateral shortage); $S_{\text{bank}} \downarrow$ (balance-sheet contraction); $S_{\text{swap}} \uparrow$ (Fed liquidity); $S_{\text{ext}} \uparrow$ (US energy exports expand later, easing the dollar)

The three phases are: Phase 1 tactical, an immediate non-linear spike; Phase 2 rebalancing, a cooling as S_{swap} and S_{ext} activate; Phase 3 strategic, long-term erosion via $T_{\text{frag}} \uparrow\uparrow$ and $\Psi_{\text{str}}(t) \downarrow$. The four scenarios are Scenario 1, full embargo with strict secondary sanctions and third parties penalised ($\uparrow\uparrow\uparrow$, "a perfect USD squeeze"); Scenario 2, embargo plus a price cap with exemptions ("moderate increase", then "near baseline", erosion "negligible"); Scenario 3, weak enforcement ("brief uptick, rapid mean reversion", then flat); Scenario 4, embargo plus kinetic disruption ("$\uparrow\uparrow\uparrow$ extreme panic bid", then a V-shape).

14.2.2 The analogue, and why this study can be scored at all

Section 12.2 treats a hypothetical as scoreable when a real episode matches the scenario's own stated conditions and the match can be established from the conditions rather than from the outcome. CS9 clears that bar, and it is the only hypothetical in the book that does so cleanly.

Table 14.10: Branch selection, from the scenarios' stated conditions only. Event facts from the public record, recorded in `data/SOURCES.md`; no dollar series was consulted in constructing this table.

Scenario	Its stated condition	The record, 2022–23
1	Third parties (China, India) penalised	**No.** Neither was sanctioned for purchasing Russian crude
2	Embargo with a price cap and exemptions, loopholes allowing flow at capped prices	**Yes.** EU seaborne crude embargo 2022-12-05; products 2023-02-05; G7 cap at $60 from 2022-12-05, with a shipping-services carve-out below the cap
3	Sanctions nominal, loopholes tolerated, no real ban	**Partly.** The EU ban was real and enforced on its own imports; evasion concerned third-country flows
4	Kinetic disruption of sea lanes, insurance or pipelines	**Before the window.** Nord Stream was destroyed 2022-09-26, ten weeks before the embargo

The realised branch is Scenario 2, with Scenario 3 a partial reading. The two are not separated by any threshold the original states, which is a defect taken up in Section 14.2.6 — but it does not obstruct the scoring, because Scenarios 2 and 3 predict the same *sign* for the dollar (a moderate rise; a brief uptick). Any result that contradicts one contradicts both.

The window is 2022-06-01 to 2023-06-30, from `registry/case_studies.md` and fixed in advance. It begins six months before the embargo so the pre-event baseline is inside the window.

14.2.3 The data

Every peak in the window falls *before* the embargo took effect. The state variable, the oil price, the geopolitical index and equity volatility all top out in June to October 2022 and decline into and through the event.

14.2.4 Scenario 2 scored, prediction by prediction

Scenario 2 makes four statements. Each is tested on the embargo's own dates.

P1: "X_{pand} and σ_{sys} modest (↑)". Comparing the three months before the embargo (2022-09-01 to 2022-12-04) with the three after (2022-12-05 to 2023-02-28), σ_{sys} fell from a

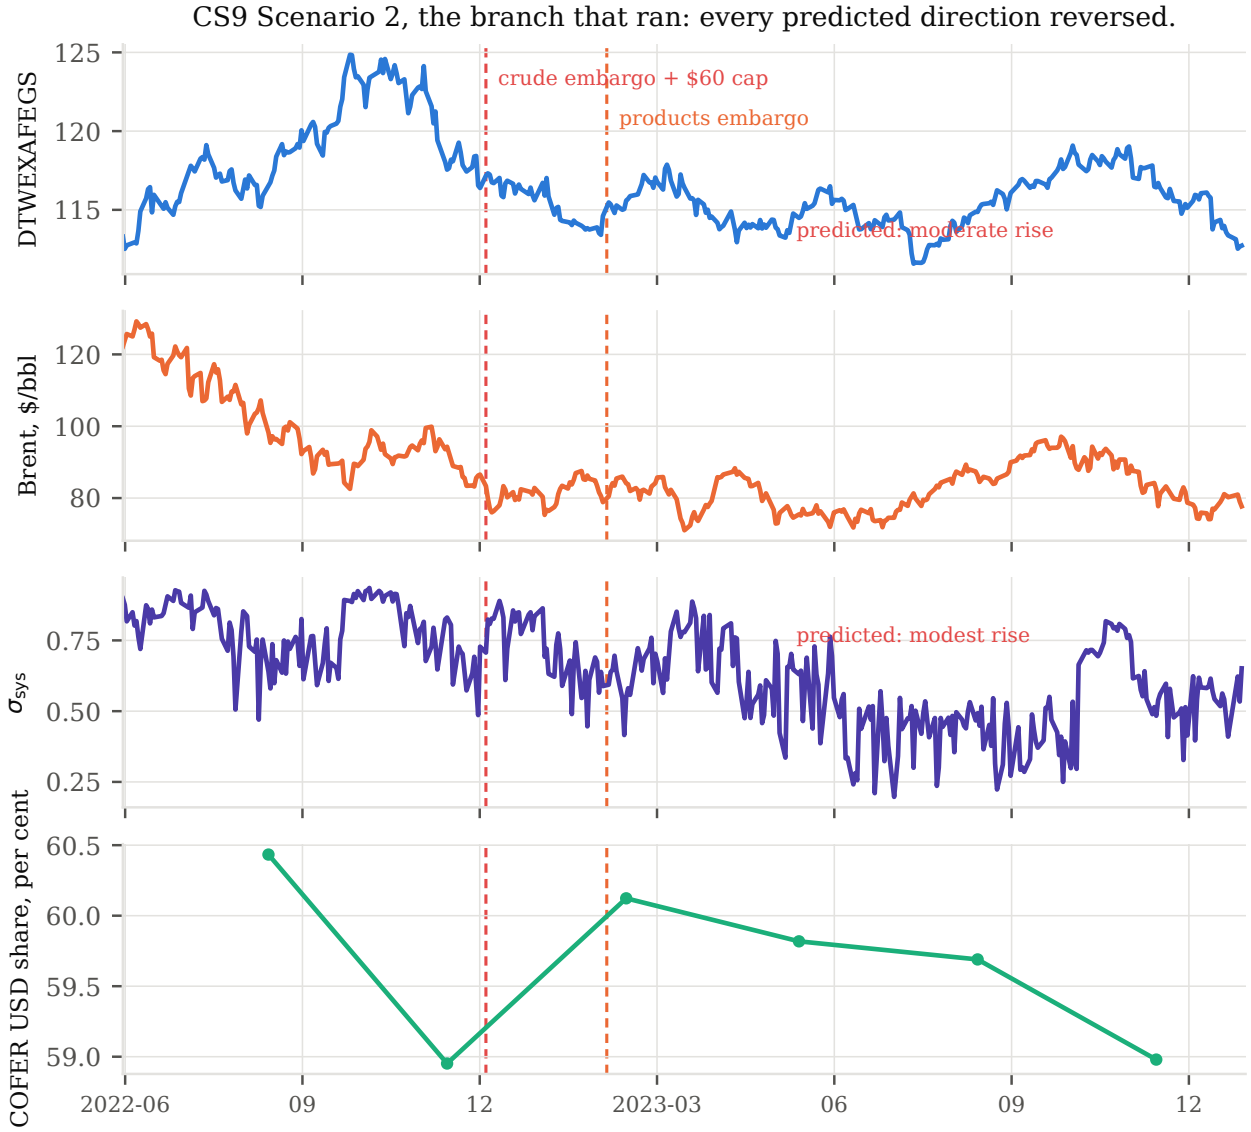


Figure 14.2: The branch the world ran, on its own dates. Scenario 2 predicts a moderate dollar rise, a modest rise in σ_{sys} and X_{pand}, a non-explosive rise in D_{trans}, and negligible erosion of $\Psi_{\text{str}}(t)$. Script: `figures/src/cs9_dashboard.py`.

Table 14.11: CS9's window measured. All series from `data/processed`.

Series	Ends of window	Change	Peak	Reads on
Dollar	113.30 → 114.33	+0.90 %	124.85 on 2022-09-26	*DXY*
Brent	122.20 → 74.51	−39.03 %	129.20 on 2022-06-08	D_{trans}
σ_{sys}	0.898 → 0.378	−57.97 %	0.935 on 2022-10-06	state
GPRD	168.3 → 102.7	−39.01 %	228.3 on 2022-10-11	X_{pand}
VIX	25.69 → 13.59	−47.10 %	34.02 on 2022-06-13	stress

mean of 0.781 to 0.695 and GPRD fell from 125.3 to 110.1. Both moved down. **Wrong sign.**

P2: "D_{trans} rises but not explosively". Brent fell 3.56 per cent across the embargo quarter, and its mean fell from $101.51 over June–November 2022 to $79.76 over December 2022–June 2023. The dollar value of the oil trade the world had to finance went down. **Wrong sign.**

P3: "DXY: moderate increase" short term, "near baseline" after. From the embargo date:

Table 14.12: The dollar from the embargo date, at four horizons.

Horizon	Dollar	Change	Verdict against "moderate increase"
one month	117.17 → 117.02	−0.12 %	wrong sign, negligible
three months	117.17 → 116.77	−0.34 %	wrong sign
to window end	117.17 → 114.33	−2.42 %	wrong sign
one year	117.17 → 115.95	−1.04 %	wrong sign

Wrong sign at every horizon. The one honest reading available to the case study is that −0.12 per cent over a month is "near baseline", which is Scenario 2's medium-term prediction rather than its short-term one; the short-term prediction of a moderate *increase* has no support at any horizon tested.

P4: "structural erosion ($\Psi_{\text{str}}(t)$) negligible". The dollar share of allocated reserves went 58.95 per cent in 2022-Q4 to 60.12 in 2023-Q1 to 58.98 in 2023-Q4 — a rise then a return, net change over the five quarters of +0.03 percentage points. **Hit**, and it is the only

one. It is also the weakest kind of hit, since "negligible" is satisfied by any outcome inside the noise of a quarterly series that moves by ±1 point routinely.

14.2.5 The supply legs, where they can be measured

Two of the five supply arrows have published series.

Table 14.13: S_{swap} and the dollar credit aggregate across the embargo. `SWPT` is Federal Reserve central-bank liquidity swaps outstanding; the BIS series is dollar credit to non-bank borrowers outside the United States.

Period	S_{swap}, mean \$m	BIS credit, \$bn	Predicted
Pre-embargo, 2022-06 to 11	966	13,298 (Q2)	—
Embargo phase, 2022-12 to 2023-02	**322**	12,821 (Q4)	$S_{\text{swap}} \uparrow$
Rebalancing, 2023-03 to 06	424	12,713 (Q1)	$S_{\text{swap}} \uparrow$, $D_{\text{debt}} \uparrow$

S_{swap} fell by two thirds into the embargo, against a predicted rise. The dollar credit aggregate contracted through the whole period, from \$13.41 trillion in 2022-Q1 to \$12.71 trillion in 2023-Q1, against a predicted $D_{\text{debt}} \uparrow$. Both are **wrong sign**, and the second is the same variable, at the same quarterly frequency, whose timing Section 11.5.6 found unable to carry a crisis mechanism.

14.2.6 Whether a four-scenario design can be wrong

CS9's four branches span the outcome space. Scenario 1 predicts a large rise, Scenario 2 a moderate rise, Scenario 3 a brief uptick and mean reversion, Scenario 4 a spike then a collapse. Any dollar path whatever is some branch's prediction. A design of that shape is falsifiable only if the branch is selected from conditions observable *before* the outcome. Table 14.10 shows CS9 clears that test where CS8 does not, and the distinction is worth stating precisely because the two look alike.

Proposition 14.2 (CS9's realised branch fails on its own terms)**.** *Of Scenario 2's four stated predictions, three have the wrong sign in the realised episode — σ_{sys} and X_{pand} fell, the oil price fell, and the dollar fell at one month, three months, seven months and one year — and the fourth ("negligible" structural erosion) is satisfied by any outcome inside quarterly noise. The two measurable supply arrows, $S_{\text{swap}} \uparrow$ and $D_{\text{debt}} \uparrow$, also have the wrong sign. This is a refutation of a stated branch, not an ambiguity: the branch was selected in Table 14.10 from enforcement conditions alone.*

Remark 14.3 (One residual defect the scoring does not remove)**.** Scenarios 2 and 3 are not separated by any threshold the original states. Here that costs nothing, because they agree

Table 14.14: Two multi-scenario studies, and why only one of them can fail.

	CS8 (Section 15.4)	CS9 (here)
Branch selector	severity of a cyber breach, defined by its consequences	severity of enforcement, defined by policy acts
Observable before the dollar?	**No** — "systemic" is defined partly by the market reaction it is meant to predict	**Yes** — a price cap either exists or does not
Verdict	unfalsifiable (Proposition 15.2)	**falsifiable, and falsified**

on sign and both are contradicted. It would cost everything in an episode where the dollar rose modestly: the study could then claim either branch after the fact. Chapter 17 records this as a specification gap in the case-study design rather than a measurement gap, in the same class as G_{pol}'s missing aggregator.

14.2.7 What the episode may not be used to confirm

CS9 encodes coercion as monotone: S_{off} up, T_{frag} up, $\Psi_{\text{str}}(t)$ down, with no return path. The 2022–23 record is consistent with all three. It confirms none of them.

The reason is Section 12.3's Rule 5, and it is a point about the sign graph rather than about the data. The functions themselves are reversible — Section 10.1.3 proves E8 is a bijection in f_{hard}, so nothing in the mathematics forbids $\Psi_{\text{str}}(t)$ from rising. What forbids it is the arrow set: no path in the dependency graph carries a fall in S_{off} into a rise in $\Psi_{\text{str}}(t)$, so a monotone path is what the specification will produce whatever the world does. An episode cannot confirm a property the specification imposes.

That this is a real constraint and not a formality is established by an episode in this same book. Section 16.1.7 records an exercise of G_{pol} that *lowered* S_{off} and *raised* D_{trans} four weeks later — the OFAC reversal on Venezuelan oil of 2026-01-29. Coercion was withdrawn and the arrows the model has no way to draw are exactly the ones the episode drew. It is the same defect as the sterling gap of Section 9.5: one defect, three dates.

14.2.8 Scorecard

Nineteen rows: **no hits**, six misses, eleven untestable, two imposed. Its single hit was withdrawn in the ISSUES #39 correction, because it had been scored on a reserve aggregate that excludes gold. **CS9 therefore scores zero of six testable predictions**, joining CS8 and CS10 — and unlike either of those, it is the most carefully specified study in the original. Those two facts belong together: the study can fail because it said enough to fail.

Table 14.15: Case Study 9 scored on its realised branch. Nineteen block arrows plus Scenario 2's four predictions; arrows with no measure are untestable by Section 12.2, not omitted.

Variable	Arrow	Verdict	Note
X_{pand}	↑↑	**miss**	GPRD mean 125.3 → 110.1 across the event
σ_{sys}	↑↑	**miss**	0.781 → 0.695; peak precedes the embargo by two months
F_{liq}	↑↑	**untestable**	no separable friction measure
G_{pol}	↑	**untestable**	Rule 4; no aggregator over the three capacities
T_{frag}	↑↑	**untestable**	no free machine-readable series
R_{US}	↑	**untestable**	unmeasured latent
Φ_{sys}	↔/↓	**untestable**	Φ_{sys} enters only as a product with λ
$\Psi_{\text{str}}(t)$	↓	**imposed**	Rule 5; may not be scored (Section 14.2.7)
D_{trans}	↑↑↑	**miss**	Brent mean 101.51 → 79.76
D_{prec}	↑↑	**untestable**	Rule 1 forbids evidencing it from VIX
D_{debt}	↑	**miss**	BIS credit 13,408 → 12,713 $bn
D_{cap}	↑	**untestable**	no window-frequency flow series
S_{off}	↑↑	**imposed**	Rule 5
S_{coll}	↓	**untestable**	the measured leg is the reverse-repo facility, which is not this
S_{bank}	↓	**untestable**	unmeasured latent
S_{swap}	↑	**miss**	mean $966m → $322m
S_{ext}	↑	**untestable**	EIA export series not obtained (Tier 3)
Sc.2 dollar	↑	**miss**	−0.12, −0.34, −2.42, −1.04 per cent
Sc.2 $\Psi_{\text{str}}(t)$	negligible	**untestable**	scored on COFER, which excludes gold (ISSUES #39). On the gold-inclusive aggregate the share falls 0.72pp over the window, and "negligible" has no stated threshold

14.2.9 Parameter implications, and the overlap with CS5

On $X_{\text{pand}} \Rightarrow \sigma_{\text{sys}} \Rightarrow DXY$. This is the third independent episode in which the chain the model draws from an energy or geopolitical shock to the dollar does not run: the trade war (Proposition 14.1), the 2026 Iran war (Section 14.3.4), and here. In this one the shock was a deliberate act of policy aimed at the world's largest oil exporter, and the dollar did not move. Table 11.6 gives the reason in a single number — the geopolitical leg's slope on the dollar is -0.08 with $t = -0.19$ over 246 months — and Chapter 17 carries the three episodes as a set.

On the phase structure. All three of CS9's phases are contradicted on their own dates: no tactical spike (-0.12 per cent), no moderation from a peak that never formed, and a Phase 3 that the specification cannot fail. The phase structure is the part of CS9 most likely to be salvageable, because it describes a plausible sequence and was tested here against a branch — Scenario 2 — whose shock was mild. Chapter 17 records it as untested rather than refuted for the strict-enforcement branch, which has no analogue.

On double counting. CS5's Ukraine window (2022-02-01 to 2022-12-31) and CS9's (2022-06-01 to 2023-06-30) overlap by seven months. The overlap carries +2.04 per cent of dollar move; CS5's exclusive period carries +4.04 and CS9's carries -2.05. The two studies score different variables — CS5 the freeze and CS9 the embargo — so no arrow is scored twice, but Chapter 17's totals treat them as one episode for the purpose of counting *independent* tests, and Section 17.2 says so.

On S_{coll}. CS9 predicts a Treasury-collateral shortage. The book's measured S_{coll} leg is the overnight reverse-repurchase facility, which measures cash placed *at* the Federal Reserve rather than collateral scarcity in the market, and Section 11.5.2 showed how badly that series behaves when its blanks are misread. The arrow is untestable here for a reason worth carrying into Chapter 17's agenda: the variable has a proxy and the proxy measures something else.

14.3 Case Study 12: The Iran War and the Strait of Hormuz, 2026

14.3.1 Why this is a new case study and not an extension of CS9

The decision is deliberate and it rests on which capacity the episode tests.

Section 14.2's Case Study 9 tests the *financial* rail of coercion: S_{off} rises, T_{frag} rises, $\Psi_{\text{str}}(t)$ erodes. That is the first of the three capacities in the G_{pol} definition — enforcing the rules of the system — exercised through payment and reserve architecture.

The 2026 Iran war tests the *second* capacity, and nothing else in this book does. G_{pol}'s second heading is **securing maritime trade routes**, which Section 6.2 calls "the physical precondition for D_{trans}: dollar-invoiced commodity trade requires that the commodity arrives". A closure of the Strait of Hormuz is that precondition failing, in the one chokepoint where it matters most. Folding this into CS9 would merge two different capacities, two different channels and two different instruments into one narrative, and the whole point of Assumption 6.1 — adopted four weeks ago in response to CS10's failure — is that G_{pol} must be scored capacity by capacity.

Three further reasons make it a separate study rather than a subsection.

It is the **first live test of** X_{pand} with a measured index. Chapter 5 added geopolitical risk to σ_{sys} and Section 5.6.5 left the weight w_X undetermined with a decision rule attached. This episode is the first in the sample where that weight has an observable stake, and Section 14.3.4 shows the two constructions disagree about it.

It is the natural **companion to the CS10 audit**. CS10 failed because G_{pol} was scored by the weight of the adversary rather than by demonstrated capability (Section 16.1.6). Here is an episode in which capability on the route-securing heading was tested directly and, for seven months and counting, not restored.

And it **contaminates every other 2026 window**, which is a methodological fact that belongs in a study of its own rather than in a footnote to someone else's.

14.3.2 Event and window

Table 14.16: The episode. Sources recorded in `data/SOURCES.md` under "2026 Iran war".

2026-02-28	Israel and the United States begin strikes on Iran, stated aims regime change and the nuclear and ballistic-missile programme. Iran's supreme leader is killed.
early March	Iran retaliates against US embassies, installations and oil infrastructure, and closes the Strait of Hormuz.
through 2026	Roughly 20 per cent of world petroleum and 20 per cent of LNG normally transit the strait. Reported traffic falls by about 95 per cent for crude and 99 per cent for LNG. A conditional ceasefire holds while talks continue; shipping remains very low; the UAE's state oil company expects full flows not to resume before 2027.
2026-09	The conflict is unresolved.

The window for scoring is **2026-02-27 to 2026-09-01**, the end of the data. It is not closed, and the study says so rather than pretending the episode is over.

14.3.3 What the price data show

Brent closed February at \$71.32 and March at \$126.69, a rise of **77.6 per cent** month-end to month-end, or 45.5 per cent on monthly averages. It peaked at **\$138.21 on 7 April** and stood at \$96.02 on 1 September, still 35 per cent above the pre-war level after six months.

A correction to a claim in circulation. It has been said that this was the largest monthly gain since the Brent series began. On the full series from 1987-05-20 (`DCOILBRENTEU`, 9,967 observations, fetched for this chapter) it is not. Month-end to month-end it ranks **second of 472 months**, behind May 2020's 88.6 per cent. On monthly averages it ranks **third**, behind May 2020 and August 1990's 58.2 per cent.

The two larger episodes are worth naming because they frame this one. May 2020 is a rebound from the pandemic collapse, not a supply interruption. August 1990 is the Iraqi

invasion of Kuwait. **So March 2026 is the largest supply-driven oil shock in the Brent record except for the invasion of Kuwait, and the comparison to make is with 1990 rather than with 2020.**

14.3.4 The model reading, channel by channel

X_{pand}: a top-five geopolitical shock, and the state variable sees it. The Caldara–Iacoviello daily index averages **325.3 in March 2026**, the **fifth highest month in the 1985–2026 series**. Above it are only October and September 2001, January 1991 and March 2003. It is above March 2022, the invasion of Ukraine, at 307.0. The index remains elevated through September at 201.3, roughly double its February level.

This is where the w_X question of Section 5.6.5 acquires a stake. The canonical three-component σ_{sys} peaks at **0.847** on 30 March; the two-component index with $w_X = 0$ peaks at **0.721** on the same day. The gap is 0.13, and it appears *before* the strikes: on 27 February the canonical index reads 0.607 against 0.410, because `GPRD` was already rising while market volatility was not.

Remark 14.4 (The first episode where w_X matters observably)**.** A financial-stress-only state variable classifies March 2026 as materially calmer than a state variable that reads geopolitical risk, and it registers the episode later. Whether that is an improvement depends on the unestimated w_X, which is exactly the point of Proposition 5.2. What the episode establishes is that the choice is not academic: the two constructions disagree by 0.13 of the state variable's range at the peak of the largest supply shock in thirty-six years.

D_{trans} and G_{pol} capacity (ii): the model's own precondition failing. The route-securing capacity is not an abstraction here. With Hormuz traffic down roughly 95 per cent for crude, dollar-invoiced commodity trade in the world's most important chokepoint did not arrive. Under Assumption 6.1 this is scored as a capability outcome, not a narrative: **capacity (ii) was tested and, as of the data cut-off, not restored**.

That gives G_{pol} a component moving *down* in the same year as CS10 recorded two components moving up. The model has no aggregator over the three capacities — G_{pol} is a single latent scalar in $[0, 1]$ — so it cannot say what the net is. Chapter 11 listed G_{pol} among the latents with no measure; this episode shows the cost is not just imprecision but *sign ambiguity* in a year when the capacities moved in opposite directions.

The dollar. `DTWEXAFEGS` is essentially unmoved across the window. That is itself a reading: the largest supply-driven oil shock since 1990, a top-five geopolitical risk month, and σ_{sys} at 0.85, produced no material move in the trade-weighted dollar. Chapter 5's finding — that the dollar's response to stress is positive but modest, and that the amplification claim is undetermined — is consistent with that, and a model predicting a large safe-haven appreciation would not be.

There is a sharper reading available and it is not flattering to σ_{sys} as constructed. Over this window the friction leg p_{STLFSI} averaged **0.228**, against 0.990 in the 2008 window and 0.887 in 2020. **Dollar funding stress never fired.** The state variable reached 0.847 almost entirely through geopolitical risk and equity volatility, and Table 11.6 shows the dollar has

no measurable relationship with the geopolitical leg at all — a slope of -0.08 with $t = -0.19$ over 246 months. So the dollar's non-response here is not evidence against a safe-haven mechanism. It is what a *funding-based* haven predicts when funding markets are calm, and it is the clearest episode in the book separating the two readings. Chapter 17 carries it as one of three independent tests of the haven question, together with the rival hypothesis that reconciles all three and the split that contradicts the rival.

14.3.5 Scorecard

Table 14.17: Case Study 12 scored. Market and index figures from `data/processed`; event facts from the sources recorded in `data/SOURCES.md`.

Channel	Verdict	Note
X_{pand} ↑↑↑	**hit**	GPRD 325.3, fifth highest month since 1985; measured, not asserted
σ_{sys} ↑↑	**hit**	0.427 to 0.847; the regime switch the model requires
D_{trans} disrupted	**hit**	95 per cent fall in crude transits
G_{pol} capacity (ii) ↓	**hit**	the only direct test of the route-securing heading in this book, and the capacity measurably fell
DXY response	**partial**	no material move: consistent with Ch.5's modest elasticity, inconsistent with a strong-haven reading
$\Psi_{\text{str}}(t)$	**untestable**	f_{hard} has no measure (Proposition 9.3)

14.3.6 The confounding this creates

From 28 February 2026 onward, every window in this book overlaps a major energy and geopolitical shock. Section 12.4 states the rule that follows and it binds every 2026 case study including this one: **no 2026 episode may claim a clean identification**. This study is not exempt. Its own σ_{sys} path, oil path and dollar path are jointly determined with the Venezuela episode of Section 16.1 and with whatever else 2026 contained.

14.3.7 One observed ordering, no attribution

OFAC lifted Venezuelan oil sanctions on **29 January 2026**. The strikes on Iran began on **28 February 2026**. The interval is thirty days, and in that interval PDVSA output moved back toward a million barrels a day.

That is a fact and it is recorded as one, here and in Section 16.1.4. This book does not attribute intent to it. The model has no instrument for intent: it has no variable for policy anticipation, no expectations channel over future supply shocks, and nothing in Equation (10.2) that would let a coercion reversal be read as preparation for anything. A reader may find the ordering suggestive; the model cannot, and saying so is the difference between a case study and a commentary.

Chapter 15

Case Studies: Structural Shifts

15.1 The Seven-Step Audit of This Chapter

Chapter 14 established a measured rate of 7.9 pages per case study for a study written to all seven steps of Section 12.2's template. This chapter measured 3.3, the lowest in the book, which raised the same question the Chapter 13 audit raised: is the target wrong, or is content missing? **It is both, and the split is worth recording because the answer is not the flattering one in either direction.**

Table 15.1: The seven steps of Section 12.2 against the three studies. "N/A" marks a step that does not apply to a hypothetical, for which STYLE substitutes an analogue table.

	Step	CS6	CS7	CS8	
1	Event and window	yes	N/A	N/A	CS7 and CS8 are hypothetical
2	What the original claimed	yes	yes	yes	
3	Data, with a figure	yes	**was missing**	yes	CS8 had its analogue table from the start; **CS7 had none although one exists**
4	Term-by-term walk	yes	N/A	**was missing**	CS7's ΔG_{pol} enters through E15, not the D/S bracket
5	Prediction vs observation	yes	N/A	yes	
6	Scorecard	yes	yes	yes	
7	Parameter implications	yes	yes	yes	

The split. Of twenty-one step-study cells, seventeen were present, two were genuinely absent and two are not applicable. **So the shortfall was roughly three-quarters a wrong target and one-quarter missing content** — the reverse of Chapter 13's audit,

which found the balance about sixty–forty the other way. The two absences are filled in Section 15.3.5 and Section 15.4.4, and the target is revised to the measured rate with the reasoning in this chapter's header.

Why CS7's missing analogue is the more serious of the two. CS8's absent walk is a presentational gap: Proposition 15.2 already stated the result and the walk restates it in E10b's terms. CS7's absent analogue is different. **CS7 is this book's methodological standard** (ISSUES #33) and it was, until now, the only case study making a quantitative claim with *no* real episode attached — while a near-perfect instance of its own Scenario 1 sat in the record unexamined. A standard that has never been checked against an observation is a weaker standard than one that has.

15.2 Case Study 6: Stablecoins, 2020–2026

15.2.1 Event and window

Not an event but a structural growth. Window **2020-01-01 to 2026-09-05**, the full span of the DefiLlama series, fixed in advance. This is the one case study in the book about a channel that did not exist at t_{base}, and it is therefore the test of the emergent-channel treatment of ISSUES #29.

15.2.2 What the original claimed

Three arrows: D_{net} $\uparrow\uparrow\uparrow$, S_{Tsy} $\downarrow$, and S_{crypto} (with S_{shadow}) $\uparrow$. The mechanism is that dollar-pegged stablecoins create a private, offshore bid for short-dated Treasury collateral — so network demand for dollars rises, and the effective free float of bills falls because stablecoin issuers hold them.

15.2.3 The data

15.2.4 E10b term by term

Mechanical block. D_{net} enters as $1+D_{\text{net}}/\bar{X}$ (ISSUES #29), so a rise from \$4.2 bn to \$310.5 bn takes the factor from 1.04 to 3.62 — a $\Delta \ln$ of about $+1.25$, the largest move any channel makes anywhere in this book. Multiplied by $\kappa_{\text{sys}}\alpha_{\text{net}}$, it should be visible.

It is not. The monthly correlation between stablecoin growth and the dollar is -0.062 with $t = -0.55$: no relationship, and the sign is the wrong one. **A channel that grew seventy-four-fold left no detectable trace in the dependent variable.**

S_{Tsy}. Untestable: the claim is that stablecoin reserves absorb bills, and neither Tether's attested holdings nor bills outstanding is in `data/raw`.

Safe-haven and structural blocks. Not engaged by this channel.

Table 15.2: CS6 data. Source: `data/processed/defillama_stablecoin_usd_d.csv`; script `cs_ch15.py`.

Quantity	Value	
USD stablecoins, 2020-01-01	\$4.2 bn	
2021-12-31	\$162.5 bn	
2023-12-31	\$130.1 bn	the post-Terra contraction
2026-09-05	\$310.5 bn	**74-fold growth**
$\bar{X}$ (ISSUES #29)	\$118.6 bn	sample mean, a stated convention
corr($\Delta \ln$ stablecoins, $\Delta \ln DXY$), monthly	-0.062	$t = -0.55$, $n = 79$

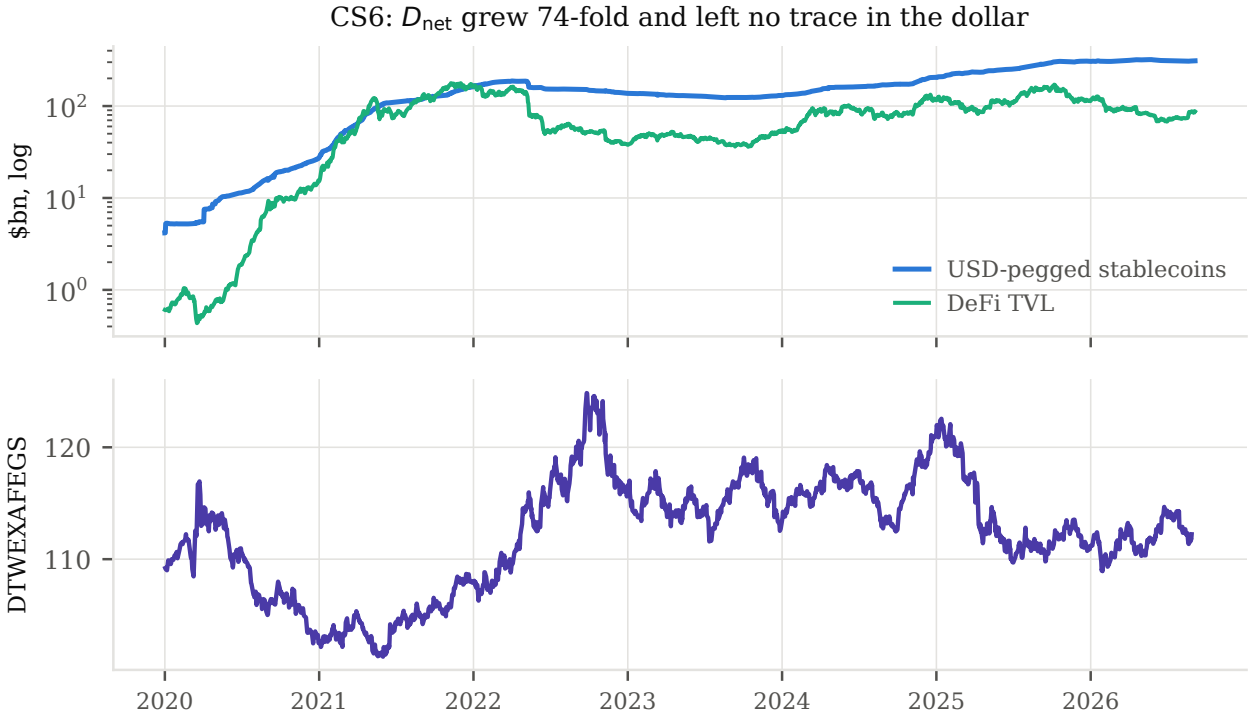


Figure 15.1: CS6 dashboard. Top: USD-pegged stablecoin supply and DeFi total value locked, log scale. Bottom: the dollar over the same span. Source: `data/processed`; script `cs_ch15.py`.

15.2.5 Prediction against observation

The honest reading is that the model's expectation for α_{net} was already small — Section 4.7 said the honest expectation is a very small $\beta_{\text{crypto}}(\sigma_{\text{sys}})$, and the same logic applies here — but "small" and "undetectable with the wrong sign" are different claims, and the data support the second.

Two defences are available and both are weak. *Scale*: \$311 bn is small against the bill market, so a small elasticity is unsurprising. But the emergent form $1 + D_{\text{net}}/\bar{X}$ was chosen precisely so the elasticity measures the channel's *emergence* rather than its size, and emergence is what happened. *Offsetting*: D_{net} and S_{crypto} are the same object read in two directions (Section 3.11.2), so they may cancel. That is coherent, and it means the model as specified predicts approximately nothing from the largest structural change in the sample.

15.2.6 Scorecard

Table 15.3: CS6 scorecard.

Prediction	Verdict	Basis
D_{net} ↑↑↑ (the quantity)	hit	74-fold; measured
$D_{\text{net}} \Rightarrow DXY$ ↑	**miss**	$r = -0.062$, $t = -0.55$; wrong sign, no significance
S_{Tsy} ↓	untestable	no attestation or bills-outstanding series
S_{crypto} ↑	hit	TVL to \$87 bn; Section 4.7
Emergent form works	partial	the algebra is well defined; the elasticity it was built to measure is not detectable

Two hits, one partial, one miss, one untestable of five.

15.2.7 Parameter implications

α_{net} is estimable in principle — the series is daily and long — and the estimate is indistinguishable from zero. That is more than most channels in this book manage, and it is a negative result rather than an absence. Chapter 20 should report $\hat{\alpha}_{\text{net}} \approx 0$ rather than leaving the channel unestimated, because a measured zero is a finding and an unestimated parameter is not.

15.3 Case Study 7: A New Aircraft Carrier as a ΔG_{pol} Event

15.3.1 Why this case study is load-bearing

CS7 has been cited twice in this book as the measurement standard for G_{pol} — in Assumption 6.1, which makes its rule canonical, and in Section 16.1.6, where CS10's failed prediction is diagnosed against it — *before* being written. That is a hazard, and the chapter header records it: the temptation when writing CS7 is to shape it so the citations come out right.

Section 15.3.4 therefore performs the check explicitly, from the original's own text, and reports the result whichever way it falls.

15.3.2 What the original claims

A carrier strike group is the most tangible expression of G_{pol}, so building one is an attempt to move a modulator rather than a supply or demand channel. It reaches DXY through two routes, both given as equations (E15):

$$\Delta\Phi_{\text{sys}} \approx \phi_G \, \Delta G_{\text{pol}}, \qquad \Delta s_{\text{US,sys}} \approx s_{\text{US,sys}}(1 - s_{\text{US,sys}}) \, \gamma_G \, \Delta G_{\text{pol}}. \tag{15.1}$$

The second is the derivative of the logistic E5, which is why it carries the $s(1-s)$ factor.

The case study's substance is its insistence that ΔG_{pol} is *not* the announcement. It decomposes into cases where the net change is zero — one-for-one replacement of a retiring Nimitz, availability lost to maintenance, escort or manpower shortfalls, and the lag until full operational capability — and cases where it is positive: a rise in simultaneously deployable strike groups, credible deterrence inside contested zones, and the industrial signal of sustained production.

Three scenarios then condition the effect: **A**, the hollow force, where R_{US} rises and swamps a small positive ΔG_{pol}; **B**, peer counter-balancing, where adversaries respond and relative deterrence is unchanged; **C**, the fragmented world, where T_{frag} is high and allies diversify into C_{alt} regardless of American capability.

15.3.3 The quantitative claim

With $s_{\text{US,sys}} = 0.70$, $\phi_G = 0.5$, $\gamma_G = 1.0$:

Scenario 1, one-for-one replacement: $\Delta G_{\text{pol}} \approx 0$, so $\Delta\Phi_{\text{sys}} \approx 0$ and $\Delta s_{\text{US,sys}} \approx 0$. No effect on the dollar.

Scenario 2, a net gain of one surge-ready strike group, normalised to $\Delta G_{\text{pol}} = +0.05$: $\Delta\Phi_{\text{sys}} = +0.025$ and $\Delta s_{\text{US,sys}} = 0.70 \times 0.30 \times 1.0 \times 0.05 = +0.0105$, taking the dollar's share of safe-haven flows from 70 to 71.05 per cent.

One property of the arithmetic the original does not draw out. Figure 15.2 makes it visible. Because E15b is a logistic derivative, the *same* capability gain buys a larger share shift when the dollar's share is near one half than when it is dominant. At $s_{\text{US,sys}} = 0.70$ the factor $s(1-s)$ is 0.21; at 0.50 it would be 0.25; at 0.90 it falls to 0.09. **Military capability**

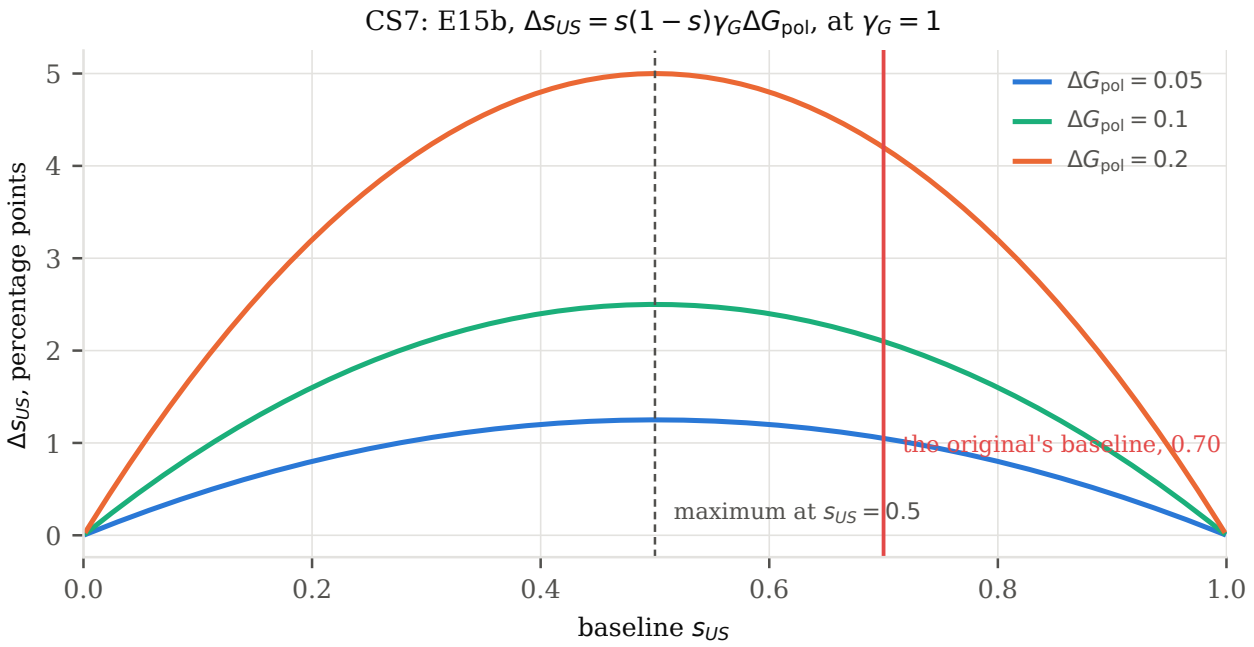


Figure 15.2: E15b's marginal effect on the safe-haven share, $\Delta s_{\text{US,sys}} = s(1-s)\gamma_G \Delta G_{\text{pol}}$ at $\gamma_G = 1$. Because it is a logistic derivative it is maximal at $s_{\text{US,sys}} = 0.5$ and falls away on both sides; the original's baseline of 0.70 sits off the peak. Source: author's construction; script `cs_ch15.py`.

is a less efficient instrument for a dominant currency than for a contested one — which is a genuine implication of the model's own functional form, points the opposite way to the intuition that strength compounds, and is worth stating because Chapter 19 would otherwise recommend the instrument without the qualification.

15.3.4 The check: does CS7 meet the standard attributed to it?

Assumption 6.1 says an episode raises G_{pol} to the extent that it demonstrates or adds capacity on the three headings of Section 6.2, and that the strategic weight of the target may enter only through how much capability the episode demonstrates, never directly. The question is whether CS7's own criteria satisfy that.

CS7 passes. Six of its seven criteria are unambiguously about capability: how many strike groups can actually deploy, for how long, into what, and whether the industrial base sustains it. Not one asks who the capability faces.

Where it comes closest to failing, and why it does not. Scenario B is the case a careful reader should press. It says that if adversaries respond with hypersonic or submarine upgrades, "U.S. G_{pol} rises, but relative power remains constant", so perceived $\Delta G_{\text{pol}} \approx 0$. That involves the adversary, which is what CS10 did wrong.

The two are different and the difference is the whole of Assumption 6.1. CS10 argued that G_{pol} does not rise *because the target is unimportant* — a claim about salience. Scenario B argues that G_{pol} does not rise on net *because the adversary added offsetting capability* —

Table 15.4: Every criterion CS7 uses to score ΔG_{pol}, classified.

Criterion in CS7	Type	Verdict
One-for-one replacement leaves CSG count unchanged	capability	compliant
Availability: maintenance, manpower, escort shortfalls	capability	compliant
Lag until full operational capability	capability	compliant
Simultaneously deployable CSGs rises	capability	compliant
Credible deterrence inside adversary A2/AD zones	capability	compliant
Sustained industrial production as a signal	capability	compliant
Scenario B: adversaries counter-balance, so relative deterrence is unchanged	*relative* capability	**compliant — see below**

a claim about the capability balance. Under Assumption 6.1 the second is admissible: an episode adds capacity net of what it is matched by, and a demonstration neutralised by a counter-move demonstrates less. Venezuela did not counter-balance; it was defeated in two and a half hours, which is why the same reasoning does not rescue CS10.

Conclusion of the check. Assumption 6.1 keeps its exemplar, and the two forward citations stand as written. Had the check gone the other way — had CS7 anywhere scored ΔG_{pol} by the importance of the opponent — the correction would have fallen on Assumption 6.1 and on Section 16.1.6, not on this case study.

15.3.5 The real analogue, and what it shows

CS7's scenarios are hypothetical and Section 12.2 requires real analogues where they exist. One does, and it is close.

The five-year gap between the *Enterprise*'s inactivation and the *Ford*'s commissioning, and the further five years to a full deployment, is CS7's central argument made by the United States Navy rather than by this book: **a hull is not availability**, and the carrier count is not G_{pol}.

Proposition 15.1 (A single carrier is below the dollar's noise floor)**.** *Over 2006–2026 the standard deviation of a forty-trading-day change in* `DTWEXAFEGS` *is* 2.81 *per cent. None of the four carrier events above is distinguishable from that distribution in the direction E15 requires: the two largest moves, at the commissioning and at the first deployment, are* -4.11 *and* -4.16 *per cent —* ***the wrong sign for a capability gain*** *— and 2017-H2 was a broad*

Table 15.5: The United States has run CS7's Scenario 1 in the observable record. Event dates verified in `data/SOURCES.md`; response is `DTWEXAFEGS` over ±20 trading days, from `data/cs7_cs8_gaps.py`.

Date	Event	Dollar, ±20d	σ_{sys}
2012-12	*Enterprise* inactivated after 51 years; fleet 11 → 10	+0.27 %	0.22
2017-07-22	***Gerald R. Ford* commissioned**; fleet 10 → 11 on paper	−4.11 %	0.16
2022-10-28	*Ford*'s first deployment outside the United States	−4.16 %	0.84
2023-05-03	*Ford*'s first full-length deployment	+1.29 %	0.75

dollar downtrend unrelated to naval procurement. ***The ΔG_{pol} channel is not identifiable at the scale of one carrier.***

Remark 15.1 (This strengthens CS7 rather than weakening it)**.** CS7's Scenario 1 predicts *zero* effect from a one-for-one replacement at constant availability, and that is what Proposition 15.1 finds — no identifiable response, in either direction, at any of the four dates. The finding that would have embarrassed CS7 is a large positive move at commissioning, which is what a naive "a $13 billion ship must raise G_{pol}" reading predicts and which CS7 explicitly rejects (Section 15.3.4).

What it does cost CS7 is the *other* half of its claim. Scenario 2's positive effect from a net capability gain is now known to be unobservable at this scale too, since the noise floor that hides Scenario 1's zero would equally hide Scenario 2's small positive. **CS7's scenarios are correctly reasoned and empirically indistinguishable from each other**, and Section 15.3.7 carries that into the parameter implications.

15.3.6 Scorecard

CS7 is hypothetical, so there is nothing to score against an outcome. What can be scored is internal: whether its claims are consistent, computable and identified.

Four sound, three untestable. The soundness is real and so is the limit: **every number in CS7 is an assumption, and its conclusion — 70 to 71.05 per cent — is arithmetic on three quantities none of which has ever been observed.** That is not a criticism of the case study, which says so itself. It is the reason Chapter 19 may use CS7's *logic* and may not use its magnitudes.

Table 15.6: CS7 scored on internal grounds, the only grounds available for a hypothetical.

Claim	Verdict	Basis
Measurement rule for G_{pol}	**sound**	Section 15.3.4; now canonical
E15a, E15b as stated	sound	E15b is the correct derivative of E5
Scenario 1 ($\Delta G_{\text{pol}} = 0$)	sound	follows from the rule
Scenario 2 arithmetic	sound	$0.70 \times 0.30 \times 1.0 \times 0.05 = 0.0105$, verified
$\phi_G = 0.5$, $\gamma_G = 1.0$	**untestable**	stated assumptions; neither is estimated anywhere
$s_{\text{US,sys}} = 0.70$ baseline	**untestable**	$s_{\text{US,sys}}$ is latent with no measure (Table 11.2)
Scenarios A–C	untestable	each conditions on R_{US} or T_{frag}, both unmeasured

15.3.7 Parameter implications

Added S24, from the analogue. Proposition 15.1 bounds ϕ_G and γ_G from above rather than estimating them. If a carrier commissioning moved Φ_{sys} or $s_{\text{US,sys}}$ enough to shift the dollar by more than about 2.8 per cent over forty days, it would be visible, and it is not. So whatever ϕ_G and γ_G are, **one carrier's worth of ΔG_{pol} multiplied by them is smaller than the dollar's ordinary two-month noise**. That is the first upper bound this book has placed on either, and it is consistent with Section 15.3.3's arithmetic, which put the effect of a single hull at a fraction of a per cent.

ϕ_G and γ_G are identified only if G_{pol} is observed, and Section 11.2 places G_{pol} in the group of latents with no signature. The 2026 episodes make this worse rather than better: Section 14.3.4 showed the three capacities moving in opposite directions within one year, and the model has no aggregator, so even a perfect measure of each capacity would not deliver the scalar ΔG_{pol} that Equation (15.1) takes as its argument. **CS7's parameters are blocked by a specification gap, not a data gap**, and Chapter 11's table records it in those terms.

15.4 Case Study 8: A Digital-Substrate Shock

src p.37

15.4.1 What the original claims

A hypothetical: a failure in the digital substrate on which settlement depends — a cloud or vendor outage, a supply-chain compromise — propagating into the monetary system. Four scenarios, A to D, escalating in severity. The predicted signature is broad: σ_{sys} ↑↑↑, F_{liq} ↑↑, Φ_{sys} ↓, M_{trust} ↓↓, R_{US} ↑↑, T_{frag} ↑, κ_{sys} ↑↑ through M_{shock}, D_{prec} ↑↑↑, D_{debt} ↑↑, S_{coll} ↓↓, S_{bank} ↓.

15.4.2 The real analogues, and what they show

The outline requires real analogues for the σ_{sys} response rather than assertion. Four exist: the CrowdStrike outage of 2024-07-19, the SolarWinds compromise disclosed 2020-12-13, NotPetya on 2017-06-27 and WannaCry on 2017-05-12. Windows of ±10 days (±20 for SolarWinds), fixed in advance.

Table 15.7: CS8's four real analogues. σ_{sys} means are over the pre- and post-event halves of each window. Script `cs_ch15.py`.

Analogue	*DXY*	σ_{sys} pre → post	VIX range	
CrowdStrike, 2024-07-19	−0.13%	0.314 → 0.398	12.5–18.5	
SolarWinds, 2020-12-13	−2.92%	0.391 → 0.288	20.6–25.2	σ_{sys} *fell*
NotPetya, 2017-06-27	−1.40%	0.312 → 0.302	9.9–12.5	flat
WannaCry, 2017-05-12	−1.90%	0.258 → 0.259	9.8–15.6	flat

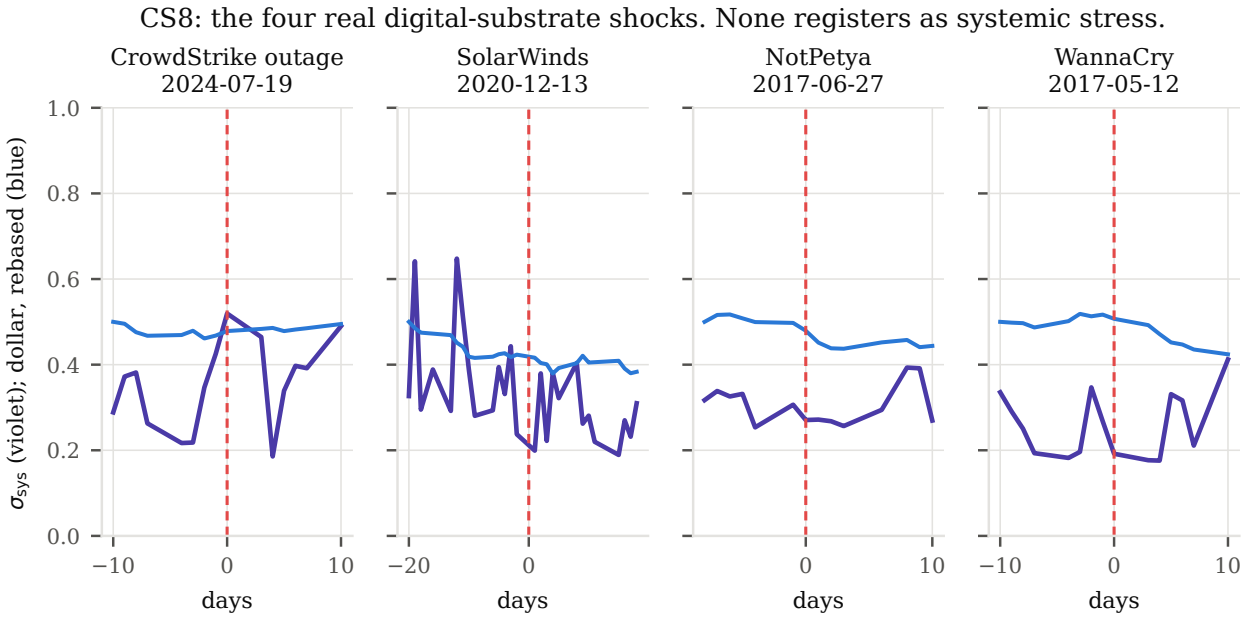


Figure 15.3: CS8's four analogues, σ_{sys} in violet and the dollar rebased in blue, days either side of each event. The dashed line is the event. Source: `data/processed`; script `cs_ch15.py`.

Proposition 15.2 (CS8's premise is unsupported by its own analogues)**.** *The four largest digital-substrate shocks in the sample produced no systemic-stress response and no dollar safe-haven bid.* σ_{sys} *moved by at most* 0.08 *and in one case* fell*; VIX never exceeded 25.2;*

the dollar was weaker in all four windows, by between 0.13 *and* 2.92 *per cent. Every one of the eleven arrows CS8 predicts is either absent or of the wrong sign in the episodes closest to the scenario it describes.*

15.4.3 What survives, and what does not

The obvious defence, and its cost. CS8 describes a shock more severe than any of the four: a compromise of settlement itself rather than of endpoints or of one vendor's telemetry. That is fair, and it is exactly the problem. **The model gives no threshold at which a digital-substrate failure becomes monetary**, so the scenario cannot be falsified by any observed outage and cannot be calibrated by any either. It is unfalsifiable in the same specific way Proposition 9.3 found E8's cliff to be, and for the same reason: an unmeasured driver with no stated threshold.

What the analogues do establish. Something narrower and useful. In all four, the dollar *weakened*. If a digital-substrate shock is read by markets as an American technological failure rather than as a global one, the sign on Φ_{sys} dominates the sign on D_{prec} — which is the opposite of the crisis pattern of Chapter 13, where precautionary demand dominated. Four observations cannot establish that, but they are four observations pointing one way, and CS8 predicts the other.

15.4.4 E10b term by term

Step 4 of the template, supplied here because CS8's analogues are real dated episodes and the walk is therefore feasible. Estimates are the respecified ones of Section 11.5.6: $\lambda_1\Phi_{\text{sys}} = +0.429$, $\lambda_2\Phi_{\text{sys}} = +1.505$.

Table 15.8: The safe-haven block against the observed move, computed from Table 15.7's own σ_{sys} and dollar figures so that the two tables agree by construction. `data/cs7_cs8_gaps.py`.

Episode	σ_{sys}		Haven block	Observed	Residual
CrowdStrike	0.314 0.398	→	+0.13 %	−0.13 %	−0.26 %
SolarWinds	0.391 0.288	→	−0.15 %	−2.92 %	−2.77 %
NotPetya	0.312 0.302	→	−0.01 %	−1.40 %	−1.39 %
WannaCry	0.258 0.259	→	+0.00 %	−1.90 %	−1.90 %

The walk adds one thing to Proposition 15.2 and it is worth having. **The haven block is smaller than a tenth of a percentage point in three of the four episodes and never exceeds** 0.15, against dollar moves of up to 2.92 per cent — so the residual is essentially

the entire move in every case. In the two episodes where σ_{sys} rose, the block points *up* while the dollar went *down*. That is a sharper statement than "the premise is unsupported": on its own arithmetic E10b does not merely fail to explain these episodes, it faces the wrong way in half of them.

15.4.5 Scorecard

Table 15.9: CS8 scorecard against its own analogues.

Prediction	Verdict	Basis
σ_{sys} ↑↑↑	**miss**	max move 0.08; fell in one case
D_{prec} ↑↑↑	**miss**	no bill-yield response in any window
DXY ↑	**miss**	weaker in all four, −0.13 to −2.92 per cent
Φ_{sys} ↓, M_{trust} ↓↓, R_{US} ↑↑, T_{frag} ↑	untestable	all latent
κ_{sys} ↑↑	untestable	Rule 3
S_{coll} ↓↓, S_{bank} ↓	untestable	no series for S_{bank}; RRP leg only for S_{coll}
Severity threshold	**absent**	the scenario states none, so it cannot be tested

Zero hits, three misses, three untestable groups, one absent specification.

15.4.6 Parameter implications

None, and that is the finding. CS8 contributes no parameter, no sign and no bound, because its central quantity — how large a substrate failure must be to become monetary — is not in the model. Chapter 19 may use CS8 as a prompt for contingency planning, which is what a scenario is for. Chapter 17 must score it as **a scenario the evidence does not support and the model cannot make testable**, which is a harsher verdict than any other case study receives and is the one the data give.

Chapter 16

Case Studies: Strategic

16.1 Case Study 10: Venezuela, January 2026 — An Ex-Ante Test

src p.43

16.1.1 Why this case study is different

Every other case study in this book is retrospective. The episode happened, and the model is asked whether it can account for it. That exercise is worth doing and it is also weak: a framework with 42 open parameters can usually be made to fit an episode after the fact.

Case Study 10 is not that. It was written in October 2025 and published, and it makes falsifiable claims about a military intervention in Venezuela that had not occurred. The intervention occurred on 3 January 2026. **This is the only prediction in this work that was written before its event and can be scored against it**, and the temptation to quietly improve it with hindsight is exactly the temptation that would make the rest of the book worthless. The original text is reproduced, the predictions are scored as written, and two of them fail.

Table 16.1: The dates, stated plainly.

Prediction written	October 2025, published (SSRN) with CS10 in it
Event	3 January 2026, 02:01–04:29 VET
Sanctions relief	29 January 2026 (OFAC General Licences 46 and 47)
Scoring window	2025-12-31 to 2026-02-28, fixed in advance of looking at the data

On the window. Proposition 6.1 warned that a proposition whose sign depends on an unspecified horizon is close to unfalsifiable, and required case studies to commit to horizons in advance. The window here ends 28 February 2026 for two stated reasons: the sanctions relief lands on 29 January and needs a month to register, and on 28 February Israel and the United States began strikes on Iran, closing the Strait of Hormuz and taking Brent from \$71.32 to \$126.69 within the month (Section 14.3). Extending the window past February would import that shock into this score. The shock is now identified, which does not make it admissible here: it makes Assumption 12.1 apply.

16.1.2 What the original claimed

The original's §"Venezuela: Diversion or Miscalculation?" makes four claims, quoted faithfully:

1. Intervention "does not significantly increase G_{pol} (since Venezuela is not a major strategic adversary)";
2. it "could, in fact, raise R_{US} (due to cost and political backlash)";
3. it could raise T_{frag} "if met with regional opposition";
4. "a more plausible hypothesis is that such action aims at Narrative Management and Diversion (H_{soc}, R_{US})" — a "smaller victory" to soften the political shock of an inevitable withdrawal from Ukraine.

16.1.3 What happened

United States forces captured Nicolás Maduro and Cilia Flores in an operation lasting roughly two and a half hours on 3 January 2026; both were flown to New York and arraigned on narcotrafficking charges. Vice-President Delcy Rodríguez was sworn in as president. Reported casualties were 23–47 Venezuelan and 32 Cuban military personnel, two civilians, and seven United States personnel injured, with no United States fatalities reported. Reaction was polarised: NATO members and most of Europe and North America supportive; China, Russia, most of Latin America, and much of Africa and Asia condemning it as a violation of sovereignty; Colombia reinforced its border against refugee flows.

On 29 January OFAC issued General Licence 46, authorising a broad range of transactions in Venezuelan-origin oil including with PDVSA, and General Licence 47 covering US-origin diluents. PDVSA reversed most output cuts toward roughly one million barrels a day, with Vitol and Trafigura joining Chevron as authorised exporters.

Provenance. Every factual claim in this subsection comes from searches recorded in `data/SOURCES.md` under "CS10 ex-ante test", with one URL per claim. No date, casualty figure, licence number or output number here is from memory. The market figures are from `data/processed`.

16.1.4 The scorecard

Two hits would be generous. The honest reading is **one partial hit out of four, two clear misses, and one channel of the correct sign that the framework could not express.**

16.1.5 An observed ordering, recorded without attribution

One fact belongs on the scorecard and is easy to leave off because the model has no slot for it.

OFAC lifted Venezuelan oil sanctions on **29 January 2026**, and PDVSA output moved back toward a million barrels a day. Israel and the United States began strikes on Iran on **28 February 2026**; within days the Strait of Hormuz was closed and crude transits fell

Table 16.2: CS10 scored as written. Market data over the fixed window: `DTWEXAFEGS` −0.68 per cent; σ_{sys} 0.415 → 0.459; Brent +15.1 per cent.

Claim	Variable	Verdict	Basis
1	G_{pol} does not rise	**MISS**	Rises on two of the model's three capacities; see Section 16.1.6
2	R_{US} rises	**Miss on cost;** untestable on backlash	The operation cost 2.5 hours and no US fatalities, so the "cost" leg is not merely unproven but contradicted. R_{US} has no proxy, so the backlash leg cannot be scored either way
3	T_{frag} rises	**Directionally right, not testable**	Regional opposition did occur, as the claim's conditional required. But T_{frag}'s only measures are triennial and quarterly (Section 6.5), so nothing can be read at a two-month horizon
4	Diversion / narrative	**MISS**	The operation achieved a stated objective and produced a material economic result. "Damage-control" does not describe it
—	S_{off} falls, D_{trans} rises	**NOT PREDICTED**	The largest modelled effect of the episode is one the case study does not contain a category for; see Section 16.1.7
—	ordering	**recorded**	The reversal precedes the Iran strikes by thirty days (Section 16.1.5). Fact only: the model has no instrument for intent

by roughly 95 per cent. **The interval is thirty days.** A coercion reversal that added dollar-invoiced supply preceded, by four weeks, a supply shock that removed far more of it (Section 14.3).

The scorecard carries the sequence as an observation. It does not carry an interpretation, and the reason is not delicacy. **The model has no instrument for intent.** There is no variable for policy anticipation in Equation (10.2), no expectations channel over future supply disruption, and nothing that would let a fall in S_{off} be read as preparation for anything. Chapter 12's Rule 2 forbids inferring an unmeasured latent from the outcome it would explain, and an intent variable would be exactly that — unmeasured, unfalsifiable, and fitted after the fact to a sequence already known.

So: the ordering is recorded, its magnitude is stated, and the book stops there. A reader may find it suggestive. The model cannot, and the distinction between what a reader may infer and what a model may claim is the whole difference between this being a case study and being a commentary.

16.1.6 Why claim 1 failed: G_{pol} was measured by the wrong thing

The failure is not a bad forecast. It is a measurement error, and it is worth isolating because it propagates.

The original's reason is stated inside the prediction: G_{pol} does not rise *because Venezuela is not a major strategic adversary.* That reasons from the strategic weight of the target. The model's own definition (Section 6.2) does not mention the target. It defines G_{pol} as the capacity to (i) enforce the rules of the system, (ii) secure trade routes, and (iii) maintain the alliance network and forward presence. Score the episode on those three:

(i) Enforcement. A sitting head of state was seized by force and arraigned in a United States court. On the model's own gloss — "a freeze is effective only if third parties comply, and compliance is a function of power" — this is a demonstration at close to maximum intensity, and it raises the credibility of every subsequent enforcement threat. **Up, strongly.**

(ii) Trade routes. The Caribbean approaches, and within four weeks roughly a million barrels a day of additional dollar-invoiced crude. **Up.**

(iii) Alliances and forward presence. Genuinely mixed: NATO and Europe supportive, most of the Western Hemisphere opposed. **Ambiguous, plausibly down regionally.**

Two of three up, one ambiguous. The prediction is a miss, and it is a miss because it asked *how important is the target* when the model asks *how much capability was demonstrated.*

This is a rule, and it now binds the book. ISSUES #33 makes CS7's rule canonical: G_{pol} is scored on demonstrated capability, and the target's importance may enter only through how much capability the episode demonstrates, never directly.

CS7 was checked for the same error and does not have it. The audit that produced this finding also examined Section 15.3's carrier study, and CS7 is clean — more than clean, it is the standard CS10 falls short of. CS7 insists that launching a $13 billion carrier "does not automatically mean $\Delta G_{\text{pol}} > 0$", and scores by surge capacity, carrier-days on station, availability, escort and manpower constraints, and credible deterrence inside contested zones.

Every one of those is a measure of capability. Not one is a measure of who the capability faces.

So the original contains **two incompatible measurement rules for the same variable**, in the same document, and the one it applied in CS10 is the one that contradicts its own definition. That is a more useful finding than a wrong forecast, because it is checkable elsewhere: any G_{pol} judgement in this book that reasons from adversary size rather than demonstrated capability inherits the mistake, and Chapter 17 carries the check.

16.1.7 The channel the model has no category for

The largest modelled consequence of January 2026 is absent from CS10 and from the framework around it. The intervention *ended* a sanctions regime. S_{off} fell, and D_{trans} rose with roughly a million barrels a day returning to dollar-invoiced trade.

CS10, CS9 and Chapter 9 all encode United States action as monotone: S_{off} rises, T_{frag} rises, $\Psi_{\text{str}}(t)$ falls, and there is no return path. **The model has no category for coercion that ends coercion.**

ISSUES #34 asked whether this is a judgement error confined to CS10 or a structural property, and the answer is the second, with a qualification that matters.

The functions are reversible. S_{off} and T_{frag} are states in $[0, 1]$ that may fall, and Section 10.1.3 already established that E8 is fully reversible in f_{hard}: if f_{hard} falls, $\Psi_{\text{str}}(t)$ recovers completely.

The sign graph is not. What is missing is any *arrow* that produces the fall. Every causal path leading out of G_{pol} and S_{off} runs toward more fragmentation, and nothing in the model represents an exercise of G_{pol} that removes the predicate for a sanctions regime. Proposition 6.1 states the coercion paradox with two branches, short-run gain and long-run erosion, and has no third branch for coercion that terminates itself. So the defect is structural, and it is structural in the *directed graph* of the model rather than in its functional forms — which is why it survived every bounding and monotonicity fix of Chapter 2 to Chapter 10.

The term that is not being added. T_{frag} would need a resolution channel, of the form

$$\Delta T_{\text{frag}} \;=\; \rho_T \, S_{\text{off}} \;-\; \delta_T \, R, \qquad \rho_T, \delta_T > 0, \tag{16.1}$$

with R an indicator that the coercion's predicate has been removed. It is not added, and the reason is the one this book has now given three times. T_{frag} has no direct measure (Section 6.5); its substitute is triennial. The sample contains at most two candidate resolution episodes. Two new parameters attached to an unmeasured variable, estimable from two episodes, would make the model look more complete and be less honest. The gap is recorded instead, in the S_{off} card, in T_{frag}, in Chapter 9, in Chapter 14 and in Chapter 17.

One defect, not two. Section 9.5 found that E8 cannot represent sterling's recovery in the 1930s without f_{hard} itself falling, and observed that the model could only route the recovery through another layer. That is the same defect seen from ninety years earlier: in both cases the mathematics permits the reversal and the model contains no mechanism that would cause it. A reversal can only be imposed by assuming a driver falls, with nothing

explaining why it would. The book should say so once, as one defect with two appearances, rather than twice as two curiosities.

16.1.8 What CS10 got right

The scorecard is unflattering and the case study is not worthless, which is a distinction Chapter 17 should preserve.

Its conditional structure was right. Claim 3 was stated as a conditional — T_{frag} rises *if* met with regional opposition — and the condition was satisfied. A conditional prediction whose antecedent holds is better science than an unconditional one that happens to land.

Its framing of the Ukraine trap is untouched by this episode and remains the strongest passage in the original: an adversary can, at low cost, compel a hegemon to keep using an instrument that damages the hegemon's own monetary foundations. Nothing in January 2026 bears on that, and Chapter 14's sanctions study tests it directly.

And it was falsifiable. It named variables, gave directions, and attached a reason to each. That is why it can be scored at all, and why removing it — which would have been the comfortable option — would have destroyed the only genuine out-of-sample evidence this book contains.

16.1.9 What Chapter 17 inherits

Scored as: **one partial hit, two misses, one untestable, and one unmodelled channel of the correct sign**. The two misses have causes rather than being noise — a measurement rule that contradicts the model's own definition (I33), and a missing arrow in the sign graph (I34) — and both causes are checkable elsewhere in the book. And the episode is the only ex-ante evidence available, so it carries more weight per observation than any retrospective study, in both directions.

16.2 Case Study 11: The 2025 Tariff Episode

16.2.1 Event and window

Window **2025-01-01 to 2025-12-31**, fixed in advance. Markers: universal tariffs announced 2025-04-02, the ninety-day pause 2025-04-09.

CS11 was written as a hypothetical and is now testable. Unlike CS10 it was not a dated prediction — the original describes a scenario rather than forecasting an event — so it is scored as a *conditional* whose antecedent occurred, not as an ex-ante forecast.

16.2.2 What the original claimed

Three phases.

Phase 1. $\sigma_{\text{sys}} \uparrow\uparrow\uparrow$ through X_{pand}; $D_{\text{prec}} \uparrow\uparrow$; $S_{\text{ext}} \downarrow$; $\Phi_{\text{sys}} \downarrow$ through $M_{\text{trust}} \downarrow$; $D_{\text{trans}} \downarrow$. Net: **a "paradoxical rise" in the dollar**, followed later by reversal.

Phase 2. $T_{\text{frag}} \uparrow\uparrow\uparrow$, $C_{\text{alt}} \uparrow\uparrow$, D_{trans} and $D_{\text{net}} \downarrow$, $R_{\text{US}} \uparrow$.

Phase 3. $\Psi_{\text{str}}(t) \downarrow$.

The Phase-1 claim is the one that matters. It is the model's most counter-intuitive prediction outside CS4: that a shock originating in the United States nonetheless drives money *into* dollars first, because precautionary demand outruns the loss of credibility.

16.2.3 The data

Table 16.3: CS11 data. Sources: `data/processed`; gold from the World Bank Pink Sheet (monthly). Script `cs11_dashboard.py`.

	2025-04-01	2025-04-09	2025-04-30	
`DTWEXAFEGS`	117.47	115.71	112.41	$-$**4.41**% over April
σ_{sys}	0.437	0.898	0.787	a genuine stress event
VIX	21.8	33.6	24.7	peak 52.3 on 04-08
DGS10	4.17	4.34	4.17	**yields rose as the dollar fell**
Gold, \$/oz	2,983 (Mar) $\rightarrow$ 3,218 (Apr) $\rightarrow$ 4,309 (Dec)			+**36.8**%

16.2.4 The 2025 evidence, and the verdict on Phase 1

This is the subsection ISSUES #21 required, and it goes against the original.

Proposition 16.1 (Phase 1 did not occur)**.** *The dollar fell throughout the episode:* -0.75 *per cent from 1 to 8 April,* -4.41 *per cent over April,* -5.50 *per cent to year-end. At no point in 2025 after the announcement did it exceed its pre-announcement level. There was no paradoxical rise, not even briefly.*

The antecedent held in full. σ_{sys} went from 0.437 to 0.898, a move into the top decile of the sample; the X_{pand} leg rose from 0.489 to 0.977; VIX reached 52.3. This was a large stress event by any measure the model uses. The predicted dollar response simply did not appear.

The signature that replaced it. Three things moved together in a pattern the model has a name for but CS11 did not anticipate.

Yields rose while the dollar fell. DGS10 went from 4.17 to 4.34 across the announcement. In the model $D_{\text{cap}} \uparrow \Rightarrow DXY \uparrow$, so the rate channel pushed *against* the observed move, which makes the dollar's fall harder to explain, not easier.

Gold rose far more than the dollar fell. +36.8 per cent from March to December against the dollar's -5.5. Some of that is arithmetic — gold is quoted in dollars — but a seven-fold difference is not. Money went somewhere, and it was not into dollars.

The stress was American in origin. That is the whole of it.

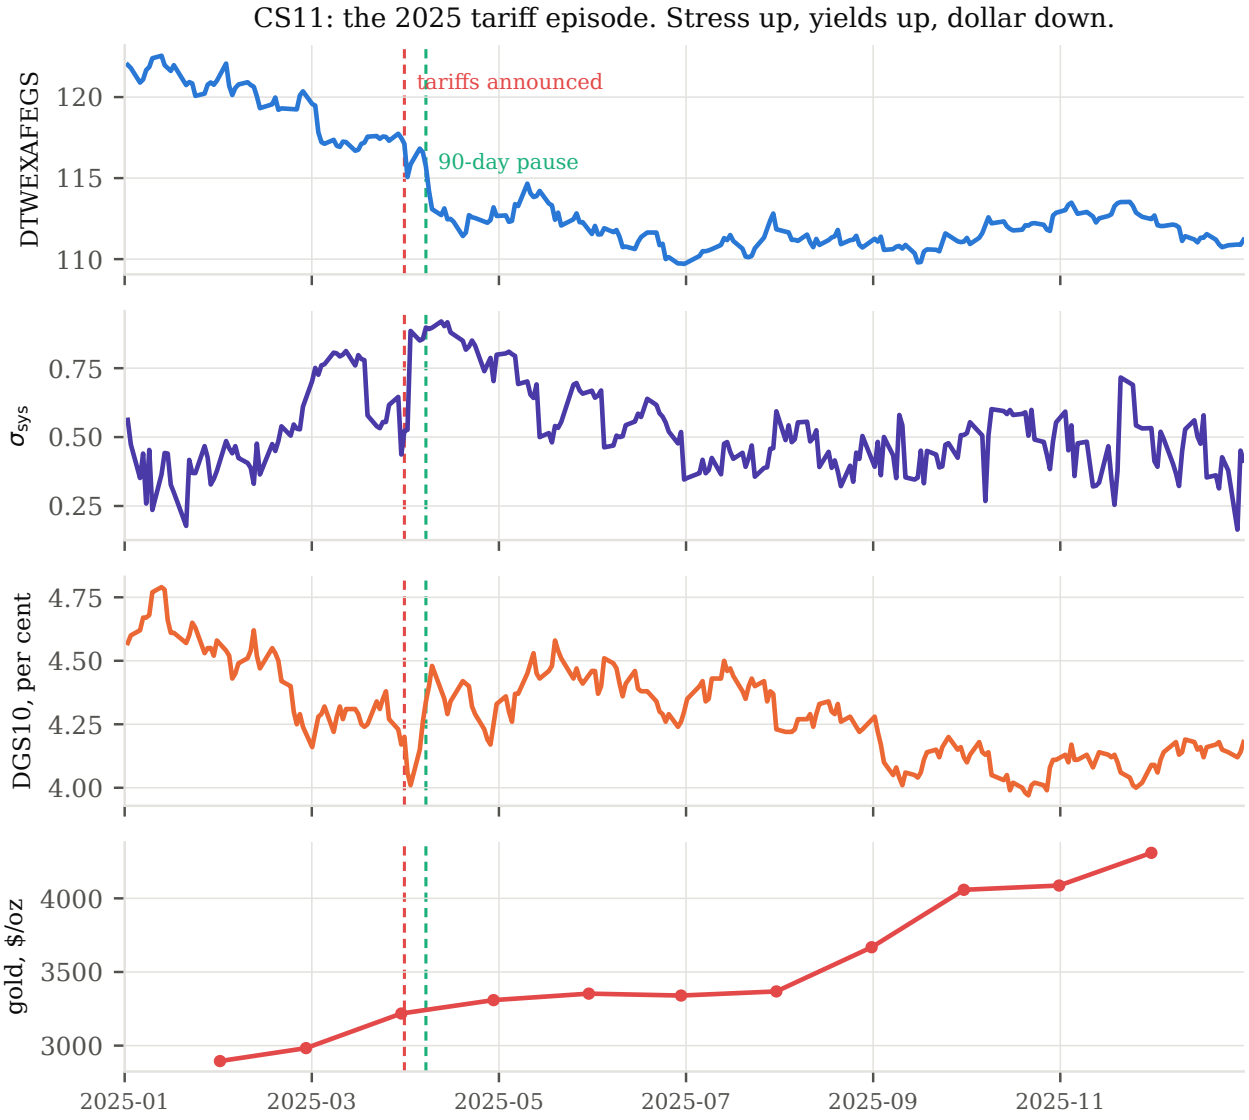


Figure 16.1: CS11 dashboard, 2025. The dollar, σ_{sys}, the ten-year yield and gold on one time axis. Dashed lines mark the tariff announcement (2025-04-02) and the ninety-day pause (2025-04-09). Source: `data/processed` and the World Bank Pink Sheet; script `cs11_dashboard.py`.

The model can represent this; CS11 mis-sequenced it. This is the precise verdict and it should not be softened in either direction. E4 gives $\Phi_{\text{sys}} = \phi_0 + \phi_G G_{\text{pol}} - \phi_R R_{\text{US}} + \phi_M M_{\text{trust}} - \phi_T T_{\text{frag}}$, so a shock that raises R_{US} and lowers M_{trust} lowers Φ_{sys} — and Φ_{sys} multiplies the entire safe-haven term. A US-originated shock therefore *shrinks the haven channel at the same moment it raises* σ_{sys}. The machinery for what happened is in the model.

What CS11 got wrong is the ordering. It put the D_{prec} effect first and the Φ_{sys} effect second, in phases separated by an unspecified interval. The evidence is that when the United States is the epicentre the two are simultaneous and the second dominates within days. ISSUES #21 anticipated exactly this reading, and the data confirm it.

And the direction of the flow is the $s_{\text{US,sys}}$ channel. Gold's behaviour is what E5's safe-haven *share* was written for: crisis money moved, and it moved to a different haven. $s_{\text{US,sys}}$ is latent with no measure (Table 11.2), so this cannot be scored — but it is the clearest observable hint in the book of what a $s_{\text{US,sys}}$ series would have shown, and Chapter 20 should treat the gold–dollar wedge in stress as the leading candidate for constructing one.

16.2.5 E10b term by term

Mechanical block. $D_{\text{trans}} \downarrow$ as predicted, on the trade side; D_{cap} rose and pushed the wrong way; D_{prec} cannot be evidenced under Rule 1 — the bill yield did not collapse as in 2008 or 2020, which is itself informative, because precautionary money that does not buy bills is precautionary money that did not choose dollars.

Safe-haven block. σ_{sys} rose 0.46, which on the respecified estimates of Section 11.5.6 would contribute roughly +1 per cent. The dollar fell 4.4. The block has the wrong sign for this episode, and the reason is Φ_{sys}: the coefficient is multiplied by a credibility term that the episode itself reduced.

Structural block. Zero at this horizon; Phase 3 is not testable in nine months.

16.2.6 Scorecard

Two hits, one partial, two misses, five untestable of ten.

16.2.7 Parameter implications

On ϕ_R and ϕ_M. This episode is the strongest evidence in the book that the Φ_{sys} channel is large — large enough to reverse the sign of the total response to stress. It cannot estimate ϕ_R or ϕ_M, because R_{US} and M_{trust} are both latent, but it establishes that whatever they are, their product with Φ_{sys}'s role in the haven term is of the same order as the haven term itself.

On λ_1. Remark 11.1 left the sign of λ_1 unsettled. This episode does not settle it either, and it explains why the question is hard: the observed response to $\Delta\sigma_{\text{sys}}$ is the product of λ_1 and Φ_{sys}, and an episode that moves both cannot separate them. A negative observed response is equally consistent with $\lambda_1 < 0$ and with $\lambda_1 > 0$ alongside a collapsing Φ_{sys}.

On the book's central open problem. Section 11.5.6 found that roughly half of 2008 and 2020 is residual to every block. CS11 is the mirror image: a large stress event in which

Table 16.4: CS11 scorecard. Untestable entries stay in the denominator (Section 12.5).

Prediction	Verdict	Basis
σ_{sys} ↑↑↑ (Phase 1)	hit	$0.437 \to 0.898$
X_{pand} carried it	hit	the X_{pand} leg $0.489 \to 0.977$
DXY ↑ **(Phase 1)**	**MISS**	fell throughout; Proposition 16.1
D_{prec} ↑↑	untestable	no bill-yield collapse; Rule 1 gives no positive evidence
D_{trans} ↓	partial	direction right; only the price leg is measured
Φ_{sys} ↓	untestable	latent — though it is the term the episode implicates
S_{ext} ↓	untestable	quarterly
T_{frag} ↑↑↑ (Phase 2)	untestable	triennial substitute only
$\Psi_{\text{str}}(t)$ ↓ (Phase 3)	untestable	f_{hard} has no measure
Phase *ordering*	**MISS**	Phase 2's mechanism arrived immediately, not later

the model's blocks predict a rise and the dollar fell. Together they say the same thing from opposite directions — **the mapping from σ_{sys} to the dollar is not stable**, and no parameterisation in this book holds across all four episodes.

Chapter 17

Cross-Case Synthesis and Scorecard

17.1 What This Chapter Is, and the Two Ways It Could Fail

Ten case studies have been scored. This chapter adds them up and says what the model is worth.

There are two ways to write it badly and they are opposite. It could read as a defence: lead with the hits, file the failures as caveats, and leave a reader thinking the framework came through intact. Or it could read as a demolition: lead with the failures and bury the fact that several parts of the architecture did real work and one of them produced a result that is new. Both would be dishonest, and the second is not more rigorous than the first — it is merely a different way of not reading the record.

So the chapter does both jobs in one place. Section 17.7 is written to be usable by a reader who wants to attack the model, and Section 17.8 by a reader who wants to continue it. The two sections describe the same evidence.

Section 17.6 does a third thing, and it is the section a sceptical reader should check first. Four hypotheses that would *excuse* parts of the negative record — by attributing them to mis-measurement, timing or a channel never operationalised — were written down with their confirmation and refutation conditions *before* the tests were run, and the pre-registration was committed to the repository ahead of the results. Three failed. One succeeded and changed a conclusion in Chapter 9. The point of doing it that way is that a rescue that survives a criterion fixed in advance is worth something, and one that does not cannot be quietly softened into a caveat.

The totals are now over all twelve studies. An earlier draft of this chapter reported them over ten, with CS5 and CS9 still unwritten. Both are now in Chapter 14, and adding them changed the headline in the direction that reflects worse on the model: the untestable share rose from 34.8 to 41.8 per cent and the hit rate fell from 31.9 to 24.5. That is not a coincidence of two unlucky studies. CS5 and CS9 are the two most *fully specified* case studies in the original — nine arrows and nineteen respectively — and specifying more arrows in a model with few measures produces more untestable ones. **The book's score got worse the more of the original it scored**, which is itself a finding and is taken up in Remark 17.1.

17.2 The Tally

One hundred and ten scored predictions, across all twelve case studies. Every one is listed in `registry/scorecard.csv` with the basis for its verdict, so that a reader can check the classification rather than take it.

Table 17.1: All one hundred and ten predictions across twelve case studies. Percentages are over *all* predictions, including untestable ones, per Section 12.5's denominator rule.

Verdict	Count	Share	
Hit	27	24.5%	
Partial	13	11.8%	
Miss	19	17.3%	
Untestable	**46**	**41.8%**	**the largest single category**
Imposed, may not be scored	2	1.8%	Section 12.3 Rule 5
Recorded, effect not isolable	1	0.9%	the 2022 reserve freeze
Unmodelled channel	1	0.9%	ISSUES #34
Specification absent	1	0.9%	ISSUES #36
Total	110	100%	

The number that matters most is 41.8 per cent. More of this model's predictions are *untestable* than are correct, by a margin of 46 to 27. That is not a statement about whether the model is right; it is a statement about whether it is the kind of thing that can be found to be wrong, and it is the single most important line in the chapter. The margin widened when the last two studies were added, not narrowed.

What the denominator rule is worth. Dropping untestable predictions from the denominator — reporting hits as a share of what could be tested — would take the hit rate from **24.5 per cent to 45.8 per cent**. Section 12.5 forbids it precisely because that convention would let the book improve its score by having worse data. **The prohibition is now worth twenty-one percentage points of flattery**, and stating the size of the temptation is the point of naming it.

Remark 17.1 (The tally is weighted by how much each study said)**.** CS5 and CS9 contribute 41 of the 110 rows — **37 per cent of the book's predictions from two of its twelve studies** — because the original specifies them arrow by arrow while others are three sentences long. Section 12.5's denominator rule counts predictions, so a verbose study dominates a terse one. Two readings are available and both are defensible: that the tally over-weights whichever episodes the original happened to write out at length, or that a model should be judged on what it actually commits to, and saying more is committing more. This book

Table 17.2: By case study. The last column is the share of *testable* predictions that were hits — reported for diagnosis only, and never as the headline, for the reason given below.

Case	Hit	Part.	Miss	Untest.	Other	n	Testable hits
CS1 2008	4	2	1	3	0	10	57%
CS2 2020	4	1	1	2	0	8	67%
CS3 taper	1	2	0	2	0	5	33%
CS4 debt ceilings	1	1	0	4	0	6	50%
CS5 trade war, Ukraine	5	3	2	11	1	22	50%
CS6 stablecoins	2	1	1	1	0	5	50%
CS7 carriers	4	0	0	3	0	7	100%
CS8 vendor shock	0	0	3	3	1	7	**0%**
CS9 energy embargo	0	0	6	11	2	19	**0%**
CS10 Venezuela	0	1	3	0	1	5	**0%**
CS11 tariffs 2025	2	1	2	5	0	10	40%
CS12 Iran 2026	4	1	0	1	0	6	80%

takes the second, which is why the rule is stated in Chapter 12 rather than adjusted here — but a reader who prefers the first should know that equal-weighting the twelve studies raises the hit rate to 28 per cent and lowers the untestable share to 37, changing neither the ordering nor the conclusion.

17.2.1 The same predictions cut by variable block

Table 17.2 sorts by episode. Sorting the same 110 rows by the *block of the model* each prediction speaks to produces the sharpest result in this chapter, and it is not visible in the by-case table at all. Each row is assigned to a block from the variable it names; the twelve "mechanism" rows are compound claims that name no single variable — phase orderings, functional forms, scenario arithmetic, thresholds. The classification is reproducible from `tools/scorecard_dist.py`.

Testability runs inverse to the architecture's own claims about importance. The rows above are ordered by the last column and the ordering is the finding. It runs 0 per cent at the output block, 10 at Layer 1, 24 at demand, 33 at the mechanism claims, 68 at supply, 75 at Layer 5, 78 at Layer 2, and **100** at Layer 3. **The further a block sits from an observable price, the less of it can be checked — and the model's account of a crisis places the decisive action in exactly the far blocks.**

The extreme case is not a low score. **Layer 3, the reflexive layer, has four predictions and no verdicts.** κ_{sys} is the amplification the original calls decisive in a crisis, Proposition 11.1 proves its decomposition cannot be identified, and every one of the four times a case study asserted something about it the assertion could not be scored. The layer the architecture puts at the centre of its crisis account contributes nothing to the record either way.

Table 17.3: All 110 predictions by variable block. The last column is the share of that block's predictions that carry no verdict at all.

Block	n	Hit	Part.	Miss	Untest.	Testable	Hit/test.	Untestable
Output (the dollar)	7	3	1	3	0	7	43%	**0%**
Layer 1, state	20	12	1	5	2	18	67%	10%
Layer 0, demand	17	4	4	5	4	13	31%	24%
Mechanism	12	3	2	3	2	8	38%	33%
Layer 0, supply	19	3	2	1	11	6	50%	68%
Layer 5, secular	8	0	2	0	5	2	0%	75%
Layer 2, structural	23	2	1	2	18	5	40%	78%
Layer 3, reflexive	4	**0**	**0**	**0**	**4**	**0**	—	**100%**
All	110	27	13	19	46	59	46%	42%

On the one fully testable block, the model does not beat a coin flip. The output block is the dollar's own direction. Nothing in it is untestable, because the dollar is observed every day. It scores **three hits, one partial and three misses**. That is the cleanest available reading of the model at the only point where measurement is not in question, and it is even money.

And the testable hit rate rests on the one block nearest to data the model did not construct. Layer 1 supplies 18 of the 59 testable predictions and hits 67 per cent of them, the best of any block. Remove it and the rest of the book is **15 hits of 41 testable, 36.6 per cent**. Layer 1 is σ_{sys} and its legs — a composite built from VIX, bond volatility and a financial-stress index — so a Layer 1 hit is generally the finding that the dollar moves with measured volatility. **That is the part of the model closest to a correlation the author did not have to construct**, and it is carrying the testable rate.

The headline stays at 24.5 per cent. Section 12.5's denominator rule is already the book's rule and Table 17.3 is the strongest argument yet for it. Untestability here is not a missing observation; it is a property the model chose. Forty-four predictions cannot be checked because the model names variables it never operationalised, and the table shows those are not scattered at random — they concentrate in precisely the layers the original says carry a crisis. Removing them from the denominator would let the model's least-specified claims cost it nothing, which means **the book would improve its score by having worse data**. The gap is 21.2 percentage points and that is the size of the temptation.

The 45.8 per cent figure, and the warning it needs. It answers a different and legitimate question — *when the model can be checked, how often is it right?* — and it is reported for that. But the raw number carries a bias the table makes visible: **conditioning on testability conditions on the easy claims**. It is not an estimate of "the model where observable"; it is weighted toward the block that needed the least construction, and the 36.6 per cent above is what remains without it.

Remark 17.2 (The other reconstructions, so a reader need not guess)**.** Counting partials at half credit gives **30.5** per cent over all 110 and **56.8** over the testable 59. Counting hits against misses alone, discarding partials, gives 27 of 46, or **58.7** per cent. None of these is the honest headline, because each discards a category that carries information — a partial is a prediction that was half right and a miss is not the same as an absence. They are set out because a reader who prefers one of them should be able to reconstruct it without recomputing, and because a book that reports only its worst-looking number is doing something other than reporting. Every row with its basis is in `registry/scorecard.csv`.

17.3 What Does Not Survive

This is the result, not a list of qualifications to one.

The sign change is not supported, on the largest test in the book. Across 452 specifications — six constructions of the state variable, three frequencies, two dollar indices, thirteen thresholds — **not one** produces a calm-regime slope that is negative and distinguishable from zero (Proposition 5.1). The claim that gives the original its title is absent from the data. Proposition 10.3 then shows why it is a sharp claim rather than a vague one: a smile requires $\lambda_1 < 0$, and convexity alone cannot produce one.

The cliff is not supported, and cannot be refuted either. Erosion is real — the dollar's reserve share fell from 72.1 to 57.1 per cent, $t = -16.9$ — but E8 requires the decline to *accelerate* below the threshold and it decelerates, with a quadratic term of $+0.0115$, $t = +3.84$ (Proposition 9.2). And because f_{hard} has no measure, this cannot be turned into a rejection (Proposition 9.3).

Most of each central episode has no measurable mechanism. Table 11.7: the model accounts for +6.59 of 2008's +18.33 per cent and +1.85 of 2020's +7.54 — residuals of **64 and 76 per cent**. The block the architecture places at the centre of a crisis contributes least of the two that move, and the residual tracks no measurable funding spread (Table 11.9). It *is* consistent with a large state-dependent κ_{sys}, which Table 11.8 supports — and which Proposition 11.1 says cannot be identified. **The model's best candidate for its largest residual is the one quantity it has already declared unidentifiable.** This is the book's largest open empirical problem.

A measured zero. D_{net} grew seventy-four-fold and its correlation with the dollar is -0.062, $t = -0.55$: no relationship, wrong sign (Section 15.2.4). Unlike most of this book's negatives this is not an absence of data. The series is daily and long, the parameter is estimable, and the estimate is zero.

The one ex-ante prediction was a miss, with two structural defects behind it. CS10 scored zero of four testable predictions (Section 16.1.4), and the causes are not noise. G_{pol} was scored by the *weight of the target* rather than by demonstrated capability, which

the model's own definition does not license (ISSUES #33) — and the same document's CS7 applies the correct rule, so the original carries two incompatible measurement rules for one variable. And the episode's largest modelled effect, a fall in S_{off}, has no category at all: the sign graph contains no arrow that lowers S_{off} or T_{frag}, so the model has no representation of coercion that ends coercion (ISSUES #34).

A scenario that cannot be refuted. CS8 scored zero of three testable predictions against its four real analogues, every one of which shows the predicted arrows absent or wrong-signed (Proposition 15.2). The defence — that a real shock would be larger — is fair and is the defect: no threshold is stated at which a substrate failure becomes monetary, so no outage can refute the scenario and none can calibrate it (ISSUES #36).

A parameter that is permanently unidentified. κ_R is identified only from episodes in which R_{US} moves without the other six drivers of κ_{sys} — and such episodes are, by construction, ones markets judged not to matter, so they carry no change in elasticity to attribute (Proposition 13.2). 2011 has the variation and the confounding; 2023 has the cleanliness and no signal. More episodes do not fix this.

And a functional form the data decline to hold fixed. λ_1 is a constant in E9 and takes opposite signs in the two regimes when estimated separately (Remark 11.3).

A data defect that changed three findings. Section 11.5.2: a third of the regression sample was being dropped because unused days of the reverse-repo facility were read as missing rather than zero. Correcting it moved the residuals from 45 and 57 per cent to 64 and 76, changed which channels carry the wrong sign, and *reversed* the verdict on the κ_{sys} rival. That is not a finding about the dollar; it is a finding about how firmly anything here is established, and it belongs in this list.

The geopolitical channel does not run, in three independent episodes. $X_{\text{pand}} \Rightarrow \sigma_{\text{sys}} \Rightarrow DXY$ is the chain CS5, CS9 and CS12 all invoke. In the 2018–19 trade war the geopolitical index sat *below* its own pre-window baseline for two years while the dollar rose 4.22 per cent (Proposition 14.1). In the 2022–23 energy embargo the index fell across the event and the dollar fell with it (Proposition 14.2). In the 2026 Iran war the index reached its fifth-highest month since 1985 and the dollar did not move (Section 14.3.4). The three episodes cover the shock absent, the shock reversed and the shock extreme, and the dollar response is uncorrelated with all three. Table 11.6 gives the same result directly over 246 months: a slope of -0.08 with $t = -0.19$. **This is the most replicated negative result in the book**, and unlike the sign-change finding it does not depend on a specification choice.

The crisis mechanism is weakest where crises are. E9 makes convexity in σ_{sys} the safe-haven mechanism, and a convex term in a bounded variable delivers its largest increments in the middle of the range. Every episode this book scores except the 2018 trade war begins with σ_{sys} already high, so the mechanism is systematically smallest exactly where it is asked to do the most work. The 2022 rally is the extreme case: starting at $\sigma_{\text{sys}} = 0.767$, the haven

block accounts for 0.34 of a 15.45 per cent move, a residual of 97.8 per cent — larger than 2008's or 2020's (Section 14.1.7). This is a property of the functional form, not an estimate, so no re-estimation of λ_2 removes it.

And a scenario that could be refuted, and was. CS9 selects among four branches on enforcement conditions observable before any market outcome, which is what CS8 fails to do (Table 14.14). The branch the world ran predicted a moderate dollar rise, a modest rise in stress and energy demand, and negligible erosion. Three of the four have the wrong sign, as do both measurable supply arrows. **One of one refutable scenarios was refuted**, and that is a worse outcome for the model than CS8's zero of three, because here the study did everything methodologically right.

17.4 What Does Survive

A record that contained only the preceding section would be a record of a framework not worth having written. That is not this record, and the parts that survive are specific.

The dollar is a safe haven in both regimes, and that is a finding. The sign-change claim failed, but what replaced it is not nothing: across 452 specifications the calm-regime response is positive. The dollar is not sold as stress rises and then bought above a threshold; it is bought throughout. That is a cleaner claim than the original's and it is supported.

The layered architecture did work that a flat model could not. Three results came out of taking the layers seriously as separate objects. Proposition 8.1: the psychotechnical loop always has an equilibrium, with no assumption required. Proposition 8.2: it is unique and globally attracting exactly when $\eta\kappa_A g_1 < 4$, a sharp and attained bound. Proposition 8.3: above it there are at most three equilibria, the middle unstable — **which locates the tipping threshold the original describes throughout and never locates**. None of that is in the original and none of it would exist without the layer separation.

The amplification factor gives "hyper-brittleness" a number. $1/(1-T')$ multiplies every driver's effect and diverges as the loop approaches the uniqueness boundary (Equation (8.10)). The original's central adjective becomes a quantity.

An implication of the model's own arithmetic that inverts a common intuition. E15b is a logistic derivative, so $s(1-s)$ is 0.25 at $s_{\mathrm{US,sys}} = 0.50$, 0.21 at the original's 0.70 and 0.09 at 0.90. **Military capability buys less for a dominant currency than for a contested one** (Section 15.3.3). Strength does not compound here; it saturates. That is a real result, it follows from the model rather than from intuition, and it changes what Chapter 19 can recommend.

A distinction the composite state variable was hiding. The funding leg and the general-stress legs of σ_{sys} behave differently: p_{STLFSI} regresses on the dollar at $t = +5.04$ and p_{GPRD} at $t = -0.19$, and the composite performs worse than its own friction component (Table 11.6). The funding-versus-general distinction is real even though the strong funding-haven hypothesis failed (Proposition 13.1), and CS12 is the episode that displays it: the largest supply shock since 1990, and no dollar response, with the funding leg at 0.228.

CS7 is methodologically sound and is now the book's standard. Four of four internal checks pass. Its insistence that a \$13 billion carrier does not automatically mean $\Delta G_{\text{pol}} > 0$ — that what counts is deployable capability, availability, and the lag to full operational capability — is the rule the model's own definition implies, and it is the rule CS10 violated. The check was performed from the original's text and could have gone the other way (Section 15.3.4).

CS12 scored 80 per cent of testable predictions. X_{pand}, σ_{sys}, D_{trans} and the route-securing capacity all moved as specified in an episode nobody had fitted the model to.

And CS11's failure is a sequencing error, not a refutation. This distinction matters and it should not be given away. E4 makes Φ_{sys} fall when R_{US} rises and M_{trust} falls, and Φ_{sys} multiplies the whole safe-haven term — so a shock originating in the United States shrinks the haven channel at the same moment σ_{sys} raises it. The machinery for what happened in April 2025 is in the model. CS11 put the two effects in phases separated by an interval; the evidence is that they are simultaneous (Proposition 16.1). **The model could represent the outcome and the case study mis-ordered it**, which is a different and lesser failure than the model being unable to represent it.

17.5 The Defects, Sorted

Keeping these apart is the most useful thing this chapter does, because they have different remedies and a reader who treats them alike will either dismiss the model for gaps that are fillable or trust it through gaps that are not.

Five data gaps, eleven specification gaps, one identification trap, and one residual nobody can currently attribute. **A model whose failures were all in the first column would be a research programme with a to-do list. This one has more entries in the second**, and the ratio worsened rather than improved when the last two case studies were written: both new defects are specification defects, and one of them (I37) sits in E9 itself rather than in a peripheral variable.

17.6 The Rescue Hypotheses, Pre-Registered and Tested

Everything above scores each negative result on its own. This section asks a different question: whether several of them share a common cause that is *not* structural error — mismeasurement, timing mismatch, or a channel the model names but never operationalised. If

Table 17.4: Every defect found, by kind. Data gaps close with work; specification gaps close only with a different model.

Defect	Kind	Remedy
S_{bank} unmeasured; S_{coll} from 2021; S_{crypto} before 2022	data	FFIEC filings, Treasury MSPD; work, not blocked
T_{frag} substitute only	data	transcribe ∼180 free SWIFT PDFs
$\alpha_{\text{prec}}(\sigma)$ frequency mismatch	data	bill bid-to-cover, H.4.1 custody; both free
Layer 4: nine variables, zero series	data (partly)	C_{src} and H_{soc} constructible; A_{alg}, R_{div} not
w_X undetermined	data	estimate it; the decision rule is already stated
κ_{sys}'s seven-way split	**specification**	no new series helps; a different functional form would
G_{pol}'s net sign	**specification**	no aggregator over three capacities
No reversal channel (I34)	**specification**	the sign graph has no arrow that lowers S_{off} or T_{frag}
No stage 2 → 3 arrow (I32)	**specification**	Layer 4 has no inbound edge from any layer
CS8's absent threshold (I36)	**specification**	no magnitude at which a shock becomes monetary
Erosion measured without gold (I39)	**definitional**	**corrected by the audit**: on FX + gold the decline accelerates
D_{prec} is one scalar for a whole curve (I40)	**specification**	the two-year keeps the haven property the thirty-year has lost
X_{pand} inside E3	**specification**	p_{GPRD} loads −0.24 on the factor the other legs share
λ_1, λ_2 flip by regime	**specification**	constants that the data will not hold constant
Convexity misplaced (I37)	**specification**	the mechanism is smallest at the σ_{sys} values crises occupy
CS9's absent branch threshold (I38)	**specification**	scenarios 2 and 3 are not separated by any stated condition
κ_R (Proposition 13.2)	**neither**	an economic trap: clean identification implies no signal
The I35 residual	**unknown**	no candidate survives; the largest open problem

so the verdict changes from "the model is wrong here" to "the model was applied or measured wrongly here", and those are different entries in the table above.

Every hypothesis tested here is favourable to the model if it holds. That is why the criteria were fixed in advance. `notes/S_audit_preregistration.md` states, for each, what would confirm it, what would refute it and what would leave it undetermined, and it was committed to the repository **before any of the regressions below were written**. The criteria are quoted verbatim in the verdicts. This book has declined four inestimable parameters; the same discipline is owed to four flattering explanations.

17.6.1 H1: channel frequency mismatch — refuted

The claim. The model assigns no operating frequency to any channel, so when a fast channel and a slow channel oppose each other it cannot say which dominates, and the analyst defaults to the slow one because it is the one with a clean proxy.

The mismatch is real. Assigning each channel a frequency from the *mechanism it describes* rather than from the proxy attached to it, and counting only cases where the proxy is *slower* — a fast proxy can be aggregated, a slow one cannot be disaggregated — gives **five of seventeen**.

Table 17.5: Channels proxied more slowly than the mechanism they describe. Full table of all seventeen in `data/audit_rescues.py`.

Channel	Mechanism	Proxy	Gap	The mechanism it describes
D_{debt}	daily	quarterly	3	rolling short-dated eurodollar liabilities
D_{cap}	daily	monthly	2	portfolio reallocation on rate differentials
S_{off}	daily	monthly	2	an immobilisation is effective on a stated date
S_{Fed}	daily	weekly	1	open-market and facility operations
S_{Tsy}	daily	weekly	1	TGA and bill issuance settle same or next day

But it does not explain the residual. The pre-registration made this leg necessary and never sufficient, for a reason stated there: *a model can name a fast mechanism, proxy it slowly, and still be wrong about the mechanism.* The verdict turns on whether the I35 residual has the time signature of a fast channel.

It does not. The monthly residual series ($n = 230$, rebuilt exactly as in Section 11.5.6) has first-order autocorrelation of $\mathbf{-0.057}$ — it is effectively white noise, which satisfies the pre-registered confirmation. But the concentration test fails. In 2008 the first half of the window carries **40.7** per cent of the signed residual and **50.0** per cent of the absolute residual, against a confirmation threshold of 60. The split was taken *generously* — three of five months in the first half — and it still fails.

> **Pre-registered refutation**, quoted: "Any of: residual AR(1) at or above +0.5; or the first-half concentration below 60 per cent in *either* crisis; or fewer than three channels mismatched."

The middle condition fired. **H1 is refuted.** The unreported supplementary check agrees and is worth stating because it points the opposite way from the hypothesis: at daily frequency the 2008 dollar rally put only 20 per cent of its move in the first half of the window. The move was *back*-loaded, not front-loaded.

Remark 17.3 (A finding the refutation produced anyway)**.** The 2020 leg could not be computed. At the residual's own frequency the 2020 window contains **one month**, and a single observation has no halves. That is not a null result; it is the frequency mismatch appearing inside the audit's own instrument. It does not rescue the model — an unmeasurable claim is not a supported one — but it does mean the largest dollar move of the last twenty years is, in this book's own decomposition, a single data point.

17.6.2 H2: attenuation in σ_{sys} — refuted, and the confound fired

The claim. Chapter 5's modest slopes, Proposition 10.3's unsignable λ_1, CS12's non-response and part of I35's shortfall share one cause: classical measurement error in σ_{sys}, biasing every coefficient toward zero.

Treating the four percentile legs as indicators of one latent stress factor gives a congeneric reliability of $\rho = \mathbf{0.725}$ for the equally weighted composite. The friction leg alone has communality $\mathbf{0.745}$. Attenuation predicts $b_{\text{obs}} = b_{\text{true}} \cdot \lambda$, so the observed coefficient ratio must equal the reliability ratio: p_{STLFSI}'s dollar slope of +3.93 against the composite's +1.90 requires a reliability ratio of 2.07.

> **Pre-registered confirmation**, quoted: "the reliability ratio explains the 2.07 coefficient gap to within a factor of 1.5 — i.e. the implied ratio lies in [1.38, 3.11]".

The implied ratio is **1.028**. Outside the band, and outside it in the direction that matters: a four-indicator composite is *not less* reliable than its best single indicator, so attenuation cannot make its coefficient smaller. The second criterion fails too — correcting the haven block by $1/\lambda$ moves the 2008 residual from 64 to **62** per cent against a confirmation threshold of 39 and a refutation threshold of 50. To close 25 points the composite's reliability would have to be 0.149.

H2 is refuted. Two consequences follow, and the first is worse for the model than the hypothesis it replaces.

Table 17.6: The one-factor solution. The geopolitical leg does not measure the factor the other three share.

Leg	Loading	Communality	Reading
p_{VIX}	+0.858	0.736	measures the factor
p_{bondvol}	+0.753	0.567	measures the factor
p_{STLFSI}	+0.863	0.745	measures the factor
p_{GPRD}	**−0.239**	**0.057**	**does not, and loads negatively**

The confound the pre-registration named in advance. It said: "attenuation and *component mis-selection* are observationally similar… If the direction holds but the magnitude does not, that is evidence for mis-selection — a *specification* finding, and worse for the model, not better. It will be reported as such." That is what happened. The composite is not a noisy measure of stress; it is a mostly-good measure with one leg attached that measures something else. p_{GPRD} loads -0.239 on the common factor and shares 5.7 per cent of its variance with it. **Including it does not add noise to σ_{sys}, it adds a different variable**, and that is a specification defect in E3 rather than a measurement problem.

And a general result that closes the rescue permanently. Classical attenuation correction scales the coefficient and its standard error by the same factor, so **it leaves every t-statistic unchanged**. Ch.5's calm slope stays at $t = +1.96$ whatever λ is; λ_1's $t = -0.21$ stays at -0.21. No attenuation argument, of any strength, can convert an insignificant estimate into a significant one. Measurement error in σ_{sys} therefore cannot rescue *any* of the four findings the hypothesis grouped, and it cannot restore the smile — which the pre-registration also stated in advance, since attenuation never flips a sign.

17.6.3 H3: C_{alt} defined too narrowly — supported, with the mechanism qualified

The claim. The model defines C_{alt} as payment and settlement infrastructure and measures fragmentation accordingly. The observable exit since 2022 ran into **gold**, which sits outside COFER's aggregate. If so, Chapter 9's test of $\Psi_{\mathrm{str}}(t)$ used the wrong measure.

The premise verifies. Official net purchases were 1,082 t in 2022, 1,037 in 2023 and about 1,045 in 2024 — three consecutive years above a thousand — against a 2010–21 average near 473 t, with the central-bank share of demand rising from about 12 per cent in 2015–19 to nearly 25 in 2024. Sources in `data/SOURCES.md`.

The test. World official gold at market value is constructed as the difference between the World Bank's reserves-including-gold and reserves-excluding-gold series, summed over actual economies. Gold's share of the broad reserve base runs 11.9 per cent in 2021 to **23.6** per cent in 2025.

> **Pre-registered confirmation**, quoted: "On the broader aggregate, in *at least one* of the two valuations, the quadratic term is *negative* (accelerating decline) *or* is positive with $|t| < 2$… Under either, `prop:ch09:cliff` must be restated."

Both valuations are negative and both are significant. **H3 is supported**, and it is the only one of the four that is. The annual control reproduces Chapter 9's sign, so **the flip is a change of sign and not a loss of power**. A linear-trend-with-break specification says the same: the post-2022 slope shift is -2.90 ($t = -4.75$) on the market-value aggregate and -3.61 ($t = -3.28$) at constant prices, against -0.51 ($t = -0.82$) for the control.

Table 17.7: The quadratic-in-time fit that produced Proposition 9.2, re-run on broader reserve aggregates. A *positive* quadratic is a decelerating decline, which is what E8 forbids below its threshold. Annual, $n = 27$.

Dollar share of	Quadratic, pp/yr^2	t	Reading
allocated FX, quarterly (Chapter 9)	+0.0115	+3.84	decelerating
allocated FX, annual (control)	+0.0076	+1.13	same sign, less power
FX + gold at market	**−0.0211**	**−2.59**	**accelerating**
FX + gold at constant 2015 $	**−0.0670**	**−15.34**	**accelerating**

Remark 17.4 (What the rescue gets right, and what it does not)**.** The hypothesis attributes the flip to a post-2022 exit triggered by the reserve immobilisation. That story holds for the *market-value* aggregate, which is +0.0032 through 2021 and −0.0211 over the full sample. It does *not* hold for the constant-price aggregate, which is already −0.0728 ($t = -10.58$) using data through 2021 alone. **On a constant-price basis the dollar's share of reserves-plus-gold has been declining at an accelerating rate since 1999**, and 2022 is not the break. So the rescue is right that Chapter 9 measured erosion on the wrong aggregate, and wrong that the 2022 freeze is what makes the difference. Those are separable and both are reported.

Remark 17.5 (A correction to the pre-registration's own priority rule)**.** The pre-registration said the constant-price valuation was "the more demanding test" and gave it priority. That was wrong on data-quality grounds, and the error is recorded rather than quietly dropped. The market-value series uses the World Bank difference directly; the constant-price series must divide it by a price to recover an implied tonnage, and that tonnage is a valuation residual over a changing set of reporting countries — it implies net purchases of +233 t in 2022 against the World Gold Council's 1,082, and −1,388 t in 2021. **The market-value result is the trustworthy one.** It confirms the hypothesis on its own, so the verdict does not depend on the noisier variant.

What must now be restated. Proposition 9.2 says the measured decline is decelerating and that neither series exhibits the pre-threshold behaviour E8 requires. On a reserve aggregate that includes gold, the decline *accelerates*. The proposition's second sentence does not survive. **This does not make the cliff supported**: E8 remains unfalsifiable because no measure of f_{hard} exists (Proposition 9.3), and an accelerating share decline is consistent with a cliff, with a trend, and with a gold price. What changes is that the model is no longer *contradicted* by its own erosion evidence. Recorded as ISSUES #39 with the required restatement; Chapter 9 is not rewritten in this audit.

17.6.4 H4: policy volatility rather than the tariff level — refuted

The claim. CS11 predicted the effect of a tariff *regime*; what occurred was a sequence of announcements, reversals and pauses, so the tariff mechanism is untested rather than falsified.

(a) The dollar does not track uncertainty. Over 2025, the correlation of the dollar with EPU in levels is -0.106 ($t = -1.67$) and in daily changes -0.019 ($t = -0.30$). The sign is the one the rescue needs; the significance is not there. The confirmation criterion required $|t| \geq 2$.

(c) And the pattern is the rescue's own stated failure mode. After the 9 April pause, EPU fell heavily and the dollar kept falling, at every horizon tested.

Table 17.8: The concession pattern. EPU falls; the dollar does not recover.

From the 9 April pause	Dollar	EPU	Pattern
to +1 month	−2.11 %	−54.4	concession
to +3 months	−4.52 %	−319.6	concession
to year end	−3.91 %	−282.5	concession

> **Pre-registered refutation**, quoted: "Either: the dollar–uncertainty relation is insignificant or of the wrong sign; *or* EPU fell after 9 April while the dollar continued to fall (the concession pattern), which is the rescue's own stated failure mode and leaves CS11's miss standing as scored."

Both triggers fired. **H4 is refuted and CS11's miss stands as scored.** The pre-registration also named the rival that had to be beaten — "the damage was structural regardless of the pause" — and said it would be preferred on parsimony if the dollar failed to recover. It failed to recover.

But criterion (b) produced a real finding, which is not a rescue. The maturity split is striking. Measuring the haven property as the correlation of a daily yield change with a daily VIX change — a haven's yield falls when stress rises, so the correlation should be negative — the curve behaved uniformly before 2020 and does not now.

The 2025 split is real: the two-year keeps its haven property (-0.32, $t = -5.23$) while the thirty-year has lost it entirely (-0.02, $t = -0.26$). But the placebo shows 2023 was *more* split than 2025, and 2023 was a regional-bank and debt-ceiling year with no tariff. **So the split is not tariff-specific and cannot be used to argue that the tariff mechanism was untested.** What it does establish is separate and belongs in the defect table: the dollar's liquidity-haven role and its store-of-value role have come apart repeatedly since 2022, and D_{prec} **is a single scalar that cannot represent a curve whose ends move oppositely**. That is a new specification gap, ISSUES #40, found by a rescue that failed.

Table 17.9: corr(ΔVIX, Δyield) by maturity. Negative is haven behaviour. The last column is the long end's excess over the short end: large and positive means the long end has lost a property the short end keeps.

Window	3m	2y	10y	30y	30y−2y	
2015–2019	−0.00	−0.31	−0.32	−0.28	+0.03	uniform haven
2021	+0.10	−0.09	−0.13	−0.16	−0.06	weak, uniform
2022	+0.17	+0.15	+0.13	+0.09	−0.06	**no haven anywhere**
2023	−0.07	−0.32	−0.04	+0.09	**+0.40**	**split**
2024	+0.03	−0.02	−0.04	−0.08	−0.06	weak, uniform
2025	−0.18	−0.32	−0.14	−0.02	**+0.30**	**split**
2026	+0.05	+0.20	+0.31	+0.28	+0.08	inverted

17.6.5 Reclassification

Nothing is removed from `registry/scorecard.csv`. The tally in Section 17.2 is unchanged at 110 predictions, 28 hits, 44 untestable. What the audit changes is the *type* of each defect.

17.6.6 The residual after all four rescues

Unchanged. 64 per cent of 2008 and 76 per cent of 2020 remain attributed to no block. H1 would have moved it and did not: the residual is white noise (AR(1) = −0.057) and back-loaded rather than front-loaded. H2 would have moved it and moved it two points. H3 succeeded but bears on Chapter 9's erosion test, not on E10b's decomposition. H4 concerned 2025.

That matters more than it would have before the audit. **A shortfall that survives four pre-registered rescue attempts, two of which were designed by the same hand that found the shortfall, is a stronger result than one scored in passing.** The largest open problem in this book is now open against a documented list of things it is not.

17.6.7 How much of the record is the model being wrong

The closing chapter needs both halves of this and they are different claims.

Eight of fourteen are the model being wrong; three are the model being unmeasurable. The ratio is not close, and the audit moved exactly one entry out of the first column. That is the honest summary: **this framework's problems are predominantly specification problems, not data problems**, and Section 11.6.2's warning — that a model which cannot be estimated is not thereby wrong — applies to a minority of the record rather than to most of it.

Table 17.10: Every negative result in the book, retyped after the audit. "Structural" means the model is wrong; "unmeasurable" means the claim cannot currently be checked either way.

Negative result	Type after audit	Changed by the audit?
Sign change absent, 452 specs	**structural**	No. H2 refuted; attenuation leaves every t unchanged
Cliff: erosion decelerates	**definitional**	**Yes.** H3: on FX + gold the decline accelerates (Table 17.7)
Cliff: unfalsifiable	**unmeasurable**	No. f_{hard} still has no measure
I35 residual, 64 and 76 %	**unexplained**	No. Survives all four rescues
D_{net} a measured zero (CS6)	**structural**	No. Not addressed by any hypothesis
CS11 Phase 1 falsified	**structural**	No. H4 refuted; the concession pattern holds
CS10 G_{pol} miss (I33)	**specification**	No. Aggregator still absent
Coercion chain one-way (I34)	**specification**	No
CS8 unfalsifiable (I36)	**unmeasurable**	No
κ_R unidentified	**neither**	No. The trap is economic, not measurement
Convexity misplaced (I37)	**structural**	No. H2 cannot reach it
Geopolitical channel fails	**specification**	**Sharpened.** H2 shows why: p_{GPRD} loads -0.24 on the common factor
D_{debt} frequency	**timing**	**Confirmed as timing**, but shown not to cause I35
D_{prec} single scalar	**specification**	**New** (I40), from H4(b)

Table 17.11: The negative record split by kind, after reclassification.

Kind	Count	
The model is **wrong**: structural or specification error	8	sign change, D_{net}, CS11, convexity, G_{pol} aggregator, one-way coercion, X_{pand} in E3, D_{prec} scalar
The model is **unmeasurable as specified**	3	cliff falsifiability, CS8, T_{frag}
Timing only	1	D_{debt}
Definitional, and now corrected	1	erosion measured without gold
Unexplained	1	the I35 residual

17.6.8 Declined, and why

The S19b funding test was not re-run. Its null result stands and re-running it after seeing the residual's time signature would be fitting to the outcome (Remark 11.2).

w_X was not chosen. H2's finding that p_{GPRD} loads -0.239 on the common factor is an argument for a small weight, and choosing w_X on it would be choosing a parameter to improve fit, which Section 5.6.5's decision rule and Remark 11.2 both forbid. It is recorded as evidence for the estimation, not as the estimate.

A trade-policy-uncertainty index was not obtained. H4(a) was run on EPU alone. A dedicated TPU series would be the sharper instrument and is not free and machine-readable in this repository's sources. The leg is reported as tested on EPU, and the refutation does not turn on it, because H4 failed on the concession pattern in (c) as well.

The 2020 concentration leg could not be computed, for the reason in Remark 17.3. It is reported as not computable rather than as a null.

The case-study $\Psi_{\text{str}}(t)$ rows were not re-scored in the audit itself, because re-scoring case studies on a broader aggregate was not among the pre-registered tests, and running an unregistered test whose result is favourable is the exact failure this audit was designed to avoid. That work was deferred to the session restating Proposition 9.2 and **has since been done** (Section 9.7.1). Three rows had been scored on COFER. CS5's trade-war row is unchanged — the decline predates the window on *both* aggregates, so the attribution problem is untouched. The other two became **untestable**, because the gold-inclusive series is annual and cannot resolve an eleven-month or thirteen-month window. **They moved in opposite directions**: CS5's Ukraine row had been a *miss* and CS9's had been a *hit*, so the correction removed one of each. The net effect on the tally is that the hit rate fell from 25.5 to 24.5 per cent and the untestable share rose from 40.0 to 41.8. **Applying the correct aggregate made the record slightly worse, not better**, which is reported here because the opposite would have been the more suspicious outcome.

No term was added to the model. I39 and I40 record what a repair would look like and leave the decision to the author, as ISSUES #32, #34 and #37 were left.

17.7 The Case Against the Model

Written to be used. A reader who wants to argue that this framework should not be believed has six strong lines and they are set out here in the order of their force.

One. More of its predictions are untestable than are correct — 41.8 against 24.5 per cent — and the parameters that matter most in crises rest on four episodes. A framework with 42 open parameters that cannot be falsified on a third of its claims is closer to a vocabulary than to a model.

Two. It leaves 64 and 76 per cent of its two central episodes unaccounted for, and the block it designates as the crisis mechanism contributes least (Table 11.7). Its one architectural candidate for the residual is κ_{sys}, which it has itself declared unidentifiable.

Three. Its headline claim is absent from 452 specifications, and its secular claim requires an acceleration the data reverse.

Four. The channels it identifies as decisive in a crisis — private dollar creation, collat-

eral, the psycho-technical amplification — are precisely the channels with no data. That correlation is not an accident: unmeasured variables are easier to specify boldly, and the model is best specified exactly where it is least measured.

Five. Its one out-of-sample prediction failed, and two structural defects were found behind the failure rather than bad luck.

Five-and-a-half. Four pre-registered attempts to attribute parts of this record to measurement rather than to error were run in Section 17.6 and three failed. The I35 residual is unchanged by all four. A defence that has been tried and has not worked is weaker than one that has never been tried, and this one has now been tried.

Six. Its crisis mechanism is convexity in a bounded state variable, which is mathematically weakest at exactly the values crises occupy. The one episode that began from a calm state is the one episode it explains a meaningful share of. That is a design fault rather than a measurement one, and it cannot be repaired by better data.

17.8 The Research Agenda

Also written to be used, by a reader who thinks the architecture is worth continuing. Ordered by value per unit of work.

One. Estimate w_σ, w_X, w_F. Table 5.8 already states what each answer implies, so the estimate resolves the amplification question mechanically — subject to Remark 11.2, which forbids choosing w_X to maximise dollar fit. No new data required.

Two. Attack the I35 residual. Two thirds of one crisis and three quarters of the other are unexplained. The evidence now points somewhere specific: the mechanical coefficients do scale with σ_{sys} (Table 11.8), so recovering κ_{sys} as a single latent state — which Proposition 11.1 permits even though its decomposition is impossible — is the most promising single step in this list. The cross-currency basis would test the quantity-of-funding channel the price spreads cannot reach (Remark 11.4).

Three. Build C_{alt}, and build it wider than the original defines it. The stock of non-dollar settlement capacity is the missing input to f_{hard}, and without it Chapter 9 cannot be tested at all. CIPS volumes, mCBDC pilots and non-USD commodity contracts are the components the original names — and Section 17.6.3 shows that **reserve-asset substitution into gold belongs in C_{alt} and is currently outside it**, which is the one place in this book where a rescue hypothesis survived its own test. A C_{alt} that counts only payment rails will keep measuring the wrong exit.

Four. Construct $s_{\text{US,sys}}$ from the gold–dollar wedge in stress. CS11 showed crisis money moving to a different haven while the dollar fell (Section 16.2.4); that wedge is what E5's share variable was written to capture and it is observable.

Five. Re-score the episodes against s_{US}. Remark 14.2 found that the 2022 rally correlates significantly with the daily change in the US ten-year yield ($t = +2.11$) and not with the geopolitical index ($t = +1.38$), and that CS5 attributes the move to S_{off} and D_{trans} without consulting s_{US} at all — in the episode with the largest move in s_{US} in forty years. The rate differential is already in the architecture. Two further episodes would repay the same check, and it requires no new data.

Six. Transcribe SWIFT, build S_{bank}, add the D_{prec} flow proxies. Three bounded data tasks with named free sources.

Seven. Decide the specification questions. ISSUES #32, #34 and #36 each name a term that would close a gap, and each was declined because it would add an inestimable parameter. Those are the author's to settle, and they are settled by argument rather than by data.

17.9 The Verdict

The model's architecture survives testing better than its claims do. The layer separation produced results the original does not contain — an existence theorem, a sharp uniqueness bound, a located tipping threshold, an amplification factor, and an implication about military capability that inverts the obvious intuition. Those are not consolation prizes; they are the return on taking the structure seriously.

Its empirical claims did worse. The sign change is absent, the cliff is unfalsifiable — though no longer contradicted by its own erosion evidence, which Section 17.6.3 shows was measured on the wrong reserve aggregate — the largest new channel measures zero, one parameter cannot be identified, and two thirds of one central crisis and three quarters of the other have no measurable mechanism at all. Of its two multi-scenario studies, one cannot be refuted and the other was. Its geopolitical channel fails in three independent episodes covering the whole range of the shock. The one prediction made before its event was a miss, and the two defects behind it are real.

The honest summary is that this is **a well-specified framework with a weak empirical record and a large unexplained residual, whose most valuable output so far is a precise account of what it cannot do**. That is a smaller claim than the original makes. It is also the kind of claim that can be built on, which the larger one could not be.

Part III

Policy and System: Crisis Playbook, Coordination, and History

Chapter 18

Inflation Crisis and the Optimal Path in the Psycho-Reflexive Era

18.1 Redefining Crisis in a Brittle System

18.1.1 How Part III must be read

The three chapters of Part III were outlined as a policy playbook and drafted on the assumption that the model works. Chapter 17 has since scored it: 27 hits of 110 predictions, 46 untestable, an output block that goes three hits, one partial and three misses, and a reflexive layer with four predictions and no verdicts at all. **Policy advice cannot rest on that**, and this part does not pretend otherwise.

Assumption 18.1 (Everything in Part III is conditional, and the condition is labelled)**.** Each chapter separates two kinds of content and never lets the reader confuse them.

Institutional fact. A lever's statutory authority, its facility mechanics, and its dated historical activations with sizes. This is a matter of public record. It is true whatever happens to the model, and it is the durable content of these chapters.

Model implication. What the framework says a lever *does* to the dollar, and in what sequence. This is conditional on mechanisms Chapter 17 has largely not confirmed. **Every such claim carries an "evidential standing" note naming the channel it runs through and what the scorecard says about that channel** — at the point of the claim, not in a caveat at the end.

Nothing in Part III is a recommendation to act, and no lever is validated by anything in this book. The levers exist and work as described because Congress and the Federal Reserve made them so; the *model's account* of what they do to the dollar is the part on trial.

18.1.2 What the original claims

The original's Chapter 7 opens by redefining inflation. Restated faithfully: within the nonlinear integrated equation, persistent high inflation is no longer a cyclical problem but *an acute structural decay event* — a destabilising catalyst released into a system that has become

Table 18.1: Evidential standing of every channel Part III uses, from `registry/scorecard.csv`. This table is the reference for the notes that follow, and it is worth reading before the argument rather than after it.

Channel	n	Verdicts	Standing
σ_{sys}	9	6 hit, 1 partial, 2 miss	best-attested channel in the book
S_{swap}	3	2 hit, 1 miss	**the only lever with hits**
D_{trans}	5	1 hit, 2 partial, 2 miss	mixed
D_{net}	2	1 hit, 1 miss	the quantity grew; the dollar did not respond
S_{Fed}	1	1 partial	one observation
S_{coll}	3	1 partial, 2 untestable	proxy measures the wrong thing
C_{alt}	2	1 partial, 1 untestable	
T_{frag}	5	1 partial, 4 untestable	no free machine-readable series
$\Psi_{\text{str}}(t)$	5	1 partial, 1 imposed, 3 untestable	
R_{US}	5	**1 miss, 4 untestable**	**no hits**
S_{off}	4	1 imposed, 1 recorded, 1 unmodelled, 1 untestable	**no scoreable verdict**
M_{trust}	1	**1 untestable**	**one prediction, no verdict**
Φ_{sys}	3	3 untestable	
$s_{\text{US,sys}}$	2	2 untestable	
S_{bank}	2	2 untestable	
S_{Tsy}	1	1 untestable	
κ_{sys}	4	**4 untestable**	**no verdicts; unidentifiable by Proposition 11.1**
H_{soc}, V_{meme}, I_{imp}	0	—	**no predictions scored at all**
D_{hcap}	0	—	**never appears in a case study**

psychologically brittle. The failure to preserve purchasing power is a severe rise in R_{US} and, in the original's words, "an existential collapse" of M_{trust}.

Under the psycho-technical layer the erosion of M_{trust} is no longer slow or linear. It is *viral*: amplified by $H_{\text{soc}} \to 1$ and accelerated by $V_{\text{meme}} \uparrow$, the story of the central bank's failure corrodes Φ_{sys} and $s_{\text{US,sys}}$ directly, and risks triggering $\Psi_{\text{str}}(t)$'s tipping point "decades earlier than expected".

That is a strong claim and it is the load-bearing one for the whole of Part III. It is also the claim Section 18.6 tests, on the only episode available.

18.2 Inflation as Structural Failure: $R_{\text{US}} \uparrow$, $M_{\text{trust}} \downarrow$ within H_{soc}

Evidential standing of this entire section. The argument runs through M_{trust} and R_{US}. Both sit in the structural block, which Table 17.3 scores as **78 per cent untestable**. Specifically: M_{trust} has **one** prediction in the whole scorecard and it is untestable; R_{US} has five, of which **one is a miss and four are untestable**, so it has never been scored right. H_{soc} and V_{meme}, which do the amplifying, have **no scored predictions at all**. **This section therefore reconstructs an argument that the book's evidence can neither support nor refute**, and it is written as reconstruction rather than as finding.

18.2.1 The trust channel

Traditional models, the original argues, failed by treating inflation as a mechanical contest between S_{Fed} and D_{cap}. This model situates it in the modulator layer, where under Layer 4 it becomes self-amplifying.

The mechanism is stated in three steps. A central bank's core mandate is to protect the currency's value. Failure generates a simple, transmissible meme — "the Fed has lost control", "the Fed printed too much money". Previously such memes propagated slowly, through academic papers and financial reports. Now, driven by A_{alg} (algorithms optimising for outrage) and P_{AI} (generative systems producing endless apparent evidence), the meme achieves $V_{\text{meme}} \uparrow\uparrow$ and infects collective consciousness almost instantly. The original's conclusion: **institutional trust collapses within months instead of years**.

Remark 18.1 (The claim has a testable form, and this book can nearly test it)**.** "Trust collapses within months" is not vague. It implies that a large inflation shock should move the market's expectation of *long-run* inflation — what inflation will be after the current episode is over — because that expectation is exactly the price of institutional credibility — the object Sargent [67] showed can be restored abruptly when a fiscal regime changes, and whose modern history is told by Bernanke [14] and Blinder [15]. M_{trust} has no measure (Chapter 11), but the five-year, five-year-forward inflation compensation is the closest observable substitute that exists, and it is free and daily. Section 18.6 uses it. **It is a substitute and not a measure**, and the conclusion drawn there is stated with that limit attached.

18.2.2 The risk channel

R_{US} rises, and the original's contribution is that it stops rising *linearly*. Inflation breeds discontent; $H_{\text{soc}} \to 1$ and A_{alg}'s echo chambers aggregate fragmented frustrations into unified movements; so R_{US} becomes *volatile* rather than merely elevated. Minor events — the original's example is fuel protests — can escalate through H_{soc} amplification into systemic crises that shock Φ_{sys} directly.

> **Standing.** R_{US}'s one scored verdict is a *miss*: CS10 predicted that R_{US} would rise on an exercise of coercive power and it did not (Section 16.1). The other four are untestable. The volatility claim specifically — that R_{US}'s variance rises with H_{soc} — requires a series for R_{US} and a series for H_{soc}, and the book has neither.

18.2.3 The dual-erosion effect

The original calls this "among the most dangerous findings" and it deserves the emphasis, because it is the one place in Chapter 7 where the psycho-technical layer reaches down into the mechanical layer rather than only modulating it.

The argument: the explosion of R_{US} attacks the system twice. Indirectly, by eroding Φ_{sys}. And *directly*, at Layer 0, by destroying the human-capital demand pillar D_{hcap}. America's talent magnet ceases to function — the world's most capable people do not migrate to a country with runaway inflation, unaffordable costs and unstable domestic risk. So inflation in the H_{soc} era causes **dual erosion**: it destroys trust and demand simultaneously.

> **Standing.** D_{hcap} **does not appear in a single case study in this book**. It has no rows in `registry/scorecard.csv`, no proxy in `data/processed`, and Chapter 3 classes its measurement as annual and indirect. The dual-erosion effect is therefore **entirely unevidenced here** — not refuted, not supported, not examined. It is recorded because it is the original's claim and because it is a good hypothesis; a reader should not mistake "interesting" for "established".

18.3 The Dilemma: Trust versus Stress, Both Amplified

With Layer 4 active, the original argues, *both* policy choices become worse than they were.

Option A — inaction, ignoring inflation to protect growth. The old structural outcome was that M_{trust} evaporates, $\Phi_{\text{sys}} \downarrow\downarrow$, and $\Psi_{\text{str}}(t)$ is triggered. The updated outcome is worse because the evaporation is accelerated by V_{meme}: the "central bank failure" narrative becomes globally dominant, and it supplies psychological fuel for other nations to adopt C_{alt}. Collective acceptance drives $\Psi_{\text{str}}(t)$ rapidly toward its tipping point. **Option A is not a slow death but an accelerated one.**

Option B — action, fighting inflation. The old short-term outcome was σ_{sys} ↑↑↑: a liquidity shock, D_{debt} ↑, S_{coll} ↓. The updated outcome is worse because that stress is now multiplied by κ_{sys} under H_{soc}. As homogeneity rises the diversity of agents vanishes and everyone panics simultaneously. The original's illustration is concrete: **a 0.75 per cent rate rise that once caused a five per cent correction may now cause a twenty per cent collapse.**

Table 18.2: The dilemma as the original states it, with each arm's evidential standing.

Option	The claim	Standing
A: inaction	M_{trust} ↓ accelerated by V_{meme}; C_{alt} adoption; $\Psi_{\text{str}}(t)$ tipping	M_{trust} and V_{meme} unscored; $\Psi_{\text{str}}(t)$'s cliff is unfalsifiable (Proposition 9.3). **Unevidenced**
B: action	σ_{sys} ↑↑↑ multiplied by $\kappa_{\text{sys}}(H_{\text{soc}})$; a 5% correction becomes 20%	σ_{sys} is well attested; the *multiplier* is not — κ_{sys} has four predictions, no verdicts, and Proposition 11.1 proves its decomposition is unidentifiable. **Half-evidenced: the shock yes, the amplification no**

The asymmetry in that table is the honest summary of the dilemma. **The stress side of Option B is the one part of this chapter's machinery the book has actually measured.** The amplification that makes it a dilemma rather than a trade-off is not.

What makes it a dilemma rather than a trade-off. The distinction is worth stating precisely, because the original relies on it and does not spell it out. An ordinary policy trade-off has an interior optimum: tighten somewhat, accept some stress, and the marginal costs balance. The original's claim is stronger — that *both* arms are *convex in the same direction*, so the cost function has no interior minimum worth choosing. Option A's cost accelerates because V_{meme} compounds; Option B's cost accelerates because κ_{sys} multiplies. If both are true, the policymaker is choosing between two accelerating losses rather than trading one linear cost against another, and the only escape is to act on a third margin — which is what the four spears are.

Remark 18.2 (The dilemma's structure depends on a convexity the book has examined)**.** Option B's convexity is λ_2, and Chapter 5 tested it: the response to stress does steepen in stressed states, in a majority of specifications, conditional on w_X. So the Option B arm has support. But ISSUES #37 found the convexity *misplaced* — largest in the middle of the σ_{sys} range and smallest at the values crises occupy — which, applied here, means the amplification is weakest exactly when the policymaker would be tightening into an already-stressed system. **That does not dissolve the dilemma; it makes Option B cheaper at high stress than the original claims.** It is a case where a finding against the model's specification is favourable to the model's policy conclusion, and those come apart more often than is comfortable.

18.3.1 The one number in the dilemma, and what it would require

Option B's statement is unusually specific for this part of the original: *a* 0.75 *per cent rate rise that once caused a five per cent correction may now cause a twenty per cent collapse.* That is a claimed amplification factor of **four**, and a claim with a number in it can be checked against something.

The check available is imperfect and the imperfection is stated first. The original's illustration is about an equity correction; this book's measured output is the dollar. What follows therefore tests the amplification on *the model's own output variable* rather than on the asset the sentence names. It is the closest test this book can run, and it is a substitute rather than the thing itself.

Table 18.3: Realised daily dollar volatility by regime, monthly, 2006–2026. Regime is set by σ_{sys} at the start of the month, on the threshold of 0.5 used throughout Chapter 5. Computed in `data/policy_evidence.py`.

Regime	Months	Mean daily s.d.	
Calm, $\sigma_{\text{sys}} < 0.5$	127	0.323	
Stressed, $\sigma_{\text{sys}} \geq 0.5$	120	0.465	
Ratio		**1.44**	**the claim requires** 4.00

Proposition 18.1 (The stated amplification is roughly three times what the dollar shows)**.** *Realised daily volatility of* `DTWEXAFEGS` *averages* 0.465 *per cent in months beginning in the stressed regime against* 0.323 *in months beginning calm, a ratio of* **1.44**. *The original's illustration requires an amplification of* 4.00. *Further, the largest single monthly dollar move in the whole* 2006–2026 *sample is* +8.8 *per cent, in October 2008, and the largest fall is* −7.0 *per cent, in May 2009.* ***A twenty per cent move has no precedent in the sample at any stress level.***

Proposition 18.1 is the only place in this book where a number the original states outright can be checked against the model's own output, which is why it is worth the space.

Remark 18.3 (Three readings, and the book cannot choose between them)**. One:** the amplification is real for equities and not for the dollar, in which case the illustration is fine and simply does not transfer to this model's output. **Two:** the amplification is a forecast about a state the system has not yet reached — H_{soc} has risen since 2008 and the sample's stressed months are mostly the old regime — in which case it is untestable now and will be testable later. **Three:** the figure is rhetorical. Nothing here separates them, because separating them requires H_{soc}.

What can be said is that the number is a factor of three away from the closest thing the book can measure, in the direction of overstatement, and that a policy argument resting on "a five per cent correction becomes twenty" should carry that. Chapter 5's regime slopes point the same way: stressed and calm slopes under the canonical construction are +1.89 and +1.73, a ratio of 1.09, not 4.

18.4 The Reflexivity Trap

Stop-go policy in this environment, the original argues, is policy suicide, and the mechanism is a loop rather than a mistake.

1. **Action.** The central bank tightens, $S_{\text{Fed}} \downarrow$.
2. **Effect.** Stress rises, and is violently amplified by $\kappa_{\text{sys}}(H_{\text{soc}})$: a small crack looks like an abyss.
3. **Reflex.** Policymakers are themselves inside H_{soc} and exposed to V_{meme}. They catch the crowd's panic.
4. **Pivot.** They loosen.
5. **Outcome.** The world, unified by H_{soc}, confirms the worst meme: the central bank is not merely powerless but *frightened*. $\Phi_{\text{sys}} \downarrow\downarrow$ and $\kappa_{\text{sys}} \uparrow\uparrow\uparrow$ together.

Remark 18.4 (The one step here that is genuinely novel)**.** Step 3 is the original's real contribution and it is not a standard result. Conventional credibility models treat the policymaker as outside the belief system being modelled — an agent choosing a rule while the public forms expectations about it. **The original puts the policymaker inside** H_{soc}, subject to the same homogeneity and the same narrative velocity as everyone else. Chapter 8's fixed-point machinery is what makes that statable: if the policymaker is one of the agents in the belief map, the equilibrium includes them, and the map's slope T' governs their reaction as much as the market's. This is a proper implication of the architecture and it survives Chapter 17 because nothing in Chapter 17 touches it — Layer 4 has no scored predictions in either direction.

18.5 The Four-Spear Strategy

The original's escape from the trap requires a coordinated multi-tier strategy. Three spears were in the earlier framework; the psycho-technical layer forces a fourth.

The original's justification for the fourth spear is worth quoting in substance because it is the sharpest sentence in Chapter 7: *in a system dominated by* H_{soc} *and* V_{meme}*, silence equals failure.* Policymakers must communicate with extreme clarity, simplicity and repetition, to provide a psychological anchor for collective consciousness.

Remark 18.5 (Spear 4 is the one recommendation with no evidential support of any kind)**.** It is also the one most likely to be right, and those two facts sit uncomfortably together. Every variable it operates on — H_{soc}, V_{meme}, B_{phase}, κ_{sys} — is unmeasured, and Remark 10.2 shows Layer 4 has *no inbound edge from any other layer*, so within the model as specified there is no channel by which communication reaches H_{soc} at all. **Spear 4 is not merely unevidenced; it acts on a part of the model that nothing else in the model can reach.** ISSUES #32 records the missing arrow and declines to add it, because doing so would introduce an eighth unidentified coefficient. That decision has a cost, and this is where the cost falls.

Table 18.4: The four spears, with the channel each runs through and its standing. Chapter 19 decomposes each into levers.

Spear	Hand	Action	Standing of the channel
1	Right hand	Tighten S_{Fed} decisively and *maintain* it, to save M_{trust}	S_{Fed}: one partial. M_{trust}: one untestable
2	Left hand	Open $S_{\text{swap}} \uparrow$, $S_{\text{coll}} \uparrow$ to extinguish stress before $\kappa_{\text{sys}}(H_{\text{soc}})$ amplifies it	S_{swap}: **2 hits, 1 miss — the best-attested lever**. S_{coll}: 1 partial, 2 untestable
3	Root	Address supply shocks and domestic instability: D_{trans}, R_{US}	D_{trans}: mixed. R_{US}: no hits
4	Narrative	Proactive narrative management — clarity, simplicity, repetition, to dominate the narrative space and dampen B_{phase}	**no scored predictions anywhere**

18.5.1 What each spear would need to become checkable

Stated because Chapter 17's research agenda is organised the same way, and because a reader who wants to take this strategy seriously should know what would have to exist first.

Two of the four are bounded data tasks and two are research programmes. That division matters more than it looks: Spears 2 and 3 could be made checkable by work, and Spears 1 and 4 could not. The strategy's two unfalsifiable legs are the two the original considers decisive.

18.6 The 2021–2023 Episode: The Chapter's Own Argument, Tested

This section is not in the original, which was written before the episode could be assessed. It is here because a chapter arguing that inflation destroys institutional trust within months has an obligation to look at the largest inflation shock in forty years, which happened, and which is over.

18.6.1 What the Right Hand did

525 basis points in seventeen months, the fastest cycle since the early 1980s. Spear 1 says to tighten decisively and maintain; that is what happened, and the rate was held above

Table 18.5: The measurement each spear is waiting on. "Tier" follows ISSUES #23.

Spear	What would have to be measured	Tier	Why it is not here
1	M_{trust}, separately from inflation expectations	—	no candidate series exists; survey-based trust indices measure something adjacent
2	the cross-currency basis, for D_{debt}	3	the single most valuable missing series in the book (Remark 11.4); commercially licensed
3	T_{frag} from SWIFT shares	2	free but PDF-only across roughly 180 documents; transcription, about an afternoon per year
4	H_{soc}, V_{meme}, B_{phase}	—	Chapter 8 proposes constructions for H_{soc} and C_{src}; none has been built, and A_{alg} and R_{div} have no route at all

Table 18.6: The 2022–23 tightening cycle. Sources: FRED `DFF`, `T5YIFR`, `MICH`; script `data/policy_evidence.py`.

Date	Event	Policy rate	5y5y fwd	Michigan 1y
2021-12-31	before the cycle	0.07	2.25	4.8
2022-03-16	first hike	0.08	2.26	5.4
2022-06-15	first 75bp move	0.83	2.40	5.3
2022-11-02	fourth consecutive 75bp	3.08	2.38	5.0
2023-07-26	final hike	5.08	2.48	3.4
2024-09-18	first cut	5.33	2.26	2.7

five per cent for a year before the first cut. As an execution of the original's Right Hand this is close to the specification.

18.6.2 Did trust collapse?

The five-year, five-year-forward inflation compensation is what the market charges for inflation *after* the current episode has passed. If institutional credibility were collapsing within months, this is where it would show.

Table 18.7: Five-year, five-year-forward inflation compensation. The comparison that matters is the second column.

Period	Mean	**Std. dev.**	Range
2015–2019, calm	1.95	**0.200**	1.41–2.35
2021–2026, the episode	2.25	**0.106**	1.91–2.67

Proposition 18.2 (Long-run inflation expectations were *more* stable during the inflation shock than before it)**.** *Over 2021–2026, encompassing the highest measured inflation in forty years and the fastest tightening cycle since 1981, five-year five-year-forward inflation compensation had a standard deviation of* 0.106 *percentage points, against* 0.200 *over the calm years 2015–2019. Its maximum was* 2.67 *and its value on the day headline inflation peaked was* 2.08*.* ***The measure was half as volatile during the episode as before it****, and never left a range consistent with the stated two per cent objective.*

Proof. Computed in `data/policy_evidence.py` from `T5YIFR`, daily, on the stated windows. Standard deviations are population standard deviations over all trading days in each window. □

What this does and does not establish. It does *not* refute the original's mechanism, and Remark 18.1 said in advance why: forward inflation compensation is a substitute for M_{trust}, not a measure of it. It contains a liquidity and term premium, it prices the market's expectation rather than the public's belief, and H_{soc} is a claim about *households and narratives* at least as much as about the bond market. A defender of the original can say, correctly, that the meme spread and the bond market ignored it.

What it does establish is narrower and still substantial. **The original's claim is that trust collapses in months, and the one large inflation episode available produced no collapse in the only credibility-linked series that exists at daily frequency.** The household series is not much help to the other side either: Michigan one-year expectations peaked at 6.6 per cent, but they peaked in **May 2025** — during the tariff episode of Section 16.2 — not during the inflation. The two candidate measures disagree about when trust was most strained, which is itself a finding about the state of M_{trust} measurement.

Remark 18.6 (The asymmetry a reader should notice)**.** Spear 1 was executed nearly to specification and the predicted trust collapse did not occur. That is compatible with two

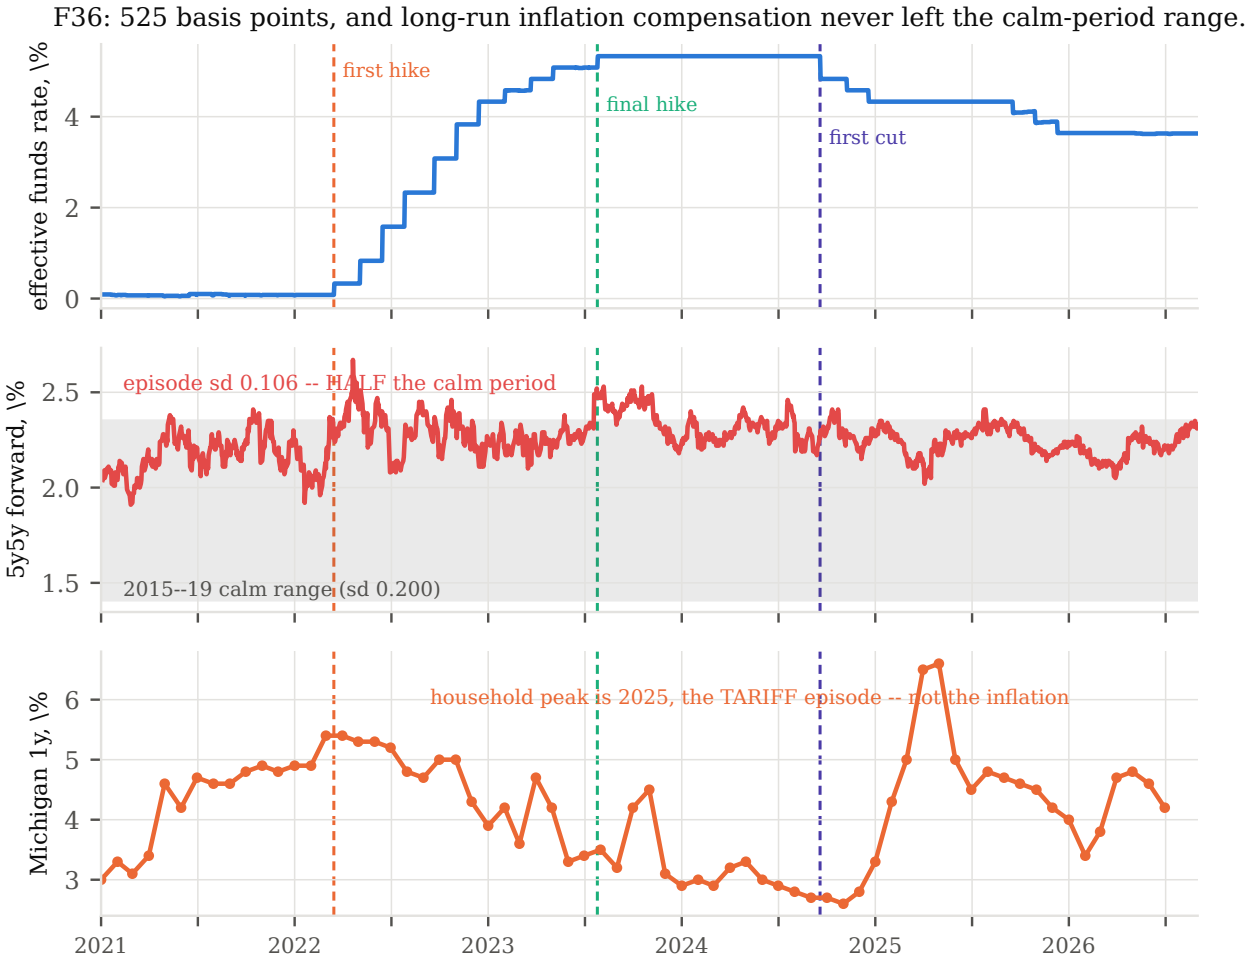


Figure 18.1: The fastest tightening cycle since 1981, and long-run inflation compensation never leaving the range it occupied in the calm years before it. The shaded band in the middle panel is the 2015–19 range. The bottom panel is the awkward one for both sides of the argument: household expectations peaked in May 2025, during the tariff episode, not during the inflation. Script: `figures/src/f36_anchoring.py`.

readings: the strategy worked, or the threat was overstated. This book cannot separate them, because separating them requires a measure of M_{trust} and a counterfactual, and it has neither. **What it can say is that the episode gives no positive evidence for the "viral collapse" premise on which the whole chapter rests**, and that a policy argument whose premise has never been observed should be read as a hypothesis about what might happen rather than as a description of what does.

18.6.3 And the Left Hand was barely used

This matters for Chapter 19 and it is stated here because it belongs with the episode.

Table 18.8: Federal Reserve central bank liquidity swaps outstanding, peak and mean, $ millions. Source: FRED `SWPT`.

Episode	Peak	Mean	
2008 crisis	583,135	348,972	2008-12-17
2020 crisis	448,946	159,811	2020-05-27
2022–23 tightening	**11,302**	**503**	2022-10-26
2023 bank stress	587	424	2023-03-22

The playbook's central instruction is that the two hands move *together*: the Right Hand induces stress deliberately and the Left Hand opens the valves so that κ_{sys} has nothing to amplify. In the one episode where the Right Hand moved hard, the Left Hand was used at **two per cent** of its 2020 peak and **one-fiftieth** of its 2008 peak. The coordinated strategy the original specifies has never been run.

Remark 18.7 (Which makes Part III a design document, not a description)**.** Neither the four-spear strategy nor the two-handed cadence of Chapter 20 has an implementation to point at. Chapter 19 is written accordingly: it documents what each lever *is* — statute, mechanics, dated activations — because that is verifiable, and labels what the model says each lever *does* as conditional, because that is not.

Chapter 19

Strategic Decomposition and Operational Playbook

Assumption 18.1 governs this chapter completely. Each lever below is documented in four parts: its **statutory basis**, its **mechanics**, its **dated activations** with sizes, and **what the model says it does**. The first three are institutional fact and are true regardless of anything in Chapter 17. The fourth is conditional, and each carries an evidential-standing note at the point it is made.

Table 19.1 is the chapter in one page, and a reader who reads nothing else should read the last column.

Two of eleven levers run through a channel the book has ever scored a hit on. That is the honest summary of this chapter's standing, and it is stated at the front rather than the back.

19.1 The Right Hand: Tightening to Rescue M_{trust}

Objective, as the original states it: reassert the central bank's commitment to protecting the currency's value, regardless of short-term economic pain — a necessary sacrifice to prevent the collapse of Φ_{sys} and $s_{\text{US,sys}}$.

19.1.1 Lever 1: the policy rate (S_{Fed})

Statutory basis. The Federal Open Market Committee directs open market operations under §12A and §14 of the Federal Reserve Act. Since 2008 the operational instrument has been administered rates — interest on reserve balances, and the overnight reverse repurchase offering rate — rather than reserve scarcity, so the target range is enforced by arbitrage rather than by quantity. The two per cent objective is not statutory: the mandate in §2A is "maximum employment, stable prices, and moderate long-term interest rates", and the numerical target is the Committee's own, set out in its Statement on Longer-Run Goals and Monetary Policy Strategy, first adopted in January 2012 and revised in August 2020.

Table 19.1: Every lever, its channel, and what Chapter 17 says about that channel. "Hits" counts scored predictions in `registry/scorecard.csv`.

	Lever	Channel	Scorecard	Read the model's claim as
Right hand — tightening to rescue M_{trust}				
1	Policy rate	S_{Fed}	1 partial	one observation
2	Communication	M_{trust}	1 untestable	**unevidenced**
3	Fiscal ceasefire	R_{US}	1 miss, 4 untestable	**no hits, one miss**
Left hand — opening the pipes to quench σ_{sys}				
4	Swap lines	S_{swap}	**2 hits**, 1 miss	**the one attested lever**
5	FIMA repo	S_{coll}	1 partial, 2 untestable	proxy measures the wrong thing
6	Bill issuance	S_{coll}	as above	as above
7	SLR relief	S_{bank}	2 untestable	**unevidenced**
8	TGA and RRP	S_{Tsy}	1 untestable	**unevidenced**
Root layer — reducing pressure on the right hand				
9	Energy and commodities	D_{trans}	1 hit, 2 partial, 2 miss	mixed
10	Coercion restraint	T_{frag}, S_{off}	1 partial, 5 untestable, 1 imposed	**unevidenced**
11	Talent and IP	D_{hcap}	**no rows at all**	**never examined**

Mechanics. The Committee sets a target range; the Desk enforces it with IORB as the upper anchor and the ON RRP offering rate as the lower. Balance-sheet policy runs alongside: caps on principal reinvestment determine the pace of runoff. The lever is therefore *two* instruments that can move in opposite directions, which matters for Section 20.2's fourth trap.

Dated activations. The 2022–23 cycle is the reference case: first increase 2022-03-16, four consecutive 75 basis-point moves from June to November 2022, final increase 2023-07-26, peak effective rate 5.33 per cent, first reduction 2024-09-18. **525 basis points in seventeen months**, the fastest since 1981.

What the model says it does. The original's instruction is precise: raise the rate above expected inflation so the *real* rate is restrictive, $r_{\text{real}} > 0$, and hold it until actual inflation slows decisively — not until expectations do. Its reasoning is that any stance leaving $r_{\text{real}} \leq 0$ is read by a homogenised market as "not serious", and M_{trust} collapses.

> **Evidential standing.** S_{Fed} has exactly one scored prediction, a *partial* (CS3, Section 13.3). The claim that $r_{\text{real}} \leq 0$ triggers a trust collapse depends on M_{trust}, which has one prediction and no verdict, and on H_{soc}, which has none. The 2022–23 cycle ran with $r_{\text{real}} < 0$ for its first fourteen months and produced no measurable loss of long-run anchoring (Proposition 18.2), which is the one piece of evidence available and it does not support the mechanism.

19.1.2 Lever 2: communication (M_{trust}, Spear 4)

Statutory basis. None required, and that is the point: this is the only lever in the chapter that needs no authority, no counterparty and no appropriation. The reporting obligations in §2B of the Federal Reserve Act — semiannual testimony and the Monetary Policy Report — are a floor, not a constraint.

Mechanics. Post-meeting statement, the Summary of Economic Projections and its dot plot, the press conference, minutes at three weeks, speeches, and the Statement on Longer-Run Goals. The instrument is repetition and consistency across all of them.

Dated activations. The clearest modern statement of the lever's intent is Powell [61], eight minutes long and deliberately so. Forward guidance as an explicit tool dates to the August 2003 "considerable period" language; calendar-based guidance to August 2011; threshold-based guidance to December 2012. The August 2020 framework revision introduced flexible average inflation targeting, which is itself a communication instrument.

What the model says it does. The original's argument is that this is no longer passive communication but an information war. Clear, repetitive communication provides a cognitive anchor for collective consciousness, breaking behavioural phase-lock B_{phase}, and *directly lowers* κ_{sys} by removing the ambiguity that algorithms amplify.

Evidential standing — the weakest in the chapter. Every variable this lever operates on is unmeasured. M_{trust}: one untestable prediction. H_{soc}, V_{meme}, B_{phase}, A_{alg}: no scored predictions anywhere in the book. κ_{sys}: four predictions, four untestable, and Proposition 11.1 proves its decomposition cannot be identified even in principle. Worse, Remark 10.2 establishes that **Layer 4 has no inbound edge from any other layer**, so within the model as specified there is no channel through which communication reaches H_{soc} at all (ISSUES #32). The lever may well work. The model cannot say that it does, and does not contain the arrow by which it would.

19.1.3 Lever 3: fiscal coordination (R_{US})

Statutory basis. The debt limit is set in 31 U.S.C. §3101 and can be raised, suspended or redefined only by statute. Appropriations lapse under the Antideficiency Act absent enacted appropriations or a continuing resolution. Neither is within the central bank's control, which is the whole difficulty: this lever is held by a body that does not answer to the institution the strategy asks to use it.

Mechanics. When the limit binds, Treasury deploys "extraordinary measures" — suspending reinvestment in the G-Fund and the Civil Service Retirement Fund, among others — which buy weeks, and then runs down the Treasury General Account until the X-date. Bill issuance is curtailed, which tightens exactly the collateral supply Lever 6 exists to expand.

Dated activations. 2011: resolution 2011-08-02, S&P downgrade 2011-08-05. 2013: October shutdown and near-miss. 2023: agreement 2023-06-03, Fitch downgrade 2023-08-01. Section 13.4 scores both 2011 and 2023 as CS4.

What the model says it does. The original argues coordination is exponentially more urgent under high H_{soc}, because a debt-ceiling fight would be amplified by V_{meme} into an informational cascade — "US government collapse" — shocking Φ_{sys} immediately.

Evidential standing. R_{US} has five scored predictions: **one miss and four untestable**. It has never been scored right. The miss is CS10's, where an exercise of coercive power was predicted to raise R_{US} and did not. CS4 is the direct test of the debt-ceiling claim and Proposition 13.2 shows why it cannot settle it: the episodes with clean identification are the ones markets judged not to matter, so they carry no signal. κ_R **is permanently unidentified**, and this lever's magnitude is exactly what κ_R would measure.

19.2 The Left Hand: Opening the Pipes to Quench σ_{sys}

Objective: while the Right Hand deliberately induces stress and κ_{sys} is at maximum amplification, the liquidity tools must open safety valves immediately. The mission is to decouple rate

policy from the plumbing's stability — in the original's phrase, "turning off the microphone before it hits the amplifier".

This is the strongest section of the original's playbook and it is also the best-attested, which is not a coincidence: these levers act on plumbing that leaves daily traces.

19.2.1 Lever 4: central bank liquidity swaps (S_{swap})

Statutory basis, and it is contested. The swap lines operate under §14 of the Federal Reserve Act, the open-market authority, pursuant to the FOMC's Authorization for Foreign Currency Operations — *not* under §13(3). That distinction is consequential: §13(3) emergency facilities require the approval of the Secretary of the Treasury under the 2010 Dodd-Frank amendments, and §14 operations do not. So the largest international liquidity instrument the United States possesses can be deployed by the FOMC alone. **Whether §14 in fact authorises them is disputed in the legal literature**, on the ground that §14 contemplates purchases and sales of cable transfers rather than reciprocal currency arrangements. The dispute is unresolved and is recorded here because a lever whose authority is contested is a lever with a political failure mode.

Mechanics. Reciprocal currency arrangements: the Federal Reserve delivers dollars to a foreign central bank against that bank's currency at the prevailing spot rate, with a simultaneous agreement to reverse at the same rate on a stated date. The foreign central bank bears the credit risk of its own counterparties; the Federal Reserve's counterparty is the central bank itself. The pricing is a spread over OIS, and the spread is the policy variable.

Dated activations. Established with the ECB and SNB 2007-12-12; expanded through 2008, with unlimited-size arrangements for the ECB, BoE, BoJ and SNB from 2008-10-13; peak outstanding **$583.1 billion on 2008-12-17**. Made **standing arrangements in October 2013** with five counterparties — Bank of Canada, Bank of England, Bank of Japan, ECB, Swiss National Bank. In 2020: pricing cut to OIS +25 basis points and an 84-day tenor added on 2020-03-15; **nine additional temporary lines on 2020-03-19** — Australia, Brazil, Denmark, Korea, Mexico, New Zealand, Norway, Singapore, Sweden — at $60 billion each for six and $30 billion each for three; peak outstanding **$448.9 billion on 2020-05-27**.

What the model says it does. As the Right Hand drives D_{debt} up and creates a global dollar squeeze, $S_{\text{swap}} \uparrow$ suppresses the shortage and prevents system-wide collapse. That the facility does compress offshore dollar funding costs is measured by Goldberg et al. [40] for the 2007–09 episode and by Bahaj and Reis [8] in general.

> **Evidential standing — the best in the chapter, and it is qualified.** S_{swap} has **two hits and one miss**, the only lever with any hits. CS1 and CS2 both score the swap expansion as a hit: it was activated, at scale, in the two episodes the model designates as crises. The miss is CS9, where the model predicted swap expansion during the 2022–23 energy embargo and outstandings *fell* by two thirds. And Table 18.8 records the larger problem for the playbook: through

the entire 2022–23 tightening the peak was **$11.3 billion**, against $449 billion in 2020. **The lever works and was not used in the episode this chapter is written about.**

F37: the Left Hand was used in both crises and not in the tightening the playbook describes.

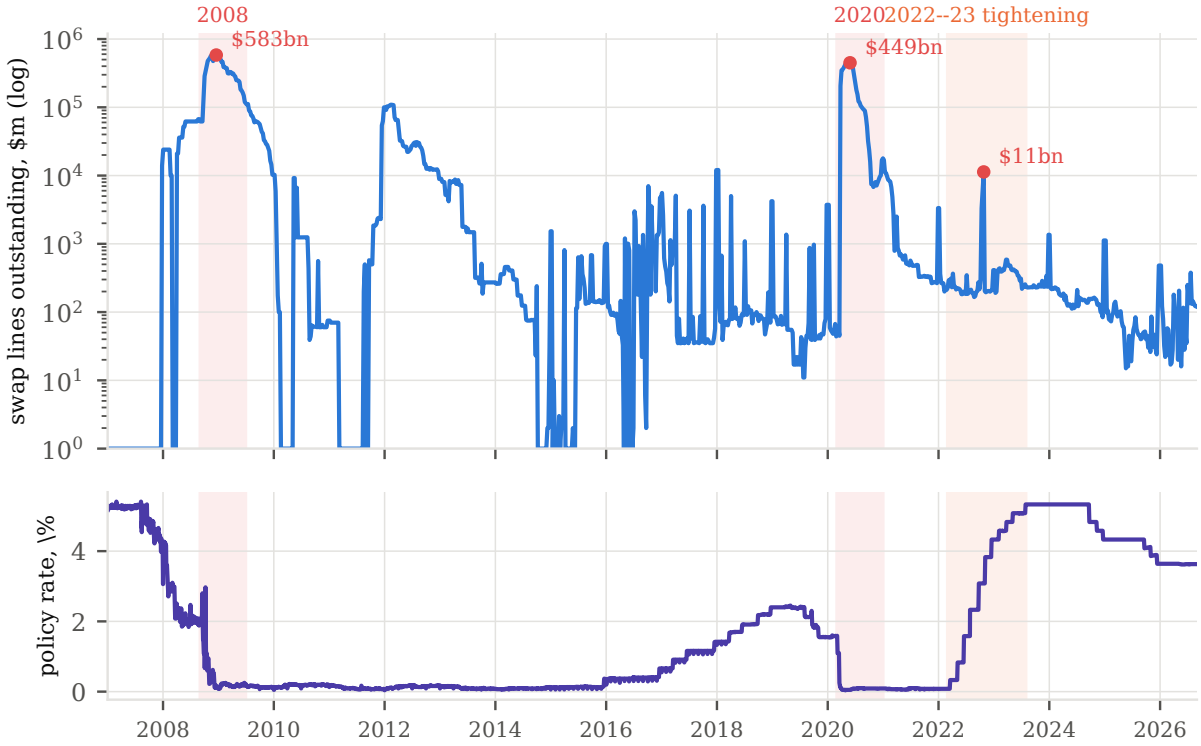


Figure 19.1: Swap-line outstandings on a log scale against the policy rate. The playbook's central instruction is that the two hands move together. In 2008 and 2020 the Left Hand opened to $583 and $449 billion. Through the 2022–23 tightening — the episode this part is written about — it peaked at $11 billion. Script: `figures/src/f37_lefthand.py`.

19.2.2 Lever 5: the FIMA repo facility

Statutory basis. Repurchase agreements with foreign official institutions against Treasury collateral held in custody at the Federal Reserve Bank of New York, under the same §14 open-market authority.

Mechanics. A foreign central bank holding Treasuries in custody at the New York Fed can borrow dollars overnight against them rather than selling them. The point is substitution: it converts a forced sale into a secured borrowing, so that a foreign reserve manager's dollar need does not become a Treasury market supply shock.

Dated activations. Announced **2020-03-31** as a temporary facility. Made a **standing facility on 2021-07-28**, alongside the domestic Standing Repo Facility, at an initial rate of 25 basis points with a $60 billion per-counterparty limit.

What the model says it does. Prevents Tier-2 holders — the original names China and Saudi Arabia — from dumping bills to obtain dollars, which would disrupt S_{coll}. They can repo quietly instead of selling.

> **Evidential standing.** The channel is S_{coll}: one partial, two untestable. And there is a measurement problem specific to this book, recorded in Section 11.5.2 and again in Section 14.2.9: the book's measured S_{coll} leg is the *overnight reverse repo facility*, which measures cash placed at the Federal Reserve, not collateral scarcity in the market. **The variable has a proxy and the proxy measures something else.** No case study has tested this lever.

19.2.3 Lever 6: bill issuance (S_{coll}, Treasury)

Statutory basis. The Secretary of the Treasury issues debt under 31 U.S.C. §3121 and chooses the maturity composition; there is no statutory bill-to-coupon ratio. Practice is governed by the quarterly refunding process and by advice from the Treasury Borrowing Advisory Committee, neither of which binds.

Mechanics. Bills are the collateral of the repo market. Increasing the bill share raises the stock of the highest-quality short-dated collateral, which supports repo intermediation and money-fund demand; reducing it does the reverse. The instrument is therefore a liquidity lever operated by an institution whose stated objective is lowest-cost-over-time financing, not financial stability.

Dated activations. The bill share is adjusted every quarter. The episode most often cited is the 2023 post-debt-ceiling rebuild, when Treasury rebuilt the General Account predominantly through bills while the ON RRP facility drained — absorbing the issuance without draining bank reserves.

What the model says it does. When $\sigma_{\text{sys}} \uparrow$ and $D_{\text{prec}} \uparrow$, bills become critical collateral. If scarce, $S_{\text{coll}} \downarrow$, the repo market freezes, σ_{sys} explodes, $\kappa_{\text{sys}}(H_{\text{soc}})$ amplifies, and the system collapses. Increasing bill supply is the lubricant.

> **Evidential standing.** Same channel and same proxy problem as Lever 5. In addition, the causal chain stated here passes through κ_{sys} at its last step, and κ_{sys} has no verdicts. **The first half of the chain — bill scarcity to repo stress — is standard market plumbing and does not need this model. The second half is the part the model adds, and it is the part with no evidence.**

19.2.4 Lever 7: leverage-ratio relief (S_{bank})

Statutory basis. The supplementary leverage ratio is a regulatory capital requirement in the Board's Regulation Q, applying to large banking organisations; the enhanced version applies to the largest holding companies. Modification is by rulemaking, which the Board can do on an interim final basis without notice and comment where it finds good cause.

Mechanics. The SLR is a non-risk-weighted ratio of tier 1 capital to total leverage exposure. Because Treasuries and reserves carry no risk weight but do enter total leverage exposure, a flight to quality that floods banks with deposits and reserves consumes leverage capacity without adding measured risk — so the constraint binds exactly when intermediation is most needed.

Dated activations. The Board announced a temporary exclusion of Treasury securities and deposits at Federal Reserve Banks from total leverage exposure on **2020-04-01**, effective through **2021-03-31**, and did not extend it. The Board's own estimate, on 2019-Q4 data, was that the change would temporarily reduce binding tier 1 requirements by roughly **$17 billion**.

What the model says it does. When $D_{\text{prec}} \uparrow$, cash floods banks; SLR prevents them absorbing it; temporary relaxation raises S_{bank}, allowing absorption of safe-haven flows and repo intermediation.

> **Evidential standing.** S_{bank} has two predictions and **both are untestable**. Table 11.10 records why: the series that would measure it — G-SIB supplementary leverage ratios and primary-dealer balance-sheet capacity — exist in FFIEC filings and have not been assembled. This is a *data* gap with a named, achievable fix, which puts it in a different class from Levers 2 and 3. The lever's mechanics above are documented fact; its model effect is untested and testable.

19.2.5 Lever 8: cash management (S_{Tsy}: TGA and ON RRP)

Statutory basis. The Treasury General Account is the government's operating account at the Federal Reserve; its balance is managed by the Treasury under its general fiscal authority. The overnight reverse repurchase facility is an FOMC instrument under §14, with counterparties including money market funds and government-sponsored enterprises.

Mechanics. Both drain reserves from the banking system when they rise. A Treasury rebuild of the TGA and a rise in ON RRP take-up are, from the banking system's point of view, the same event as balance-sheet runoff: reserves fall. Three institutions can therefore drain liquidity simultaneously without any of them intending to.

Dated activations. Through the 2022–23 tightening the TGA rose from $212 billion at end-2021 to $657 billion in June 2022, and ON RRP take-up peaked at **$2.55 trillion on 2022-12-30**. By 2026-09-01 the RRP was effectively empty at $1 billion and the TGA stood at $951 billion.

What the model says it does. Avoid a double liquidity shock that $\kappa_{\text{sys}}(H_{\text{soc}})$ could amplify into catastrophe. This is Trap 4 of Section 20.2, stated as a lever.

Evidential standing, and this one is a direct hit against the model. S_{Tsy} has one prediction and it is untestable. But the episode is informative anyway: during 2022–23 *every* condition the trap describes was met at once — the Fed raised rates 525 basis points and ran off its balance sheet, the TGA rose by $445 billion, and the RRP absorbed over $2 trillion. **The predicted catastrophe did not occur.** The RRP drain was in fact the shock absorber: it fell rather than reserves. That is a mechanism the model does not contain, and it is recorded in Section 20.2 rather than smoothed over.

19.3 The Root Layer: Reducing Pressure on the Right Hand

Objective: address non-monetary sources of inflation, so the Right Hand need not tighten to destructive levels.

19.3.1 Lever 9: energy and commodity supply (D_{trans})

Statutory basis. The Strategic Petroleum Reserve was established by the Energy Policy and Conservation Act of 1975, enacted 1975-12-22 in response to the 1973–74 embargo; drawdown authority rests with the President through the Secretary of Energy. Separately, the Defense Production Act supports domestic capacity, and tariff and export-control authorities sit with the Executive under delegated trade statutes.

Mechanics. An SPR release is a physical sale into the market, executed by competitive bid, with delivery over weeks. Its effect on price is a flow effect against a stock, so it is bounded by the size and duration of the release relative to consumption.

Dated activations. The largest in the reserve's history was announced **2022-03-31**: one million barrels a day for six months, **180 million barrels** announced, of which the Department of Energy subsequently put the figure at 172 million barrels released. **The constraint now binding is the reserve itself**: SPR holdings fell below 300 million barrels in August 2026, the lowest since 1983. A lever can be documented and simultaneously be nearly spent, and this one is.

What the model says it does. If $D_{\text{trans}} \downarrow$, price pressure eases naturally, allowing less aggressive tightening and mitigating σ_{sys} risk.

Evidential standing. D_{trans} has five predictions: one hit, two partials, two misses. It is the most-tested Layer 0 channel and the record is genuinely mixed. CS12 (Section 14.3) scores D_{trans} disruption as a hit; CS5's trade war and CS9's embargo both score it a miss, in the latter case because the oil price *fell* across an embargo. **The lever's direction is right and its magnitude is not established.**

19.3.2 Lever 10: restraint in coercion (T_{frag}, S_{off}, $\Psi_{\text{str}}(t)$)

Statutory basis. Sanctions authority runs principally through the International Emergency Economic Powers Act and its implementing executive orders, administered by the Office of Foreign Assets Control; asset immobilisation of a foreign central bank proceeds by directive under a declared national emergency. Congress can and does mandate sanctions by statute, which removes executive discretion to unwind them.

Mechanics. Primary sanctions bind US persons; secondary sanctions threaten third parties with loss of US market access. The second is the instrument that makes the dollar a weapon, and it is also the one that gives non-US parties a reason to build alternatives.

Dated activations. G7 statement on Central Bank of Russia assets 2022-02-26, US and EU measures effective 2022-02-28; approximately $300 billion immobilised. EU seaborne crude embargo and the G7 $60 price cap effective 2022-12-05; refined products 2023-02-05. OFAC reversal on Venezuelan oil 2026-01-29 (Section 16.1).

What the model says it does. This is the original's "largest structural trap": overuse of S_{off} creates the meme "the dollar is unsafe", which spreads through H_{soc} and V_{meme} and accelerates $\Psi_{\text{str}}(t)$ toward its tipping point. It must, in the original's words, be avoided at all costs.

> **Evidential standing — and here the book has something to say.** T_{frag}: one partial, four untestable. S_{off}: no scoreable verdict at all — one imposed under Rule 5, one recorded, one unmodelled, one untestable. $\Psi_{\text{str}}(t)$: one partial, one imposed, three untestable, and Proposition 9.3 makes E8 unfalsifiable. **But Section 9.7.1 is a partial vindication of this lever's premise**: measured on a reserve aggregate that includes gold rather than on COFER alone, the dollar's share is declining at an *accelerating* rate, and official gold buying tripled from 2022. The exit the lever warns about is visible once the right aggregate is used. What remains unestablished is the causal step — that coercion is what produced it — and Remark 9.2 shows the constant-price series was already accelerating before 2022.

19.3.3 Lever 11: talent and intellectual property (D_{hcap})

Statutory basis. Immigration categories and caps are set by the Immigration and Nationality Act as amended; employment-based preference numbers and per-country limits are statutory and require Congress to change. Federal research funding is appropriated annually. This lever is almost entirely legislative.

Mechanics. The channel is the stock of high-productivity human capital and the intellectual property it generates, which shows up in the balance of payments as licensing and royalty receipts and in the capital account as inward direct investment.

Dated activations. There is no discrete activation to date. The instrument is a policy stance rather than a facility, which is why it has no history of the kind the other ten levers have.

What the model says it does. It is the strategic defence against the dual-erosion effect of Section 18.2.3: while R_{US} pushes talent away, investment in D_{hcap} pulls it back, protecting a long-term demand pillar.

> **Evidential standing — the weakest possible.** D_{hcap} has **no rows in the scorecard**. It appears in no case study, has no proxy in `data/processed`, and is measured in Chapter 3 only at annual frequency and indirectly. This lever is a hypothesis attached to a variable the book never examined. It is retained because the original states it and because it is plausible; it is not evidence of anything.

19.4 Levers That Exist and the Playbook Does Not Name

The original's Chapter 7 was written in 2025 and its lever set is essentially the 2008–2020 one. Three instruments created since are absent from it, and one of them was the principal tool of the most recent episode of bank stress. They are documented here because Assumption 18.1 makes institutional completeness this part's job, and because a playbook that omits the facility actually used in 2023 is incomplete as a playbook whatever one thinks of the model.

19.4.1 The Standing Repo Facility

Basis and mechanics. Established by the FOMC on **2021-07-28**, the same day as the standing FIMA facility, under §14. It offers overnight repurchase agreements against Treasury, agency and agency-MBS collateral to primary dealers and eligible depository institutions, at an administered minimum bid rate, with an aggregate operation limit.

Why it matters to this playbook. It converts Lever 6 from a *quantity* instrument into a *price* instrument. Before the SRF, collateral scarcity had to be addressed by issuing more bills, which is Treasury's decision on Treasury's timetable. With the SRF, the Federal Reserve can put a ceiling on repo rates directly, without asking anyone. **That is a material change to the coordination problem of Table 19.2** and the original's playbook does not reflect it.

19.4.2 The Bank Term Funding Program

Basis and mechanics. Announced **2023-03-12** under §13(3), following the failures of Silicon Valley Bank and Signature Bank, and therefore *with* the approval of the Secretary of the Treasury, backstopped by up to $25 billion from the Exchange Stabilization Fund.

Advances of up to one year against Treasury, agency debt and agency-MBS collateral — valued at **par**, with margin at 100 per cent of par. No new advances after **2024-03-11**.

Why par valuation is the whole point. The 2023 problem was not credit but duration: banks held long securities at large unrealised losses, and any facility lending against *market* value would have lent too little to matter. Valuing collateral at par lends against the security's redemption value rather than its current price, which is precisely the constraint an ordinary discount-window advance imposes.

> **Model mapping, and its standing.** BTFP is a direct S_{bank} intervention: it restores balance-sheet capacity without requiring a capital-rule change, so it does what Lever 7 does by a different route and without rulemaking. S_{bank} has two predictions and **both are untestable**. The model has no term that distinguishes a par-valued advance from a market-valued one, and that distinction is the entire content of the facility.

19.4.3 The discount window, and why it is not a lever

Primary credit is available to sound depository institutions on demand against a broad collateral schedule, and it is the oldest liquidity facility the Federal Reserve has. It is listed here as an *absent* lever because of stigma: borrowing is disclosed with a lag and market participants read it as distress, so institutions avoid it precisely when it would help. The 2023 episode is the demonstration — BTFP existed in part because the window's stigma made it insufficient.

Remark 19.1 (Stigma is a Layer 4 phenomenon and the model cannot express it)**.** Discount-window stigma is a belief about what *other* agents will infer from an action — exactly the object H_{soc} was introduced to carry. A facility's capacity is not its usable capacity, and the gap is a function of collective belief. **The model has the right variable for this and no way to connect it**: S_{bank} sits in Layer 0 and H_{soc} in Layer 4, and Remark 10.2 shows Layer 4 has no inbound edge — so there is no term in which belief reduces a supply channel's effective size. This is a fifth instance of ISSUES #32's missing arrow, and it is the most concrete one, because the stigma effect is large, well documented, and cost the 2023 response a facility.

19.5 Who Actually Holds Each Lever

The playbook asks for coordinated action. The levers are held by four institutions with different mandates, different appointment structures and, in two cases, a statutory independence from each other.

Remark 19.2 (The coordination problem is the real content of this chapter)**.** The original treats coordination as a matter of will — policymakers must abandon single-mandate thinking. The institutional record says it is a matter of structure. The Fed cannot make Treasury issue bills; Treasury cannot make Congress raise the debt limit; and nothing obliges any of

Table 19.2: Institutional allocation. The last column is why coordination is hard, and it is institutional fact rather than model implication.

Holder	Levers	Constraint on using them together
FOMC	1 policy rate, 4 swaps, 5 FIMA, 8 ON RRP	Statutorily independent in monetary policy; swap lines need no Treasury sign-off under §14, unlike §13(3) facilities
Federal Reserve Board	2 communication, 7 SLR	Rulemaking on capital is a Board function, not an FOMC one — a different body from the one setting rates
Treasury	6 bill issuance, 8 TGA	Objective is lowest cost of financing over time, not financial stability; no statutory duty to consider the Fed's stance
Congress	3 debt limit and appropriations, 11 immigration and research	Cannot be committed in advance; the debt limit is a recurring scheduled hazard
Executive	9 SPR and trade, 10 sanctions	Sanctions are increasingly mandated by statute, which removes the discretion to unwind them

them to sequence their actions. In 2022–23 all three tightened at once, and Section 19.2.5 shows the predicted catastrophe did not follow — which means the coordination failure the model treats as fatal occurred, in full, without the consequence. **That is the single most important observation in this chapter and it is not favourable to the framework.**

19.6 Expected Impacts as the Original States Them

For completeness, and marked as what they are.

Table 19.3: The original's expected impacts (E16), restated. None is established by this book.

Spear	Expected impact	Status
Right hand	Short term $\sigma_{\text{sys}} \uparrow$ times $\kappa_{\text{sys}}(H_{\text{soc}}) \uparrow\uparrow \Rightarrow DXY$ rises non-linearly. Medium term $M_{\text{trust}} \uparrow \Rightarrow \Phi_{\text{sys}} \uparrow$, $s_{\text{US,sys}} \uparrow$	the short-term half is the smile claim, absent from 452 specifications (Chapter 5); the medium-term half is unmeasured
Left hand	$S_{\text{swap}} \uparrow$, $S_{\text{coll}} \uparrow$, $S_{\text{bank}} \uparrow \Rightarrow$ $\sigma_{\text{sys}} \downarrow$, $\kappa_{\text{sys}} \downarrow$, $F_{\text{liq}} \downarrow$	directionally standard; the κ_{sys} step is unidentifiable
Root layer	$D_{\text{trans}} \downarrow$ eases inflation; $T_{\text{frag}} \downarrow$ avoids the $\Psi_{\text{str}}(t)$ tipping point; $D_{\text{hcap}} \uparrow$ strengthens Φ_{sys}	one channel mixed, one unfalsifiable, one never examined

Chapter 20

Execution Rhythm, Structural Traps, and the Expected DXY Trajectory

Assumption 18.1 continues to govern. This chapter has a further advantage the other two lack: three of its four traps and all three of its trajectory phases have a dated episode to check them against, so it is the part of Part III that can be scored rather than only labelled.

20.1 Implementation Cadence: Hour, Week, Quarter, Year

The original's structure is four nested frames. Each is restated with its watchlist mapped to named series, and each watchlist entry marked with whether this book actually holds the series — because a monitoring plan whose instruments do not exist is a different object from one whose instruments do.

20.1.1 Hour–day: extinguish σ_{sys} and hold the narrative

Objective: quench liquidity fires caused by the Right Hand before $\kappa_{\text{sys}}(H_{\text{soc}})$ amplifies them into a non-linear explosion.

Priority 1, Spear 4. Before the market opens, a joint statement emphasising *decoupling*: rates are being raised to fight inflation, and liquidity is being provided without limit to keep the plumbing stable, and *these are separate actions*. The original's logic is that this dominates the narrative space and denies V_{meme} its material.

Priority 2, Left Hand. Activate swap lines with the five standing counterparties, extend to Tier-2 allies, cut spreads to zero, and publicise the FIMA facility so foreign official holders repo rather than sell.

Remark 20.1 (The watchlist is better than the book's own instrument set)**.** Three of the four series the original names for emergency monitoring are ones this book does not hold, and in two cases the substitute is known to measure something else. That is not a criticism of the

Table 20.1: The hour–day watchlist, mapped to series. "Held" means the series is in `data/processed` or `data/raw` and is used somewhere in this book.

Variable	Series named	Held?	Note
σ_{sys}	VIX (equity), MOVE (bond)	partly	VIX held; MOVE is proprietary, and Section 5.3 substitutes realised Treasury volatility
F_{liq}	FRA–OIS, LIBOR–OIS	no	LIBOR ceased publication for most tenors in 2023; the successor spread is SOFR-based and the book uses `STLFSI4` instead
D_{debt}	cross-currency basis (EUR/USD)	**no**	Tier 3 under ISSUES #23, and Remark 11.4 identifies it as the single most valuable missing series in the book
S_{coll}	bill specialness, GCF–SOFR	no	the book substitutes the reverse-repo facility, which measures something else (Section 11.5.2)

watchlist — **it is the correct list**, and the cross-currency basis in particular is exactly the instrument Chapter 17's research agenda asks for. It is a statement about the gap between what the framework specifies and what has been assembled to test it.

20.1.2 Weekly: stabilise domestic plumbing

Objective: ensure S_{coll} and S_{bank} do not rupture while the Right Hand keeps tightening. Treasury shifts issuance toward bills; the Board signals conditional, temporary capital relief; Treasury manages the TGA and RRP to avoid a simultaneous drain.

The original observes that *signalling* SLR relief already raises perceived S_{bank} capacity. That is a claim about announcement effects, and it is untestable in this book for the reason Section 19.2.4 gives: S_{bank} has no series.

20.1.3 Quarterly: restore M_{trust} and ease D_{trans}

Assess expected inflation and commodity data; if expectations remain high, continue tightening to reaffirm credibility; deploy industrial and energy policy to expand supply. The original's instruction here is a shift of attention: from liquidity indicators to structural ones — five-year five-year-forward inflation compensation, supply-chain price indices.

> **Standing.** This is the one frame whose primary indicator the book holds and uses. Proposition 18.2 is computed on exactly the series this frame names, and it finds that the indicator was *more* stable during the inflation episode than before it.

Table 20.2: The weekly watchlist. The original does not state one for this frame; this is the set implied by the actions it specifies.

Variable	Series	Held?	Note
S_{Tsy}	TGA balance, H.4.1 weekly (`WTREGEN`)	yes	used in Table 20.5
S_{Tsy}	ON RRP take-up (`RRPONTSYD`)	yes	and it is the series whose blanks caused the defect of Section 11.5.2
S_{Fed}	Reserve balances and total assets (`WALCL`)	yes	
S_{bank}	G-SIB supplementary leverage ratios	**no**	FFIEC filings; Tier 2, achievable, not assembled
S_{coll}	bill share of marketable debt	no	Treasury quarterly refunding documents

Table 20.3: The quarterly watchlist. This is the best-provisioned frame in the chapter.

Variable	Series	Held?	Note
E_{inf}	5y5y forward inflation compensation (`T5YIFR`)	**yes**	the basis of Proposition 18.2
E_{inf}	Michigan one-year expectations (`MICH`)	yes	disagrees with the above about when trust was strained
D_{trans}	Brent (`DCOILBRENTEU`)	yes	used throughout Chapter 13–16
D_{trans}	supply-chain price indices	no	not obtained

20.1.4 Annual: consolidate Φ_{sys} and $\Psi_{\text{str}}(t)$

Three actions: strengthen D_{hcap} against dual erosion; convert D_{net} into a controlled demand pillar by giving dollar stablecoins a clear legal framework rather than leaving them as unregulated S_{crypto} risk; and manage G_{pol} diplomatically rather than through excessive S_{off}, to slow $\Psi_{\text{str}}(t)$ erosion.

> **Standing, and one of these has been tested.** The D_{net} recommendation is the only annual-frame item with a scored verdict, and it is split: CS6 (Section 15.2) records that the stablecoin stock grew seventy-four-fold — a *hit* on the quantity — and that its correlation with the dollar is a measured zero, a *miss* on the channel. **Building the pillar is achievable and there is no evidence it does anything to the dollar.** D_{hcap} has never been examined. The S_{off} restraint item is Section 19.3.2's Lever 10.

Table 20.4: The annual watchlist. Two of five instruments exist.

Variable	Series	Held?	Note
D_{net}	stablecoin market capitalisation (DefiLlama)	**yes**	the quantity is a hit and the channel is a measured zero
$\Psi_{\text{str}}(t)$	COFER dollar share, and **plus gold**	**yes**	the gold-inclusive aggregate of Section 9.7.1 is the one that shows acceleration
T_{frag}	SWIFT currency shares	no	free, PDF-only, untranscribed
C_{alt}	CIPS volumes, mCBDC pilots	no	Tier 3
D_{hcap}	IP receipts, skilled migration	**no**	never examined anywhere in this book

Remark 20.2 (The cadence is best equipped where the horizon is longest)**.** Counting across the four frames: the hour–day watchlist has one usable instrument of four, the weekly three of five, the quarterly three of four, the annual two of five. **The emergency frame — the one that decides whether the strategy survives its first week — is the worst provisioned**, and the single instrument it most needs, the cross-currency basis, is the same one Chapter 17's agenda ranks first. A monitoring plan is only as good as its fastest series, and this one's fastest series is the one it does not have.

20.2 Amplified Structural Traps

The original identifies four self-sabotaging actions. Unusually for Part III, **three of the four occurred in 2022–23**, which makes this section testable. Each is stated as the original states it, then checked.

20.2.1 Trap 1: mixed messaging

The trap. The central bank announces tightening, then panics when stress rises and pivots prematurely. Impact: $M_{\text{trust}} \downarrow\downarrow$, markets confirm "the Fed has failed", H_{soc} and V_{meme} spread it, $\kappa_{\text{sys}} \uparrow\uparrow\uparrow$ as expectations unanchor.

Did it occur? **No.** The 2022–23 cycle raised 525 basis points and held above five per cent for a year, through the March 2023 regional-bank failures, before the first reduction in September 2024. This is the one trap the record shows being *avoided*, and it is therefore the one place where the strategy and the outcome coincide.

> **What that is worth.** Less than it appears. The trap was avoided *and* expectations stayed anchored, which is consistent with the strategy working — and equally consistent with Remark 18.6's alternative, that the unanchoring threat was overstated. One episode with no counterfactual cannot separate them, and M_{trust} has no measure that would.

20.2.2 Trap 2: political brinkmanship

The trap. Congress triggers a debt-ceiling or budget crisis; $R_{\text{US}} \uparrow\uparrow\uparrow$ amplified by H_{soc}; the "US government incompetence" meme erodes Φ_{sys} and nullifies the Right Hand.

Did it occur? **Yes.** The 2023 debt-ceiling standoff ran to an agreement on 2023-06-03, and Fitch downgraded the United States on 2023-08-01 — during the tightening cycle, precisely the conjunction the trap describes.

> **Standing.** Section 13.4 scores this as CS4 and the result is one hit, one partial, four untestable. The dollar rose over the 2023 window. What the model requires — that $R_{\text{US}} \uparrow$ transmits to $\Phi_{\text{sys}} \downarrow$ — is untestable, because neither variable has a series. And Proposition 13.2 establishes that this class of episode can never identify the coefficient: episodes clean enough to isolate R_{US} are ones markets judged unimportant, so they carry no signal. **Trap 2 occurred, and the model has no way to say whether it mattered.**

20.2.3 Trap 3: over-weaponisation

The trap. To fight energy inflation the United States escalates broad secondary sanctions. Impact: the meme "the dollar is unsafe" spreads through H_{soc} and V_{meme}, accelerating $\Psi_{\text{str}}(t)$ to its tipping point — an outright strategic failure.

Did it occur? **Yes, and this is the one trap with supporting evidence.** The immobilisation of Central Bank of Russia assets took effect 2022-02-28; the embargo and price cap followed on 2022-12-05.

> **Standing, and it cuts the model's way for once.** Section 9.7.1 finds that on a reserve aggregate including monetary gold — which COFER excludes — the dollar's share is declining at an *accelerating* rate, -0.0211 pp/yr^2 with $t = -2.59$, against $+0.0076$ ($t = +1.13$) for the same fit on COFER alone. Official gold purchases tripled from 2022 and gold rose from 12.3 to 23.6 per cent of the broad

reserve base. The exit this trap predicts is visible once the right instrument is used. **Two things stop this being a confirmation.** The causal step — that coercion produced it — is not established, and Remark 9.2 shows the constant-price series was already accelerating before 2022. And $\Psi_{\text{str}}(t)$'s tipping point remains unfalsifiable (Proposition 9.3). This is the strongest evidence in Part III and it is still only consistency.

20.2.4 Trap 4: plumbing neglect

The trap. Poor coordination — the Fed tightening and running off the balance sheet while the Treasury simultaneously drains liquidity through under-issuance of bills or a TGA rebuild. Impact: an unnecessary double liquidity shock, σ_{sys} surges, $\kappa_{\text{sys}}(H_{\text{soc}})$ amplifies it into a crash, which forces the pivot of Trap 1 and destroys M_{trust}.

Did it occur? **Yes, in full, and more completely than the trap describes.**

Table 20.5: Every condition of Trap 4, met simultaneously in 2022–23.

Condition	What happened	
Fed tightening	+525 bp, 2022-03 to 2023-07	fastest since 1981
Balance-sheet runoff	concurrent from June 2022	
TGA rebuild	\$212bn → \$657bn	a \$445bn drain
Third drain not in the trap	ON RRP to **\$2.55tn** on 2022-12-30	over two trillion absorbed from the system
Predicted consequence: a market crash forcing a policy pivot.		
Observed consequence: none. No crash, no pivot, no unanchoring.		

Proposition 20.1 (Trap 4's conditions were met and its consequence did not follow)**.** *Over 2022–2023 the Federal Reserve raised its target range by 525 basis points while running off its balance sheet; the Treasury General Account rose by \$445 billion; and the overnight reverse repurchase facility absorbed a further \$2.55 trillion at its peak. All four drains operated at once. The dollar rose* 13.4 *per cent and then fell* 9.7*;* σ_{sys} *declined from* 0.91 *to* 0.65 *across the period; no crash occurred and no pivot was forced.*

Remark 20.3 (And the reason is a mechanism the model does not contain)**.** The reverse-repo facility acted as a *buffer* rather than a drain: as reserves would otherwise have fallen, money funds withdrew from the RRP instead, so the balance-sheet runoff was absorbed by the facility rather than by bank reserves. In the model's terms an increase in S_{Tsy} is unambiguously a reduction in effective supply — there is no term for a substitution between two liabilities of the central bank, and no state variable that records whether a drain lands on reserves or on the facility. **The model cannot represent the thing that made the episode survivable**, which is a specification gap of the same class as ISSUES #40's scalar

D_{prec}. It is recorded here rather than repaired, because repairing it means adding a term with no identification — the ground on which ISSUES #32, #34 and #37 were all declined.

20.3 The Optimal DXY Trajectory and the Updated Equation

20.3.1 What "optimal" means here

The original's framing deserves to be preserved intact because it is the most useful sentence in Part III and it is independent of everything the scorecard disputes:

> The optimal path does not aim for a strong or a weak dollar. **A sharp rise driven by $\sigma_{\text{sys}} \uparrow$ is a bad rise.** The goal is a stable and credible dollar in a brittle system.

That is a statement about *which* dollar movements are informative, and it follows from the decomposition rather than from any contested elasticity. A dollar that rises because stress rose is a symptom; a dollar that is stable while stress is high is the objective. This survives Chapter 17 entirely, because it is a definition rather than a prediction.

20.3.2 The three phases, and the episode

The original predicts three phases with amplitudes larger than historical: Phase 1 a violent surge as the tightening shock is amplified by $\kappa_{\text{sys}}(H_{\text{soc}})$; Phase 2 a rapid cooling as the Left Hand opens the valves and deprives κ_{sys} of fuel; Phase 3 a flattening if M_{trust} is restored and T_{frag} contained.

Table 20.6: The predicted trajectory against the 2022–26 record. Dollar is `DTWEXAFEGS`; script `data/policy_evidence.py`.

Phase	Window	Dollar	σ_{sys}	Predicted
1	2022-03-16 to 2022-09-27	$+\mathbf{13.4}\,\%$	$0.91 \to 0.89$	violent surge
2	2022-09-27 to 2023-12-29	$-\mathbf{9.7}\,\%$	$0.89 \to 0.65$	rapid cooling
3	2024-01-02 to 2026-09-01	$-\mathbf{1.3}\,\%$	$0.57 \to 0.44$	flat

Proposition 20.2 (The predicted shape is present; the predicted mechanism is not)**.** *The dollar's path across the 2022–26 tightening and its aftermath matches the original's three-phase prediction in sign and in relative magnitude: a large rise, a large fall of roughly two-thirds the rise, then a flat period, with σ_{sys} declining monotonically across the three. But the original attributes Phase 2's cooling to the Left Hand opening liquidity valves, and Table 18.8*

shows swap lines peaked at $11.3 *billion in 2022–23 against* $449 *billion in 2020.* ***The valves were never opened. The shape occurred without the mechanism assigned to it.***

Remark 20.4 (Which is the strongest thing that can be said for Part III, and the weakest)**.** Strongest: the trajectory is right, and it was written before the episode. A three-phase prediction that matches in sign and rough magnitude is not nothing, and it is more than Chapter 16's CS11 managed on the tariff episode.

Weakest: a prediction that comes true through a mechanism the predictor did not use is a prediction that was overdetermined. Rates rose, then the tightening cycle ended and the dollar gave back part of the move — which is what a rate-differential account predicts without any of Part III's apparatus, and Remark 14.2 already found the rate channel outperforming the model's preferred channel on daily data in 2022. **The shape is consistent with the model and does not discriminate between it and the simplest available rival.**

20.3.3 The updated equation, E10c

The original closes with a differenced form in which the modulators become explicit functions of the psycho-technical layer:

$$\begin{aligned}\Delta \ln DXY \approx \kappa_{\text{sys}}(H_{\text{soc}},\cdot)\Big[\sum_i \alpha_i(\sigma_{\text{sys}})\Delta \ln D_i - \sum_j \beta_j(\sigma_{\text{sys}})\Delta \ln S_j\Big] \\ + \Big(\lambda_1 \Delta\sigma_{\text{sys}} + \lambda_2 \Delta\sigma_{\text{sys}}^2\Big)\,\Phi_{\text{sys}}\Big(R_{\text{US}}(H_{\text{soc}})\Big).\end{aligned} \tag{20.1}$$

Equation (20.1) differs from Equation (10.5) in two places and both are arguments rather than terms: κ_{sys} is written as an explicit function of H_{soc}, and Φ_{sys} is written as a function of R_{US} which is itself a function of H_{soc}. The original's text is clipped at this point in the source (ISSUES #19); the reading above is the one consistent with Chapter 7's surrounding argument and it is recorded as a reconstruction.

Remark 20.5 (E10c adds no estimable content over E10b)**.** Writing $\kappa_{\text{sys}}(H_{\text{soc}},\cdot)$ rather than κ_{sys} makes a dependency explicit without making it measurable: Proposition 11.1 shows κ_{sys}'s seven-way decomposition is unidentifiable, and H_{soc} has no series. The same is true of $\Phi_{\text{sys}}(R_{\text{US}}(H_{\text{soc}}))$ — three nested latents, none measured, entering multiplicatively so that only their product with λ is ever recoverable. **E10c is a statement of what the model believes, not a specification that can be taken to data**, and it is the clearest single illustration of Section 11.6.1's finding that the model is best specified exactly where it is least measured. It is presented because the original states it and because the concordance requires nothing to be lost.

20.3.4 The strategic battlegrounds

The original's closing division, restated with standing attached.

Remark 20.6 (The closing statement of Part III)**.** The division in Table 20.7 is the summary of these three chapters. The half of the strategy that operates on plumbing — swap lines,

Table 20.7: Where the original says the fight is, and what this book can say about each.

Sector	The instruction	Standing
Real	Minimise the physical shock $\Delta\sigma_{\text{sys}}$ via the Left Hand; reduce $\Delta \ln D_{\text{trans}}$ via the Root Layer	σ_{sys} is the best-attested channel in the book; S_{swap} is the one lever with hits. **This half is supported**
Narrative	Maximise Φ_{sys} via M_{trust}; defend $\Psi_{\text{str}}(t)$ by controlling T_{frag}; counter H_{soc}	Φ_{sys}, M_{trust}, H_{soc} all unmeasured; T_{frag} one partial and four untestable. **This half is unevidenced**

collateral, bank capacity, the timing of drains — is documented, dated, legally grounded and, where the book could test it, partly supported. The half that operates on belief — trust, homogeneity, narrative velocity, structural integrity — is the half the original considers decisive, and it is supported by nothing in this book in either direction.

A reader should take from Part III the institutional map and not the strategy. The levers are real, their authorities are real, their activations are dated and their sizes are public. The claim that pulling them in the specified order produces the specified dollar path rests on a model that Chapter 17 scores at 27 hits of 110, and on a psycho-technical layer with no scored predictions at all. Those are not the same kind of statement and this part has tried never to let them look alike.

Chapter 21

The Decentralized System Coordination Mechanism

21.1 The Hypothesis, and Why It Is the Least Testable Claim in the Book

Assumption 18.1 governs this chapter as it governs Part III, and here the separation it demands is sharper than anywhere else, because the chapter's central claim is about *intent* — or about something that behaves as intent without being it — and the model has no instrument for intent.

21.1.1 What the original claims

Restated faithfully. The long-term stability of the dollar system is not the outcome of individual policy decisions but an *emergent property* of a network of interests, institutional constraints and non-linear feedback. Viewed through the model, a striking pattern appears: although policy actions look fragmented and are driven by different motives, the system's overall trajectory converges around a **local attractor basin** in which structural risk to the dollar is minimised.

The original is careful, and the care should be preserved. This is explicitly *not* a claim about a conscious central actor or a "deep state". It is a claim about self-organisation: in the state space defined by G_{pol}, R_{US} and Φ_{sys}, there may exist a narrow region toward which policy trajectories are pulled back regardless of their starting point, sustaining the dollar's structural stability over the long run.

21.1.2 Evidential standing, stated before the argument rather than after it

This is the least testable claim in the book, and it fails in two independent ways.

First, the state space has no measurable axes. The attractor is posited in $(G_{\text{pol}}, R_{\text{US}}, \Phi_{\text{sys}})$. G_{pol} has six scored predictions — two hits, one miss, three

untestable — and worse, ISSUES #33 established that the book had been scoring it by the strategic weight of the target rather than by demonstrated capability, while Section 14.3.4 showed its three capacities moved in *opposite directions* during 2026 with no aggregator to combine them. G_{pol} **is not a well-defined scalar**, so one axis of the space is not merely unmeasured but ill-posed. R_{US} has five predictions, one miss and four untestable: **it has never been scored right**. Φ_{sys} has three predictions and **all three are untestable**, because it enters the equations only as a product with λ and cannot be separated from it. **No axis of the attractor's state space can be located.**

Second, the model has no instrument for intent. Section 14.3.7 settled this in a different context and the ruling binds here: the model has no variable for policy anticipation, no expectations channel over future policy, and nothing in Equation (10.2) that would let a pattern of actions be read as purposive rather than coincidental. A convergence and a coordination are the same data.

Remark 21.1 (What would make the claim testable, and why it is not a small gap)**.** An attractor is a statement with content: trajectories that start away from the basin should return to it, and the return should be faster the further away they start. That is measurable *if the state can be located*. So the requirement is not exotic — it is a series for each of three variables, at any frequency, for any sample. The book has none, and Chapter 11 classes two of the three as structural rather than data problems. Section 21.7 does the one thing that can be done instead: it tests the claim on the *output* rather than on the state, which is a far weaker question, and reports the answer.

21.1.3 The claim drawn, and the axes labelled

Panel (a) is included because the outline for this chapter asks for a phase diagram rather than a proof, and because drawing a claim is a good way to see what it would take to test it. Trajectories starting anywhere return to a basin; the basin's depth is Φ_{sys}. That is a statement with real content *if the axes exist*. **The figure's own axis labels are the chapter's verdict.**

21.2 The Other Two Pieces of Part 0's Evidence

The original offers four items. Plaza and the GFC are treated at length in Sections 21.3 and 21.4; the remaining two are shorter and are dealt with here, because each has a distinct evidential character and neither should be folded into the others.

21.2.1 Contradictory actions as evidence of imperfection

The claim. Frequent spikes in R_{US} — debt-ceiling crises above all — do not disprove self-regulation. They show its imperfections and lags. Short-term political dynamics push the system away from the optimal zone; structural forces pull it back. A pendulum disturbed by turbulence, always oscillating back toward equilibrium.

F38: an attractor asserted in a space with no measurable axis, and the one test that could be run.

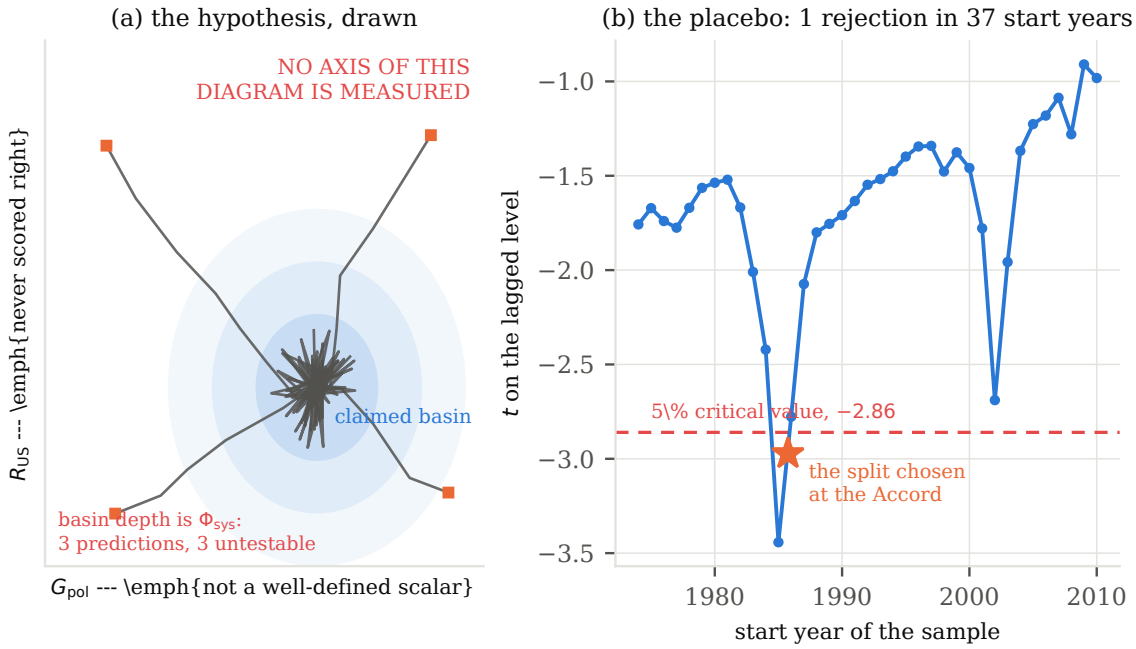


Figure 21.1: Panel (a) is a phase portrait of the *hypothesis*, not of data — Chapter 11 and Chapter 17 between them establish that no axis of it can be located. Panel (b) is data: the mean-reversion test at every start year from 1974 to 2010. One rejection in thirty-seven, and it is the start chosen at the Accord. Script: `figures/src/f38_attractor.py`.

Standing. The pendulum image is the whole difficulty in one metaphor. A pendulum that always returns is a testable claim about a restoring force; a pendulum that returns *eventually*, with unspecified lags and imperfections, is compatible with any path whatever. Proposition 21.2 tests the observable form and finds no restoring force in the dollar. The unobservable form — a restoring force in R_{US} — would need a series for R_{US}, which does not exist, and Proposition 13.2 shows why one episode type that would help, the debt ceiling, is structurally unable to: the clean episodes are the ones markets judged unimportant.

21.2.2 Post-1971 persistence as evidence

The claim. Abandoning gold created a state space in which self-regulation became necessary. The floating, debt- and trust-based system is inherently unstable, and its persistence to this day may itself be evidence that institutions, implicit rules and feedback loops evolved organically to sustain it.

Standing — and this is a survivorship argument. "The system has not collapsed, therefore something is stabilising it" is an inference from a single realised history with no comparison group. The same reasoning would have supported the stability of the gold standard at any point before 1971, and Chapter 22 is the chapter about the moment it stopped being true. **Persistence is what a stable system and a lucky one look like from inside.** This item is recorded because the concordance requires nothing to be lost; it is not evidence.

21.3 Historical Evidence I: The GFC as Systemic Self-Reinforcement

21.3.1 The documented record

A crisis originating in United States mortgage credit produced a *rise* in the dollar. That is the fact, and it is not in dispute. Section 13.1 scores it: the trade-weighted dollar rose 18.33 per cent between 1 August and 21 November 2008, with σ_{sys} at its sample maximum of 0.961. The policy response is equally documented: swap lines expanded to a peak of $583.1 billion outstanding on 2008-12-17 (Section 19.2.1).

21.3.2 The model reading, and its standing

The original's interpretation is that the system's own non-linearities — D_{debt} leverage and S_{bank} fragility — caused a shock *originating in* the United States to *automatically reinforce* the dollar's centrality, through safe-haven demand and a short squeeze. The subsequent bailouts and swap lines are then read as efforts to stabilise the new attractor state rather than as an intentional reset.

Standing. The *direction* is a hit and the *mechanism* is not. CS1 scores $\sigma_{\text{sys}} \uparrow$, $F_{\text{liq}} \uparrow$, $D_{\text{prec}} \uparrow$ and $S_{\text{swap}} \uparrow$ as hits. But Section 11.5.6 found that the safe-haven

> block accounts for +0.80 of the +18.33 per cent move, leaving **64 per cent residual to every block in the model**, and the audit of Section 17.6 failed four separate attempts to attribute that residual. **The episode the chapter cites as evidence of self-reinforcement is the episode the book cannot explain.** That does not refute the reading; it means the reading is being supported by a move whose cause is unknown.

Remark 21.2 (The self-reinforcement claim is nearly circular)**.** "A shock to the dollar system strengthens the dollar" is either a finding or a definition, depending on how D_{prec} is measured. Section 12.3's Rule 1 exists for exactly this reason: D_{prec} may not be evidenced from VIX, because VIX is already the base leg of σ_{sys}, so "stress rose and safe-haven demand rose" would be one series reported twice. The GFC evidence survives Rule 1 — CS1 evidences D_{prec} from bill yields and custody holdings, not from VIX — but the general claim does not, and the chapter's use of it here is looser than the case study's.

21.4 Historical Evidence II: The Plaza Accord

This is the chapter's strongest example and it is worth stating in full, because the documented record is unambiguous and the model reading is where the difficulty lies.

21.4.1 The documented record

Table 21.1: The Plaza episode. Dates from the public record (`data/SOURCES.md`); the index is `DTWEXM`, the Federal Reserve's major-currencies trade-weighted dollar, whose first observation is 1973-01-02.

Date	Event	`DTWEXM`	
1980-07-01	trough before the rise	93.09	
1985-02-25	**the peak**	**148.12**	+59 % from the trough
1985-09-20	last session before the Accord	132.22	
1985-09-22	**Plaza Accord**, G5, New York	—	
1985-09-23	first session after	128.55	−2.8 % in a day
1986-09-22	one year on	105.10	−20.5 % from the Accord
1987-02-22	Louvre Accord, G7, Paris	100.35	signed to *halt* the fall
1991-02-11	the eventual trough	82.53	−44.3 % from the peak

The G5 — the United States, Japan, West Germany, France and the United Kingdom — agreed at the Plaza Hotel on 1985-09-22 to cooperate through concerted intervention to lower the dollar. The dollar fell 2.8 per cent in the first session and 20.5 per cent over the following year. Seventeen months later the same governments, now the G7, signed the Louvre Accord to *stop* the decline they had started. All of that is institutional fact, and the negotiating record is Funabashi [35], with Frankel [34] the thirty-year reassessment.

21.4.2 The model reading

The original asks the right question: why would a dominant power deliberately weaken its own currency? Its answer is that Plaza was a structural self-correction. A persistently strong dollar strangled domestic manufacturing (D_{hcap} weakened), produced a large trade deficit (S_{ext} deeply negative), and damaged allied economies. Sustained, that would have forced allies toward regional blocs and alternative settlement — $T_{\text{frag}} \uparrow$ and $\Psi_{\text{str}}(t) \downarrow$. The least-bad option was to sacrifice short-term dollar strength to preserve systemic integrity. And the decision need not have come from a single actor: it may reflect convergence among the Treasury, exporters and allied governments who each saw the same risk.

21.4.3 The ordering problem

Proposition 21.1 (The correction preceded the coordination)**.** *The dollar peaked on* ***1985-02-25*** *at* 148.12 *on* `DTWEXM`*. The Plaza Accord was signed on* ***1985-09-22****,* ***seven months later****, by which time the index stood at* 132.22 — ***already*** 10.7 ***per cent below its peak****. The reversal the Accord is credited with had begun and was well advanced before the Accord existed.*

Proof. Peak located by maximum of `DTWEXM` over 1980-01-01 to 1990-01-01; values on the stated dates read from the same series. `data/history_1971_1985.py`. □

This does not dispose of the reading, and it is important to say why not. Three interpretations survive Proposition 21.1 and the chapter cannot choose between them.

Table 21.2: Three readings of the same seven months.

Reading	What it says	What would distinguish it
Attractor	The system was already returning to the basin; Plaza *ratified* and accelerated a correction already under way, which is what a decentralised mechanism looks like	a measure of the state, so that "already returning" is observable rather than inferred from the outcome
Anticipation	Markets priced the Accord in advance; the seven months are the negotiation becoming known	the negotiating record, and a measure of expectations over policy the model does not have
Coincidence	The dollar was extended after a 59 per cent rise and mean-reverting; the Accord arrived during a fall that would have happened anyway	**whether the dollar mean-reverts at all**, which Section 21.7 tests

Remark 21.3 (The ordering is a problem for the strong form only)**.** If the claim is that the system *acted* to correct itself, the ordering is awkward: the correction started first. If the claim is the weaker one the original actually makes — that the system converges on a basin without anyone steering it — then a correction beginning before the meeting is *consistent*, and arguably better evidence than a correction beginning after it. **The weak form survives Proposition 21.1. It survives by being harder to distinguish from no mechanism at all**, which is the difficulty this whole chapter has.

21.5 Information Asymmetry and Personal Calculations

The original anticipates the obvious objection — that policy is frequently self-destructive — and answers it with two mechanisms.

Information asymmetry. If a coordination mechanism exists, its logic must prioritise long-horizon structural variables, $\Psi_{\text{str}}(t)$ and Φ_{sys}. Those are precisely the variables no participant observes directly. So the mechanism cannot be executed by anyone; it can only be approximated by actors reacting to proxies.

Personal calculations. Political actors optimise over horizons far shorter than $\Psi_{\text{str}}(t)$'s. A debt-ceiling confrontation is rational for its participants on a two-year electoral horizon and damaging on a twenty-year structural one. The contradiction is not a failure of the mechanism but a feature of its substrate.

> **Standing.** This section is unfalsifiable as stated, and the reason is worth naming precisely rather than dismissing. It explains *both* coordinated action and uncoordinated action — convergence is the mechanism working, divergence is the mechanism's known imperfection. Table 14.14 set out the test that separates a legitimate multi-branch account from an unfalsifiable one: *can the branch be selected from conditions observable before the outcome?* Here it cannot, because the selector is the degree of information asymmetry, which is not observed. **This is the same defect as CS8's** (ISSUES #36), in a more sophisticated form.

21.6 The Self-Contradictory System

The original's conclusion is that the picture is not a perfect conspiracy but a continuous struggle between long-horizon systemic logic and short-horizon political dynamics, producing a trajectory that alternates between optimisation and self-inflicted disruption. The mechanism, it argues, is not omnipotent, is constrained by information, and operates through feedback loops rather than commands — bond markets punishing instability and forcing politicians to retreat.

Remark 21.4 (The feedback-loop version is the testable one, and the book has tested it twice)**.** "Bond markets punish instability and force a retreat" is a claim about prices and dates, not about intent, and it can be checked. The book contains two direct tests and

they disagree. Section 13.4's CS4 finds the 2011 and 2023 debt-ceiling episodes produce no identifiable punishment: Proposition 13.2 shows that episodes clean enough to isolate R_{US} are ones markets judged unimportant, so the punishment cannot be measured even in principle. Against that, Table 17.9 finds the thirty-year Treasury losing its haven property in 2023 and 2025 while the two-year kept it — which *is* a long-horizon market repricing of something, and ISSUES #40 records that the model cannot represent it with a scalar D_{prec}. **So the one mechanism in this chapter that could be observed shows a signal the model has no term for.**

21.7 The One Observable Shadow of the Attractor Claim

The state space cannot be located, so the state cannot be tested. What can be tested is whether the *output* behaves as a system with an attractor would: does the dollar level return toward a central value, or does it wander? This is a much weaker question than the one the chapter asks — a system can have an attractor in its state space and a random-walk output — and it is the only one available.

Table 21.3: Mean-reversion test on the monthly log level of `DTWEXM`. The five per cent critical value for this statistic is about −2.86, not −1.96, because the regressor is a lagged level.

Sample	n	Coefficient	t	Verdict
Full, 1973-01 to 2019-12	563	−0.0111	−1.84	cannot reject a random walk
Post-Plaza, from 1985-10	410	−0.0258	−2.97	rejects — *but see below*

The second row looks like support and is not, for a reason this book has ruled on elsewhere: the subsample was chosen at the event whose effect it is being used to assess, which is the selection Remark 11.2 forbids. The placebo settles it.

Proposition 21.2 (The dollar does not mean-revert, and the post-Plaza result is a selection artefact)**.** *On the full 1973–2019 monthly sample the mean-reversion coefficient on the lagged log level of* `DTWEXM` *is* −0.0111 *with* $t = -1.84$*, which does not reject a random walk at five per cent. Repeating the test at every start year from 1974 to 2010 —* 37 *samples — rejects in* ***one****, and that one is the 1985-10 start chosen at the Accord. A start in 1986 gives* $t = -2.78$ *and does not reject.* ***A test that rejects in*** 3 ***per cent of arbitrary start dates is not evidence of an attractor.***

Proof. Dickey–Fuller-form regression of $\Delta \log \texttt{DTWEXM}_t$ on $\log \texttt{DTWEXM}_{t-1}$ and a constant, monthly, in `data/history_1971_1985.py`. Placebo over start years 1974–2010 retaining samples of at least 120 months. □

Remark 21.5 (What this does and does not settle)**.** It does *not* refute the attractor hypothesis. A system can have a strong attractor in $(G_{\text{pol}}, R_{\text{US}}, \Phi_{\text{sys}})$ and an output that is a near-random walk, if the mapping from state to price is noisy — and this book's own estimate of that mapping leaves 64 and 76 per cent of its two central episodes unexplained, which is more than enough noise to hide a state-space attractor.

What it settles is narrower and worth having. **The single most-cited surface evidence for the hypothesis — "the dollar always comes back" — does not survive a placebo.** The implied half-life on the full sample is 5.2 years with a t that does not reject zero, so even the point estimate describes a drift rather than a pull. Anyone arguing for the attractor must argue it in the state space, where nothing can currently be measured, and not from the dollar's own path, where it is not there.

21.8 Verdict on This Chapter

Stated plainly, because a chapter this speculative should not end ambiguously.

Table 21.4: The chapter's four claims, sorted.

Claim	Status	
Plaza and the GFC happened as described	**documented fact**	dates, magnitudes and policy responses are public record and are not in question
The system converges on an attractor in $(G_{\text{pol}}, R_{\text{US}}, \Phi_{\text{sys}})$	**not testable**	no axis of the space can be located; one is not a well-defined scalar
The convergence reflects self-organisation rather than direction	**not testable**	the model has no instrument for intent (Section 14.3.7)
Contradictions are the mechanism's imperfections	**unfalsifiable**	explains convergence and divergence alike; the CS8 defect in a more sophisticated form
The dollar itself returns to a central value	**tested, and refuted**	Proposition 21.2

A reader should take the history and leave the hypothesis. The Plaza Accord is a remarkable episode and the question the original asks about it — why a dominant power would weaken its own currency — is a good question that the model answers coherently. Coherent is not the same as supported. Chapter 17 lists no scored prediction anywhere in this book that bears on the attractor claim, and this chapter has now added the one test that could be run, which came back negative.

Chapter 22

The Inflection Point: Vietnam, the Nixon Shock, and the Birth of the Modern System

This chapter has the opposite character to Chapter 21. There the claim was about intent in a state space with no measurable axes; here the events are documented history with dates, and part of the dollar's response is observable. The conditional device of Assumption 18.1 still applies, but it applies to less of the chapter, and where it applies it is because of the model's variables rather than the historical record.

22.1 What Is Observable and What Is Not

Stated first, because the chapter's evidential position is unusual and a reader should know it before the argument rather than discover it inside one.

Remark 22.1 (A correction to an assumption about this chapter's data)**.** It has been supposed that a dollar index running from 1967 was available here. It is not. `raw/stooq_DXY_d.csv` was removed from the repository on 2026-09-05 because the download returned a Cloudflare bot-check page rather than observations, and that removal is recorded in `data/SOURCES.md`. The earliest trade-weighted dollar series held is `DTWEXM`, whose first observation is **1973-01-02** — after the shock of 1971-08-15, after the Smithsonian realignment of 1971-12-18, and two months before the generalised float of March 1973. **The index does not cover the event.** Bilateral rates from 1971-01 do, and this chapter uses those, as bilateral rates and not as an index.

22.2 Breakdown of the G_{pol}–R_{US} Equilibrium

22.2.1 What the original claims

From the end of the Second World War to Vietnam, the United States maintained a balance: the projection of external power — the Marshall Plan, containment, Cuba — was matched by internal consensus and low domestic risk. External strength rested on internal cohesion.

Table 22.1: The chapter's instrument set. "Held" means the series is in this repository and covers the stated period.

Object	Held?	Coverage	Note
Trade-weighted dollar	**partly**	from 1973-01-02	`DTWEXM`. **Begins seventeen months after the Nixon shock**
Bilateral dollar rates	**yes**	from 1971-01	sterling, yen, Swiss franc, D-mark, Canadian dollar; monthly
Gold price	**yes**	from 1960-01	World Bank pink sheet, monthly
CPI, unemployment, policy rate	**yes**	from 1947–54	
US reserves excluding gold	**yes**	from 1950-12	
US gold stock, in ounces	**no**	—	the series the chapter's central narrative needs. Attempts on FRED returned `404`, not a timeout, so this is a confirmed absence rather than a failed download
Net exports	**yes**	from 1947-Q1	`NETEXP`; Table 22.5
G_{pol}, R_{US}, M_{trust}	**no**	—	no series for any period

Vietnam broke it, in two steps. First, a *failure to maximise* G_{pol}: the attempt to demonstrate absolute military supremacy produced not a swift victory but a protracted war of attrition that eroded credibility and diminished global perceptions of American power. Second, and the original marks this as the more important, an *explosion* of R_{US}: mass protest, urban unrest and a collapse of trust in government pushed domestic risk to extremes. For the first time external ambition collided with, and damaged, internal stability. The war abroad became a war within.

> **Standing.** Both variables are in the structural block, which Table 17.3 scores at **78 per cent untestable**, and Section 21.1.2 sets out their condition in full: G_{pol} is not a well-defined scalar and R_{US} has never been scored right. What is specific to this chapter is narrower and worse. **Neither variable has a series for 1964–1971 — or for any other period**, and unlike the modern episodes there is not even a contemporaneous proxy to argue over. So the claim that external ambition and domestic risk moved in opposite directions during Vietnam is a historical judgement stated in the model's vocabulary. It may well be right. It is not a measurement, and nothing in this book makes it one.

22.2.2 What can be measured about the strain

The fiscal consequence of the contradiction is measurable even though the two variables are not, and it is worth having on the page.

Inflation nearly quadrupled between 1964 and 1970 while unemployment stayed below

Table 22.2: The macroeconomic record across the Vietnam escalation. Sources: FRED `CPIAUCSL`, `UNRATE`, `FEDFUNDS`, `TRESEGUSM052N`; script `data/history_1971_1985.py`.

Year	CPI, y/y %	Unemployment %	Policy rate %	
1964	1.6	5.6	3.48	before the escalation
1966	1.9	4.0	4.42	
1968	3.6	3.7	4.61	Tet; the escalation's peak year
1970	**6.2**	3.9	**8.98**	**the strain visible**
1972	3.3	5.8	3.51	after the shock; controls in force
1974	**9.6**	5.1	9.65	the first oil shock
1978	6.8	6.4	6.70	

four per cent, and the policy rate rose from 3.48 to 8.98 per cent. That is the "guns and butter" strain the original describes, in the series that exist. **What it does not show is the gold outflow**, which is the mechanism the chapter turns on, and for which this repository has no series in ounces. The absence is stated rather than papered over with the reserves-excluding-gold series, which measures a different thing and moves the other way.

22.3 Collapse of M_{trust} and the 1971 Structural Break

22.3.1 The documented timeline

The historiography is settled enough to be cited rather than argued. Garber [37] is the standard account of the collapse, in Bordo and Eichengreen [16]'s volume; Gowa [45] is the domestic-politics reading of the decision itself; Eichengreen [29] places it in the long history of monetary regimes. For the petrodollar claim of Section 22.4, Spiro [71] is the source, and its claims are treated with the care Section 14.3.7's ruling on intent requires.

Table 22.3: The end of Bretton Woods. Event record verified and sourced in `data/SOURCES.md`.

Date	Event
1944-07	Bretton Woods establishes the dollar–gold parity at \$35 an ounce
1971-08-15	**Nixon suspends convertibility of the dollar into gold for foreign central banks**, and imposes a ten per cent surcharge on dutiable imports, explicitly to force other countries to revalue against the dollar
1971-12-18	**Smithsonian Agreement**: the G10 agree a new set of fixed rates around a devalued dollar; the official gold price is raised from \$35 to **\$38**
1973-03	The Smithsonian parities are abandoned after fifteen months; generalised floating begins

22.3.2 What the dollar did

Table 22.4: Bilateral dollar rates across the break, monthly. These are bilateral rates and are not combined into an index; no trade-weighted index covers this window.

Month	$/GBP	JPY/$	CHF/$	DEM/$	gold $/oz
1971-01	2.406	358.0	4.305	3.637	38.00
1971-08	2.435	355.8	4.031	3.416	43.00
1971-12	2.527	320.1	3.904	3.269	43.00
1973-03	2.472	261.9	3.217	2.813	84.00
1975-01	2.362	299.7	2.528	2.365	176.00
1980-01	2.264	237.9	1.595	1.725	675.00

Proposition 22.1 (The devaluation was large, and it was not uniform)**.** *Between 1971-01 and 1975-01 the dollar fell* ***41.3 per cent*** *against the Swiss franc,* ***35.0 per cent*** *against the D-mark and* ***16.3 per cent*** *against the yen, while gold rose from $38 to $176, a rise of* ***363 per cent****. Against sterling it fell* ***1.8 per cent****. The dispersion across the four is larger than the move against three of them.*

Remark 22.2 (The sterling exception is not noise, and it matters for the chapter)**.** Sterling was in its own crisis — the 1967 devaluation, and by 1976 an IMF programme — so the dollar's near-flat performance against it reflects the counterparty rather than the dollar. This is the general hazard in reading bilateral rates as a currency's strength, and it is exactly why a trade-weighted index exists. **The chapter cannot construct one for this window and does not pretend to.** What the four rates jointly support is that the dollar fell substantially against strong-currency creditors and did not fall against a weak one, which is the pattern a forced realignment produces and not a general loss of confidence in every direction at once. Section 9.5's sterling material makes the same point from the other side.

22.3.3 The external balance, which is measurable

The original's fiscal argument has a counterpart in the external accounts, and unlike G_{pol} and R_{US} that one has a series back to 1947.

The external position turned negative in the quarters immediately before the gold window closed, which is the fact the original's argument needs and the one piece of its mechanism this book can actually show. It does not establish causation — a deficit is also what a domestic boom produces — but the ordering is right, and it is the only part of the 1971 narrative supported by a series rather than by a variable name.

22.3.4 The model reading

The original's claim is that this was not a policy change but a *structural break*: it destroyed the institutional anchor established in 1944 and collapsed M_{trust} in the United States as the

F39: no index covers the shock. The bilateral rates do, and they disagree about how far the dollar fell.

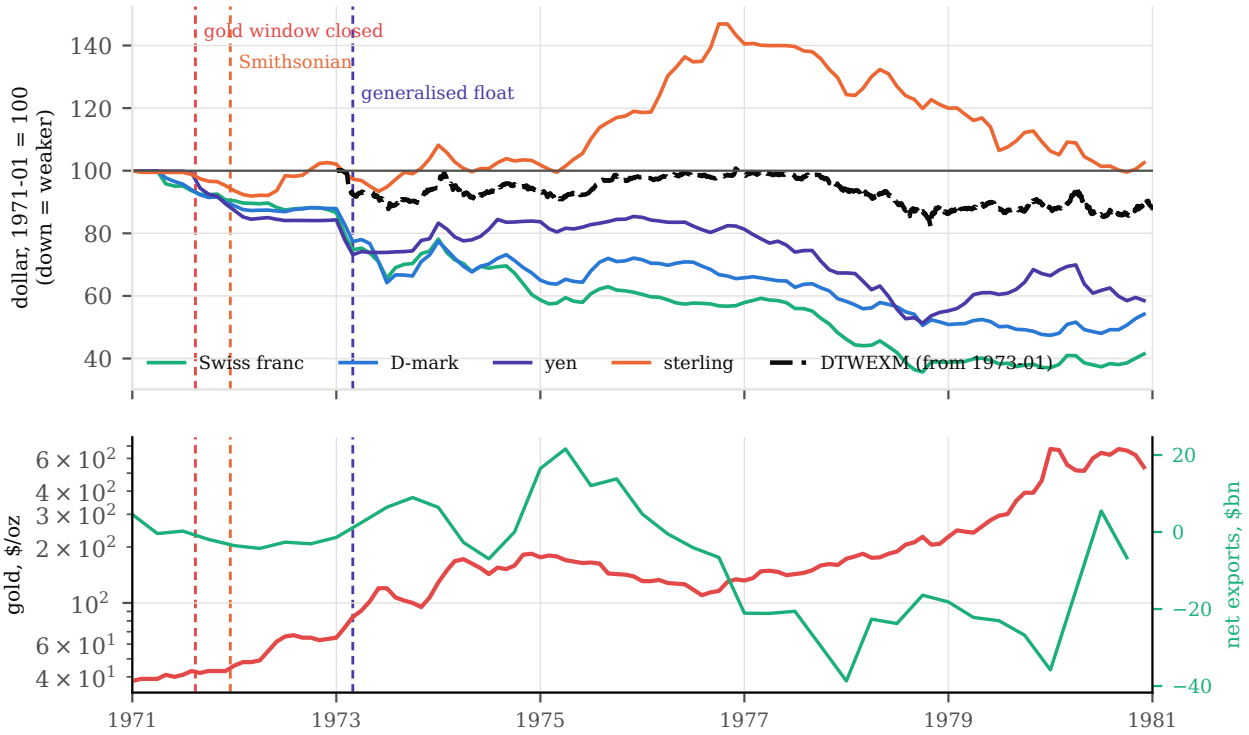


Figure 22.1: The top panel is the point of the figure: four bilateral rates indexed to 1971-01, oriented so that down is a weaker dollar, with `DTWEXM` overlaid in black from the date it begins so a reader can see where the index picks up — seventeen months after the event. The four disagree substantially about the size of the devaluation. Script: `figures/src/f39_nixon.py`.

Table 22.5: Net exports of goods and services, quarterly, $bn (FRED `NETEXP`). The Bretton Woods parity was defended against a deteriorating external position.

Quarter	Net exports	
1964-Q1	+7.21	the surplus era
1968-Q1	+1.05	Tet, and the surplus nearly gone
1971-Q2	**−0.40**	**the first deficit; the shock is that August**
1972-Q4	−3.06	after the Smithsonian
1975-Q1	+16.48	the devaluation working, with a recession
1980-Q1	−35.80	

system's guarantor. The dollar was no longer backed by a tangible asset but by the full faith and credit of the United States — an abstraction that this model quantifies as M_{trust} and Φ_{sys}.

> **Standing.** M_{trust} has **one** prediction in the entire scorecard and it is untestable; Φ_{sys} has three and all three are untestable. **The chapter's central claim — that 1971 collapsed institutional trust — is stated in two variables the book has never once scored.** The gold price is the closest observable, and it rose 363 per cent over four years, which is consistent with a collapse in the credibility of the old parity and equally consistent with the end of a price control, since $35 had been an administered price since 1934. **Nothing here separates a repricing from a loss of trust**, and the model contains no term that would.

Remark 22.3 (The one place this book can speak to 1971 with authority)**.** Section 9.7.1 is the exception and it runs the other way. It shows that measuring the dollar's reserve position on foreign exchange alone — excluding monetary gold — has concealed an accelerating decline in the dollar's share since 1999. The 1971 break is the moment gold ceased to be the numeraire and became an *alternative asset*; the audit of Section 17.6.3 shows that in the 2020s it is behaving as one again, at 23.6 per cent of world official reserves in 2025 against 8.3 in 2015. **The chapter's subject and the book's one confirmed rescue hypothesis are the same asset, fifty years apart**, and that connection is the strongest thing in this chapter.

22.4 The Modern System and Its Inherent Instability

The original identifies three consequences of abandoning the gold anchor. Each is restated with its standing.

Debt explosion (D_{debt}). Removing the constraint unleashed exponential expansion of dollar credit, especially eurodollar credit, making D_{debt} a core but unstable pillar.

> **Standing.** D_{debt} has three predictions: one miss, one partial, one untestable. And Section 17.6.1 found it to be the worst frequency mismatch in the model — a mechanism that operates daily, proxied by a BIS series that is quarterly, a gap of three steps on the frequency ladder. **The variable this section makes central to the modern system is the one the book measures least well relative to its speed.**

Exchange-rate volatility (κ_{sys}, σ_{sys}). Floating made the dollar highly volatile and sensitive to speculative flows and herd behaviour.

> **Standing, and here the record is split in a useful way.** σ_{sys} is the best-attested channel in the book: six hits, one partial, two misses. κ_{sys} has four predictions and **no verdicts at all**, and Proposition 11.1 proves its decomposition cannot be identified. So the claim that floating raised *sensitivity* has support and the claim that it raised *herd amplification* has none.

The petrodollar (D_{trans}). The arrangement to price oil in dollars, reinforced after 1973, is read as a strategic attempt to manufacture enforced demand after the loss of the gold peg.

> **Standing.** D_{trans} has five predictions: one hit, two partials, two misses — the most-tested Layer 0 channel and genuinely mixed. The specific claim here is about *intent* in creating the arrangement, and Section 14.3.7's ruling binds: the model has no instrument for intent. The arrangement is documented; the purpose attributed to it is not something this book can assess.

22.5 Epilogue: Trust as the Only Anchor

The original closes by calling the post-1971 transformation a civilisational one as much as a financial one: the birth of a self-referential monetary ecosystem, untethered and reflexive, whose persistence depends less on gold than on the continuous regeneration of trust.

That is the book's thesis in a sentence, and it is worth saying what has happened to it over the preceding twenty-one chapters.

Table 22.6: The epilogue's claim, decomposed.

Component	Status	
The system became self-referential in 1971	**documented**	convertibility ended on a dated day by a stated act
It is more volatile than what preceded it	**supported**	σ_{sys} is the best-attested channel in the book
Its persistence depends on trust	**not testable**	M_{trust}: one prediction, no verdict. Φ_{sys}: three, all untestable
Trust must be continuously regenerated	**not testable**	and Proposition 18.2 found the one measurable proxy *more* stable through the largest inflation in forty years than before it

Remark 22.4 (Closing the historical part)**.** The 1971 break is the best-documented event in this book and the one whose model reading is least checkable, and those two facts are the same fact. What makes the event legible — that it was a discrete institutional act with a date — is what makes its consequences hard to attribute, because everything else changed at the same moment and nothing about the state before it was measured.

The honest position is that this chapter is history with a model-shaped commentary, and the history is the part that will last. The dates, the parities, the realignment and the subsequent devaluation against strong-currency creditors are established. That the mechanism was a collapse of M_{trust}, that G_{pol} and R_{US} had diverged, and that the modern system therefore rests on continuously regenerated trust are readings this book has no instrument to confirm — as Table 17.3 shows for the structural block generally, and as Chapter 17's agenda proposes to change.

Appendix A

Notation and Symbol Table

Generated from `registry/variables.csv` by `tools/gen_notation_table.py`. Arrow scale used in tables: ↑ small, ↑↑ large, ↑↑↑ extreme (same for ↓); ↔ no change.

Table A.1: Notation. Layer, kind and original location follow `registry/variables.csv`.

Symbol	Name	Layer / kind	Original section; status
DXY	U.S. Dollar Index	0 output; observable	Ch.I 2.1 (p.4); defined
DXY_0	Baseline / integrating constant	0 output; parameter	Ch.I 2.1 (p.4); defined; identification resolved (ISSUES #12 S0-DECIDED), and preserved term by term for emergent channels by ISSUES #29 S0-DECIDED; t_base provisional
ϵ	Exogenous random shock	0 output; error	Ch.I 2.1 (p.4); defined
D^*	Global effective demand for USD	0 mechanical; composite	Abstract; Ch.II (p.3); defined
S^*	Global effective supply of USD	0 mechanical; composite	Abstract; Ch.III (p.3); defined
D_{trans}	Transactional demand (trade & commodities)	0 mechanical (demand); flow	Ch.II 2.1 (p.6); defined
D_{debt}	Debt-servicing demand (Eurodollar system)	0 mechanical (demand); flow	Ch.II 2.2 (p.6); defined
D_{cap}	Capital-asset demand (growth & rate differentials)	0 mechanical (demand); flow	Ch.II 2.3 (p.6); defined
D_{prec}	Safe-haven demand (flight-to-quality)	0 mechanical (demand); flow	Ch.II 2.4 (p.7); defined

Symbol	Name	Layer / kind	Original section; status
D_{res}	Reserve-management demand	0 mechanical (demand); flow	Ch.II 2.5 (p.7); defined
D_{hedge}	Hedging demand (FX swap costs)	0 mechanical (demand); flow	Ch.II 2.6 (p.7); defined
D_{gamma}	Derivative demand (dealer gamma)	0 mechanical (demand); flow	Ch.II 2.7 (p.7); defined
D_{net}	Network demand (payments & stablecoins)	0 mechanical (demand); flow	Ch.II 2.8 (p.8); defined; emergent-channel form (1 + X/Xbar) (ISSUES #29 S0-DECIDED)
$\bar{X}$	Scale constant for emergent channels	0 mechanical; convention	5.1; 5.3 (p.27); defined; value to be STATED in S14 (ISSUES #29 S0-DECIDED)
D_{micro}	Microstructural demand (seasonal)	0 mechanical (demand); flow	Ch.II 2.9 (p.8); defined
D_{hcap}	Human-capital & IP demand FLOW	0 mechanical (demand); flow	Ch.II 2.10; Ch.I 2.3 (p.8); defined; separated from the stock H_cap (ISSUES #24 S0-DECIDED)
$h(H_{\text{cap}})$	Stock-to-flow mapping for human capital	0 mechanical (demand); function	Ch.II 2.10; Ch.I 2.3 (p.8); defined; introduced by ISSUES #24 S0-DECIDED
S_{Fed}	Federal Reserve actions	0 mechanical (supply); flow	Ch.III 3.1 (p.9); defined
S_{Tsy}	U.S. Treasury actions	0 mechanical (supply); flow	Ch.III 3.2 (p.9); defined
S_{coll}	Collateral availability (repo)	0 mechanical (supply); flow	Ch.III 3.3 (p.9); defined
S_{bank}	Bank balance-sheet capacity	0 mechanical (supply); flow	Ch.III 3.4 (p.10); defined
S_{swap}	Fed USD swap lines	0 mechanical (supply); flow	Ch.III 3.5 (p.10); defined; emergent-channel form (1 + X/Xbar) (ISSUES #29 S0-DECIDED)
S_{ext}	External balance-of-payments flows	0 mechanical (supply); flow	Ch.III 3.6 (p.10); defined; effective form carries beta_ext (ISSUES #4 S0-DECIDED)
$S_{\text{ext}}^{\text{eff}}$	Effective external supply (attenuated by S_off)	0 mechanical (supply); composite	Ch.III 3.6; 4.2.1 (p.10); defined; log-safe form (ISSUES #28 S0-DECIDED)

Symbol	Name	Layer / kind	Original section; status
S_{crypto}	Synthetic USD supply (DeFi)	0 mechanical (supply); flow	Ch.III 3.7 (p.10); defined; emergent-channel form (1 + X/Xbar) (ISSUES #29 S0-DECIDED)
S_{off}	Sanctions / asset-immobilisation intensity	0 mechanical (supply); state in [0,1]	4.2.1; CS5; CS9; Ch.7; Ch.8 (p.13); defined; placement (ISSUES #4) and log-safe form (ISSUES #28) both S0-DECIDED
δ_{off}	Sanctions attenuation coefficient	0 mechanical (supply); parameter	4.2.1; CS5; CS9 (p.13); OPEN value; introduced by ISSUES #28 S0-DECIDED
S_{shadow}	Shadow-bank private USD creation (component of S_coll)	0 mechanical (supply); flow	CS1; CS6 (p.31); defined; placement resolved (ISSUES #5 S0-DECIDED)
P_i	Policy rate (price sub-component of S_Fed)	0 mechanical (supply); sub-component	Ch.III 3.1 (p.9); defined
M_L	Fed balance sheet (quantity sub-component of S_Fed)	0 mechanical (supply); sub-component	Ch.III 3.1 (p.9); defined
D_{div}	Policy-divergence demand (ABSORBED into D_cap)	0 mechanical (demand) — placement OPEN; flow	CS3 (p.32); ABSORBED into D_cap (ISSUES #6 S0-DECIDED) - retained for concordance only
$\alpha_{\text{trans}}(\sigma_{\text{sys}})$	State-dependent demand elasticities (name-indexed)	1 state; function	Ch.I 2.1-2.2; Ch.5 5.1 (p.4); defined; index set and functional form (ISSUES #7 S0-DECIDED)
$\beta_{\text{Fed}}(\sigma_{\text{sys}})$	State-dependent supply elasticities (name-indexed)	1 state; function	Ch.I 2.1-2.2; Ch.5 5.1 (p.4); defined; index set and functional form (ISSUES #7 S0-DECIDED)
$\alpha_{\text{prec}}(\sigma)$	Safe-haven elasticity (on the USD-captured base s_US D_prec)	1 state; function	Ch.5 5.1 (p.27); defined; share moved inside the base (ISSUES #8 S0-DECIDED)
β_j^{crisis}	Crisis-regime supply elasticity	1 state; parameter	Ch.I 2.1-2.2; Ch.5 5.1 (p.27); OPEN value; introduced by ISSUES #7 S0-DECIDED
β_j^{calm}	Calm-regime supply elasticity	1 state; parameter	Ch.I 2.1-2.2; Ch.5 5.1 (p.27); OPEN value; introduced by ISSUES #7 S0-DECIDED
α_i^{crisis}	Crisis-regime demand elasticity	1 state; parameter	Ch.I 2.1-2.2; Ch.5 5.1 (p.27); OPEN value; introduced by ISSUES #7 S0-DECIDED

Symbol	Name	Layer / kind	Original section; status
α_i^{calm}	Calm-regime demand elasticity	1 state; parameter	Ch.I 2.1-2.2; Ch.5 5.1 (p.27); OPEN value; introduced by ISSUES #7 S0-DECIDED
σ^*	Shared regime threshold in sigma_sys	1 state; parameter	Ch.I 2.1-2.2; Ch.5 5.1 (p.27); OPEN value; introduced by ISSUES #7 S0-DECIDED
θ	Shared steepness of the regime logistic	1 state; parameter	Ch.I 2.1-2.2; Ch.5 5.1 (p.27); OPEN value; introduced by ISSUES #7 S0-DECIDED
$\Lambda(\sigma_{\text{sys}})$	Regime-weight logistic in sigma_sys	1 state; parameter	Ch.I 2.1-2.2; Ch.5 5.1 (p.27); OPEN value; introduced by ISSUES #7 S0-DECIDED
λ_1	Linear safe-haven coefficient	1 state; parameter	Ch.I 2.1 (p.4); defined; value OPEN; bounded because sigma_sys in [0,1] (ISSUES #9 S0-DECIDED)
λ_2	Quadratic (convexity) safe-haven coefficient	1 state; parameter	Ch.I 2.1-2.2 (p.4); defined; value OPEN; bounded because sigma_sys in [0,1] (ISSUES #9 S0-DECIDED)
$\kappa(\sigma, G_{\text{pol}}, R_{\text{US}})$	Convexity term (original)	3 reflexive; function	Ch.I 2.2-2.3 (p.5); superseded by kappa_system (E6); R_US/G_pol channel dropped in E6 — ISSUES #3
σ	Systemic stress (original)	1 state; state	Ch.I 1-2 (p.4); superseded
σ_{sys}	Systemic instability index	1 state; composite state	4.1; 5.2.1 (p.11); defined; bounded construction (ISSUES #9 S0-DECIDED); weights w OPEN
$\sigma_{\text{VIX/MOVE}}$	Endogenous financial stress	1 state; observable	4.1.1 (p.11); defined
X_{pand}	Exogenous (non-financial) shock trigger	1 state; state	4.1.2 (p.12); defined
F_{liq}	Inter-market liquidity friction	1 state; state	4.1.3 (p.12); defined; free proxy needed
w_σ	sigma_system weight (weight on sigma_VIX/MOVE)	1 state; parameter	4.1.3; 5.2.1 (p.27); OPEN value
w_X	sigma_system weight (weight on X_pand)	1 state; parameter	4.1.3; 5.2.1 (p.27); OPEN value
w_F	sigma_system weight (weight on F_liq)	1 state; parameter	4.1.3; 5.2.1 (p.27); OPEN value
$\Phi(G_{\text{pol}}, R_{\text{US}})$	Safe-haven trust (original)	2 structural; function	Ch.I 2.1; 2.3 (p.4); superseded

Symbol	Name	Layer / kind	Original section; status
Φ_{sys}	Systemic shelter credibility	2 structural; composite state	4.2; 5.2.2 (p.12); defined; E4 sign on M_trust is POSITIVE (ISSUES #1 S0-DECIDED)
s_{US}	Safe-flow market share (original)	2 structural; function	Ch.I 2.1; 2.3 (p.5); superseded
$s_{\text{US,sys}}$	Safe-flow market share (logistic)	2 structural; composite state	4.2; 5.2.2 (p.28); defined; coefficients OPEN
G_{pol}	Geopolitical power	2 structural; latent	Ch.I 2.3; 4.2.1 (p.13); defined; measurement protocol OPEN
R_{US}	Internal U.S. risk	2 structural; latent	Ch.I 2.3; 4.2.2; 4.2.5 (p.13); defined; scale fixed to [0,1] (ISSUES #26 S0-DECIDED); measurement protocol OPEN
R_{USD}	R_USD (CS5 token - typo for R_US)	2 structural; unknown	CS5 (p.33); TYPO for R_US (ISSUES #6 S0-DECIDED) - retained for concordance only
M_{trust}	Institutional trust	2 structural; latent	4.2.3 (p.13); defined; sign in E4 positive (ISSUES #1 S0-DECIDED); inflation channel specified (ISSUES #25 S0-DECIDED)
H_{cap}	Human-capital and IP stock	2 structural; latent	Ch.I 2.3; Ch.II 2.10 (p.8); defined; introduced by ISSUES #24 S0-DECIDED
m_E	E_inf coefficient in M_trust	2 structural; parameter	4.2.3; Ch.7 (p.13); OPEN value; introduced by ISSUES #25 S0-DECIDED
m_0	Intercept of M_trust	2 structural; parameter	4.2.3; Ch.7 (p.13); OPEN value; introduced by ISSUES #25 S0-DECIDED
T_{frag}	Geopolitical fragmentation index	2 structural / 5 secular; latent	4.2.4; 4.5.2 (p.14); defined; measurement protocol OPEN
E_{inf}	Inflation expectations	2 structural; observable	5.2.2; Ch.7 (E_infl) (p.28); defined; spelling unified; dual channel specified (ISSUES #25 S0-DECIDED)
ϕ_0	intercept of Phi_system	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
ϕ_G	G_pol coefficient in Phi_system (0.5 assumed in CS7)	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
ϕ_R	R_US coefficient in Phi_system	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value

Symbol	Name	Layer / kind	Original section; status
ϕ_M	M_trust coefficient in Phi_system (positive sign)	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value; sign fixed positive (ISSUES #1 S0-DECIDED)
ϕ_T	T_frag coefficient in Phi_system	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
γ_0	intercept of s_US logistic	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
γ_G	G_pol coefficient in s_US (1.0 assumed in CS7)	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
γ_R	R_US coefficient in s_US	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
γ_E	E_inf coefficient in s_US	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
γ_T	T_frag coefficient in s_US	2 structural; parameter	5.2.2; CS4; CS7 (p.28); OPEN value
κ_R	R_US flattening coefficient on kappa_sys	3 reflexive; parameter	5.2.2; CS4; CS7 (p.28); OPEN value; restored to E6 (ISSUES #3 S0-DECIDED)
κ_{sys}	Reflexive acceleration / convexity	3 reflexive; composite state	4.3; 5.2.3 (p.15); defined; explicit multiplicative form (ISSUES #3 S0-DECIDED)
κ_{base}	Baseline convexity	3 reflexive; latent	4.3.1 (p.15); defined
V_{USD}	Velocity of shock transmission	3 reflexive; latent	4.3.1 (p.16); defined
M_{shock}	Crisis memory	3 reflexive; state	4.3.1; Ch.8 'New Update' E11-E14 (p.16); defined; formulas live at end of Ch.8 — move to Layer 3
A_{meta}	Meta-reflexivity (Soros)	3 reflexive; latent	4.3.1 (p.16); defined
a_σ	sigma_sys coefficient in kappa_sys	3 reflexive; parameter	4.3; 5.2.3 (p.15); OPEN value; introduced by ISSUES #3 S0-DECIDED
a_A	A_meta coefficient in kappa_sys	3 reflexive; parameter	4.3; 5.2.3 (p.15); OPEN value; introduced by ISSUES #3 S0-DECIDED
a_M	M_shock coefficient in kappa_sys	3 reflexive; parameter	4.3; 5.2.3 (p.15); OPEN value; introduced by ISSUES #3 S0-DECIDED
a_V	V_USD coefficient in kappa_sys	3 reflexive; parameter	4.3; 5.2.3 (p.15); OPEN value; introduced by ISSUES #3 S0-DECIDED
a_H	H_soc coefficient in kappa_sys	3 reflexive; parameter	4.3; 5.2.3 (p.15); OPEN value; introduced by ISSUES #3a

Symbol	Name	Layer / kind	Original section; status
a_I	I_imp coefficient in kappa_sys	3 reflexive; parameter	4.3; 5.2.3 (p.15); OPEN value; introduced by ISSUES #3a
B_{phase}	Behavioral phase-lock index	3/4; consequence	4.3.2; 4.4.1 (p.17); defined with an equation (ISSUES #10 S0-DECIDED)
Ω	set of major past crises (e.g. 2008, 2020)	3 reflexive (memory); parameter/function	Ch.8 E11 (p.67); defined; symbol clashes resolved by the ISSUES #11 renames (S0-DECIDED)
S_k	severity of crisis k (2008 = 10/10)	3 reflexive (memory); parameter/function	Ch.8 E11 (p.67); defined; symbol clashes resolved by the ISSUES #11 renames (S0-DECIDED)
t_k	date of crisis k	3 reflexive (memory); parameter/function	Ch.8 E11 (p.67); defined; symbol clashes resolved by the ISSUES #11 renames (S0-DECIDED)
λ_f	Psychological forgetting rate	3 reflexive (memory); parameter/function	Ch.8 E12 (p.67); defined; renamed lambda -> lambda_f (ISSUES #11 S0-DECIDED)
$\rho_{\text{gen}}(t, t_k)$	generational decay: share of capital run by people who lived t_k; max(0, 1-(t-t_k)/T_career)	3 reflexive (memory); parameter/function	Ch.8 E13 (p.68); defined; symbol clashes resolved by the ISSUES #11 renames (S0-DECIDED)
T_{career}	average investor career span (20-25 y)	3 reflexive (memory); parameter/function	Ch.8 E13 (p.68); defined; symbol clashes resolved by the ISSUES #11 renames (S0-DECIDED)
α_{res}	Resonance coefficient in M_shockˆeff	3 reflexive (memory); parameter/function	Ch.8 E14 (p.68); defined; renamed alpha -> alpha_res (ISSUES #11 S0-DECIDED)
$\text{Sim}(t, t_k)$	similarity of current event to crisis k	3 reflexive (memory); parameter/function	Ch.8 E14 (p.68); defined; symbol clashes resolved by the ISSUES #11 renames (S0-DECIDED)
$M^{\text{eff}}_{\text{shock}}$	re-activated memory = M_shock (1 + alpha A_alg Sim)	3 reflexive (memory); parameter/function	Ch.8 E14 (p.68); defined; symbol clashes resolved by the ISSUES #11 renames (S0-DECIDED)

Symbol	Name	Layer / kind	Original section; status
H_{soc}	Homogeneity of collective consciousness	4 psycho-technical; state in [0,1]	4.4.1-4.4.2; 5.2.4 (p.19); defined; two-way feedback with I_imp closed via C_loop (ISSUES #13 S0-DECIDED); fixed point characterised (ISSUES #30 S0-DECIDED)
I_{imp}	Impulse efficiency index	4 psycho-technical; state in [0,1]	4.4.1-4.4.2; 5.2.5 (p.19); defined; two-way feedback with H_soc (ISSUES #13 S0-DECIDED); fixed point characterised (ISSUES #30 S0-DECIDED)
A_{alg}	Algorithmic reinforcement strength	4 psycho-technical; latent driver	4.4.1 (p.18); defined
C_{src}	Content source concentration	4 psycho-technical; observable driver	4.4.1 (p.19); defined
P_{AI}	Generative-AI penetration ratio	4 psycho-technical; driver	4.4.1 (p.19); defined
T_{exp}	Exposure time index	4 psycho-technical; driver	4.4.1 (p.19); defined
R_{div}	Residual diversity resistance	4 psycho-technical; defensive	4.4.1 (p.19); defined
V_{meme}	Viral meme transmission velocity	4 psycho-technical; consequence	4.4.1; 4.5.2; 5.2.6 (p.19); defined with an equation (ISSUES #10 S0-DECIDED)
b_P	P_AI coefficient in V_meme	4 psycho-technical; parameter	4.4.1; 4.5.2 (p.19); OPEN value; introduced by ISSUES #10 S0-DECIDED
b_A	A_alg coefficient in V_meme	4 psycho-technical; parameter	4.4.1; 4.5.2 (p.19); OPEN value; introduced by ISSUES #10 S0-DECIDED
v_0	Baseline meme velocity	4 psycho-technical; parameter	4.4.1; 4.5.2 (p.19); OPEN value; introduced by ISSUES #10 S0-DECIDED
S_{access}	speed of access (think -> click)	4 psycho-technical; driver	4.4.2; 5.2.5 (p.20); defined
R_{fric}	resistance friction (doubt, UX steps, confirmations)	4 psycho-technical; driver	4.4.2; 5.2.5 (p.20); defined

Symbol	Name	Layer / kind	Original section; status
C_{loop}	cognitive-loop shortening (think-decide-act)	4 psycho-technical; driver	4.4.2; 5.2.5 (p.20); defined; feedback closed (ISSUES #13 S0-DECIDED); admissibility g_0 >= 0, g_0 + g_1 <= 1 required (ISSUES #30 S0-DECIDED)
g_1	Slope of the reduced impulse map	4 psycho-technical; composite	4.4.1-4.4.2; 5.2.4-5.2.5 (p.20); defined; introduced by ISSUES #30 S0-DECIDED; with eta < 4 maintained (ISSUES #30 AMENDED) the contraction bound binds for every admissible g_1
g_0	Intercept of the reduced impulse map	4 psycho-technical; composite	4.4.1-4.4.2; 5.2.4-5.2.5 (p.20); defined; introduced by ISSUES #30 S0-DECIDED
K	Impulse scale constant	4 psycho-technical; composite	4.4.1-4.4.2; 5.2.4-5.2.5 (p.20); defined; introduced by ISSUES #30 S0-DECIDED
κ_A	Psycho-technical driver product	4 psycho-technical; composite	4.4.1-4.4.2; 5.2.4-5.2.5 (p.20); defined; introduced by ISSUES #30 S0-DECIDED
c_{L1}	H_soc coefficient in C_loop	4 psycho-technical; parameter	4.4.2; 5.2.5 (p.20); OPEN value; introduced by ISSUES #13 S0-DECIDED
c_{L0}	Intercept of C_loop	4 psycho-technical; parameter	4.4.2; 5.2.5 (p.20); OPEN value; introduced by ISSUES #13 S0-DECIDED
P_{habit}	programmed habit strength	4 psycho-technical; driver	4.4.2; 5.2.5 (p.20); defined
E_{sys}	embedded system coupling (social + wallet + commerce + finance)	4 psycho-technical; driver	4.4.2; 5.2.5 (p.20); defined
η	Steepness of the H_soc logistic	4 psycho-technical; parameter	4.4.2; 5.2.4 (p.20); OPEN value; uniqueness threshold characterised (ISSUES #30 S0-DECIDED); eta < 4 adopted as a maintained working assumption throughout Parts I-III (ISSUES #30 AMENDED S0), multiplicity confined to the standalone Ch.8 section

Symbol	Name	Layer / kind	Original section; status
$\Psi_{\text{str}}(t)$	Structural integrity coefficient	5 secular; state in [1-L,1]	4.5; 5.2.6 (p.23); defined; direction (ISSUES #2), L semantics (ISSUES #14) and normalisation (ISSUES #12) all S0-DECIDED
$\Psi_{\min}$	Lower limit of structural integrity	5 secular; parameter	4.5.4; 5.2.6 (p.25); defined; introduced by ISSUES #14 S0-DECIDED; value OPEN
C_{alt}	Alternative infrastructure	5 secular; latent driver	4.5.2 (p.24); defined
$L_{\max}$	Maximum loss of structural integrity	5 secular; parameter/function	4.5.4; 5.2.6 (p.25); defined; renamed L -> L_max, Psi_min = 1 - L_max (ISSUES #14 S0-DECIDED); value OPEN
k_Ψ	Erosion speed (dynamic)	5 secular; parameter/function	4.5.4 (p.25); calibrated
$k_{\Psi,\text{base}}$	Baseline erosion speed (old form of Psi)	5 secular; parameter/function	4.5.4 (p.25); defined; renamed (ISSUES #11 S0-DECIDED); value OPEN
θ_0	Erosion threshold (renamed from t_0)	5 secular; parameter/function	4.5.4 (p.25); defined; renamed t_0 -> theta_0 (ISSUES #11 S0-DECIDED)
f_{hard}	cumulative hard-driver buildup	5 secular; parameter/function	4.5.4 (p.25); OPEN value / naming (ISSUES #11, #14)
c_0	Log-linear intercept (= ln DXY_0)	synthesis; parameter	5.3 (p.29); defined; identification resolved (ISSUES #12 S0-DECIDED)
r_{real}	Real policy rate (Ch.7 lever)	policy; observable	Ch.7 Part II 2.1 (p.50); defined

Appendix B

Equation Compendium of the Original Paper

This appendix reproduces every displayed equation of the October 2025 paper in the canonical notation, with the defects noted where they were found. The corrected forms live in Chapters 2–10.

Equation compendium of the original paper

E1 — Foundational equation (Ch. I, §2.1, p. 4). The right edge of this equation is clipped in the original PDF (the bracket's closing brace and its exponent overflow the text width). The exponent $\kappa(\sigma, G_{\text{pol}}, R_{\text{US}})$ on the bracket is implied by §2.2 pillar 3 and confirmed by E9.

$$DXY = DXY_0 \left\{ \left[\frac{\mathcal{N}(\sigma)}{\mathcal{S}(\sigma)} \right]^{\kappa(\sigma, G_{\text{pol}}, R_{\text{US}})} \times \exp\Big[(\lambda_1 \sigma + \lambda_2 \sigma^2)\, \Phi(G_{\text{pol}}, R_{\text{US}}) \Big] \right\} + \epsilon, \tag{B.1}$$

where the numerator and denominator products are written out in the original as

$$\mathcal{N}(\sigma) = D_{\text{trans}}^{\alpha_1(\sigma)}\, D_{\text{debt}}^{\alpha_2(\sigma)}\, D_{\text{cap}}^{\alpha_3(\sigma)}\, (D_{\text{prec}})^{\alpha_4(\sigma)\, s_{\text{US}}(G_{\text{pol}}, R_{\text{US}})}\, D_{\text{res}}^{\alpha_5(\sigma)}\, D_{\text{hedge}}^{\alpha_6(\sigma)}\, D_{\text{gamma}}^{\alpha_7(\sigma)}\, D_{\text{net}}^{\alpha_8(\sigma)}\, D_{\text{micro}}^{\alpha_9(\sigma)}\, D_{\text{hcap}}^{\alpha_H(\sigma)}, \tag{B.2}$$

$$\mathcal{S}(\sigma) = S_{\text{Fed}}^{\beta_1(\sigma)}\, S_{\text{Tsy}}^{\beta_2(\sigma)}\, S_{\text{coll}}^{\beta_3(\sigma)}\, S_{\text{bank}}^{\beta_4(\sigma)}\, S_{\text{swap}}^{\beta_5(\sigma)}\, S_{\text{ext}}^{\beta_6(\sigma)}\, S_{\text{crypto}}^{\beta_7(\sigma)}. \tag{B.3}$$

Original gloss: DXY_0 is the baseline (integrating constant); ϵ an exogenous random shock; σ the stress state (VIX/MOVE); $\kappa(\sigma, G_{\text{pol}}, R_{\text{US}}) = \kappa(\sigma, G_{\text{pol}}, R_{\text{US}})$ with $\kappa'(\sigma) > 0$; $\lambda_2 \sigma^2$ guarantees convexity.

E2 — Seasonal proxy (Ch. II, §2.9, p. 8).

$$D_{\text{micro}} = 1.02 \quad \text{on the final day of each quarter (calendar dummy); 1 otherwise.} \tag{B.4}$$

E3 — Systemic instability index (§4.1, p. 11–12; §5.2.1, p. 27). Two forms are given. §4.1.3:

$$\sigma_{\text{sys}} \approx \Big(\sigma_{\text{VIX/MOVE}} + w_X\, X_{\text{pand}} \Big)\, (1 + w_F\, F_{\text{liq}}) \tag{B.5}$$

§5.2.1 (adds a weight on the base term):

$$\sigma_{\text{sys}} = \Big(w_\sigma\, \sigma_{\text{VIX/MOVE}} + w_X\, X_{\text{pand}}\Big) \times (1 + w_F\, F_{\text{liq}}) \tag{B.6}$$

with the generic statement $\sigma_{\text{sys}} = f(\sigma_{\text{VIX/MOVE}}, X_{\text{pand}}, F_{\text{liq}})$ in §4.1.

E4 — Systemic shelter credibility (§5.2.2, p. 28).

$$\Phi_{\text{sys}} = \phi_0 + \phi_G\, G_{\text{pol}} - \phi_R\, R_{\text{US}} - \phi_M\, M_{\text{trust}} - \phi_T\, T_{\text{frag}} \tag{B.7}$$

Note the sign on M_{trust}: as written, higher institutional trust *lowers* credibility, contradicting §4.2.3 ("M_{trust} forms the bedrock of Φ_{sys}"). See ISSUES #1. Qualitative statements in §4.2.2 and §4.2.4: $R_{\text{US}} \uparrow\Rightarrow \Phi_{\text{sys}} \downarrow\downarrow$, $s_{\text{US,sys}} \downarrow$, $\kappa_{\text{sys}} \downarrow$; $T_{\text{frag}} \uparrow\Rightarrow G_{\text{pol}}$ effectiveness $\downarrow$, $\Phi_{\text{sys}}, s_{\text{US,sys}} \downarrow$, and "$\Psi_{\text{str}}(t) \uparrow$ (erosion coefficient activation)" — the last is a direction error, see ISSUES #2.

E5 — Safe-flow market share (§5.2.2, p. 28).

$$s_{\text{US,sys}} = \frac{1}{1 + \exp\left(-\gamma_0 + \gamma_G\, G_{\text{pol}} - \gamma_R\, R_{\text{US}} - \gamma_E\, E_{\text{inf}} - \gamma_T\, T_{\text{frag}}\right)} \tag{B.8}$$

E6 — Reflexive acceleration, dual-amplified (§4.3.3, p. 17; §5.2.3, p. 28).

$$\kappa_{\text{sys}} = \underbrace{f(\kappa_{\text{base}}, V_{\text{USD}}, M_{\text{shock}}, A_{\text{meta}})}_{\text{baseline reflexive sensitivity}} \times \underbrace{H_{\text{soc}}}_{\text{homogeneity amplifier}} \times \underbrace{I_{\text{imp}}}_{\text{immediacy amplifier}} \tag{B.9}$$

$f(\cdot)$ is never specified; the original $\kappa(\sigma, G_{\text{pol}}, R_{\text{US}})$ dependence of E1 (and the "$-\kappa_R R_{\text{US}}$ term" invoked in Case Study 4) does not appear here. See ISSUES #3.

E7 — Homogeneity convergence (§4.4.2, p. 20; §5.2.4, p. 28).

$$H_{\text{soc}}(t) = \frac{1}{1 + \exp\left(-\eta\left((A_{\text{alg}} \cdot C_{\text{src}} \cdot P_{\text{AI}} \cdot T_{\text{exp}} \cdot I_{\text{imp}}) - R_{\text{div}}\right)\right)} \tag{B.10}$$

E7b — Impulse function (§4.4.2, p. 20–21; §5.2.5, p. 28).

$$I_{\text{imp}} = f(S_{\text{access}}, R_{\text{fric}}, C_{\text{loop}}, P_{\text{habit}}, E_{\text{sys}}) \qquad I_{\text{imp}} \approx \frac{S_{\text{access}} \cdot C_{\text{loop}} \cdot P_{\text{habit}} \cdot E_{\text{sys}}}{1 + R_{\text{fric}}} \tag{B.11}$$

The text says H_{soc} and I_{imp} "mutually reinforce one another", but E7 depends on I_{imp} while E7b does not depend on H_{soc} (one-way). ISSUES #13.

E8 — Structural erosion / cliff function (§4.5.4, p. 25; §5.2.6, p. 29). Old form:

$$\Psi_{\text{str}}(t) \approx 1 - L_{\max} \cdot f(C_{\text{alt}}, T_{\text{frag}}, k_{\Psi,\text{base}}, \theta_0) \tag{B.12}$$

Updated form (§4.5.4 writes $\approx$, §5.2.6 writes $=$):

$$\Psi_{\text{str}}(t) = 1 - L_{\max}\left[\frac{1}{1 + \exp\left(-k_\Psi \cdot (f_{\text{hard}} - \theta_0)\right)}\right] \tag{B.13}$$

Original gloss: f_{hard} = cumulative build-up of hard drivers; θ_0 (written t_0) = threshold at which they begin to bite; $k(\cdot)$ = erosion rate, a function of $H_{\text{soc}}, V_{\text{meme}}, I_{\text{imp}}$; "$L$: the lower limit of Ψ" — as written the lower limit is $1 - L$, not L (ISSUES #14). $\Psi = 1$ intact, $\Psi \to 0$ collapse (§4.5).

E9 — Final integrated equation (§5.1, p. 27).

$$DXY = \Psi_{\text{str}}(t)(t,\cdot)\cdot DXY_0 \cdot \left[\frac{\prod_i D_i^{\alpha_i(\sigma_{\text{sys}})}\cdot (D_{\text{prec}})^{\alpha_{\text{prec}}(\sigma)\cdot s_{\text{US,sys}}(\cdot)}}{\prod_j S_j^{\beta_j(\sigma_{\text{sys}})}}\right]^{\kappa_{\text{sys}}(\cdot)} \times \exp\Big[(\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma_{\text{sys}}^2)\,\Phi_{\text{sys}}(\cdot)\Big] + \epsilon \tag{B.14}$$

Note D_{prec} appears both inside $\prod_i$ (if i ranges over all ten) and as a separate factor; the index set must exclude *prec* from the product (ISSUES #7).

E10 — Log-linear (estimable) form (§5.3, p. 29).

$$\ln DXY = \ln\Psi_{\text{str}}(t)(\cdot) + c_0 + \kappa_{\text{sys}}(\cdot)\left[\sum_i \alpha_i(\sigma_{\text{sys}})\ln D_i - \sum_j \beta_j(\sigma_{\text{sys}})\ln S_j\right] + \Big(\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma_{\text{sys}}^2\Big)\,\Phi_{\text{sys}}(\cdot) + \epsilon \tag{B.15}$$

The $\alpha_{\text{prec}}(\sigma)\cdot s_{\text{US,sys}}$ structure of E9 is dropped here (ISSUES #8). Differenced versions used in the case studies (CS9 §3, p. 42; Ch. 7 §3.3, p. 57):

$$\Delta\ln DXY \approx \kappa_{\text{sys}}\left[\sum_i \alpha_i(\sigma)\,\Delta\ln D_i - \sum_j \beta_j(\sigma)\,\Delta\ln S_j\right] + \Big(\lambda_1\Delta\sigma_{\text{sys}} + \lambda_2\Delta\sigma_{\text{sys}}^2\Big)\,\Phi_{\text{sys}} \tag{B.16}$$

$$\Delta\ln DXY \approx \kappa_{\text{sys}}(H_{\text{soc}},\cdot)\left[\sum_i \alpha_i(\sigma_{\text{sys}})\,\Delta\ln D_i - \sum_j \beta_j(\sigma_{\text{sys}})\,\Delta\ln S_j\right] + \Big(\lambda_1\Delta\sigma_{\text{sys}} + \lambda_2\Delta\sigma_{\text{sys}}^2\Big)\,\Phi_{\text{sys}}(R_{\text{US}}(H_{\text{soc}})) \tag{B.17}$$

(E10c is clipped at the right margin in the original; the closing of $\Phi_{\text{sys}}(R_{\text{US}}(H_{\text{soc}}\ldots$ is not visible.) Proportional form used in CS8 §5 (p. 39): $DXY \propto (D^*/S^*)^{\kappa_{\text{sys}}} \times \exp[(\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma_{\text{sys}}^2)\Phi_{\text{sys}}]$.

E11–E14 — Crisis-memory formulas (Ch. 8 "New Update", p. 67–68; numbered (1)–(4) in the original).

$$M_{\text{shock}}(t) = \sum_{k\in\Omega} S_k\cdot e^{-\lambda_f(t-t_k)}\cdot\rho_{\text{gen}}(t,t_k) \tag{B.18}$$

$$\text{psychological decay:}\quad e^{-\lambda_f(t-t_k)} \tag{B.19}$$

$$\rho_{\text{gen}}(t,t_k) \approx \max\left(0,\ 1-\frac{t-t_k}{T_{\text{career}}}\right) \tag{B.20}$$

$$M_{\text{shock}}^{\text{eff}}(t) = M_{\text{shock}}(t)\times(1+\alpha_{\text{res}}\cdot A_{\text{alg}}\cdot\text{Sim}(t,t_k)) \tag{B.21}$$

Original gloss: Ω = set of major past crises ("e.g. {2008, 2020, 2029, …}" — 2029 is hypothetical); S_k severity (2008 = 10/10, 2001 = 6/10); λ_f forgetting rate (higher when V_{meme} is high); $T_{\text{career}} \approx$ 20–25 years; worked example $\rho_{\text{gen}}(t,t_k) \approx 0.6$ in 2018 and ≈ 0 in 2038 for $t_k = 2008$. Symbol clashes: the forgetting rate is written λ (vs. λ_1,λ_2) and the resonance coefficient α (vs. α_i) — renamed here; ISSUES #11.

E15 — Marginal effects of ΔG_{pol} (Case Study 7, p. 34–36).

$$\Delta\Phi_{\text{sys}} \approx \phi_G \cdot \Delta G_{\text{pol}} \tag{B.22}$$

$$\Delta s_{\text{US,sys}} \approx s_{\text{US}}(1 - s_{\text{US}}) \cdot \gamma_G \cdot \Delta G_{\text{pol}} \qquad \text{(derivative of the logistic)} \tag{B.23}$$

Worked numbers (assumptions): $s_{\text{US,sys}} = 0.70$, $\phi_G = 0.5$, $\gamma_G = 1.0$; Scenario 2: $\Delta G_{\text{pol}} = +0.05 \Rightarrow \Delta\Phi_{\text{sys}} = +0.025$, $\Delta s_{\text{US,sys}} = 0.70 \cdot 0.30 \cdot 1.0 \cdot 0.05 = +0.0105$ (70% → 71.05%).

E16 — Policy-impact statements written as relations (Ch. 7 Part II, p. 53–54).

$$S_{\text{swap}} \uparrow,\ S_{\text{coll}} \uparrow,\ S_{\text{bank}} \uparrow \Rightarrow\ \sigma_{\text{sys}} \downarrow,\ \kappa_{\text{sys}} \downarrow,\ F_{\text{liq}} \downarrow \tag{B.24}$$

$$D_{\text{trans}} \downarrow \Rightarrow \text{inflation eases}; \quad T_{\text{frag}} \downarrow \Rightarrow \text{no } \Psi_{\text{str}}(t) \text{ tipping}; \quad D_{\text{hcap}} \uparrow \Rightarrow \Phi_{\text{sys}} \uparrow$$

Layer map as stated in the original. §4.0 announces a "four-layer regulatory hierarchy" (4.1–4.4) but Chapter 4 has five sections (4.1–4.5); §5.0 and §5.4 use six strata: Layer 0 mechanical (D^*/S^*), Layer 1 state (σ_{sys}), Layer 2 structural ($\Phi_{\text{sys}}, s_{\text{US,sys}}$), Layer 3 reflexive ($\kappa_{\text{sys}}$), Layer 4 psycho-technical ($H_{\text{soc}}, I_{\text{imp}}$), Layer 5 secular ($\Psi_{\text{str}}(t)$). Case Study 8 calls the mechanical block "Layer 4". The six-stratum numbering of §5.4 is adopted as canonical (ISSUES #15).

Appendix C

Data Appendix

C.1 How to Reproduce This Book

Every number in this edition comes from a file in `data/` or from a web source recorded in `data/SOURCES.md` with its URL and retrieval date. Nothing is quoted from memory. `data/raw/` holds files exactly as downloaded and is never edited; `data/processed/` holds scripted transforms only.

Table C.1: The pipeline. Each script is runnable on its own and prints what it computed.

Script	What it does
`data/fetch_tier1.py`	downloads the FRED series of Tier 1
`data/fetch_audit.py`	the rescue audit's series: World Bank reserves with and without gold, `DGS2`, `DGS30`
`data/fetch_1971.py`	the pre-1973 series: bilateral rates from 1971-01, CPI, unemployment, policy rate, net exports
`data/extract_tier1.py`	tidies raw to processed. **Carries the `blank_is_zero` flag** added after the defect of Section 11.5.2
`data/build_sigma.py`	constructs σ_{sys} from its four percentile legs
`tools/gen_macros.py`	`registry/variables.csv` → `latex/macros.tex`
`tools/check_notation.py`	enforces macro-only notation; fails hard on a broken cross-reference in the built PDF
`tools/concordance_audit.py`	Chapter E's three passes
`tools/scorecard_dist.py`	Table 17.3's distribution

C.2 The Series This Book Uses

C.3 The Absences, and Whose They Are

This is the section a reader reproducing the book needs most, and it makes a distinction the rest of the literature usually does not: **some gaps are the world's and some are ours**. A series that

Table C.2: Held series, by the layer they serve. All are free and machine-readable.

Series	Serves	Coverage	Note
`DTWEXAFEGS`	DXY	2006-01–	the book's dollar throughout
`DTWEXBGS`	DXY	2006-01–	broad index, robustness only
`DTWEXM`	DXY	1973-01–2019	the only pre-2006 index held
`VIXCLS`	σ_{sys} base	1990-01–	
`DGS10`, realised vol	σ_{sys} base	1962–	MOVE substitute, Section 5.3
`gpr_daily` (GPRD)	X_{pand}	1985-01–	Caldara–Iacoviello
`STLFSI4`	F_{liq}	1993-12–	the friction leg
`DCOILBRENTEU`	D_{trans}	1987-05–	
`us_de_10y_spread`	D_{cap}	1960–	
`SWPT`	S_{swap}	2002-12–	
`WALCL`, `WTREGEN`	S_{Fed}, S_{Tsy}	2002–	
`RRPONTSYD`	S_{coll} leg	2003-02–	**see Section C.4**
`SOFR`, `IORB`	S_{coll}	2018–	
`TEDRATE`, `DTB3`	funding	1986–	
IMF COFER	D_{res}, $\Psi_{\text{str}}(t)$	1999-Q1–	levels and shares
BIS USD credit	D_{debt}	2000-Q1–	quarterly; see Section C.4
BIS FX turnover	T_{frag} subst.	triennial	
DefiLlama	D_{net}, S_{crypto}	2019–	
World Bank gold price	s_{US}	1960-01–	monthly
World Bank reserves	$\Psi_{\text{str}}(t)$	1960–	incl. and excl. gold
`T5YIFR`, `MICH`	E_{inf}	2003, 1978–	Proposition 18.2
`EXUSUK` *et al.*	1971	1971-01–	bilateral, Chapter 22
`NETEXP`	S_{ext}	1947-Q1–	

does not exist is a different problem from one that exists behind a licence, and both are different from one nobody has yet transcribed.

The count that matters. Of the twelve absences above, **five are ours** — SWIFT, weekly MMF, S_{bank}, world trade volume, stablecoins by issuer — and every one has a free, named source and a bounded amount of work. **Seven are the world's**: licensed, or nonexistent, or a construction nobody has performed. Section 17.8's research agenda is ordered on exactly that division, because the five are a to-do list and the seven are not.

C.4 One Defect Found, and What It Changed

Recorded here because a reader reproducing this book will meet the same trap.

`RRPONTSYD` is published with **blank entries on days the reverse-repo facility was not used** — 1,744 of 5,052 rows, 35 per cent of the series. The extraction pipeline read blanks as missing values and dropped them. **A day on which nobody used the facility is a day of zero usage.** Because every regression in Chapter 11 needs a complete row, one blank removed the whole month: 59 of 230 months, dropped *non-randomly*, since the blanks are precisely the calm stretches. Among them June to August 2008, the first half of the largest dollar rally in the sample.

Three conclusions changed when it was fixed. The I35 residuals moved from 45 and 57 per cent to **64 and 76**. The set of wrong-signed channels changed membership. And the test of whether a state-dependent κ_{sys} could carry the residual **reversed from fail to pass**.

Remark C.1 (How it was found, which is the part worth copying)**.** Not by inspection, and not by any check in `tools/`. A later test asked why the 2008 crisis window contained four months of data, and the answer was this. **It is the class of defect that produces confident wrong answers rather than visible failures**, and the only defence is to check sample sizes against calendars whenever a window looks thin. `extract_tier1.py` now carries a documented `blank_is_zero` flag so regeneration cannot reintroduce it.

C.5 Tiering

ISSUES #23 sorts every series into three tiers and the book holds to it. **Tier 1**: free, machine-readable, fetched — everything in Table C.2. **Tier 2**: free but requiring construction or transcription — SWIFT, S_{bank}, weekly D_{prec} proxies, world trade volume. **Tier 3**: licensed or nonexistent — MOVE, the cross-currency basis, dealer gamma, G_{pol}.

No Tier 3 series is substituted silently anywhere in this book. Where one is missing, the text names it, names the substitute, and states what the substitute cannot do. That rule is the reason Chapter 17's untestable share is 41.8 per cent rather than lower: the alternative was to proxy the unmeasurable and report a better score.

Table C.3: Every material absence, its kind, what was substituted, and what the substitution costs. "Ours" means the series is obtainable and this book has not obtained it.

What is missing	Whose	Substitute in place	What the substitution costs
US gold stock, oz	**world's**	none; reserves *excluding* gold is *not* used as one	Chapter 22's central mechanism — the gold outflow that closed the window — cannot be shown. `GOLDRESV` and `Q1568BUSQ027SNBR` both return HTTP **404**: a confirmed absence, not a failed download
ICE DXY	**world's** (licensed)	`DTWEXAFEGS` from 2006, `DTWEXM` from 1973	no index covers the Nixon shock, so Chapter 22 uses bilateral rates and says so. `raw/stooq_DXY_d.csv` was removed 2026-09-05: the response body was a Cloudflare bot-check page, three lines, no observations
SWIFT currency shares	**ours**	BIS triennial turnover + COFER, per ISSUES #23	the substitute is *adjacent, not approximate*: turnover is trading not settlement, and triennial not monthly. Free, but locked in roughly 180 monthly PDFs. **About an afternoon per year of data.** T_{frag} scores 1 partial and 4 untestable because of this
Dealer net gamma	**world's**	none	D_{gamma}'s card entirely. No adequate free substitute exists, and the book says so rather than proxying it
Weekly institutional MMF	**ours**	`IMFSL` monthly	D_{prec} at weekly-or-better frequency after 2021-02, and a proxy that does not invert in 2008. Section 11.3.4 records that `IMFSL` does invert in 2008, which is why $\alpha_{\text{prec}}(\sigma)$ is weakly identified and worse in the episode that matters most
Cross-currency basis	world's (licensed)	TED, bill–fed-funds, SOFR–IORB	**the single most valuable missing series in the book** (Remark 11.4). Every substitute prices *funding*; the mechanism is a *quantity*. The I35 residual test returned "neither" partly for this reason
MOVE	world's (licensed)	realised Treasury volatility	the bond leg of σ_{sys} is a construction, not the index the original names
S_{bank} (G-SIB SLR)	**ours**	none	the card entirely; 2 predictions, both untestable. FFIEC filings, must be assembled
G_{pol}, S_{off}, M_{trust}, H_{soc} family	world's	none	no provider; these are construction tasks. Between them they account for most of Table 17.3's untestable share
World trade volume	ours	Brent only	F3 shows the price leg of D_{trans}, not the volume leg. CPB World Trade Monitor is

Appendix D

Derivations and Proofs

Coverage

Every proposition in this book that requires a formal argument is proved in one of two places, and the division is deliberate. Chapter 8's three fixed-point results are stated there with *proof sketches* and proved in full in Section D.1 below; Chapter 10's three comparative statics are stated there without proof and proved in Section D.3. **Nothing is proved twice.** The propositions of Chapters 5, 9, 11, 13 to 15, 18 and 20 to 22 are empirical: they are established by a computation in `data/`, and each carries its own one-paragraph proof naming the script, which is the right place for it.

D.1 The Psycho-Technical Fixed Point

This section proves Proposition 8.1, Proposition 8.2 and Proposition 8.3, which Chapter 8 states with sketches. The setting is ISSUES #13 and ISSUES #30.

Setup. Write $\Lambda(x) = (1 + e^{-x})^{-1}$ for the standard logistic, and let

$$G(H) = g_0 + g_1 H, \qquad F(I) = \Lambda\big(\eta(\kappa_A I - R_{\text{div}})\big), \qquad T = F \circ G, \tag{D.1}$$

with $\kappa_A = A_{\text{alg}} C_{\text{src}} P_{\text{AI}} T_{\text{exp}} \in [0,1]$, $\eta > 0$, $R_{\text{div}} \in [0,1]$, and the admissibility restriction of Assumption 8.3,

$$g_0 \geq 0, \qquad g_0 + g_1 \leq 1. \tag{D.2}$$

Note two immediate consequences of Equation (D.2) used throughout: $g_1 \leq 1$, and $G([0,1]) \subseteq [0,1]$.

Three standard facts about the logistic are used without further comment: Λ is smooth and strictly increasing with range $(0,1)$; $\Lambda' = \Lambda(1-\Lambda)$; and Λ' is strictly unimodal, increasing on $(-\infty, 0)$, decreasing on $(0, \infty)$, even about 0, with maximum $\Lambda'(0) = \frac{1}{4}$.

D.1.1 Existence

Proof of Proposition 8.1. By Equation (D.2), G maps $[0,1]$ into $[0,1]$, and G is affine hence continuous. F is a composition of an affine map with Λ, hence continuous, and its range is contained in $(0,1)$ because Λ attains neither bound. Therefore $T = F \circ G$ is continuous on $[0,1]$ with $T([0,1]) \subseteq (0,1)$.

Define $\psi(H) = T(H) - H$ on $[0,1]$. Then ψ is continuous, $\psi(0) = T(0) > 0$ and $\psi(1) = T(1) - 1 < 0$, the strict inequalities holding because T takes values strictly inside $(0,1)$. By the intermediate value theorem ψ has a zero $H^* \in (0,1)$, which is a fixed point of T. No restriction on η, κ_A or g_1 was used. □

D.1.2 Uniqueness

Proof of Proposition 8.2. T is differentiable on $[0,1]$ and, by the chain rule,

$$T'(H) = F'\big(G(H)\big)\, g_1 = \eta\, \kappa_A\, g_1\, \Lambda'\big(u(H)\big), \qquad u(H) \equiv \eta(\kappa_A G(H) - R_{\text{div}}). \tag{D.3}$$

Every factor is non-negative, so $T' \geq 0$: T is non-decreasing and the iteration $H_{n+1} = T(H_n)$ is monotone, never oscillatory.

Since $\Lambda' \leq \frac{1}{4}$ with equality at $u = 0$,

$$\sup_{H\in[0,1]} \big|T'(H)\big| \;=\; \tfrac{1}{4}\,\eta\,\kappa_A\, g_1, \tag{D.4}$$

where the supremum is attained whenever $\kappa_A G(H) = R_{\text{div}}$ has a solution $H \in [0,1]$, and is otherwise approached at the nearer endpoint; in either case Equation (D.4) is the least upper bound, so the constant is sharp.

If $\eta\kappa_A g_1 < 4$ then $\sup|T'| < 1$, so by the mean value theorem $|T(H) - T(\tilde{H})| \leq \sup|T'| \cdot |H - \tilde{H}|$ for all $H, \tilde{H} \in [0,1]$: T is a contraction on the complete metric space $[0,1]$. The Banach fixed-point theorem gives a unique fixed point H^*, global attraction from any $H_0 \in [0,1]$, and the geometric error bound $|H_n - H^*| \leq (\sup|T'|)^n |H_0 - H^*|$.

Finally, $\kappa_A \leq 1$ by Assumption 8.2 and $g_1 \leq 1$ by Equation (D.2), so $\eta < 4$ implies $\eta\kappa_A g_1 < 4$. This is the maintained assumption Assumption 8.4. □

D.1.3 Multiplicity

Proof of Proposition 8.3. Assume $g_1 > 0$ and set

$$a \equiv \eta\kappa_A g_1 > 0, \qquad b \equiv \eta(\kappa_A g_0 - R_{\text{div}}), \tag{D.5}$$

so that $u(H) = aH + b$ is an increasing affine bijection from $[0,1]$ onto $[b, a+b]$, with inverse $H = (u-b)/a$. A point H is fixed for T if and only if $H = \Lambda(u(H))$, i.e. if and only if $u = u(H)$ satisfies

$$\varphi(u) \;\equiv\; b + a\,\Lambda(u) - u \;=\; 0. \tag{D.6}$$

Fixed points of T in $[0,1]$ correspond bijectively to zeros of φ in $[b, a+b]$.

At most three. $\varphi'(u) = a\Lambda'(u) - 1$. Because Λ' is strictly unimodal with maximum $\frac{1}{4}$, the equation $\Lambda'(u) = 1/a$ has no solution when $a < 4$, one solution when $a = 4$, and exactly two solutions $u_- < 0 < u_+$ when $a > 4$. Hence φ' has at most two sign changes, so φ has at most two turning points on $\mathbb{R}$. A continuously differentiable function of one variable with at most two turning points has at most three zeros, by Rolle's theorem: four zeros would force three interior turning points.

Exactly one when $a \leq 4$. Then $\varphi'(u) = a\Lambda'(u) - 1 \leq a/4 - 1 \leq 0$, with equality only at the single point $u = 0$ when $a = 4$, so φ is strictly decreasing; together with Proposition 8.1 it has exactly one zero. This recovers Proposition 8.2 by a second route.

Generically one or three. For $a > 4$, φ decreases on $(-\infty, u_-)$, increases on (u_-, u_+) and decreases on (u_+, ∞), with $\varphi(-\infty) = +\infty$ and $\varphi(+\infty) = -\infty$. The number of zeros is one if

$\varphi(u_-) > 0$ or $\varphi(u_+) < 0$, three if $\varphi(u_-) < 0 < \varphi(u_+)$, and two in the non-generic case where one of $\varphi(u_\pm)$ vanishes — a codimension-one condition on (a, b), i.e. a saddle-node bifurcation.

Stability classification. From Equation (D.3), $T'(H^*) = a\Lambda'(u^*)$, and from Equation (D.6), $\varphi'(u^*) = a\Lambda'(u^*) - 1 = T'(H^*) - 1$. So $T'(H^*) < 1$ exactly when $\varphi'(u^*) < 0$, i.e. exactly when φ crosses zero downward at u^*. With three zeros $u_1 < u_2 < u_3$ and the shape established above, φ crosses downward at u_1 and u_3 and upward at u_2. Hence $T' < 1$ at the outer two fixed points — which, with $T' \geq 0$, makes them locally asymptotically stable — and $T' > 1$ at the middle one, which makes it unstable. Ordering is preserved because $H \mapsto u$ is increasing. □

D.1.4 The amplification factor

Derivation of Equation (8.10). Let θ be any parameter entering T smoothly, and suppose $\eta\kappa_A g_1 < 4$ so that the fixed point $H^*(\theta)$ is unique. Since T is smooth and $\partial(T(H) - H)/\partial H = T' - 1 \neq 0$ at H^*, the implicit function theorem applies to $H - T(H; \theta) = 0$ and gives

$$\frac{\partial H^*}{\partial \theta} = \frac{\partial T/\partial \theta}{1 - T'(H^*)}. \tag{D.7}$$

Under the contraction condition $0 \leq T'(H^*) < 1$, so the multiplier $M = 1/(1 - T'(H^*))$ satisfies $M \geq 1$, with $M = 1$ only in the degenerate case $g_1 = 0$ (no feedback), and $M \to \infty$ as $T'(H^*) \to 1$. Because $M > 0$, the sign of $\partial H^*/\partial\theta$ equals the sign of the direct effect $\partial T/\partial\theta$, which is what makes the comparative statics of Equation (8.11) readable off E7 alone. □

D.2 From E9 to E10: the Log-Linearisation

Chapter 10 states Equation (10.4) and Equation (10.5) without deriving them. Both follow in two lines, and the value of writing them out is not the algebra but the three approximations that get buried in the "≈".

Derivation of Equation (10.4). Take logarithms of Equation (10.2), setting $\epsilon = 0$. Because the bracket enters as a power κ_{sys} and the exponential is already an exponential,

$$\begin{aligned} \ln DXY = {} & \ln \Psi_{\text{str}}(t) + \ln DXY_0 \\ & + \kappa_{\text{sys}}\Big[\sum_i \alpha_i(\sigma_{\text{sys}}) \ln D_i + \alpha_{\text{prec}}(\sigma) s_{\text{US,sys}} \ln D_{\text{prec}} - \sum_j \beta_j(\sigma_{\text{sys}}) \ln S_j\Big] \\ & + (\lambda_1 \sigma_{\text{sys}} + \lambda_2 \sigma_{\text{sys}}^2)\Phi_{\text{sys}}. \end{aligned} \tag{D.8}$$

The multiplicative architecture of Section 2.1 is exactly what makes this exact rather than approximate: every modulator enters as a product or a power, so logarithms separate them additively and no term is linearised. Equation (10.4) is Equation (D.8). □

Derivation of Equation (10.5)*, and its three approximations.* Difference Equation (D.8) between t and $t - 1$. The constant $\ln DXY_0$ drops. The remaining terms give

$$\begin{aligned} \Delta \ln DXY = {} & \Delta \ln \Psi_{\text{str}}(t) + \kappa_{\text{sys}}\Big[\sum_i \alpha_i(\sigma_{\text{sys}}) \Delta \ln D_i - \sum_j \beta_j(\sigma_{\text{sys}}) \Delta \ln S_j\Big] \\ & + \Delta\kappa_{\text{sys}} \cdot \ln\left(\frac{D^*}{S^*}\right) + \Delta\big[(\lambda_1 \sigma_{\text{sys}} + \lambda_2 \sigma_{\text{sys}}^2)\Phi_{\text{sys}}\big] + (\text{elasticity-change terms}). \end{aligned} \tag{D.9}$$

Equation (10.5) is Equation (D.9) with three things done to it.

One. The elasticities $\alpha_i(\sigma_{\text{sys}})$ and $\beta_j(\sigma_{\text{sys}})$ are held fixed across the window, which kills the last term. Defensible over a week; not over a year.

Two. The cross-term $\Delta\kappa_{\text{sys}} \cdot \ln(D^*/S^*)$ is dropped as second order. It is second order *only if* κ_{sys} *moves little* — and Chapter 7 argues κ_{sys} moves most in exactly the crises the case studies examine. **This is the approximation most likely to be doing damage**, and Section 17.6.2's finding that the κ_{sys} rival passes its regime test is evidence that it is.

Three. Φ_{sys} is held fixed, so $\Delta[(\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma^2_{\text{sys}})\Phi_{\text{sys}}]$ becomes $(\lambda_1\Delta\sigma_{\text{sys}} + \lambda_2\Delta\sigma^2_{\text{sys}})\Phi_{\text{sys}}$, which is why every case study recovers only the products $\lambda_1\Phi_{\text{sys}}$ and $\lambda_2\Phi_{\text{sys}}$ and never λ or Φ_{sys} separately. Φ_{sys} **is unidentified by construction, not by bad luck**, and Table 17.3's three-of-three untestable verdict on it is this fact reported downstream.

Finally, $\Delta\sigma^2_{\text{sys}}$ denotes Δ applied to the squared term, $\sigma_{\text{sys}}[t]^2 - \sigma_{\text{sys}}[t-1]^2$, and *not* $(\Delta\sigma_{\text{sys}})^2$. The two differ in sign when stress falls, and the case studies use the former. □

D.3 Comparative Statics of the Assembled Equation

Chapter 10 states three propositions without proof. They are proved here.

D.3.1 The sign-change condition

Proof of Proposition 10.3. The dollar's response to stress, holding everything else fixed, is the derivative of the safe-haven exponent in Equation (D.8):

$$\frac{\partial \ln DXY}{\partial \sigma_{\text{sys}}} = (\lambda_1 + 2\lambda_2\sigma_{\text{sys}})\,\Phi_{\text{sys}}. \tag{D.10}$$

$\Phi_{\text{sys}} > 0$ by Assumption 2.5, so the sign is the sign of $\lambda_1 + 2\lambda_2\sigma_{\text{sys}}$. This is affine and increasing in σ_{sys} when $\lambda_2 > 0$. An affine increasing function changes sign on $[0, 1]$ if and only if it is negative at 0 and positive at 1, that is if and only if

$$\lambda_1 < 0 \qquad \text{and} \qquad \lambda_1 + 2\lambda_2 > 0, \tag{D.11}$$

with the turning point at $\sigma^\dagger_{\text{sys}} = -\lambda_1/(2\lambda_2)$. Hence $\lambda_1 < 0$ is *necessary*. Convexity alone — $\lambda_2 > 0$ with $\lambda_1 \geq 0$ — makes Equation (D.10) strictly increasing but never negative, so it steepens an already-positive response and cannot reverse it. □

Remark D.1 (What the proof costs the model empirically)**.** The condition is a *sign* restriction on one coefficient, which is why Chapter 5 could test it directly and why the test is decisive rather than a matter of degree: across 452 specifications, none produced a calm-regime slope negative and distinguishable from zero. Section 17.6.2 adds the second nail. Attenuation correction scales a coefficient and its standard error by the same factor and so leaves every t-statistic unchanged, and it never flips a sign — so no measurement-error argument, of any strength, can supply the $\lambda_1 < 0$ that Equation (D.11) requires.

D.3.2 The discontinuity Layer 4 induces in Layer 5

Proof of Proposition 10.2. $\Psi_{\text{str}}(t)$ is smooth in each of its own arguments. The discontinuity is inherited, not intrinsic, and it arrives through the composition

$$\text{drivers} \longmapsto H^*_{\text{soc}} \longmapsto k_\Psi \longmapsto \Psi_{\text{str}}(t). \tag{D.12}$$

By Proposition 8.3 the fixed-point map T admits up to three solutions, and the outer two are stable. At a saddle-node, two of the three merge and annihilate, so the stable selection H^*_{soc} *jumps* to the surviving branch: H^*_{soc} is a discontinuous function of the parameters even though T is smooth in all of them. Since k_Ψ is strictly increasing in H_{soc} and $\Psi_{\text{str}}(t)$ is strictly monotone in k_Ψ for $f_{\text{hard}} \neq \theta_0$, both compositions transmit the jump.

The direction follows from Equation (10.1). Writing $\Psi_{\text{str}}(t) = 1 - L_{\max}\Lambda(k_\Psi(f_{\text{hard}} - \theta_0))$, a rise in k_Ψ moves $\Lambda(k_\Psi(f_{\text{hard}} - \theta_0))$ *away* from $\frac{1}{2}$ in the direction of $\operatorname{sign}(f_{\text{hard}} - \theta_0)$. So $\Psi_{\text{str}}(t)$ jumps *up* when $f_{\text{hard}} < \theta_0$, jumps *down* when $f_{\text{hard}} > \theta_0$, and is unmoved at $f_{\text{hard}} = \theta_0$, where the argument is zero for every k_Ψ. □

Remark D.2 (The counter-intuitive half)**.** Below the threshold, a *rise* in belief homogeneity *raises* structural integrity. That is not a modelling error: k_Ψ scales the steepness of a logistic evaluated below its midpoint, and a steeper logistic below its midpoint is closer to its upper asymptote. In the model's terms, a homogeneous world that has not yet accumulated hard drivers converges on the incumbent faster. It is worth flagging because it is the one comparative static in Chapter 10 whose sign most readers will guess wrong.

D.3.3 The two thresholds are distinct objects

Proof of Proposition 10.1. Three claims, each immediate from the definitions.

Disjoint drivers. The Chapter 8 fold is a property of $T = F \circ G$ and depends on η, $\kappa_A = A_{\text{alg}} C_{\text{src}} P_{\text{AI}} T_{\text{exp}}$, R_{div}, g_0 and g_1. The Chapter 9 threshold is the location θ_0 at which $f_{\text{hard}}(C_{\text{alt}}, T_{\text{frag}})$ crosses, and depends on C_{alt}, T_{frag} and θ_0. The two lists share no symbol.

One-way coupling. k_Ψ takes H_{soc}, V_{meme} and I_{imp} as arguments, so Layer 4 enters Layer 5 through the erosion rate. Nothing in T takes C_{alt}, T_{frag}, θ_0 or $\Psi_{\text{str}}(t)$ as an argument, so there is no return path. The coupling is a directed edge, not a loop.

Only one is a tipping point. A tipping point requires that the crossing be irreversible under a reversal of the driver. At the Chapter 8 fold this holds: after a saddle-node the system sits on a different stable branch, and restoring the parameter does not restore the branch until the *other* fold is reached — hysteresis. At the Chapter 9 threshold it does not: θ_0 is an inflection of a static, strictly monotone logistic in f_{hard}, so $\Psi_{\text{str}}(t)$ is a bijection of f_{hard} and retracing f_{hard} retraces $\Psi_{\text{str}}(t)$ exactly. □

Remark D.3 (The finding this proof produced)**.** The third claim is the one that matters and it is not what the original supposes. The original places the irreversibility in Layer 5, in the erosion cliff. The proof puts it in Layer 4, in the fold — **the only irreversible object in the model is the belief equilibrium**, and Layer 5's cliff, for all its rhetorical weight, is a reversible monotone map. Section 10.1 states this in the assembly and Chapter 17 carries it.

D.4 The Logistic Derivative Behind CS7

Chapter 15's CS7 turns on a fact about logistic functions that inverts the obvious intuition, and Chapter 6 states it without derivation.

Derivation of Equation (6.9). E5 writes the safe-flow share as a logistic in the structural drivers, $s_{\text{US,sys}} = \Lambda(z)$ with $z = \gamma_G G_{\text{pol}} + \gamma_T T_{\text{frag}} + \gamma_R R_{\text{US}} + \gamma_E \dots$ collecting them. Then

$$\frac{\partial s_{\text{US,sys}}}{\partial G_{\text{pol}}} = \Lambda'(z)\,\gamma_G = s_{\text{US,sys}}(1 - s_{\text{US,sys}})\gamma_G. \tag{D.13}$$

The factor $s(1-s)$ is maximised at $s=\frac{1}{2}$, where it equals $\frac{1}{4}$, and falls to zero at both $s=0$ and $s=1$. □

Remark D.4 (Why this is CS7's whole argument)**.** Equation (D.13) says that an increment of geopolitical capability moves the safe-flow share *least* for the power whose share is already largest. At the baseline $s_{\mathrm{US,sys}}=0.70$ the factor is 0.21; at $s_{\mathrm{US,sys}}=0.50$ it would be 0.25; at 0.95 it is 0.0475, a fifth of its value at the midpoint. **The incumbent gets the smallest return on the marginal carrier**, which is the opposite of the intuition that dominance compounds. It is a proper implication of the model's own functional form, it is one of the five results Section 1.5.1 lists as surviving, and Proposition 15.1 now bounds it from above empirically: whatever γ_G is, one carrier's worth of ΔG_{pol} times $s(1-s)\gamma_G$ is smaller than the dollar's ordinary two-month noise.

D.5 Properties of the Cliff Function

Chapter 9's test rests on a curvature property of E8 that the chapter asserts and does not derive.

Derivation of Equation (9.4). Write $\Psi_{\mathrm{str}}(t)=1-L_{\max}\Lambda(u)$ with $u=k_\Psi(f_{\mathrm{hard}}-\theta_0)$ and $L_{\max}\in(0,1]$. Then

$$\frac{\partial\Psi_{\mathrm{str}}(t)}{\partial f_{\mathrm{hard}}}=-L_{\max}k_\Psi\,\Lambda(u)\big(1-\Lambda(u)\big)<0, \tag{D.14}$$

so $\Psi_{\mathrm{str}}(t)$ is strictly decreasing in the hard drivers for every $k_\Psi>0$: erosion is monotone, always. Differentiating again and using $\Lambda''=\Lambda(1-\Lambda)(1-2\Lambda)$,

$$\frac{\partial^2\Psi_{\mathrm{str}}(t)}{\partial f_{\mathrm{hard}}^2}=-L_{\max}k_\Psi^2\,\Lambda(u)\big(1-\Lambda(u)\big)\big(1-2\Lambda(u)\big). \tag{D.15}$$

Below the threshold $f_{\mathrm{hard}}<\theta_0$ gives $u<0$, hence $\Lambda(u)<\frac{1}{2}$ and $1-2\Lambda(u)>0$, so Equation (D.15) is **strictly negative**: the decline *accelerates*. Above the threshold the sign reverses and the decline decelerates toward the floor $1-L_{\max}$. □

Remark D.5 (This is the whole content of the cliff test)**.** Equation (D.15) is why Proposition 9.2 takes the form it does. E8 does not merely predict that the dollar's reserve share falls — monotonicity in Equation (D.14) makes that automatic and therefore uninformative. It predicts a *negative second derivative below the threshold*, and that is a testable curvature restriction on an observable series.

Measured on COFER the quadratic term is $+0.0115$ pp/yr^2 with $t=+3.84$: positive, so the wrong sign. Measured on a reserve aggregate that includes monetary gold it is -0.0211 ($t=-2.59$) at market value and -0.0670 ($t=-15.34$) at constant 2015 prices — the sign Equation (D.15) requires (Proposition 9.4). **The model was being tested on the wrong aggregate**, which Section 17.6.3 established and Section 9.7.1 carries. It remains unfalsifiable for the separate reason that f_{hard} has no measure, so no reading of the curvature can distinguish "below the threshold" from "far from it".

D.6 Stability of the Safe-Haven Exponential

One loose end, recorded because Chapter 10 raises it and nowhere resolves it.

Boundedness of the safe-haven term. $\sigma_{\text{sys}} \in [0,1]$ by Assumption 2.2 and ISSUES #9's logistic construction, and $\Phi_{\text{sys}} \in [0,1]$ by Assumption 2.5. The exponent is therefore a continuous function on a compact set and attains bounds

$$\exp\big[\min(0,\lambda_1)\big] \;\le\; \exp\big[(\lambda_1\sigma_{\text{sys}} + \lambda_2\sigma_{\text{sys}}^2)\Phi_{\text{sys}}\big] \;\le\; \exp\big[\max(0,\lambda_1+\lambda_2)\big]. \tag{D.16}$$

The term cannot diverge for any finite λ_1, λ_2. □

Remark D.6 (The boundedness is doing less work than it appears)**.** It is worth stating what Equation (D.16) does *not* rule out. The bound depends on σ_{sys} being bounded, which is a construction choice — ISSUES #9 puts σ_{sys} through a logistic precisely so that it is. **The model is stable because it was built to be, not because the world is**, and the substantive question is whether a bounded state variable can represent a crisis at all. Section 14.1.7 is where that bites: an episode beginning at $\sigma_{\text{sys}} = 0.767$ has little room left in the convex term, so the mechanism the architecture designates for crises is weakest at the values crises occupy (ISSUES #37). Boundedness buys mathematical safety and costs explanatory range, and Equation (D.16) is where the trade is visible.

Appendix E

Concordance: Original Sections to This Edition

E.1 What This Appendix Is For

Chapter 1 states the rule this book was written under: every sentence of the October 2025 paper must survive, possibly reworded, never dropped. This appendix is the check on that rule, and it is the only place where the check is performed rather than asserted.

The concordance was drafted at the start of the project. The book has since been restructured across thirty working sessions — chapters split, targets revised twice, case studies moved between chapters, one proposition withdrawn and restated. **A mapping written before that work cannot be assumed to have held through it.** So the table below is not the original plan reprinted; it is the output of `tools/concordance_audit.py`, which extracts every section heading of the source, assigns it to a row, and checks that the target chapter exists and carries the content.

E.2 Method

Three passes, each stricter than the last.

Pass 1 — coverage. Every heading of the original is located by page and matched to a concordance row. **This is the weakest of the three** and is reported as such: the rows tile pages 1–68 contiguously, so every heading lands in one by construction. A pass here means only that no page was forgotten.

Pass 2 — target exists. For each row, the target chapter file must exist, must not carry `STATUS: STUB`, and must be non-trivial.

Pass 3 — content. The real test. Sixty-one of the original's headings name a model variable. For each, the corresponding registry macro must appear in the chapter the concordance sends it to. A row can pass Passes 1 and 2 and fail here, and that is the failure mode the audit exists to catch: a heading assigned to a chapter that never discusses it.

Remark E.1 (A correction the audit forced on itself)**.** Pass 3 initially reported three failures — Φ_{sys} and $s_{\text{US,sys}}$ absent from Chapter 5, and D_{net} absent from Chapter 14. All three were artefacts

of the audit, not of the book. Two concordance rows overlap wherever one page of the original carries two sections: page 12 holds both §4.1 and §4.2, and page 33 holds both CS5 and CS6. The first version of the script took the *first* matching row and so sent §4.2 to Chapter 5 instead of Chapter 6. Corrected to test against every matching row. Verified by hand afterwards: Φ_{sys} appears 42 times in Chapter 6 and D_{net} six times in Chapter 15. **The failures were in the checker.**

E.3 Result

Table E.1: The three passes.

Pass	Checked	Failed	
1. Every heading has a row	115	**0**	
2. Every target is written	33 rows	**0**	no stubs remain in the book
3. Every named variable is in its target	61	**0**	after Remark E.1

Nothing in the original is orphaned. That is the finding, and it is a finding rather than an assumption because the third pass could have failed and did not.

E.4 The Concordance

E.5 Where Content Moved, and Why

Four rows are not one-to-one, and a reader chasing a passage should know which.

E.6 Two Equations Without a Standalone Home

The heading audit passes, but equations need their own check, and two of the sixteen in `registry/equations.md` have no separately numbered display in a chapter. Both are accounted for and neither is lost.

E.7 What This Appendix Does Not Certify

Stated plainly, because a concordance that passes is easy to over-read.

It certifies location, not treatment. That G_{pol} appears in Chapter 6 does not mean the original's account of G_{pol} survived intact — ISSUES #33 records that this book changed how G_{pol} is scored, and Section 14.3.4 records that its three capacities were found to move in opposite directions. The claim is preserved; the assessment of it is not the original's.

It certifies headings and named variables, not every sentence. Sixty-one of 115 headings name a variable and are checked mechanically. The remainder are narrative and were checked by the session that wrote the chapter, which is a weaker guarantee.

Table E.2: Original to this edition. "Headings" counts the source's own section headings in that page range.

Original	Pages	Now at	Head.	
Disclaimer §1–6	1–3	front matter	2	verbatim, unaltered
Ch.I Abstract	3	front matter	—	
Ch.I §1 Introduction	3–4	Chapter 1	4	
Ch.I §2, 2.1–2.3 architecture	4–6	Chapter 2	7	
Ch.II §2.1–2.10 demand	6–8	Chapter 3	12	ten D channels
Ch.III §3.1–3.7 supply	8–10	Chapter 4	11	seven S channels
§4.0 intro, four-layer list	11	Chapter 2	4	the layer map
§4.1–4.1.3 state	11–12	Chapter 5	7	
§4.2–4.2.6 structural	12–15	Chapter 6	11	
§4.3–4.3.4 reflexive	15–18	Chapter 7	9	
§4.4–4.4.3 psycho-technical	18–22	Chapter 8	9	incl. the p.22 impact table
§4.5–4.5.4 secular	23–26	Chapter 9	9	
§5.0, 5.1 (E9), 5.2.1–5.2.6	26–29	Chapter 10	12	functions also in Ch.5–9
§5.3, 5.3.1–5.3.2	29–30	Chapter 11	5	
§5.4 six strata	30	Chapter 2, Chapter 10	2	
CS1, CS2	31–32	Chapter 13 §1–2	9	
CS3, CS4	32	Chapter 13 §3–4	5	
CS5	33	Chapter 14 §1	4	
CS6	33	Chapter 15 §1	4	
CS7 (§1–6, scenarios A–C, 1–2, E15)	34–37	Chapter 15 §2	4	
CS8 (§1–5, scenarios A–D)	37–40	Chapter 15 §3	2	
CS9 (§1–3, block table, scenarios 1–4)	40–43	Chapter 14 §2	2	
CS10	43–44	Chapter 16 §1	3	
CS11 (phases 1–3)	44–47	Chapter 16 §2	3	
Ch.7 Part I	47–50	Chapter 18	3	
Ch.7 Part II §2.1–2.3	50–54	Chapter 19	6	
Ch.7 Part III §3.1–3.3	54–57	Chapter 20	4	
Ch.8 Part 0	58–59	Chapter 21 §0–1	3	
Ch.8 Part 1 §1.1–1.2	59–62	Chapter 21 §2–3	5	
Ch.8 Part 2 §2.1–2.3	62–64	Chapter 21 §4–5	4	
Ch.8 Part 3 §3.1–3.3	65–67	Chapter 22	5	
Ch.8 "New Update" M	67–68	Chapter 7 §4	1	and Chapter B

Table E.3: The four non-trivial mappings.

Original material	Where it went, and the decision behind it
§4.0's four-layer list	Moved *forward* to Chapter 2, because the architecture chapter needs the layer map before any layer is built. The original states it at the head of its Chapter 4; this edition states it at the head of the book
§5.4's six strata	Split: the enumeration is in Chapter 2, the assembled equation and its comparative statics in Chapter 10 §7
"New Update" M_{shock}	Moved *backward* from the original's Chapter 8 to Chapter 7 §4, where κ_{sys} is built — M_{shock} is one of its seven drivers and belongs with them (ISSUES #16). The original appends it as an afterthought at the end of the paper
CS5 and CS9	Both to Chapter 14 although the original separates them by seven pages, because both are geo-economic coercion episodes. CS6, CS7 and CS8 go to Chapter 15 as structural shifts. The grouping is thematic and is the one place the edition reorders case studies

Table E.4: Equations whose chapter treatment differs from the original's.

Equation	Original form	This edition
E3a, E3b	two variants of the σ_{sys} composite, with and without w_σ (pp. 12 and 27)	**consolidated** into one display, Equation (5.1), under ISSUES #19. Section 5.2 states both variants and says which is adopted as canonical and why. Both survive in Chapter B as `eq:orig:E3a` and `eq:orig:E3b`
E12	the psychological-decay factor $e^{-\lambda_f(t-t_k)}$, displayed separately	**a factor inside** Equation (7.2) rather than a display of its own, with its own paragraph of discussion in Chapter 7 — including the rename from the original's λ, which collided with λ_1 and λ_2 (ISSUES #11)
E16	policy relations, written as arrows	rendered as Table 19.3, a table rather than an equation, because arrows between unmeasured variables are a relation and not a formula

It says nothing about whether any claim is true. Chapter 17 is where that is settled, and it settles a good deal of it against the model. **A complete concordance and a** 24.5 **per cent hit rate are consistent, and this book contains both.**

Bibliography

[1] Tobias Adrian and Hyun Song Shin. Liquidity and leverage. *Journal of Financial Intermediation*, 19(3):418–437, 2010.

[2] Gara Afonso, Marco Cipriani, Adam Copeland, Anna Kovner, Gabriele La Spada, and Antoine Martin. The market events of mid-september 2019. *Federal Reserve Bank of New York Economic Policy Review*, 27(2), 2021.

[3] Shekhar Aiyar, Jiaqian Chen, Christian H. Ebeke, Roberto Garcia-Saltos, Tryggvi Gudmundsson, Anna Ilyina, Alvar Kangur, Tansaya Kunaratskul, Sergio L. Rodriguez, Michele Ruta, Tatjana Schulze, Gabriel Soderberg, and Juan P. Trevino. Geo-economic fragmentation and the future of multilateralism. IMF Staff Discussion Note SDN/2023/001, International Monetary Fund, 2023.

[4] Iñaki Aldasoro and Torsten Ehlers. The geography of dollar funding of non-US banks. *BIS Quarterly Review*, pages 15–26, December 2018.

[5] Serkan Arslanalp, Barry Eichengreen, and Chima Simpson-Bell. The stealth erosion of dollar dominance and the rise of nontraditional reserve currencies. *Journal of International Economics*, 138, 2022.

[6] Stefan Avdjiev, Wenxin Du, Cathérine Koch, and Hyun Song Shin. The dollar, bank leverage, and deviations from covered interest parity. *American Economic Review: Insights*, 1(2):193–208, 2019. doi: 10.1257/aeri.20180322.

[7] Marina Azzimonti. Partisan conflict and private investment. *Journal of Monetary Economics*, 93:114–131, 2018. doi: 10.1016/j.jmoneco.2017.10.007.

[8] Saleem Bahaj and Ricardo Reis. Central bank swap lines: Evidence on the effects of the lender of last resort. *Review of Economic Studies*, 89(4):1654–1693, 2022. doi: 10.1093/restud/rdab074.

[9] Scott R. Baker, Nicholas Bloom, and Steven J. Davis. Measuring economic policy uncertainty. *Quarterly Journal of Economics*, 131(4):1593–1636, 2016. doi: 10.1093/qje/qjw024.

[10] Eytan Bakshy, Solomon Messing, and Lada A. Adamic. Exposure to ideologically diverse news and opinion on facebook. *Science*, 348(6239):1130–1132, 2015. doi: 10.1126/science.aaa1160.

[11] Bank for International Settlements. Triennial central bank survey of foreign exchange and over-the-counter derivatives markets in 2025. OTC foreign exchange turnover in April 2025, 2025.

[12] Brad M. Barber, Xing Huang, Terrance Odean, and Christopher Schwarz. Attention-induced trading and returns: Evidence from Robinhood users. *Journal of Finance*, 77(6):3141–3190, 2022. doi: 10.1111/jofi.13183.

[13] Nicholas Barberis, Andrei Shleifer, and Robert Vishny. A model of investor sentiment. *Journal of Financial Economics*, 49(3):307–343, 1998. doi: 10.1016/S0304-405X(98)00027-0.

[14] Ben S. Bernanke. *21st Century Monetary Policy: The Federal Reserve from the Great Inflation to COVID-19*. W. W. Norton, New York, 2022. ISBN 9781324020462.

[15] Alan S. Blinder. *A Monetary and Fiscal History of the United States, 1961–2021*. Princeton University Press, Princeton, 2022. ISBN 9780691238388.

[16] Michael D. Bordo and Barry Eichengreen, editors. *A Retrospective on the Bretton Woods System: Lessons for International Monetary Reform*. University of Chicago Press, Chicago, 1993.

[17] Scott A. Brave and R. Andrew Butters. Monitoring financial stability: A financial conditions index approach. *Economic Perspectives, Federal Reserve Bank of Chicago*, 35(1), 2011.

[18] Scott A. Brave and R. Andrew Butters. Diagnosing the financial system: Financial conditions and financial stress. *International Journal of Central Banking*, 8(2), 2012.

[19] Markus K. Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009. doi: 10.1093/rfs/hhn098.

[20] Valentina Bruno and Hyun Song Shin. Cross-border banking and global liquidity. *Review of Economic Studies*, 82(2):535–564, 2015. doi: 10.1093/restud/rdu042.

[21] Ricardo J. Caballero, Emmanuel Farhi, and Pierre-Olivier Gourinchas. The safe assets shortage conundrum. *Journal of Economic Perspectives*, 31(3):29–46, 2017. doi: 10.1257/jep.31.3.29.

[22] Dario Caldara and Matteo Iacoviello. Measuring geopolitical risk. *American Economic Review*, 112(4):1194–1225, 2022. doi: 10.1257/aer.20191823.

[23] Alex Chinco, Adam D. Clark-Joseph, and Mao Ye. Sparse signals in the cross-section of returns. *Journal of Finance*, 74(1):449–492, 2019. doi: 10.1111/jofi.12733.

[24] Adam Copeland, Darrell Duffie, and Yilin Yang. Reserves were not so ample after all. *Quarterly Journal of Economics*, 140:239–281, 2025.

[25] Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from covered interest rate parity. *Journal of Finance*, 73(3):915–957, 2018. doi: 10.1111/jofi.12620.

[26] Darrell Duffie. Still the world's safe haven? redesigning the U.S. treasury market after the COVID-19 crisis. Hutchins Center Working Paper 62, Hutchins Center on Fiscal and Monetary Policy, Brookings Institution, 2020.

[27] James Durbin and Siem Jan Koopman. *Time Series Analysis by State Space Methods*. Oxford University Press, Oxford, 2 edition, 2012. ISBN 9780199641178.

[28] Barry Eichengreen. *Exorbitant Privilege: The Rise and Fall of the Dollar and the Future of the International Monetary System*. Oxford University Press, Oxford, 2011.

[29] Barry Eichengreen. *Globalizing Capital: A History of the International Monetary System*. Princeton University Press, Princeton, 3 edition, 2019. ISBN 9780691193908.

[30] Barry Eichengreen, Arnaud Mehl, and Livia Chiţu. *How Global Currencies Work: Past, Present, and Future*. Princeton University Press, Princeton, NJ, 2018.

[31] Emmanuel Farhi and Xavier Gabaix. Rare disasters and exchange rates. *Quarterly Journal of Economics*, 131(1):1–52, 2016. doi: 10.1093/qje/qjv040.

[32] Emmanuel Farhi and Matteo Maggiori. A model of the international monetary system. *Quarterly Journal of Economics*, 133(1):295–355, 2018. doi: 10.1093/qje/qjx031.

[33] Henry Farrell and Abraham L. Newman. Weaponized interdependence: How global economic networks shape state coercion. *International Security*, 44(1):42–79, 2019. doi: 10.1162/isec_a_00351.

[34] Jeffrey Frankel. The plaza accord, 30 years later. NBER Working Paper 21813, National Bureau of Economic Research, 2015.

[35] Yoichi Funabashi. *Managing the Dollar: From the Plaza to the Louvre*. Institute for International Economics, Washington, DC, 1988.

[36] Xavier Gabaix and Matteo Maggiori. International liquidity and exchange rate dynamics. *Quarterly Journal of Economics*, 130(3):1369–1420, 2015.

[37] Peter M. Garber. The collapse of the Bretton Woods fixed exchange rate system. In Michael D. Bordo and Barry Eichengreen, editors, *A Retrospective on the Bretton Woods System: Lessons for International Monetary Reform*, pages 461–494. University of Chicago Press, Chicago, 1993.

[38] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman and Hall/CRC, Boca Raton, 3 edition, 2013. ISBN 9781439840955.

[39] Georgios Georgiadis, Gernot J. Müller, and Ben Schumann. Global risk and the dollar. *Journal of Monetary Economics*, 144:103549, 2024. doi: 10.1016/j.jmoneco.2024.01.002.

[40] Linda S. Goldberg, Craig Kennedy, and Jason Miu. Central bank dollar swap lines and overseas dollar funding costs. *Federal Reserve Bank of New York Economic Policy Review*, 17(1), 2011.

[41] Gita Gopinath, Emine Boz, Camila Casas, Federico J. Díez, Pierre-Olivier Gourinchas, and Mikkel Plagborg-Møller. Dominant currency paradigm. *American Economic Review*, 110(3):677–719, 2020. doi: 10.1257/aer.20171201.

[42] Gary B. Gorton and Jeffery Y. Zhang. Taming wildcat stablecoins. *University of Chicago Law Review*, 90(3):909–972, 2023.

[43] Pierre-Olivier Gourinchas and Hélène Rey. From world banker to world venture capitalist: U.S. external adjustment and the exorbitant privilege. In Richard H. Clarida, editor, *G7 Current Account Imbalances: Sustainability and Adjustment*, pages 11–66. University of Chicago Press, Chicago, 2007.

[44] Pierre-Olivier Gourinchas, Hélène Rey, and Nicolas Govillot. Exorbitant privilege and exorbitant duty. IMES Discussion Paper 2010-E-20, Institute for Monetary and Economic Studies, Bank of Japan, 2010.

[45] Joanne Gowa. *Closing the Gold Window: Domestic Politics and the End of Bretton Woods*. Cornell University Press, Ithaca, 1983. ISBN 9780801416224.

[46] Andrew G. Haldane and Robert M. May. Systemic risk in banking ecosystems. *Nature*, 469: 351–355, 2011. doi: 10.1038/nature09659.

[47] James D. Hamilton. A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384, 1989.

[48] Ethan Ilzetzki, Carmen M. Reinhart, and Kenneth S. Rogoff. Exchange arrangements entering the twenty-first century: Which anchor will hold? *Quarterly Journal of Economics*, 134(2): 599–646, 2019.

[49] Zhengyang Jiang, Arvind Krishnamurthy, and Hanno Lustig. Foreign safe asset demand and the dollar exchange rate. *Journal of Finance*, 76(3):1049–1089, 2021. doi: 10.1111/jofi.13003.

[50] Kevin L. Kliesen and Douglas C. Smith. Measuring financial market stress. Economic synopses, Federal Reserve Bank of St. Louis, 2010.

[51] Arvind Krishnamurthy and Hanno Lustig. Mind the gap in sovereign debt markets: The U.S. treasury basis and the dollar risk factor. In *Proceedings of the Economic Policy Symposium, Jackson Hole*. Federal Reserve Bank of Kansas City, 2019.

[52] Andrew Lilley, Matteo Maggiori, Brent Neiman, and Jesse Schreger. Exchange rate reconnect. *Review of Economics and Statistics*, 104(4):845–855, 2022. doi: 10.1162/rest_a_00978.

[53] Matteo Maggiori. Financial intermediation, international risk sharing, and reserve currencies. *American Economic Review*, 107(10):3038–3071, 2017. doi: 10.1257/aer.20130479.

[54] Ulrike Malmendier and Stefan Nagel. Depression babies: Do macroeconomic experiences affect risk taking? *Quarterly Journal of Economics*, 126(1):373–416, 2011. doi: 10.1093/qje/qjq004.

[55] Ulrike Malmendier and Stefan Nagel. Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87, 2016. doi: 10.1093/qje/qjv037.

[56] Robert N. McCauley, Patrick McGuire, and Vladyslav Sushko. Global dollar credit: Links to US monetary policy and leverage. *Economic Policy*, 30(82):187–229, 2015.

[57] Richard A. Meese and Kenneth Rogoff. Empirical exchange rate models of the seventies: Do they fit out of sample? *Journal of International Economics*, 14(1–2):3–24, 1983. doi: 10.1016/0022-1996(83)90017-X.

[58] Silvia Miranda-Agrippino and Hélène Rey. U.S. monetary policy and the global financial cycle. *Review of Economic Studies*, 87(6):2754–2776, 2020. doi: 10.1093/restud/rdaa019.

[59] Nicholas Mulder. *The Economic Weapon: The Rise of Sanctions as a Tool of Modern War*. Yale University Press, New Haven, CT, 2022.

[60] Maurice Obstfeld and Kenneth Rogoff. The six major puzzles in international macroeconomics: Is there a common cause? In Ben S. Bernanke and Kenneth Rogoff, editors, *NBER Macroeconomics Annual 2000*, pages 339–390. MIT Press, Cambridge, MA, 2000.

[61] Jerome H. Powell. Monetary policy and price stability. Speech at "Reassessing Constraints on the Economy and Policy", symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 2022. 26 August 2022.

[62] Zoltan Pozsar. Shadow banking: The money view. OFR Working Paper 14-04, Office of Financial Research, U.S. Department of the Treasury, 2014.

[63] Eswar S. Prasad. *The Dollar Trap: How the U.S. Dollar Tightened Its Grip on Global Finance*. Princeton University Press, Princeton, NJ, 2014.

[64] Giorgio E. Primiceri. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852, 2005. doi: 10.1111/j.1467-937X.2005.00353.x.

[65] Hélène Rey. Dilemma not trilemma: The global financial cycle and monetary policy independence. In *Proceedings of the Economic Policy Symposium, Jackson Hole*, pages 285–333. Federal Reserve Bank of Kansas City, 2013.

[66] Kenneth Rogoff. The purchasing power parity puzzle. *Journal of Economic Literature*, 34(2): 647–668, 1996.

[67] Thomas J. Sargent. The ends of four big inflations. In Robert E. Hall, editor, *Inflation: Causes and Effects*, pages 41–98. University of Chicago Press, Chicago, 1982.

[68] Robert J. Shiller. Narrative economics. *American Economic Review*, 107(4):967–1004, 2017. doi: 10.1257/aer.107.4.967.

[69] Robert J. Shiller. *Narrative Economics: How Stories Go Viral and Drive Major Economic Events*. Princeton University Press, Princeton, 2019.

[70] George Soros. *The Alchemy of Finance*. Simon & Schuster, New York, 1987. Reissued as a Wiley Investment Classic.

[71] David E. Spiro. *The Hidden Hand of American Hegemony: Petrodollar Recycling and International Markets*. Cornell University Press, Ithaca, 1999. ISBN 9780801428845.

[72] Cass R. Sunstein. *#Republic: Divided Democracy in the Age of Social Media*. Princeton University Press, Princeton, 2017. ISBN 9780691175515.

[73] Timo Teräsvirta. Specification, estimation, and evaluation of smooth transition autoregressive models. *Journal of the American Statistical Association*, 89(425):208–218, 1994.

[74] Annette Vissing-Jorgensen. The treasury market in spring 2020 and the response of the federal reserve. *Journal of Monetary Economics*, 124:19–47, 2021.

[75] World Gold Council. Gold demand trends: Full year 2024. World Gold Council, Goldhub, 2025.