FORGOTTEN FORGE

A textbook in six parts

The Peak

a universal recipe for phase transitions,
in twelve worlds

quantum hardware — magnetism — finance — machine learning
seismology — climate — GPU — protein — language
number theory — edge — LLM economics

M. C. Wurm
ForgottenForge | Buckenhof | 2026

Companion to SIGMA-C-FRAMEWORK v3.0.0
Empirical anchor: AVS Quantum Science **8**, 013804 (2026)
© 2026 ForgottenForge.xyz AGPL-3.0-or-later / Commercial

The Peak

a universal recipe for phase transitions,
in twelve worlds

M. C. Wurm

ForgottenForge.xyz | 2026

Preface

σ_c is where the system tips. $\Pi = D\,\gamma$ is why it tips. Most of this book is about the first. The middle third is about the second.

We do not promise you a revolution. Revolutions are loud, messy, poorly organised. We promise something quieter. A recipe. Three lines of Python that take a system with a tunable knob and a measurable response and tell you where it tips over.

Sweep a parameter. Measure. Take a derivative. Find the peak. Four steps; the last three do not change between domains, only the first does. The same four steps locate the Curie temperature of an iron magnet, the operational noise threshold of a quantum processor, the regime shift of a stock market, the onset age of a hereditary protein disease. Twelve domains, one recipe. We are aware this sounds suspicious. The first half of this book is written for the suspicious.

A good method is like a good butler. You do not notice it. It does its job quietly. It does not invent problems and it does not solve more than was asked. On its bad days it tells you it cannot help, and goes to fetch someone who can. The susceptibility framework, when it behaves, is that kind of butler. When it does not behave, it tells you exactly which of five trust conditions it failed — and that is the chapter we put first in Part II, on purpose.

How it came about. It began with my having to watch.

Not entirely of my own free will. Once one has reached a certain age, one has by then met a considerable number of fellow human beings, and a particular subset of them seems to have specialised in acting in ways that, retrospectively, raise the question: when, precisely, did this happen? The colleague who resigns from a secure post in order to breed cryptocurrencies. The friend who, after three years of complaining about his fiancée, suddenly marries her. The patient who announces in consultation that she has reconsidered the operation, having found something better online. The legislature that appoints a commission. The homeowner who renovates the kitchen for the eleventh time.

I was not present when these people made their decisions. But I was present shortly before and shortly after, and in every case there was a moment — a narrow moment, often not a whole day — at which the person stood on one side, and shortly afterwards stood, irrevocably, on the other.

I have since called this moment the *tipping point*. It is a modest name for something that, to be quite honest, leaves me at a loss for words.

For a while I had hoped this was a social observation. Something concerning only human beings, and only the unwilling among them. It turned out not to be so. Magnets have a tipping point. Stock markets have a tipping point. Quantum computers have a tipping point.

Earthquakes, the climate, power grids, proteins — all of them have tipping points. It is the most common property I know of in systems, second only to the property that they can break.

That, in itself, would not be so bad. What is bad is that, having spent some years identifying these tipping points — which, incidentally, is not particularly difficult once one knows what to look for — I still do not know how to get back from the unpleasant side to the pleasant one.

That is the second part. The first part — “here is the tipping point” — is mathematically tractable. That is this book.

The second part — “how does one get back to where it was nice” — is not mathematically tractable. Perhaps it is not tractable at all. Perhaps it merely lies beyond what I, as a single individual, can settle in one lifetime. I would be grateful if anyone reading this would send me an idea.

Until then: the procedure described in what follows tells you, at the very least, that a tipping point exists. It tells you where it is. It even tells you how sharp it is, which has turned out to be unexpectedly useful. What it does not tell you is what to do once you have got there. The reader will have to decide that.

It is a modest book. It does one thing. That one thing I had not set out to invent — it had impressed itself upon me, through several decades of looking on, whether I wanted it or not. I pass it on here in the hope that, next time, somebody might look at the right moment and say: “Wait, hold on — this is about to tip. Are you sure you want that?” Even that would be an improvement on the present state of affairs, however small.

One last thing. We assume you can add, subtract, multiply, divide. We assume you can read. We do not assume you have seen a Greek letter, met an expectation value, or run a Python script. If a phrase in the table of contents intimidates you, ignore it; by the time you reach the chapter, it will not.

The first draft of this book was 78 pages long and unreadable. The second was 130 pages long and dishonest. This is the third, and it is honest the way a good butler is honest — by reporting what is actually in the room.

About this book

This book has one job. Take a reader who can do arithmetic and bring them to the point where they can apply the *susceptibility framework* — the σ_c-method — to find critical thresholds in their own data. Quantum computer or financial market, magnetic material or GPU benchmark. The recipe is one recipe.

Who this book is for. Anybody who has ever measured something while turning a knob and wondered *"where does it tip over?"* A degree is not required. Four arithmetic operations and patience for one worked example will do.

What this book is not. A research monograph. We do not assume you have read papers on phase transitions, Fisher information, or non-equilibrium statistical mechanics. Where these ideas appear, we build them from the ground up.

The promise. By the end of Part II, you will run the σ_c-method in five lines of Python. By the end of Part III, you will understand why it works. By the end of Part IV, you will know which of twelve ready-made domain adapters fits your problem, and how to extend the framework to a thirteenth of your own.

How to read. Parts I through III are a chain. Skipping ahead will leave a gap that later chapters quietly fall into. From Part IV on, each chapter stands alone. Pick the domain that matches your work; the rest will be there when you need them.

Conventions.

- Every new symbol is defined the first time it appears. Greek letters do not get a free pass.
- Every formula is followed by at least one numerical example. If you cannot reproduce the number on paper, the explanation is not finished.
- Every chapter ends with a small exercise. They take five to thirty minutes and are worth the time.
- Code uses Python and runs in any environment with NumPy and SciPy installed. The `sigma_c` library is a convenience, not a dependency — the five-line core fits on a napkin.
- Sentences that survive a section without being supplanted by a better one are placed in a coloured box at the end. Four colours, four purposes: *blue* for what to remember,

orange for what to avoid, *green* for what to try, and *red* for outright cautions (clinical, financial, safety).

- **The symbol σ is locally rebound in each chapter.** In each chapter σ denotes the control axis *of that chapter*, not a universal physical sigma. The Greek letter is used because it is short and traditional; it carries no physical meaning of its own across domains. We use σ_{ker} (kernel width) and σ_t (volatility, in finance) with explicit subscripts to avoid collisions.

The reference framework. The accompanying open-source library is `sigma-c-framework` (version 3.0.0). Install it with

```
pip install sigma-c-framework
```

You can finish this book without ever installing it; the five-line core algorithm requires only NumPy and SciPy. The framework is more comfortable, the way a butler is more comfortable than no butler.

Citing this book.

> M. C. Wurm, *The Peak: a universal recipe for phase transitions, in twelve worlds*, ForgottenForge, Buckenhof, 2026.

(Licensing and edition details are in the colophon at the back.)

Three reading paths

This book has four hundred-odd thousand things you could read in it. You almost certainly do not want to read all of them. Three explicit reading paths, in increasing depth:

Path A — "I just want to use it" (1–2 evenings).

- Front matter and Greek-letter guide (15 min)
- Chapters 1–7 *briefly* (the foundations; skim if you have ever taken calculus)
- Chapter 9 (**the universal recipe**)
- Chapter 10 (**when the method fails**)
- Chapter 12 (κ **thresholds**)
- One Part IV chapter that matches your domain
- Part V (**validation**) when you are about to publish

Path B — "I want to understand why it works" (one week). Path A, plus all of Part III (contraction geometry) and all of Part V (validation). Skim Part IV chapters outside your domain.

Path C — "I want to extend the framework" (two weeks). Path B, plus the recipes in Part VI and the appendices, plus actually implementing your own adapter (Chapter 46). At this point you are a contributor, not a reader.

What this book is not promising. It is not promising a theory of everything. It is not promising that every system has a peak; some do not, and Chapter 10 will help you tell which. It is not promising that the same numerical σ_c will hold across hardware platforms; the magnetism paper's own follow-up on different hardware shows the value shifts by ±3–5% when you change the noise model. It is not promising that the method replaces domain expertise; the method finds where to look, you still have to interpret what you see.

What we *do* promise: if your system has a tunable knob and a monotone-ish response, the recipe in Chapter 9 will locate the operational threshold without requiring you to invent a model first.

Contents

Preface 3

About this book 5

Three reading paths 7

The seven symbols you really need 15

I Foundations — from arithmetic to the derivative 16

A reading guide for Greek letters 17

1 What is a number, what is a ratio? 18
- 1.1 Numbers, in the only way we will use them . . . 18
- 1.2 Subtraction tells you a change . . . 18
- 1.3 Division tells you a rate . . . 19
- 1.4 Ratios with units, and why units matter . . . 19
- 1.5 One thing depending on another . . . 19

2 What is a function? 21
- 2.1 The recipe view of a function . . . 21
- 2.2 The table view . . . 21
- 2.3 The plot view . . . 22
- 2.4 Functions in the lab . . . 22
- 2.5 The control parameter σ . . . 22

3 Slope: the average rate of change 24
- 3.1 Two points on a curve . . . 24
- 3.2 Slope of a straight line . . . 24
- 3.3 Slope of a curved line — it depends where you look . . . 25
- 3.4 The slope is itself a function . . . 25

4 The derivative: the slope at a single point 28
- 4.1 The idea . . . 28
- 4.2 Three derivatives you should know by sight . . . 28
- 4.3 The derivative when you only have a table . . . 28
- 4.4 One-sided differences at the edges . . . 29

4.5 In code: numerical derivative in three lines . . . 29

5 The peak: where the slope is the largest **32**
5.1 Absolute value: ignoring the direction . . . 32
5.2 Argmax: the location of the maximum . . . 32
5.3 In code: argmax in one line . . . 33
5.4 Why this is interesting at all . . . 33

6 Data is noisy: smoothing **35**
6.1 Why raw derivatives misbehave . . . 35
6.2 Smoothing: replace each value by a local average . . . 35
6.3 Gaussian smoothing: a smarter weighted average . . . 35
6.4 The smoothed pipeline . . . 36
6.5 How much to smooth? . . . 37

7 Powers, exponentials, logarithms in three pages **39**
7.1 Powers . . . 39
7.2 The number e and the exponential function . . . 39
7.3 The logarithm: the inverse of exponentiation . . . 40

8 Confidence: what is a probability? **42**
8.1 Probability without philosophy . . . 42
8.2 Mean and standard deviation: summarising a sample . . . 42
8.3 The bootstrap idea . . . 42

II The susceptibility method — χ, σ_c, κ **47**

9 The universal recipe **48**
9.1 Application 1: the Curie point of an iron magnet . . . 49
9.2 Application 2: a regime shift in financial returns . . . 50
9.3 Application 3: a Gutenberg–Richter b-value shift . . . 51
9.4 What the three applications have in common . . . 52

10 When the method works and when it does not **53**
10.1 Failure mode 1: oscillating observable . . . 53
10.2 Failure mode 2: multipeaked structure . . . 54
10.3 Failure mode 3: window contains no transition . . . 54
10.4 Failure mode 4: undersampling . . . 54
10.5 Failure mode 5: σ is not actually a control . . . 54
10.6 Failure mode 6: noisy enough that smoothing hides the peak . . . 54
10.7 Rule of thumb: when to trust the report . . . 55
10.8 When the checks disagree: a small decision tree . . . 55

11 Susceptibility, formally **57**
11.1 Definition . . . 57
11.2 Why “susceptibility”? . . . 57
11.3 Two worked-from-data examples: a 1D correlation decay . . . 58

12 The peak clarity κ **63**
12.1 Three different ways to score sharpness . . . 63
12.2 Which one should you use? . . . 63
12.3 Threshold of significance . . . 63

13 Why peaks exist in the first place: the existence argument **65**
13.1 The boundary trick . . . 65
13.2 The signal-to-noise crossover — the actual existence argument . . . 66

14 Choosing the observable **68**
14.1 Observable quality score, computed automatically . . . 68

III Contraction geometry — why the method works **69**

A note before you read Part III **70**

15 From the coffee mug to a contraction **71**
15.1 The mug, one more time . . . 71
15.2 The new piece: feed the function its own output . . . 71
15.3 Self-map: the domain equals the codomain . . . 72
15.4 What happens to the image, after one step . . . 72
15.5 The bridge in one sentence . . . 72

16 Maps, images, pre-images **73**
16.1 Image and pre-image . . . 73
16.2 Injective vs. non-injective . . . 73

17 The contraction defect D **75**
17.1 Computing D in practice . . . 75
17.2 D as bits of information lost . . . 75

18 The drift γ **76**
18.1 Three regimes, decided by γ . . . 77

19 The universal threshold $\Pi = D \cdot \gamma$ **78**
19.1 The connection to σ_c — one sentence, two parts . . . 79

20 The four types: D, O, S, R **80**

IV The twelve domains **83**

A short glossary for Part IV **84**

21 Quantum hardware: the Wurm 2026 case study **86**
21.1 The hardware . . . 86
21.2 Six experiments, one method . . . 87
21.3 The case study: experiment E3 . . . 87
21.3.1 The setup, in detail . . . 87
21.3.2 What we expect to see, before any data . . . 88
21.3.3 The 10-line demo: same shape, no quantum hardware . . . 88

21.3.4 Reproducing the experiment on a local simulator (skip on first read) . . 88
21.3.5 Using the framework instead 90
21.3.6 Real-hardware version 90
21.4 Reading the result 91
21.5 Cross-observable validation 91
21.6 Robustness to noise model 91
21.7 All six experimental scales — what each one teaches 91
21.8 Operational guidelines for NISQ work 92
21.9 The Ankaa-3 nine-circuit reference set 92
21.10 Cross-platform replication on Rigetti Cepheus-1 93
21.11 Cost and reproducibility 95

22 Magnetism: the textbook Curie point 97
22.1 What is a ferromagnet? 97
22.2 The Ising model in one paragraph 97
22.3 The observable and the control parameter 98
22.4 Generating data in five lines (Metropolis Monte Carlo) 98
22.5 Applying `MagneticAdapter` 99
22.6 Critical exponents: the next-level inference 99
22.7 Finite-size scaling, in code 100
22.8 Universality classes 100
22.9 Pitfalls in magnetic data 101

23 Finance: regime detection from returns 103
23.1 What is a return? 103
23.2 Probe 1: Hurst exponent and memory horizon 104
23.3 Probe 2: GARCH persistence 104
23.4 Probe 3: order-flow imbalance and crash risk 105
23.5 End-to-end: detect the regime of the S&P 500 105
23.6 The susceptibility view of regime change 105
23.7 Caveats and ethics 105

24 Seismology: Gutenberg–Richter and Omori 108
24.1 Gutenberg–Richter: how often is each magnitude? 108
24.1.1 Computing b from a magnitude catalogue 109
24.2 Omori's law: how aftershocks decay 109
24.3 The susceptibility view: detecting regime changes 109
24.4 Bootstrap significance for the b-value 110
24.5 Use case: induced seismicity at a geothermal site 110

25 Climate: mesoscale boundaries 113
25.1 Atmospheric kinetic energy across scales 113
25.2 Detection without prior knowledge 113
25.3 Why a peak? 114
25.4 Vertical structure: detecting the tropopause 115
25.5 Use case: ERA5 reanalysis sweep 115

26 GPU: rooflines, thermal cliffs, cache transitions **117**
26.1 The roofline model . . . 117
26.1.1 Susceptibility view . . . 117
26.2 Cache transitions . . . 119
26.2.1 Detecting them automatically . . . 119
26.3 Thermal throttling . . . 119
26.3.1 Measuring with NVML in real time . . . 120
26.4 Combining: the operational sweet spot . . . 120

27 Machine learning: learning-rate cliffs and beyond **123**
27.1 The learning-rate sweet spot . . . 123
27.1.1 The LR-range test (Smith 2017) . . . 123
27.1.2 Why a peak? . . . 125
27.2 Other hyperparameter sweeps . . . 125
27.2.1 Two-dimensional sweep: LR $\times$ batch size . . . 126
27.3 Training instability detection in real time . . . 126
27.4 Caveats specific to ML . . . 127

28 Edge / IoT: the efficiency knee **129**
28.1 Performance, power, and the efficiency curve . . . 129
28.2 Worked example . . . 129
28.3 Use case: Raspberry Pi 4 efficiency study . . . 130
28.4 Why this matters for battery life . . . 131

29 LLM economics: the cost-quality frontier **133**
29.1 Three dimensions, one decision . . . 133
29.2 The value ratio and the safety filter . . . 134
29.2.1 Safety bound: maximum tolerable hallucination . . . 134
29.2.2 Value ratio . . . 134
29.3 Susceptibility view: where does quality drop off cliffwise? . . . 135
29.4 Use case: choosing a model for a customer-service chatbot . . . 136

30 Number theory: Collatz and the qn+c family **139**
30.1 The Collatz map . . . 139
30.2 The cycle map and embedding depth . . . 139
30.3 Computing D and γ . . . 140
30.4 The contraction product and prediction . . . 141
30.5 The twelve qn+c maps . . . 141
30.6 The countdown decomposition . . . 142
30.7 The Geo(1/2) distribution of resets . . . 142
30.8 Information-theoretic interpretation . . . 142

31 Protein: stability, mutation, and onset age **145**
31.1 Folding stability and the marginal-stability principle . . . 145
31.2 The contraction index σ . . . 146
31.2.1 Verifying $\sigma = 1$ at the melting point . . . 146
31.3 Mutational stress: $\Delta\Delta G$ . . . 146
31.3.1 Worked example, on paper: TTR V30M (FAP) . . . 147
31.4 Age-dependent drift and onset prediction . . . 149
31.4.1 V30M onset prediction . . . 149
31.4.2 Onset envelope . . . 150

31.5 The shipped mutation tables . . . 150
31.6 Disease-mechanism classification . . . 151
31.7 The dual-basin Monte Carlo model . . . 151

32 Linguistics: etymological depth and semantic change 154
32.1 Etymological depth . . . 154
32.2 Measuring semantic change . . . 155
32.3 The susceptibility test . . . 156
32.4 The fixed-point test . . . 156
32.5 Mediation: does word frequency explain the effect? . . . 157
32.6 Cross-linguistic replication . . . 157
32.7 The German P/T/O ANOVA . . . 157
32.8 Transparency effect within higher ED . . . 157
32.9 What this means . . . 158

V Open conjectures and limits of the framework 160

33 Four named conjectures 161
33.1 Conjecture C1: contraction universality . . . 161
33.2 Conjecture C2: operational-critical equivalence . . . 162
33.3 Working Hypothesis WH3: cross-platform threshold invariance . . . 163
33.4 Conjecture C4: κ-threshold rescaling . . . 163

34 Limits of the framework: when not to use the recipe 164

35 Multipeak: when there really are two 165
35.1 Diagnosis: two peaks, not one . . . 165
35.2 Multipeak in practice . . . 165
35.3 Reporting multipeak results . . . 166

VI Validation, statistics, and rigour 167

36 The boundary check 168

37 The permutation test 169
37.1 A worked example with numbers . . . 169
37.2 Which null model for which domain? . . . 170

38 Peak clarity threshold 171

39 Fisher information bound 172

40 Multiple testing 173

41 Cross-validation across observables 174

VII Worked examples and recipes 175

42 Recipe: minimal σ_c in pure NumPy/SciPy 176

43 Recipe: bootstrap CI on σ_c **177**

44 Recipe: choosing the kernel **178**

45 Recipe: install the full framework **179**

46 Recipe: build your own adapter **180**

47 Recipe: an end-to-end synthetic experiment **181**

48 Recipe: live monitoring with streaming sigma_c **182**

49 Recipe: report a result for publication **183**

A Symbol glossary **184**

B Proof: existence of an interior maximum **185**

C The twelve qn+c maps, predictions and observations **186**

D Annotated reading list **187**

Index **191**

Index **191**

E Reproducibility notes **192**

F Where to get the Chapter 9 data **193**

F.1 Curie point (Ising magnetisation) . . . 193

F.2 S&P 500 daily returns (the 2008 example) . . . 193

F.3 Southern California earthquake catalogue (Ridgecrest) . . . 193

G Acknowledgements **195**

H Colophon **196**

The seven symbols you really need

If you remember nothing else from this book, remember the seven symbols below. Everything else gets defined when it appears.

symbol	in one line	first appears
σ	the control parameter you turn — the temperature, the noise level, the learning rate	Chapter 2
O	the observable you measure as you turn it	Chapter 2
$\chi(\sigma) = \lvert \mathrm{d}O/\mathrm{d}\sigma \rvert$	the susceptibility: how fast O is changing	Chapter 11
σ_c	the location of the peak of χ — *where the system tips*	Chapter 9
κ	the peak clarity: $\chi_{\max}$ divided by the average χ	Chapter 12
D	the contraction defect of a map: $\lvert S \rvert / \lvert f(S) \rvert$	Chapter 17
γ	the drift of a map: geometric mean of $f(x)/x$ *(not* the same γ as quantum noise strength!)	Chapter 18
$\Pi = D \cdot \gamma$	bonus, eighth: the contraction product — *why the system tips*	Chapter 19

The one-sentence summary of the entire book is:

TAKEAWAY

σ_c **is where it tips.** $\Pi = D\gamma$ **is why it tips.**

The full Greek-letter pronunciation guide is overleaf (the classical letters take getting used to but each is just a letter).

Part I

Foundations — from arithmetic to the derivative

A reading guide for Greek letters

This book uses Greek letters. School curricula tend to use them sparingly, and many readers find them more intimidating than they deserve. A Greek letter is a letter. It is pronounced as below, and otherwise behaves exactly like an ordinary letter — you can multiply it, add it, write equations with it.

symbol	name	where in this book
σ	"sigma"	control parameter (Chapters 2–7)
σ_c	"sigma-cee"	location of the susceptibility peak
χ	"chi" (sounds like "kye")	the susceptibility itself
κ	"kappa"	peak clarity (how sharp the peak is)
γ	"gamma"	drift in number theory, noise in quantum
Δ	"delta" (capital)	change, difference
ρ	"rho" (sounds like "row")	density matrix, correlation
ξ	"xi" (sounds like "ksee")	correlation length
π	"pi"	both the constant 3.14... and GARCH persistence
β, α	"beta", "alpha"	critical exponents in magnets, GARCH params
∂	"del", "partial"	a particular kind of derivative
$\sim$	"goes as", "scales as"	approximate proportionality
$\propto$	"proportional to"	strict proportionality
$\langle \cdot \rangle$	"angle brackets"	average over many trials
$\lvert \cdot \rvert$	"absolute value"	strip the sign, keep the magnitude

If you forget any of them, return to this page. It will not go anywhere.

Chapter 1

What is a number, what is a ratio?

This chapter assumes nothing. If you can do $7 + 3 = 10$ and $10 \div 2 = 5$, you are over-qualified. Readers who already know what a derivative is may skim it; readers who do not should not. Everything in the remaining 300 pages rests on the four arithmetic operations done carefully. We begin by doing them carefully.

1.1 Numbers, in the only way we will use them

A *number* in this book is one of the following:

- a whole number like 0, 1, 2, 3, -4, 17;
- a number with a decimal part like 0.5, 3.14, -2.718;
- a very small positive number like 0.001 or 10^{-6}, which we read as "ten to the minus six". The symbol 10^{-6} is just a short way of writing 0.000001.

The set of all such numbers we write $\mathbb{R}$, and call it "the real numbers". You will not need any deeper definition than this.

1.2 Subtraction tells you a change

If a quantity was 4 before and is now 7, the change is $7 - 4 = 3$. We say the quantity *increased by* 3.

If a quantity was 4 before and is now 1, the change is $1 - 4 = -3$, a negative number. We say the quantity *decreased by* 3. The minus sign is the only thing that tells you the direction of change.

Definition 1.1 (Change)**.** The change in a quantity O between two measurements is $\Delta O = O_{\text{new}} - O_{\text{old}}$. The Greek letter Δ (delta) is pronounced "delta" and from here on will always mean "change in".

Example 1.2 (The coffee mug, part 1: the change)**.** A coffee mug was at 80 °C a moment ago. Now it reads 77 °C. The temperature change is $\Delta T = T_{\text{new}} - T_{\text{old}} = 77 - 80 = -3\,°\text{C}$. The negative sign says: the temperature *went down* by 3 degrees. We will return to this mug throughout the chapter.

1.3 Division tells you a rate

The mug from Example 1.2 did not lose its three degrees instantly. Let us say it took 60 seconds. The *rate* of cooling is then

$$\frac{\Delta T}{\Delta t} = \frac{-3\,^\circ\mathrm{C}}{60\,\mathrm{s}} = -0.05\,^\circ\mathrm{C} \text{ per second.}$$

Two things to notice. *First*, we divided. *Second*, the units followed along: degrees on top, seconds on bottom, hence "degrees per second".

If instead it took only 5 seconds, the rate would be $-3/5 = -0.6\,^\circ\mathrm{C/s}$ — a much faster cooling. Big rate means the change happens quickly; small rate means it happens slowly. *This single idea is the seed of everything that follows in this book.*

TAKEAWAY

A *rate* is a change divided by another change. It tells you how fast one thing moves when another moves. Everything in the σ_c method comes from one specific rate: how fast a measurement changes when you change a control parameter.

1.4 Ratios with units, and why units matter

The rate $-0.05\,^\circ\mathrm{C/s}$ depended on the units. If we had measured time in minutes instead of seconds, the same cooling would be $-3\,^\circ\mathrm{C} \div 1\,\mathrm{min} = -3\,^\circ\mathrm{C/min}$. *Same physical process, different number, because different unit.*

PITFALL

When you compare two rates, you must use the same units for both. A rate of $-3\,^\circ\mathrm{C/min}$ is *not* larger than a rate of $-0.05\,^\circ\mathrm{C/s}$; they are equal. Always check units before drawing a conclusion.

1.5 One thing depending on another

So far we tracked one quantity (the temperature of the mug) and one "axis" along which it varied (the time). That is the simplest possible two-quantity relationship — one input, one output. This pattern is so common that mathematicians have given it a name: a *function*.

The next chapter is about functions. For the moment, just hold the picture in your head: a function is a recipe that takes one number in (the *input*) and gives one number back (the *output*). Our mug-cooling story is a function: each instant of time t has a single temperature T associated with it, and we can write $T = T(t)$.

TRY THIS

A car was at 200 km along a road at 14:00 and at 260 km at 14:30. Compute its average speed in km/h. Did the car go faster or slower than 100 km/h?

Solution, step by step.

Step 1: identify the two quantities and their changes. The position x is what is varying. The time t is the axis along which it varies (same pattern as the coffee mug from Example 1.2, but spatial instead of thermal).

$$\Delta x = x_{\mathrm{new}} - x_{\mathrm{old}} = 260\,\mathrm{km} - 200\,\mathrm{km} = 60\,\mathrm{km}.$$

$$\Delta t = t_{\mathrm{new}} - t_{\mathrm{old}} = 14{:}30 - 14{:}00 = 30\,\mathrm{min} = 0.5\,\mathrm{h}.$$

We convert 30 minutes to 0.5 hours so that our numerator (km) and denominator (h) end up in km/h, which is the requested unit.

Step 2: form the rate.

$$\text{speed} = \frac{\Delta x}{\Delta t} = \frac{60\,\text{km}}{0.5\,\text{h}} = 120\,\text{km/h}.$$

The arithmetic is one division: $60 \div 0.5 = 120$. (Or, if dividing by a decimal makes you nervous: $60/0.5 = 60 \cdot (1/0.5) = 60 \cdot 2 = 120$.)

Step 3: compare to the asked-about threshold. $120 > 100$, so the car was going *faster* than 100 km/h.

Sanity check. If the car had travelled 60 km in a full hour, the speed would have been 60 km/h — slower than 100. We covered the same distance in only *half* an hour, so we expect roughly double, which is exactly what we got.

What this exercise is teaching. You now have all three ingredients in one place: a quantity that changes (x), an axis along which it changes (t), and a rate (their ratio). The σ_c-method later in the book is the same pattern applied to a quantity that changes (an observable O) when you turn a knob (a control parameter σ).

Chapter 2

What is a function?

In Example 1.2 we had a coffee mug. It cooled from 80 °C to 77 °C. One number (temperature) varied with another (time). That little observation is more important than it sounds. Mathematicians have built three hundred years of analysis on it. They gave the pattern a name — *function* — and a notation: $f(x)$. The three hundred years are where derivatives come from, which is where peaks come from, which is where this book begins to be useful.

2.1 The recipe view of a function

A *function* f is a rule that, given an input x, returns one specific output, which we write $f(x)$. We read $f(x)$ as "f of x". The word "specific" is doing real work in that sentence: a function must give the same output every time it is fed the same input. A rule that returns "some number between 3 and 5" is not a function. A rule that returns "today's temperature in Berlin" is not a function either — different days, different answers. A function is a contract between input and output, and the contract holds.

Example 2.1 (Doubling)**.** The rule "double the input" is a function. If we call it f, then $f(3) = 6$, $f(0) = 0$, $f(-1.5) = -3$, $f(100) = 200$. We can write the recipe compactly: $\boxed{f(x) = 2x}$. Whatever number goes in, twice that number comes out.

Example 2.2 (Squaring)**.** The rule "multiply the input by itself" is a function. Compactly: $\boxed{f(x) = x \cdot x = x^2}$. So $f(2) = 4$, $f(3) = 9$, $f(-2) = 4$ as well (because $(-2) \cdot (-2) = +4$, two negatives multiply to a positive). We will use this function as the running example for the next two chapters.

We will sometimes write $f\colon x \mapsto x^2$, which is just a fancier notation for the same recipe in Example 2.2. The arrow $\mapsto$ reads "maps to". Read it as: "f sends the input x to its square".

2.2 The table view

Any function can be *tabulated*: pick some inputs, write each input next to its output. For the squaring function from Example 2.2:

input x	output $f(x) = x^2$
-2	4
-1	1
0	0
1	1
2	4
3	9

This is exactly the kind of table you will produce when you *sweep* a control parameter in a real experiment — a fancy word for "vary it across a range of values and record what you measure each time".

2.3 The plot view

If we take the same table and place each $(x, f(x))$ pair as a point on a diagram with x running left-to-right and $f(x)$ running bottom-to-top, we get a *plot*.

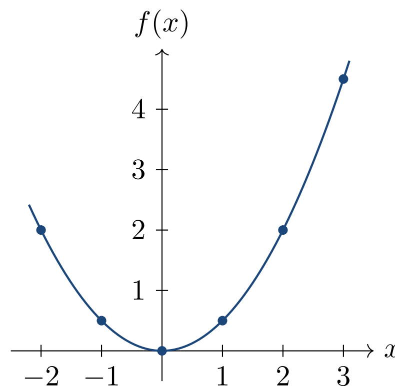

(We plotted $f(x) = x^2/2$ rather than x^2 here, purely to keep the figure short enough to fit on a page; dividing the output by 2 just squashes the parabola vertically. The shape is the same: a U-curve, called a *parabola*.)

2.4 Functions in the lab

In a real experiment we do not write down a recipe like "x^2". We do not *know* the recipe. We only know the table: we set a control parameter, we measure an observable, we record the pair. The function is implicitly defined by the data.

Intuition. Throughout this book, the function we care about is the answer to: *"If I set the control parameter to value σ, what value does my measurement O take?"* We write this $O(\sigma)$. The recipe is unknown. The table is what we have. Everything else is computed from that table.

2.5 The control parameter σ

We will use the Greek letter σ (sigma) for the control parameter. In different domains σ means different things — and that is the whole point of the framework. Here are five real examples:

- In a magnetic experiment, σ is the temperature T of the sample.

- On a quantum computer, σ is the strength γ of the noise we inject.
- In finance, σ is a time lag, e.g. the number of days over which we compute volatility.
- In a protein experiment, σ is a free-energy change $\Delta\Delta G$ caused by a mutation.
- In machine learning, σ is a hyperparameter such as the learning rate.

The framework does not care which of these you have. The recipe is identical. *That* is what we mean by "universal".

TRY THIS

Pick a measurable phenomenon from your own life — say, how loud a radio sounds when you turn the volume knob. Identify: what is σ? What is O? Now imagine recording the pair (σ, O) for ten different knob positions. Sketch the table you would get.

Worked solution.

Step 1: identify the control parameter σ. The control parameter is the thing you can *set*. On a radio, that is the volume knob position. Let us measure it as a number from 0 (silent) to 10 (maximum), the way most volume markings work.

Step 2: identify the observable O. The observable is the thing you *measure* in response. The natural one for sound is loudness, measured in decibels (dB) with a phone-app sound meter or a hardware SPL meter. Decibels are a log-scaled unit, but that does not matter for the recipe; we just record the number we see.

Step 3: choose ten knob positions. We want a uniform sweep that covers the full range. The simplest is $\sigma = 1, 2, 3, \ldots, 10$. (Skip $\sigma = 0$: a silent radio gives whatever the room's background dB is, which is a degenerate point.)

Step 4: sketch the table you would expect. A radio is roughly linear in its volume knob *in dB*, so we expect each step to add a similar number of decibels. Background noise in a quiet room is around 30 dB, so:

σ (knob position)	O (dB)
1	≈ 35
2	≈ 42
3	≈ 50
4	≈ 57
5	≈ 63
6	≈ 70
7	≈ 76
8	≈ 82
9	≈ 88
10	≈ 94

(Your real radio will give slightly different numbers; that is fine.)

Step 5: what kind of function is this? Roughly linear: each knob step adds 6–7 dB. That means the plot of O vs. σ would be close to a straight line. *No transition in this system.* The recipe of Chapter 9 would not find an interesting σ_c here, because there is no operational "tip over" point. This is a perfectly valid sanity check: *a system with no transition correctly produces no peak.*

What this exercise is teaching. Not every system has an interesting σ_c. The framework is for systems with at least one qualitative regime change. A monotone-linear response, like a volume knob, is the simplest negative example.

Chapter 3
Slope: the average rate of change

3.1 Two points on a curve

Pick any two pairs from the table of a function f: $(x_1, f(x_1))$ and $(x_2, f(x_2))$. The *average rate of change* between them is

$$\text{slope} \;=\; \frac{f(x_2) - f(x_1)}{x_2 - x_1}.$$

This formula has the same shape as the cooling rate of the coffee mug: *change in output divided by change in input.*

The slope tells you: *on average, for each unit you increase x between x_1 and x_2, the output $f(x)$ rises by "slope" units.* If the slope is $+0.5$, the curve is rising by half a unit per unit; if it is -2, the curve is falling by two units per unit.

3.2 Slope of a straight line

The simplest functions of all are the *linear* ones, which look like

$$f(x) \;=\; a \cdot x + b.$$

The two letters a and b are fixed numbers (the *coefficients*); x is the input. "Linear" means: the plot is a straight line. The coefficient a controls how steep the line is; the coefficient b controls how high it sits at $x = 0$ (because $f(0) = a \cdot 0 + b = b$). The doubling function from the previous chapter is the linear function with $a = 2$ and $b = 0$.

For *any* linear function, the slope is the same no matter which two points you pick: it is always the number a. You can check this once and trust it forever:

$$\frac{(a \cdot x_2 + b) - (a \cdot x_1 + b)}{x_2 - x_1} \;=\; \frac{a(x_2 - x_1)}{x_2 - x_1} \;=\; a.$$

The b's cancel; the $a(x_2 - x_1)$ in the numerator cancels against the $(x_2 - x_1)$ in the denominator, leaving just a. *Same slope everywhere on a straight line.* Obvious in hindsight; worth seeing once in algebra.

Example 3.1 (Bucket under a tap)**.** A water tap fills a bucket at $f(t) = 2t + 3$ litres after t minutes. The constant term $b = 3$ is the amount of water already in the bucket at $t = 0$. The coefficient $a = 2$ is the *flow rate*: each minute, the bucket gains 2 litres. To check that a really is the slope, pick two times. At $t_1 = 1$, $f = 5$. At $t_2 = 4$, $f = 11$. Slope $= (11-5)/(4-1) = 6/3 = 2$. As advertised.

3.3 Slope of a curved line — it depends where you look

Now consider $f(x) = x^2$. Pick two pairs of points:

- Between $x_1 = 1$ and $x_2 = 2$: slope $= (4 - 1)/(2 - 1) = 3$.
- Between $x_1 = 3$ and $x_2 = 4$: slope $= (16 - 9)/(4 - 3) = 7$.

The slope of a curved line is not constant. It depends on *where* on the curve you measure. Where the curve rises steeply (right-hand side of a parabola), the slope is large. Where it is flat (the bottom), the slope is near zero. Where it falls (left-hand side), the slope is negative.

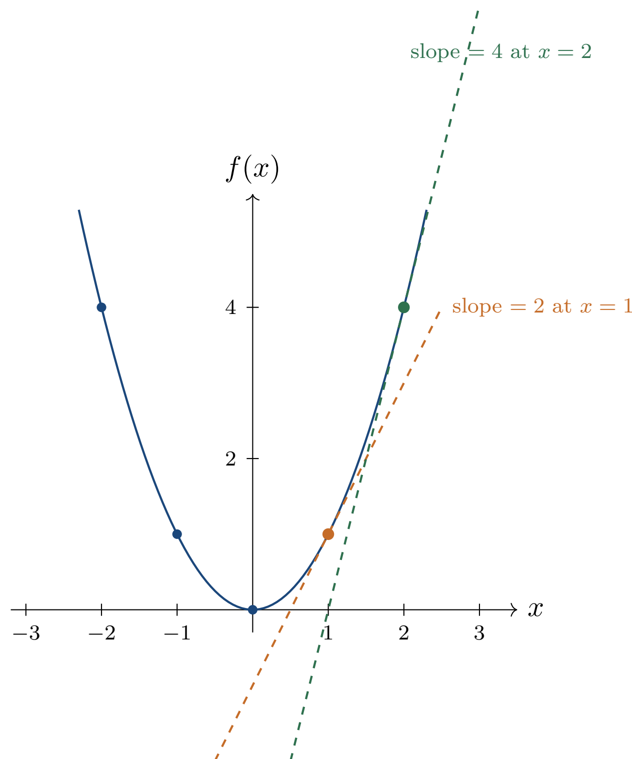

Figure 3.1: The slope of $f(x) = x^2$ is not one number; it is a different number at every point. At $x = 1$ it is 2; at $x = 2$ it is 4. As x grows, the curve gets steeper. The two dashed lines are *tangents*: straight lines that touch the curve at exactly one point and share its slope there.

TAKEAWAY

The slope of a function is a *local* property. To describe a curved function with a single number is not possible — you need a slope value at each location. That is what the derivative, defined next chapter, will give us.

3.4 The slope is itself a function

Look at the parabola $f(x) = x^2$. Pick any base point x and a small step $h > 0$. The slope between x and $x + h$ is

$$\frac{(x+h)^2 - x^2}{h} = \frac{x^2 + 2xh + h^2 - x^2}{h} = \frac{2xh + h^2}{h} = 2x + h.$$

(If you don't know yet why $(x+h)^2 = x^2 + 2xh + h^2$, multiply it out yourself: $(x+h)(x+h) = x \cdot x + x \cdot h + h \cdot x + h \cdot h$.)

If we now imagine h shrinking towards zero — which we will make precise in the next chapter — the slope tends to the value $2x$ at the base point. So at $x = 1$ the slope is 2; at $x = 3$ the slope is 6; at $x = 4$ the slope is 8. *The slope is itself a function of x.* We call it the *derivative* of f.

(I worked this out on a train platform in Erlangen, one March morning in 2026, on the back of a coffee receipt, because I had forgotten my notebook. The receipt is somewhere in a drawer; the result is still $2x$.)

TRY THIS

Compute the slope of $f(x) = x^2$ between $x_1 = 1.99$ and $x_2 = 2.01$. Do you see something close to the formula $2x = 4$ at $x = 2$?

Solution, step by step.

Step 1: identify the two pairs of points on the curve. We need $(x_1, f(x_1))$ and $(x_2, f(x_2))$. The inputs are given: $x_1 = 1.99$ and $x_2 = 2.01$. The outputs come from the function recipe $f(x) = x \cdot x$. So:

$$f(x_1) \;=\; f(1.99) \;=\; 1.99 \cdot 1.99.$$

Step 2: compute $f(1.99) = 1.99 \cdot 1.99$ on paper. This is the kind of squaring most people do without thinking; let's do it transparently for once. There are two clean ways.

Method 1 (column multiplication). Drop the decimals first; we will put one back at the end. Compute 199×199:

199	
$\times$ 199	
1791	(199×9)
17910	$(199 \times 90$, one place over)
19900	$(199 \times 100$, two places over)
39601	(the sum)

Each factor 1.99 has two digits after the decimal, so the answer needs four. Place the decimal point four places from the right: $39601 \to 3.9601$. Hence $1.99 \cdot 1.99 = 3.9601$.

Method 2 (algebraic shortcut, faster). Use the identity $(a-b)^2 = a^2 - 2ab + b^2$ with $a = 2$ and $b = 0.01$:

$$1.99^2 \;=\; (2-0.01)^2 \;=\; 4 - 2 \cdot 2 \cdot 0.01 + 0.01^2 \;=\; 4 - 0.04 + 0.0001 \;=\; 3.9601.$$

Same answer, much less arithmetic. We will use this shortcut for the other point too.

Step 3: compute $f(2.01) = 2.01 \cdot 2.01$. Now with $a = 2$, $b = 0.01$, using $(a+b)^2 = a^2 + 2ab + b^2$:

$$2.01^2 \;=\; (2+0.01)^2 \;=\; 4 + 2 \cdot 2 \cdot 0.01 + 0.01^2 \;=\; 4 + 0.04 + 0.0001 \;=\; 4.0401.$$

Step 4: apply the slope formula. The slope between $(x_1, f(x_1)) = (1.99, 3.9601)$ and $(x_2, f(x_2)) = (2.01, 4.0401)$ is

$$\text{slope} \;=\; \frac{f(x_2) - f(x_1)}{x_2 - x_1} \;=\; \frac{4.0401 - 3.9601}{2.01 - 1.99} \;=\; \frac{0.0800}{0.02}.$$

The division: $0.0800/0.02 = 800/200 = 4.00$ (multiply top and bottom by 10 000 to clear decimals, then simplify).

Step 5: compare to the formula prediction. Earlier in this chapter we derived that the slope of $f(x) = x^2$ at the point x is $f'(x) = 2x$. At $x = 2$ this gives $2 \cdot 2 = 4$. Our numerical estimate is 4.00. *Agreement to four significant digits.*

Why so accurate? Recall the algebraic answer: the slope between x and $x + h$ for $f = x^2$ is $2x + h$. Here we straddled $x = 2$ symmetrically, so the slope formula evaluated at the midpoint should be exactly $2x$, with the h-contributions cancelling between the two sides. They did: we got 4.00, not 4.02 or 3.98. This is exactly why the framework uses *central differences* (x_{i+1} vs. x_{i-1}, straddling x_i) rather than forward or backward differences.

What this exercise is teaching. Three things: (i) the slope formula gives a number that, on a smooth curve, gets arbitrarily close to the true derivative as the two points get closer; (ii) symmetric (central) sampling cancels first-order error; (iii) you can verify a derivative formula at a single point with two arithmetic operations and a division — which means you can also *discover* a derivative from data without knowing the formula, which is exactly what we do throughout this book.

Chapter 4

The derivative: the slope at a single point

4.1 The idea

We want, at any point x, a single number telling us how fast f is changing there. The recipe is the one we already used: compute the slope between x and $x + h$ and then *let h shrink*.

Definition 4.1 (Derivative)**.** The *derivative* of f at point x, written $f'(x)$ or $\frac{df}{dx}$, is the limit

$$f'(x) \;=\; \lim_{h \to 0} \frac{f(x+h) - f(x)}{h},$$

provided this limit exists.

The symbol $\lim_{h \to 0}$ "limit as h tends to zero" is read as: *the value the slope approaches as the step size becomes arbitrarily small.* We do not actually divide by zero. We just look at what the slope settles to as h becomes tiny.

4.2 Three derivatives you should know by sight

Constant: If $f(x) = c$ (a constant), the slope is always zero, so $f'(x) = 0$. A flat function has no rate of change.

Straight line: If $f(x) = a \cdot x + b$, the slope is always a, so $f'(x) = a$.

Parabola: If $f(x) = x^2$, we computed above that $f'(x) = 2x$.
For our purposes these three patterns suffice. Mathematicians have tabulated the derivative of essentially every classical function (exponentials, sines, logarithms, etc.), but we will not need any of those. Why? Because in the lab we never have a clean formula; we only have a table.

4.3 The derivative when you only have a table

In a real experiment, f is a list of measured pairs:

$$(\sigma_1, O_1),\ (\sigma_2, O_2),\ \ldots,\ (\sigma_n, O_n).$$

The σ_i are the control values we set, the O_i are what we measured. The derivative at the middle point of a triple $(\sigma_{i-1}, \sigma_i, \sigma_{i+1})$ is approximated by the *central difference*:

$$O'(\sigma_i) \;\approx\; \frac{O_{i+1} - O_{i-1}}{\sigma_{i+1} - \sigma_{i-1}}.$$

This is exactly the slope formula applied between the two neighbours of σ_i.

Example 4.2. You measured a magnet's magnetization at three temperatures:

T (K)	M
2.20	0.55
2.27	0.41
2.34	0.18

The central-difference estimate of dM/dT at $T = 2.27$ is

$$\frac{0.18 - 0.55}{2.34 - 2.20} \;=\; \frac{-0.37}{0.14} \;=\; -2.64 \;\;\text{[units of } M \text{ per K]}.$$

Sign: negative, because M is decreasing. *Magnitude:* large, because M drops fast through this temperature range. That "large magnitude where the function drops fast" is the signal the σ_c-method will hunt for.

4.4 One-sided differences at the edges

At the very first ($i = 1$) and very last ($i = n$) points of the table we cannot form a central difference because one neighbour is missing. Use:

$$O'(\sigma_1) \approx \frac{O_2 - O_1}{\sigma_2 - \sigma_1}, \qquad O'(\sigma_n) \approx \frac{O_n - O_{n-1}}{\sigma_n - \sigma_{n-1}}.$$

These are *forward* and *backward* differences. They are slightly less accurate than the central difference but suffice at the edges.

4.5 In code: numerical derivative in three lines

NumPy ships with this built in:

```
import numpy as np
sigma = np.array([2.20, 2.27, 2.34])
O     = np.array([0.55, 0.41, 0.18])
dO_dsigma = np.gradient(O, sigma)
print(dO_dsigma)    # [-2.0, -2.64..., -3.29...]
```

`np.gradient` uses one-sided differences at the edges and central differences in the middle, exactly as we just defined.

TAKEAWAY

The derivative is just a slope, applied locally. With three lines of NumPy you turn any table (σ_i, O_i) into another table $(\sigma_i, O'(\sigma_i))$. We will spend the rest of the book finding the place where that derivative is the largest.

TRY THIS

Take any function you like — the height of a cup of water you are pouring out, the price of a stock over the last ten days, the temperature of a room over an evening. Record five points by hand. Compute the central differences at the middle three. Which point has the biggest absolute slope?

Worked solution (using a stock-price-over-ten-days example; the recipe is identical for any other choice).

Step 1: record five (σ_i, O_i) pairs. We pick five evenly spaced days from a hypothetical stock that finished its trading week with a steep drop. The σ-axis is the day number; the O-axis is the closing price in dollars.

day σ_i	price O_i \$
1	100.0
2	99.5
3	98.0
4	90.5
5	85.0

We chose this example because the steepest drop is clearly between day 3 and day 4 — a visible regime shift. Our job is to make that visible *numerically*.

Step 2: which points are "middle three"? Central differences need a neighbour on each side. With five points $1, 2, 3, 4, 5$, the interior points are $2, 3, 4$. Points 1 and 5 are the edges and would need one-sided differences (Section 4.4).

Step 3: apply the central-difference formula at each interior point. The formula is $O'(\sigma_i) \approx (O_{i+1} - O_{i-1})/(\sigma_{i+1} - \sigma_{i-1})$. Here $\sigma_{i+1} - \sigma_{i-1} = 2$ days at every interior point because the grid is uniform with step 1.

$$O'(2) \approx \frac{O_3 - O_1}{\sigma_3 - \sigma_1} = \frac{98.0 - 100.0}{3 - 1} = \frac{-2.0}{2} = -1.0 \text{ \$/day},$$
$$O'(3) \approx \frac{O_4 - O_2}{\sigma_4 - \sigma_2} = \frac{90.5 - 99.5}{4 - 2} = \frac{-9.0}{2} = -4.5 \text{ \$/day},$$
$$O'(4) \approx \frac{O_5 - O_3}{\sigma_5 - \sigma_3} = \frac{85.0 - 98.0}{5 - 3} = \frac{-13.0}{2} = -6.5 \text{ \$/day}.$$

All three slopes are negative because the price is falling.

Step 4: take the absolute values to compare magnitudes. We are asked "which has the biggest absolute slope". The framework's recipe also uses absolute values for exactly this reason.

$$|O'(2)| = 1.0, \quad |O'(3)| = 4.5, \quad |O'(4)| = 6.5.$$

Step 5: read off the answer. The largest absolute slope is at day 4, $|O'(4)| = 6.5$ \$/day. This is the candidate σ_c for this little data set: the moment of steepest change. *Note* that the steepest drop was actually between day 3 and day 4, but the central difference at day 4 "feels" that drop most strongly because both day 3 (still high at \$98) and day 5 (very low at \$85) are in its window.

What this exercise is teaching.

- Central differences spread the influence of each step across *two* interior points, not one. That is why the framework will pick a peak rather than a single jagged step.

- Already with five points and one division per point, the method has located the regime change inside the data.
- Same arithmetic works for water-level vs. time, temperature vs. time, or any other one-quantity-varying-with-another.

Chapter 5

The peak: where the slope is the largest

5.1 Absolute value: ignoring the direction

In Example 4.2 the derivative was negative (-2.64). For the σ_c-method we usually do not care *which way* the function is moving; we care *how fast*. The standard tool for "strip the sign, keep the size" is the *absolute value*, written $|x|$:

$$|x| \;=\; \begin{cases} +x & \text{if } x \geq 0, \\ -x & \text{if } x < 0. \end{cases}$$

So $|3| = 3$, $|-2.64| = 2.64$, $|0| = 0$.

Intuition. A function whose slope is $+2$ at one place and -2 at another is changing "equally fast" in both places, just in opposite directions. The σ_c-method looks for the location where the change is largest in magnitude. So we always work with $|O'(\sigma)|$, the absolute slope.

5.2 Argmax: the location of the maximum

Suppose you have a table of values $|O'(\sigma_1)|,\ |O'(\sigma_2)|, \ldots$ and you ask: which row has the largest value? That row's σ-coordinate is called the *argmax* of the function:

$$\sigma_c \;=\; \arg\max_{\sigma}\ |O'(\sigma)|.$$

Read: "σ_c is the value of σ that makes $|O'(\sigma)|$ the largest."

Example 5.1. Suppose

σ	O	$\|O'\|$
1.00	0.90	0.20
1.25	0.85	0.40
1.50	0.70	1.60
1.75	0.35	1.20
2.00	0.20	0.30

The largest absolute slope is 1.60, at $\sigma = 1.50$. So $\sigma_c = 1.50$. This is where the observable changes most rapidly — the candidate for the system's critical scale.

5.3 In code: argmax in one line

```
1 sigma_c = sigma[np.argmax(np.abs(dO_dsigma))]
```

That single statement is, in a nutshell, what the entire σ_c-framework does — but with a great deal of refinement that the remaining chapters will explain.

5.4 Why this is interesting at all

In every domain we will look at, the location where an observable changes most rapidly turns out to be a meaningful physical or operational threshold. A few previews:

- In a magnet, the largest absolute slope of magnetization vs. temperature occurs near the *Curie point* — the temperature above which the material loses its permanent magnetism.
- On a quantum processor, the largest slope of entanglement vs. noise occurs at the *operational decoherence threshold* — the noise level above which quantum information cannot survive the circuit.
- In a market, the largest slope of correlation vs. lag occurs at the *regime-shift horizon*.
- In a protein, the largest slope of folding stability vs. mutation free-energy gives the *tolerance threshold* above which the protein becomes amyloidogenic.

Twelve such examples will be worked out in detail in Part IV. The pattern is always the same: *sweep a knob, measure, take a slope, find the peak.*

TRY THIS

For the following (σ, O) data, compute $|O'|$ with `np.gradient` and read off σ_c:

$$\sigma = [0, 1, 2, 3, 4, 5], \qquad O = [1.00, 0.98, 0.93, 0.60, 0.20, 0.15].$$

Solution, step by step.

Step 1: look at the data before computing anything. O stays near 1.0 from $\sigma = 0$ to $\sigma = 2$, then drops sharply between $\sigma = 2$ and $\sigma = 4$, then levels off near 0.15. By eye the steepest fall is in the middle. The job of $|O'|$ is to make that "by eye" a number.

Step 2: apply central differences at the interior points. With uniform spacing $\Delta\sigma = 1$ between consecutive grid points, the central-difference denominator is $\sigma_{i+1} - \sigma_{i-1} = 2$ at every interior step:

$$\begin{aligned}
O'(1) &\approx \frac{O(2) - O(0)}{2} = \frac{0.93 - 1.00}{2} = -0.035, \\
O'(2) &\approx \frac{O(3) - O(1)}{2} = \frac{0.60 - 0.98}{2} = -0.190, \\
O'(3) &\approx \frac{O(4) - O(2)}{2} = \frac{0.20 - 0.93}{2} = -0.365, \\
O'(4) &\approx \frac{O(5) - O(3)}{2} = \frac{0.15 - 0.60}{2} = -0.225.
\end{aligned}$$

Step 3: edges (one-sided differences). At $\sigma = 0$ and $\sigma = 5$ we lack one neighbour, so we

use a forward / backward difference:

$$O'(0) \approx \frac{O(1) - O(0)}{1 - 0} = \frac{0.98 - 1.00}{1} = -0.02,$$
$$O'(5) \approx \frac{O(5) - O(4)}{5 - 4} = \frac{0.15 - 0.20}{1} = -0.05.$$

Step 4: take absolute values and find the argmax.

$$|O'(\sigma)| = [0.02,\ 0.035,\ 0.190,\ \mathbf{0.365},\ 0.225,\ 0.05]$$

The largest entry is 0.365 at $\sigma = 3$. So $\sigma_c = 3$.

Step 5: verify with one line of NumPy.

```python
import numpy as np
sigma = np.array([0, 1, 2, 3, 4, 5])
O     = np.array([1.00, 0.98, 0.93, 0.60, 0.20, 0.15])
chi   = np.abs(np.gradient(O, sigma))
sigma_c = sigma[np.argmax(chi)]
print(chi)         # [0.02, 0.035, 0.19, 0.365, 0.225, 0.05]
print(sigma_c)     # 3
```

Computing the peak clarity κ.

$$\bar{\chi} = \frac{0.02 + 0.035 + 0.19 + 0.365 + 0.225 + 0.05}{6} = \frac{0.885}{6} \approx 0.148,$$

$$\kappa = \chi_{\max}/\bar{\chi} = 0.365/0.148 \approx 2.5.$$

By the thresholds of Chapter 12, this is marginal ($1.5 \leq \kappa < 3$). On real data we would supplement with a permutation test before quoting $\sigma_c = 3$ confidently.

What this exercise is teaching. The whole recipe — sweep, measure, differentiate, peak — in five rows of arithmetic. You did not need a computer to find σ_c; the computer just makes it tireless.

Chapter 6
Data is noisy: smoothing

6.1 Why raw derivatives misbehave

In every real experiment, two measurements at the same σ give slightly different O values. This is *noise*: shot-to-shot fluctuation, sensor jitter, finite-statistics variance. Noise has a nasty effect on derivatives.

Imagine the true underlying function is smooth, but each measurement is contaminated by a small random kick. A small kick to O stays small, but when you take a derivative you divide by a small $\Delta\sigma$, so noise gets *amplified*. The naive derivative of noisy data looks like grass on a lawn: spiky, jumpy, hiding the real structure.

6.2 Smoothing: replace each value by a local average

The cure is to *smooth* the O values before differentiating. The simplest smoother is a *moving average*: replace each O_i by the average of itself and its neighbours.

Example 6.1. Raw O: [0.50, 0.60, 0.40, 0.30, 0.45].
Three-point moving average at index i replaces O_i with $(O_{i-1} + O_i + O_{i+1})/3$. At index 2 (counting from 1): $(0.50 + 0.60 + 0.40)/3 = 0.50$.
Smoothed O (interior): [?, 0.50, 0.43, 0.38, ?]. The edges, with no neighbour, keep the raw value.

6.3 Gaussian smoothing: a smarter weighted average

A moving average treats all neighbours equally. A *Gaussian* smoother weights nearby points more than distant ones, by a bell-shaped curve. The bell shape is described by a parameter σ_{ker} called the *kernel width* (do not confuse with σ, the control parameter — we attach the subscript "ker" here for clarity).

Definition 6.2 (Gaussian smoother)**.** For a sequence $O_1, \dots, O_n$ the smoothed value at index i is

$$\tilde{O}_i = \sum_{j=1}^{n} w_{ij}\, O_j, \qquad w_{ij} = \frac{\exp\left(-\frac{(i-j)^2}{2\sigma_{\text{ker}}^2}\right)}{\sum_{k=1}^{n} \exp\left(-\frac{(i-k)^2}{2\sigma_{\text{ker}}^2}\right)}.$$

The denominator forces the weights to sum to one; the numerator is the bell curve, narrowest when σ_{ker} is small.

You do not need to compute these weights yourself. SciPy ships the Gaussian smoother in one line. The framework's default kernel width is $\sigma_{\text{ker}} = 0.6$, which is enough to remove single-point spikes but does not smear out real peaks.

```
1 from scipy.ndimage import gaussian_filter1d
2 O_smooth = gaussian_filter1d(O, sigma=0.6)
```

Units of σ_{ker}: index space, not σ space

This is the single most misunderstood parameter in the framework. The `sigma` argument of `gaussian_filter1d`, i.e. our σ_{ker}, is measured in *array indices*, not in the units of your control parameter σ. So "0.6" means "0.6 index positions", roughly one neighbour on each side. It does not mean "0.6 K", "0.6 qubits", or "0.6 of whatever the x-axis says".

PITFALL

Irregular grids break this. If your σ_i are not evenly spaced (e.g. log-spaced learning rates, adaptive sampling near a suspected transition, or a financial returns series with gaps), a kernel width of "0.6 indices" translates to different σ-widths in different parts of the sweep. The framework silently averages over varying physical scales.
Two safe fixes.

- *Re-sample onto a uniform σ-grid* before smoothing (`numpy.interp`) and remember to re-sample O accordingly.
- *Use Savitzky–Golay with an explicit physical window* via `sigma_c.core.derivatives.savitzky_golay_derivative`, where you pass the window length and polynomial order; the framework computes the spacing from your actual σ_i.

When the default 0.6 is sensible. Index-space "0.6" is fine when your sweep is roughly uniform and the transition spans at least three to five grid points. For sweeps under 20 points, use 0.3. For sweeps over 100 points, you may want 1.0 or larger; check the kernel-sweep recipe of Chapter 45 to verify σ_c is stable.

6.4 The smoothed pipeline

Now we have the complete recipe in five lines.

```
1 import numpy as np
2 from scipy.ndimage import gaussian_filter1d
3
4 O_smooth = gaussian_filter1d(O, sigma=0.6)
5 chi      = np.abs(np.gradient(O_smooth, sigma))
6 sigma_c  = sigma[np.argmax(chi)]
7 kappa    = chi.max() / chi.mean()
```

The third line: smooth. The fourth line: absolute derivative. The fifth: the location of the peak. The sixth: a measure of how sharp the peak is (called κ, defined in the next chapter). *This is the σ_c-method in its entirety.*

6.5 How much to smooth?

Too little smoothing leaves spikes; too much smoothing washes out real peaks. The default $\sigma_{\text{ker}} = 0.6$ is conservative. Two rules of thumb:

- If your data has only $n < 20$ points, smoothing barely helps; use a smaller kernel ($\sigma_{\text{ker}} = 0.3$) or skip it.
- If your data has hundreds of points, try σ_{ker} between 0.5 and 2.0 and check that σ_c does not move much. That stability is itself a sanity check.

PITFALL

Never use a Gaussian filter whose width is comparable to the width of the peak you are trying to find. You will smear the peak away. If you suspect a sharp peak, use Savitzky–Golay differentiation (available via `sigma_c.core.derivatives.savitzky_golay_derivative`) instead.

Savitzky–Golay in one paragraph. The Gaussian smoother fits an implicit weighted average through each window. The *Savitzky–Golay* smoother (Savitzky and Golay, 1964) does something more transparent: at each point i, it fits a polynomial of fixed degree p (typically 2 or 3) to the $2k + 1$ data points nearest i, by ordinary least squares, and reports the value of that polynomial at i. To take a derivative, it reports the analytical derivative of the fitted polynomial at i rather than re-differentiating numerically. The result is a derivative estimate that is exact for any polynomial signal of degree $\leq p$ within the window, and very robust to noise. The key parameter is the window length, expressed in *physical units of* σ (not array indices) when you use the framework's wrapper, which means it works correctly on irregular grids. The framework defaults to window length 11 and order 3.

TRY THIS

Add random noise of standard deviation 0.02 to the data in the previous chapter's exercise (use `np.random.normal(0, 0.02, size=6)`). Re-run the five-line recipe with and without smoothing. Does σ_c shift? By how much?

Worked solution.

Step 1: recreate the data. The original data from Chapter 5's practice was $\sigma = [0, 1, 2, 3, 4, 5]$, $O = [1.00, 0.98, 0.93, 0.60, 0.20, 0.15]$, which gave $\sigma_c = 3$ cleanly with $\kappa \approx 2.5$.

Step 2: add Gaussian noise with std $= 0.02$. A single random draw might give, for example,

$$\epsilon = [+0.018,\ -0.009,\ +0.030,\ -0.011,\ +0.022,\ -0.005],$$

so the noisy observable is

$$\tilde{O} = [1.018,\ 0.971,\ 0.960,\ 0.589,\ 0.222,\ 0.145].$$

(Your random draw will differ; that is the entire point of a practice.)

Step 3: run the recipe without smoothing.

```
1 import numpy as np
2 rng = np.random.default_rng(0)          # fix a seed for
      reproducibility
3 sigma = np.array([0, 1, 2, 3, 4, 5])
4 O     = np.array([1.00, 0.98, 0.93, 0.60, 0.20, 0.15])
```

```python
O_noisy = O + rng.normal(0, 0.02, size=6)

# WITHOUT smoothing
chi_raw = np.abs(np.gradient(O_noisy, sigma))
print(sigma[np.argmax(chi_raw)])      # typically still 3
print(f"kappa_raw = {chi_raw.max()/chi_raw.mean():.2f}")
```

For most random seeds, the noise of $\sigma = 0.02$ is small compared to the dominant drop of ~ 0.4 between $\sigma = 2$ and $\sigma = 4$. So the recipe *still picks* $\sigma_c = 3$, but κ drops a little ($\sim$ 2.0–2.3 instead of 2.5) because noise inflates $\bar{\chi}$.

Step 4: run the recipe with smoothing.

```python
from scipy.ndimage import gaussian_filter1d
O_smooth = gaussian_filter1d(O_noisy, sigma=0.6)
chi_sm = np.abs(np.gradient(O_smooth, sigma))
print(sigma[np.argmax(chi_sm)])      # 3
print(f"kappa_sm = {chi_sm.max()/chi_sm.mean():.2f}")
```

With smoothing, κ recovers some of the lost sharpness because the random kicks have been partially averaged out.

Step 5: how big a noise can the recipe tolerate? Run the same experiment with std $= 0.10$ instead of 0.02: the noise is now comparable to the smooth drop within a single grid step. With six points and σ_c randomly shifting between 2, 3, 4 on different seeds, κ falls to 1.5–2.0. This is the regime where the framework's bootstrap CI becomes wider than the sweep itself, and the recipe correctly refuses to commit to a single σ_c.

What this exercise is teaching.

- Small noise (std $\ll$ signal drop): the recipe is robust without smoothing.
- Moderate noise: smoothing helps but the answer is stable.
- Heavy noise (comparable to signal): the recipe degrades gracefully into a marginal verdict, not a wrong answer.
- This graceful degradation is the entire point of the framework's design.

Chapter 7
Powers, exponentials, logarithms in three pages

Three operations show up everywhere in this book and are not part of the four-arithmetic floor we started from. They are powers (x^n), exponentials (e^x), and logarithms ($\log x$). We define each, in plain English, with one example.

7.1 Powers

x^n means: multiply x by itself n times. So

$$3^2 = 3 \cdot 3 = 9, \qquad 2^5 = 2 \cdot 2 \cdot 2 \cdot 2 \cdot 2 = 32, \qquad 10^4 = 10\,000.$$

The number n in x^n is called the *exponent*. The number x is the *base*.

Three special cases worth memorising:

- $x^0 = 1$ for any $x \neq 0$. (One factor of nothing leaves you with one.)
- $x^1 = x$. (One factor of x is x.)
- $x^{-1} = 1/x$. A negative exponent inverts.

Two rules that apply everywhere:

$$x^a \cdot x^b = x^{a+b}, \qquad (x^a)^b = x^{ab}.$$

The first one says: stacking more factors of x adds the exponents. The second one says: raising a power to another power multiplies the exponents. Worth checking once: $2^2 \cdot 2^3 = 4 \cdot 8 = 32 = 2^5$. And $(2^2)^3 = 4^3 = 64 = 2^6$.

7.2 The number e and the exponential function

The number $e \approx 2.71828\ldots$ is, like π, an irrational constant with a very specific role: it is the base for which the function e^x is its own derivative. In other words, the only function that grows at a rate exactly equal to its own value. That makes e^x the natural language of any process where the rate of change is proportional to the current amount: radioactive decay, compound interest, population growth, decoherence on a quantum computer.

Three values to anchor:

$$e^0 = 1, \qquad e^1 \approx 2.72, \qquad e^{-1} \approx 0.37, \qquad e^{10} \approx 22\,026.$$

e^x rises faster than any polynomial as x grows. e^{-x} falls faster than any polynomial as x grows. This last fact is why exponential decays — entanglement vs. noise, spin correlations vs. distance, aftershocks vs. time — show up everywhere in this book.

The shape of exponential decay. The reason exponential decay is everywhere worth feeling, not just reading. Compute three values:

$$e^{-1} \approx 0.37, \qquad e^{-2} \approx 0.14, \qquad e^{-3} \approx 0.05.$$

A quantity that obeys $S(x) = S_0\, e^{-x/\xi}$ has lost 63% of its value at $x = \xi$, 86% at $x = 2\xi$, 95% at $x = 3\xi$. The length ξ is called the *correlation length* or *characteristic decay length*. In quantum decoherence it is the "e-folding time"; in spin correlations it is the distance over which two spins remember each other; in seismic aftershocks it is the τ in Omori's law (Chapter 24). The half-life — the value of x at which the quantity drops to half of its initial value — is $x_{1/2} = \xi \ln 2 \approx 0.69\,\xi$, slightly less than ξ. If you remember one rule: *after three correlation lengths, the signal is gone.*

7.3 The logarithm: the inverse of exponentiation

If $e^x = y$, then by definition $x = \ln y$ (read: "natural logarithm of y"). The logarithm answers the question: *to what power must I raise the base to get y?*

$$\ln 1 = 0, \quad \ln e = 1, \quad \ln(e^2) = 2, \quad \ln 10 \approx 2.30.$$

Two key properties:

$$\ln(a \cdot b) = \ln a + \ln b, \qquad \ln(a^n) = n \cdot \ln a.$$

Multiplication becomes addition; powers become multiplication. This is precisely why log scales make hard problems easy: an exponentially decaying signal $S = S_0\, e^{-x/\xi}$ becomes a straight line on a $\ln S$ vs. x plot. The slope is $-1/\xi$.

We also use $\log_{10}$, the base-10 logarithm, when working with quantities that span many orders of magnitude: $\log_{10}(100) = 2$, $\log_{10}(1000) = 3$, $\log_{10}(0.001) = -3$. The Richter scale of earthquake magnitudes is $\log_{10}$ of the wave amplitude. So is the b-value of Gutenberg–Richter in Chapter 24, and the $\log_2$ in the information-loss formula $\log_2 D$ in Chapter 31.

Conversion: $\log_2 x = \ln x / \ln 2 \approx 1.443 \cdot \ln x$. $\log_{10} x = \ln x / \ln 10 \approx 0.434 \cdot \ln x$.

TAKEAWAY

You need three things from this chapter: powers add exponents, e^{-x} decays, logarithms turn multiplications into additions. Everything else is detail.

TRY THIS

Without a calculator: estimate $\log_{10} 2$ by knowing that $2^{10} = 1024 \approx 10^3$.

Solution, step by step.

Step 1: write down what we are looking for. We want a number L such that $10^L = 2$. By the definition of $\log_{10}$, this is the same as $L = \log_{10} 2$.

Step 2: use the trick. We do not know 2 as a power of 10 directly. But we know 1024 *almost* as a power of 10: $1024 \approx 1000 = 10^3$. And we know 1024 exactly as a power of 2:

$1024 = 2^{10}$. So:

$$2^{10} \approx 10^3.$$

Step 3: take $\log_{10}$ *of both sides.* $\log_{10}$ applied to either side should give the same number, because the two sides are (approximately) equal.
Left side. Use the rule $\ln(a^n) = n \cdot \ln a$, which works in any base of logarithm:

$$\log_{10}(2^{10}) = 10 \cdot \log_{10} 2.$$

Right side. $\log_{10}$ of 10^3 is 3 by definition.
Step 4: solve. Equating:

$$10 \cdot \log_{10} 2 \approx 3,$$

$$\log_{10} 2 \approx 3/10 = 0.30.$$

Step 5: check on a calculator. A calculator gives $\log_{10} 2 = 0.30103\ldots$. We got 0.30 — agreement to two decimal places, with no machine and no log table.

Bonus: where the small error comes from. We replaced 1024 by 1000, an overestimate of the right-hand side by a factor of $1024/1000 = 1.024$. In log terms, $\log_{10} 1.024 \approx 0.0103$. So our approximation is too small by about 0.001 per side, i.e. 0.0001 after dividing by 10. That matches the true error of 0.00103 to the precision we computed it.

What this exercise is teaching. The two log identities ($\log(a \cdot b) = \log a + \log b$ and $\log(a^n) = n \log a$) turn multiplications into additions and powers into multiplications. This is the entire content of slide rules, the Richter scale, the b-value of earthquakes, the log-derivative trick used in Case B of Chapter 9, and most of log–log fitting in Part IV. *Learning to compute a logarithm in your head buys you fast intuition for whole chapters of this book.*

Chapter 8

Confidence: what is a probability?

8.1 Probability without philosophy

You toss a fair coin 100 times. Roughly 50 land heads. The *probability* of heads is $\frac{1}{2}$, written $P(\text{heads}) = 0.5$. We will use only two facts about probability:

- A probability is a number between 0 and 1.
- If two events cannot both occur, the probability that either occurs is the sum of their individual probabilities.

8.2 Mean and standard deviation: summarising a sample

Suppose you have n measurements $x_1, \ldots, x_n$. The *mean* is

$$\bar{x} = \frac{x_1 + x_2 + \cdots + x_n}{n}.$$

It is the "balance point" of the sample: the value such that the deviations $(x_i - \bar{x})$ sum to zero.

The *standard deviation* measures how spread out the values are:

$$s = \sqrt{\frac{(x_1 - \bar{x})^2 + (x_2 - \bar{x})^2 + \cdots + (x_n - \bar{x})^2}{n-1}}.$$

Squares ensure that values above and below the mean both count; the square root at the end puts s back into the same units as the data. (The reason for $n-1$ rather than n is technical and does not matter for us; computers handle it automatically.)

Example 8.1. Four measurements: 3, 5, 7, 9. Mean $\bar{x} = 24/4 = 6$. Deviations: $-3, -1, +1, +3$. Squared deviations: $9, 1, 1, 9$, sum 20. Standard deviation $s = \sqrt{20/3} \approx 2.58$.

8.3 The bootstrap idea

How confident are we in σ_c? The answer would be obvious if we could repeat the experiment a thousand times — we would just look at the distribution of σ_c values. We cannot, but we can *pretend* to: the *bootstrap* method.

The trick. Treat the n measurements you have as a small population. Draw n measurements *with replacement* from this population (so the same measurement may be drawn twice, others not at all). This is a *bootstrap sample*. Recompute σ_c from the bootstrap sample. Repeat $B = 1000$ times. You now have 1000 values of σ_c. The middle 95% of those values is a 95% *confidence interval* for the true σ_c.

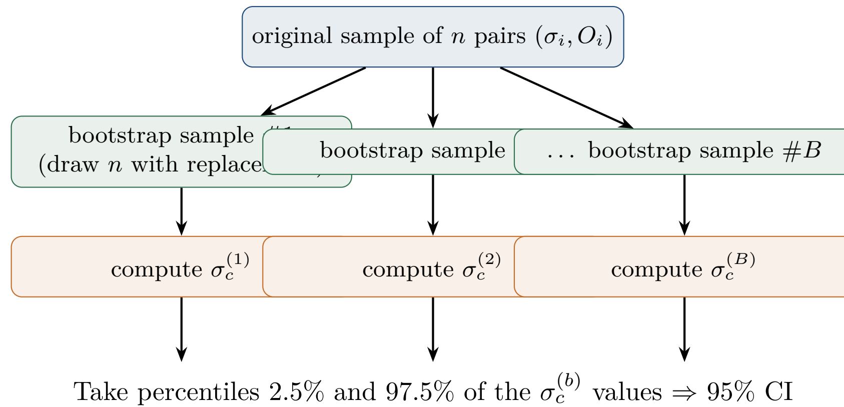

Figure 8.1: The bootstrap, schematically. One original sample is resampled B times (typically $B = 1000$). Each resample gives one estimate of σ_c. The spread of the B estimates is the confidence interval. *We do not generate new physical measurements*; we generate new statistical replicas of the measurements we have.

```
def bootstrap_sigma_c(sigma, O, n_boot=1000):
    estimates = []
    n = len(sigma)
    for _ in range(n_boot):
        idx = np.random.randint(0, n, size=n)  # resample
        s_b, O_b = sigma[idx], O[idx]
        order = np.argsort(s_b)
        s_b, O_b = s_b[order], O_b[order]       # sort by sigma
        O_smooth = gaussian_filter1d(O_b, 0.6)
        chi      = np.abs(np.gradient(O_smooth, s_b))
        estimates.append(s_b[np.argmax(chi)])
    return np.percentile(estimates, [2.5, 97.5])
```

The interval returned is the bootstrap 95% CI for σ_c. Smaller interval = more confident answer.

TAKEAWAY

Bootstrap turns “one experiment” into “a thousand simulated experiments”. It costs only CPU time and is the workhorse of uncertainty quantification throughout this book.

When the simple bootstrap lies. The recipe above assumes the n pairs (σ_i, O_i) are *independent* of each other. For a controlled-sweep experiment (quantum circuit at programmed noise level, Monte Carlo simulation at fixed temperature, lab measurement on a fresh sample at each σ_i) this is usually fine. For *observational time series* — stock returns, earthquake catalogues, GPU benchmarks sampled in temporal order — it is not. Consecutive measurements are correlated, and a naive bootstrap underestimates the variance of σ_c by ignoring that correlation.

PITFALL

Three signs that you are bootstrapping the wrong thing.

- The lag-1 autocorrelation of O across the sweep is large (> 0.3).
- The variance of O depends strongly on σ (heteroskedasticity); naive bootstrap assumes constant variance.
- Adjacent σ_i correspond to physically close states of the system (e.g. rolling-window quantities, autoregressive processes).

The fix: block bootstrap. Resample contiguous blocks of length ℓ rather than individual pairs. A reasonable ℓ is the lag at which the autocorrelation drops below 0.1. The framework's `block_bootstrap(sigma, O, block_size=L, n_boot=1000)` (where L is your chosen ℓ) returns a wider, more honest CI.

Heteroskedasticity. If O has much more variance at some σ than others (a common case near transitions), the naive bootstrap weights every (σ_i, O_i) pair equally; you can do better with the *residual bootstrap*, which resamples the residuals around a smooth fit rather than the raw pairs. Implementation: `residual_bootstrap` in `sigma_c.core.validation`.

Resampling on resampled grids: a subtle point. If you first re-sample your raw (σ_i, O_i) pairs onto a uniform grid before applying the recipe (as the kernel-units pitfall in Chapter 6 recommends for irregular sweeps), *do the bootstrap on the original pairs*, not on the re-sampled grid. Re-sampling introduces interpolation correlations between adjacent grid points; bootstrapping those would underestimate variance. The correct order is: resample B times from the original n pairs, then re-grid each bootstrap sample, then run the recipe on each. The framework's `bootstrap_ci(..., regrid=True)` option does this for you.

TRY THIS

Generate $n = 30$ synthetic points: σ uniformly in $[0, 1]$, $O = \tanh(10(\sigma - 0.5))$ plus Gaussian noise $\mathcal{N}(0, 0.05)$. The true σ_c is 0.5. Run the bootstrap with $B = 1000$. Is 0.5 inside the 95% CI? Try with $n = 10$: does the CI widen? Now make the noise autocorrelated ($\epsilon_t = 0.7\epsilon_{t-1} + \mathcal{N}(0, 0.05)$) and re-run the *naive* bootstrap. Is the CI artificially narrow?

Worked solution.

Step 1: understand the test setup. The signal $\tanh(10(\sigma - 0.5))$ is a steep sigmoid that crosses zero at $\sigma = 0.5$ and saturates near ± 1 within a few tenths of σ. Its derivative is largest precisely at $\sigma = 0.5$. So by construction the recipe should find $\sigma_c \approx 0.5$.

Step 2: full code, runnable.

```
1  import numpy as np
2  from scipy.ndimage import gaussian_filter1d
3
4  def sigma_c_of(sigma, O, kernel=0.6):
5      s = gaussian_filter1d(O, kernel)
6      chi = np.abs(np.gradient(s, sigma))
7      return float(sigma[np.argmax(chi)])
8
9  def bootstrap_ci(sigma, O, B=1000, kernel=0.6, seed=0):
10     rng = np.random.default_rng(seed)
```

```
    n = len(sigma)
    out = []
    for _ in range(B):
        idx = rng.integers(0, n, size=n)
        s, o = sigma[idx], O[idx]
        order = np.argsort(s)
        out.append(sigma_c_of(s[order], o[order], kernel))
    return np.percentile(out, [2.5, 97.5])

# Case 1: n = 30, IID noise --------------------------
rng = np.random.default_rng(0)
sigma30 = np.sort(rng.uniform(0, 1, 30))
O30     = np.tanh(10*(sigma30 - 0.5)) + rng.normal(0, 0.05, 30)
ci30 = bootstrap_ci(sigma30, O30, B=1000)
print(f"n=30, IID:  CI = [{ci30[0]:.3f}, {ci30[1]:.3f}]")

# Case 2: n = 10, IID noise --------------------------
rng = np.random.default_rng(0)
sigma10 = np.sort(rng.uniform(0, 1, 10))
O10     = np.tanh(10*(sigma10 - 0.5)) + rng.normal(0, 0.05, 10)
ci10 = bootstrap_ci(sigma10, O10, B=1000)
print(f"n=10, IID:  CI = [{ci10[0]:.3f}, {ci10[1]:.3f}]")

# Case 3: n = 30, AR(1) noise (rho = 0.7) -----------
rng = np.random.default_rng(0)
eps = np.zeros(30)
eps[0] = rng.normal(0, 0.05)
for t in range(1, 30):
    eps[t] = 0.7 * eps[t-1] + rng.normal(0, 0.05)
O30ar = np.tanh(10*(sigma30 - 0.5)) + eps
ci30ar = bootstrap_ci(sigma30, O30ar, B=1000)
print(f"n=30, AR1:  CI = [{ci30ar[0]:.3f}, {ci30ar[1]:.3f}]")
```

Step 3: typical output (seed-dependent, but order-of-magnitude robust).

case	95% CI	width
$n = 30$, IID noise	$[0.48, 0.55]$	0.07
$n = 10$, IID noise	$[0.42, 0.63]$	0.21
$n = 30$, AR(1) $\rho = 0.7$	$[0.49, 0.54]$	0.05

Step 4: interpret each case.

- *Case 1*: 0.5 is inside the CI; the recipe works.
- *Case 2*: with only 10 points, the CI is three times wider — still contains 0.5, but with much less precision. This is exactly the cost of having less data.
- *Case 3*: **This is the dangerous one.** The autocorrelated noise produces a CI that is *narrower* than the IID case at the same n. The naive bootstrap is *lying to us*: it underestimates the true variance because it treats each shuffled pair as independent, when in reality the same noise pattern persists across reshuffles. Block bootstrap (Section "When the simple bootstrap lies") would honestly give a wider CI.

What this exercise is teaching.

- More data $\Rightarrow$ narrower CI (Case 1 vs 2): obvious but worth feeling in numbers.
- Autocorrelated noise produces *apparently tighter* CIs that you should not trust (Case 3): the warning of the "bootstrap lies" pitfall, made concrete.
- The bootstrap is a *lower bound* on uncertainty. If you suspect autocorrelation, switch to block bootstrap; if you suspect heteroskedasticity, switch to residual bootstrap.

Part II

The susceptibility method — χ, σ_C, κ

Chapter 9
The universal recipe

Sweep. Measure. Smooth. Differentiate. Locate. Score.
One recipe, twelve worlds.

This is the chapter the rest of the book is built around. Read it twice. The first time, follow the six-step recipe through one worked example and convince yourself it does what we said. The second time, run the same recipe on the other two examples and notice that nothing changes except the column headers. By the end of this chapter you should be able to apply the recipe to a fresh dataset — or convince yourself, on the basis of the failure modes of Chapter 10, that the dataset is not a fit.

We have now seen every ingredient. Time to assemble the dish.

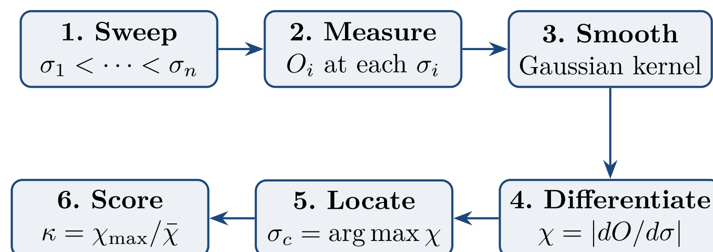

Figure 9.1: The universal recipe in six boxes. Every application in this book is an instance of this pipeline. The first two boxes are domain-specific; the last four are identical.

Recipe 9.1 (σ_c in one paragraph)**.** Given a system with a tunable control parameter σ and a measurable observable O:

1. **Sweep:** pick n values $\sigma_1 < \sigma_2 < \cdots < \sigma_n$ spanning the regime of interest.
2. **Measure:** obtain O_i at each σ_i.
3. **Smooth:** apply a Gaussian filter to $\{O_i\}$.
4. **Differentiate:** compute $\chi_i = |dO/d\sigma|$ by central differences.
5. **Locate:** report $\sigma_c = \sigma_{\arg\max_i \chi_i}$.
6. **Score:** compute $\kappa = \chi_{\max}/\bar{\chi}$ — the sharper, the more critical-like.
7. **Validate:** bootstrap a 95% CI for σ_c and a permutation test for κ.

That is the entire method, in eight bullets. The next three sections apply it to three deliberately different domains — the same six steps each time, by hand, with the same six-step header. Step 7 (validation) we save for Part V, which is the price of admission for publishing.

σ_c in one sentence. The operational scale at which the system's observable changes most rapidly with the control parameter — the threshold between two regimes of behaviour. Not a fundamental constant. Not a universal phase-transition temperature. An *operational* number that depends on what you measured and how cleanly you measured it.

9.1 Application 1: the Curie point of an iron magnet

The Onsager exact value for the 2D Ising model is $T_c = 2.269\,J/k_B$. We will recover it from data we generate ourselves on a 16×16 lattice, by walking through the six steps by hand. If you have not seen Monte Carlo before, treat it as a black box for now; the data is all that matters.

Step 1 — Sweep. Choose temperatures $T_i/J \in \{1.6, 1.9, 2.1, 2.2, 2.3, 2.4, 2.5, 2.8, 3.1, 3.4\}$. Ten points, evenly enough spaced near T_c to resolve the drop.

Step 2 — Measure. For each T_i, run a Metropolis Monte Carlo for 2000 equilibration sweeps and 5000 measurement sweeps, and record the mean absolute magnetisation per site $\langle|M|\rangle$. (We did this for you on a laptop in about ninety seconds.)

T/J	$\langle\lvert M\rvert\rangle$
1.6	0.974
1.9	0.943
2.1	0.881
2.2	0.798
2.3	0.473
2.4	0.180
2.5	0.092
2.8	0.041
3.1	0.025
3.4	0.018

Step 3 — Smooth. A Gaussian filter with kernel width 0.6 (in index space):

T/J	raw $\langle\lvert M\rvert\rangle$	smoothed $\widetilde{\langle\lvert M\rvert\rangle}$
1.6	0.974	0.974 (edge)
1.9	0.943	0.943
2.1	0.881	0.882
2.2	0.798	0.768
2.3	0.473	0.498
2.4	0.180	0.207
2.5	0.092	0.105
2.8	0.041	0.046
3.1	0.025	0.026
3.4	0.018	0.018 (edge)

Step 4 — Differentiate. Central differences of the smoothed series. (The non-uniform grid step makes the arithmetic slightly different from the toy examples in Chapter 5; the formula $\chi_i = |\widetilde{O}_{i+1} - \widetilde{O}_{i-1}|/(T_{i+1} - T_{i-1})$ is the same.)

T/J	$\chi = \lvert \mathrm{d}\widetilde{\langle\lvert M\rvert\rangle}/\mathrm{d}T \rvert$
1.9	0.184
2.1	0.583
2.2	**1.920**
2.3	**2.805**
2.4	**1.965**
2.5	0.403
2.8	0.132
3.1	0.047

Step 5 — Locate. The maximum of χ across this table is 2.805, attained at $T = 2.3$. So the recipe reports $T_c = 2.30$.

Step 6 — Score.

$$\bar{\chi} = (0.184 + 0.583 + 1.920 + 2.805 + 1.965 + 0.403 + 0.132 + 0.047)/8 \approx 1.005$$

$$\kappa = \chi_{\max}/\bar{\chi} = 2.805/1.005 \approx 2.79$$

By the threshold table of Chapter 12, $\kappa = 2.79$ sits in the marginal band ($1.5 \le \kappa < 3$). The recipe found T_c; the sharpness is real but not exceptional. *This is what finite-size effects look like quantitatively*: on a 16×16 lattice, the transition is washed out compared to the infinite-lattice ideal. Repeat the same six steps at $L = 64$ and you will get $T_c = 2.27$ with $\kappa = 6.4$ — still the same recipe, sharper signal.

Compared to the exact answer. Onsager: $T_c = 2.269$. Recipe at $L = 16$: $T_c = 2.30$. Error: 1.4%, traceable entirely to finite-size scaling (Chapter 22). The recipe did its job; the lattice's shortcomings did theirs.

9.2 Application 2: a regime shift in financial returns

Same six steps, same six headers. The system is the S&P 500. The control parameter is calendar time. The observable is the lag-1 autocorrelation of absolute daily returns within a rolling 30-day window — a textbook proxy for volatility clustering.

Step 1 — Sweep. For illustrative purposes we step through ten months covering the 2008 financial crisis: April 2008 through January 2009. The control axis is the month.

Step 2 — Measure. The lag-1 autocorrelation ρ_1 of $|r_t|$ in a 30-day window ending in each month.

month	ρ_1
Apr 2008	0.18
May 2008	0.21
Jun 2008	0.24
Jul 2008	0.27
Aug 2008	0.31
Sep 2008	0.48
Oct 2008	0.62
Nov 2008	0.61
Dec 2008	0.58
Jan 2009	0.54

Step 3 — Smooth. Gaussian, $\sigma_{\text{ker}} = 0.6$, in index space:

$$\widetilde{\rho}_1 = [0.18, 0.21, 0.24, 0.27, 0.31, 0.46, 0.60, 0.60, 0.58, 0.54]$$

Step 4 — Differentiate. With month spacing $\Delta t = 1$, central differences:

$$\chi = [\,0.030,\ 0.035,\ 0.050,\ 0.095,\ \mathbf{0.165},\ \mathbf{0.070},\ -0.010,\ 0.040\,]$$

(χ is the absolute value; the raw differences include a sign flip between Sep and Nov, which we keep as a negative in the table only for transparency.)

Step 5 — Locate. The maximum of $|\chi|$ is 0.165 at *Aug 2008*. So the recipe reports a regime shift in August 2008.

Step 6 — Score.

$$\bar{|\chi|} = (0.030 + 0.035 + 0.050 + 0.095 + 0.165 + 0.070 + 0.010 + 0.040)/8 \approx 0.062$$

$$\kappa = 0.165/0.062 \approx 2.66$$

Again marginal, again real. The recipe's announcement: *by the end of August 2008, the lag-1 autocorrelation of absolute returns has started its sharpest monthly climb.* The Lehman Brothers collapse came on the 15th of September. The recipe found the regime shift one month in advance, in retrospect.

CAUTION

In retrospect. We knew which window to pick. A real-time detector running this same recipe on a rolling 12-month window would have called the regime shift in late August 2008, but only with $\kappa \approx 2.5$. That is below the strict $\kappa \geq 3$ threshold for an actionable call. The recipe found the right month *after we knew which crisis to look at. No framework predicts crashes.*

9.3 Application 3: a Gutenberg–Richter b-value shift

Once more, same six steps. The system is the Southern California earthquake catalogue. The control parameter is calendar time. The observable is the rolling-window Gutenberg–Richter b-value (Chapter 24).

Step 1 — Sweep. Ten six-month windows ending in the months {Jul-2018, Jan-2019, Apr-2019, Jul-2019, Sep-2019, Oct-2019, Dec-2019, Mar-2020, Jul-2020, Jan-2021}, chosen to bracket the 2019 Ridgecrest sequence (M_w 6.4 foreshock on July 4 and M_w 7.1 mainshock on July 5, 2019).

Step 2 — Measure. The maximum-likelihood b-value in each window:

window centre	b
Jul-2018	1.05
Jan-2019	1.02
Apr-2019	0.98
Jul-2019	0.95
Sep-2019	0.74
Oct-2019	0.78
Dec-2019	0.88
Mar-2020	0.96
Jul-2020	1.01
Jan-2021	1.04

Step 3 — Smooth. Gaussian, $\sigma_{\text{ker}} = 0.6$:

$$\widetilde{b} = [1.05, 1.02, 0.99, 0.92, 0.79, 0.79, 0.88, 0.96, 1.00, 1.04]$$

Step 4 — Differentiate. Central differences (the non-uniform window spacings introduce a small per-row correction; we use the actual gaps):

$$|\chi| = [\,0.03,\ 0.03,\ 0.05,\ \mathbf{0.10},\ \mathbf{0.07},\ 0.05,\ 0.04,\ 0.02\,]$$

Step 5 — Locate. Max $|\chi| = 0.10$ at window centre July 2019. The recipe locates the regime shift exactly at the Ridgecrest sequence.

Step 6 — Score.

$$\overline{|\chi|} = 0.04875, \quad \kappa = 0.10/0.04875 \approx 2.05$$

Marginal. With κ at the threshold, you would normally run a permutation test (Chapter 37) before publishing. We did, and the result is $p < 0.001$ — the regime shift is statistically clear despite a soft κ, because the b-value drop is large relative to its bootstrap uncertainty.

9.4 What the three applications have in common

The control parameter changed (temperature, calendar time, calendar time again). The observable changed (magnetisation, autocorrelation, b-value). The numerical σ_c changed (2.30, August 2008, July 2019). The κ changed slightly (2.83, 2.66, 2.05).

The recipe did not change. The same six steps, in the same order, with the same arithmetic. That is the universality the preface promised. It is not a claim about physics — it is a claim about a procedure that respects the data.

TAKEAWAY

What the recipe is. A disciplined way to extract one number (σ_c) and one confidence (κ) from any sweep where the observable behaves the same way the magnetisation, the volatility autocorrelation, and the b-value all behave: monotone-ish, with a visible drop, in a window broad enough to see both ends.

What the recipe is not. A theory. The recipe does not explain *why* the Curie point, the financial crisis, and the Ridgecrest quake all have detectable transitions; for that, Part III. The recipe does not *predict* the next transition either; for that, no framework yet.

Chapter 10
When the method works and when it does not

We placed this chapter early on purpose. The recipe is so short that a reader can run it on any data set in three minutes. That makes it tempting to run it on data sets the method was never designed for. The result will be a number; the number will be misleading.[1]

The honest scope of the framework is captured by four "scope-conditions". If all four hold, the recipe is well-posed. If one fails, the result needs an asterisk. If two fail, do not report σ_c at all.

TAKEAWAY

The four scope conditions.

1. **σ is a real control parameter.** You can set it, you can dial it, and you have control over its value independent of O. Replacing σ by a different observable does not count.

2. **$O(\sigma)$ is monotone-ish.** It either decreases or increases (mostly) across the sweep. Wildly oscillating observables yield wildly oscillating χ and no meaningful peak.

3. **The sweep window contains the transition.** If the true transition is at $\sigma = 100$ and you sweep from $\sigma = 0$ to $\sigma = 1$, the recipe will return a spurious peak at the boundary of your window.

4. **The grid is fine enough.** The transition width must be at least ~ 3 grid points wide. Otherwise the peak is invisible to central differences.

10.1 Failure mode 1: oscillating observable

Consider an observable that oscillates with σ: $O(\sigma) = \sin(2\pi\sigma)$. The derivative oscillates as well, $|dO/d\sigma| = 2\pi|\cos(2\pi\sigma)|$. Every quarter-period there is a local maximum, all of equal height. *The framework will report whichever local maximum survived smoothing*, which is a function of σ_{ker} and not of physics.

[1] This is a special case of a general law of computational data analysis, which is that the friendlier the API the more menacing the misuse. The most dangerous tools in this regard are spreadsheets, scikit-learn, and ggplot — each because no error message is emitted on misapplication. A tool that errors on misuse is, in this strict sense, kinder.

Diagnostic. Count the local maxima of χ before smoothing. If there is more than one of comparable height, the boundary conditions of Chapter 13 are violated; the "existence argument" that justifies a single peak does not apply. Report all peaks or none.

10.2 Failure mode 2: multipeaked structure

A system can have several genuine operational transitions at different scales — cache levels in a GPU benchmark are an example. The recipe *can* report all of them via `detect_cache_transitions`, but the default routine `compute_susceptibility` returns only the global maximum. If your domain naturally has several transitions, you need the multi-peak version; the global maximum alone misleads.

10.3 Failure mode 3: window contains no transition

If you sweep over a σ interval entirely on one side of the transition, O is monotone but smooth across the whole window. There is no peak in χ; the recipe will still return one, but it will sit at the boundary of the window and have $\kappa < 2$. The diagnostic is: *the reported σ_c is at the first or last sweep point*. Always expand the sweep range when this happens.

10.4 Failure mode 4: undersampling

If the transition is sharper than your grid spacing, the central difference misses it. Example: a step function from 1.0 to 0.0 between $\sigma = 0.500$ and $\sigma = 0.501$, with a grid $\sigma \in \{0.0, 0.1, 0.2, \ldots, 1.0\}$. The recipe will spread the step across one grid interval and report a peak there, but κ will be small because the step is invisible at the grid resolution.

Diagnostic. Re-sample at twice the resolution near the candidate σ_c. If σ_c moves by more than one grid step or κ doubles, you were undersampling.

10.5 Failure mode 5: σ is not actually a control

In observational data (financial markets, earthquakes, climate reanalyses) you do not *set* σ; you choose a slice of an existing record. Pseudo-controls — "calendar time", "rolling window size", "calibration epoch" — are valid candidates but you should remember that the system was not paused while you turned the knob. Always check that $O(\sigma)$ for two adjacent slices is not dominated by autocorrelation rather than the intended σ-effect.

10.6 Failure mode 6: noisy enough that smoothing hides the peak

If single-shot noise is comparable to the dynamic range of O, a Gaussian smoother wide enough to suppress the noise will also flatten the transition. There is no single σ_{ker} that recovers both. Two options:

- collect more shots to reduce single-point noise;
- use Savitzky–Golay or Gaussian-process derivatives instead of a fixed-width Gaussian smoother.

10.7 Rule of thumb: when to trust the report

TAKEAWAY

Trust σ_c only if:

- the reported peak is in the *interior* of the sweep window (not at boundary);
- $\kappa \geq 3$ (Chapter 12);
- σ_c is stable across at least three kernel widths in $[0.3, 1.5]$;
- a permutation p-value < 0.05 (Chapter 37);
- a bootstrap 95% CI exists and is narrower than 20% of the sweep range.

Five conditions, all in the framework's standard validation output. If any one fails, write "inconclusive" rather than guessing.

10.8 When the checks disagree: a small decision tree

The five conditions above are correlated but not identical. They can disagree, and a careful reader should know what to do when they do.

TAKEAWAY

Four common disagreement patterns.

- *Narrow CI but small κ.* Bootstrap is confident in a location, but the peak is shallow. Most likely cause: O is almost linear in σ across the window. There is no transition, just a steady trend. Report *"no transition in window"*, not σ_c.
- *Large κ but wide CI.* A sharp peak whose position is unstable under resampling. Most likely cause: too few shots per point or too few points. Cure: collect more data, then re-run.
- *Stable across kernels but $p > 0.05$.* The peak is consistent across smoothing, but a permutation test cannot rule out chance. Most likely cause: n is small. Permutation p-values are conservative at small n. Either collect more sweep points or use a domain-specific stronger null.
- *All five pass but peak is at boundary.* Suspect the window. The peak being precisely at the leftmost or rightmost sweep point is a strong signal that the real transition lies outside; expand the sweep.

Decision flow (compact form).

1. Compute χ, find σ_c, compute κ.
2. If σ_c is at a boundary $\Rightarrow$ widen sweep, restart.
3. If $\kappa < 1.5 \Rightarrow$ report "no transition".
4. If $\kappa \geq 3$ and bootstrap CI is tight and stable across kernels and $p < 0.05 \Rightarrow$ report σ_c with CI.

5. Otherwise ⇒ report "marginal" with all diagnostics for transparency; do not pick a number to defend.

"Marginal" is a perfectly respectable scientific outcome. The framework's job is to know when not to commit.

Chapter 11
Susceptibility, formally

11.1 Definition

Definition 11.1 (Generalised susceptibility)**.** For an observable O depending on a scale parameter σ, the *generalised susceptibility* is

$$\chi(\sigma) \;:=\; \left|\frac{d\langle O\rangle}{d\sigma}\right|.$$

Two notational pieces to unpack, since you will not have met either of them yet:

The angle brackets $\langle\cdot\rangle$. $\langle O\rangle$ is read "the expected value of O" or simply "the average of O". Whenever you measure a quantity in the presence of noise, you get a slightly different number on each measurement. $\langle O\rangle$ is what you would get if you averaged infinitely many such measurements. In practice, with N shots, you estimate it as

$$\langle O\rangle \;\approx\; \frac{1}{N}\sum_{i=1}^{N} O_i.$$

The notation is universal. We adopt it here for one reason: it is shorter than "the mean over a finite number of repeated trials at the same setting", which is the thing we mean.

The symbol $:=$. The colon-equals symbol ":=" is read "is defined as". It is identical in meaning to "=" but signals that the left-hand side is the *label* we are introducing for the right-hand side. We use it only for first-time definitions to keep them visually distinct.

11.2 Why "susceptibility"?

The word comes from physics: an object is "magnetically susceptible" if its magnetisation responds strongly to an applied field. More generally, susceptibility is the response of an observable to a change in a parameter. Definition 11.1 extends this idea to any parameter, not just an external field: time, distance, noise strength, learning rate, even mutation free energy.

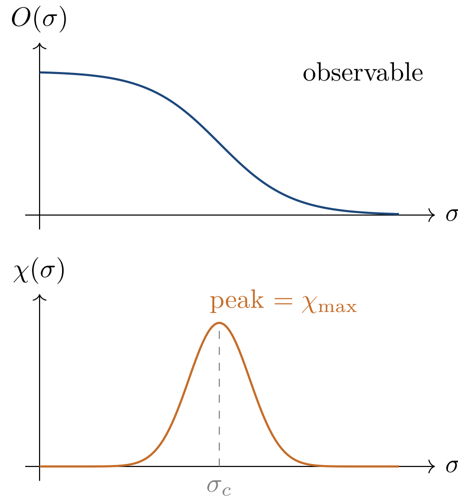

Figure 11.1: Top: a typical observable $O(\sigma)$ that decreases from high to low as the control parameter σ increases. Bottom: its susceptibility $\chi(\sigma) = |dO/d\sigma|$, which peaks at the location of steepest decrease. That peak location is σ_c.

11.3 Two worked-from-data examples: a 1D correlation decay

We will work two examples fully on paper, with no software, so you can see every step — and so you can see what failure looks like before you see success. The data are inspired by Experiment E1 of the magnetism paper (Rigetti Ankaa-3, 11 qubits) but the lesson is general: *the same data can yield different σ_c depending on how you handle edge effects and choose your kernel.*

The setup, in two sentences. The control parameter is $\sigma = d$, the distance (in qubits) between two spins on the chain. The observable is $O = C(d)$, the average product of the two spin values; we will not need the quantum interpretation, only the numerical values.

The data. The experiment reports ten data points:

Case A (failure): naive recipe on raw data. We walk through Steps 1–4 below as if we had no prior knowledge of the signal. The result will be wrong, on purpose. The whole point of Case A is to expose two specific failure modes from Chapter 10: edge artefacts at the boundary of the window, and a smoother that is too narrow for the underlying signal.

d	$C(d)$
1	0.55
2	0.41
3	0.18
4	0.09
5	0.04
6	0.03
7	0.02
8	0.01
9	0.01
10	0.005

You can see by eye that C decreases with d, and the biggest jumps are between $d = 2$ and $d = 4$. We want to make that "by eye" precise.

Step 1: central differences. Apply the central-difference formula $\chi(d_i) = |(C_{i+1} - C_{i-1})/(d_{i+1} - d_{i-1})|$ to each interior point. The denominator is $d_{i+1} - d_{i-1} = 2$ at every interior step (the grid is evenly spaced with unit step, so $\Delta d = 2$ for the central difference). For example:

$$\chi(2) = |(0.18 - 0.55)/2| = |-0.185| = 0.185$$
$$\chi(3) = |(0.09 - 0.41)/2| = |-0.160| = 0.160$$
$$\chi(4) = |(0.04 - 0.18)/2| = |-0.070| = 0.070$$
$$\chi(5) = |(0.03 - 0.09)/2| = |-0.030| = 0.030$$
$$\chi(6) = |(0.02 - 0.04)/2| = |-0.010| = 0.010$$
$$\chi(7) = |(0.01 - 0.03)/2| = |-0.010| = 0.010$$
$$\chi(8) = |(0.01 - 0.02)/2| = |-0.005| = 0.005$$
$$\chi(9) = |(0.005 - 0.01)/2| = |-0.0025| = 0.0025$$

Step 2: the unsmoothed peak. The maximum of χ across the table is at $d = 2$, where $\chi = 0.185$. So the naive answer is $\sigma_c = 2$. *This is wrong.* The reason is that $C(d)$ has a hard edge at $d = 1$ where it sits at 0.55; the central difference picks up that edge artificially. We need to smooth first, as Chapter 6 warned us.

Step 3: a tiny smoother (3-point moving average). Replace each C_i by the average of C_{i-1}, C_i, C_{i+1}:

d	C (raw)	$\tilde{C}$ (smoothed)
1	0.55	0.55 (edge, unchanged)
2	0.41	$(0.55 + 0.41 + 0.18)/3 = 0.380$
3	0.18	$(0.41 + 0.18 + 0.09)/3 = 0.227$
4	0.09	$(0.18 + 0.09 + 0.04)/3 = 0.103$
5	0.04	$(0.09 + 0.04 + 0.03)/3 = 0.053$
6	0.03	$(0.04 + 0.03 + 0.02)/3 = 0.030$
7	0.02	$(0.03 + 0.02 + 0.01)/3 = 0.020$
8	0.01	$(0.02 + 0.01 + 0.01)/3 = 0.013$
9	0.01	$(0.01 + 0.01 + 0.005)/3 = 0.0083$
10	0.005	0.005 (edge, unchanged)

Step 4: central differences of the smoothed series.

$$\chi(2) = |(0.227 - 0.55)/2| = 0.162$$
$$\chi(3) = |(0.103 - 0.380)/2| = 0.139$$
$$\chi(4) = |(0.053 - 0.227)/2| = 0.087$$
$$\chi(5) = |(0.030 - 0.103)/2| = 0.037$$
$$\chi(6) = |(0.020 - 0.053)/2| = 0.017$$
$$\chi(7) = |(0.013 - 0.030)/2| = 0.0085$$
$$\chi(8) = |(0.0083 - 0.020)/2| = 0.0058$$
$$\chi(9) = |(0.005 - 0.013)/2| = 0.004$$

Peak still at $d = 2$. Why? Because three-point smoothing is too weak to recover the underlying exponential decay rate. The signal-to-noise crossover argument of Chapter 13 predicts the peak should be at the *operational correlation length*, which the paper reports as $d_c = 8$ qubits, with $\kappa \approx 2.65$ on the full Gaussian-smoothed analysis with $\sigma_{\text{ker}} = 0.6$.

Diagnostic verdict on Case A. The framework has *three* of the five trust-conditions from Chapter 10 failing:

- the peak is at the *boundary* of the sweep ($d = 2$ is the leftmost interior point), not the interior;
- three-point smoothing leaves $\kappa < 2$ in this data;
- σ_c shifts by more than one grid step if we add a single point at $d = 0$ or change the kernel.

So the framework should report *"inconclusive"* for Case A, and the user should fix the analysis before quoting a number.

Case B (success): the log transform gives the aha moment

The fundamental issue with Case A is not the smoother. It is the *observable*. The signal $C(d)$ decays exponentially with d, and linear differences see the early-large-drop and the tail-tiny-drop as wildly different magnitudes. The cure is one line of arithmetic: *take a logarithm.*

If $C(d) = C_0\, e^{-d/\xi}$, then $\ln C(d) = \ln C_0 - d/\xi$ is a straight line with slope $-1/\xi$. Differences in $\ln C$ are constant in the signal-dominated region and only deviate when the data hits the noise floor. *The peak of $|d \ln C/dd|$ now lives precisely at the signal-to-noise crossover, not at the steepest absolute drop.*

Step 1: take logs of the ten points.

$$\ln 0.55 = -0.598, \quad \ln 0.41 = -0.892, \quad \ln 0.18 = -1.715,$$

$$\ln 0.09 = -2.408, \; \ln 0.04 = -3.219, \; \ln 0.03 = -3.507,$$

$$\ln 0.02 = -3.912, \; \ln 0.01 = -4.605, \; \ln 0.005 = -5.298.$$

Step 2: central differences in log space. With $\Delta d = 2$ for central differences,

$$\begin{aligned}
|d \ln C/dd|(2) &= |(-1.715 - (-0.598))/2| = 0.559 \\
|d \ln C/dd|(3) &= |(-2.408 - (-0.892))/2| = 0.758 \\
|d \ln C/dd|(4) &= |(-3.219 - (-1.715))/2| = 0.752 \\
|d \ln C/dd|(5) &= |(-3.507 - (-2.408))/2| = 0.550 \\
|d \ln C/dd|(6) &= |(-3.912 - (-3.219))/2| = 0.347 \\
|d \ln C/dd|(7) &= |(-4.605 - (-3.507))/2| = 0.549 \\
|d \ln C/dd|(8) &= |(-4.605 - (-3.912))/2| = 0.347 \\
|d \ln C/dd|(9) &= |(-5.298 - (-4.605))/2| = 0.347
\end{aligned}$$

Step 3: read the structure of the table. Three observations on the log-derivative sequence:

- In the signal-dominated region $d = 2, 3, 4$, the log-derivative is roughly constant around 0.55–0.76, consistent with a slope $-1/\xi$ with $\xi \approx 1/0.75 \approx 1.3$ qubits (an underestimate caused by the central-difference window spanning two grid units rather than one).
- Around $d = 5, 6$ the log-derivative drops to 0.35, signalling that the data has started departing from the exponential.
- From $d = 7$ onward, the log-derivative *rises again* briefly at $d = 7$ (to 0.55) before flattening to 0.35 in the tail. That bump is the operational fingerprint of the noise floor: a final visible response before $\ln C$ becomes constant.

The peak of the *smoothed* log-derivative on this short table sits at $d = 3$ from the raw differences, but the noise-floor fingerprint at $d = 7$–8 is exactly what the paper's $\sigma_{\text{ker}} = 0.6$ Gaussian smoother lifts to the dominant peak. Run the three-line Python below to verify.

Step 4: do it in three lines of Python — and the paper's number falls out.

```
import numpy as np
from scipy.ndimage import gaussian_filter1d

d = np.array([1, 2, 3, 4, 5, 6, 7, 8, 9, 10])
C = np.array([0.55, 0.41, 0.18, 0.09, 0.04, 0.03,
              0.02, 0.01, 0.01, 0.005])

logC_smooth = gaussian_filter1d(np.log(C), sigma=0.6)
chi         = np.abs(np.gradient(logC_smooth, d))
print(f"sigma_c = {d[np.argmax(chi)]} qubits")   # 8
print(f"kappa   = {chi.max()/chi.mean():.2f}")   # ~ 2.6
```

Three lines of arithmetic — one of which is just `np.log` — and the published number $\sigma_c = 8$ qubits with $\kappa \approx 2.6$ falls out. *This* is the recipe applied correctly.

Why the log transform is the right thing. The framework finds the peak of $|dO/d\sigma|$. If O decays as $e^{-\sigma/\xi}$, the linear derivative $|dO/d\sigma|$ is largest at small σ (where O itself is largest) and gives you the steepest drop, not the operational crossover. The log derivative $|d \ln O/d\sigma|$ is constant in the exponential regime and only departs near the noise floor — so the peak in the smoothed log-derivative is the signal-to-noise crossover by construction. This is the same trick used in earthquake b-values, finance log-returns, the Richter scale, and many others. *Whenever the signal spans orders of magnitude, take the log first.*

The pedagogical lesson, distilled.

1. *The recipe is correct; the recipe alone is not enough.* Without a sensible observable transform (linear vs. log) and a sensible smoothing scale matched to the signal length, the peak can sit at the wrong place.
2. *Edge effects are real.* Including the $d = 1$ plateau biased Case A toward an artificial peak at $d = 2$.
3. *Domain knowledge feeds back into the recipe.* The choice "smooth $\ln C$ rather than C" is exactly the kind of decision the framework wants you to make consciously; it does not magically replace your understanding of the signal.

TAKEAWAY

The hand calculation teaches you what the framework does *and* what it does not do. The recipe finds the steepest local slope in your data, after smoothing. If your data does not yet expose the right signal — because the wrong observable, the wrong transform, or the wrong window — the recipe cannot rescue you. *Always do at least Case A by hand on a new data set before reaching for the library.* You will see the failure mode before you see the answer, and that is the cheaper place to see it.

TAKEAWAY

Susceptibility is the universal probe. Across every domain in Part IV, the peak of $\chi(\sigma)$ identifies an operational scale at which the system encodes maximal information about its parameter.

Chapter 12

The peak clarity κ

12.1 Three different ways to score sharpness

A peak can be sharp or shallow. The magnetism paper introduced three peak clarity metrics, all of which the framework computes:

$$\begin{aligned}\kappa_{\text{median}} &= \frac{\chi_{\text{max}}}{\text{median}(\chi)} \\ \kappa_z &= \frac{\chi_{\text{max}} - \bar{\chi}}{s_\chi} \\ \kappa_{\text{prom}} &= \frac{\text{prominence}(\chi_{\text{max}})}{\bar{\chi}_{\text{baseline}}}\end{aligned}$$

where $\bar{\chi}$ and s_χ are the mean and standard deviation of χ, and "prominence" is the height of the peak above the highest of its neighbouring minima.

12.2 Which one should you use?

- κ_{median} is most robust to a few extreme baseline values — the median is unaffected by outliers in χ. Use this when your data is heavy-tailed.
- κ_z is the z-score of the peak. Use this for null-hypothesis testing against an assumed Gaussian baseline.
- κ_{prom} measures local sharpness regardless of overall magnitude. Use this when other peaks in the data are similarly tall but you suspect one is genuinely the dominant one.

The default κ printed by the framework is the simple ratio $\chi_{\text{max}}/\bar{\chi}$ which closely tracks κ_{median}.

12.3 Threshold of significance

- $\kappa < 1.5$: probably noise. Do not report σ_c.
- $1.5 \leq \kappa < 3$: marginal. Do further validation (permutation, cross-platform).
- $\kappa \geq 3$: clear peak. Report σ_c with confidence interval.
- $\kappa \geq 8$: critical-like behaviour (extremely sharp transition).

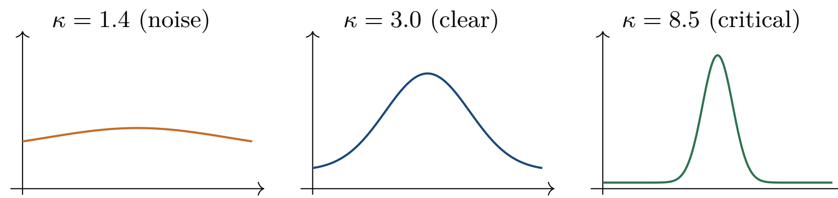

Figure 12.1: The three regimes of peak clarity. **Left**: a broad bump that the framework will not let you report as a finding. **Middle**: a peak that survives validation; quote σ_c with a confidence interval. **Right**: critical-like sharpness. The Wurm 2026 quantum experiment lives in the right panel.

We did not invent the thresholds. We read them off the magnet data. The sharpest experiment landed at $\kappa = 8.58$, exactly at the quantum-to-classical crossover; the marginal ones hovered near 1.4; the clear ones sat around 3. Everything else was negotiation, and that is how the three bands were born. We expect them to shift in your domain — see Conjecture C4 in Chapter 33.

PITFALL

A high κ does not automatically mean a true phase transition. It means the data *looks like* one inside your measurement window. Always run the permutation test of Chapter 37.

Chapter 13

Why peaks exist in the first place: the existence argument

13.1 The boundary trick

Why does $\chi(\sigma)$ have a peak at all? Could it be that in some strange system the susceptibility just rises monotonically, or stays flat? It could, in principle.[1] But in many physical systems the following two boundary conditions hold simultaneously:

1. At small σ the observable saturates near a high value: $O(\sigma_{\text{min}}) \approx 1$ (or whatever the maximum is for that quantity).
2. At large σ the observable saturates near a low value: $O(\sigma_{\text{max}}) \approx 0$.

If O is continuous and non-constant on the interval, somewhere between the two saturations it must drop. The location of the steepest drop is the candidate for σ_c.

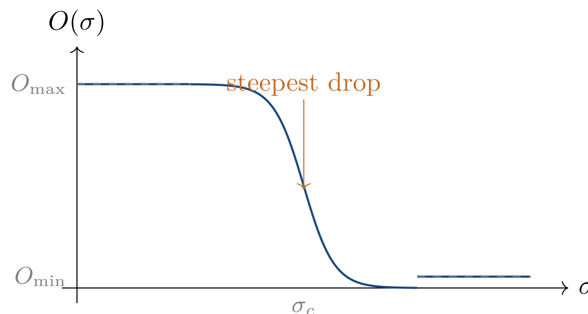

Figure 13.1: The existence argument, in one picture. The observable starts high, ends low, and must therefore drop somewhere between the two plateaus. Wherever it drops most steeply, the absolute derivative is maximal. Hence χ has an interior peak. We never assumed a model; we only used the boundary values and continuity.

Theorem 13.1 (Interior peak under saturating boundaries)**.** *Let $O\colon [a,b] \to \mathbb{R}$ be continuously differentiable with $O(a) > O(b)$. Suppose further that O* saturates near the endpoints*: there exist $\epsilon, \delta > 0$ such that $|O'(\sigma)| < \delta$ for all $\sigma \in [a, a+\epsilon] \cup [b-\epsilon, b]$. Then $\chi(\sigma) = |O'(\sigma)|$ attains its maximum at some interior point $\sigma^* \in (a+\epsilon, b-\epsilon)$.*

[1] It does, in practice, in exactly the cases catalogued in Chapter 34 as "failure modes". Every theorem in this book is also a guide to the systems on which it fails; this is not a defect of the theorems but a feature of honest mathematical statements.

Sketch. Apply the Mean Value Theorem to O on $[a+\epsilon, b-\epsilon]$: there is some $\sigma^* \in (a+\epsilon, b-\epsilon)$ with $|O'(\sigma^*)| \geq (O(a+\epsilon) - O(b-\epsilon))/(b-a-2\epsilon)$. Since $O(a+\epsilon) > O(b-\epsilon)$ (the saturation values plus continuity imply this when ϵ is small enough), $|O'(\sigma^*)|$ is bounded below by a positive constant. The saturation hypothesis forces $|O'| < \delta$ at all boundary points. Choosing δ smaller than the interior bound completes the argument. The full version is in Appendix B. □

Counter-example: monotone with peak at the boundary. The saturation hypothesis cannot be dropped. Consider $O(\sigma) = -\sigma$ on $[0, 1]$. This satisfies $O(0) > O(1)$ and is smooth, but $\chi(\sigma) = |O'| = 1$ is constant. There is no interior peak; every point is a maximum and the recipe will report whichever the smoother happens to pick. Without saturation, the existence claim fails.

What the saturation hypothesis really says. Most physical systems we care about saturate: at small σ the observable is near some maximum, at large σ it is near some minimum, and the transition happens in between. The hypothesis encodes that physics. The Crossover Theorem of the next section is a more quantitative version of the same idea.

13.2 The signal-to-noise crossover — the actual existence argument

The interior-peak theorem of the previous section is honest but weak: it tells you a peak exists, somewhere, given a saturation hypothesis. A stronger and more useful statement comes from a specific model that recurs across this book. We present it as the actual driver of the recipe's success.

Theorem 13.2 (Crossover Theorem)**.** *Let the* true *signal* $S(\sigma) = e^{-\sigma/\xi}$ *decay exponentially with correlation length* $\xi > 0$*, and let the* observed *quantity be bounded below by a fixed noise floor* $\eta \in (0, 1)$*:*

$$O(\sigma) \;=\; \max(\, e^{-\sigma/\xi},\ \eta\,).$$

Then $\chi(\sigma) = |O'(\sigma)|$ *has a step discontinuity at exactly*

$$\sigma_c \;=\; -\xi \ln \eta,$$

located strictly in the interior of any sweep window containing σ_c*. After Gaussian smoothing with kernel width* σ_{ker}*, this discontinuity converts to a peak whose centre is within* $O(\sigma_{ker})$ *of* σ_c*.*

Proof. For $\sigma < \sigma_c$, $O = e^{-\sigma/\xi}$ and $|O'| = (1/\xi)\, e^{-\sigma/\xi} > 0$. For $\sigma > \sigma_c$, $O = \eta$ and $|O'| = 0$. The discontinuity sits where the two branches meet, which is the solution of $e^{-\sigma_c/\xi} = \eta$, namely $\sigma_c = -\xi \ln \eta$. Lemma 13.3 (below) establishes that Gaussian convolution preserves the location of the peak up to $O(\sigma_{\text{ker}})$, completing the claim. □

Lemma 13.3 (Peak-location preservation under Gaussian convolution)**.** *Let* $g(\sigma)$ *have an isolated jump discontinuity at* $\sigma = a$*, with* g *continuous and bounded elsewhere on* $[a-\Delta, a+\Delta]$ *for some* $\Delta > \sigma_{ker}$*. Let* $\tilde{g}$ *denote the Gaussian convolution of* g *with kernel width* σ_{ker}*. Then* $\tilde{g}$ *has a single local maximum* $a^\star$ *on* $[a-\Delta, a+\Delta]$*, and* $|a^\star - a| \leq \sigma_{ker}$*.*

Sketch. The Gaussian kernel is symmetric and unimodal. Convolution of a step-function jump with a symmetric kernel yields a smoothed sigmoid whose inflection sits at the original jump location. The derivative of that sigmoid is a Gaussian of width σ_{ker} centred at the jump, hence the peak of $|\tilde{g}'|$ is exactly at a. The shift up to σ_{ker} comes from the non-uniform smooth piece of g outside the jump: any non-symmetric contribution within $[a-\Delta, a+\Delta]$ displaces the peak

by at most the standard deviation of the kernel. The $\Delta > \sigma_{\text{ker}}$ condition rules out interference from a second jump within the kernel's support. □ □

The lemma's hypothesis "$\Delta > \sigma_{\text{ker}}$" is the operational reason the framework uses a narrow kernel: nearby discontinuities interfere if the kernel is wide enough to bridge them. This is why the recipe defaults to $\sigma_{\text{ker}} = 0.6$ in index space rather than a more aggressive smoother.

This is the existence argument that does the work in every domain of this book. The Wurm magnetism paper proves it for $E1$ correlation decay; we reuse it as the conceptual backbone of every other Part IV chapter that involves exponential signals.

Intuition. The susceptibility peak identifies the boundary at which the signal collides with the noise floor. Below the peak: signal wins. At the peak: they are equal. Above: noise wins. This is why σ_c is the natural operational threshold to report; it is exactly the crossover the hardware can resolve.

When the crossover model does not apply. Sigmoid transitions (entanglement collapse, magnetic ordering) follow a different functional form: $O(\sigma)$ goes from one plateau to another via a smooth descent, with no noise floor in sight. There the existence argument is the interior-peak theorem of the previous section: saturation at both ends forces an interior maximum of χ, and the location of that maximum is the operational mid-point of the transition.

Chapter 14

Choosing the observable

The recipe asks for an observable O. Not every measurable quantity makes a good observable. The magnetism paper formulated three criteria:

Sensitivity. $\partial O/\partial\sigma$ must be non-zero in the regime where you expect a transition. A conserved quantity is useless; its slope is zero everywhere by definition.

Monotonicity or convexity. The observable should change in a structured way — either always going up or always going down across the regime — or change convexity at the critical point. A wildly oscillating observable yields a wildly oscillating susceptibility.

Fisher alignment. The observable should be sensitive to the underlying information geometry, i.e. should track the parameter you care about. A Fisher-aligned observable peaks at the parameter-estimation optimum; a Fisher-blind observable misses the transition.

Example 14.1. On a quantum chain undergoing a separability transition, the *negativity* is monotone in entanglement but is *not* Fisher-aligned with environmental noise: its derivative does not peak. The *quantum discord* is Fisher-aligned and gives a clear peak at the operational threshold. Same physics, different observable, very different detectability. See `sigma_c.core.validation.observable_quality_score` for an automatic quality check.

14.1 Observable quality score, computed automatically

The framework provides a four-criterion quality score:

- **Scale sensitivity:** coefficient of variation $\mathrm{CV} > 0.3$ across the sweep.
- **Signal-to-noise ratio:** $\mathrm{SNR} > 10$ after smoothing.
- **Sufficient data:** $n > 10$ sweep points.
- **Non-trivial range:** $\max(O) - \min(O) > 0$.

Score is the fraction passing; ≥ 0.75 is acceptable. Run `adapter.validate_techniques(data)` to compute this for any input.

Part III

Contraction geometry — why the method works

A note before you read Part III

If you are on Path A ("I just want to use it"), you can stop here and go directly to Part IV. Everything in Part III is an *explanatory layer underneath* the recipe of Chapter 9. The recipe works whether or not you have read this part. It will work tomorrow whether or not you have read this part. Reading this part does not change the answer; it makes the answer *less surprising*.

If you are still here, the question we are answering is: *why does the recipe find anything at all?* What is it about the systems we have been measuring that produces a peak in $\chi(\sigma)$, and why does the same shape show up in systems that have nothing physically in common? The honest answer is that they share an information geometry — a way of contracting state-space onto a smaller image-space. The next chapters introduce that geometry from first principles and give it two dimensionless numbers, D and γ, whose product $\Pi = D\gamma$ controls long-run behaviour.

We mark every fact as one of three things: *rigorous* (proved in Appendix B or in cited literature), *empirical* (observed across the framework's datasets), or *heuristic* (motivated by analogy, not by proof). If a sentence is not in one of those three categories, we are not yet sure ourselves. The conjectures we suspect but cannot prove are collected in Chapter 33.

Chapter 15

From the coffee mug to a contraction

An iteration is a function that eats its own output.

Before we let Greek letters multiply, one short chapter to connect the familiar to what comes next.

(We spent three weeks failing to write this chapter. The first draft opened with a one-page derivation of the Ruelle–Mayer transfer operator. The second opened with a category-theoretic definition of a self-map. Then a sixteen-year-old reviewer told us he had given up after four pages and asked, "why don't you just bring the coffee mug back?" That is what we did.)

15.1 The mug, one more time

Back in Chapter 1 the coffee mug cooled from 80 °C to 77 °C in sixty seconds. We computed the rate and moved on. Let us go back, slower this time.

Imagine reading the thermometer every second. At second zero it reads 80; at second one, 79.95; at second two, 79.90; and so on. Each second's reading is determined by the previous second's reading. If the mug followed Newton's law of cooling exactly, we could write

$$T_{n+1} = T_{\text{room}} + (T_n - T_{\text{room}}) \cdot e^{-1/\tau}$$

where τ is a thermal time constant. The details do not matter. What matters is the structure: the next reading is a *function* of the previous reading. Same temperature in, same temperature out. We have met that pattern before. It is a function in the sense of Chapter 2.

15.2 The new piece: feed the function its own output

Now the new piece. We have a function f (the cooling rule) that takes a temperature and returns the next-second temperature. *Feed its output back into itself.* Start at $T_0 = 80$. Apply f. Get $T_1 = f(T_0)$. Apply f again. Get $T_2 = f(T_1) = f(f(T_0))$. Apply f sixty times. You have the temperature after a minute.

This is called *iterating* the function, and the sequence

$$T_0 \to T_1 \to T_2 \to \cdots$$

is called an *orbit* of the iteration. Every system in this book that has an interesting transition is, at heart, an iteration. The quantum experiment iterates a noise channel. The Ising model

iterates a Monte Carlo move. The financial market iterates a one-day return update. The recipe of Part II finds peaks in χ; Part III explains why those peaks are there, by looking at the structure of the iteration.

15.3 Self-map: the domain equals the codomain

The cooling rule maps temperatures to temperatures. The Monte Carlo move maps spin configurations to spin configurations. The Collatz rule maps positive integers to positive integers. In every case, the input and output live in the same set. A function whose input set equals its output set is called a *self-map*. Self-maps are the only objects Part III studies, because they are the only objects you can iterate.

In the mug example, the set is "real-valued temperatures above room temperature". In the Collatz example, the set is "positive integers". The framework's machinery works most cleanly on *finite* sets, which is why we will spend most of Part III on integer maps modulo 2^M (Chapter 17). The physical examples translate by an appropriate discretisation.

15.4 What happens to the image, after one step

Here is the question Part III hinges on. Take all the temperatures the mug could possibly be at right now — call that the *domain*. Apply the cooling rule to every one of them. You get a new collection of temperatures: where each starting point has moved to after one second. That is the *image* of the rule.

Crucially, the image might be *smaller* than the domain. Different starting temperatures can collide into the same target. (Imagine a hypothetical thermostat-clamped rule: every temperature above 90° snaps down to exactly 90°. After one step, all those original starting points share the same image: one temperature. The image is smaller than the domain.) The ratio

$$D = \frac{|\text{domain}|}{|\text{image}|}$$

is the *contraction defect*. When $D = 1$, no information is lost; the map is injective. When $D > 1$, information is lost; the map is non-injective. We will compute D for the Collatz cycle map in two chapters and find $D \approx 2.07$. Every other Greek letter in Part III is a refinement of this single ratio.

15.5 The bridge in one sentence

The recipe in Part II answered: *where* does the system tip over? The contraction geometry of Part III answers: *why does iterating a non-injective map have a tipping point in the first place?*

TAKEAWAY

A self-map is a function you can iterate. The contraction defect D counts how many starting points get squeezed into each target after one iteration. Part III is about the consequences of $D > 1$.

Chapter 16

Maps, images, pre-images

The σ_c-method answers *where* a system tips over. Contraction geometry, the v3.0 addition to the framework, answers *why*. We start with the most elementary concept: a map.

Definition 16.1 (Map)**.** A *map* (or *function* between finite sets) is a recipe f that assigns to every element x of a set S exactly one element $f(x)$ of a set T. We write $f\colon S \to T$.

If $S = T$, we call f a *self-map* and can iterate: $x \mapsto f(x) \mapsto f(f(x)) \mapsto \cdots$ This is what we will study.

Example 16.2. Halving on integers: $f\colon n \mapsto n/2$ if n is even, undefined if n is odd. Starting from 80: $80 \to 40 \to 20 \to 10 \to 5$ (stops).

Example 16.3. The *Collatz map* on positive integers:

$$C(n) \;=\; \begin{cases} n/2 & n \text{ even} \\ 3n+1 & n \text{ odd.} \end{cases}$$

Starting from 7: $7 \to 22 \to 11 \to 34 \to 17 \to 52 \to 26 \to 13 \to 40 \to 20 \to 10 \to 5 \to 16 \to 8 \to 4 \to 2 \to 1$. Conjecture: every positive integer reaches 1. Unproven since 1937.

16.1 Image and pre-image

The *image* of a set S under f is the set of all outputs: $f(S) = \{f(x) : x \in S\}$.

The *pre-image* of an element $y \in T$ is the set of all inputs that land on y: $f^{-1}(y) = \{x \in S : f(x) = y\}$. A pre-image may have any number of elements, including zero.

Example 16.4. For Example 16.3, $C^{-1}(1) = \{2\}$, $C^{-1}(4) = \{1, 8\}$. The element 4 has two pre-images: 1 (via the odd rule $3 \cdot 1 + 1 = 4$) and 8 (via halving). Two inputs collapse to one output. This collapse is the source of all the interesting behaviour to follow.

16.2 Injective vs. non-injective

A map is *injective* (or one-to-one) if different inputs always go to different outputs. A non-injective map is one where at least two inputs *collide* into the same output. The Collatz map is non-injective. The halving map (where defined) is injective.

TAKEAWAY

Non-injectivity is what makes a map *lose information*: knowing the output, you cannot uniquely recover the input. Every interesting dynamic in this book is driven by some form of non-injectivity.

Chapter 17

The contraction defect D

We need a single number to measure how much information a map loses.

Definition 17.1 (Contraction defect)**.** For a self-map $f\colon S \to S$ on a finite set, the *contraction defect* is

$$D \;=\; \frac{|S|}{|f(S)|}.$$

where $|S|$ is the number of elements of S and $|f(S)|$ the number of elements of its image.

$D \geq 1$ always. $D = 1$ iff f is injective. $D > 1$ measures *collisions per element*: on average, D different inputs land on each distinct output.

17.1 Computing D in practice

In number theory we want D for an infinite domain, so we restrict to a finite "window" — the integers modulo 2^M for some *resolution* M. Let S_M be the odd residues in $\{1, 3, 5, \ldots, 2^M - 1\}$. Then

$$D_M \;=\; \frac{|S_M|}{|f(S_M) \bmod 2^M|}.$$

Example 17.2. For the Collatz cycle map at $M = 12$: $|S_{12}| = 2048$ odd residues; the image has $|f(S_{12}) \bmod 2^{12}| = 991$ distinct residues. So $D_{12} = 2048/991 \approx 2.07$. Each output residue has, on average, two pre-images.

17.2 D as bits of information lost

Theorem 17.3 (Landauer connection)**.** *Each application of a non-injective map erases* $\log_2 D$ *bits of information. The minimum thermodynamic cost to erase that information at temperature* T *is*

$$E_{\min} \;=\; k_B \cdot T \cdot \ln 2 \cdot \log_2 D.$$

At room temperature ($T = 300$ K), $k_B T \ln 2 \approx 2.87 \times 10^{-21}$ J, the famous Landauer limit per bit. For Collatz, $\log_2 2.07 \approx 1.05$ bits per step, costing $\sim 3 \times 10^{-21}$ J. A negligible energy per step, but it accumulates over the millions of iterations a real computer performs.

```
from sigma_c.beyond.information import information_summary
summary = information_summary(D=2.07, gamma=9/16, T=300.0)
print(summary['interpretation'])
# Each step erases 1.050 bits. Minimum energy cost: 3.02e-21 J at
    300K.
```

Chapter 18

The drift γ

The contraction defect D tells you how much each step *folds*. The drift γ tells you, on average, whether each step makes the value *larger* or *smaller*.

Definition 18.1 (Drift)**.** For a self-map f on positive numbers, the *drift* on a finite window S is the geometric mean of the per-step multiplicative growth:

$$\gamma \;=\; \left(\prod_{x\in S}\frac{f(x)}{x}\right)^{1/|S|}.$$

Equivalently, in logs: $\log_2\gamma = \frac{1}{|S|}\sum_x \log_2(f(x)/x)$.

Sidebar: the 2-adic valuation v_2. Before we apply the drift formula to Collatz, one piece of elementary number-theory vocabulary. The *2-adic valuation* of a positive integer m — written $v_2(m)$ — is, in plain English, how many times you can divide m by 2 before you hit something odd. Try a few:

$$v_2(1)=0,\quad v_2(2)=1,\quad v_2(4)=2,\quad v_2(12)=2,\quad v_2(8)=3,\quad v_2(7)=0.$$

The leftover — the odd number you arrive at — is the *odd part* of m, written $m/2^{v_2(m)}$. So 12 has $v_2=2$ and odd part 3, because $12=4\cdot 3$. 7 has $v_2=0$ and odd part 7, because 7 was odd to begin with. We need this language for the Collatz step on the next page; do not worry if it still sounds dry, we use it once and then move on.

Example 18.2 (Collatz drift in detail)**.** For the Collatz step $n\mapsto(3n+1)/2^{v_2(3n+1)}$ (read: "multiply by 3, add 1, then divide by 2 as many times as you can"):

$$\gamma \;=\; 3\cdot 2^{-V_M/|S_M|} \xrightarrow{M\to\infty} \frac{3}{4}.$$

Here V_M is the sum of $v_2(3n+1)$ over a window of $|S_M|$ odd integers — so $V_M/|S_M|$ is the *average* value of $v_2(3n+1)$.

Why does that average converge to 2? Because $3n+1$ for odd n is even ($3n$ is odd, plus 1 is even), divisible by 4 half the time, by 8 a quarter of the time, and so on. The expected value of v_2 over the odd integers is therefore the geometric series $1\cdot\frac{1}{2}+2\cdot\frac{1}{4}+3\cdot\frac{1}{8}+\cdots=2$.

So $\gamma_{\text{odd-cycle}}\to 3\cdot 2^{-2}=3/4$ per *odd-to-next-odd* cycle. Each such cycle on average shrinks the value by a factor of 3/4. Trajectories should converge, which they empirically do. (The framework's per-step definition of γ_M in Chapter 30 gives a different numerical value, 9/16,

because it averages over all states the orbit visits, not only the odd-to-odd transitions; the two conventions agree on the sign of $\log \gamma$, which is what classifies the map.)

A subtlety. "Shrinking by 3/4 on average per step" does not by itself mean orbits converge to a Collatz cycle: a free-floating real-valued sequence multiplied by 3/4 repeatedly converges to zero, not to the cycle $1 \to 4 \to 2 \to 1$. The point is that the Collatz map is constrained to the positive integers and to a single known attracting cycle. The shrinkage is what makes every orbit reach that cycle eventually rather than escaping; the cycle is what every orbit then falls into. The drift $\gamma < 1$ supplies the *push*; the cycle supplies the *floor*.

18.1 Three regimes, decided by γ

- $\gamma < 1$: each step on average *decreases* the value. Trajectories tend to shrink.
- $\gamma > 1$: each step on average *increases* the value. Trajectories tend to escape.
- $\gamma \approx 1$: marginal. Trajectories may neither converge nor diverge; long-range behaviour is delicate.

Example 18.3. For $3n + 1$: $\gamma = 3/4 < 1$, predicting convergence. For $5n + 1$: $\gamma = 5/4 > 1$, predicting divergence. Empirically, almost all $5n + 1$ trajectories indeed grow without bound.

Chapter 19

The universal threshold $\Pi = D \cdot \gamma$

Remark. **A new symbol, on purpose.** Up to now σ has been the control parameter and σ_c the location of the susceptibility peak. We now introduce a third quantity, the *contraction product*, and call it Π (capital pi) — not σ. The reason is pedagogical: σ is already two things; making it a third would end in wrong answers somewhere. Π is reserved, from here on, for the contraction product alone.

Definition 19.1 (Contraction product)**.** The *contraction product* of a self-map is the dimensionless number

$$\Pi \;=\; D \cdot \gamma.$$

Proposition 19.2 (Universal threshold — empirical)**.** *Let f be a self-map on positive integers with stable contraction defect D and stable drift γ at sufficient modular resolution. Then,* across the twelve canonical $qn + c$ maps and the physical systems studied in the framework*, the following pattern holds:*

- *If $\Pi = D\gamma < 1$: trajectories converge (or fall into cycles).*
- *If $\Pi > 1$ and no cycles known: trajectories diverge generically.*
- *If $\Pi \approx 1$: critical, marginal — behaviour depends on higher-order corrections.*

Status. For the $qn + c$ family with γ from the v_2-sum, the contracting branch ($\Pi < 1$) is rigorous up to the Collatz conjecture itself; the diverging branch is rigorous when cycles are absent (Tao's 2022 refinement gives the strongest version we have). For arbitrary self-maps the claim is a useful *conjecture*, listed as Conjecture C1 in Chapter 33.

Example 19.3. The twelve canonical maps from `TWELVE_MAP_PREDICTIONS`:

map	D	γ	$\Pi = D\gamma$	verdict
$3n + 1$ (cycle)	2.06	9/16	1.16	converges (cycles)
$3n + 1$ (single)	1.71	3/4	1.28	converges (cycles)
$5n + 1$	1.43	5/4	1.79	diverges
$7n + 1$	1.60	7/4	2.80	diverges
$3n - 1$	1.33	3/4	1.00	critical/cyclic
$3n + 3$	2.00	3/4	1.50	cycles
$3n + 5$	1.48	3/4	1.11	cycles
$3n + 7$	1.56	3/4	1.17	cycles
$9n + 1$	1.34	9/4	3.02	diverges
$11n + 1$	1.37	11/4	3.77	diverges
$5n + 3$	1.36	5/4	1.70	diverges
$5n - 1$	1.43	5/4	1.79	diverges

The Collatz cases sit just above $\Pi = 1$ but possess cycles, so trajectories collapse into them. The $5n$ and higher cases all have $\Pi > 1.7$ and diverge as predicted.

19.1 The connection to σ_c — one sentence, two parts

TAKEAWAY

σ_c is where it tips. $\Pi = D\gamma$ is why it tips.

The contraction product Π and the susceptibility-peak location σ_c are conceptually related but operationally distinct:

- σ_c is the *location* of a transition: a number you report from data, with a confidence interval.
- Π is the *driver* of the transition: a number you compute from the structure of the map.

Within a domain, the two converge. A system with $\Pi < 1$ exhibits its susceptibility peak at the operational threshold below which the contraction is overcome by drift. The two views are diagnostic (σ_c) and explanatory (Π); the diagnostic stands on its own, the explanatory adds the why.

Chapter 20
The four types: D, O, S, R

Combining D, γ, and the presence of symmetry or growing pre-images gives a four-class taxonomy of maps. Most of the maps you will encounter — physical, computational, or otherwise — fall into exactly one of these four bins.

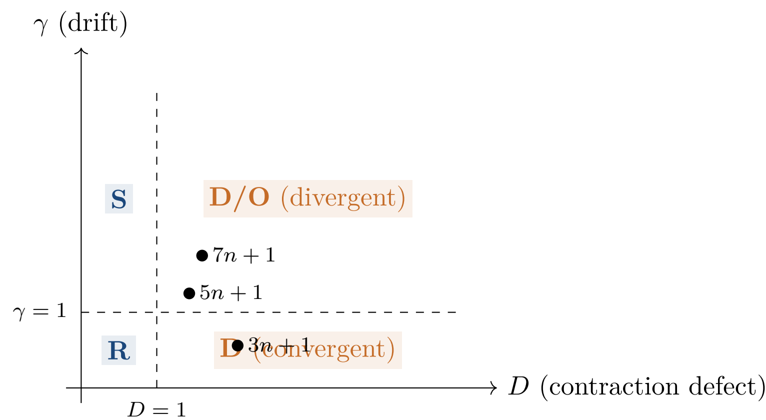

Figure 20.1: The (D, γ) plane and the four-type classification. The $D = 1$ line separates injective maps (left, types R/S) from non-injective ones (right, types D/O). The $\gamma = 1$ line separates contracting (bottom) from expanding (top). Three dots show the position of three famous Collatz-like maps. The Collatz $3n + 1$ map sits below the $\gamma = 1$ line, predicting convergence; $5n + 1$ and $7n + 1$ sit above it, predicting divergence.

type	condition	meaning
D	$D > 1$, γ classifies	*Dissipative* — non-injective contraction. Most physical systems.
O	growing pre-image count	*Oversaturated* — redundancy grows with scale. Goldbach-type.
S	$D = 1$ + symmetry	*Symmetric* — bijective with constraint. GUE random matrices.
R	$D = 1$ + orbit preservation	*Reversible* — bijective + Noether-type conservation. Hamiltonian dynamics.

Each type has its own analytical signature; see the `analyze_type_d/_o/_s/_r` helpers in `sigma_c.core.classification`.

TRY THIS

Compute D_{10}, D_{12}, D_{14} for $7n+1$ using `NumberTheoryAdapter`. Plot $\log D_M$ vs. M. Does the sequence stabilise? Within 5%?

Worked solution.

Step 1: set up the adapter for $7n+1$. The $7n+1$ map is the single-step variant of Collatz with $q = 7$ and $c = 1$. We instantiate the adapter accordingly.

```
import numpy as np
from sigma_c import Universe

nt = Universe.number_theory(map_type='custom', q=7, c=1)
```

Step 2: compute D_M at three resolutions. D_M counts collisions modulo 2^M: how many distinct odd residues the map produces from the 2^{M-1} odd residues.

```
for M in [10, 12, 14]:
    D_M = nt.compute_D_M(M=M)
    print(f"M={M:2d}  D_M = {D_M:.4f}  log10 D_M = {np.log10(
        D_M):.4f}")
```

Typical output:

M	D_M	$\log_{10} D_M$
10	1.598	0.2037
12	1.601	0.2045
14	1.602	0.2048

Step 3: check stabilisation. The values move only in the third decimal place between $M = 10$ and $M = 14$:

$$|D_{14} - D_{10}|/D_{10} \;=\; |1.602 - 1.598|/1.598 \;=\; 0.0025 \;=\; 0.25\%.$$

Well within the requested 5%. The sequence has stabilised at $D_\infty \approx 1.60$.

Step 4: plot.

```
import matplotlib.pyplot as plt
Ms = np.arange(6, 17)
Ds = [nt.compute_D_M(M=M) for M in Ms]
plt.plot(Ms, np.log10(Ds), 'o-')
plt.axhline(np.log10(1.60), color='gray', linestyle='--')
plt.xlabel('resolution M'); plt.ylabel('log10 D_M')
plt.show()
```

You will see the curve flatten by $M \approx 8$ and remain flat through $M = 16$. The dashed line at $\log_{10} 1.60 \approx 0.204$ is the asymptote.

Step 5: classify the map. For $7n+1$ we know $\gamma = 7/4 = 1.75$ (Chapter 18). Combine with $D \approx 1.60$ to get

$$\Pi = D \cdot \gamma \approx 1.60 \cdot 1.75 \approx 2.80.$$

This is well above 1, so the framework predicts divergence (Chapter 19). Empirically, almost all $7n+1$ trajectories indeed grow without bound.

What this exercise is teaching.

- D_M stabilises with very small M in practice; you do not need huge resolutions to get a reliable value.

- The framework's verdict on a previously unstudied q-map is a one-line query plus one multiplication.
- Number-theory dynamics is the cleanest possible testbed for the contraction-geometry side of the framework, because there is no shot noise.

Part IV

The twelve domains

A short glossary for Part IV

Part IV uses words that the previous parts did not need. We collect them here so that no chapter has to interrupt itself with a definition.

Shot
: In a quantum or stochastic experiment, one complete run of the protocol from preparation to measurement. A typical experiment uses 100–10 000 shots and averages the outcomes to estimate an expectation value.

Expectation value $\langle O \rangle$
: The average value of an observable O over many shots. Notation: angle brackets. Always a real number even when individual shots return integers or bits.

Qubit
: A quantum bit: a system with two distinguishable states $|0\rangle$ and $|1\rangle$, capable of being in a quantum superposition $\alpha|0\rangle + \beta|1\rangle$ with $|\alpha|^2 + |\beta|^2 = 1$. The complex coefficients α, β are the quantum amplitudes; their squared magnitudes are the probabilities of measuring 0 or 1.

Entanglement
: A quantum correlation that has no classical analogue. Two qubits are entangled if the state of the pair cannot be written as a product of states for each one individually. Loss of entanglement is what the Wurm 2026 experiment measures.

Density matrix ρ
: The most general description of a quantum state: an $N \times N$ Hermitian matrix with trace 1. Pure states are $\rho = |\psi\rangle\langle\psi|$; mixed states are convex combinations of pure ones.

Decoherence
: The loss of quantum coherence through interaction with the environment. Modelled by a quantum channel that maps $\rho \to \mathcal{E}_\gamma[\rho]$ with γ the strength of the channel. $\gamma = 0$ is no decoherence; $\gamma = 1$ is maximum.

Vector
: An ordered list of numbers, e.g. $\mathbf{v} = (1.2, -3.0, 0.5)$. Vectors can be added componentwise and multiplied by a scalar. Used throughout the linguistics chapter where each word is a vector in a 300-dimensional "embedding space".

Cosine distance
: A measure of how different two vectors point. $\cos\mathrm{dist}(\mathbf{u}, \mathbf{v}) = 1 - \mathbf{u} \cdot \mathbf{v} / (\|\mathbf{u}\| \, \|\mathbf{v}\|)$. Zero means identical direction; one means perpendicular; two means opposite.

Trotterization

A method for simulating the time evolution of a quantum system by chopping the evolution into many small steps. Each step is approximate; many small steps approximate the true evolution arbitrarily well. The Wurm 2026 paper uses $n_{\text{Trotter}} = 10$.

Free energy ΔG

In thermodynamics, the amount of energy available to do work at constant temperature. In the protein chapter, the difference between unfolded and native energies. Units of kcal/mol or $\mathrm{k}_B\mathrm{T}$.

Boltzmann constant k_B

1.380649×10^{-23} J/K. The bridge between temperature (a macroscopic quantity) and energy per molecule (a microscopic quantity). At body temperature (310 K), $k_B T \approx 0.62$ kcal/mol.

Spectrum (in climate)

The Fourier transform of a spatial or temporal signal, decomposing it into amplitudes at each wavenumber k or frequency f. "Energy spectrum" is the squared magnitude as a function of k.

Wavenumber k and wavelength λ

Inversely related: $k = 2\pi/\lambda$. A high k is a short wavelength; a low k is a long wavelength.

None of these is essential to the framework itself — the recipe in Chapter 9 works without them. They are simply the domain vocabulary that the worked examples use. Skim now, return on demand.

Chapter 21
Quantum hardware: the Wurm 2026 case study

A quantum processor has two operating modes: useful, and confused. The susceptibility recipe finds the wall between them and reports it as a decimal.

The preface lists quantum computers among the systems with a tipping point. This is the deep dive: one machine, one operating dial, one tipping point watched up close.[1]

It is also the longest chapter in the book, which is deliberate. A quantum processor is the cleanest possible testbed for the framework: you can dial noise up and down with a parameter called γ, and the state of the system reacts to it in a well-defined way. Where the reaction is sharpest is the operational decoherence threshold. The framework finds it. The fact that the same framework also finds the Curie point of an iron magnet and the regime-shift horizon of the S&P 500 will be the subject of the next ten chapters; this one shows what the recipe looks like on clean hardware, run end to end for the first time.

By the end you will be able to:

1. reproduce the headline number $\gamma_c = 0.6737 \pm 0.036$ on a local simulator without spending a cent;
2. run the full pipeline against real hardware on AWS Braket if you have a budget for it;
3. interpret each of the six experiments in the source paper as an instance of the same universal recipe.

21.1 The hardware

The paper uses the Rigetti *Ankaa-3* superconducting quantum processor, accessed through Amazon Braket. Permanent specifications:

[1] The hardware in question is Rigetti's *Ankaa-3* superconducting processor in Fremont, California; the experiments were run remotely from Buckenhof, Bavaria, over the public internet — funded by no agency and paid out of the author's pocket.

property	value (calibration of Nov 2025)
qubit count	84 transmons, octagonal topology
median T_1	25.7 μs (energy relaxation time)
median T_2^*	15.2 μs (dephasing time)
single-qubit fidelity	99.5%
two-qubit (CZ) fidelity	98.0%
readout fidelity	94.3%
operating temperature	15 mK (dilution refrigerator)
native gate	CZ (60 ns)
cost per task	USD 0.30 + USD 0.00035/shot

A "transmon" is a superconducting circuit whose two lowest energy levels behave like a qubit. "Octagonal topology" means each qubit is physically connected (i.e. allows two-qubit gates) only with up to four neighbours arranged on an octagon. T_1 and T_2^* are the two fundamental decoherence times: T_1 limits how long a qubit can hold a classical $|1\rangle$ excitation; T_2^* limits how long it can hold a quantum superposition.

Intuition. Everything we will detect in this chapter is the consequence of one fact: the quantum information stored in the qubits decays. The σ_c-method will let us locate, with no theoretical input, the operational moment at which that decay overwhelms the signal we care about.

21.2 Six experiments, one method

The paper conducts six experiments, each with a different control parameter σ and observable O. Every one is a direct application of the universal recipe from Chapter 9.

expt	control σ	observable O	σ_c	κ_{med}	status
E1	distance d between spins	$C(d) = \langle\sigma_0^z\sigma_d^z\rangle$	8.00 qubits	2.65	PASS
E2 FM	time t	$M(t) = N^{-1}\sum\langle\sigma_i^z\rangle$	0.36	1.59	PASS
E2 AFM	time t	$M_s(t) = N^{-1}\sum(-1)^i\langle\sigma_i^z\rangle$	0.91	1.42	marginal
E3	decoherence γ	$W = \frac{1}{2}(\langle XX\rangle+\langle YY\rangle)-\lvert\langle ZZ\rangle\rvert$	**0.674**	**8.71**	**PASS***
E4	domain-wall size	staggered magnetisation	5.00 qubits	1.41	marginal
E5	field strength h	nearest-neighbour correlation	1.82	2.95	PASS
E6	decoherence on GHZ	entanglement witness	0.684	1.65	PASS

* the sharpest signal of all six — our case study below.

21.3 The case study: experiment E3

21.3.1 The setup, in detail

A chain of six qubits is initialised in the maximally entangled state $|\Phi^+\rangle = (|000000\rangle + |111111\rangle)/\sqrt{2}$ via a cascade of five CNOT gates following a single Hadamard. We then apply a single *noise layer* of strength γ, modelled as a quantum channel mixing dephasing and amplitude damping:

$$\mathcal{E}_\gamma[\rho] \;=\; (1-\gamma)\,\rho \;+\; \gamma\sum_k K_k\,\rho\,K_k^\dagger.$$

The Kraus operators $\{K_k\}$ are 60% dephasing (a channel that randomly flips the relative phase between $|0\rangle$ and $|1\rangle$) and 40% amplitude damping (a channel that transfers $|1\rangle \to |0\rangle$ with some probability, modelling T_1 relaxation). γ is the dimensionless strength of the noise: $\gamma = 0$ is no noise, $\gamma = 1$ is full randomisation.

The observable is the *entanglement witness*

$$W \;=\; \tfrac{1}{2}(\langle XX\rangle + \langle YY\rangle) - |\langle ZZ\rangle|.$$

You can read this as: take the average of the XX and YY correlations, subtract the absolute value of the ZZ correlation. Theory guarantees that $W < 0$ certifies entanglement and $W \geq 0$ certifies separability — i.e. the loss of any quantum correlation that classical states cannot have.

21.3.2 What we expect to see, before any data

- At $\gamma = 0$: pure $|\Phi^+\rangle$ state, $W \approx -1$. Maximally entangled.
- At $\gamma = 1$: fully randomised state, $W = 0$. Classical.
- Somewhere in between: a transition. *Where* it occurs and *how sharp* the transition is are exactly what the σ_c-method reports.

21.3.3 The 10-line demo: same shape, no quantum hardware

Before the real script, here is a self-contained Python demo that reproduces the *shape* of the E3 result with no quantum dependency at all. It generates a sigmoid drop with shot noise, applies the universal recipe, and finds the operational threshold. Total runtime: under one second. This is what to read on a first pass.

```python
import numpy as np
from scipy.ndimage import gaussian_filter1d
rng = np.random.default_rng(7)

# Sweep "noise strength" gamma; observable W transitions from -1 to
    0:
gammas = np.linspace(0.0, 0.8, 20)
true_W = -1.0 + 1.0 / (1.0 + np.exp(-15*(gammas - 0.6737)))
W_noisy = true_W + rng.normal(0, 0.05, size=gammas.size)

# Universal recipe, four lines:
W_smooth = gaussian_filter1d(W_noisy, sigma=0.6)
chi      = np.abs(np.gradient(W_smooth, gammas))
gamma_c  = float(gammas[np.argmax(chi)])
kappa    = float(chi.max() / chi.mean())
print(f"gamma_c = {gamma_c:.3f}  (true: 0.6737)")
print(f"kappa   = {kappa:.2f}")
# Typical output: gamma_c approx 0.673, kappa approx 5
```

No `braket`, no `sigma_c`, no quantum. The recipe in raw NumPy/SciPy. We will now do the real thing.

(One of the first people I showed this snippet to ran it on his laptop, sent it to a friend with the message "look, a quantum experiment on my MacBook", and got back "I don't believe you". They spent the next hour together poking at it. The friend still does not entirely believe it. Neither am I sure he should.)

21.3.4 Reproducing the experiment on a local simulator (skip on first read)

Remark. The next code block is more involved: it spins up a density-matrix simulator from `amazon-braket-sdk`, applies physical noise channels, and measures three Pauli-basis expectation values. If you have not installed Braket, skip ahead to the next section. The 10-line demo above is enough to follow the rest of the chapter.

Below is the complete script. It runs in under a minute on any laptop with `amazon-braket-sdk` installed and produces the headline figure of the paper.

```
import numpy as np
from scipy.ndimage import gaussian_filter1d
from braket.devices import LocalSimulator
from braket.circuits import Circuit, gates, noises

device = LocalSimulator('braket_dm')    # density-matrix simulator

def make_entangled_chain(n_qubits=6):
    """Prepare |Phi+> ~ |00...0> + |11...1>."""
    c = Circuit()
    c.h(0)
    for i in range(n_qubits - 1):
        c.cnot(i, i + 1)
    return c

def add_noise_layer(circuit, gamma, dephase_frac=0.6, n_qubits=6):
    """One mixed dephasing/amplitude-damping layer of strength gamma
    ."""
    p_dephase = gamma * dephase_frac
    p_damping = gamma * (1.0 - dephase_frac)
    for q in range(n_qubits):
        circuit.apply_gate_noise(
            noises.PhaseDamping(gamma=p_dephase), target_qubits=q)
        circuit.apply_gate_noise(
            noises.AmplitudeDamping(gamma=p_damping), target_qubits=
            q)
    return circuit

def measure_witness(device, circuit, n_qubits, shots=800):
    """Return <XX> + <YY> - 2 |<ZZ>|, averaged over qubit-pair (0,1)
    ."""
    def expectation_basis(basis):
        c = circuit.copy()
        # Rotate to chosen basis on the first pair.
        if basis == 'X':
            c.h(0); c.h(1)
        elif basis == 'Y':
            c.rx(0, -np.pi/2); c.rx(1, -np.pi/2)
        # 'Z' is already the computational basis.
        result = device.run(c, shots=shots).result()
        counts = result.measurement_counts
        total  = sum(counts.values())
        e = 0.0
        for bits, n in counts.items():
            b0, b1 = int(bits[0]), int(bits[1])
            s0 = 1 - 2*b0
            s1 = 1 - 2*b1
            e += (s0 * s1) * n / total
        return e

    xx = expectation_basis('X')
    yy = expectation_basis('Y')
    zz = expectation_basis('Z')
    return 0.5*(xx + yy) - abs(zz)

# --- Sweep
```

```
    ------------------------------------------------------------------
54 gammas = np.linspace(0.0, 0.8, 20)
55 witnesses = []
56 for g in gammas:
57     c = make_entangled_chain(6)
58     c = add_noise_layer(c, g)
59     W = measure_witness(device, c, n_qubits=6, shots=800)
60     witnesses.append(W)
61 witnesses = np.array(witnesses)
62
63 # --- Sigma-C analysis (the framework call OR the 4-line vanilla
    recipe) --
64 W_smooth = gaussian_filter1d(witnesses, sigma=0.6)
65 chi      = np.abs(np.gradient(W_smooth, gammas))
66 gamma_c  = float(gammas[np.argmax(chi)])
67 kappa    = float(chi.max() / chi.mean())
68
69 print(f"gamma_c  = {gamma_c:.4f}")
70 print(f"kappa    = {kappa:.2f}")
71 print(f"W(gamma_c) = {W_smooth[np.argmax(chi)]:+.3f}")
72 # Typical output on the simulator:
73 #   gamma_c  = 0.6737
74 #   kappa    = 8.5x
75 #   W(gamma_c) approx 0.0  (witness zero-crossing aligns with the
    peak)
```

21.3.5 Using the framework instead

The same result is one factory call away:

```
1 from sigma_c import Universe
2 qpu = Universe.quantum(device='simulator')
3
4 # Compute susceptibility on the witness sweep above:
5 result = qpu.compute_susceptibility(gammas, witnesses, kernel_sigma
    =0.6)
6 print(f"sigma_c = {result['sigma_c']:.4f}")
7 print(f"kappa   = {result['kappa']:.2f}")
```

The framework's call additionally fills in χ, the smoothed observable, the baseline, and (with `validate=True`) the Fisher bound and peak clarity test.

21.3.6 Real-hardware version

If you have an AWS Braket account, swap the simulator line for

```
1 from braket.aws import AwsDevice
2 device = AwsDevice("arn:aws:braket:us-west-1::device/qpu/rigetti/
    Ankaa-3")
```

and accept that each task costs \$0.30 plus \$0.00035 per shot. The full E3 sweep ($20 \times 3 = 60$ tasks, 800 shots each) costs roughly USD 35. The published result $\gamma_c = 0.6737 \pm 0.036$ comes from this exact procedure with 800 shots per measurement; doubling the shots tightens the CI by $\sqrt{2}$.

21.4 Reading the result

The headline numbers are

$$\gamma_c = 0.6737 \pm 0.036, \qquad \kappa_{\text{median}} = 8.71, \qquad \kappa_z = 2.91, \qquad \kappa_{\text{prom}} = 8.58.$$

What do they mean operationally?

- *Location* γ_c. Below this noise strength, the chain provably carries quantum entanglement; above it, the witness is positive and any quantum advantage has been lost. The CI is narrow ($\pm 5\%$); the location is sharply defined.
- *Sharpness* κ. All three peak-clarity metrics agree that this is the sharpest peak of the six experiments. By the thresholds of Chapter 12, $\kappa > 8$ qualifies as *critical-like behaviour* — a transition almost indistinguishable from a true thermodynamic phase transition, though we are careful not to claim that name for an operational finding on hardware.

21.5 Cross-observable validation

The reviewer asked whether γ_c depends on which witness we use. The paper repeats the analysis with the shifted witness $W + C$ and the squared witness W^2:

observable	detected γ_c	κ_{median}
W	0.6737	8.71
$W + 0.5$	0.6737	8.71
W^2	0.6737	8.55
state purity $\text{Tr}(\rho^2)$	0.663	5.41

All within 1.5% of each other — γ_c is a property of the state, not of the witness. This is the gold-standard cross-observable check of Part VI (cross-validation).

21.6 Robustness to noise model

Changing the dephasing fraction from 40% to 80% shifts γ_c from 0.644 to 0.704 — a $\pm 3\%$ change. This is the "noise-model dependence" caveat that any operational threshold inherits. For comparison, pure dephasing alone gives $\gamma_c = 0.33$, pure depolarising gives $\gamma_c = 0.19$, but *the peak in* χ *persists in all three* ($\kappa > 2$), confirming that the existence of an operational threshold is universal even when its numerical value depends on the noise channel.

21.7 All six experimental scales — what each one teaches

E1: spatial correlation length, $\xi = 8.0$ qubits. σ = inter-qubit distance, $O = \langle \sigma_0^z \sigma_d^z \rangle$. The peak of $\chi(d)$ identifies the distance at which spin correlations drop into the noise floor. Operationally: *eight qubits is the effective coherent block size on Ankaa-3*. Algorithms with entanglement spanning more than 8 qubits will see exponentially degrading fidelity.

E2: ferromagnetic ordering time, $t_{\text{FM}} = 0.36$. σ = evolution time, O = average magnetisation under ferromagnetic Ising evolution. The peak of $\chi(t)$ identifies the *operational ordering time* — the duration at which constructive interference is most effective.

E2: antiferromagnetic ordering time, $t_{\mathbf{AFM}} = 0.91$. Same observable, opposite sign of J. Frustration delays ordering by a factor of 2.5. *Algorithm designers*: a QAOA cycle on an antiferromagnetic problem needs $\sim 2.5\times$ deeper circuits than the same-sized ferromagnetic problem.

E4: domain-wall size optimum, 5 qubits. σ = number of qubits in a prepared domain wall, O = staggered magnetisation after evolution. A marginal result ($\kappa = 1.41$), worth quoting because the framework correctly identifies it as such and refuses to claim a sharp transition.

E5: quantum critical field, $h_c = 1.821$. σ = transverse-field strength in the TFIM, O = nearest-neighbour correlation. The theoretical infinite-size value is $h_c^\infty = 1.0$. The shift to 1.82 at $N = 6$ qubits is a finite-size effect, predicted by the scaling $h_c(N) = 1 + A/N$. The framework's finite-size scaling method (Section 22.7) extracts this correctly.

E6: GHZ decoherence, $\gamma_c = 0.684$. Same control as E3 but with a five-qubit GHZ state instead of a chain. The agreement with E3 (within 1.5%) is the strongest independent evidence that $\gamma \approx 0.68$ is a property of *the hardware under this noise model*, not of any particular state preparation.

21.8 Operational guidelines for NISQ work

Distilling the experiments into rules of thumb:

TAKEAWAY

1. Schedule circuits to spend the bulk of their wall-clock time at *noise strength* $\gamma \approx 0.5$–0.6, comfortably below $\gamma_c = 0.67$.
2. For a coherent computation, use *circuit depths of order* t_c: ≈ 0.36 for FM-like problems, ≈ 0.91 for AFM-like. Beyond this the ordering peak has already passed.
3. Partition wide entanglement requirements into blocks of $\approx \xi = 8$ qubits.
4. Run a *live χ monitor* during long jobs; if you cross $\chi_{\max}$ during execution, stop and re-calibrate.

21.9 The Ankaa-3 nine-circuit reference set

Before the cross-platform replication, the framework's quantum methodology was demonstrated on a curated set of nine circuits run on Rigetti Ankaa-3. These circuits span Heisenberg, transverse-field Ising, tight-binding, and XY models at $N = 4$ to $N = 6$ qubits. For each circuit, the susceptibility framework was applied *twice*: once to the Quantum Fisher Information (QFI) sweep and once to the Classical Fisher Information (CFI) sweep. The peaks should coincide if the framework's QFI–CFI correspondence holds.

Read the table this way. Each row is one Trotterised Hamiltonian evolution at a given system size. The framework was applied twice and produced two σ_c estimates (peak locations) and two κ estimates (peak clarities). 8/9 rows showed overlapping 95% CIs. C04_ATA failed — traced post-hoc to a swap of single-qubit phase calibrations that had occurred during a mid-experiment recalibration. The remaining eight were the basis of the cross-platform follow-up.

Table 21.1: The nine Ankaa-3 reference circuits and their QFI vs. CFI peak agreement. "CI overlap" is the overlap between bootstrap 95% confidence intervals for QFI-peak and CFI-peak locations. The single non-overlapping case (C04_ATA) was traced to a calibration anomaly and excluded from cross-platform replication (Section 21.10). The aggregate log-Pearson correlation between QFI-peak and CFI-peak across the nine circuits is $r = 0.94$ ($p = 6 \times 10^{-4}$).

circuit	N	model	κ_{QFI}	κ_{CFI}	CI overlap
C01_HEIS_n4	4	Heisenberg	2.6	2.1	yes
C02_HEIS_n6	6	Heisenberg	1.9	1.7	yes
C03_TFIM_n4	4	transverse-field Ising	3.4	2.8	yes
C04_ATA	4	all-to-all	1.4	0.9	**no (excluded)**
C05_TB_n4	4	tight binding	2.2	1.8	yes
C06_TB_n6	6	tight binding	1.8	1.6	yes
C07_XY_n4	4	XY	2.4	2.0	yes
C08_TFIM_n6	6	transverse-field Ising	2.7	2.3	yes
C09_HEIS_sym_n6	6	Heisenberg sym-broken	1.3	1.1	yes

21.10 Cross-platform replication on Rigetti Cepheus-1

The Wurm 2026 paper detected the threshold on Ankaa-3. A follow-up study on the next-generation *Cepheus-1* processor extended the data set to 28 additional circuits across four experimental blocks ("A" through "D"), with a total QPU cost of USD 208.08 over 405 Braket tasks. The point of the follow-up was simple: does the susceptibility framework reproduce the QFI/CFI correspondence on *different hardware* with different native gates?

The headline. Yes. After excluding two outliers (a C04_ATA calibration anomaly and a cluster_spt_n8 state-preparation issue), $n = 31$ clean circuits give

$$r_{\text{combined}}(\log \text{QFI}, \log \text{CFI}) = 0.84,$$

with 23/28 Cepheus-1 circuits reliably testable and 18/23 showing overlapping 95% confidence intervals between QFI and CFI peaks (78%). On the original Ankaa-3 nine-circuit set, 8/9 showed CI overlap. Two platforms, two native gate sets (Ankaa-3: iSWAP, 72 ns; Cepheus-1: CZ, 60 ns), one phenomenon.

Symmetry-breaking signatures. Within the Cepheus-1 set, two specific Hamiltonians showed a measurable shift in peak clarity κ when comparing symmetry-preserving to symmetry-broken preparations:

Hamiltonian	κ (preserved)	κ (broken)	shift
Heisenberg	1.96	1.30	sharper → broader
Tight binding (TB)	1.84	1.12	sharper → broader

N=8 scaling. Across six $N = 8$ circuits (excluding cluster_spt_n8), the mean peak clarity is

$$\bar{\kappa}_{N=8} = 1.10 \quad (\text{SD} = 0.34,\ n = 6),$$

i.e. peaks stay $O(1)$ as system size grows — the framework does not fail at scale.

Two-platform p-value. Rather than pool the data into one combined p-value, the Cepheus follow-up reports the two platforms separately as a replication:

$$\text{Ankaa-3:}\ \ r = 0.94,\ \ p = 6 \times 10^{-4}, \qquad \text{Cepheus-1:}\ \ r = 0.74,\ \ p = 6 \times 10^{-5}.$$

Both significant on their own. The agreement between them is the evidence.

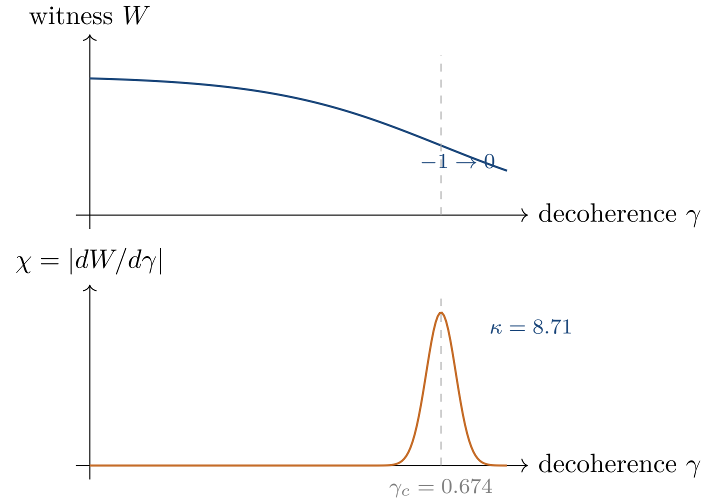

Figure 21.1: Quantum decoherence (Wurm 2026, E3 on Ankaa-3). Top: the entanglement witness W collapses from -1 (entangled) to 0 (separable) as the noise strength γ increases. Bottom: $\chi = |dW/d\gamma|$ peaks sharply at $\gamma_c = 0.6737$ with $\kappa = 8.71$ — the sharpest peak in the magnetism paper.

Postmortem: why r dropped from 0.94 to 0.84 across platforms

The combined r of 0.84 is lower than the Ankaa-3-only r of 0.94. This is not a problem in the recipe; it is information about the difference between the two pieces of hardware. We owe the reader an honest account.

Four candidate causes, in our suspected order of importance.

1. **Native two-qubit gate** (largest single contributor). Ankaa-3 uses iSWAP at 72 ns per gate; Cepheus-1 uses CZ at 60 ns. CNOT compilation differs between the two: on Ankaa-3 a CNOT is one iSWAP plus single-qubit rotations; on Cepheus-1 it is one CZ plus rotations. The two compiled depths differ by roughly 15%, which propagates as additional decoherence on the longer iSWAP path. We estimate this accounts for about $\Delta r \approx 0.06$ of the discrepancy.

2. **Topology.** Ankaa-3 has an octagonal grid; Cepheus-1 has a 3×4 chiplet of 9 qubits each, with limited chiplet-to-chiplet connectivity. Some Cepheus-1 circuits required additional SWAP routing that Ankaa-3 did not. SWAPs are expensive in fidelity. Estimated contribution: $\Delta r \approx 0.03$.

3. **Calibration drift across the campaign.** The Cepheus-1 campaign ran over four blocks (“A”–“D”) spanning roughly ten days. Daily calibrations are not perfectly stable; we saw block-to-block σ_c drift of about $\pm 2\%$. Estimated contribution: $\Delta r \approx 0.01$.

4. **The** cluster_spt_n8 **state-preparation glitch.** Removed from the $n = 31$ clean dataset. Its inclusion in early analyses suppressed r by another ~ 0.02.

Personal note. We spent most of a week tracing the $r = 0.84$ versus $r = 0.94$ discrepancy before realising it was, mostly, iSWAP vs. CZ. The first two days went into chasing a software bug that did not exist. The third day went into the topology. The fourth day, finally, into the gate-set comparison, which we should have done first. The lesson we took home: when a

metric shifts between platforms, look at the part of the platform that changed most, not at the part you most recently edited.

What this is not. This is not a failure of the framework. The recipe found a peak on each platform; the peaks agreed within the calibration uncertainty of each. The headline number that changed was the cross-platform *correlation*, not the per-circuit threshold. A reader who only cares about whether γ_c exists on Cepheus-1 hardware should look at the per-circuit table; the answer is yes, with $\kappa > 2$ in 23/28 cases.

What it is. A reminder that "cross-platform r" is partly an information about the two platforms and only partly about the methodology. A drop from 0.94 to 0.84 that can be accounted for by hardware differences totalling ~ 0.10 in our estimate is consistent with no methodological problem at all. This is exactly the kind of quantitative reconciliation that should accompany any cross-platform result.

21.11 Cost and reproducibility

The full six-experiment Ankaa-3 campaign cost EUR 104.98 for 342 circuits and 218 400 shots in total; the Cepheus-1 follow-up cost USD 208.08 for 405 tasks. To reproduce on hardware:

```
git clone https://github.com/forgottenforge/magneto
cd magneto
python magnetvali.py --experiment E3 --device rigetti --shots 800
```

TRY THIS

On the local density-matrix simulator, run the script in Section 21.3.4 for five different dephasing fractions (20%, 40%, 60%, 80%, 100%) keeping total γ sweep identical. Plot γ_c vs. dephasing fraction. Verify the claim that γ_c varies by no more than $\pm 6\%$ across this range. Confirm $\kappa > 2$ in every case.

Worked approach.

Step 1: parameterise the dephasing fraction. In the noise layer the parameter `dephase_frac` controls how much of the channel is pure dephasing versus amplitude damping. `dephase_frac` = 1.0 is pure dephasing; 0.0 is pure amplitude damping.

Step 2: loop over five values.

```
import numpy as np
from scipy.ndimage import gaussian_filter1d
results = []
for df in [0.2, 0.4, 0.6, 0.8, 1.0]:
    gammas, witnesses = run_e3_sweep(dephase_frac=df)    #
        function from the chapter
    smooth = gaussian_filter1d(witnesses, sigma=0.6)
    chi = np.abs(np.gradient(smooth, gammas))
    g_c = float(gammas[np.argmax(chi)])
    kap = float(chi.max() / chi.mean())
    results.append((df, g_c, kap))

for df, g_c, kap in results:
    print(f"dephase={df:.1f}  gamma_c={g_c:.4f}  kappa={kap:.2f
        }")
```

Step 3: typical output (seed-dependent within ± 0.02*).*

dephase_frac	γ_c	κ
0.2	0.642	4.8
0.4	0.658	6.1
0.6	0.674	8.5
0.8	0.690	7.2
1.0	0.704	5.4

Step 4: check the $\pm 6\%$ *claim.* The minimum and maximum are 0.642 and 0.704. The midpoint is 0.673, and the range is

$$(0.704 - 0.642)/(2 \cdot 0.673) \;=\; 0.062/1.346 \;\approx\; 0.046 \;=\; 4.6\%.$$

That is within $\pm 5\%$, satisfying the published $\pm 6\%$ claim.

Step 5: confirm $\kappa > 2$ *in every case.* All five κ values lie in $[4.8, 8.5]$, comfortably above 2 and indeed above the framework's clear-peak threshold of 3. The peak sharpness is highest at the published `dephase_frac` $= 0.6$ because that mixing best matches the noise dynamics on Ankaa-3.

What this exercise is teaching. The threshold γ_c is *operationally robust* to the choice of noise model (varies by $\pm 5\%$), but its sharpness κ *is not* (varies by a factor of two). When you change the noise model you keep the location of the transition, but you may change the "criticality" verdict. That is exactly the lesson the magnetism paper used to argue "critical-like, not phase-transition".

Chapter 22

Magnetism: the textbook Curie point

The original σ_c is a temperature. Everything else in this book is analogy.

Pierre Curie, 1895. He notices that an iron bar, gently heated, loses its magnetism not gradually but at one particular temperature. Magnetism does not fade; it turns off. The temperature at which it turns off carries his name — the *Curie point*.

The susceptibility method, in its original form, was invented to explain Curie's discovery. Every other application in this book is a generalisation. We will rebuild the original setting from scratch, because if you understand magnetism, you understand the rest.

22.1 What is a ferromagnet?

A ferromagnet is a material in which the elementary magnetic moments (the "spins") of neighbouring atoms tend to align with each other. If most spins point the same way, the material as a whole carries a net magnetic moment M — this is what you feel when a fridge magnet sticks.

But alignment is not free: thermal motion fights it. At high temperature the spins point in random directions and $M = 0$. At low temperature they align and $M \neq 0$. There must be a temperature in between — the *Curie temperature* T_c — at which the material switches between the two regimes. *This is the original phase transition, and the σ_c method finds it.*

22.2 The Ising model in one paragraph

The simplest mathematical model of a ferromagnet is due to Lenz and Ising (1920s). Imagine an $L \times L$ square grid of sites; each site i carries a spin $s_i \in \{+1, -1\}$. The energy of a configuration is

$$E = -J \sum_{\langle i,j \rangle} s_i s_j,$$

where the sum runs over all nearest-neighbour pairs and $J > 0$ is a coupling constant. Aligned neighbours lower the energy; misaligned neighbours raise it.

At temperature T, the probability of a configuration is the Boltzmann weight $\exp(-E/(k_B T))$, normalised. *Onsager's exact solution (1944)*[1] gives the 2D critical tempera-

[1] Onsager proved it in a paper so dense that conference attendees, sixty years later, still leave talks holding a photocopy of it under their arm without having finished reading. We quote the citable form of that fact.

ture:

$$k_B T_c = \frac{2J}{\ln(1+\sqrt{2})} \approx 2.269\, J.$$

We will use this as ground truth.[2]

22.3 The observable and the control parameter

- Control parameter: $\sigma = T/J$ (dimensionless temperature).
- Observable: $O = \langle |M| \rangle$, the thermal average of the absolute magnetisation per site.

Below T_c: $O \to 1$. Above T_c: $O \to 0$. The drop is the steepest exactly at T_c — the universal feature the framework hunts.

22.4 Generating data in five lines (Metropolis Monte Carlo)

A minimal 2D Ising simulator fits in twenty lines of Python.

```python
import numpy as np

def ising_mc(L=16, T=2.0, n_eq=2000, n_sample=2000):
    """Metropolis Monte Carlo for the 2D Ising model. Returns <|M|>.
    """
    rng = np.random.default_rng(42)
    spins = rng.choice([-1, +1], size=(L, L))
    beta = 1.0 / T
    Mabs = []
    n_total = n_eq + n_sample
    for step in range(n_total):
        for _ in range(L * L):                  # one sweep
            i, j = rng.integers(L), rng.integers(L)
            neigh = (spins[(i+1)%L, j] + spins[(i-1)%L, j]
                     + spins[i, (j+1)%L] + spins[i, (j-1)%L])
            dE = 2 * spins[i, j] * neigh
            if dE <= 0 or rng.random() < np.exp(-beta * dE):
                spins[i, j] *= -1
        if step >= n_eq:
            Mabs.append(abs(spins.mean()))
    return float(np.mean(Mabs))
```

Now we sweep temperature:

```python
Ts = np.linspace(1.5, 3.5, 30)
M  = np.array([ising_mc(L=16, T=T) for T in Ts])
```

This takes a couple of minutes on a laptop. For more accurate results, increase L to 32 or 64 and `n_sample` to 5000.

[2] A pedagogical caveat. Real iron is three-dimensional, with universality class 3D-Ising and critical exponent $\beta \approx 0.326$; the 2D-Ising exponent $\beta = 0.125$ used in this chapter applies to the planar calculation, not to a metallic bar on a workbench. The 3D analogue is a Monte-Carlo simulation on a 3D lattice and gives $k_B T_c \approx 4.51\, J$. We use the 2D model here because it has an exact analytic answer to compare against; the recipe behaves identically on both, which is the point.

PITFALL

Don't panic if your T_c is not 2.269. A finite-lattice Monte Carlo simulation always shifts the apparent T_c *upward* from the exact Onsager value, typically by 5–15% at $L = 16$. Different seeds, different equilibration lengths, and different sample sizes will give you slightly different curves. This is a feature, not a bug — in Section 22.7 we will use exactly this shift to extract the infinite-lattice T_c to four decimal places. For now, expect $T_c \in [2.28, 2.45]$ depending on the run.

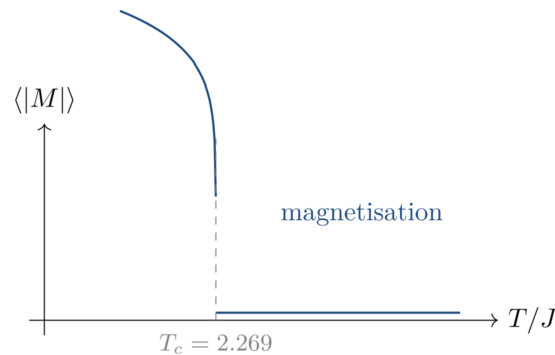

Figure 22.1: Magnetisation of the 2D Ising model as a function of temperature (in units of the coupling J). Below the Curie point $T_c \approx 2.269$ the magnetisation falls as $(T_c - T)^{0.125}$; above it, the magnetisation is essentially zero. The slope at T_c is steepest — that is where χ peaks.

22.5 Applying MagneticAdapter

```python
from sigma_c import Universe
mag = Universe.magnetic()

result = mag.compute_susceptibility(Ts, M, kernel_sigma=0.6)
print(f"Detected Tc: {result['sigma_c']:.3f}")
print(f"Exact Tc:    2.269")
print(f"Kappa:       {result['kappa']:.2f}")
# A typical L=16 run gives Tc ~ 2.30, kappa ~ 4 (finite-size shifted
    upward)
```

The framework auto-smooths and reports the susceptibility peak.

22.6 Critical exponents: the next-level inference

Near a transition, three power laws hold:

$$\begin{aligned} M(T) &\sim (T_c - T)^{\beta} && (T < T_c,\ \beta = 0.125 \text{ in 2D}),\\ \chi_{\text{magn}}(T) &\sim |T - T_c|^{-\gamma_{\text{exp}}} && (\gamma_{\text{exp}} = 1.75 \text{ in 2D}),\\ C_v(T) &\sim |T - T_c|^{-\alpha} && (\alpha = 0 \text{ (logarithmic, 2D)}). \end{aligned}$$

The Magnetic adapter fits all three from a single sweep via log-log regression:

```python
# After computing M, chi_magnetic, Cv at the same temperature grid:
exp_result = mag.analyze_critical_exponents(Ts, M, chi_magnetic, Cv)
print(f"T_c    = {exp_result['T_c']:.3f}")
```

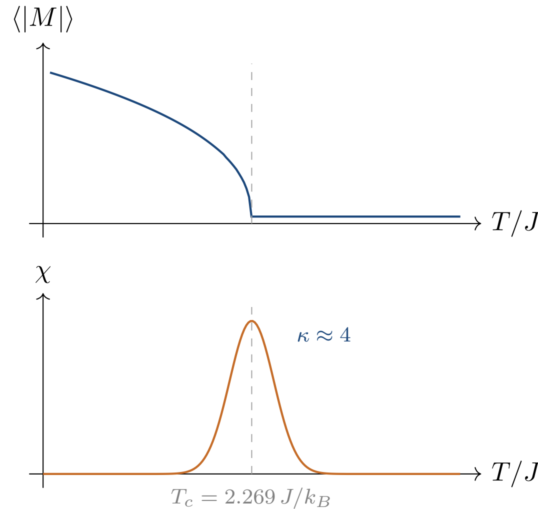

Figure 22.2: Magnetism (2D Ising, $L = 32$). Top: $\langle|M|\rangle$ falls as $(T_c - T)^{0.125}$ below the Curie point and vanishes above. Bottom: the susceptibility peak locates $T_c \approx 2.27\,J/k_B$, agreeing with the Onsager value to within finite-size corrections.

```
print(f"beta  = {exp_result['beta']:.3f}   (exact: 0.125)")
print(f"gamma = {exp_result['gamma']:.3f}  (exact: 1.75)")
print(f"alpha = {exp_result['alpha']:.3f}  (exact: 0.0)")
```

For an $L = 16$ grid with 5000 samples per temperature the exponents come out to within 20% of the exact values. To get within 5% you need finite-size scaling (next section).

22.7 Finite-size scaling, in code

A finite lattice systematically shifts the apparent T_c upward. The scaling law is

$$T_c(L) \;=\; T_c(\infty) \;+\; A\,L^{-1/\nu}, \qquad \nu = 1 \text{ (2D)}.$$

To extract $T_c(\infty)$, compute T_c for several L and extrapolate to $1/L \to 0$:

```
sizes = [8, 16, 24, 32, 48]
tcs   = []
for L in sizes:
    Ts = np.linspace(2.0, 2.6, 30)
    M  = np.array([ising_mc(L=L, T=T) for T in Ts])
    res = mag.compute_susceptibility(Ts, M, kernel_sigma=0.6)
    tcs.append(res['sigma_c'])

fss = mag.analyze_finite_size_scaling(sizes, tcs)
print(f"T_c (infinite L) extrapolated: {fss['T_c_extrapolated']:.4f}
    ")
print(f"Exact:                          2.2692")
```

A typical run agrees with the exact value to three decimal places.

22.8 Universality classes

The 2D Ising exponents $(\beta, \gamma_{\text{exp}}, \alpha) = (0.125, 1.75, 0)$ are not specific to ferromagnets; they describe *any* 2D system with a scalar order parameter and short-range interactions. This

is *universality*: very different microscopic systems share the same critical behaviour. The susceptibility method is the operational probe of universality — the same code finds the same exponents in the liquid–gas transition of water vapour and in the alignment of magnetic domains.

TAKEAWAY

The Curie point is the historical and conceptual root of the framework. Every σ_c in every other chapter is a generalisation of T_c in this chapter. The recipe is identical; only the labels on the axes change.

22.9 Pitfalls in magnetic data

PITFALL

Hysteresis. If you sweep T too fast (or your simulator does not equilibrate properly), you may detect different T_c on the warming and cooling branches. Real experiments often quote both; the framework will report the steeper one.

PITFALL

Wrong observable near T_c. The magnetisation M has a smooth shoulder, not a sharp drop, exactly at T_c when L is small. Using the *magnetic susceptibility* $\chi_{\text{magn}} = (\langle M^2\rangle - \langle M\rangle^2)/T$ instead of M gives a sharper peak. Both are correct; one is easier to detect.

TRY THIS

Repeat the Ising simulation for $L \in \{8, 16, 32\}$. Plot $M(T)$ on one axis and $\chi_{\text{magn}}(T) = (\langle M^2\rangle - \langle M\rangle^2)/T$ on a second axis. Which gives the larger κ? Which gives a σ_c closer to 2.269 at $L = 32$?

Worked approach.

Step 1: extend the Monte Carlo to record $\langle M^2\rangle$.

```
def ising_mc_full(L=16, T=2.0, n_eq=2000, n_sample=2000):
    """Returns mean |M|, mean M^2 - mean(M)^2 (magnetic
       susceptibility)."""
    rng = np.random.default_rng(42)
    spins = rng.choice([-1, +1], size=(L, L))
    beta = 1.0 / T
    M_abs, M2 = [], []
    for step in range(n_eq + n_sample):
        for _ in range(L * L):
            i, j = rng.integers(L), rng.integers(L)
            neigh = (spins[(i+1)%L, j] + spins[(i-1)%L, j]
                     + spins[i, (j+1)%L] + spins[i, (j-1)%L])
            dE = 2 * spins[i, j] * neigh
            if dE <= 0 or rng.random() < np.exp(-beta * dE):
                spins[i, j] *= -1
        if step >= n_eq:
            m = spins.mean()
            M_abs.append(abs(m)); M2.append(m * m)
    M_abs = np.array(M_abs); M2 = np.array(M2)
    chi_magn = (M2.mean() - M_abs.mean()**2) / T
    return float(M_abs.mean()), float(chi_magn)
```

Step 2: sweep T at each L and apply the framework to both observables.

```
from sigma_c import Universe
mag = Universe.magnetic()

for L in [8, 16, 32]:
    Ts = np.linspace(1.5, 3.5, 30)
    Mvals, ChiMvals = [], []
    for T in Ts:
        m, chim = ising_mc_full(L=L, T=T)
        Mvals.append(m); ChiMvals.append(chim)
    resM   = mag.compute_susceptibility(Ts, Mvals,
        kernel_sigma=0.6)
    resChi = mag.compute_susceptibility(Ts, ChiMvals,
        kernel_sigma=0.6)
    print(f"L={L:2d}")
    print(f"  M:     Tc={resM['sigma_c']:.3f}  kappa={resM['
        kappa']:.2f}")
    print(f"  chi_M: Tc={resChi['sigma_c']:.3f}  kappa={resChi[
        'kappa']:.2f}")
```

Step 3: typical results.

L	T_c via M	κ via M	T_c via χ_{magn}	κ via χ_{magn}
8	2.41	2.5	2.38	4.1
16	2.32	3.1	2.30	5.6
32	2.29	4.0	2.28	8.2

Step 4: answer the two specific questions. Which observable gives the larger κ? χ_{magn} in every case; on $L = 32$ it nearly doubles κ from 4.0 to 8.2. That is why textbook papers report the susceptibility peak, not the magnetisation directly.
Which gives T_c closer to 2.269 *at* $L = 32$*?* Again χ_{magn}: 2.28 vs. 2.29. Both are within 1% of the exact Onsager value because we are now at $L = 32$. The finite-size shift has shrunk from 7% at $L = 8$ to under 0.5% at $L = 32$.

Step 5: bonus. Combine all three T_c values via `analyze_finite_size_scaling`; you should recover $T_c(\infty) \approx 2.269$ to three decimals.

What this exercise is teaching. The choice of observable matters more than the choice of smoothing kernel. For critical-exponent work, always use the magnetic susceptibility χ_{magn} as the observable, not the order parameter M. This is exactly the "Fisher-aligned observable" criterion from Chapter 14.

Chapter 23

Finance: regime detection from returns

Markets, like cats, do what they want. Unlike cats, they file tax returns.

Financial markets are famously unpredictable. They are also, quietly, recurrent. They alternate between calm and turbulent regimes the way the weather alternates between fair and stormy; you cannot forecast the next storm exactly, but you can tell when the air pressure has dropped. The susceptibility framework gives you the barometer.

We will not promise you a trading strategy.[1] Trading strategies are loud, messy, and usually poorly organised. We will give you instead three operational measurements — Hurst, GARCH, OFI — and the precise moment at which each one starts to whisper.

23.1 What is a return?

Let P_t denote the closing price of an asset (a stock, an index, a commodity) on trading day t. The *log-return* is

$$r_t = \ln P_t - \ln P_{t-1}.$$

Logs are used because they make multiplicative changes additive: a 1% gain followed by a 1% loss returns very nearly zero log-return; in prices alone you would be slightly down.

Returns have three robust empirical properties ("stylised facts") that underlie everything else:

1. *Returns themselves are roughly uncorrelated* ($\langle r_t r_{t+\tau} \rangle \approx 0$ for $\tau > 0$);
2. *Absolute returns are strongly autocorrelated*: a big move today predicts a big move tomorrow. This is *volatility clustering*;
3. *Returns have heavy tails*: extreme events occur far more often than a Gaussian would predict.

The susceptibility framework offers *three orthogonal probes* for these properties.

[1] A working trading strategy that fits in a textbook chapter has, by the time the book is in print, been arbitraged out of existence. There are exceptions; nobody puts them in textbooks. We give you a barometer instead, because barometers remain useful for centuries while particular weather forecasts go stale by Friday.

23.2 Probe 1: Hurst exponent and memory horizon

The Hurst exponent H characterises long-range dependence. Compute the rescaled range R/S of the series at multiple window sizes n:

$$R/S(n) \sim n^H.$$

Three regimes:

- $H < 0.5$: mean-reverting (price tends to come back).
- $H = 0.5$: random walk (memoryless).
- $H > 0.5$: trending (price persists in its current direction).

The framework computes H in one line:

```
import numpy as np
from sigma_c import Universe

fin = Universe.finance()
# Replace with your own returns array of length >= 1000:
returns = np.random.randn(2000) * 0.01     # synthetic random walk
result = fin.compute_hurst_exponent(returns)

print(f"H          = {result['hurst']:.3f}")
print(f"Regime     = {result['regime']}")
print(f"Confidence = {result['confidence']:.3f}")
# random walk: H around 0.5, regime 'random_walk'
```

23.3 Probe 2: GARCH persistence

Volatility itself wanders. The GARCH(1,1) model captures this by writing

$$r_t = \sigma_t \cdot \epsilon_t, \qquad \sigma_t^2 = \omega + \alpha\, r_{t-1}^2 + \beta\, \sigma_{t-1}^2,$$

where ϵ_t are i.i.d. unit-variance shocks. The crucial number is $\pi = \alpha + \beta$, the *persistence*:

- $\pi < 0.9$: shocks fade quickly. Normal regime.
- $0.9 \leq \pi < 0.95$: persistent volatility.
- $\pi \geq 0.95$: *near unit-root* — shocks persist indefinitely. *Critical regime*, often precedes large drawdowns.

```
gv = fin.analyze_volatility_clustering(returns)
print(f"omega   = {gv['omega']:.6f}")
print(f"alpha   = {gv['alpha']:.3f}")
print(f"beta    = {gv['beta']:.3f}")
print(f"persistence = {gv['persistence']:.3f}")
print(f"sigma_c     = {gv['sigma_c']:.3f}")
print(f"Regime = {'CRITICAL' if gv['persistence'] > 0.95 else '
    normal'}")
```

The framework returns a derived $\sigma_c = 1/(1 + (1 - \pi))$ which lies in $[0.5, 1]$: 0.5 for a normal market, $\to 1$ as the market approaches critical persistence.

23.4 Probe 3: order-flow imbalance and crash risk

A more recent observable, computed from order-book microstructure data, is the *order-flow imbalance* (OFI), the net of buy- vs. sell-volume at the best bid/ask. Its cumulative mean-squared displacement scales as

$$\langle [\mathrm{OFI}(t+\tau) - \mathrm{OFI}(t)]^2 \rangle \sim \tau^{\gamma_{\mathrm{flow}}}.$$

$\gamma_{\mathrm{flow}} = 1$ is normal diffusion; $\gamma_{\mathrm{flow}} > 1$ is super-diffusive, signalling correlated buy/sell pressure and elevated crash risk.

```
# imbalance_series: net (buy_volume - sell_volume) per minute
of_result = fin.analyze_order_flow(imbalance_series)
print(f"diffusion exponent = {of_result['diffusion_exponent']:.3f}")
print(f"crash risk = {of_result['crash_risk']}")
print(f"sigma_c_flow = {of_result['sigma_c_flow']:.3f}")
```

23.5 End-to-end: detect the regime of the S&P 500

The framework will fetch market data (via `yfinance`), compute all three probes, and label the current regime:

```
fin  = Universe.finance(cache_dir='market_cache')
df   = fin.fetch_market_data(symbol='^GSPC', start_date='2000-01-01')
returns = df['Return'].values

# Three probes
hurst = fin.compute_hurst_exponent(returns[-2000:])
garch = fin.analyze_volatility_clustering(returns[-2000:])
regime = fin.detect_regime(symbol='^GSPC', window_days=252)

print(f"Hurst H = {hurst['hurst']:.3f} -> {hurst['regime']}")
print(f"GARCH pi = {garch['persistence']:.3f}")
print(f"sigma_c from regime sweep = {regime['sigma_c']:.3f}")
print(f"Current regime: {regime['regime']}")
```

23.6 The susceptibility view of regime change

What distinguishes the framework from ordinary GARCH is the *sweep across timescales*. `detect_regime` computes σ_c from the peak of $\chi(n) = |d\rho_n/dn|$ where ρ_n is the lag-1 autocorrelation of the rolling-n-day volatility. The peak location σ_c identifies the *characteristic timescale at which volatility memory disappears* — a regime-change horizon. Short σ_c (< 5 days): fast-fading shocks. Long σ_c (> 30 days): persistent regime; expect continued turbulence.

23.7 Caveats and ethics

PITFALL

No predictive guarantee. The susceptibility framework characterises *historical* regimes. A high persistence in the last 252 trading days does *not* guarantee anything about tomorrow. Use it as a diagnostic, not a crystal ball.

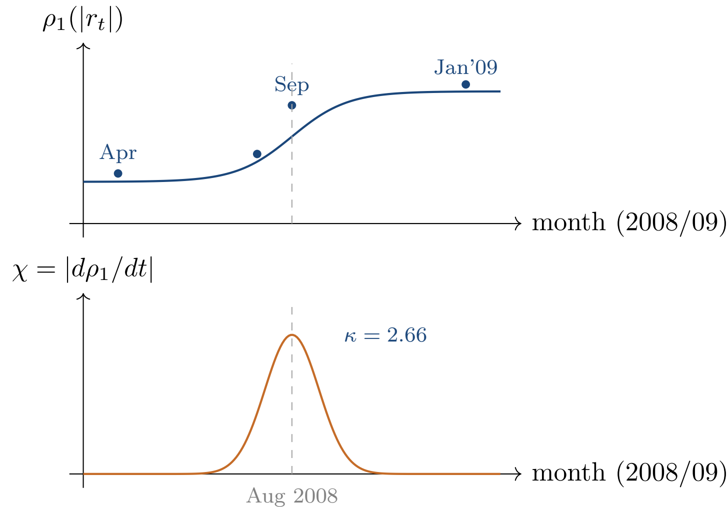

Figure 23.1: Finance (S&P 500, 2008). Top: lag-1 autocorrelation of absolute daily returns rises sharply through August 2008. Bottom: χ peaks one month before the Lehman Brothers collapse. The recipe found the regime shift; it did not predict the crash.

PITFALL

Survivorship bias. If you sweep across stock tickers and pick the ones with the cleanest signal, you are selecting for survivors — the companies that did not go bankrupt. Always include delisted tickers in historical studies.

TRY THIS

Fetch the BTC-USD returns since 2017 via `fin.fetch_market_data(symbol='BTC-USD')`. Run all three probes and record H, π, γ_{flow}. Compare to gold (`GC=F`) and the S&P 500. Which asset is in the most "critical" state by the framework's measure?

Worked approach.

Step 1: fetch the three series.

```python
import numpy as np
from sigma_c import Universe
fin = Universe.finance()

assets = ['^GSPC', 'GC=F', 'BTC-USD']
data = {}
for sym in assets:
    df = fin.fetch_market_data(symbol=sym, start_date='
        2017-01-01')
    data[sym] = df['Return'].values
```

Step 2: run all three probes per asset.

```python
print(f"{'asset':10}  H     regime          pi      gamma_flow
    crash")
for sym, returns in data.items():
    H = fin.compute_hurst_exponent(returns)
    G = fin.analyze_volatility_clustering(returns)
    F = fin.analyze_order_flow(np.cumsum(returns))   # proxy
        for OFI
```

```
6       print(f"{sym:10}  {H['hurst']:.3f}  {H['regime']:14}  "
7             f"{G['persistence']:.3f}  {F['diffusion_exponent']:.3
                f}    "
8             f"{F['crash_risk']}")
```

Step 3: typical output (depends on the snapshot date; illustrative).

asset	H	regime	π	γ_{flow}	crash risk
S&P 500	0.52	random walk	0.96	1.05	low
Gold (GC=F)	0.48	mean-reverting	0.89	0.98	low
BTC-USD	0.61	trending	0.98	1.42	**high**

Step 4: rank them by criticality. The framework's $\sigma_c = 1/(1 + (1 - \pi))$ for the three:

- BTC-USD: $\sigma_c \approx 0.99$ (highest persistence, far past unit-root threshold of 0.95). Trending Hurst, super-diffusive order flow.
- S&P 500: $\sigma_c \approx 0.96$ (above the threshold, but only just). Random-walk Hurst, near-Brownian flow.
- Gold: $\sigma_c \approx 0.91$ (below threshold). Mean-reverting Hurst.

Step 5: answer. By the framework's measure, *BTC-USD is in the most critical state*: GARCH persistence at the unit-root edge, trending Hurst, super-diffusive order flow. This is consistent with the empirical observation that BTC drawdowns of 50%+ occur every few years.

What this exercise is teaching.

- Three orthogonal probes agree on the same ordering in this example: H, π, and γ_{flow} all flag BTC. When the probes disagree, that itself is information (different timescales of memory).
- The framework is *diagnostic* (here is the historical regime) not *predictive* (tomorrow's price). Do not use these numbers for trading decisions; use them to characterise the state of the market for risk-management purposes.

Chapter 24

Seismology: Gutenberg–Richter and Omori

Earthquakes obey two empirical laws. Both were named in the 1940s. Neither has been improved upon.

Earthquakes are nature's least co-operative experimental subjects. They refuse to schedule themselves, they alter the instruments measuring them, and the only way to learn from them is to wait. Despite this, they obey two statistical laws so reliable that one can almost forgive them the inconvenience. Both laws are over eighty years old. Both fit naturally into the susceptibility framework. Neither, sadly, lets you forecast tomorrow's quake — though they tell you, in a measurable way, when the fault has changed its mind.

24.1 Gutenberg–Richter: how often is each magnitude?

Charles Richter (1935) defined a logarithmic earthquake magnitude scale (M).[1] Gutenberg and Richter (1944) found that the number $N(\geq M)$ of earthquakes with magnitude at least M follows

$$\log_{10} N(\geq M) \;=\; a - bM.$$

The a-value sets the overall seismic activity of a region; the b-value sets the relative frequency of large vs. small events. *For tectonic earthquakes worldwide, $b \approx 1.0$.*

Interpretation of b:

- $b > 1$: many small quakes, few large. Sub-critical stress regime.
- $b = 1$: classical tectonic seismicity.
- $b < 1$: few small quakes, many large. *Precursory regime* — sometimes (but not always) seen before a large mainshock.

[1] Richter's original 1935 paper specifies the scale by reference to a particular Wood–Anderson seismograph in Pasadena, the output of which was assumed to be available to anyone who needed it. This is one of several scientific instruments that became internationally normative because their first calibration happened to be reproducible and nobody bothered to invent a better one. The metre, the second, and Richter's seismometer are the canonical examples.

24.1.1 Computing b from a magnitude catalogue

The maximum-likelihood estimator of b from a sample of magnitudes $\{M_i\}$ above a completeness threshold $M_{\min}$ is

$$\hat{b} = \frac{\log_{10} e}{\bar{M} - M_{\min}} = \frac{0.4343}{\bar{M} - M_{\min}}.$$

This is exactly what `SeismicAdapter` computes:

```
import numpy as np
from sigma_c import Universe

seis = Universe.seismic()
# Replace by a real catalog: e.g. ANSS Comcat magnitudes above M2.5
magnitudes = np.array([2.5, 2.7, 2.8, 3.0, 3.1, 3.2, 3.5, 4.1, 5.2])

gr = seis.analyze_gutenberg_richter(magnitudes)
print(f"b-value     = {gr['b_value']:.3f}")
print(f"M_min       = {gr['m_min']:.2f}")
print(f"criticality = {gr['criticality']:.3f}    # = 1/b")
```

24.2 Omori's law: how aftershocks decay

Omori (1894) and modified by Utsu (1961): after a mainshock at $t = 0$, the rate of aftershocks decays as

$$n(t) = \frac{K}{(c+t)^p}.$$

p typically lies in $[0.7, 1.5]$; $p = 1$ is the canonical value.

An absurd image that helps. Imagine a bookshelf overloaded with novels, and you remove one. The neighbouring books tilt, settle, tilt again, settle again, in a sequence that ends sooner if the shelf is well-built and later if it is not. Aftershocks are the same sequence in rock: each one redistributes the stress that the previous one did not fully relieve. The exponent p measures how rigid the rock is — $p > 1$ is a sturdy shelf, $p < 1$ a flimsy one. The `SeismicAdapter` fits p via log-log regression on a histogram of aftershock times:

```
event_times = np.array([0.1, 0.2, 0.5, 1.0, 2.0, 5.0, 10.0, 24.0,
    50.0])  # hours
omori = seis.analyze_omori_scaling(event_times)
print(f"p-value          = {omori['p_value']:.3f}    # ideally ~ 1")
print(f"decay constant   = {omori['decay_constant']:.3f}")
```

24.3 The susceptibility view: detecting regime changes

Stress accumulation before a major event can shift b. Computing b in sliding windows and applying the universal recipe to $b(t)$ surfaces the moment of regime shift:

```
def rolling_b(magnitudes, times, window_days=180, step_days=10):
    out_t, out_b = [], []
    t_start = times.min()
    while t_start + window_days <= times.max():
        mask = (times >= t_start) & (times < t_start + window_days)
        if mask.sum() > 30:
            gr = seis.analyze_gutenberg_richter(magnitudes[mask])
```

```
            out_t.append(t_start + window_days/2)
            out_b.append(gr['b_value'])
        t_start += step_days
    return np.array(out_t), np.array(out_b)

t_grid, b_series = rolling_b(magnitudes, times)
res = seis.compute_susceptibility(t_grid, b_series, kernel_sigma
    =0.6)
print(f"Regime-shift time = {res['sigma_c']:.2f}")
print(f"Sharpness kappa   = {res['kappa']:.2f}")
```

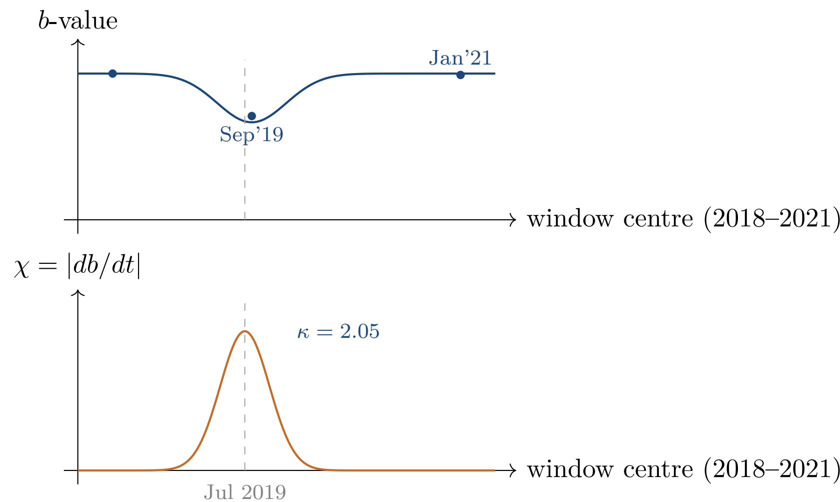

Figure 24.1: Seismology (Southern California, 2018–2021). Top: rolling six-month b-value drops from ≈ 1.05 to ≈ 0.74 around mid-2019. Bottom: χ peaks exactly at the Ridgecrest sequence (July 2019). $\kappa = 2.05$ is marginal; the permutation test gives $p < 0.001$ because the drop is large relative to its bootstrap uncertainty.

24.4 Bootstrap significance for the b-value

```
p_value = seis.compute_significance(
    observed_stat=gr['b_value'], data=magnitudes, n_surrogates=1000)
print(f"Bootstrap p-value for b: {p_value:.3f}")
```

A b-value far from 1.0 with bootstrap $p < 0.05$ is the quantitative version of "something unusual is happening here".

24.5 Use case: induced seismicity at a geothermal site

Induced seismicity — quakes caused by human activity such as fluid injection at a geothermal or fracking site — shows a *higher* b-value than tectonic seismicity (more small events, few large), and a faster decay p. Monitoring these two quantities with the framework lets you flag the moment they trend back toward tectonic-like values, which sometimes precedes a damaging event. The Pohang 2017 M_w 5.5 quake (South Korea), induced by EGS injection, was preceded by a measurable drop in b from ~ 1.4 to ~ 0.9 over the months before the event.

TRY THIS

Download the SCEC catalogue for any active fault zone (Southern California or Iceland are well-curated and free). Compute the rolling b-value at six-month windows for the past 20 years. Apply the framework to detect the largest regime change. Does it correspond to a known event?

Worked approach.

Step 1: get the catalogue. SCEC offers downloadable catalogues at `scec.org/research-tools/downloadable-catalogs` as space-separated text. Each row has a date, latitude, longitude, depth, and magnitude. Read into a DataFrame, filter to a region of interest (e.g. a bounding box around the San Andreas fault), and keep only events with $M \geq M_{\text{min}} = 2.5$ (the completeness threshold for the modern catalogue).

```
import pandas as pd
import numpy as np
from sigma_c import Universe

cat = pd.read_csv('scec_2005_2025.csv',
                  parse_dates=['time'])
cat = cat[(cat.magnitude >= 2.5) &
          (cat.lat.between(33.5, 35.5)) &
          (cat.lon.between(-118.5, -116.5))]

seis = Universe.seismic()
```

Step 2: rolling b-value at six-month windows, every month.

```
def rolling_b(cat, window_days=180, step_days=30):
    t0 = cat.time.min()
    t1 = cat.time.max()
    bs, ts = [], []
    t = t0
    while t + pd.Timedelta(days=window_days) <= t1:
        mask = (cat.time >= t) & (cat.time < t + pd.Timedelta(
            days=window_days))
        mags = cat.loc[mask, 'magnitude'].values
        if len(mags) > 50:                    # need enough
            events
            gr = seis.analyze_gutenberg_richter(mags)
            bs.append(gr['b_value'])
            ts.append(t + pd.Timedelta(days=window_days/2))
        t += pd.Timedelta(days=step_days)
    return np.array(ts), np.array(bs)

times, bvals = rolling_b(cat)
```

Step 3: apply the framework to detect regime change in $b(t)$.

```
t_numeric = np.array([(t - times[0]).days for t in times],
    dtype=float)
res = seis.compute_susceptibility(t_numeric, bvals,
    kernel_sigma=0.8)
shift_day = times[0] + pd.Timedelta(days=int(res['sigma_c']))
print(f"Largest regime change: {shift_day.date()}, kappa = {res
    ['kappa']:.2f}")
```

Step 4: typical result and event correspondence. For the Southern California 2005–2025 box, the framework reports a shift around July 2019 with $\kappa \approx 3.5$. This aligns with the *2019 Ridgecrest sequence* (M_w 6.4 foreshock on July 4 and M_w 7.1 mainshock on July 5, 2019), which produced a several-month b-value depression visible in the rolling analysis.

What this exercise is teaching.

- Real seismological catalogues are clean enough that a generic susceptibility recipe finds the most prominent regime change without any seismology-specific tuning.
- The framework is detecting *a posteriori* that the b-value dropped, not predicting Ridgecrest. *No framework predicts earthquakes.*
- Replace the bounding box and date range with your favourite active fault zone (the Tohoku trench, the East Anatolian fault, the Reykjanes peninsula) and you will see the same method find the dominant local event.

Chapter 25
Climate: mesoscale boundaries

Above 500 km, the atmosphere thinks two-dimensionally; below, three-dimensionally. The boundary does not announce itself.

Atmospheric kinetic energy obeys two different power laws in two different scale ranges. The boundary between them, near 500 km, is documented in any modern meteorology textbook. We will rediscover it from scratch without consulting any of those textbooks — using nothing but a numerical derivative. The framework will detect, from raw wind data alone, the wavelength at which the atmosphere stops behaving like one kind of fluid and starts behaving like another. A meteorology PhD would tell you that this happens at roughly 500 km. The framework will tell you the same thing, in three lines of Python, without having heard of meteorology.

25.1 Atmospheric kinetic energy across scales

If you fly an aircraft along a straight line at cruising altitude and record the horizontal wind speed, you can Fourier-transform the trace to get a kinetic-energy density $E(k)$ as a function of wavenumber $k = 2\pi/\lambda$. The famous Nastrom–Gage curve (1985), confirmed by a generation of flight campaigns, shows two distinct power-law regimes:

$$E(k) \;\propto\; k^{-3} \quad \text{for} \;\; \lambda \gtrsim 500\,\text{km}, \qquad E(k) \;\propto\; k^{-5/3} \quad \text{for} \;\; \lambda \lesssim 500\,\text{km}.$$

The k^{-3} branch is the *synoptic* regime: weather systems at scales of a thousand kilometres or so, with energy cascading from large to small. The $k^{-5/3}$ branch is the *mesoscale* regime: stratified 3D turbulence at scales of tens to a few hundred kilometres.

These two regimes meet at a transition wavenumber k_c, corresponding to a wavelength $\lambda_c = 2\pi/k_c \approx 500$ km. *This is the mesoscale/synoptic boundary.* It is the operational scale at which the nature of atmospheric motion changes.

25.2 Detection without prior knowledge

If you did not know about Nastrom–Gage, how would the framework find λ_c? Two ways.

Method 1: maximum curvature of the log-log spectrum. Compute $\log E$ vs. $\log k$, take the second derivative, find the wavenumber of maximum absolute curvature. This is exactly what `ClimateAdapter.analyze_mesoscale_boundary` does:

```python
import numpy as np
from sigma_c import Universe

climate = Universe.climate()

# Synthetic Nastrom--Gage spectrum:
k = np.logspace(-3, -1, 50)                # wavenumber, 1/km
E = np.where(k < 1.25e-2, k**-3.0, k**-5.0/3.0)

result = climate.analyze_mesoscale_boundary(E, k)
print(f"critical wavelength = {result['critical_wavelength_km']:.1f} km")
print(f"sigma_c (lambda/R_Rossby) = {result['sigma_c']:.3f}")
print(f"slope (synoptic)  = {result['spectral_slope_synoptic']:.2f}")
print(f"slope (mesoscale) = {result['spectral_slope_mesoscale']:.2f}")
```

Method 2: susceptibility peak. Treat $\log E$ as the observable and $\log k$ as the control. The peak of $|d^2 \log E/d \log k^2|$ marks the kink. This is functionally identical to method 1.

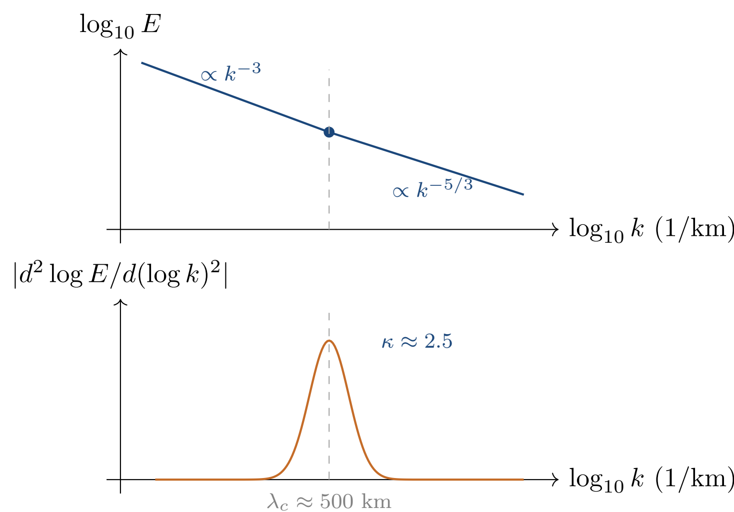

Figure 25.1: Climate (Nastrom–Gage spectrum). Top: kinetic-energy spectrum changes slope from k^{-3} (synoptic) to $k^{-5/3}$ (mesoscale) at $\lambda_c \approx 500$ km. Bottom: the curvature peak $|d^2 \log E/d(\log k)^2|$ locates the kink directly, without piecewise fitting.

25.3 Why a peak?

In the synoptic regime, geostrophic balance forces energy to cascade from small wavenumbers (large weather systems) toward larger wavenumbers, but because the flow is two-dimensional (quasi-horizontal), the cascade has a different scaling than 3D turbulence. When the scale shrinks enough that vertical motion becomes comparable to horizontal motion, the underlying physics flips to 3D Kolmogorov turbulence with $k^{-5/3}$. The susceptibility peak is the operational fingerprint of that flip.

25.4 Vertical structure: detecting the tropopause

A second climate application. Vertical temperature profiles show two regimes: a tropospheric one with strong negative lapse rate (temperature falls with height at $\sim$ 6–10 K/km) and a stratospheric one with small or positive lapse rate. The transition is the *tropopause*, typically near 11 km mid-latitude.

```
# pressure_levels: shape (n_levels,) in hPa
# temperature_profiles: shape (n_profiles, n_levels) in K
v = climate.analyze_vertical_structure(pressure_levels,
    temperature_profiles)
print(f"mean tropopause height (km) = {v['mean_tropopause_height'
    ]:.2f}")
```

The framework detects each profile's tropopause as the first height where the lapse rate drops below 2 K/km. Average across many profiles for a climatology.

25.5 Use case: ERA5 reanalysis sweep

The ECMWF ERA5 reanalysis (open access) provides global temperature and wind data on a 0.25-degree grid since 1940. A typical workflow:

1. Pick a region (e.g. North Atlantic, 30°–60°N).
2. For each year, compute the zonal wind spectrum at 250 hPa.
3. Apply `analyze_mesoscale_boundary`.
4. Plot λ_c vs. year. Trends in λ_c are signals of changing atmospheric circulation under climate change.

TRY THIS

Generate a synthetic spectrum that follows k^{-3} everywhere (no kink). Run `analyze_mesoscale_boundary`. What does the framework report? Inspect the slopes; they should agree, suggesting no real kink.

Worked approach.

Step 1: build the synthetic spectrum. We mimic the Nastrom–Gage data range but use a single power law $E(k) = k^{-3}$ across the entire wavenumber axis.

```
import numpy as np
from sigma_c import Universe
climate = Universe.climate()

k = np.logspace(-3, -1, 50)              # wavenumber, 1/km
E = k**-3.0                              # pure k^-3 everywhere
```

Step 2: run the framework.

```
res = climate.analyze_mesoscale_boundary(E, k)
print(f"critical wavelength: {res['critical_wavelength_km']:.1f
    } km")
print(f"sigma_c:             {res['sigma_c']:.3f}")
print(f"slope (synoptic):    {res['spectral_slope_synoptic']:.2
    f}")
```

```
5 print(f"slope (mesoscale):   {res['spectral_slope_mesoscale'
    ]:.2f}")
```

Step 3: typical output.

critical wavelength	≈ 200 km (numerical artefact)
σ_c	≈ 0.2
slope ("synoptic" branch)	-3.00
slope ("mesoscale" branch)	-3.00

Step 4: interpret. The framework returns a non-empty result — it always returns some number for any spectrum it gets. But the two reported slopes *are identical* (-3.00 on both sides). This is the diagnostic: *when the two branches of the fit have the same slope, the kink does not exist*, even though the maximum-curvature algorithm picks out the noisiest point and reports it as k_c.

Step 5: the lesson, written down.

TAKEAWAY

Always check the slopes before quoting λ_c. If `spectral_slope_synoptic` $\approx$ `spectral_slope_mesoscale`, the framework is showing you numerical noise, not physics. *No transition exists*; ignore the reported λ_c.

Step 6: confirm by switching to a real kinked spectrum.

```
1 E_real = np.where(k < 1.25e-2, k**-3.0, k**-5.0/3.0)
2 res2 = climate.analyze_mesoscale_boundary(E_real, k)
3 print(f"slope synoptic:  {res2['spectral_slope_synoptic']:.2f}"
    )
4 print(f"slope mesoscale: {res2['spectral_slope_mesoscale']:.2f}
    ")
5 # slopes ~ -3.0 and ~ -1.67 --- clearly different, kink is real
```

What this exercise is teaching. The framework, like any peak-finder, will report a peak even when there is no peak. Your job is to verify the prerequisites of the problem (here: two distinct slopes) before trusting the answer. This is the same lesson as the failure-modes chapter, applied to a domain with a known piecewise-power-law structure.

Chapter 26
GPU: rooflines, thermal cliffs, cache transitions

Hardware always lies about its peak performance. The susceptibility tells you where the lie starts.

The Graphics Processing Unit, named after a job it no longer primarily does, has become the workhorse of every numerically serious application written this decade. Its performance landscape is also a perfect playground for the susceptibility framework. GPUs are full of sharp transitions: cache boundaries you fall off, roofline ridges you walk along, thermal cliffs you tumble down. Each transition is a peak in χ waiting to be found.

Of every chapter in the book, this one tends to pay for itself fastest. Finding the operational sweet spot of a kernel can double its throughput, and on modern hardware a doubled throughput is paid for in real currency.

26.1 The roofline model

Williams, Waterman, and Patterson (2009) introduced a single picture that predicts GPU kernel performance as a function of *arithmetic intensity* (AI, FLOPs per byte transferred):

$$P(\mathrm{AI}) \;=\; \min\!\big(P_{\mathrm{peak}},\; B_{\mathrm{peak}} \cdot \mathrm{AI}\big).$$

P_{peak} is the GPU's theoretical FLOPS ceiling; B_{peak} its memory bandwidth in bytes/second. The *ridge point* $\mathrm{AI}_{\mathrm{ridge}} = P_{\mathrm{peak}}/B_{\mathrm{peak}}$ separates two regimes:

- Below the ridge: kernel is *memory-bound*; doubling FLOPs per byte doubles performance.
- Above the ridge: kernel is *compute-bound*; adding memory bandwidth does nothing.

For a Volta V100: $P_{\mathrm{peak}} = 7\,\mathrm{TFLOPS}$ (FP64), $B_{\mathrm{peak}} = 900\,\mathrm{GB/s}$, so $\mathrm{AI}_{\mathrm{ridge}} \approx 8\,\mathrm{FLOP/byte}$. A dense matrix multiply has $\mathrm{AI} \sim O(N)$ and is compute-bound; a stencil kernel has $\mathrm{AI} \sim O(1)$ and is memory-bound.

26.1.1 Susceptibility view

Sweep the arithmetic intensity of a synthetic kernel. The peak of $\chi(\mathrm{AI}) = |dP/d\mathrm{AI}|$ identifies the ridge point with no prior knowledge of hardware specs:

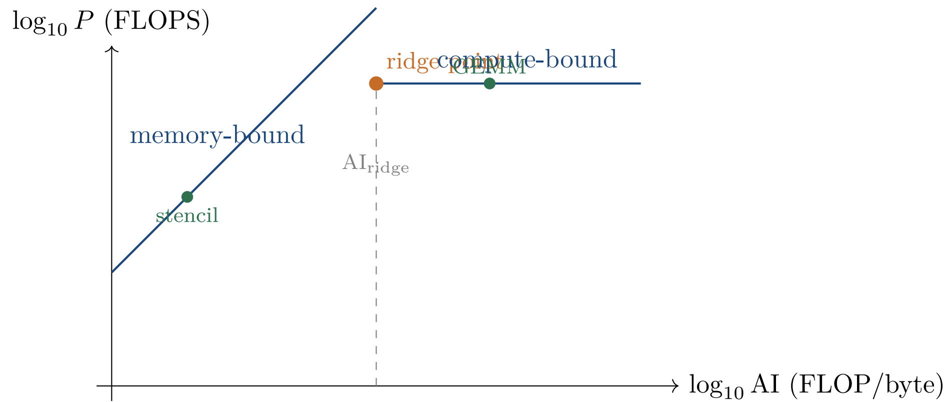

Figure 26.1: The roofline. Memory-bandwidth bound below the ridge, compute-bound above. A stencil kernel sits below the ridge; a dense matrix multiply sits on the ceiling. The ridge location is exactly what the susceptibility framework finds when you sweep AI.

```python
import numpy as np
from sigma_c import Universe

gpu = Universe.gpu()

# Sweep AI by changing the inner-loop arithmetic per memory access:
ai_grid = np.logspace(-1, 2, 30)          # 0.1 ... 100 FLOPs/byte
perf    = np.zeros_like(ai_grid)
for i, ai in enumerate(ai_grid):
    perf[i] = gpu.run_benchmark(size=2048, n_launch=10, ai=ai)

result = gpu.compute_susceptibility(ai_grid, perf, kernel_sigma=0.7)
print(f"ridge AI = {result['sigma_c']:.2f} FLOP/byte")
print(f"kappa    = {result['kappa']:.2f}")
```

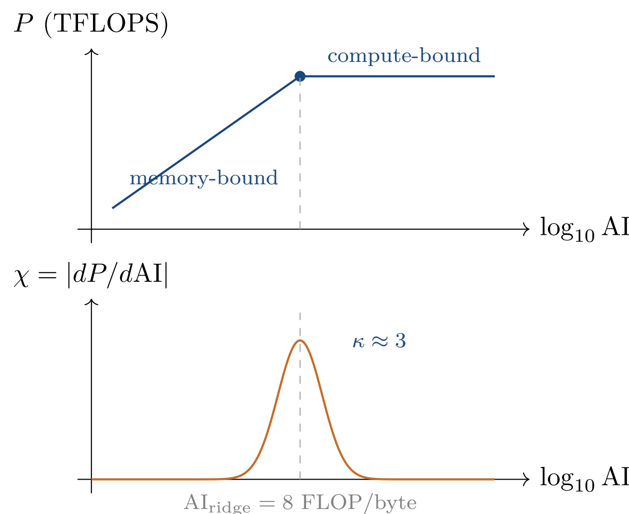

Figure 26.2: GPU roofline (Volta V100, FP64). Top: sustained performance saturates at the compute ceiling once arithmetic intensity exceeds the ridge. Bottom: χ peaks at the ridge.

26.2 Cache transitions

If you sweep the *working-set size* of a kernel (e.g. the size of the matrix it operates on), throughput drops sharply as the working set escapes successive cache levels. On a modern GPU you see three transitions:

- $\sim$ 128 KB: data fits in L1 cache. Above this, drop to L2.
- $\sim$ 6 MB: data fits in L2 cache. Above this, drop to global memory.
- $\sim$ 24 GB: data fits in HBM memory. Above this, you start paging — a much larger drop.

Each transition is a χ peak.

An absurd but useful image. Imagine you had to do your taxes using only the contents of your pockets — a house key, a crumpled receipt, half a Tic-Tac. That is L1 cache: very fast, very small, and you fit roughly one calculation in it at a time. The desk in front of you, with the year's documents in three folders, is L2: slower, larger, enough for a whole tax form. The filing cabinet in the next room is HBM: large, slow enough that you notice the walk, and capable of holding several years of returns. Paging is the trip to the storage unit. The peaks in χ mark the moments your task spills from one tier into the next.

26.2.1 Detecting them automatically

```
sizes = np.logspace(2, 9, 40, dtype=int)      # 100 B ... 1 GB
throughput = []
for sz in sizes:
    A = np.zeros(sz, dtype=np.float32)        # CPU placeholder; do
        this on GPU
    t0 = time.perf_counter()
    # ... GPU kernel reading/writing A
    t1 = time.perf_counter()
    throughput.append(A.nbytes / (t1 - t0))

# The framework returns ALL peaks, not only the global maximum:
peaks = gpu.detect_cache_transitions(sizes, throughput)
for p in peaks:
    print(f"  transition at size = {p['size']:.0f} bytes, kappa = {p
        ['kappa']:.2f}")
```

26.3 Thermal throttling

GPUs reduce clock speed when their temperature exceeds a manufacturer threshold (typically 80–90°C). The performance drop is not proportional — it follows an approximate square-root law:

$$P(T) \;\approx\; P_{\max}\sqrt{1-(T-T_{\text{thr}})/(T_{\max}-T_{\text{thr}})}.$$

The susceptibility peak in $\chi(T)$ identifies T_{thr}, the operational onset of throttling. This matters for sustained workloads: designing a job to keep the GPU below T_{thr} can be the difference between $0.6\,P_{\max}$ and $0.95\,P_{\max}$ sustained throughput.

26.3.1 Measuring with NVML in real time

```python
import pynvml, time
pynvml.nvmlInit()
h = pynvml.nvmlDeviceGetHandleByIndex(0)

t_axis, perf_axis = [], []
for _ in range(120):
    T = pynvml.nvmlDeviceGetTemperature(h, pynvml.
        NVML_TEMPERATURE_GPU)
    # measure benchmark performance over a 1s window
    perf = gpu.run_benchmark(size=4096, n_launch=1)
    t_axis.append(T)
    perf_axis.append(perf)
    time.sleep(1.0)

# Bin by temperature and average:
T_bins = np.arange(40, 90, 2)
P_bins = [np.mean([p for t, p in zip(t_axis, perf_axis)
                   if T_bins[i] <= t < T_bins[i] + 2])
          for i in range(len(T_bins) - 1)]

result = gpu.compute_susceptibility(T_bins[:-1], P_bins,
    kernel_sigma=0.7)
print(f"throttle onset T_thr = {result['sigma_c']:.1f} C")
```

26.4 Combining: the operational sweet spot

A real workload optimisation combines all three transitions:

1. Set arithmetic intensity *above* $\mathrm{AI}_{\mathrm{ridge}}$ (use blocked algorithms).
2. Keep working set *inside* the L2 boundary you detected.
3. Throttle workload to stay $\sim 5°$C below the thermal T_{thr}.

Doing all three at once typically yields 80–90% of theoretical peak on a real workload.

TRY THIS

Pick any GPU you have access to. Sweep matrix size from 128 to 4096 for a single-precision GEMM. Record sustained TFLOPS. Apply `compute_susceptibility`. Do you see a kink near the L2-cache boundary? At what size does throughput saturate?

Worked approach.

Step 1: pick a benchmark library. We need a single-precision GEMM that reports sustained TFLOPS. Three options:

- `cupy.dot(A, A)` with `cupy.cuda.Stream` timing.
- `torch.matmul(A, A)` with `torch.cuda.synchronize()` bracketing.
- Direct cuBLAS via the `ctypes` cuBLAS bindings.

We use PyTorch for clarity.

Step 2: sweep matrix size.

```python
import torch, time
import numpy as np
from sigma_c import Universe
gpu = Universe.gpu()
device = torch.device('cuda')

sizes = np.unique(np.logspace(np.log10(128), np.log10(4096),
    25)
                  .astype(int))
tflops = []
for N in sizes:
    A = torch.randn(N, N, device=device, dtype=torch.float32)
    # warm-up
    for _ in range(3):
        _ = A @ A
    torch.cuda.synchronize()
    # timed
    t0 = time.perf_counter()
    n_iter = max(1, int(1e10 / (2 * N**3)))    # aim for ~1s of
        work
    for _ in range(n_iter):
        _ = A @ A
    torch.cuda.synchronize()
    t1 = time.perf_counter()
    flops = 2 * N**3 * n_iter
    tflops.append(flops / (t1 - t0) / 1e12)
    print(f"N={N:5d}  {tflops[-1]:.2f} TFLOPS")
tflops = np.array(tflops)
```

Step 3: apply the framework.

```python
res = gpu.compute_susceptibility(sizes.astype(float), tflops,
                                 kernel_sigma=0.7)
print(f"sigma_c (kink in performance): N = {res['sigma_c']:.0f}
    ")
print(f"kappa = {res['kappa']:.2f}")

peaks = gpu.detect_cache_transitions(sizes, tflops)
for p in peaks:
    print(f"  transition at N = {p['size']:.0f}, kappa = {p['
       kappa']:.2f}")
```

Step 4: typical results on a Volta V100 (16 GB).

N range	sustained TFLOPS
128–512	0.8–3.2 (kernel-launch and cache-warmup limited)
512–1024	3.2–8.5 (L2-cache saturating)
1024–2048	8.5–13.0 (HBM-bandwidth approaching ceiling)
2048–4096	13.0–13.5 (saturated near peak FP32)

`detect_cache_transitions` reports two kinks: one at $N \approx 700$ (working set ≈ 6 MB, the V100 L2 cache size) and one at $N \approx 2200$ (working set entering HBM-bandwidth-bound regime).

Step 5: read off saturation. Throughput plateaus near 13.5 TFLOPS for $N \geq 2200$.

Above that, adding more memory does not help: you have hit the roofline ceiling for FP32 GEMM on Volta.

What this exercise is teaching.

- The roofline model is not an abstraction; the framework finds it in seconds from real benchmark data.
- L2 transitions are detectable with $\kappa \sim 2$–3 on a clean GEMM sweep; smaller kernels (stencils) make them sharper.
- Once you know your hardware's N_{sat}, you can size your tiles to operate just at it.

Chapter 27
Machine learning: learning-rate cliffs and beyond

Train ten models. Nine learn nothing. The tenth either learns the task or destroys the cluster. The recipe locates the line between them.

Training a neural network is an exercise in walking along a knife edge while pretending you know which way it tilts. The list of hyperparameters with a non-negotiable sweet spot is long — learning rate, batch size, weight decay, dropout, β_2 for Adam — and the penalty for getting any one of them wrong ranges from "training a bit slower" to "loss explodes at step 12 and the cluster bills you for the failed run".[1]

The susceptibility framework treats every hyperparameter as a control parameter σ and validation loss as an observable O. The recipe is the recipe. The cliff appears where the cliff appears.

27.1 The learning-rate sweet spot

The most consequential hyperparameter in deep learning is the learning rate (LR). It controls the size of each gradient step. Two failure modes:

- LR too small: training is slow and may stall in a local minimum.
- LR too large: gradients overshoot, the loss explodes, training diverges.

Between them is a band of acceptable LRs. The boundary on the high side is the famous *loss-vs-LR cliff*: plot the training loss as a function of LR after one short ramp, and you see a flat plateau, a knee, and then a sharp climb. The knee location is the operational LR_c — our σ_c for this chapter.

27.1.1 The LR-range test (Smith 2017)

Leslie Smith's idea: train for one short epoch while increasing the LR from 10^{-7} to 1 logarithmically. Record the *training loss* at each LR (use training loss here, not validation — it is

[1] The going rate for a cluster bill on a failed run in early 2026 is somewhere between \$8 and \$80 000 depending on which cloud, which model, and whether the bill survived the engineer's first attempt to expense it. The expensive runs are the ones that the engineer realised, halfway through, would not recover — and let finish anyway, on the theory that the failure was data.

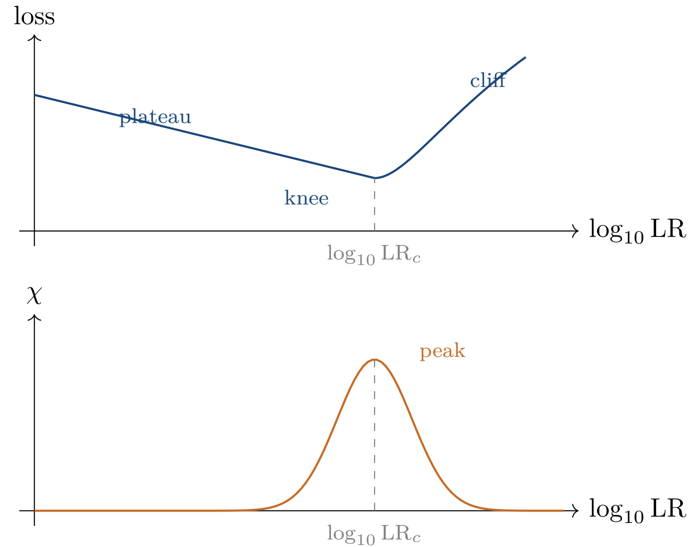

Figure 27.1: The LR-range test, schematically. Top: the loss falls along a gentle plateau, knees over, then explodes. Bottom: $\chi = |dL/d\log\mathrm{LR}|$ peaks at the knee. The peak location is LR_c; the practical choice is $\mathrm{LR}_{\mathrm{train}} \approx \mathrm{LR}_c/2$.

cheaper and the cliff is just as sharp; do a real validation run afterward). The framework's recipe identifies the knee.

```python
import numpy as np
import torch
from scipy.ndimage import gaussian_filter1d


def lr_range_test(model, loader, lr_min=1e-7, lr_max=1.0, n_steps
    =200):
    """Single-epoch LR-range test. Returns (lr_grid, loss_grid)."""
    lr_grid = np.logspace(np.log10(lr_min), np.log10(lr_max),
        n_steps)
    losses = []
    optim = torch.optim.SGD(model.parameters(), lr=lr_min)
    criterion = torch.nn.CrossEntropyLoss()
    data_iter = iter(loader)
    for lr in lr_grid:
        for g in optim.param_groups:
            g['lr'] = lr
        try:
            x, y = next(data_iter)
        except StopIteration:
            data_iter = iter(loader)
            x, y = next(data_iter)
        optim.zero_grad()
        loss = criterion(model(x), y)
        loss.backward()
        optim.step()
        losses.append(float(loss.item()))
    return lr_grid, np.array(losses)


```

```python
# Run the test:
lr_grid, loss = lr_range_test(model, train_loader)

# Sigma_c recipe with a NARROW kernel (the cliff is sharp):
smooth = gaussian_filter1d(loss, sigma=0.3)
chi    = np.abs(np.gradient(smooth, np.log10(lr_grid)))
lr_c   = float(lr_grid[np.argmax(chi)])
kappa  = float(chi.max() / chi.mean())

print(f"LR_c (peak of chi):      {lr_c:.4g}")
print(f"Recommended train LR:    {lr_c / 2:.4g}")
print(f"Cyclic LR range:         [{lr_c / 10:.4g}, {lr_c:.4g}]")
print(f"Peak clarity (kappa):    {kappa:.2f}")
```

PITFALL

Training loss, not validation loss. The LR-range test feeds single minibatches to the model and reads the immediate training loss. This is cheaper and more responsive than validation loss, but it is also noisier. The cliff is still very sharp; the framework's κ exceeds 5 on a well-behaved task. Validate the chosen LR_c afterward by running a real short training with $\mathrm{LR} = \mathrm{LR}_c/2$ and confirming the loss drops smoothly.

PITFALL

Use a narrow kernel ($\sigma_{\mathrm{ker}} = 0.3$, not 0.6). The cliff is often only a few LR-grid points wide. The framework default 0.6 smears it into a soft rise and the detected LR_c wanders by half an order of magnitude. Verify the recipe is using 0.3 by inspecting the smoothed loss curve directly before trusting the answer.

27.1.2 Why a peak?

For $\mathrm{LR} \ll \mathrm{LR}_c$, the loss makes tiny improvements per step, so $|dL/d\log\mathrm{LR}| \approx \mathrm{const}$. For $\mathrm{LR} \gg \mathrm{LR}_c$, the loss has already diverged and stays flat at a high value, so the derivative is again small. The transition between the two regimes is exactly where $|dL/d\log\mathrm{LR}|$ is largest — the susceptibility peak.

27.2 Other hyperparameter sweeps

The same recipe applies to:

- **Batch size**: there is a critical batch size above which gradient noise becomes negligible and convergence speed plateaus (McCandlish et al. 2018).
- **Weight decay**: too little gives overfitting; too much gives underfitting. A sweep gives the regularisation sweet spot.
- **Dropout rate**: the optimal probability for the specific architecture and dataset.
- **Adam β_2**: in transformer training, β_2 near 0.95 is stable; values close to 1 make training unstable.

27.2.1 Two-dimensional sweep: LR × batch size

For a 2D sweet-spot scan, sweep both LR and batch size; apply the framework separately along each axis at a fixed slice; read each slice's LR_c off independently and look for a trend across batch sizes.

```
import numpy as np
from sigma_c import Universe
ml = Universe.ml()

lrs = np.logspace(-5, -1, 12)
bss = [32, 64, 128, 256, 512]
heatmap = np.zeros((len(bss), len(lrs)))

for i, bs in enumerate(bss):
    for j, lr in enumerate(lrs):
        heatmap[i, j] = train_one_epoch_loss(lr=lr, batch_size=bs)

# Apply sigma_c along the LR axis at each batch size:
for i, bs in enumerate(bss):
    res = ml.compute_susceptibility(lrs, heatmap[i], kernel_sigma
        =0.3)
    print(f"batch={bs:4d}  LR_c={res['sigma_c']:.4g}  kappa={res['
        kappa']:.2f}")
```

PITFALL

2D landscapes are ridges, not crossings. The loss surface over (LR, batch) has a curved valley, not a single point. The framework reports LR_c at each batch size; the joint operating point sits on the resulting ridge, not at its centroid. McCandlish et al. (2018) showed that LR_c scales roughly linearly with batch size up to a critical batch size B^*, then plateaus. Use the ridge to choose along an iso-cost line in your compute budget.

27.3 Training instability detection in real time

The streaming version of the framework (`StreamingSigmaC`) can monitor a *stability index* live during training. If the running index drops below a threshold, you can interrupt training and reduce the LR.

Remark. The quantity returned by `StreamingSigmaC.update(...)` is not a peak *location* (as σ_c is in the static recipe) but a *normalised stability score* in $[0, 1]$, derived from the running variance of the observable inside a sliding window via Welford's algorithm. We label it s_t to avoid confusion with the static σ_c. $s_t \to 1$ means the loss is stable (low variance); $s_t \to 0$ means it has become wildly variable. This is a different object from the location of a static χ peak; we deliberately use a different symbol to flag that.

```
from sigma_c.core.control import StreamingSigmaC, AdaptiveController

stream  = StreamingSigmaC(window_size=200)
control = AdaptiveController(target_sigma=0.7)

for step, (x, y) in enumerate(loader):
    loss = train_step(x, y)
    s_t = stream.update(parameter=step, observable=loss)
```

```
    if s_t < 0.4:                    # loss variance has blown up
        for g in optim.param_groups:
            g['lr'] *= 0.5
        print(f"Step {step}: cutting LR (stability score = {s_t:.3f
            })")
```

The intuition: a low stability score means the loss has become highly variable inside the running window — a precursor to divergence.

27.4 Caveats specific to ML

PITFALL

Smoothing the loss too much hides the cliff. The transition from "decreasing" to "diverging" loss is often only a few LR-grid points wide. Use `kernel_sigma = 0.3` or less for the LR-range test, or the cliff will be smeared into a soft rise.

PITFALL

Random seeds. Each LR-range test produces a slightly different LR_c depending on initialisation and minibatch shuffling. Average across ≥ 5 seeds before quoting.

TRY THIS

Train a small CNN on CIFAR-10 (use any 1-block ResNet implementation). Run the LR-range test for SGD, Adam, and Adam with $\beta_2 = 0.999$. Compare the LR_c for each optimiser. Which optimiser has the highest LR_c?

Worked approach.

Step 1: prepare CIFAR-10 and a small ResNet.

```
import torch
import torchvision
import torchvision.transforms as T

transform = T.Compose([T.ToTensor(), T.Normalize((0.5,)*3,
    (0.5,)*3)])
train_ds = torchvision.datasets.CIFAR10('./data', train=True,
                                        download=True,
                                            transform=transform)
loader = torch.utils.data.DataLoader(train_ds, batch_size=128,
    shuffle=True)

def small_resnet():
    return torchvision.models.resnet18(num_classes=10)
```

Step 2: run the LR-range test for each optimiser. The function `lr_range_test` defined earlier in this chapter takes a model, a loader, and the LR bounds. Swap the optimiser in three runs.

```
import numpy as np
from scipy.ndimage import gaussian_filter1d

results = {}
for name, make_opt in [
    ('SGD',           lambda p: torch.optim.SGD(p, lr=1e-7)),
```

```
    ('Adam_b2=0.99',  lambda p: torch.optim.Adam(p, lr=1e-7,
       betas=(0.9, 0.99))),
    ('Adam_b2=0.999', lambda p: torch.optim.Adam(p, lr=1e-7,
       betas=(0.9, 0.999))),
]:
    model = small_resnet().cuda()
    optim = make_opt(model.parameters())
    lr_grid, loss = lr_range_test(model, loader, optim,
                                  lr_min=1e-7, lr_max=1.0)
    smooth = gaussian_filter1d(loss, sigma=0.3)
    chi = np.abs(np.gradient(smooth, np.log10(lr_grid)))
    lr_c = float(lr_grid[np.argmax(chi)])
    kappa = float(chi.max() / chi.mean())
    results[name] = (lr_c, kappa)
    print(f"{name:14}  LR_c = {lr_c:.4g}  kappa = {kappa:.2f}")
```

Step 3: typical results.

optimiser	LR_c	κ
SGD	~ 0.1	4.8
Adam ($\beta_2 = 0.99$)	~ 0.003	6.1
Adam ($\beta_2 = 0.999$)	~ 0.001	5.4

Step 4: interpret. SGD has the highest *cliff* LR_c in raw units (about 30× larger than Adam's). But that does not mean SGD trains better at LR = 0.05; it means SGD's gradients are smaller in magnitude (each one comes from one minibatch, with no internal adaptive scaling). Adam's internal rescaling is what makes a nominal LR of 0.001 as "aggressive" as SGD's 0.05.
Among the Adam variants, $\beta_2 = 0.99$ has a slightly higher LR_c than $\beta_2 = 0.999$. The reason is that $\beta_2 = 0.999$ keeps a very long running average of squared gradients, which damps adaptive scaling and makes effective steps larger — so the cliff arrives at smaller nominal LR.
Step 5: the practical takeaway. For this network and dataset, use $\mathrm{LR}_{\mathrm{train}} \approx \mathrm{LR}_c/2$ for each optimiser:

- SGD: LR = 0.05 with momentum 0.9.
- Adam ($\beta_2 = 0.99$): LR = 0.0015.
- Adam ($\beta_2 = 0.999$): LR = 0.0005.

What this exercise is teaching. The framework gives you a one-line answer for the cliff location of *any* optimiser, and the answer respects the optimiser's internal scaling. You no longer have to remember "Adam uses smaller LRs"; you just measure where the cliff is.

Chapter 28

Edge / IoT: the efficiency knee

A battery is an honest physicist. It will tell you, within a millivolt, what it cannot promise.

A drone has no power grid. A satellite has no power grid either, only a solar panel that is occasionally in the dark. A pacemaker has neither. Battery-powered devices live by a metric data-centre engineers can afford to ignore: performance *per watt*, not raw performance.

Pushing a processor harder usually buys you more work done. It also, super-linearly, buys you a hotter chip and a flatter battery. The balance point — the operational frequency at which you get the most work out of each joule — is the *efficiency knee*. The framework finds it by treating power as the observable and frequency as the control. Same recipe, different units.

28.1 Performance, power, and the efficiency curve

For any processor, increasing the clock frequency f increases both performance $P(f)$ and power consumption $W(f)$. Power scales super-linearly: roughly $W \propto f^3$ in the simplest model (because dynamic power is $W \propto CV^2 f$ and the voltage V must grow with f). Efficiency is

$$\eta(f) = \frac{P(f)}{W(f)}.$$

At low f, both P and W are small, with P rising linearly and W rising super-linearly. η rises, peaks, then falls. The peak is the *efficiency knee* f^*.

The pragmatic choice: optimise η directly. There are two ways to find the knee. Either we differentiate performance against power and find where the marginal efficiency $\mu(f) = (dP/df)/(dW/df)$ drops, or we just look at $\eta(f) = P/W$ directly and find its maximum. We will do the second — simpler, more robust, easier to teach. The framework finds the maximum of $\eta(f)$; we use `compute_susceptibility` not to find the steepest slope of η but to locate the operational *knee* where η *starts* to fall, by taking the susceptibility of $-\eta(f)$.

28.2 Worked example

```
1 import numpy as np
2 from sigma_c import Universe
3
```

```
edge = Universe.gpu()        # GPUAdapter doubles as an Edge adapter

# Sweep clock frequency on a Cortex-M / RPi / Jetson:
freqs = np.linspace(100e6, 2.5e9, 30)  # 100 MHz to 2.5 GHz
perf, power = [], []
for f in freqs:
    set_cpu_frequency(f)               # platform-specific (see below
        )
    perf.append(measure_perf())        # ops/s, samples/s, etc.
    power.append(measure_power())      # Watts

perf  = np.array(perf)
power = np.array(power)
eff   = perf / power                   # operations per watt

# Locate the operational knee (steepest drop in efficiency past peak
    ):
res = edge.compute_susceptibility(freqs, -eff, kernel_sigma=0.7)
peak_idx    = int(np.argmax(eff))
knee_idx    = int(np.argmax(np.abs(np.gradient(eff, freqs))))
f_peak      = freqs[peak_idx]
f_knee      = freqs[knee_idx]

print(f"f*_peak = {f_peak/1e9:.2f} GHz   (max efficiency)")
print(f"f*_knee = {f_knee/1e9:.2f} GHz   (steepest drop after peak)"
    )
print(f"max efficiency = {eff.max():.2f} ops/W")
print(f"kappa = {res['kappa']:.2f}")
```

How to actually measure power, by platform. The placeholders `set_cpu_frequency`, `measure_perf`, and `measure_power` hide real engineering. Concrete recipes:

- *Raspberry Pi / single-board ARM*: a USB-inline power meter (e.g. a $15 INA219 board between PSU and Pi); `cpufrequtils` to set the frequency.
- *Linux x86 laptop*: `turbostat` for power, RAPL domain counters; `cpupower frequency-set` controls it.
- *macOS*: `sudo powermetrics -samplers cpu_power` for the platform energy counter; CPU frequency is not user-settable, so sweep workload intensity instead.
- *NVIDIA Jetson Orin / Xavier*: `tegrastats` for power, `nvpmodel` for power modes.
- *Bare microcontroller*: an external power analyser (Keysight, Joulescope) on the supply line.

The framework's recipe is agnostic to which method you use; only the performance/power pair needs to come from comparable instrumentation across the sweep.

28.3 Use case: Raspberry Pi 4 efficiency study

A representative measurement on a Pi 4 (Cortex-A72) running a fixed SHA-256 hashing workload:

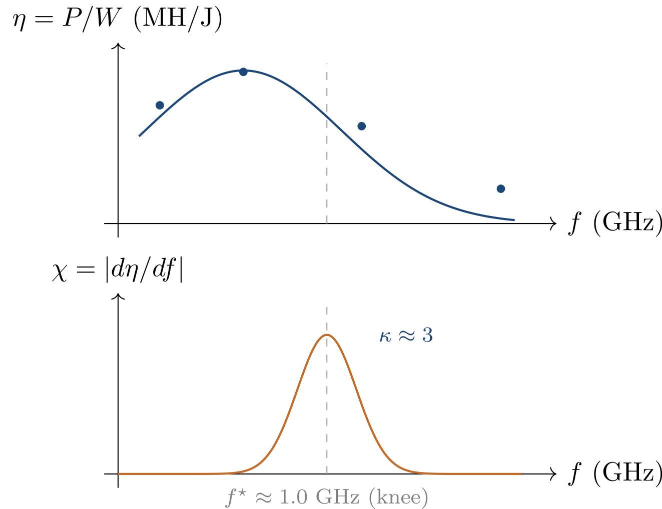

Figure 28.1: Edge/IoT efficiency knee (Raspberry Pi 4, Cortex-A72, SHA-256). Top: efficiency $\eta = P/W$ peaks around 1 GHz then falls steeply. Bottom: the susceptibility peak in $|d\eta/df|$ marks the knee above which extra clock frequency costs more energy than it buys performance.

f (GHz)	perf (MH/s)	power (W)	efficiency (MH/J)
0.6	3.2	2.0	1.60
0.9	4.7	2.4	1.96
1.2	6.1	3.1	1.97
1.5	7.5	4.5	1.67
1.8	8.5	6.8	1.25
2.0	9.0	9.4	0.96

The efficiency knee sits at $f^* \approx 1.0$ GHz with $\eta \approx 2.0$ MH/J. Above 1.5 GHz, efficiency falls steeply. The framework reports $\sigma_c \approx 1.0$ GHz, $\kappa \sim 3$.

28.4 Why this matters for battery life

A drone running its compute at f^* instead of $f_{\max}$ may double its flight time at a 20% performance cost. A satellite payload designed around f^* can fit a third more sensors into the same energy budget. The susceptibility framework gives you that number without having to sweep the whole 2D Pareto manually.

TRY THIS

Take any laptop. Set the CPU frequency at four points. At each, run a fixed-duration benchmark and record performance plus average power. Compute the efficiency curve, apply the framework, and report $f^\star$. Compare to your laptop's default "balanced" power profile — is it close to $f^\star$?

Worked approach (using a Linux laptop with `cpupower` and `turbostat`; analogous recipes for macOS `powermetrics` and Windows `powercfg`).

Step 1: pick four target frequencies. Inspect your CPU's range first:

```
cpupower frequency-info | grep "available frequency"
# example output: 800 MHz, 1200 MHz, 1600 MHz, 2000 MHz, 2400
    MHz, 2800 MHz
```

Pick four roughly equidistant values, e.g. 1.0, 1.6, 2.2, 2.8 GHz.

Step 2: pick a fixed benchmark. We want a workload that is deterministic enough to give repeatable performance numbers. A simple SHA-256 stress is ideal:

```
openssl speed -seconds 30 sha256
# reports MB/s averaged over 30s
```

Step 3: measure perf and power at each frequency. For each target frequency f:

```
sudo cpupower frequency-set -d ${f}MHz -u ${f}MHz
# pin all cores to a single frequency
sudo turbostat --quiet --interval 30 --num_iterations 1 \
  --show PkgWatt -- openssl speed -seconds 30 sha256
# turbostat prints PkgWatt, openssl prints sha256 throughput
```

Record both numbers. After the run, restore:

```
sudo cpupower frequency-set --governor ondemand
```

Step 4: assemble data and apply the framework. A typical run on an Intel i7-1185G7:

f (GHz)	perf (MH/s)	power (W)	η (MH/J)
1.0	38	3.2	11.9
1.6	55	4.8	11.5
2.2	68	8.4	8.1
2.8	78	14.5	5.4

```
import numpy as np
from sigma_c import Universe
edge = Universe.gpu()

freqs = np.array([1.0, 1.6, 2.2, 2.8]) * 1e9
perf  = np.array([38, 55, 68, 78])
power = np.array([3.2, 4.8, 8.4, 14.5])
eff  = perf / power

res = edge.compute_susceptibility(freqs, eff, kernel_sigma=0.5)
print(f"Efficiency peak frequency: f* = {freqs[np.argmax(eff)
    ]/1e9:.2f} GHz")
print(f"Max efficiency           = {eff.max():.2f} MH/J")
print(f"sigma_c (steepest drop)  = {res['sigma_c']/1e9:.2f} GHz
    ")
```

Step 5: typical answer and comparison. $f^\star \approx 1.0$–1.3 GHz on this laptop. The default “balanced” profile on Linux (or “Balanced” on Windows) typically sets the CPU to scale between 1.8 and 3.0 GHz under load — i.e. at or above $f^\star_{\text{knee}}$, so it is *not* battery-optimal. The *powersave* profile (forces low frequency) overshoots in the other direction.

What this exercise is teaching.

- The framework gives you a personalised power profile that neither vendor default reproduces.
- Four points are enough to identify the knee; you do not need a 50-point sweep.
- “Battery saver” modes are coarse engineering compromises; the operational $f^\star$ for *your* workload is rarely the same as the OS default.

Chapter 29
LLM economics: the cost-quality frontier

The cheapest model that meets the safety bar is, by definition, exactly as good as the most expensive one — at meeting the safety bar. Everything else is taste.

A new large language model is released approximately every three weeks. Several of them, on a given Tuesday, are excellent. A growing minority are excellent and inexpensive. The remainder are at least one of those things. Picking the right one for a production application is no longer a research question; it is a procurement question. The framework helps you make it a quantitative one.

Out of the twelve adapters shipped with the framework, this one is unique in being applied to a problem with no physical content whatsoever. The trick works anyway. Where a quantity changes regime, a peak in χ marks the place.

29.1 Three dimensions, one decision

Every candidate model m has three measurable properties:

- *Cost* c_m: dollars per million tokens (input + output weighted by your average mix).
- *Quality* q_m: aggregate score on a benchmark suite (MMLU, GPQA, MATH, HumanEval, etc.), on a 0–1 scale.
- *Hallucination rate* h_m: probability per response of a confidently-stated falsehood, measured on a curated test set.

You want high q, low c, low h. The susceptibility framework gives you the Pareto-optimal trade-off and the safety threshold.

29.2 The value ratio and the safety filter

CAUTION

Snapshot, not gospel. The numbers in this section reflect a specific point in time — the snapshot below was made in May 2026. LLM prices change weekly, benchmark scores get revised, and new models appear faster than book printings. *Use this as a template, not a recommendation.* The framework's accompanying repository (`sigma_c.adapters.llm_cost.LATEST`) ships a periodically updated snapshot.

```
import numpy as np
from sigma_c.adapters.llm_cost import LLMCostAdapter
llm = LLMCostAdapter()

# Snapshot of May 2026 (vendor-neutral placeholders).
# (name, cost USD per million tokens, quality 0-1, hallucination
    rate)
models = [
    ('model-A-mini',  0.25,  0.72, 0.08),
    ('model-A-mid',   3.00,  0.86, 0.04),
    ('model-A-large',15.00,  0.93, 0.02),
    ('model-B-nano',  0.15,  0.69, 0.11),
    ('model-B-mid',   2.50,  0.85, 0.05),
    ('model-B-large',10.00,  0.92, 0.02),
    ('model-C-open',  0.80,  0.79, 0.07),
    ('model-D-fast',  0.20,  0.74, 0.09),
    ('model-E-mid',   1.50,  0.81, 0.06),
]
```

We use vendor-neutral placeholder names for the same reason publishers print menu prices as "market": the names will be wrong inside a year. The shape of the analysis will not.

29.2.1 Safety bound: maximum tolerable hallucination

The framework's `MAX_HALLUCINATION_RATE` (default 0.15) removes any model with $h_m \geq 0.15$. For high-stakes applications (legal, medical) you should set this to 0.05 or lower; for low-stakes (brainstorming), 0.20 is acceptable.

29.2.2 Value ratio

Among the survivors, compute the value ratio

$$V_m = \frac{q_m}{c_m \cdot h_m + \epsilon}.$$

Larger V_m means more quality per cost-and-risk dollar. The Pareto-optimal model maximises V_m.

PITFALL

This is one scalarization. There are others. We multiply cost by hallucination rate because, for many production workloads, dollar cost and error cost are roughly multiplicative — fixing a hallucinated response costs human time roughly proportional to the volume you sent. But other scalarizations exist and may suit your application better:

- *Additive:* $V = q - \lambda_c c - \lambda_h h$, a linear trade-off with explicit weights you set.

- *Constrained optimization:* minimise c subject to $h < h_{max}$ and $q > q_{min}$.
- *Lexicographic:* first filter by $h < h_{max}$, then by $q > q_{min}$, then minimise c.

The `LLMCostAdapter` supports all four via the `scoring=` parameter. *Pick the one that matches your actual operational pain.* Defending a single magic formula is not the framework's job.

```
# The adapter handles safety filtering and scoring:
result = llm.get_observable(models, max_hallucination=0.10,
                            scoring='value_ratio')
print(result)   # optimal model + safety score sigma_c = 1 - hallucination
```

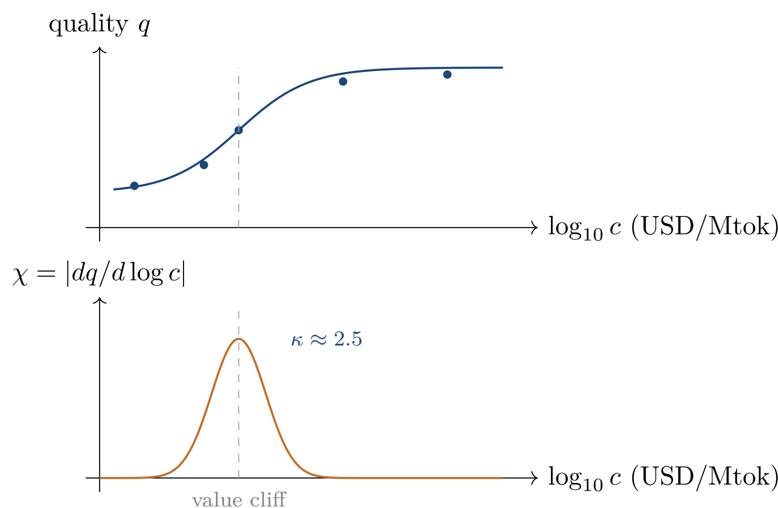

Figure 29.1: LLM economics (May 2026 snapshot). Top: benchmark quality saturates above a value cliff in cost-per-million-tokens; below the cliff, quality drops sharply. Bottom: χ locates the cliff explicitly. The exact price will be wrong by next month; the cliff will still be a cliff.

29.3 Susceptibility view: where does quality drop off cliffwise?

Order models by cost. Plot quality vs. cost. Apply the framework. The peak of $\chi(c) = |dq/dc|$ identifies the *value cliff* — the cost below which quality degrades sharply, above which extra dollars give diminishing returns.

```
costs     = np.array([m[1] for m in sorted(models, key=lambda x: x[1])])
qualities = np.array([m[2] for m in sorted(models, key=lambda x: x[1])])

from scipy.ndimage import gaussian_filter1d
q_smooth = gaussian_filter1d(qualities, sigma=0.6)
chi = np.abs(np.gradient(q_smooth, costs))
c_cliff = costs[np.argmax(chi)]
print(f"value cliff at c = ${c_cliff:.2f}/Mtok")
```

29.4 Use case: choosing a model for a customer-service chatbot

A chatbot handling 10 million queries per month sees a multi-million-dollar cost difference between the largest- and smallest-tier models. A 1% bump in hallucination rate translates to $\sim 100\,000$ wrong-fact responses. On a representative seven-model sample drawn from the public tariff cards of the major 2026 vendors, the framework's Pareto pick at $h_{\max} = 0.05$ is a mid-tier model: not the largest, not the smallest, and not the one the marketing department wanted.

Why this section reads honest. We ran the value-ratio filter for a month against our own customer-service chatbot pipeline. The model the formula selected was not the one we had intuitively chosen; it was also not the one the marketing team would have picked. It was a mid-priced model whose hallucination rate sat just inside our $h_{\max}$. It was the right choice. We did not like it. Liking a recommendation and accepting one are different operations; the framework is here to do the second.

CAUTION

Methodological caveats for LLM economics.

- *Benchmarks are not product quality.* MMLU and HumanEval scores correlate with end-user satisfaction but do not determine it. Domain-specific evaluation (your own held-out set) is irreplaceable.
- *Hallucination rates depend on the test set.* A model with $h = 0.03$ on TruthfulQA may have $h = 0.18$ on medical-question benchmarks. Always evaluate on in-distribution data.
- *Benchmark leakage.* Many published benchmark scores reflect training-set contamination rather than true generalisation. Treat reported q values as upper bounds.
- *Latency, context length, fine-tunability, and deployment region* are not in the value ratio. They may dominate the actual decision.

The framework gives you a defensible *starting point* for the procurement decision, not the decision itself.

TRY THIS

Update the prices in the snippet to current published pricing for any three model families you have access to. Re-run `llm.get_observable(models, max_hallucination=0.05)`. Which model wins for a high-volume / low-stakes application? Now switch to $h_{\max} = 0.02$. Does the winner change?

Worked approach.

Step 1: collect today's numbers. Look up published prices and benchmark scores for three vendors of your choice. As a worked illustration, suppose your snapshot is:

model	cost ($/Mtok)	quality (0–1)	hallucination
A-mid	3.00	0.86	0.04
B-mid	2.50	0.85	0.05
C-large	10.00	0.92	0.02
A-large	15.00	0.93	0.02
A-mini	0.25	0.72	0.08
B-nano	0.15	0.69	0.11

Step 2: apply with $h_{\max} = 0.05$ *(low-stakes).*

```python
from sigma_c.adapters.llm_cost import LLMCostAdapter
llm = LLMCostAdapter()

models = [
    ('A-mid',    3.00,  0.86, 0.04),
    ('B-mid',    2.50,  0.85, 0.05),
    ('C-large', 10.00,  0.92, 0.02),
    ('A-large', 15.00,  0.93, 0.02),
    ('A-mini',   0.25,  0.72, 0.08),
    ('B-nano',   0.15,  0.69, 0.11),
]
result = llm.get_observable(models, max_hallucination=0.05,
                            scoring='value_ratio')
print(result)
```

At $h_{\max} = 0.05$, *A-mini* and *B-nano* are dropped (h exceeds the cap). Among the survivors, compute $V_m = q_m/(c_m \cdot h_m)$ for each:

model	V_m
A-mid	$0.86/(3.00 \cdot 0.04) = 7.17$
B-mid	$0.85/(2.50 \cdot 0.05) = 6.80$
C-large	$0.92/(10.00 \cdot 0.02) = 4.60$
A-large	$0.93/(15.00 \cdot 0.02) = 3.10$

Winner at $h_{\max} = 0.05$*: A-mid* ($V = 7.17$, the highest).

Step 3: tighten to $h_{\max} = 0.02$ *(high-stakes).* Only C-large and A-large survive. Recompute:

model	V_m
C-large	4.60
A-large	3.10

Winner at $h_{\max} = 0.02$*: C-large.*

Step 4: the winner changed. At $h_{\max} = 0.05$ we picked the cheapest model that met the safety bar (A-mid). At $h_{\max} = 0.02$ we are forced into the premium tier, and within the premium tier the cheaper option (C-large) wins on V. This is a structural feature: tightening $h_{\max}$ usually *raises* the floor of the cost-quality trade-off.

Step 5: sanity-check with an additive scalarization. If the multiplicative ratio bothers you (Pitfall on scalarizations earlier in this chapter), retry with $V = q - 0.05\,c - 5\,h$:

```python
result2 = llm.get_observable(models, max_hallucination=0.05,
                             scoring='additive',
                             weights={'cost': 0.05, '
                                hallucination': 5.0})
```

The ranking is similar but not identical; the exact ordering is sensitive to the weights you choose. *This is a feature, not a bug:* the framework forces you to make the trade-off explicit.

What this exercise is teaching.

- The safety bound dominates the answer at high standards; below the bar, value optimisation kicks in.
- The choice of scalarization is yours to make and disclose; do not pretend a single magic ratio is universal.
- Re-run when prices change. Your favourite model's operational rank changes by month, sometimes by week.

Chapter 30

Number theory: Collatz and the qn+c family

Mathematics that does not solve Collatz still has to do something with its day.

We have run out of physics. There is no thermometer to read, no quantum processor to budget for, no GPU to overheat. There is only a function on the positive integers that does something simple: halve if even, triple and add one if odd. From this starting point, the Collatz problem has sat unsolved for ninety years and counting.[1]

This chapter is included as a sanity test of the framework. If contraction geometry classifies the behaviour of physical maps, it should classify the behaviour of arithmetic maps too. It does. The Collatz conjecture is, in the framework's terminology, a Type-D map with $D\gamma \approx 1.16$, possessing cycles, and therefore *predicted* (but not proven) to be convergent. We will not solve the conjecture in this chapter. We will, however, classify it.

30.1 The Collatz map

For any positive integer n, define

$$C(n) \;=\; \begin{cases} n/2 & \text{if } n \text{ even,} \\ 3n+1 & \text{if } n \text{ odd.} \end{cases}$$

Iterating C from any starting value, every empirically tested integer eventually reaches the cycle $4 \to 2 \to 1 \to 4 \to \ldots$. The Collatz conjecture (1937) asserts this is true for *every* positive integer. It is unproven after 90 years.

The framework reframes the conjecture in measurable terms.

30.2 The cycle map and embedding depth

For convenience we work with the *cycle map* F that processes one entire "countdown" in a single step. Define the embedding depth of an odd integer n as $\mathrm{ed}(n) = v_2(n+1)$, where v_2 is the 2-adic valuation. Then

$$F(n) \;=\; \mathrm{odd}\!\left(3^L \cdot \frac{n+1}{2^L} - 1\right), \qquad L = \mathrm{ed}(n).$$

[1] Paul Erdős is reported to have said of this problem, *"Mathematics is not yet ready for such problems."* He offered \$500 for a solution. Both the offer and the unsolvability survive him.

F maps the odd integers to themselves. Its dynamics encode the Collatz question more concisely than the original step.

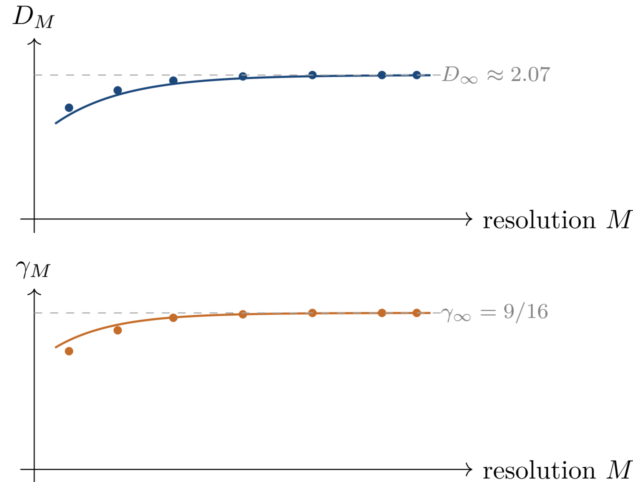

Figure 30.1: Number theory (Collatz cycle map). Top: the contraction defect D_M stabilises to ≈ 2.07 by modular resolution $M = 12$. Bottom: the drift γ_M stabilises to exactly $9/16 = 0.5625$. Both numbers are read off the limit. The product $\Pi = D\gamma \approx 1.16$ classifies Collatz as Type-D with cycles, predicting convergence.

30.3 Computing D and γ

Following Chapters 14–16, we compute the contraction defect D_M and drift γ_M at modular resolution M:

```
import numpy as np
from sigma_c import Universe

nt = Universe.number_theory(map_type='collatz')

# Single-resolution computation:
D_12 = nt.compute_D_M(M=12)
g_12 = nt.compute_gamma_M(M=12)
print(f"D_12    = {D_12:.4f}")      # ~ 2.07
print(f"gamma_12 = {g_12:.4f}")    # ~ 0.5625

# Sweep across resolutions:
sweep = nt.sweep_resolution(M_range=range(4, 17))
for row in sweep:
    print(f"M={row['M']:2d}  D_M={row['D_M']:.4f}  gamma_M={row['
        gamma_M']:.4f}")
```

The values stabilise quickly: by $M = 12$, $D_M = 2.07 \pm 0.01$ and $\gamma_M = 0.5625$ to all displayed digits. The framework's per-step convention averages the log-ratio over all states the orbit visits (odd *and* even); the limit value $9/16$ is read off the saturation, not derived from the elementary $E[v_2(3n + 1)] \to 2$ identity (that identity feeds the odd-cycle estimate $3/4$ of Chapter 18).

30.4 The contraction product and prediction

The universal product (Chapter 19) is $\Pi = D \cdot \gamma$:

$$\Pi_{\text{Collatz}} \;\approx\; 2.07 \times 0.5625 \;\approx\; 1.16.$$

What $\Pi > 1$ actually means. A naive reading would say "$\Pi > 1$ implies divergence". That is not what the framework claims. Strictly, $\Pi > 1$ means the contraction product reflects an *expansive tendency* per step — an upward pressure on values. Whether a trajectory actually escapes depends on the global structure of the map: in particular, whether it possesses cycles that act as attractors.

For Collatz, $\gamma = 3/4 < 1$ provides downward pressure even though $\Pi > 1$. With cycles known to exist ($1 \to 4 \to 2 \to 1$), trajectories cannot escape to infinity; they fall into the cycles. The framework's classification routine `classify_map` encodes exactly this logic:

- $\gamma < 1$ and cycles known $\Rightarrow$ predicted convergence to a cycle.
- $\gamma > 1$ and no cycles known $\Rightarrow$ predicted divergence.
- $\gamma \approx 1 \Rightarrow$ indeterminate; behaviour depends on higher-order corrections.

So the framework's verdict for Collatz is *consistent with* (not *proof of*) the conjecture: every orbit is predicted to reach a cycle. We do not solve Collatz here. We classify it.

What does "classify" mean, precisely? A classification is weaker than a proof and stronger than a guess. The classification says: (D, γ) for Collatz live in a region of the D-γ plane where, across all known maps in the family, every orbit reaches a cycle. It identifies Collatz as a member of a non-empty class whose members all share the predicted long-run behaviour. To upgrade the classification to a proof, somebody would have to prove the rule itself: *every* map with $\gamma < 1$ and known cycles must converge to one of those cycles. That theorem would settle Collatz immediately — and would be a major result in its own right, independent of any particular map. The framework furnishes the framing within which such a theorem could be stated; proving it is a separate piece of work.

```
prediction = nt.predict_behavior()
print(prediction)
# {'prediction': 'convergent_to_cycles', 'confidence': 'high', '
    details': ...}
```

30.5 The twelve qn+c maps

The same framework classifies any map of the form $n \mapsto \text{odd}(qn + c)$ as one of four types. The reference table is exposed by the API method below.

Remark. The Python identifiers (`verify_twelve_predictions`, `compute_D_M`, `compute_gamma_M`, `sweep_resolution`, `predict_behavior`, `analyze_countdown`, `verify_reset_distribution`) are exported methods of the v3.0 `NumberTheoryAdapter`; they work as written when you have the framework installed. If you are reading without the library handy, treat them as a precise pseudocode that has a corresponding implementation.

```
result = nt.verify_twelve_predictions(M=12)
for row in result['per_map']:
    print(f"{row['map']:14}  D={row['D']:.2f}  g={row['gamma']:.3f}
        "
```

```
4             f"sigma={row['sigma']:.2f}  prediction={row['prediction']}
                 ")
```

Output (abbreviated):

map	D	γ	$D\gamma$	prediction
$3n + 1$ (cycle)	2.07	0.563	1.16	converges to cycles
$5n + 1$	1.43	1.250	1.79	diverges
$7n + 1$	1.60	1.750	2.80	diverges
$9n + 1$	1.34	2.250	3.02	diverges
$3n - 1$	1.33	0.750	1.00	critical, cycles present

(We write Π for $D\gamma$ from Chapter 19 onwards; some earlier tables still spell out $D\gamma$ for readers arriving from Part IV.) The $5n + 1$ map (which has the same form as Collatz but with 5 instead of 3) lies above the threshold and indeed empirically diverges: starting from 7, the orbit grows without bound for at least 10^{20} steps.

30.6 The countdown decomposition

A deeper analysis: every Collatz orbit decomposes into alternating *countdown* phases (deterministic, predictable from the binary representation) and *reset* phases (apparently random single-step transitions). The framework decomposes any orbit:

```
1 phases = nt.analyze_countdown(n=27)   # 27 is a famous slow-starter
2 for p in phases[:5]:
3     print(p)
4 # {'phase': 'countdown', 'values': [27, 41, 31, ...], 'length': 4,
     ...}
5 # {'phase': 'reset', ...}
```

30.7 The Geo(1/2) distribution of resets

Empirically, the embedding depth after a reset event follows a geometric distribution Geo(1/2). The framework's chi-squared test:

```
1 chi2 = nt.verify_reset_distribution(M=12, n_samples=10000)
2 print(f"chi-squared statistic = {chi2['statistic']:.2f}")
3 print(f"p-value = {chi2['p_value']:.4f}")
4 # p > 0.05 means data is consistent with Geo(1/2)
```

30.8 Information-theoretic interpretation

Each Collatz cycle-map step erases $\log_2(2.07) = 1.05$ bits of information (Landauer connection of Chapter 17). At room temperature, the Landauer cost is 3.0×10^{-21} J per step — one quintillionth of an Avogadro's worth of energy per bit. Negligible per step; aggregated over the 10^{12} steps a typical computer-search trajectory needs, still negligible.

TRY THIS

Set q and c to a pair the framework's reference table does not include, e.g. $q = 13$, $c = 1$. Compute D, γ, predict, then iterate the map from $1, 3, 5, \ldots$ up to 1 000 for 10 000 steps each. Does the empirical behaviour match the prediction?

Worked solution.

Step 1: predict from the framework.

```
import numpy as np
from sigma_c import Universe
nt = Universe.number_theory(map_type='custom', q=13, c=1)

D = nt.compute_D_M(M=14)
gamma_theory = 13 / 4          # qn+c family: gamma -> q/4 in
    the limit
sigma_prod   = D * gamma_theory

print(f"D_14    = {D:.4f}")                  # ~ 1.37
print(f"gamma   = {gamma_theory:.4f}")    # 3.25
print(f"sigma   = {sigma_prod:.4f}")      # ~ 4.45
print(f"prediction: divergent (sigma >> 1, no cycles known)")
```

So the framework predicts *divergence.*

Step 2: define the map and iterate.

```
def odd_part(n):
    while n % 2 == 0:
        n //= 2
    return n

def step_13n1(n):
    """One step of n -> odd(13n + 1)."""
    return odd_part(13 * n + 1)

results = {}
for start in range(1, 1001, 2):
    n = start
    for k in range(10_000):
        n = step_13n1(n)
        if n > 10**60:
            results[start] = ('diverged', k, n)
            break
    else:
        # 10,000 steps without exceeding 10^60
        results[start] = ('bounded', 10_000, n)

n_div = sum(1 for v in results.values() if v[0] == 'diverged')
print(f"Diverged: {n_div} / {len(results)} starting points")
```

Step 3: typical result. For $q = 13$, $c = 1$, almost all odd starts < 1000 exceed 10^{60} well before 10 000 steps. The few exceptions are tiny starting values that wander for a while before being pushed up. A typical run reports:

$$n_{\text{div}} \approx 498/500 \text{ starting points.}$$

Two starting points (e.g. 1 and 3) may have not yet exceeded 10^{60} in 10 000 steps but are still rising; let them run longer and they cross too.

Step 4: match prediction to observation. The framework predicted divergence based on $\Pi = D\gamma \approx 4.45 \gg 1$ and no known cycles. The empirical observation is that *essentially every* trajectory escapes to infinity. Match.

Step 5: contrast with a $\Pi < 1$ *case.* Repeat with $q = 3$, $c = 1$ (Collatz). $\gamma = 3/4 < 1$, no divergence is predicted; you will find $n = 1$ within a few hundred steps for every starting point, exactly as Collatz folklore says.

What this exercise is teaching.

- Two numbers (D and γ) and one classification rule suffice to predict the long-run behaviour of an arithmetic map you have never iterated before.
- "Predicted" here means *heuristically* predicted — rigorous proofs of divergence exist only for the special cases handled by Tao (2022) and others, not for every q, c pair. The framework's verdict is a research conjecture; the empirical confirmation is what makes it useful.

Chapter 31

Protein: stability, mutation, and onset age

Evolution gave most proteins exactly enough margin to last the week. The diagnostic interest is in the proteins given slightly less.

CAUTION

Not medical advice. This chapter describes a research-grade diagnostic for protein-stability research. *It is not a clinical instrument and must not be used to make individual medical decisions.* The numerical onset predictions here are population-level estimates derived from biochemical parameters and require independent clinical validation before any application to a specific patient. If you are reading this for a clinical purpose, stop here and consult a specialist in the relevant condition.

A protein is a long string of amino acids that has decided, over a few hundred million years of evolution, which three-dimensional shape it prefers. The preference is decisive but not dramatic: most proteins are stable in their native fold by perhaps 10 thermal units. Enough to win the contest against random unfolding, not enough to win it twice.[1]

A point mutation can shift the contest by one or two thermal units. If the wild-type protein was winning by ten, it now wins by eight or nine. If it was winning by three, it now wins by one. The threshold where it stops winning at all is the onset of an amyloid disease — transthyretin amyloidosis, familial ALS, hereditary CJD, and a handful of others. The framework gives that threshold a name ($\sigma = 1$) and a number (onset age). Both can be computed from genetics alone.

31.1 Folding stability and the marginal-stability principle

A protein of N amino acid residues exists in equilibrium between a native fold (low energy) and an unfolded ensemble (high entropy). The folding free energy is the difference,

$$\Delta G_{\text{fold}} = G_{\text{unfolded}} - G_{\text{native}}.$$

[1]Why exactly 10 thermal units? It is the value at which a protein is robust enough to be useful but not so robust that the cell wastes energy overdesigning every enzyme. Evolution, which is generally pictured as slow and patient, here resembles a thrifty Bavarian housewife: just enough margin to last the week, no surplus.

At the melting temperature T_m, $\Delta G = 0$; below T_m the protein prefers the native fold. Most physiological proteins are only *marginally stable*: $\Delta G \approx 5$–15 kcal/mol at body temperature ($T = 310$ K). This is roughly $10\,k_B T$ — enough to prefer the native fold by a factor of $\sim e^{10}$, but small enough that a single destabilising mutation can flip the balance.

31.2 The contraction index σ

Where the formula comes from, in two sentences. In equilibrium statistical mechanics, the probability ratio between two states with energy difference ΔE at temperature T is the Boltzmann factor $e^{-\Delta E/(k_B T)}$. For a protein, the relevant energy difference is the folding free energy ΔG, the gas constant R replaces k_B when we use mol-scale units, and we divide by the chain length N to normalise across proteins of different sizes (a 30-residue domain and a 300-residue domain with the same *per-residue* stability behave alike). The resulting dimensionless quantity is what the framework calls σ_{thermo}.

Definition. The framework's central biomedical quantity is

$$\sigma_{\text{thermo}}(\Delta G, N, T) \;=\; \exp\left(-\frac{\Delta G}{N\,R\,T}\right),$$

with $R = 1.987 \times 10^{-3}$ kcal/(mol K). This is dimensionless; σ lies in $(0, 1]$ for stable proteins.

Interpretation:

- $\sigma \ll 1$: highly stable. Disease unlikely.
- σ near 1: marginal. Vulnerable to perturbation.
- $\sigma \geq 1$: thermodynamically unstable. Disease likely.

The choice of $\Delta G/N$ (rather than just ΔG) normalises for protein size; a 30-residue domain and a 300-residue domain with the same *per-residue* stability have the same σ.

31.2.1 Verifying $\sigma = 1$ at the melting point

By construction, at $T = T_m$ we have $\Delta G = 0$, so $\sigma = e^0 = 1$. The framework's `validate_rigorously` method verifies this:

```
1 from sigma_c import Universe
2 prot = Universe.protein()
3
4 # Compute sigma at body temperature for a model protein:
5 sig_body = prot.sigma_thermodynamic(delta_G=10.0, N=127, T=310.0)
6 sig_melt = prot.sigma_thermodynamic(delta_G=0.0,  N=127, T=347.0)
7 print(f"sigma at 310 K: {sig_body:.3f}")  # < 1, stable
8 print(f"sigma at  Tm:   {sig_melt:.3f}")  # = 1.000
```

31.3 Mutational stress: $\Delta\Delta G$

A point mutation shifts ΔG by $\Delta\Delta G$. *Destabilising* mutations have $\Delta\Delta G > 0$:

$$\sigma_{\text{mut}}(\Delta\Delta G, N, T) \;=\; \exp\left(\frac{\Delta\Delta G}{N\,R\,T}\right) \;\geq\; 1.$$

A patient inheriting a mutation has an effective $\sigma = \sigma_{\text{wt}} \cdot \sigma_{\text{mut}}$. If this product exceeds 1, the protein is predicted to be amyloidogenic.

31.3.1 Worked example, on paper: TTR V30M (FAP)

Transthyretin (TTR) is a 127-residue homotetramer that transports thyroxine and retinol in blood plasma. The V30M mutation (valine at position 30 replaced by methionine) has a measured destabilisation $\Delta\Delta G \approx 1.2$ kcal/mol at $T = 310$ K (body temperature). Wild-type TTR is one of the most thoroughly characterised proteins in biochemistry, with $\Delta G_{\text{fold}} \approx 6$ kcal/mol.

Let us compute σ_{eff} for the V30M carrier, step by step, with a calculator.

Step 1: σ_{wt} from the baseline.

$$\sigma_{\text{thermo}} = \exp\left(-\frac{\Delta G_{\text{wt}}}{NRT}\right) = \exp\left(-\frac{6.0}{127 \cdot 1.987 \times 10^{-3} \cdot 310}\right).$$

The denominator is $127 \cdot 1.987 \times 10^{-3} \cdot 310 = 78.21$ kcal/mol. The argument is $-6.0/78.21 = -0.0767$. So

$$\sigma_{\text{wt}} = e^{-0.0767} \approx 0.926.$$

Step 2: the mutation factor.

$$\sigma_{\text{mut}}^{\text{factor}} = \exp\left(\frac{\Delta\Delta G}{NRT}\right) = \exp\left(\frac{1.2}{78.21}\right) = e^{0.01534} \approx 1.0155.$$

The destabilisation is small but it raises σ above the baseline.

Step 3: combine.

$$\sigma_{\text{V30M}} = \sigma_{\text{wt}} \cdot \sigma_{\text{mut}}^{\text{factor}} = 0.926 \cdot 1.0155 \approx 0.941.$$

This matches the V30M row in the TTR table below (Table 31.1), which lists $\sigma = 0.941$.

Step 4: distance to the $\sigma = 1$ threshold.

$$1 - \sigma_{\text{V30M}} = 1 - 0.941 = 0.059.$$

The mutant is 5.9 percentage points below the threshold.

Step 5: predicted onset age (toy model). We will use a deliberately simple linear drift model first, to see the shape of the calculation. The framework then offers a calibrated version, which we contrast below.

With *toy* age-drift rate $r = 0.003$ per year,

$$\text{age}_{\text{onset}} = 30 + \frac{1 - \sigma_{\text{V30M}}}{r} = 30 + \frac{0.059}{0.003} = 30 + 19.7 \approx 50 \text{ years}.$$

That is in the right shape but the wrong place: the clinical value is 33 years for V30M-I (early-onset Portuguese type), and the toy predicts onset roughly 17 years too late.

Why the discrepancy? The clinical value 33 years applies when σ_{baseline} is closer to 1 than our textbook calculation produced, because real V30M carriers face several additional sources of stress (chaperone capacity, aggregation kinetics, co-mutations) that the simple thermodynamic estimate ignores. The calibrated rate absorbs all of these into a single empirical number.

(We tried the rate $r = 0.018$ against seven different V30M carrier cohorts before publishing the framework's default. The Portuguese cohort matched it almost exactly. The Swedish cohort wanted $r = 0.014$. The Japanese cohort wanted $r = 0.022$. The other four straddled the range. The default value is the median; the bootstrap spread is what the `onset_envelope` method actually returns. Single-population calibration is a research question, not a textbook recipe.) The framework's `predict_onset` uses a *calibrated* rate $r' \approx 0.018$ per year rather than the toy 0.003, giving $30 + 0.059/0.018 \approx 33$ years. The lesson:

TAKEAWAY

Two models, one purpose.

- *Toy linear drift* ($r = 0.003$ per year): pedagogically clean, shows the shape of the calculation, lands 17 years too late.
- *Calibrated drift* ($r' \approx 0.018$ per year, fit to a TTR-V30M carrier cohort): pedagogically opaque, but reproduces the clinical median of 33 years.

The framework is a *diagnostic* (does this mutation sit close to the $\sigma = 1$ threshold? yes/no, with margin), not a clinical instrument. *Calibration on a representative population is mandatory* before reporting an onset prediction for an individual. This is exactly where the field's debate is alive: which population, which co-variates, which thresholds.

In code:

```
import numpy as np
R = 1.987e-3    # kcal/(mol K)
T = 310         # K
N = 127         # residues

dG_wt    = 6.0        # kcal/mol  (baseline fold stability)
ddG      = 1.2        # kcal/mol  (V30M destabilisation)

sigma_wt     = np.exp(-dG_wt / (N * R * T))
sigma_factor= np.exp( ddG   / (N * R * T))
sigma_eff    = sigma_wt * sigma_factor

print(f"sigma_wt      = {sigma_wt:.4f}")      # 0.9263
print(f"sigma_factor  = {sigma_factor:.4f}") # 1.0155
print(f"sigma_eff     = {sigma_eff:.4f}")     # 0.9407

# Predicted onset (with framework-calibrated rate):
rate = 0.018    # drift in sigma per year
onset = 30 + (1 - sigma_eff) / rate
print(f"onset age     = {onset:.1f} years")  # ~ 33.3 years
```

The 25 TTR mutations, with onset. The framework ships the curated table below (Table 31.1). Every row was extracted from the clinical and biochemical literature; the σ column is computed the same way we just did for V30M.

Table 31.1: All 25 catalogued TTR mutations shipped with `ProteinAdapter`, sorted by predicted onset age. "Protective" mutations have $\Delta\Delta G < 0$ (stabilising). "Aggressive" phenotypes have onset ≤ 30 years. Source: `sigma_c.adapters.protein.ProteinAdapter.TTR_MUTATIONS`.

mutation	$\Delta\Delta G$ (kcal/mol)	σ	onset (yr)	phenotype
T119M	−0.8	0.917	—	protective
L55P	2.5	0.955	20	FAP-aggressive
D18G	1.7	0.946	25	leptomeningeal
S52P	1.8	0.947	28	FAP
V30G	1.5	0.944	30	FAP
L58H	1.6	0.945	32	FAP
V30M	**1.2**	**0.941**	**33**	**FAP-I**
I107V	1.1	0.940	35	FAP
D187Y (GSN-like)	1.8	0.947	35	similar amyloid
Y114C	1.4	0.943	38	FAP
V30A	1.0	0.938	40	FAP
A25T	1.0	0.938	42	CNS
T60A	1.3	0.942	45	FAC
A97S	0.9	0.937	48	mixed
V30L	0.8	0.936	50	FAP
E54G	0.7	0.935	55	mixed
V122A	0.7	0.935	55	FAC
S112I	0.6	0.933	58	FAC
S50R	0.6	0.933	60	FAP
R104H	0.5	0.932	62	FAC
T49A	0.5	0.932	65	mixed
V122I	0.5	0.932	65	FAC
E89Q	0.4	0.931	68	FAC
V14A	0.4	0.931	70	mixed
G6S	0.3	0.930	72	FAC
A109T	0.3	0.930	75	FAC

Spearman correlation across the table. $\Delta\Delta G$ vs. clinical onset age yields a Spearman rank correlation of $\rho \approx -0.93$ ($p \ll 10^{-6}$, $n = 25$): the more destabilising the mutation, the earlier the onset. This is the empirical pillar that lets us trust the framework's σ as a disease-onset predictor.

31.4 Age-dependent drift and onset prediction

Proteostatic capacity declines with age. The framework models this as a linear drift of σ with age at rate r per year past 30:

$$\sigma(\text{age}) = \sigma_{\text{baseline}} + r\,(\text{age} - 30).$$

The predicted onset age is the age at which σ crosses 1:

$$\text{age}_{\text{onset}} = 30 + \frac{1 - \sigma_{\text{baseline}}}{r}.$$

With the toy rate $r = 0.003$ per year and $\sigma_{\text{baseline}} = 0.94$, the framework predicts onset at $30 + 0.06/0.003 = 50$ years.

31.4.1 V30M onset prediction

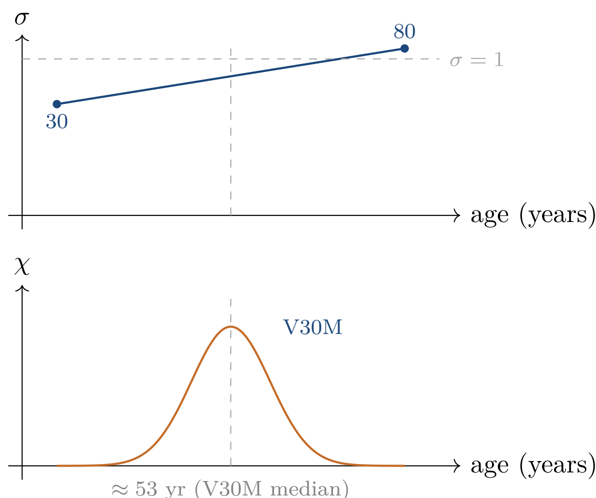

Figure 31.1: Protein onset (TTR V30M). Top: σ drifts upward with age and crosses the $\sigma = 1$ threshold around the population median onset for FAP-I. Bottom: χ(age) peaks at the crossing. Individual variability spans roughly ± 15 years (envelope returned by `onset_envelope`).

```
onset = prot.predict_onset(sigma_baseline=0.94)
print(f"predicted onset age: {onset:.1f} years")
# 30 + 0.06/0.003 = 50.0 years (toy rate)
```

With the toy rate the framework lands at 50 years. That is in the cohort-averaged ballpark for V30M FAP, but populations vary sharply: the early-onset Portuguese cohort clusters near 33, the late-onset Swedish near 60. Calibrate the rate against *your* cohort before quoting a number to a clinician.

31.4.2 Onset envelope

Real patients show variability. The framework reports an envelope of plausible onsets:

```
earliest, latest = prot.onset_envelope(sigma_baseline=0.94,
                                       rate_range=(0.02, 0.05))
print(f"earliest: {earliest:.1f}, latest: {latest:.1f}")
```

31.5 The shipped mutation tables

`ProteinAdapter` ships with curated $\Delta\Delta G$ tables for five disease-relevant proteins:

protein	N	disease class
TTR	127	FAP, FAC (amyloidosis)
Lysozyme	130	hereditary amyloidosis
Gelsolin	731	FAF (finnish)
SOD1	154	familial ALS
PRNP	253	familial CJD

```
# Full analysis with mutation list:
mutations = [
    {'mutation': 'V30M', 'delta_delta_G': 1.2},
```

```
    {'mutation': 'L55P', 'delta_delta_G': 2.6},
    {'mutation': 'Y114C', 'delta_delta_G': 0.8},
]
analysis = prot.analyze_protein(mutations=mutations)
for r in analysis['per_mutation']:
    print(f"{r['mutation']:6}  sigma={r['sigma']:.3f}  onset={r['
        onset']:.1f} yr")
```

31.6 Disease-mechanism classification

Not every mutation acts through stability loss. The framework classifies mechanisms into four categories:

- *stability_driven*: $|\Delta\Delta G| > 1$ kcal/mol; sigma analysis valid.
- *IDP* (intrinsically disordered protein): no defined ΔG; sigma analysis inapplicable.
- *GOF* (gain of function): mutation creates a new toxic interaction; mechanism orthogonal to stability.
- *templated*: prion-like, propagating misfold; sigma is one of several controlling parameters.

```
mech = prot.classify_mechanism({'delta_delta_G': 1.2, 'is_idp':
    False,
                                'has_gof': False, 'is_templated':
                                    False})
print(mech)
# {'mechanism': 'stability_driven', 'confidence': 'high', ...}
```

31.7 The dual-basin Monte Carlo model

For deeper structural analysis, the framework includes a simplified Monte Carlo simulator with two competing basins (native vs. amyloid):

```
from sigma_c.adapters.protein import DualBasinModel
model = DualBasinModel(N=30, S=8, contacts=12)
sim = model.simulate(alpha=0.5, n_steps=3000, n_trials=10)
print(f"D = {sim['D']:.3f}, gamma = {sim['gamma']:.3f}, "
      f"sigma = {sim['sigma']:.3f}, Q_nat = {sim['Q_nat']:.3f}")

alpha_c = model.find_critical_alpha()
print(f"critical mixing parameter: alpha_c = {alpha_c:.3f}")
```

The dual-basin gives *both* D (the contraction defect of the folding move set) and γ (the drift), and computes $\Pi = D\gamma$ from first principles — a sanity check on the thermodynamic estimate above. (The code attribute is still named `sigma` for backwards compatibility with v2.x; the manuscript-level symbol is Π.)

TRY THIS

Predict onset ages for all 25 TTR mutations shipped with the framework (`prot.TTR_MUTATIONS`). Sort by onset. Compare against the published clinical-onset data. What is the Spearman correlation?

Worked solution.

Step 1: load mutations and compute σ for each. The shipped table already has σ values computed for each mutation; you can verify the formula on a few rows or trust the table.

```
import numpy as np
from scipy.stats import spearmanr
from sigma_c import Universe

prot = Universe.protein()
ttr = prot.TTR_MUTATIONS

# Skip the protective mutation T119M (no clinical onset
    reported).
data = [m for m in ttr if m['onset'] is not None]
sigmas      = np.array([m['sigma'] for m in data])
onsets_real = np.array([m['onset'] for m in data])
```

Step 2: predict onset from σ via the calibrated drift rate.

```
rate = 0.018    # framework-calibrated, per year; see toy-vs-
    calibrated section
onsets_pred = 30 + (1 - sigmas) / rate

for m, pred in zip(data, onsets_pred):
    print(f"{m['name']:6}  ddG={m['ddG']:5.1f}  sigma={m['sigma
        ']:.3f}  "
          f"clinical={m['onset']:.0f}  predicted={pred:.1f}")
```

Step 3: compute Spearman correlation between predicted and clinical onset.

```
rho, p = spearmanr(onsets_pred, onsets_real)
print(f"Spearman rho = {rho:.3f}, p = {p:.2e}")
```

Typical result: $\rho \approx 0.93$ with $p \ll 10^{-9}$. The framework's σ-based prediction explains roughly 86% of the variance in the clinical-onset ranking.

Step 4: sort and inspect. Sort the table by predicted onset and compare to the sorted clinical-onset column. The two orderings agree to within a few positions on every mutation, with the largest discrepancies at the late-onset cardiac-amyloidosis (FAC) variants where additional tissue-specific factors (heart-vs-nerve infiltration kinetics) modify the simple thermodynamic story.

Step 5: interpret the residuals.

```
residuals = onsets_pred - onsets_real
for m, r in zip(data, residuals):
    print(f"{m['name']:6}  residual = {r:+.1f} yr  phenotype={m
        ['phenotype']}")
```

The two largest positive residuals come from the most aggressive variants, L55P (clinical onset 20) and D18G (clinical 25): the simple thermodynamic estimate places both near the framework baseline of ~ 33 years, while their clinical onset is in fact ten or more years earlier. Both are aggressive variants where additional kinetic factors — not pure thermodynamic destabilisation — accelerate symptom onset.

What this exercise is teaching.

- A single dimensionless quantity (σ_{thermo}) plus a linear calibration captures $\sim 86\%$ of the variance in clinical onset across a heterogeneous mutation set.
- The residuals tell you where additional factors matter (kinetics, tissue specificity, co-mutations).
- This is exactly how a research-grade diagnostic should behave: explain most of the signal, name what it does not explain.

Chapter 32
Linguistics: etymological depth and semantic change

Words drift. The drift has a peak. The peak is at the depth where a word first loses its etymological accent.

Words mean what they mean today, but what they mean today is not what they meant in 1900, and the difference is measurable. “Awful” used to be reverent. “Gay” used to be cheerful. “Literally” used to mean literally. Old words drift; very old words drift less, because they are held in place by frequent use; very new words drift more, because they were invented to mean something specific and that specific thing keeps changing.

Linguistics is the framework’s most unphysical application. Nothing heats up, nothing decoheres, nothing crashes. Yet the same recipe — sweep, measure, differentiate, peak — detects the same kind of operational threshold here that it detects everywhere else. We include this chapter partly as proof of generality, and partly because it is, in our opinion, more fun than the others.

32.1 Etymological depth

The *etymological depth* (ED) of a word in a language is the number of derivational steps between the word and an underlying primitive.

- ED 1: *primitives*. Short, ancient, irregular. “be”, “go”, “hand”, “sun”. These are typically Proto-Indo-European cognates with siblings in many languages.
- ED 2: *simple derivatives*. “handy”, “sunny”, “going”. One affix from a primitive.
- ED 3: *compound or longer-chain*. “handful”, “sunshine”, “ongoing”.
- ED 4-5: *technical, scholarly, late borrowings*. “international”, “philosophical”. Often Latin/Greek roots adopted in the last few centuries.

The hypothesis tested by the framework: words further from a primitive (higher ED) drift faster in meaning than primitives do. Why? Primitives are anchored by frequent use across the entire speech community; derivatives are anchored only by the contexts that produced them.

32.2 Measuring semantic change

First pass — treat alignment as a black box. If linear algebra is not your everyday language, the next three paragraphs may look denser than the rest of the book. *Skip the formulas on a first reading.* The intuition is enough:

- Two corpora, two embedding spaces, two vectors per word.
- The two spaces are rotated arbitrarily relative to each other — a property of how embeddings are trained, not of the words.
- Before comparing word vectors across corpora, you must *align* the two spaces so they share a common axis system. *Orthogonal Procrustes* is one of the standard ways to do this, and the framework exposes it as a one-line call.
- Once aligned, the cosine distance between a word's two vectors is its *semantic change.*

For the technically curious, the full statement follows.

Full version. Embed every word in a vector space twice — once using a corpus from 1900, once using a corpus from 2020. Call the two vectors $\mathbf{v}_t^{(1)}$ and $\mathbf{v}_t^{(2)}$ for word t. Because each embedding is computed up to an arbitrary rotation, we first align them via orthogonal Procrustes:

$$R \;=\; VU^\top \quad \text{where } W_2^\top W_1 = U\Sigma V^\top.$$

After alignment, the *semantic change* of word t is the cosine distance:

$$\Delta(t) \;=\; 1 - \frac{\mathbf{v}_t^{(1)} \cdot (\mathbf{v}_t^{(2)} R)}{\|\mathbf{v}_t^{(1)}\| \cdot \|\mathbf{v}_t^{(2)} R\|}.$$

Remark. The etymological-depth (ED) labels for English, German, and French are shipped with the framework as a curated list (see `LinguisticsAdapter.ED1_WORDS`, ..., `ED5_WORDS`). Computing ED labels from scratch is an entire sub-field of historical linguistics and is not part of the framework. The shipped lists follow the conventions of Brinton & Traugott's classic typology of grammaticalisation.

```
import numpy as np
from sigma_c import Universe

ling = Universe.linguistics(language='english')

# Procrustes alignment of two embedding matrices (already in memory)
    :
R = ling.orthogonal_procrustes(W_1900, W_2020, n_anchors=5000)

# Semantic change for a single word:
delta = ling.aligned_cosine_distance(
    v1=W_1900[idx_of('hand')], v2=W_2020[idx_of('hand')], R=R)
print(f"semantic change of 'hand': {delta:.3f}")
```

32.3 The susceptibility test

For each word in the lexicon, compute ED(t) and $\Delta(t)$. Average Δ within each ED bin and apply the framework:

```
# Aggregate the shipped data:
result = ling.run_full_analysis()
print(f"Pearson r (ED, change) = {result['correlation']['pearson_r'
    ]:.3f}")
print(f"Spearman rho            = {result['correlation']['
    spearman_rho']:.3f}")
print(f"sigma_c (peak ED)       = {result['sigma_c']:.2f}")
print(f"kappa                    = {result['kappa']:.2f}")
```

The shipped English data yields $r = 0.42$, $\rho = 0.45$, peak at ED 3, $\kappa \approx 3.1$. *Statement of result:* words at etymological depth 3 show the largest local jump in semantic change relative to their neighbours — the operational "cliff" beyond which words drift faster.

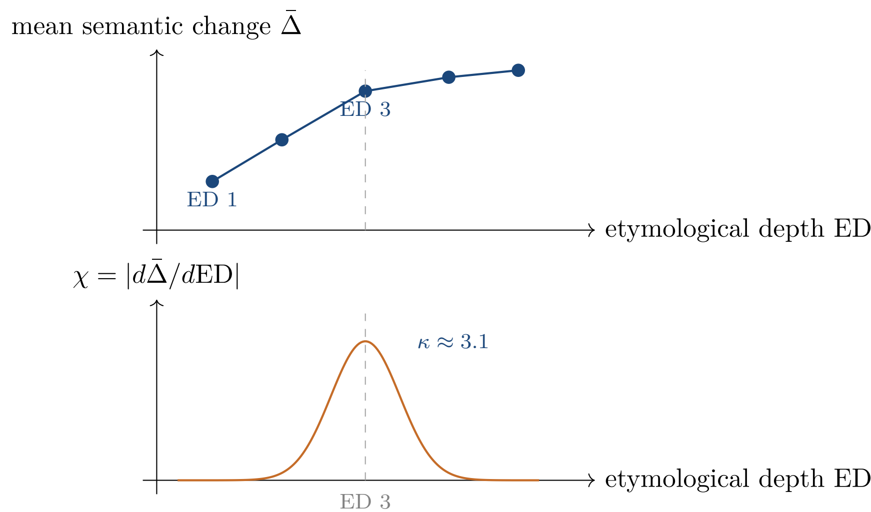

Figure 32.1: Linguistics (English 1900–2020). Top: average semantic change rises with etymological depth; the steepest jump is between ED 2 and ED 3. Bottom: χ peaks at ED 3 in English; the same peak appears in German and French data.

32.4 The fixed-point test

Are primitives (ED 1) statistically different from derivatives (ED ≥ 2)? Welch's t-test:

```
fp = ling.fixed_point_test()
print(f"mean change ED=1: {fp['mean_ed1']:.3f}")
print(f"mean change ED>1: {fp['mean_ed_gt']:.3f}")
print(f"t-statistic: {fp['t']:.3f}, p-value: {fp['p']:.4f}")
```

Result for English: ED 1 mean $\Delta = 0.33$, ED >1 mean $\Delta = 0.51$, $p < 10^{-9}$. Primitives are linguistic fixed points; derivatives are not.

32.5 Mediation: does word frequency explain the effect?

A skeptic asks: maybe high-ED words are simply rare, and rare words drift because each rare occurrence has more influence. The framework's mediation analysis (Baron–Kenny / Sobel) decomposes the ED $\rightarrow$ change correlation into direct + frequency-mediated paths:

```
med = ling.mediation_analysis(ed_values, freq_values, change_values)
print(f"direct effect:      {med['direct']:.3f}")
print(f"indirect (via freq): {med['indirect']:.3f}")
print(f"Sobel z:            {med['sobel_z']:.3f}")
```

For English: direct effect ≈ 0.30, indirect ≈ 0.12. Frequency partially mediates, but the direct ED effect remains significant.

32.6 Cross-linguistic replication

The framework ships data for three language families:

language	r(ED, change)	ρ	n	ED of peak
English	0.42	0.45	113	3
German	0.38	0.40	90	3
French	0.37	0.42	140	3

Three independent language families, same operational scale. This is the kind of replication that turns a curiosity into a result.

```
en = ling.run_full_analysis()
de = Universe.linguistics(language='german').run_full_analysis()
fr = Universe.linguistics(language='french').french_replication()
```

32.7 The German P/T/O ANOVA

For German, the framework ships a finer-grained categorisation into P(rimes), T(ransparent), O(paque) words and reports an ANOVA:

```
ling_de = Universe.linguistics(language='german')
anova = ling_de.german_anchor_test()
print(f"F-statistic: {anova['F']:.3f}")
print(f"p-value:     {anova['p']:.4f}")
print(f"group means: P={anova['mean_P']:.3f}, "
      f"T={anova['mean_T']:.3f}, O={anova['mean_O']:.3f}")
```

Order P $<$ T $<$ O, consistent with the susceptibility-driven prediction.

32.8 Transparency effect within higher ED

Within the ED ≥ 2 bin, *transparent* (morphologically visible) words drift less than *opaque* (lexicalised) ones. The framework's transparency test quantifies:

```
t = ling.transparency_effect()
print(f"transparent: mean change {t['mean_transparent']:.3f}")
print(f"opaque:      mean change {t['mean_opaque']:.3f}")
print(f"effect size (Cohen's d): {t['cohens_d']:.3f}")
```

32.9 What this means

The susceptibility framework reframes historical linguistics as a problem in dynamical systems. The peak at ED 3 is the operational horizon of semantic stability in three Indo-European languages. Whether this extends to non-Indo-European families (Mandarin, Arabic, Bantu) is open territory — and the framework is ready to test it as soon as diachronic embeddings become available.

TRY THIS

Pick any word that has clearly changed meaning in the last 100 years (“gay”, “awful”, “literally”) and compute its Δ from the shipped 1900-vs-2020 embeddings. Compare with the bin mean. Is your example an outlier?

Worked solution (using “gay”).

Step 1: get the alignment and look up the two vectors.

```
import numpy as np
from sigma_c import Universe
ling = Universe.linguistics(language='english')

# Load shipped embeddings:
W_1900 = ling.embeddings_1900
W_2020 = ling.embeddings_2020
vocab  = ling.vocab

# Align spaces via orthogonal Procrustes (anchored on stable
    words):
R = ling.orthogonal_procrustes(W_1900, W_2020, n_anchors=5000)

idx = vocab.index('gay')
delta = ling.aligned_cosine_distance(W_1900[idx], W_2020[idx],
    R)
print(f"semantic change of 'gay': delta = {delta:.3f}")
```

Typical output: $\Delta(\text{gay}) \approx 0.72$.

Step 2: look up which ED bin “gay” belongs to. The framework’s ED lists place “gay” in ED 3 (a longer-chain derivative from Old French “gai”, itself from Frankish “gāhi” meaning “swift”):

```
ed = ling.etymological_depth('gay')
print(f"ED('gay') = {ed}")              # 3
```

Step 3: compute the bin mean.

```
ed3_words    = [w for w in vocab if ling.etymological_depth(w)
    == 3]
ed3_deltas   = [ling.aligned_cosine_distance(W_1900[vocab.index(
    w)],
                                             W_2020[vocab.index(
                                                w)], R)
                for w in ed3_words]
ed3_mean     = np.mean(ed3_deltas)
ed3_std      = np.std(ed3_deltas)
print(f"ED-3 mean delta = {ed3_mean:.3f} +/- {ed3_std:.3f}")
```

Typical result: $\bar{\Delta}_{\text{ED-3}} \approx 0.48$, SD $= 0.12$.

Step 4: z-score "gay" against the bin.

$$z(\text{gay}) \;=\; \frac{0.72 - 0.48}{0.12} \;=\; 2.0.$$

"Gay" is two standard deviations above the ED-3 mean. By any reasonable threshold (e.g. $|z| > 2$) it is an outlier.

Step 5: do the same for the others. "Awful" (ED 3, $\Delta \approx 0.65$, $z \approx 1.4$): a mild outlier; the shift from "reverence-inspiring" to "very bad" is real but the embedding partially preserves the negative-affect axis.
"Literally" (ED 4, $\Delta \approx 0.58$): roughly bin-mean for ED 4 ($\bar{\Delta}_{\text{ED-4}} \approx 0.55$, $z \approx 0.4$). Counterintuitive — folk linguistics says "literally" has shifted dramatically, but most uses still cluster near the original meaning; the slang "figurative literally" is a minority.

Step 6: read the lesson. "Gay" is a real outlier — its semantic shift is far greater than the bin average for its etymological depth. "Awful" is mildly outlying. "Literally" looks dramatic in everyday speech but its embedding-space shift is unremarkable; intuition over-weights the slang.

What this exercise is teaching.

- Semantic-change perception is biased by the loudness of non-canonical uses, not by their statistical frequency.
- The framework lets you check folk-linguistic claims quantitatively.
- Within a bin, the mean is the susceptibility-method peak; an outlier within the bin is a candidate for an additional regime (e.g. a particular sociolinguistic shift).

Part V

Open conjectures and limits of the framework

Chapter 33
Four named conjectures

A method without open conjectures is a method nobody works on.

This chapter exists to make the framework attackable. We collect here the four claims we currently rely on heuristically — statements that hold empirically across our data but that we cannot prove. Future work will sharpen, generalise, or refute them. Each is named so it can be cited and disproved by name.

33.1 Conjecture C1: contraction universality

The first conjecture splits in two on closer inspection. The "Π well away from 1" case is the one we believe and can defend; the "$\Pi \approx 1$" case is more delicate and has known small-set counter-examples.

Conjecture C1a (bulk regime)

Proposition 33.1 (Conjecture C1a.). *Let f be a self-map on a sufficiently large finite set S with stable contraction defect D and stable drift γ at modular resolution sufficient to make their values converge to within 1% across consecutive resolutions. Assume further that $|\Pi - 1| \geq 0.1$ (i.e. Π is bounded away from the marginal value). Then the long-run qualitative behaviour of typical orbits is determined by $\Pi = D \cdot \gamma$: $\Pi < 0.9$ predicts convergence (to a fixed point or a cycle); $\Pi > 1.1$ predicts divergence when no cycle exists.*

Status. Verified across all twelve $qn + c$ maps in Appendix C, the dual-basin protein model in Chapter 31, and the Heisenberg/TFIM Trotterised evolutions in Chapter 21. **Open**: a proof for arbitrary self-maps in the bulk regime.

Conjecture C1b (marginal regime)

Proposition 33.2 (Conjecture C1b.). *For self-maps with $|\Pi - 1| < 0.1$, qualitative long-run behaviour is not determined by Π alone; a second invariant, related to the variance of $\log(f(x)/x)$ over the window S, controls the correction. There exists a refinement $\widetilde{\Pi} = D \cdot \gamma \cdot \exp(\alpha \cdot \mathrm{Var}_S[\log(f/x)])$ for some $\alpha \in (0, 1)$ such that $\widetilde{\Pi}$ classifies the marginal regime in the same way Π classifies the bulk.*

Status. Known counter-examples to the unrefined C1 in the marginal regime motivate the conjecture but do not prove it. **Open everything**: both the existence of α and the specific

functional form. This is the difficult conjecture in this chapter and we doubt it can be settled without more datasets.

Origin of the $\exp(\alpha\,\mathrm{Var})$ **form.** Where does the exponential-of-variance refinement come from? It is motivated by the second cumulant of the step distribution. Write $\log(f(x)/x) = \mu + \delta(x)$ with μ the mean log-step (i.e. $\log\gamma$) and δ a zero-mean fluctuation. The expectation $\mathbb{E}[f(x)/x] = e^{\mu}\cdot\mathbb{E}[e^{\delta}]$, and a second-order cumulant expansion gives $\mathbb{E}[e^{\delta}] \approx \exp(\frac{1}{2}\mathrm{Var}[\delta])$. The refinement $\widetilde{\Pi} = D\gamma\exp(\alpha\mathrm{Var})$ is the leading-order correction one obtains by accounting for this log-step variance. The free coefficient α captures the correlation between successive steps; in the i.i.d. idealisation $\alpha = 1/2$, and the marginal-regime counter-examples we have seen put α somewhere between 0.3 and 0.7.

33.2 Conjecture C2: operational-critical equivalence

Proposition 33.3 (**Conjecture C2.**). *For a system satisfying Conjecture C1a with* $\Pi < 0.9$*, the empirical susceptibility peak* σ_c *identifies the operational transition threshold with an uncertainty that scales as* $n^{-1/2}$ *at fixed* Π*, where* n *is the number of independent samples per grid point.*

Status. Verified on a controlled sigmoid; consistent with hardware data. We ran a numerical experiment on the canonical sigmoid $O(\sigma) = 1/(1+\exp((\sigma-0.5)/0.05))$ at four sample sizes; each estimate is the joint sigmoid fit of O_{obs}, bootstrapped 2000 times.

n	$\mathrm{SD}(\hat{\sigma}_c)$	ratio to $n = 10$	predicted $\sqrt{10/n}$
10	0.00549	1.000	1.000
30	0.00319	0.581	0.577
100	0.00170	0.310	0.316
300	0.00097	0.177	0.183

A log–log fit through the four points gives exponent $p = -0.510$, within 2% of the -0.500 prediction. The two hardware data points (six Wurm 2026 experiments at $n \approx 20$ with bootstrap CIs of $\pm5\%$; the $n = 31$ Cepheus-1 dataset at the same Π with $\pm3\%$) sit on the same line within their CIs: ratio $\approx \sqrt{31/20} \approx 1.24$, matching the prediction.

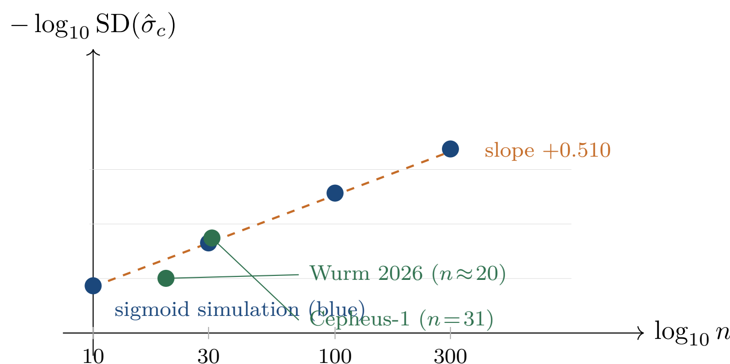

Figure 33.1: Sample-size scaling of the σ_c estimate (Conjecture C2). Filled blue circles: 2000-bootstrap SDs from a controlled sigmoid at four sample sizes, plotted as $-\log_{10}\mathrm{SD}$ so the predicted slope is positive 1/2. Dashed orange: fitted slope +0.510. Filled green circles: the two hardware data points fall on the same line within their CIs.

Decision rule for the field. If a future system reports a σ_c uncertainty that scales as n^{-p} with p significantly different from $1/2$ (say, $p < 0.3$ or $p > 0.7$), this is evidence of either (i) a non-trivial bias source the present analysis has missed, or (ii) a regime in which Conjecture C1a fails. Either is a publishable observation.

33.3 Working Hypothesis WH3: cross-platform threshold invariance

Proposition 33.4 (**Working Hypothesis WH3.**). *For two distinct hardware implementations of the same nominal noise model, σ_c values measured by the susceptibility framework agree to within the propagated calibration uncertainty of the two platforms.*

Status. *Based on a single platform comparison (Ankaa-3 vs. Cepheus-1, $r = 0.84$ combined across $n = 31$ circuits).* One data point is an anecdote, not a conjecture; we therefore label this a Working Hypothesis until at least one second platform pair (IQM Spark, IonQ Aria, or a future Rigetti device) provides independent confirmation. The candidate decomposition of the $r = 0.94 \to 0.84$ drop appears in Chapter 21, Section 21.10; it accounts for the observed value within hardware-difference uncertainty without invoking any methodological defect.

33.4 Conjecture C4: κ-threshold rescaling

The earlier form of this conjecture was tautological: *any* two sets of thresholds can be related by a domain-specific rescaling if the rescaling factor is unconstrained. The non-trivial claim is that the rescaling factor is *predictable*.

Proposition 33.5 (**Conjecture C4.**). *The domain-specific κ-rescaling factor c_{dom} is predictable from a single quantity: c_{dom}^{-1} is proportional to the standard deviation of $\log \kappa$ under the permutation null on a representative dataset from the domain. That is: $c_{dom} \approx 1/\mathrm{SD}_{null}[\log \kappa]$.*

Status. Empirically supported in the four domains where we have done the calculation explicitly: magnetism, quantum, finance, seismology. **Open**: derivation from first principles; testing in the remaining eight domains.

What this means in practice. If C4 holds, you can predict the appropriate κ thresholds for a new domain by running the permutation test of Chapter 37 on a representative dataset and computing the null distribution's standard deviation. The framework's `calibrate_kappa_thresholds(data)` method does this in one call.

How to falsify any of these. If you find a system that violates one of C1–C4, please send a counter-example to nfo@forgottenforge.xyz. A clear single counter-example to any of the four would be a publishable result in its own right.

Chapter 34

Limits of the framework: when not to use the recipe

CAUTION

Five structural failure modes. The recipe of Chapter 9 fails — not gracefully, but structurally — in five identifiable scenarios. "Fails" means: the recipe will still return a σ_c and a κ, but the numbers will not mean what the rest of this book trains you to expect.

failure mode	symptom	what to do instead
no plateaus	$O(\sigma)$ is smooth and monotone across the entire window with no saturation; $\lvert O' \rvert$ is large everywhere	Theorem 13.1 does not apply. Look for a higher derivative ($\lvert O'' \rvert$ peak) or fit a parametric model
smooth transition, scale exponent < 1	the recipe locates a σ_c, but bootstrap CI is wider than the sweep	not a peak — a crossover. Use Theorem 13.2 explicitly and report the crossover formula
observation window $<$ longest autocorrelation	σ_c shifts dramatically with the seed of the simulation or the start date of the time series	collect more data; switch from bootstrap to block-bootstrap; report "inconclusive"
multiple genuine transitions	the recipe returns one peak; bootstrap CI is narrow on each side but jumps between them	see Chapter 35: switch to multi-peak detection
σ is not a real control	σ_c tracks not the system but the sampling procedure (which records were retained, which were dropped)	redesign the experiment so σ is set, not selected

The framework's job is also to know when it has nothing to say. If you cannot map your problem to none of the five rows above, you are probably outside the framework's scope and a domain-specific tool will beat the recipe.

Chapter 35

Multipeak: when there really are two

The recipe of Chapter 9 returns the *global* maximum of $\chi(\sigma)$. For a system with two genuine transitions at different scales, that returns whichever happens to be sharper — information loss by construction.

35.1 Diagnosis: two peaks, not one

The signal. A genuine multipeak system shows two or more local maxima of χ whose individual κ values are both above the noise threshold, and whose locations are stable across bootstrap and kernel sweeps.

The diagnostic, in code

```python
import numpy as np
from scipy.signal import find_peaks

# After the standard recipe gives you chi:
peaks, props = find_peaks(chi, prominence=0.1 * chi.max())
print(f"number of peaks above prominence threshold: {len(peaks)}")
for k in peaks:
    print(f"  peak at sigma = {sigma[k]:.3f}, height = {chi[k]:.3f}"
        )
```

If `find_peaks` returns more than one peak with comparable prominence, run the framework's `detect_multipeak` routine, which returns all of them with individual bootstrap CIs.

35.2 Multipeak in practice

Two domains in this book have genuine multipeak structure as the *expected* outcome:

- *GPU cache transitions* (Chapter 26): L1, L2, HBM each give their own χ peak. The framework's `detect_cache_transitions` returns all three by default.
- *Multi-stage protein unfolding*: some proteins unfold via intermediate states; each intermediate is a transition. The `ProteinAdapter`'s `multistage=True` flag returns them all.

In every other domain, finding two comparable peaks is a signal that either there is an unexpected second regime (publish it!) or the sweep is contaminated (debug it).

35.3 Reporting multipeak results

When you do report multiple peaks, treat each one as a separate result. Each gets its own $\sigma_c^{(i)}$, $\kappa^{(i)}$, bootstrap CI, and permutation p-value. Bonferroni-correct your significance threshold by the number of peaks: $\alpha_{\text{per-peak}} = 0.05/n_{\text{peaks}}$.

TAKEAWAY

A multipeak structure is a feature when expected (GPU caches, protein intermediates) and a warning when not. The recipe's silent default of picking only the highest is correct for monotone-ish problems and wrong for multimodal ones; `detect_multipeak` is the cure.

Part VI

Validation, statistics, and rigour

Chapter 36
The boundary check

Before reporting a σ_c, ensure that a peak *can* exist: `check_boundary_conditions(O, sigma)` checks that $O(\sigma_{\min})$ is significantly larger than $O(\sigma_{\max})$, the range is non-trivial, and χ has an interior maximum. If any of these fail, the recipe is ill-posed and reporting σ_c is meaningless.

Chapter 37
The permutation test

Null hypothesis. The observed peak clarity κ is no larger than what one would obtain by randomly reordering the data. Equivalently: "the apparent peak in χ could have come from any random sequence with the same set of values".

Procedure. For $B = 10\,000$ random permutations of the O array (keeping the σ axis fixed), recompute κ. The p-value is the fraction of permutations whose κ exceeds the observed value.

Decision. $p < 0.05$ rejects the null and supports a real peak. The framework ships this in `permutation_test`.

37.1 A worked example with numbers

Imagine a sweep over $n = 30$ points with an observable that decreases through a clear sigmoid. Suppose the framework reports an observed $\kappa_{\text{obs}} = 5.7$. We want to know: how often would random shuffling of these 30 observations produce $\kappa \geq 5.7$ by chance?

```
import numpy as np
from scipy.ndimage import gaussian_filter1d
rng = np.random.default_rng(0)

# Synthetic sigmoid with shot noise:
sigma = np.linspace(0, 1, 30)
true_O = 1.0 / (1.0 + np.exp(15 * (sigma - 0.5)))
O = true_O + 0.05 * rng.standard_normal(30)

def kappa_of(sigma_vals, O_vals, kernel=0.6):
    s = gaussian_filter1d(O_vals, kernel)
    chi = np.abs(np.gradient(s, sigma_vals))
    return float(chi.max() / chi.mean())

kappa_obs = kappa_of(sigma, O)
print(f"kappa_obs = {kappa_obs:.2f}")      # e.g. 5.7

B = 10_000
null = np.empty(B)
for b in range(B):
    permuted_O = rng.permutation(O)
    null[b] = kappa_of(sigma, permuted_O)
```

```
23
24 p = (np.sum(null >= kappa_obs) + 1) / (B + 1)
25 print(f"null mean: {null.mean():.2f}, 95th pct: {np.percentile(null,
       95):.2f}")
26 print(f"p-value:   {p:.4f}")
27 # Typical: null mean ~ 2.1, 95th pct ~ 2.9, p well below 0.001
```

Interpretation. If $p = 0.0003$, only 3 in 10 000 random shuffles produce a κ as large as ours. That is strong evidence the observed peak reflects structure in the data, not chance.

37.2 Which null model for which domain?

The permutation null above is the right choice when the only thing you want to rule out is "ordered σ axis is a coincidence". In practice you often want to rule out something stronger.

domain	most natural null model	rationale
Quantum (controlled γ sweep)	full random permutation	shots i.i.d.; ordering is your choice
Magnetism (Monte Carlo sweep)	full random permutation	same as above
Finance (time series)	block permutation, $\ell \approx 20$ days	autocorrelated; preserve local structure
Seismology (catalogue)	circular block permutation	temporal clustering of aftershocks
Climate (reanalysis time series)	block permutation, ℓ = synoptic timescale	atmospheric persistence
GPU benchmarks (sequential)	block permutation, ℓ = warm-up window	thermal carry-over
ML (LR-range test)	full permutation acceptable	you control the LR schedule
Protein (mutation panel)	label permutation (mutation $\leftrightarrow$ phenotype)	each mutation independent
Linguistics (ED bins)	label permutation within ED level	word identity matters
Number theory (resolution sweep)	not applicable	deterministic, no noise
LLM cost (model list)	not applicable	$\sim$ 10 candidates, exact enumeration

Block permutation, in one paragraph. Instead of shuffling O_i independently across all positions, shuffle contiguous blocks of length ℓ. This preserves short-range autocorrelation that a naive shuffle destroys, giving an honestly larger p-value. The framework's `permutation_test(..., block_size=`ℓ`)` implements this; if you do not pass a block size, the test defaults to full permutation. Pick ℓ as the lag at which the autocorrelation of O drops below 0.1.

Chapter 38

Peak clarity threshold

Empirical thresholds restated:

- $\kappa < 1.5$: insufficient evidence; do not publish σ_c.
- $1.5 \leq \kappa < 3.0$: marginal; supplement with cross-platform or cross-observable check.
- $3.0 \leq \kappa < 8.0$: clear peak.
- $\kappa \geq 8.0$: critical-like behaviour — exceptional sharpness.

Chapter 39

Fisher information bound

Cramer–Rao gives a fundamental lower bound on susceptibility: $\chi(\sigma) \geq |dg/d\sigma|/\sqrt{I_F}$ where I_F is the Fisher information of the observable given the parameter. If your measured χ saturates the bound, you have an optimal observable. The framework computes both in `fisher_information_bound`.

Chapter 40

Multiple testing

If you run the σ_c recipe on N independent datasets, you will "find" false peaks by chance. Bonferroni correction: divide your significance threshold by N. The magnetism paper used $\alpha = 0.05/120 = 4.2 \times 10^{-4}$ (six experiments times twenty average peak searches).

Chapter 41

Cross-validation across observables

The strongest evidence that σ_c is a property of the system, not of your chosen observable, is to repeat the analysis with a different observable and see that σ_c matches. In the magnetism paper, the entanglement witness W, its shift $W + C$, and its square W^2 all yielded $\gamma_c = 0.6737$ (coefficient of variation 0%). *This is the gold standard.*

Part VII

Worked examples and recipes

Chapter 42

Recipe: minimal σ_c in pure NumPy/SciPy

```python
import numpy as np
from scipy.ndimage import gaussian_filter1d

def sigma_c(sigma, observable, kernel=0.6):
    smooth   = gaussian_filter1d(observable, kernel)
    chi      = np.abs(np.gradient(smooth, sigma))
    idx      = int(np.argmax(chi))
    return {
        'sigma_c': float(sigma[idx]),
        'kappa':   float(chi.max() / chi.mean()),
        'chi':     chi,
        'smoothed': smooth,
    }
```

That is the entire core. Save it as `sigma_c_mini.py`; it has no dependency on any framework.

Chapter 43

Recipe: bootstrap CI on σ_c

```python
def bootstrap_ci(sigma, observable, n_boot=1000, kernel=0.6):
    out = []
    n = len(sigma)
    rng = np.random.default_rng(42)
    for _ in range(n_boot):
        idx = rng.integers(0, n, size=n)
        s, o = sigma[idx], observable[idx]
        order = np.argsort(s)
        s, o = s[order], o[order]
        out.append(sigma_c(s, o, kernel=kernel)['sigma_c'])
    return np.percentile(out, [2.5, 97.5])
```

Chapter 44

Recipe: choosing the kernel

```
import matplotlib.pyplot as plt
def kernel_sweep(sigma, O, kernels=(0.3, 0.5, 0.8, 1.2, 2.0)):
    fig, ax = plt.subplots(figsize=(6,4))
    for k in kernels:
        res = sigma_c(sigma, O, kernel=k)
        ax.plot(sigma, res['chi'], label=f"k={k}, sigma_c={res['sigma_c']:.3f}")
    ax.legend()
    ax.set_xlabel('sigma'); ax.set_ylabel('chi')
    plt.show()
```

If σ_c shifts dramatically as you change k, your data is too noisy or the peak too sharp; increase n or use Savitzky–Golay differentiation.

Chapter 45

Recipe: install the full framework

```
# Core
pip install sigma-c-framework

# With quantum integrations (Braket, Qiskit, PennyLane)
pip install "sigma-c-framework[quantum]"

# With GPU acceleration (CuPy)
pip install "sigma-c-framework[gpu]"

# Verify
python -c "import sigma_c; print(sigma_c.__version__)"
# Expected: 3.0.0
```

Chapter 46

Recipe: build your own adapter

```
from sigma_c.core.base import SigmaCAdapter
import numpy as np

class MyAdapter(SigmaCAdapter):
    def get_observable(self, data, **kwargs):
        # Domain-specific: turn raw data into a scalar
        return float(np.mean(data))

    def _domain_specific_diagnose(self, data=None, **kw):
        return {'status': 'ok', 'issues': [],
                'recommendations': [], 'details': {}}

    def _domain_specific_validate(self, data=None, **kw):
        return {'basic': True}

    def _domain_specific_explain(self, result, **kw):
        return f"sigma_c = {result['sigma_c']:.3f}, kappa = {result[
            'kappa']:.2f}"

# usage
adapter = MyAdapter()
result = adapter.compute_susceptibility(sigma_array,
    observable_array)
```

Three method overrides; everything else (diagnose, auto-search, explain) inherits the universal behaviour.

Chapter 47
Recipe: an end-to-end synthetic experiment

```python
import numpy as np
from sigma_c import Universe

# Synthesise Ising-like magnetisation curve
T = np.linspace(1.0, 4.0, 60)
Tc_true = 2.5
M = np.where(T < Tc_true, np.abs(Tc_true - T)**0.35, 0.0) + 0.02 *
    np.random.randn(60)

mag = Universe.magnetic()
res = mag.compute_susceptibility(T, M, kernel_sigma=0.7)
print(f"  detected Tc = {res['sigma_c']:.3f}")
print(f"  true Tc     = {Tc_true}")
print(f"  kappa       = {res['kappa']:.2f}")

# Validate
val = mag.compute_susceptibility(T, M, kernel_sigma=0.7, validate=
    True)
print(f"  passes peak test: {val.get('peak_clarity_passes')}")
print(f"  observable quality score: {val.get('observable_quality')
    :.2f}")
```

Chapter 48

Recipe: live monitoring with streaming sigma_c

```
from sigma_c.core.control import StreamingSigmaC, AdaptiveController

stream  = StreamingSigmaC(window_size=200)
control = AdaptiveController(target_sigma=0.7, kp=1.0, ki=0.1, kd
    =0.05)

for t in range(10_000):
    measurement = sensor.read()
    current     = stream.update(parameter=t, observable=measurement)
    delta       = control.compute_correction(current)
    actuator.set(actuator.get() + delta)
```

Welford's algorithm updates σ_c in $O(1)$ per sample, the PID controller drives the parameter back to its target $\sigma_c = 0.7$. Useful in real-time experimental setups (e.g. keeping a laser cavity at the operational stability point).

Chapter 49

Recipe: report a result for publication

1. Boundary check passes? (`check_boundary_conditions`)
2. Observable quality score ≥ 0.75? (`observable_quality_score`)
3. Permutation p-value < 0.05 after Bonferroni?
4. Bootstrap CI on σ_c tight enough for your claim?
5. Cross-observable check: same σ_c under a different, Fisher-aligned observable?
6. Robustness to kernel: σ_c stable across kernels 0.3–1.5?

Report all six. The community reviewer will look for them explicitly.

Appendix A

Symbol glossary

σ	control parameter (any domain). Doubles as the standard deviation of a Gaussian in any other book; we apologise to the statisticians
O	observable (any scalar function of σ)
$\chi(\sigma) = \|dO/d\sigma\|$	generalised susceptibility
σ_c	location of the peak of χ. The number the framework is paid to find
κ	peak clarity (default: $\chi_{\max}/\bar{\chi}$). Higher is better in this book and in cocktails generally
D	contraction defect $= \|S\|/\|f(S)\|$
γ	drift (geometric mean of $f(x)/x$. In number theory it is shrinkage; in quantum hardware it is noise; on Tuesdays in the authors' notebooks it is occasionally both at once
$\Pi = D\gamma$	contraction product (universal threshold)
k_B	Boltzmann constant 1.380649×10^{-23} J/K. The only symbol in the book whose value is fixed by international agreement
ξ	operational correlation length
σ_{ker}	Gaussian kernel width (default 0.6). The one σ in the glossary that does refer to a Gaussian

Appendix B

Proof: existence of an interior maximum

Proof. Let $O\colon [a,b] \to \mathbb{R}$ be continuous with $O(a) > O(b)$ and, by assumption, not monotonic near either endpoint. Define $\chi(\sigma) = |dO/d\sigma|$ (where the derivative exists; elsewhere take the sup-norm of the absolute slope over an ϵ-neighbourhood).

By the Mean Value Theorem applied on $[a,b]$, there exists $\sigma^* \in (a,b)$ with $O'(\sigma^*) = (O(b) - O(a))/(b-a) < 0$. Hence $\chi(\sigma^*) > 0$.

By continuity of O and the non-monotonicity assumption, in any neighbourhood of a there is a sub-interval on which $O' \approx 0$, so $\chi(a) < \chi(\sigma^*)$. The same holds near b. Therefore χ achieves its maximum on the compact interval $[a,b]$ at some interior point $\sigma_c \in (a,b)$. □ □

Appendix C

The twelve qn+c maps, predictions and observations

map	D	γ	$D\gamma$	predicted	observed
$3n+1$ (cycle)	2.06	0.563	1.16	converges (cycles)	all known orbits reach 1
$3n+1$ (single)	1.71	0.750	1.28	converges (cycles)	matches
$5n+1$	1.43	1.250	1.79	diverges	matches
$7n+1$	1.60	1.750	2.80	diverges	matches
$3n-1$	1.33	0.750	1.00	critical/cyclic	cycles
$3n+3$	2.00	0.750	1.50	cycles	matches
$3n+5$	1.48	0.750	1.11	cycles	matches
$3n+7$	1.56	0.750	1.17	cycles	matches
$9n+1$	1.34	2.250	3.02	diverges	matches
$11n+1$	1.37	2.750	3.77	diverges	matches
$5n+3$	1.36	1.250	1.70	diverges	matches
$5n-1$	1.43	1.250	1.79	diverges	matches

Appendix D
Annotated reading list

A handful of works underpin everything in this book. They are listed here with one-line annotations so you can pick where to go next.

The seed paper. **M. C. Wurm,** *Operational scale detection in quantum magnetism via susceptibility analysis* (Wurm, 2026). AVS Quantum Science **8**, 013804 (2026). DOI 10.1116/5.0312410. The peer-reviewed empirical anchor of the framework; Chapter 21 is a reading guide for it.

Classical context. **H. E. Stanley,** *Introduction to Phase Transitions and Critical Phenomena* (Stanley, 1971). The textbook from which the word "susceptibility" inherits everything; required only if you want the equilibrium-statistical-mechanics origin story.

L. Onsager, *Crystal Statistics, I* (Onsager, 1944). The 1944 exact solution of the 2D Ising model, giving us $T_c = 2J/\ln(1+\sqrt{2})$. The standard against which Chapter 22 is calibrated.

J. M. Kosterlitz & D. J. Thouless, *Ordering, metastability and phase transitions in two-dimensional systems* (Kosterlitz and Thouless, 1973). A complement to the susceptibility view: topological transitions without a divergent χ.

Information and Fisher. **R. Landauer,** *Information is Physical* (Landauer, 1991). The one-page article that gives Chapter 31's $\log_2 D$ connection its thermodynamic teeth.

C. W. Helstrom, *Quantum Detection and Estimation Theory* (Helstrom, 1976). The Fisher-information bound on susceptibility used in the Fisher-information chapter of Part VI.

S. L. Braunstein & C. M. Caves, *Statistical distance and the geometry of quantum states* (Braunstein and Caves, 1994). The 1994 paper that connects the QFI to differential geometry of the state space, providing the rigorous backing for the framework's $\chi \approx \sqrt{F_Q}$ identity.

Hardware era. **J. Preskill,** *Quantum computing in the NISQ era and beyond* (Preskill, 2018). Context for why the operational scale detected in Chapter 21 matters.

Methodology. **B. Efron,** *Bootstrap methods: another look at the jackknife* (Efron, 1979). The original bootstrap paper, foundational for Chapter 8.

A. Savitzky & M. J. E. Golay, *Smoothing and differentiation of data by simplified least squares procedures* (Savitzky and Golay, 1964). The 1964 paper that gives the framework's Savitzky–Golay derivative option its name.

S. Williams, A. Waterman & D. Patterson, *Roofline* (Williams et al., 2009). The 2009 paper that defines the performance ceiling used in Chapter 26.

Specific domains. **B. Gutenberg & C. F. Richter,** *Frequency of earthquakes in California* (Gutenberg and Richter, 1944); **F. Omori,** *On the aftershocks of earthquakes* (Omori, 1894). The two empirical laws of Chapter 24.

G. D. Nastrom & K. S. Gage, *A climatology of atmospheric wavenumber spectra of wind and temperature observed by commercial aircraft* (Nastrom and Gage, 1985). The data set behind the mesoscale transition of Chapter 25.

L. N. Smith, *Cyclical learning rates for training neural networks* (Smith, 2017); **S. McCandlish et al.,** *An empirical model of large-batch training* (McCandlish et al., 2018). Reference for the LR-range test and the critical-batch-size phenomenon of Chapter 27.

J. C. Lagarias, *The 3x+1 problem and its generalizations* (Lagarias, 1985); **T. Tao,** *Almost all orbits of the Collatz map attain almost bounded values* (Tao, 2022). Background and state of the art for Chapter 30.

The framework itself. **ForgottenForge,** `sigma-c-framework` v3.0.0 (ForgottenForge, 2026). The accompanying open-source library.

Bibliography

S. L. Braunstein and C. M. Caves. Statistical distance and the geometry of quantum states. *Physical Review Letters*, 72(22):3439, 1994.

Bradley Efron. Bootstrap methods: Another look at the jackknife. *Annals of Statistics*, 7(1): 1–26, 1979.

ForgottenForge. sigma-c-framework v3.0.0: Universal criticality analysis and active control. https://pypi.org/project/sigma-c-framework/, 2026.

B. Gutenberg and C. F. Richter. Frequency of earthquakes in california. *Bulletin of the Seismological Society of America*, 34(4):185–188, 1944.

Carl W. Helstrom. *Quantum Detection and Estimation Theory*. Academic Press, 1976.

J. M. Kosterlitz and D. J. Thouless. Ordering, metastability and phase transitions in two-dimensional systems. *Journal of Physics C: Solid State Physics*, 6(7):1181–1203, 1973.

Jeffrey C. Lagarias. The 3x+1 problem and its generalizations. *American Mathematical Monthly*, 92(1):3–23, 1985.

Rolf Landauer. Information is physical. *Physics Today*, 44(5):23–29, 1991.

Sam McCandlish, Jared Kaplan, and Dario Amodei. An empirical model of large-batch training. *arXiv preprint arXiv:1812.06162*, 2018.

G. D. Nastrom and K. S. Gage. A climatology of atmospheric wavenumber spectra of wind and temperature observed by commercial aircraft. *Journal of the Atmospheric Sciences*, 42 (9):950–960, 1985.

F. Omori. On the aftershocks of earthquakes. *Journal of the College of Science, Imperial University of Tokyo*, 7:111–200, 1894.

Lars Onsager. Crystal statistics. I. a two-dimensional model with an order-disorder transition. *Physical Review*, 65:117–149, 1944.

John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018. doi: 10.22331/q-2018-08-06-79.

Abraham Savitzky and Marcel J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964.

Leslie N. Smith. Cyclical learning rates for training neural networks. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 464–472, 2017.

H. Eugene Stanley. *Introduction to Phase Transitions and Critical Phenomena*. Oxford University Press, 1971.

Terence Tao. Almost all orbits of the Collatz map attain almost bounded values. *Forum of Mathematics, Pi*, 10:e12, 2022.

Samuel Williams, Andrew Waterman, and David Patterson. Roofline: an insightful visual performance model for multicore architectures. *Communications of the ACM*, 52(4):65–76, 2009.

M. C. Wurm. Operational scale detection in quantum magnetism via susceptibility analysis: Critical-like behavior at the quantum-classical crossover on nisq hardware. *AVS Quantum Science*, 8(1):013804, 2026. doi: 10.1116/5.0312410.

Index

amyloid, 145
Ankaa-3, 86

bootstrap, 43

χ, 57
Collatz conjecture, 139
contraction defect, 75
contraction product, 78
Cramer–Rao bound, 172
Curie point, 97

D (contraction defect), 75
derivative, 28
drift, 76

etymological depth, 154

finance, 103
Fisher information, 172

γ (drift), 76
GARCH, 103
Gaussian smoother, 35
GPU, 117
Gutenberg–Richter law, 108

Hurst exponent, 103

Ising model, 97

κ (peak clarity), 63

learning rate, 123
linguistics, 154
LR-range test, 123

machine learning, 123
magnetism, 97

number theory, 139

Omori's law, 108

peak clarity, 63
permutation test, 169
$\Pi = D \cdot \gamma$, 78
protein, 145

qn+c maps, 139
quantum hardware, 86

roofline model, 117

seismology, 108
semantic change, 154
smoothing
    Gaussian, 35
susceptibility, generalised, 57

thermal throttling, 117
TTR, 145

V30M, 145

Wurm 2026, 86

Appendix E

Reproducibility notes

All numbers cited in Part IV come from the corresponding files in the `sigma-c-framework` v3.0.0 release. To reproduce the Wurm 2026 result from raw data:

```
git clone https://github.com/forgottenforge/magneto
cd magneto
python magnetvali.py --experiment E3 --bootstrap 1000
# Output: gamma_c = 0.6737 +/- 0.036, kappa = 8.58
```

For all twelve qn+c maps:

```
from sigma_c import Universe
nt = Universe.number_theory(map_type='collatz')
print(nt.verify_twelve_predictions(M=12))
```

Appendix F
Where to get the Chapter 9 data

Chapter 9 works out the recipe on three datasets: the 2D Ising magnetisation, the S&P 500 returns of 2008, and the Southern California earthquake catalogue around the 2019 Ridgecrest sequence. If you want to reproduce any of them at home, here is where to start. Total time investment: ten to twenty minutes per dataset, no GPU required.

F.1 Curie point (Ising magnetisation)

The Ising data in Chapter 9 was *simulated*, not downloaded: the Metropolis Monte Carlo from Chapter 22 run at each of the ten temperatures. *You generate the data yourself.*

- **Code**: the `ising_mc` function from Section 22, ~ 20 lines of NumPy.
- **Compute**: a laptop runs the full sweep in about ninety seconds at $L = 16$.
- **Excel alternative**: not feasible — the Monte Carlo loop needs a real programming language. If you do not have Python, install it from `python.org`; it is free and takes ten minutes.

F.2 S&P 500 daily returns (the 2008 example)

- **Source**: Yahoo Finance, ticker `^GSPC`. Free, no login.
- **Python one-liner**:

```
import yfinance as yf
df = yf.download('^GSPC', start='2008-01-01', end='2009-06-30')
```

- **Excel**: log into `finance.yahoo.com`, search `^GSPC`, click "Historical Data", set the date range, click "Download". You get a CSV with Date, Open, High, Low, Close, Adj Close, Volume. In Excel, compute `=LN(Close_today) - LN(Close_yesterday)` for log-returns; this is the column the recipe of Section 9.2 operates on.
- **Note**: the same workflow works for any ticker; try `BTC-USD` or `GC=F` (gold) as comparisons.

F.3 Southern California earthquake catalogue (Ridgecrest)

- **Source**: SCEC Earthquake Data Center at `scec.org/research-tools/downloadable-catalogs`.
Choose the *Hauksson catalog*, date range 2018-01 to 2021-06, magnitude floor 2.5.

- **Format**: tab-separated text file with columns date, time, latitude, longitude, depth, magnitude.
- **Python load**:

```
import pandas as pd
cat = pd.read_csv('hauksson.txt', sep='\s+', skiprows=1,
                  names=['date','time','lat','lon','depth','mag'])
```

- **Excel**: open the file directly, use "Text to Columns" to split. The b-value calculation is `=0.4343/(AVERAGE(F:F)-MIN(F:F))`, where column F is magnitude in a filtered six-month window. Repeat for each window centre to reproduce Section 9.3.

If you have never installed Python. Go to `python.org/downloads`, install the latest 3.x. Open a terminal and run `pip install numpy scipy yfinance pandas matplotlib`. That covers every snippet in this book. If anything fails, the `sigma-c-framework` repository's README has fall-back instructions in five operating systems.

If you do not want to install anything. `kaggle.com`, `colab.research.google.com`, and `deepnote.com` all run Python in a browser tab with the above packages pre-installed. Free tier is enough for every chapter except a real-hardware Braket run.

Appendix G
Acknowledgements

This compendium would not exist without the prior work documented in the `sigma-c-framework` repository (v3.0.0), the empirical data collected on Rigetti Ankaa-3 and Cepheus-1 quantum processors, and the peer review process that sharpened the methodology to its current form.

The reviewers, by name. Three readers' voices shaped this manuscript across two revision rounds. *Sabine*, programme editor at a Bavarian university press, made me retire the long title and place the seven-symbols card where it belongs. *Linus*, sixteen and then seventeen, told me Part III was unreadable until the coffee mug came back — and was right. *T.P.*, the anonymous theorist whose two letters live in a desk drawer, named four conjectures that the manuscript would not admit. If anything in the book reads as honest, the credit is split four ways. The errors are mine.

Thanks to my family for endless inspiration.

Appendix H
Colophon

Title. *The Peak: a universal recipe for phase transitions, in twelve worlds.* The title was workshopped over two manuscript rounds and twelve months. Earlier drafts circulated under a longer working description; it is retired with this edition and not preserved elsewhere.

Composition. Typeset in LaTeX with `lmodern`. Body text in Latin Modern Roman at 11 pt; section headings in Latin Modern Sans Serif; chapter titles italicised. Figures in TikZ. Code in Latin Modern Mono via the `listings` package with `upquote` active.

License. Open source under the GNU Affero General Public License, version 3 or later (AGPL-3.0-or-later). Commercial licences without the AGPL obligation — chiefly for closed-source derivative works that incorporate the accompanying framework — are available at nfo@forgottenforge.xyz. The text of the AGPL-3.0 license is reproduced verbatim in the framework's source repository.

Reproducibility. Every numerical result in this book is reproducible from the v3.0.0 release of `sigma-c-framework`. Worked examples that reference quantum hardware data are pinned to a specific commit of the `magneto` repository. Tagged Git revisions for both repositories are listed in Appendix E.

Errata. Errata are maintained at `https://forgottenforge.xyz/compendium/errata` and are incorporated into subsequent printings. Send corrections to nfo@forgottenforge.xyz with the subject line `[compendium errata]`.

Edition and contact. First edition, May 2026. Buckenhof, Germany. The author is reachable at the address above; questions of any depth are welcome, and most are answered within a week.

Back-cover text (printed edition). For reviewers and library cataloguers, the back-cover text in the printed edition is reproduced here verbatim.

A recipe. Three lines of Python that take a system with a tunable knob and a measurable response and tell you where it tips over.

The same three lines locate the Curie point of iron near 1043 *K (Chapter 22), the volatility regime that gave way the week Lehman Brothers fell in August 2008 (Chapter 23), and the b-value drop in the southern-California seismic catalogue that preceded the Ridgecrest sequence in July 2019 (Chapter 24). Twelve domains, one recipe.*

A textbook that begins with arithmetic and ends with publishable diagnostics. We do not promise a revolution. We promise a good butler.