



A New Perspective on Value

Quan-Hoang Vuong

Minh-Hoang Nguyen

A New Perspective on Value

Quan-Hoang Vuong^{1,2,3}

Minh-Hoang Nguyen¹

¹The Centre for Interdisciplinary Social Research, Phenikaa University, Hanoi, Vietnam

²The University College, Korea University, Seoul, South Korea

³Centre Emile Bernheim de Recherche Interdisciplinaire en Gestion, Université Libre de Bruxelles, Brussels, Belgium

Manuscript v31

September 4, 2026

This material is also available from:

SSRN: <https://papers.ssrn.com/abstract=7386118>



© 2026 AISDL

Preface

“Sir Kingfisher, you told me earlier that the Nightingales have the right argument. How can we now lose?”

“True. The Nightingales have the right argument. But the Magpies have multiple larger scale right arguments. Understand?”

— in “Righteous Judge-bird,” *Wild Wise Weird* (2022)

This book grows out of a single short paper, *Informational Entropy-Based Value Formation: A New Paradigm for a Deeper Understanding of Value* (Vuong, La, and Nguyen, 2025), hereafter referred to throughout as the source paper. That paper introduces, largely in discursive and qualitative terms, a way of thinking about value as something that emerges from how a mind or a society distributes probability across competing pieces of information, rather than something residing in objects, labor, or utility functions taken in isolation. It draws its vocabulary from quantum mechanics, from Shannon’s information theory, and from the authors’ own mindsponge theory, and it argues, convincingly in our view, that a genuinely interdisciplinary account of value needs some such information-processing foundation.

What the source paper does not do, and does not claim to do, is derive its central images from mathematics. Statements such as “entropy is highest when all information is equally important” are asserted rather than proved. This is not a defect in that paper, which is explicitly a short communication aimed at motivating a new perspective rather than a technical monograph. But it leaves open a natural question: if the qualitative claims are taken seriously, what is the smallest, most standard body of mathematics that would make them precise, and what, if anything, would that mathematics let us say that the qualitative version cannot?

This book is, for the great majority of its length, an attempt to answer that question in full, and is best understood there as a formal reconstruction, not a work of independent discovery. Every substantial mathematical tool used in that reconstruction, the Gibbs distribution, the maximum entropy principle, Kullback-Leibler divergence, cluster expansions and their diagrammatic bookkeeping, is a standard result from probability theory, information theory, or statistical mechanics, proved here from first principles for the reader’s convenience rather than presented as new. The contribution, through that point, lies entirely in the mapping: in showing which piece of elementary mathematics corresponds to which qualitative claim in the source paper, and in following that mapping far enough to say something concrete about a genuinely difficult applied problem, the valuation of an asset, Bitcoin, that classical economic theories struggle to explain.

One part of this book is not a reconstruction, and we want to say so plainly here rather than let the reader discover it unannounced partway through. Chapters 13 and 14, gathered as Part III, *A New Perspective on Environmental Values*, develop original material. They build, from

Brownian motion up, a continuous-time stochastic model of irreversible ecological collapse, culminating in a closed-form valuation theorem in the same technical family as the Black-Scholes equation. We want to be clear about what this borrowing of financial vocabulary is for. Part III is not an attempt to found a new economic theory, or to argue that ecological value should literally be traded as a financial instrument; option-pricing mathematics is used here because it is a familiar, precise language for reasoning about irreversibility and collapse risk, and borrowing that language is a means of sharpening ecological intuition, not an end in itself. This work does not reconstruct anything in Vuong et al. (2025a); it formalizes a different, related idea, the semiconducting principle of monetary and environmental value exchange due to Vuong (2021), that money converts to environmental value only in one direction. The result is a theorem, proved without qualification, that any valuation ignoring an ecological asset's collapse risk necessarily overstates its worth. Its title deliberately echoes this book's own: where the rest of the book asks what a new perspective on value in general would look like, mostly by reconstruction, Part III asks the same question specifically for environmental value, and answers it with a result that is this book's own. A reader interested primarily in that answer, rather than in the reconstruction project surrounding it, may reasonably treat Part III as the book's central contribution and read the rest as the context that motivated it; we say more about how to read the book that way in Chapter 1. We flag the distinction between reconstruction and original contribution clearly again at the start of Part III itself, so that a reader can tell at every point in the book which kind of claim they are reading.

We have tried throughout to keep two commitments in tension. The first is mathematical honesty: nothing is asserted here without a proof, mostly relegated to the appendix-length chapters of Part II so that the main argument in Part III and the original material of Part III can be read without constant interruption. The second is interpretive honesty: at every point where we translate a piece of mathematics into a claim about minds, markets, or social systems, we have tried to say plainly that this translation is an interpretive choice on our part, not a theorem, and that reasonable readers might translate differently. Readers who want the mathematics without the interpretation, or the interpretation without full derivations, should find both cleanly separable.

The book assumes no background beyond calculus and elementary probability, with one exception: Part III's original material additionally builds up the stochastic calculus it needs, Brownian motion and Itô's lemma, from the same starting point, so no prior exposure to that material is assumed either. Part II builds every tool needed for the reconstruction from that starting point, including the statistical mechanics and combinatorics that a physics or econophysics reader would already have. Readers who already know this material may prefer to treat Part II as reference and move directly to Part III; readers interested primarily in the book's original contribution may move directly to Part III after Part III, since Part III does not depend on the applications gathered in Chapter 12.

Beyond the mathematics, we hope this book also contributes some thought and perspective, however modest, to the grave and mounting problems of environmental crises and ecological degradation. That hope is not confined to any single section; it runs through several of this book's central chapters, from the irreversible collapse and overvaluation results of Part III, to the biodiversity-as-cultural-information argument of Chapter 15, to the bird conservation application of Chapter 12 and the self-contained Extraction Limit of Appendix A. This book does not purport to resolve the ecological crisis; what it offers, at most, is a precise language for examining why ecosystem degradation and biodiversity loss cannot always be treated as reversible or compensable, a distinction the semiconducting principle and its γ and δ parameters give exact, estimable content rather than leaving as a slogan. We offer this hope with the same humility as everything else in this preface: as a mathematical perspective that might sharpen how these problems are discussed and support researchers, policymakers, and communities in developing more rigorous approaches to prevention and sustainability, not as a substitute for the ecological, political, and economic work of addressing them.

Notation

The table below collects the symbols used throughout the book, together with the section in which each is first introduced. Symbols local to a single worked example are not included here and are defined where they occur.

Symbol	Meaning	First used
n	number of informational quanta in the mind, finite	Ch. 7
x_i	the i -th informational state or quantum	Ch. 2
P $(p_1, \dots, p_n)$	= probability distribution over quanta; $p_i = P(X = x_i)$	Ch. 2
Δ^{n-1}	probability simplex on n outcomes	Ch. 7
U	uniform distribution on n outcomes	Ch. 2
$H(P)$	Shannon entropy of P	Ch. 2
$D_{KL}(P\ Q)$	Kullback–Leibler divergence of P from Q	Ch. 2
E_i	informational energy of quantum i	Ch. 4
J_{ij}, J_{ijk}	pairwise, triple-wise interaction couplings	Ch. 8
T	openness / temperature parameter	Ch. 4
β	inverse temperature, $\beta = 1/T$	Ch. 4
Z	partition function	Ch. 5
F	free energy, $F = -T \log Z$	Ch. 5
Mind, Environment	sets of informational quanta with Mind $\subseteq$ Environment	Ch. 9
Q_{val}	a value quantum, a quantum retained with high probability	Ch. 11
$\partial P / \partial E_i$	elasticity of the distribution with respect to quantum i 's energy	Ch. 10
SOV, AUT, SSV, HSV	Store-of-Value, Autonomy, Social-Signal Value, Hedonic-Sentiment Value (Bitcoin case study)	Ch. 12

Contents

Preface	i
Notation	iii
1 Why Value Needs a New Paradigm	1
1.1 Granular Interaction Thinking Theory	1
1.2 What Is Missing	2
1.3 How This Book Is Built: Scaffolding Toward Part III	3
I Mathematical Foundations	4
2 Probability, Entropy, and Information	5
2.1 A Toy Mind	5
2.2 Discrete Probability Distributions	5
2.3 Shannon Entropy	6
2.4 Kullback-Leibler Divergence and the Proof of Proposition 2.1	6
2.5 Summary	7
3 Constrained Optimization and the Maximum Entropy Principle	8
3.1 Lagrange Multipliers, from Scratch	8
3.2 Maximizing Entropy Subject to Normalization Alone	9
3.3 Adding a Second Constraint	9
3.4 Summary	11
4 The Gibbs Distribution	12
4.1 Definition and Basic Properties	12
4.2 Monotonicity in Temperature	13
4.3 Summary	14
5 Partition Functions and Free Energy	15
5.1 Free Energy	15
5.2 The Deferred Proof: Entropy Increases with Temperature	16
5.3 Summary	17
6 Generating Functions and Diagrammatic Expansions	18
6.1 Ordinary Generating Functions, Briefly	18
6.2 Exponential Generating Functions	18
6.3 The Exponential Formula	19
6.4 Application: The Partition Function of Interacting Quanta	19
6.5 Diagrams	20

6.6	Summary	21
II	The Formal Model	22
7	Granularity as a Finite Probability Simplex	23
7.1	The Mind as a Point on the Simplex	23
7.2	A Capacity Bound	23
7.3	Interpretation: Scarcity as a Mathematical Necessity, Not a Psychological Fact .	24
7.4	Granularity Under the Gibbs Distribution	24
7.5	Summary	25
8	Relationality via Interaction Diagrams	26
8.1	The General Interaction Energy	26
8.2	The Full Partition Function and Its Diagrammatic Decomposition	26
8.3	What Relationality Buys, Formally	27
8.4	Summary	28
9	Indeterminacy as Stochastic Filtering	29
9.1	The Environment as a Larger Candidate Set	29
9.2	Retention as a Bernoulli Process	29
9.3	Tying Retention Probability to Energy	30
9.4	Correlated Retention	30
9.5	Summary	31
10	Elasticity and Belief as Sensitivity	32
10.1	Elasticity as a Derivative	32
10.2	Recovering the Four Qualitative Factors	33
10.3	Elasticity with Respect to Temperature	33
10.4	Summary	34
11	Value Formation and Entropy Reduction	35
11.1	Value Formation, Defined	35
11.2	The Value Formation Theorem	35
11.3	KL Divergence as a Value Intensity Index	36
11.4	Connecting Back to Herding and Borrowed Conviction	36
11.5	Summary	37
12	Applications	38
12.1	Bitcoin Valuation	38
12.2	Phase Transitions: Bubbles and Crashes	41
12.3	Bird Conservation and Cultural Belief Drift	43
III	A New Perspective on Environmental Values	45
13	Irreversibility: A Stochastic Model of Ecological Value	46
13.1	Why Symmetric Valuation Fails	46
13.2	Brownian Motion and Stochastic Differential Equations, from Scratch	48
13.3	Itô's Lemma	49
13.4	A Model of Ecological Value with an Absorbing Barrier	50
13.5	The Valuation Equation	50
13.6	Beyond Geometric Brownian Motion: A General Class of Models	51
13.7	An Endogenous Tipping Point: The Allee Model	52
13.8	Summary	53

14 The Cost of Ignoring Irreversibility	55
14.1 Solving via the Reflection Principle	55
14.2 The Closed-Form Valuation	56
14.3 The Overvaluation Theorem	56
14.4 Probability of Collapse	57
14.5 Comparative Statics: Collapse Vega	57
14.6 Barrier Uncertainty	61
14.7 Partial Recovery: A Falsifiable Semiconducting Principle	62
14.8 A Worked Example	63
14.9 An Optimal Conservation Investment Policy	64
14.10A Correction Factor for Conventional Valuation	67
14.11Connecting Back to Entropy	68
14.12Implications for the Climate and Ecological Crisis	68
14.13Limitations	69
14.14Summary	69
15 Biodiversity as an Irreplaceable Source of Cultural Information	70
15.1 Biodiversity as Informational Input	70
15.2 A Shrinking Environment	71
15.3 The Cultural Entropy Ceiling Theorem	71
15.4 Combinatorial Diversity Collapse	72
15.5 Beyond Richness: Hill Numbers	74
15.6 The Semiconducting Principle, Informational Version	75
15.7 Coupling to Ecological Value	76
15.8 The Missing Direction: Does Richness Affect Volatility?	78
15.9 A Worked Example	80
15.10Limitations	80
15.11Summary	80
16 Sufficiency as an Ecological Strategy	82
16.1 The Earth Sufficiency Paradigm	82
16.2 A Capacity Bound for Conservation Priorities	83
16.3 Diminishing Returns to Ecological Prioritization	83
16.4 Hysteresis Between Eco-Deficit and Eco-Surplus	84
16.5 Summary	86
17 Predictions, Scope, and Non-Claims	87
17.1 Testable Predictions	87
17.2 Scope and Non-Claims	89
Bibliography	93
A Self-Contained Notes on Detection, Drift, Belief, and Extraction	96
A.1 Setup	96
A.2 The Result	97
A.3 Why the Threshold Cannot Be Fixed	98
A.4 A Second Channel: Suspending the Comparison Itself	98
A.5 A Distributional Generalization	99
A.6 The Extraction Limit	100
A.7 Connection to Section 12.2	101
A.8 Summary	101
B A Central Limit Theorem for Extinction Counts Under Regime-Switching Stress	103
B.1 Setup	103
B.2 The Central Limit Theorem	104

B.3	From Intensity to Extinction Count	104
B.4	Relation to Chapter 15's Coupling	105
B.5	Summary	105
C	Self-Contained Notes on Prioritization and Institutional Hysteresis	107
C.1	Diminishing Returns	107
C.2	Institutional Hysteresis	108
C.3	Summary	108

List of Figures

4.1	The Gibbs distribution over the toy mind’s four quanta ($E = (0, 1, 1, 3)$) at three temperatures. As T falls, probability mass concentrates entirely on the lowest-energy quantum x_1 , illustrating Propositions 4.1 and 4.2 directly.	13
6.1	Diagrams for the exponential formula. Top: the two-quantum example, where the single-vertex diagrams and the connecting edge are the only two diagram types, matching $\log(1 + z_1) + \log(1 + z_2)$ plus the interaction term derived above. Bottom: a four-quantum system with two separate interacting pairs; the connected diagrams $\{x_1, x_2\}$ and $\{x_3, x_4\}$ each contribute separately, but their disconnected union, spanning all four vertices as a single diagram, does not appear in $\log Z$ at all.	21
12.1	Posterior mean and 95% HPDI for each of the four value dimensions’ association with next-day Bitcoin returns, redrawn from the estimates of Vuong et al. (2025b). SSV is the only interval excluding zero.	40
14.1	Naive versus barrier-adjusted valuation of an ecological asset, $B = 60$, $\mu = 0.01$, $\sigma = 0.30$, $r = 0.03$, $T = 10$. The vertical gap between the curves, guaranteed strictly positive by Proposition 14.4, is the overvaluation from ignoring irreversible collapse risk.	64
14.2	The semiconducting principle of value exchange, drawn as a diode. Environmental value converts to monetary value under forward bias (top); the reverse conversion is blocked (bottom), with only a thin, partial leakage governed by the restoration parameter γ	67

Chapter 1

Why Value Needs a New Paradigm

Every serious attempt to address the major crises of the present moment, climate change and ecological overshoot, widening inequality, and the rapid, poorly governed advance of artificial intelligence, eventually runs into a question that sounds simple and is not: what, exactly, is value, and how does it form? Environmental economists ask how to price a wetland. Financial economists ask why an asset with no cash flow can be worth a trillion dollars. Sociologists ask why some norms persist for centuries while others collapse in a decade. Each discipline has an answer, and the answers do not agree with one another, not even in spirit.

Classical economics located value in the cost of production. Neoclassical economics re-located it to marginal utility and relative scarcity. Evolutionary economics emphasized procedural uncertainty over static optimization. Bioeconomics tied value to low-entropy matter and the biological costs of sustaining life. Econophysics borrowed the vocabulary of statistical mechanics directly. Each of these traditions captures something real; none was built to talk to the others, and a theory of value that only knows how to talk about exchange prices will always struggle with things that are valued but never traded.

Nowhere is this fragmentation more visible than in the valuation of Bitcoin, which by the time of this writing has sustained a market capitalization in the trillions of dollars while generating no dividend, no interest, no productive output, and conferring no legal ownership claim on anything at all. Classical labor-value reasoning fails immediately: Bitcoin is not produced by labor in any ordinary sense, but by an algorithmic proof-of-work process. Discounted cash flow models fail for the more basic reason that there is no cash flow to discount. We return to Bitcoin at length in Section 12.1, where it serves as this book's central worked application of Parts I and II; we introduce it here because it illustrates, more sharply than any other example available at present, exactly the gap this book's early chapters try to fill.

1.1 Granular Interaction Thinking Theory

The source paper this book formalizes, Vuong, La, and Nguyen's *Informational Entropy-Based Value Formation*, proposes a way out of this fragmentation. Its central move is to stop asking what value is made of and to ask instead what a mind or a social system is doing when it assigns value to something. Its answer, stated qualitatively, is that valuing is a form of information processing: a mind faced with a large space of possible things to attend to, believe, or act on assigns unequal probabilities to different possibilities, and the possibilities that receive comparatively high probability are, by definition, its values. A system in which nothing is prioritized over anything else is a system that has not yet formed any values at all.

Granular interaction thinking theory (GITT) begins from an analogy with quantum mechanics (for the physics itself, see Griffiths & Schroeter (2018)): minds and social systems are

proposed to be composed of finite informational quanta whose interactions generate the beliefs, values, and cultural structures observed at the level of individuals and societies. Three features of the quantum picture are carried over by analogy. *Granularity* holds that information, like energy, comes in finite units; a mind operates on some bounded set of informational quanta at any given time. *Relationality* holds that meaning and value do not reside in an informational quantum taken alone, but emerge from its interactions with other quanta. *Indeterminacy* holds that the future state of a mind or system is only probabilistically constrained by its past state, since the process by which information is absorbed and retained is itself probabilistic.

GITT distinguishes the mind, an information collection-and-processing system, from the environment, a broader system, social or ecological, that encompasses it, formalizing their relationship loosely as a set-theoretic containment: the mind is a subset of the environment, with membership probabilistic rather than guaranteed. An informational quantum assigned comparatively high retention probability is what the source paper calls a value quantum, and this move is what lets it connect value to entropy: a mind in which no piece of information is more important than any other is, in informational terms, a maximum-entropy mind, one that has not yet formed values, while assigning value to some subset of information is asserted, though never proved, to lower the entropy of the mind's informational state. The source paper extends this account to markets directly: a decentralized market's convergence on a shared valuation is treated as a form of collective entropy reduction, achieved without any centralized authority dictating the outcome, exactly the qualitative picture Section 12.1 eventually formalizes for Bitcoin.

Everything just stated is presented, in the source paper, in prose rather than equations. The claim that entropy falls as values are formed is not derived from the entropy formula it is paired with; the claim that granularity, relationality, and indeterminacy are genuine structural features, rather than suggestive labels borrowed from physics, is asserted rather than shown. This is not a criticism specific to one short paper aimed at motivating a perspective rather than proving it; it is precisely the gap the rest of this book addresses.

1.2 What Is Missing

The next several parts of this book go through the source paper's claims one at a time and ask what would be needed to state each as mathematics. Granularity is, on inspection, the least demanding to formalize: a finite probability distribution over a finite set of outcomes already assumes it, since a probability simplex on n points has no content unless n is finite, though the source paper never exploits this, drawing no bound on n and no consequence of finiteness. Relationality is intuitively compelling and mathematically absent: nothing in the source paper specifies what an interaction between two informational quanta is, how strong it can be, or how the interactions among many quanta combine. Indeterminacy amounts to a single set-theoretic statement, mind as a subset of environment, with no specification of how retention probability is determined. The discussion of belief elasticity draws an extended analogy to elastic and inelastic scattering in quantum mechanics, listing four qualitative factors, motion, position, interaction, and capacity, with no equation given for any of them. And the most important gap, the one the paper's own title foregrounds, is the claim that entropy falls as values form, asserted rather than proved. Finally, the source paper never applies its own framework to a single case in quantitative depth.

Table 1.1 summarizes where each gap is closed. No entry represents a novel mathematical result; each is a standard tool from probability, information theory, or statistical mechanics, applied here, so far as we are aware, for the first time to this particular set of claims.

Original claim	Formal treatment	Chapter
Granularity	Finite probability simplex	7
Relationality	Interaction couplings, cluster expansion	8
Indeterminacy	Stochastic filtering process	9
Elasticity	Gibbs distribution, sensitivity derivative	10
Entropy reduction	KL divergence, Gibbs' inequality	11
Application (Bitcoin)	Energy/temperature mapping, published evidence	12

Table 1.1: Correspondence between the original paper's qualitative claims and their formalization in this book.

1.3 How This Book Is Built: Scaffolding Toward Part III

We want to state the shape of the whole book plainly before beginning it, rather than let the reader discover it piecemeal. Parts I and II build a single toolkit, entropy, the maximum entropy principle, the Gibbs distribution, and the combinatorics of interacting quanta, and Chapter 12 spends that toolkit on a reconstruction of the source paper's own claims, closing Table 1.1's gaps one row at a time and culminating in the Bitcoin application. That reconstruction is real work, but it is not this book's final destination. It is scaffolding, in two distinct senses, for Part III, *A New Perspective on Environmental Values*, which we regard as this book's central and most original contribution.

The first sense is technical and direct: Chapter 15's results, the entropy ceiling on cultural possibility and the combinatorial collapse of achievable cultural configurations, are not analogies to Part I's apparatus but direct applications of it, the same maximum-entropy theorem and the same combinatorics, proved once and reused rather than rederived. A reader who has followed Chapters 2 and 6 carefully has already done most of the work Chapter 15 requires. The second sense is thematic rather than technical: Chapters 13 and 14 build their own, self-contained stochastic calculus and do not draw on the Gibbs-distribution apparatus at all, but the directionality they formalize, that value formation and its collapse are not symmetric, that entropy reduction runs one way, is the same directionality Chapter 11's value formation theorem establishes for minds. We say clearly, and repeat at the relevant points, which of these two relationships holds where; we do not want the word scaffolding to imply a technical dependency where only a thematic echo exists.

A reader who wants the shortest route to this book's original contribution can take it directly: read this chapter for motivation, then move to Part III, which is self-contained given only calculus and elementary probability, the same prerequisites as the rest of the book. A reader who wants to see the full apparatus built, tested against a genuinely difficult case, Bitcoin, and only then extended to new ground, should read the parts in order. Either way, we want the destination visible from the outset: Parts I and II are not proposing a new economic theory, nor is their mathematics itself novel, every substantial tool used, the Gibbs distribution, the maximum entropy principle, the cluster expansion, is standard and well established. What they offer is a mapping, tested against Bitcoin, and a toolkit, spent again in Part III. Part III is different in kind: its models and closed-form results are this book's own construction, motivated by, but not derived from, the source paper or any other single prior work.

Part I

Mathematical Foundations

Chapter 2

Probability, Entropy, and Information

This chapter, and the four that follow it, build the mathematical toolkit used throughout the rest of the book. Nothing here is assumed beyond calculus and the barest familiarity with what a probability is; readers wanting a fuller foundational treatment can consult Feller (1968) for probability, Rudin (1976) for the underlying real analysis, or Cover & Thomas (2006) for information theory specifically, though none is needed to follow what is developed here. Readers who already know this material, entropy, Lagrange multipliers, the Gibbs distribution, partition functions, generating functions, are encouraged to skim and move to Part II, where the toolkit is put to use.

2.1 A Toy Mind

Throughout this chapter we carry along a single running example: a toy mind with just four informational quanta, labeled x_1, x_2, x_3, x_4 . The reader should picture these as four competing pieces of information the mind might retain, four candidate beliefs, four possible interpretations of some event, anything with that structure will do. The point of keeping $n = 4$ fixed and small is that every abstract definition below can be checked by hand on this example before it is stated in general.

2.2 Discrete Probability Distributions

Definition 2.1. A discrete probability distribution over n outcomes $x_1, \dots, x_n$ is a vector $P = (p_1, \dots, p_n)$ with $p_i \geq 0$ for every i and $\sum_{i=1}^n p_i = 1$.

Each p_i is interpreted, in this book, as the probability that quantum x_i is the one retained, or attended to, or acted upon, by the mind at a given moment. In our toy mind, a distribution such as $P = (0.4, 0.3, 0.2, 0.1)$ says that x_1 is the most likely quantum to be retained, and x_4 the least, with the four probabilities summing, as they must, to one.

The set of all such distributions on n outcomes is called the probability simplex, written Δ^{n-1} . Chapter 7 returns to this object directly; for now it is enough to know that every distribution we discuss is a single point in this set.

A special point of the simplex is the uniform distribution $U = (1/n, \dots, 1/n)$, in which every outcome is equally likely. For our toy mind, $U = (0.25, 0.25, 0.25, 0.25)$. This is the distribution of a mind that has not yet distinguished among its four candidate quanta at all, and it will recur constantly below.

2.3 Shannon Entropy

Definition 2.2. The Shannon entropy (Shannon 1948) of a distribution $P = (p_1, \dots, p_n)$, measured in bits, is

$$H(P) = - \sum_{i=1}^n p_i \log_2 p_i,$$

with the convention that $0 \log_2 0 = 0$.

Entropy measures how spread out, or how concentrated, a distribution is. Two extreme cases make this concrete on our toy mind. If $P = (1, 0, 0, 0)$, one quantum is certain and the rest are impossible; every term in the sum is either $1 \cdot \log_2 1 = 0$ or $0 \cdot \log_2 0 = 0$, so $H(P) = 0$. This is the lowest entropy a distribution can have: complete certainty, complete prioritization. At the other extreme, the uniform distribution $U = (0.25, 0.25, 0.25, 0.25)$ gives

$$H(U) = -4 \times (0.25 \times \log_2 0.25) = -4 \times (0.25 \times -2) = 2 \text{ bits}.$$

This turns out to be the largest entropy any distribution on four outcomes can have, a fact we prove in general below.

Proposition 2.1 (Maximum entropy of the uniform distribution). *Among all distributions P on n outcomes, entropy $H(P)$ is maximized uniquely by the uniform distribution U , with $H(U) = \log_2 n$.*

We defer the proof of Proposition 2.1 to Section 2.4 below, since the cleanest proof uses the Kullback-Leibler divergence introduced there. It is worth pausing, however, on what this proposition means for the source paper's claim, restated in Chapter 1, that entropy is highest when no information is prioritized over any other. Proposition 2.1 makes that claim exactly true, for the first time in this book, as a theorem rather than an assertion.

2.4 Kullback-Leibler Divergence and the Proof of Proposition 2.1

Definition 2.3. Given two distributions P and Q on the same n outcomes, with Q having no zero entries where P is positive, the Kullback-Leibler divergence of P from Q is

$$D_{KL}(P\|Q) = \sum_{i=1}^n p_i \log_2 \frac{p_i}{q_i}.$$

$D_{KL}(P\|Q)$ measures how different P is from Q , in a specific information-theoretic sense; it is not a distance in the ordinary sense, since in general $D_{KL}(P\|Q) \neq D_{KL}(Q\|P)$, but it has the property that will do all the work in this section.

Lemma 2.1 (Gibbs' inequality). *For any two distributions P, Q on n outcomes, $D_{KL}(P\|Q) \geq 0$, with equality if and only if $P = Q$.*

Proof. This follows from the concavity of the logarithm via Jensen's inequality. Write

$$-D_{KL}(P\|Q) = \sum_{i=1}^n p_i \log_2 \frac{q_i}{p_i}.$$

Since $\log_2$ is concave, Jensen's inequality gives

$$\sum_{i=1}^n p_i \log_2 \frac{q_i}{p_i} \leq \log_2 \left(\sum_{i=1}^n p_i \cdot \frac{q_i}{p_i} \right) = \log_2 \left(\sum_{i=1}^n q_i \right) = \log_2 1 = 0,$$

where the sum is taken only over i with $p_i > 0$. Hence $-D_{KL}(P\|Q) \leq 0$, giving $D_{KL}(P\|Q) \geq 0$. Equality in Jensen's inequality for a strictly concave function holds if and only if the quantity being averaged, here q_i/p_i , is constant across all i with $p_i > 0$; combined with both P and Q summing to one, this forces $p_i = q_i$ for every i . $\square$

Proposition 2.1 now follows quickly.

Proof of Proposition 2.1. Take $Q = U$ in the definition of D_{KL} . Then

$$\begin{aligned} D_{KL}(P\|U) &= \sum_{i=1}^n p_i \log_2 \frac{p_i}{1/n} = \sum_{i=1}^n p_i \log_2 p_i + \sum_{i=1}^n p_i \log_2 n \\ &= -H(P) + \log_2 n, \end{aligned}$$

using $\sum_i p_i = 1$ for the second sum. Rearranging, $H(P) = \log_2 n - D_{KL}(P\|U)$. By Lemma 2.1, $D_{KL}(P\|U) \geq 0$ with equality only at $P = U$, so $H(P) \leq \log_2 n$, with equality only at $P = U$. $\square$

This identity, $H(P) = \log_2 n - D_{KL}(P\|U)$, is worth keeping in view; it says that entropy and divergence from uniform are simply two ways of describing the same quantity, one measuring how spread out a distribution is directly, the other measuring how far it has moved from the point of maximal spread. Chapter 11 returns to this identity to give the source paper's central claim, that forming values lowers entropy, its full proof.

2.5 Summary

We have defined discrete probability distributions, Shannon entropy, and Kullback-Leibler divergence, and proved, from Gibbs' inequality alone, that the uniform distribution uniquely maximizes entropy. Every step used only algebra and the concavity of the logarithm; nothing here depended on any claim from the source paper, nor on any external authority beyond Gibbs' inequality itself, which we proved rather than cited. The next chapter turns to a different question: given some external constraint on a distribution, short of pinning it down entirely, what distribution best respects that constraint while assuming as little else as possible? The answer, the maximum entropy principle, is what eventually produces the Gibbs distribution central to Chapters 4 and 10.

Chapter 3

Constrained Optimization and the Maximum Entropy Principle

Chapter 2 showed that, among all distributions on n outcomes, the uniform distribution has the largest entropy. That is a statement about an unconstrained maximum: no requirement was placed on P beyond being a valid distribution. This chapter asks a more useful question, the one Jaynes (2003) places at the center of an entire approach to probability and inference: suppose we know something more about the mind's informational state, not merely that it is a distribution, but that it satisfies some additional constraint, some average cost, or level of engagement, or informational load, is held fixed. Which distribution, among all those satisfying that constraint, has the largest entropy? The answer to this question, worked out in full generality below, is what produces the Gibbs distribution central to Chapters 4 and 10.

3.1 Lagrange Multipliers, from Scratch

We first recall the method of Lagrange multipliers in its simplest setting, a standard technique of multivariable calculus (Apostol 1967), since everything that follows is a repeated application of the same idea.

Worked Example

Suppose we want to maximize $f(x, y) = xy$ subject to the constraint $x + y = 10$. One way to solve this is by substitution: set $y = 10 - x$, so $f = x(10 - x) = 10x - x^2$, and maximize over x by setting the derivative to zero: $10 - 2x = 0$, giving $x = 5$, $y = 5$, and a maximum value of 25.

The method of Lagrange multipliers reaches the same answer without solving for y explicitly, and it is this version that generalizes cleanly to the constraints we need below. Form the Lagrangian

$$\mathcal{L}(x, y, \lambda) = xy - \lambda(x + y - 10),$$

where λ is a new variable, the multiplier, introduced to enforce the constraint. Setting all three partial derivatives to zero,

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x} &= y - \lambda = 0, \\ \frac{\partial \mathcal{L}}{\partial y} &= x - \lambda = 0, \\ \frac{\partial \mathcal{L}}{\partial \lambda} &= -(x + y - 10) = 0,\end{aligned}$$

gives $y = \lambda$, $x = \lambda$, and $x + y = 10$, so $x = y = 5$ as before. The multiplier λ has a value, 5 in this case, but its main role here was procedural: it converted a constrained problem into an unconstrained one by adding one new variable for each constraint.

The general recipe: to maximize f subject to $g = c$, form $\mathcal{L} = f - \lambda(g - c)$, and solve $\nabla \mathcal{L} = 0$, treating λ as an additional unknown to be solved for alongside the original variables.

3.2 Maximizing Entropy Subject to Normalization Alone

We now apply this method to entropy itself, first recovering Proposition 2.1 by this route, as a check that the machinery works before adding further constraints.

We want to maximize $H(P) = -\sum_{i=1}^n p_i \log_2 p_i$ over $p_1, \dots, p_n$, subject only to $\sum_i p_i = 1$. Form the Lagrangian

$$\begin{aligned}\mathcal{L}(p_1, \dots, p_n, \lambda) &= -\sum_{i=1}^n p_i \log_2 p_i \\ &\quad - \lambda \left(\sum_{i=1}^n p_i - 1 \right).\end{aligned}$$

Taking the partial derivative with respect to a single p_i , using

$$\frac{d}{dp}(p \log_2 p) = \log_2 p + \frac{1}{\ln 2},$$

and setting it to zero,

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial p_i} &= -\log_2 p_i - \frac{1}{\ln 2} - \lambda = 0 \\ \implies \log_2 p_i &= -\frac{1}{\ln 2} - \lambda,\end{aligned}$$

which does not depend on i at all. Hence every p_i takes the same value, and since the p_i must sum to one, $p_i = 1/n$ for all i : the uniform distribution. This matches Proposition 2.1, now reached by a method that generalizes.

3.3 Adding a Second Constraint

Now suppose the mind's distribution is not free to be uniform, because retaining different quanta carries different costs, an energy E_i associated with quantum x_i , perhaps reflecting how much cognitive effort, social risk, or contradiction with prior belief that quantum entails. Suppose further that the mind's expected cost is constrained to some fixed value $\bar{E}$:

$$\sum_{i=1}^n p_i E_i = \bar{E}, \quad \sum_{i=1}^n p_i = 1.$$

Among all distributions satisfying both constraints, which has the largest entropy? This is precisely the question of the least biased distribution consistent with a known average cost: it assumes nothing about P beyond what the two constraints require.

Form the Lagrangian with two multipliers, λ for normalization and β for the energy constraint:

$$\mathcal{L} = -\sum_i p_i \log_2 p_i - \lambda \left(\sum_i p_i - 1 \right) - \beta \left(\sum_i p_i E_i - \bar{E} \right).$$

Differentiating with respect to p_i and setting the result to zero,

$$-\log_2 p_i - \frac{1}{\ln 2} - \lambda - \beta E_i = 0 \implies \log_2 p_i = -\frac{1}{\ln 2} - \lambda - \beta E_i,$$

so

$$p_i = 2^{-1/\ln 2 - \lambda} \cdot 2^{-\beta E_i} = c \cdot 2^{-\beta E_i}$$

for a constant c that does not depend on i . Converting to natural-log form by absorbing constants, this is

$$p_i \propto e^{-\beta' E_i}$$

for some constant β' related to β by a change of base. We suppress the distinction between β and β' from here on and write simply

$$p_i = \frac{e^{-\beta E_i}}{\sum_{j=1}^n e^{-\beta E_j}},$$

where the denominator is exactly the normalizing constant that makes the p_i sum to one.

Proposition 3.1 (Maximum entropy under an energy constraint). *Among all distributions P on n outcomes satisfying $\sum_i p_i E_i = \bar{E}$, the entropy $H(P)$ is maximized by*

$$p_i = \frac{e^{-\beta E_i}}{Z}, \quad Z = \sum_{j=1}^n e^{-\beta E_j},$$

for the value of β that makes the constraint hold.

This is, up to notation, the Gibbs distribution, and Chapter 4 takes up its properties, its interpretation, and its normalizing constant Z in full. For now we record only the shape of the result: entropy maximization under an energy constraint does not produce a uniform distribution, but an exponentially weighted one, in which low-energy quanta receive disproportionately high probability whenever $\beta > 0$.

Worked Example

Return to the toy mind with four quanta, and suppose their energies are $E_1 = 0$, $E_2 = 1$, $E_3 = 1$, $E_4 = 3$, in some convenient units, and suppose $\beta = 1$. Then

$$Z = e^0 + e^{-1} + e^{-1} + e^{-3} \approx 1 + 0.368 + 0.368 + 0.050 = 1.786,$$

giving $p_1 \approx 0.560$, $p_2 = p_3 \approx 0.206$, $p_4 \approx 0.028$. Compare this to the uniform distribution $(0.25, 0.25, 0.25, 0.25)$: the low-energy quantum x_1 has more than doubled its probability, at the expense of the high-energy quantum x_4 , which has fallen to roughly a tenth of its uniform share. This is the qualitative pattern behind everything Chapter 10 says about belief and value: low-energy quanta, whatever energy is eventually taken to represent, are disproportionately retained.

3.4 Summary

Starting from the ordinary calculus of Lagrange multipliers, we derived the exponential, Gibbs-shaped distribution as the unique maximum-entropy solution under a single linear constraint on expected energy, and we recovered the uniform distribution as the special case of no constraint beyond normalization. Nothing in this derivation referred to physics; it is a fact about constrained optimization applied to the entropy functional, true regardless of what E_i is eventually interpreted to mean. Chapter 4 takes up that interpretation directly, connecting β to a notion of openness or temperature, and connecting Z to the partition function whose properties drive much of the remainder of this book.

Chapter 4

The Gibbs Distribution

Chapter 3 derived, from the maximum entropy principle alone, a distribution of the shape $p_i \propto e^{-\beta E_i}$. This chapter takes that result and develops it properly: naming it, studying how it behaves as β varies, and building the interpretive bridge between β and what the source paper calls openness, the disposition of a mind to let new information in rather than defending its existing informational state.

4.1 Definition and Basic Properties

Definition 4.1 (Gibbs distribution). Given energies $E_1, \dots, E_n$ and a parameter $\beta > 0$, the Gibbs distribution is

$$p_i = \frac{e^{-\beta E_i}}{Z(\beta)}, \quad Z(\beta) = \sum_{j=1}^n e^{-\beta E_j}.$$

$Z(\beta)$, the normalizing constant, is given its own name, the partition function, and its properties are developed fully in Chapter 5; here we treat it only as the constant that makes the p_i sum to one.

It is standard, following the statistical-mechanical literature this construction is borrowed from (see Kittel & Kroemer (1980) for a standard introductory treatment, or Pathria & Beale (2011) and Reif (1965) for a fuller one), to write $\beta = 1/T$, where T is called the temperature. This book uses T as a formal stand-in for what the source paper calls openness: the willingness of a mind to spread its probability across many competing quanta rather than concentrating it on a few. The next two propositions justify that reading.

Proposition 4.1 (High-temperature limit). As $T \rightarrow \infty$ (equivalently $\beta \rightarrow 0$), the Gibbs distribution converges to the uniform distribution U .

Proof. As $\beta \rightarrow 0$, each $e^{-\beta E_i} \rightarrow e^0 = 1$, so $Z(\beta) \rightarrow n$ and $p_i \rightarrow 1/n$ for every i . $\square$

Proposition 4.2 (Low-temperature limit). As $T \rightarrow 0^+$ (equivalently $\beta \rightarrow \infty$), the Gibbs distribution converges to a point mass on the quantum, or quanta, of lowest energy.

Proof. Let $E_{\min} = \min_i E_i$ and let $S = \{i : E_i = E_{\min}\}$ be the set of quanta achieving it. Write $p_i = e^{-\beta(E_i - E_{\min})} / \sum_j e^{-\beta(E_j - E_{\min})}$, which is the same distribution as before, since both numerator and denominator have been multiplied by the constant $e^{\beta E_{\min}}$. For $i \in S$, $E_i - E_{\min} = 0$, so the numerator equals one for every β ; for $i \notin S$, $E_i - E_{\min} > 0$, so $e^{-\beta(E_i - E_{\min})} \rightarrow 0$ as $\beta \rightarrow \infty$. Hence the denominator converges to $|S|$, and $p_i \rightarrow 1/|S|$ for $i \in S$, $p_i \rightarrow 0$ for $i \notin S$.

otherwise: a uniform distribution over only the lowest-energy quanta, degenerating to a single point mass when the minimum is achieved uniquely. $\square$

Together, Propositions 4.1 and 4.2 justify reading T as an openness parameter exactly in the sense Chapter 1 attributes to the source paper, though only as an interpretation, not as something the source paper itself derives. At high T , the mind is maximally open: every quantum is equally likely to be retained, regardless of its cost, and by Proposition 2.1 this is also the state of maximum entropy. At low T , the mind is maximally closed: only the cheapest, least effortful, or most familiar quantum survives, and the distribution collapses toward a single fixed belief.

Worked Example

Return once more to the toy mind, with energies $E_1 = 0, E_2 = 1, E_3 = 1, E_4 = 3$. Table 4.1 shows the Gibbs distribution at three temperatures.

T	p_1	p_2	p_3	p_4
$T = 10$ ($\beta = 0.1$)	0.282	0.255	0.255	0.209
$T = 1$ ($\beta = 1$)	0.560	0.206	0.206	0.028
$T = 0.1$ ($\beta = 10$)	1.000	0.000	0.000	0.000

As T falls from 10 to 0.1, the distribution moves visibly from close to uniform toward complete concentration on the single lowest-energy quantum x_1 , exactly as Propositions 4.1 and 4.2 predict.

Figure 4.1 shows the same three distributions as a grouped bar chart, making the collapse toward x_1 visually immediate in a way the table's numbers require a moment to parse.




Figure 4.1: The Gibbs distribution over the toy mind's four quanta ($E = (0, 1, 1, 3)$) at three temperatures. As T falls, probability mass concentrates entirely on the lowest-energy quantum x_1 , illustrating Propositions 4.1 and 4.2 directly.

4.2 Monotonicity in Temperature

One further property is worth establishing here, since Chapter 10 relies on it directly: entropy itself is monotonic in T .

Proposition 4.3 (Entropy increases with temperature). *Let $H(\beta)$ denote the entropy of the Gibbs distribution at inverse temperature β , for fixed energies $E_1, \dots, E_n$ not all equal. Then $H(\beta)$ is a strictly decreasing function of β on $\beta > 0$, and hence a strictly increasing function of T .*

We defer the proof to Chapter 5, where it follows quickly and transparently from a property of the partition function $Z(\beta)$ established there, namely that the second derivative of $\log Z(\beta)$ with respect to β equals the variance of the energy under the Gibbs distribution, a non-negative quantity. Stating Proposition 4.3 here, ahead of its proof, is deliberate: it is the fact that ultimately justifies treating T as a single, well-behaved dial that moves a mind smoothly between full prioritization and full indifference, with entropy tracking that movement monotonically rather than erratically.

4.3 Summary

We have named the exponential distribution derived in Chapter 3, shown that it interpolates smoothly between the uniform distribution and a point mass as its temperature parameter varies, and stated, ahead of proof, that entropy varies monotonically along the same path. Chapter 10 uses this apparatus to give the source paper's account of belief elasticity a precise, differentiable meaning; Chapter 5 develops the partition function's own properties in full, including the deferred proof of Proposition 4.3.

Chapter 5

Partition Functions and Free Energy

The partition function $Z(\beta)$ has appeared twice already, first as the normalizing constant of the Gibbs distribution in Chapter 3, then as a function whose behavior was promised, but not yet shown, to control the monotonicity of entropy in Chapter 4. This chapter makes good on that promise. It turns out that $Z(\beta)$ is not merely a normalizing constant; its logarithm encodes essentially everything about the Gibbs distribution's statistical behavior, entropy, expected energy, and their dependence on temperature, all follow from $\log Z(\beta)$ by differentiation. This is the single fact that makes the partition function worth naming rather than leaving as an unlabeled denominator.

5.1 Free Energy

Definition 5.1 (Free energy). Given a partition function $Z(\beta)$, the free energy is

$$F(\beta) = -\frac{1}{\beta} \log Z(\beta) = -T \log Z(\beta).$$

The name is inherited from thermodynamics, and nothing in this book depends on the physical connotations of the word; it is used here purely as an accounting device. Its usefulness is that expected energy and entropy can both be recovered from it by differentiation, a fact we now establish.

Proposition 5.1 (Expected energy from the partition function). Let $\bar{E}(\beta) = \sum_i p_i E_i$ denote the expected energy under the Gibbs distribution at inverse temperature β . Then

$$\bar{E}(\beta) = -\frac{d}{d\beta} \log Z(\beta).$$

Proof. Differentiate $\log Z(\beta) = \log \sum_j e^{-\beta E_j}$ directly:

$$\begin{aligned} \frac{d}{d\beta} \log Z(\beta) &= \frac{1}{Z(\beta)} \sum_j (-E_j) e^{-\beta E_j} = -\sum_j E_j \cdot \frac{e^{-\beta E_j}}{Z(\beta)} \\ &= -\sum_j E_j p_j = -\bar{E}(\beta). \end{aligned}$$

□

Proposition 5.2 (Entropy from the partition function). The entropy of the Gibbs distribution

(in natural-log units, $H(P) = -\sum_i p_i \ln p_i$) satisfies

$$H(\beta) = \beta \bar{E}(\beta) + \log Z(\beta).$$

Proof. Substitute $p_i = e^{-\beta E_i} / Z(\beta)$ directly into the entropy formula:

$$\begin{aligned} H(\beta) &= -\sum_i p_i \ln p_i = -\sum_i p_i (-\beta E_i - \log Z(\beta)) \\ &= \beta \sum_i p_i E_i + \log Z(\beta) \sum_i p_i = \beta \bar{E}(\beta) + \log Z(\beta), \end{aligned}$$

using $\sum_i p_i = 1$ in the last step. □

5.2 The Deferred Proof: Entropy Increases with Temperature

We can now complete the proof of Proposition 4.3, stated without proof in Chapter 4.

Proof of Proposition 4.3. Differentiate the identity of Proposition 5.2 with respect to β :

$$\frac{dH}{d\beta} = \bar{E}(\beta) + \beta \frac{d\bar{E}}{d\beta} + \frac{d}{d\beta} \log Z(\beta) = \bar{E}(\beta) + \beta \frac{d\bar{E}}{d\beta} - \bar{E}(\beta) = \beta \frac{d\bar{E}}{d\beta},$$

using Proposition 5.1 to replace $\frac{d}{d\beta} \log Z(\beta)$ with $-\bar{E}(\beta)$, which cancels the first term. It remains to determine the sign of $d\bar{E}/d\beta$. Differentiating $\bar{E}(\beta) = -\frac{d}{d\beta} \log Z(\beta)$ once more,

$$\frac{d\bar{E}}{d\beta} = -\frac{d^2}{d\beta^2} \log Z(\beta).$$

A direct computation (the variance identity for exponential families, included here for completeness) gives $\frac{d^2}{d\beta^2} \log Z(\beta) = \text{Var}(E)$, the variance of the energy under the Gibbs distribution, which is non-negative, and strictly positive whenever the E_i are not all equal. Hence $d\bar{E}/d\beta = -\text{Var}(E) \leq 0$, so $\frac{dH}{d\beta} = \beta \cdot (-\text{Var}(E)) \leq 0$ for $\beta > 0$, strictly negative whenever the energies are not all equal. Since $T = 1/\beta$ is a strictly decreasing function of β , H is strictly increasing in T . □

For completeness, the variance identity used above follows from one further differentiation of $\log Z(\beta)$:

$$\begin{aligned} \frac{d^2}{d\beta^2} \log Z(\beta) &= \frac{d}{d\beta} (-\bar{E}(\beta)) \\ &= \sum_i E_i^2 p_i - \left(\sum_i E_i p_i \right)^2 = \text{Var}(E), \end{aligned}$$

where the middle equality is a direct, if slightly tedious, differentiation of

$$-\bar{E}(\beta) = -\sum_i E_i \frac{e^{-\beta E_i}}{Z(\beta)}$$

using the quotient rule, which we do not reproduce term by term here since it contributes no further insight beyond the result itself.

5.3 Summary

The free energy $F(\beta) = -T \log Z(\beta)$ is not merely a bookkeeping device; its logarithm's derivatives recover expected energy and entropy directly, and the second derivative of $\log Z$ is exactly the variance of energy under the Gibbs distribution, which is what forces entropy to increase monotonically with temperature. This closes the deferred proof from Chapter 4 and completes Part I's toolkit for single, isolated quanta. What remains missing, and what Chapter 6 takes up next, is the case where quanta are not independent, where the energy of retaining one quantum depends on which others are also retained. This is where relationality finally receives a mathematical home.

Generating Functions and Diagrammatic Expansions

Every result so far has treated the mind's informational quanta as independent: the energy E_i of retaining quantum i did not depend on which other quanta were also retained. This chapter removes that restriction, which is necessary before relationality, the claim that value emerges from interaction rather than from isolated quanta, can be given any mathematical content at all. The tool required, a classical combinatorial result usually called the exponential formula, or in its physics dress the linked cluster theorem, is built up here from ordinary generating functions, following the standard treatment in Wilf (2006), assuming no prior exposure to either.

6.1 Ordinary Generating Functions, Briefly

Definition 6.1. Given a sequence of numbers $a_0, a_1, a_2, \dots$, its ordinary generating function is the formal power series $A(x) = \sum_{k=0}^{\infty} a_k x^k$.

The word formal is important: we do not ask whether this series converges for any particular x , only that it correctly encodes the sequence a_k as coefficients. The value of this bookkeeping device is that operations on sequences correspond to simple operations on the generating functions. For example, if $A(x)$ generates a_k and $B(x)$ generates b_k , then $A(x)B(x)$ generates the convolution $c_k = \sum_{j=0}^k a_j b_{k-j}$. We will not need this fact in its full generality, only its exponential cousin, developed next.

6.2 Exponential Generating Functions

Definition 6.2. Given a sequence $a_0, a_1, a_2, \dots$, where a_k counts some kind of labeled structure on a set of k elements, its exponential generating function (EGF) is $A(x) = \sum_{k=0}^{\infty} a_k \frac{x^k}{k!}$.

The division by $k!$ is what makes EGFs the right tool for counting labeled structures, structures built on a specific, distinguishable set of k labeled elements (in our setting, the elements will be informational quanta), where the labels matter but their order of construction does not.

Worked Example

Let a_k count the number of ways to partition a set of k labeled elements into any number of nonempty, unlabeled parts (this is the k -th Bell number). It is a classical fact, which we now derive, that the EGF of the Bell numbers is $\exp(e^x - 1)$.

The key observation is this: a partition of a k -element set into parts is the same thing as choosing, for the single part containing a particular fixed element, some subset of the remaining elements to join it, and then independently partitioning the rest. This recursive, per-element independence is exactly what the exponential of an EGF captures in general, stated as the exponential formula below.

6.3 The Exponential Formula

Proposition 6.1 (Exponential formula, combinatorial form). *Let c_k count labeled connected structures on k elements (for instance, connected graphs on k labeled vertices), with EGF $C(x) = \sum_{k \geq 1} c_k \frac{x^k}{k!}$. Let a_k count arbitrary labeled structures on k elements formed as disjoint unions of connected structures (a labeled structure being partitioned into connected components, each itself a connected structure on its own labels), with EGF $A(x)$. Then*

$$A(x) = \exp(C(x)).$$

We omit the fully general combinatorial proof of Proposition 6.1, which is standard (see, for instance, any treatment of species or labeled enumeration in enumerative combinatorics), and instead verify it directly in the setting we actually need, small systems of interacting quanta, in Section 6.4 below, where the claim reduces to an identity that can be checked by direct expansion.

The consequence that matters for this book is the following restatement: taking logarithms of both sides, $\log A(x) = C(x)$. In words, the logarithm of the generating function for all structures equals the generating function for connected structures alone. Disconnected structures, those splitting into two or more separate components, contribute to $A(x)$, but their contribution is fully accounted for by products of connected pieces, and taking the logarithm strips those products back down to sums over connected pieces only.

6.4 Application: The Partition Function of Interacting Quanta

We now apply this to the mind's informational quanta directly. Suppose, in addition to each quantum's own energy E_i , there is a pairwise interaction energy J_{ij} between quanta i and j , representing the added or reduced cost of retaining both simultaneously, exactly the relational structure the source paper gestures toward without specifying. The total energy of retaining a subset $S \subseteq \{1, \dots, n\}$ of quanta together is taken to be

$$E(S) = \sum_{i \in S} E_i + \sum_{\{i,j\} \subseteq S} J_{ij},$$

and the partition function over all subsets is

$$Z = \sum_{S \subseteq \{1, \dots, n\}} e^{-\beta E(S)}.$$

For $n = 2$, this is simple enough to verify Proposition 6.1 directly. Write $z_i = e^{-\beta E_i}$ and $w_{ij} = e^{-\beta J_{ij}}$. Summing over the four subsets of $\{1, 2\}$ (empty, $\{1\}$, $\{2\}$, $\{1, 2\}$):

$$Z = 1 + z_1 + z_2 + z_1 z_2 w_{12}.$$

If $J_{12} = 0$ (no interaction), $w_{12} = 1$ and $Z = 1 + z_1 + z_2 + z_1 z_2 = (1 + z_1)(1 + z_2)$ factors exactly, giving $\log Z = \log(1 + z_1) + \log(1 + z_2)$, a sum of independent, single-quantum contributions, with no trace of any interaction between them, since there is none.

With $J_{12} \neq 0$, the extra factor w_{12} prevents this clean factorization, and

$$\log Z = \log \left((1 + z_1)(1 + z_2) + z_1 z_2 (w_{12} - 1) \right),$$

which does not split into a sum of single-quantum terms. Write $x_0 = (1 + z_1)(1 + z_2)$ for the value $\log Z$ would take with no interaction, and $\delta = z_1 z_2 (w_{12} - 1)$ for the extra term the interaction contributes, so that $\log Z = \log(x_0 + \delta)$. For δ small relative to x_0 , the first-order approximation $\log(x_0 + \delta) \approx \log x_0 + \delta/x_0$ applies, and using $w_{12} - 1 = e^{-\beta J_{12}} - 1 \approx -\beta J_{12}$ for small βJ_{12} gives, to leading order,

$$\log Z \approx \log(1 + z_1) + \log(1 + z_2) + \frac{z_1 z_2}{(1 + z_1)(1 + z_2)} \cdot (-\beta J_{12}),$$

The first two terms are exactly the two disconnected, single-quantum contributions found above; the third, new term is the connected, two-quantum contribution, present only because $J_{12} \neq 0$. This is Proposition 6.1 made concrete: $\log Z$ decomposes into a sum over connected pieces, here a sum over single quanta plus one two-quantum connected term, with no three-or-more-body disconnected cross-terms appearing anywhere, exactly as the exponential formula predicts.

6.5 Diagrams

It is standard, and useful, to represent this decomposition pictorially. Draw one vertex for each informational quantum; draw an edge between i and j whenever $J_{ij} \neq 0$. The contribution of a connected diagram, a connected subgraph of these vertices and edges, to $\log Z$ is exactly the connected term the exponential formula assigns to that subset of quanta. In the two-quantum example above, there are exactly two diagrams: the single vertex (appearing twice, once for each quantum) and the single edge connecting both. No larger, disconnected diagram, two separate edges, for instance, ever appears in $\log Z$; those contributions were already accounted for, and exactly cancelled, in passing from Z to $\log Z$. This is the precise sense in which, as promised in Chapter 1, relationality is formalized here: the free energy and entropy of an interacting system of quanta are built entirely out of the connected clusters of quanta among them, and larger systems of independent, non-interacting quanta contribute nothing beyond the sum of their individual, disconnected parts.




Figure 6.1: Diagrams for the exponential formula. Top: the two-quantum example, where the single-vertex diagrams and the connecting edge are the only two diagram types, matching $\log(1 + z_1) + \log(1 + z_2)$ plus the interaction term derived above. Bottom: a four-quantum system with two separate interacting pairs; the connected diagrams $\{x_1, x_2\}$ and $\{x_3, x_4\}$ each contribute separately, but their disconnected union, spanning all four vertices as a single diagram, does not appear in $\log Z$ at all.

6.6 Summary

Starting from ordinary and exponential generating functions, we stated the exponential formula relating disconnected to connected labeled structures, verified it directly for two interacting quanta, and showed that it reproduces exactly the intuitive claim that free energy from independent quanta is additive, while free energy from interacting quanta requires an additional, connected correction term. Chapter 8 takes this apparatus and applies it to the mind's full set of n quanta, including three-body and higher interactions J_{ijk} , to give relationality its complete formal treatment.

Part II

The Formal Model

Granularity as a Finite Probability Simplex

Part I built the mathematical toolkit; this chapter begins Part II's task of applying it to the mind itself, one qualitative claim of the source paper at a time. We start with granularity, both because it is logically prior to the rest, relationality and indeterminacy presuppose a finite set of quanta to relate and filter, and because, as Chapter 1 noted, it is the easiest of the source paper's three structural claims to formalize and the one the source paper itself does the least with. This chapter closes that gap.

7.1 The Mind as a Point on the Simplex

Setup

Throughout this and the following four chapters, the mind's informational state at a given moment is represented by a distribution $P = (p_1, \dots, p_n)$ on the probability simplex Δ^{n-1} , where n is finite, $x_1, \dots, x_n$ are the mind's informational quanta at that moment, and p_i is the probability that quantum x_i is the one retained, attended to, or acted upon.

Granularity, in the source paper's own words, is the claim that a cognitive or socio-economic system can process only limited grains of information at any given time. Representing the mind's state as a point on Δ^{n-1} already builds this in: n is a fixed, finite number, not a continuum, and the state is fully specified by finitely many numbers $p_1, \dots, p_n$ summing to one. This is a modest formal step, but it is also, quite literally, all that the word finite requires; the substance of granularity, as a claim with actual consequences, comes from what finiteness forces, which is the subject of the rest of this chapter.

7.2 A Capacity Bound

The source paper calls the high-probability quanta values, but never asks how many quanta can simultaneously qualify as values under any fixed standard of what counts as high. This is a natural question, and it has an exact, elementary answer.

Proposition 7.1 (Capacity bound). *Fix a threshold $\theta \in (0, 1]$, and call a quantum x_i a value quantum if $p_i \geq \theta$. Then, for any distribution P on n outcomes, the number of value quanta is at*

most $\lfloor 1/\theta \rfloor$.

Proof. Let k be the number of quanta with $p_i \geq \theta$. Since all $p_i \geq 0$ and the k value quanta each contribute at least θ to the sum $\sum_i p_i = 1$,

$$1 = \sum_{i=1}^n p_i \geq \sum_{i: p_i \geq \theta} p_i \geq k\theta,$$

so $k \leq 1/\theta$. Since k is an integer, $k \leq \lfloor 1/\theta \rfloor$. $\square$

The proof uses nothing beyond the requirement that probabilities are non-negative and sum to one, which is why Chapter 1 described this as the least demanding of the source paper's claims to formalize. It is, nonetheless, a genuine mathematical necessity rather than a contingent empirical fact about minds: however many quanta a mind happens to contain, and however it happens to be structured, no threshold-based notion of value can be satisfied by more than $\lfloor 1/\theta \rfloor$ quanta simultaneously, simply because probability mass is conserved.

Worked Example

If a value is defined as any quantum receiving at least 20% of the mind's attention, $\theta = 0.2$, then at most $\lfloor 1/0.2 \rfloor = 5$ quanta can be values at once, regardless of how large n is. Demanding a stricter standard, say $\theta = 0.5$, meaning a quantum must claim a majority of the mind's probability to count as a value, immediately implies that at most one quantum can qualify at any given moment, since two quanta at 0.5 each would exhaust the entire distribution, leaving no room for a strict inequality to be violated by a third.

7.3 Interpretation: Scarcity as a Mathematical Necessity, Not a Psychological Fact

This gives quantitative grounding to a claim the source paper makes only in passing, that minds must evaluate, distinguish, compare, and combine information rather than treat all of it as equally salient. Proposition 7.1 shows that this is not merely a claim about limited cognitive capacity, plausible as that is on independent psychological grounds, but a mathematical consequence of representing value as bounded probability mass on a distribution that must sum to one. A mind, or a market, could in principle have unlimited attention, unlimited memory, unlimited processing capacity, and it would still be true that no more than $\lfloor 1/\theta \rfloor$ things could simultaneously clear a θ -threshold of prioritization. Scarcity of values, in this specific sense, is not a fact about minds; it is a fact about probability distributions, and it would hold even for an idealized, unboundedly capable mind.

7.4 Granularity Under the Gibbs Distribution

Chapter 4 showed that as temperature T falls, the Gibbs distribution concentrates on the lowest-energy quanta. Proposition 7.1 lets us say something sharper: it places an upper bound on how many quanta can simultaneously exceed any fixed threshold under the Gibbs distribution as well, since the Gibbs distribution is, after all, only a particular choice of P on the same simplex Δ^{n-1} .

Corollary 7.1. *Under the Gibbs distribution at any temperature $T > 0$, at most $\lfloor 1/\theta \rfloor$ quanta can satisfy $p_i \geq \theta$.*

This corollary requires no separate proof; it is Proposition 7.1 applied to a specific P , and it is worth stating explicitly only because it connects the two chapters directly: whatever a change in temperature does to the shape of the mind's distribution, spreading it out or concentrating it, it can never violate the capacity bound, since that bound holds for every point on the simplex, Gibbs-distributed or not.

7.5 Summary

Granularity, formalized as membership on a finite probability simplex, forces a genuine and previously unstated mathematical consequence: a hard cap on how many informational quanta can simultaneously count as values under any fixed importance threshold, following from nothing more than the conservation of probability mass. This is a small proposition, proved in a few lines, and we do not claim it as a deep result; it is included because it is exactly the kind of consequence the source paper's own language gestures toward without deriving, and because it sets the pattern the rest of Part II follows, take a qualitative claim, find the smallest formal statement that captures it, and prove that statement rather than asserting it. Chapter 8 turns next to relationality, where the corresponding formal statement is considerably less immediate.

Chapter 8

Relationality via Interaction Diagrams

Chapter 6 developed the cluster expansion in the simplest nontrivial case, two interacting quanta. This chapter extends that machinery to the mind's full set of n quanta, including three-body interactions, and states the general result that Chapter 1 promised: relationality, formalized, means that the mind's free energy and entropy decompose entirely into contributions from connected clusters of interacting quanta.

8.1 The General Interaction Energy

Setup

For a subset $S \subseteq \{1, \dots, n\}$ of jointly retained quanta, the total energy is

$$E(S) = \sum_{i \in S} E_i + \sum_{\{i,j\} \subseteq S} J_{ij} + \sum_{\{i,j,k\} \subseteq S} J_{ijk} + \dots,$$

where E_i is the isolated energy of quantum i , J_{ij} the pairwise interaction between i and j , J_{ijk} the triple-wise interaction among i , j , and k , and so on for higher-order terms. In practice we truncate this expansion at some fixed order (pairwise, or pairwise plus triple-wise), since the source paper gives no reason to expect interactions beyond low order to matter, and since higher-order terms multiply the number of free parameters rapidly.

This is the direct formalization of relationality: the cost, and hence the retention probability, of a quantum is no longer determined by E_i alone, but by which other quanta accompany it. Two quanta that are individually costly, high E_i , but jointly cheap, strongly negative J_{ij} , can be more likely to be retained together than either is likely to be retained alone; this is exactly the kind of interaction effect the source paper's discussion of use value, where the value of a mushroom depends on its relation to scarcity, demand, and substitutes rather than any property in isolation, gestures toward without stating formally.

8.2 The Full Partition Function and Its Diagrammatic Decomposition

The partition function over all subsets is $Z = \sum_{S \subseteq \{1, \dots, n\}} e^{-\beta E(S)}$, exactly as in Chapter 6, now with $E(S)$ given by the general expansion above.

Proposition 8.1 (Diagrammatic decomposition of free energy). *$\log Z$ decomposes as a sum over connected diagrams on subsets of $\{1, \dots, n\}$, where a diagram is connected if its vertices (quanta) cannot be split into two nonempty groups with no interaction term, pairwise or higher, linking a vertex in one group to a vertex in the other. Diagrams corresponding to disconnected subsets of quanta, those splitting into two or more mutually non-interacting pieces, contribute nothing to $\log Z$ beyond the sum of their pieces' own connected contributions.*

This is the exponential formula, Proposition 6.1, applied at full generality: connected structures here are connected diagrams of interacting quanta, and the exponential formula's content is precisely that $\log Z$, the free energy up to the constant factor $-T$, is built entirely out of these connected pieces. We verified the two-quantum case of this by direct computation in Chapter 6; the general proof proceeds by the same combinatorial argument (the standard cluster-expansion argument of statistical mechanics, following from Proposition 6.1 applied to the labeled structures given by connected subsets under the interaction graph), and we do not reproduce every step of the general case, since it is a direct, if notationally heavier, repetition of the $n = 2$ computation already carried out in full.

Worked Example

Consider three quanta with pairwise couplings J_{12}, J_{13}, J_{23} , all nonzero, and no three-body term ($J_{123} = 0$). The diagrams contributing to $\log Z$ are: three single-vertex diagrams (one per quantum, from E_1, E_2, E_3 alone), and three two-vertex, single-edge diagrams (one per pairwise interaction, from J_{12}, J_{13}, J_{23} respectively). No three-vertex connected diagram contributes, since $J_{123} = 0$ means there is no term linking all three simultaneously in a way that could not already be captured by the three pairwise edges; the three edges alone do not form a single irreducible three-body term, only a triangle built from independently accounted pairwise pieces. If instead $J_{123} \neq 0$, a genuine three-body connected diagram appears, contributing a term to $\log Z$ that cannot be decomposed into any combination of the pairwise and single-vertex terms already present; this is the mathematical content of a three-way interaction that is not reducible to any combination of two-way interactions, exactly analogous to a mushroom whose value depends jointly on scarcity, cultural context, and available substitutes in a way no pairwise comparison among the three alone would capture.

8.3 What Relationality Buys, Formally

Proposition 8.1 makes precise a distinction the source paper's own language leaves implicit: between quanta whose value is additive, contributing independently to the mind's overall free energy and entropy regardless of what else is retained, and quanta whose value is genuinely relational, contributing a term to $\log Z$ that exists only because of a specific joint configuration and vanishes if that configuration is broken up. A mind, or a market, in which every $J_{ij} = J_{ijk} = \dots = 0$ is one in which relationality, in the technical sense developed here, plays no role at all, and the mind's entropy and value structure reduce entirely to the single-quantum case of Chapters 3 through 5. Everything interesting that relationality adds, in this formalization, is carried by the connected, multi-quantum diagrams, and nothing else.

We flag, with the same honesty as Chapter 1 promised, that truncating the interaction expansion at pairwise or triple-wise order is a modeling choice, not something the source paper specifies or that any general principle forces. Higher-order interactions can always be added at the cost of more free parameters; we do not pursue them further here, since nothing in this book's applications, developed in Chapter 12, requires them.

8.4 Summary

Relationality, as the source paper states it only qualitatively, is now a specific mathematical object: an interaction energy expansion in E_i, J_{ij}, J_{ijk} , together with a diagrammatic decomposition of the resulting free energy into connected clusters, following directly from the exponential formula developed in Chapter 6. Chapter 9 turns next to the source paper's third structural claim, indeterminacy, treating the relationship between mind and environment as a stochastic filtering process built on top of the interaction structure developed here.

Chapter 9

Indeterminacy as Stochastic Filtering

The source paper's third structural claim, indeterminacy, is stated in Chapter 1 as a single set-theoretic relation: the mind is a subset of the environment, and membership is probabilistic rather than guaranteed. This chapter turns that relation into an explicit process, using the same energy structure introduced in Chapters 4 and 8, so that indeterminacy is not a separate mechanism bolted onto the rest of the model, but a direct consequence of it.

9.1 The Environment as a Larger Candidate Set

Setup

Let $\text{Environment} = \{y_1, \dots, y_m\}$ be a finite set of candidate informational quanta available to the mind, with $m \geq n$. Each y_j carries an energy E_j and, where relevant, interaction terms J_{jk} with other candidates, exactly as in Chapter 8. The mind's eventual quanta $x_1, \dots, x_n$ are precisely those y_j that are retained by the filtering process defined below; Mind is realized as a random subset of Environment, not fixed in advance.

9.2 Retention as a Bernoulli Process

The simplest filtering process, and the one we develop first, treats retention as independent across candidates: each y_j is retained with some probability $r_j \in (0, 1)$, independently of whether any other y_k is retained. This is a genuine simplification, since Chapter 8's whole point was that quanta interact, and we return to the correlated case in Section 9.4 below; the independent case is developed first because it already reproduces the source paper's set-theoretic claim exactly, and because it is the natural baseline against which the correlated case's added structure can be measured.

Proposition 9.1 (Independent retention reproduces the source paper's claim). *Under independent Bernoulli retention with retention probabilities $r_1, \dots, r_m$, for any candidate y_j : if $r_j > 0$, then y_j has positive probability of appearing in Mind; if $r_j = 0$, then y_j never appears in Mind.*

Proof. Immediate from the definition of a Bernoulli trial with parameter r_j : the probability that y_j is retained is r_j by construction, positive whenever $r_j > 0$ and zero whenever $r_j = 0$. $\square$

Proposition 9.1 is deliberately modest; it is very close to a restatement of the source paper's own claim, that a piece of environmental information has some probability of appearing

in Mind if it exists in Environment, and zero probability otherwise. What it adds is not conceptual novelty but a specific mechanism, a Bernoulli process with an explicit parameter r_j , in place of an unspecified probability. This mechanism becomes useful once r_j is tied to the energy structure already developed, which we do next.

9.3 Tying Retention Probability to Energy

The most natural choice, consistent with everything built in Chapters 3 through 4, is to let low-energy candidates be retained with higher probability, exactly as the Gibbs distribution favors low-energy quanta once retained. A convenient and standard choice is the logistic form

$$r_j = \frac{1}{1 + e^{\beta(E_j - \mu)}},$$

where μ is a reference energy level (a candidate at exactly $E_j = \mu$ has retention probability $1/2$) and β plays the same role as in the Gibbs distribution: at high temperature ($\beta \rightarrow 0$), $r_j \rightarrow 1/2$ for every candidate regardless of its energy, and the filter becomes uninformative; at low temperature ($\beta \rightarrow \infty$), $r_j \rightarrow 1$ for candidates with $E_j < \mu$ and $r_j \rightarrow 0$ for candidates with $E_j > \mu$, and the filter becomes a sharp cutoff.

Proposition 9.2 (Conditional distribution given retention). *Suppose retention probabilities are logistic in energy as above, and condition on the event that a candidate y_j is retained. Then, restricted to the retained set and among candidates whose energy sits above the cutoff μ by a comparable margin, the relative odds of retaining y_j over y_k recover the Gibbs ratio $e^{-\beta(E_j - E_k)}$ in the low-temperature limit.*

Proof. For β large and both E_j, E_k above μ by a comparable margin, so that both candidates sit in the logistic function's low-retention tail, the retention probabilities can be approximated as $r_j \approx e^{-\beta(E_j - \mu)}$ and $r_k \approx e^{-\beta(E_k - \mu)}$: when $E_j - \mu \gg 0$, the denominator $1 + e^{\beta(E_j - \mu)}$ is dominated by its second term, so $r_j = 1/(1 + e^{\beta(E_j - \mu)}) \approx e^{-\beta(E_j - \mu)}$. The ratio $r_j/r_k \approx e^{-\beta(E_j - E_k)}$ then matches the Gibbs distribution's own odds ratio $p_j/p_k = e^{-\beta(E_j - E_k)}$ exactly, established already in Chapter 4. $\square$

This proposition is what justifies treating the filtering process and the Gibbs distribution as two views of the same underlying mechanism, rather than as two independent, potentially conflicting models: a mind that filters its environment according to a logistic retention rule tied to energy will, at low temperature, retain candidates in almost exactly the same relative proportions that a Gibbs distribution over those same candidates would assign directly.

9.4 Correlated Retention

The independent case above ignores exactly the structure Chapter 8 introduced: interaction terms J_{jk} that make the joint retention of two candidates more or less costly than retaining either alone. A correlated filter, consistent with that structure, replaces the independent retention probabilities with the full joint distribution over subsets $S \subseteq \text{Environment}$ given by the Gibbs distribution on subsets developed in Chapter 8 directly:

$$P(\text{Mind} = S) = \frac{e^{-\beta E(S)}}{Z}, \quad Z = \sum_{S' \subseteq \text{Environment}} e^{-\beta E(S')}.$$

Under this formulation, indeterminacy and relationality are not two separate mechanisms but a single one: which subset of Environment becomes Mind is exactly as indeterminate,

and exactly as shaped by interaction terms, as the diagrammatic decomposition of Chapter 8 already describes. We do not develop the correlated case further in this chapter, since its full mathematical content is already contained in Chapters 8 and 6; what this section adds is only the reinterpretation of that same content as an account of how Mind comes to be a (random) subset of Environment, closing the loop the source paper leaves as an unspecified set relation.

9.5 Summary

Indeterminacy, formalized, is a Bernoulli or Gibbs-weighted filtering process by which candidates from a larger environmental set are retained into the mind's eventual distribution, with retention probability tied to the same energy and interaction structure already developed. In its correlated form, indeterminacy is not a fourth independent piece of machinery, but the same relational, energy-based apparatus of Chapters 4 through 8, reinterpreted as a statement about which quanta a mind comes to contain rather than merely how it weights the quanta it already has. Chapter 10 turns next to the source paper's account of belief and elasticity, the last of the qualitative claims requiring its own formal chapter before value formation itself, Chapter 11, can be addressed directly.

Elasticity and Belief as Sensitivity

The source paper's discussion of belief change draws an extended analogy to elastic and inelastic scattering: some informational quanta, described as more elastic, are more prone to have their retention probability change under interaction, while others, described as beliefs, resist such change. Four qualitative factors are listed as governing elasticity: the quantum's own motion within the mind, its position relative to other quanta, the interaction potential arising between quanta, and the mind's overall capacity limit. This chapter gives that analogy a literal mathematical form, using the machinery already built.

10.1 Elasticity as a Derivative

Definition 10.1 (Elasticity). The elasticity of quantum i with respect to its own energy is the derivative $\partial p_i / \partial E_i$, evaluated at the current distribution P . The elasticity of quantum i with respect to temperature is $\partial p_i / \partial T$.

This definition captures exactly what elasticity ought to mean in an informational setting: how sensitive a quantum's retention probability is to a small change in its cost, or to a small change in the mind's overall openness. A highly elastic quantum is one whose probability shifts substantially under a small perturbation; a highly inelastic quantum, a belief in the source paper's terminology, is one whose probability barely moves even under a sizable perturbation.

Proposition 10.1 (Elasticity under the Gibbs distribution). *Under the Gibbs distribution $p_i = e^{-\beta E_i} / Z$, the elasticity with respect to E_i is*

$$\frac{\partial p_i}{\partial E_i} = -\beta p_i (1 - p_i).$$

Proof. Write $p_i = e^{-\beta E_i} / Z(\beta)$, where Z itself depends on E_i through the term $e^{-\beta E_i}$ it contains. By the quotient rule,

$$\begin{aligned} \frac{\partial p_i}{\partial E_i} &= \frac{-\beta e^{-\beta E_i} \cdot Z - e^{-\beta E_i} \cdot (-\beta e^{-\beta E_i})}{Z^2} = -\beta \frac{e^{-\beta E_i}}{Z} + \beta \left(\frac{e^{-\beta E_i}}{Z} \right)^2 \\ &= -\beta p_i + \beta p_i^2 = -\beta p_i (1 - p_i). \end{aligned} \quad \square$$

This single formula reproduces, in a precise sense, all four of the source paper's qualitative elasticity factors at once, which is worth spelling out term by term.

10.2 Recovering the Four Qualitative Factors

The magnitude $\beta p_i(1 - p_i)$ depends on exactly three ingredients: the inverse temperature β , and the quantity $p_i(1 - p_i)$, which is itself determined by where p_i sits in $[0, 1]$.

The dependence on β recovers the source paper's fourth factor, the mind's capacity or openness: at high temperature (small β), every quantum's elasticity shrinks toward zero, meaning that in a maximally open mind, no single perturbation to one quantum's energy moves its probability much, precisely because attention is already spread thin across many competitors. At low temperature (large β), elasticity is amplified, meaning that in a closed, highly committed mind, a change in a quantum's energy can produce an outsized shift in its retention probability, at least for quanta not already pinned at $p_i \approx 0$ or $p_i \approx 1$.

The dependence on $p_i(1 - p_i)$ recovers the source paper's first three factors together, since p_i itself, under the Gibbs distribution, already encodes the quantum's energy relative to the rest of the system (its position, in the source paper's language), and the interaction terms of Chapter 8, once folded into E_i as an effective energy, are precisely what determine p_i in the first place. The function $p_i(1 - p_i)$ is maximized at $p_i = 1/2$ and falls to zero at both $p_i = 0$ and $p_i = 1$.

Proposition 10.2 (Beliefs are inelastic). *As $p_i \rightarrow 1$ or $p_i \rightarrow 0$, the elasticity $\partial p_i / \partial E_i \rightarrow 0$.*

Proof. Immediate from Proposition 10.1: $p_i(1 - p_i) \rightarrow 0$ as $p_i \rightarrow 0$ or $p_i \rightarrow 1$, and β is fixed, so the product vanishes in either limit. $\square$

This is the precise mathematical counterpart of the source paper's claim that beliefs, quanta already assigned very high retention probability, are comparatively inelastic: Proposition 10.2 shows that a quantum near certainty, p_i close to 0 or 1, is almost immune to small perturbations in its own energy, regardless of how large β is, simply because there is very little remaining probability mass left to move. A quantum in the middle of the distribution, p_i near $1/2$, genuinely undecided between being retained and not, is maximally sensitive to exactly such a perturbation.

Worked Example

Return to the toy mind of Chapter 4, at $T = 1$ ($\beta = 1$), where $p_1 \approx 0.560$. The elasticity of quantum 1 is $-\beta p_1(1 - p_1) \approx -1 \times 0.560 \times 0.440 \approx -0.246$: a moderate sensitivity, reflecting that x_1 , while the most probable quantum, is not yet so dominant as to be immune to a further increase in its own cost. Quantum 4, with $p_4 \approx 0.028$, has elasticity $\approx -1 \times 0.028 \times 0.972 \approx -0.027$, an order of magnitude smaller: x_4 is already so marginal that a further increase in its energy has little further probability left to remove.

10.3 Elasticity with Respect to Temperature

The same style of computation applies to sensitivity with respect to T rather than E_i , connecting elasticity directly to the entropy-monotonicity result of Chapter 5.

Proposition 10.3. *Under the Gibbs distribution, $\partial p_i / \partial \beta = -p_i(E_i - \bar{E})$, where $\bar{E} = \sum_j p_j E_j$ is the expected energy of Proposition 5.1.*

Proof. Take logarithms of $p_i = e^{-\beta E_i} / Z(\beta)$: $\ln p_i = -\beta E_i - \ln Z(\beta)$. Differentiating both sides with respect to β ,

$$\frac{1}{p_i} \frac{\partial p_i}{\partial \beta} = -E_i - \frac{d}{d\beta} \ln Z(\beta),$$

so $\frac{\partial p_i}{\partial \beta} = p_i \left(-E_i - \frac{d}{d\beta} \ln Z(\beta) \right)$. Using Proposition 5.1, $\frac{d}{d\beta} \log Z = -\bar{E}$, this becomes

$$\frac{\partial p_i}{\partial \beta} = p_i(-E_i + \bar{E}) = -p_i(E_i - \bar{E}). \quad \square$$

This says that as the mind cools (as β increases), quanta with below-average energy gain probability ($E_i < \bar{E}$ gives $\partial p_i / \partial \beta > 0$) and quanta with above-average energy lose it, exactly the redistribution that Proposition 4.2 showed converges, in the limit, to a point mass on the single lowest-energy quantum.

10.4 Summary

The source paper's four qualitative elasticity factors, motion, position, interaction, and capacity, are recovered here as the two ingredients of a single closed-form derivative, $\partial p_i / \partial E_i = -\beta p_i(1 - p_i)$: the capacity or openness factor entering through β , and the remaining three factors entering jointly through p_i itself, which already encodes energy, relative position, and interaction structure by construction. Beliefs, quanta near certainty, are shown to be genuinely and provably inelastic, not merely described as such. Chapter 11 now has every piece needed to state and prove the source paper's central, still-unproven claim: that the formation of values lowers the entropy of the mind's informational state.

Chapter 11

Value Formation and Entropy Reduction

Every chapter of Part II has been preparation for this one. The source paper's central claim, the one its own title foregrounds, is that value formation is a process of entropy reduction: a mind that distinguishes and prioritizes some subset of its informational quanta has, by that very act, lowered the entropy of its informational state relative to a mind that has not yet made any such distinction. Chapter 1 identified this as the single most important gap in the source paper, asserted but never derived. This chapter closes it.

11.1 Value Formation, Defined

Definition 11.1 (Value formation). A mind is said to have formed values, relative to a baseline of complete indifference, if its distribution P departs from the uniform distribution U , that is, if $P \neq U$.

This definition is deliberately the simplest one consistent with the source paper's own language: a mind with no basis for distinguishing among its informational quanta is, by definition, uniform, since nothing favors any quantum over any other; any departure from uniformity is, by the same logic, evidence that some distinguishing process, some assignment of differential cost or interaction structure, has taken place.

11.2 The Value Formation Theorem

Proposition 11.1 (Value formation strictly lowers entropy). *Let P be the Gibbs distribution at inverse temperature $\beta > 0$ over energies $E_1, \dots, E_n$, not all equal. Then*

$$H(P) < H(U) = \log_2 n,$$

and the exact size of the gap is

$$H(U) - H(P) = D_{KL}(P\|U) \geq 0,$$

with $D_{KL}(P\|U) > 0$ strictly whenever $P \neq U$.

Proof. The identity $H(U) - H(P) = D_{KL}(P\|U)$ was established in the proof of Proposition 2.1 (Chapter 2), applied there in the other direction to show $H(P) \leq H(U)$. By Lemma 2.1 (Gibbs' inequality), $D_{KL}(P\|U) \geq 0$, with equality only when $P = U$. It remains to show $P \neq U$

whenever $\beta > 0$ and the energies are not all equal, which is immediate: if the E_i are not all equal, then $e^{-\beta E_i}$ is not constant across i for any $\beta > 0$, so the Gibbs weights $p_i = e^{-\beta E_i} / Z$ are not all equal, and $P \neq U$. Hence $D_{KL}(P\|U) > 0$ strictly, giving $H(P) < H(U)$ strictly. $\square$

This is the source paper's central claim, proved rather than asserted, and it is worth being precise about exactly what has and has not been shown. What has been shown is that under the specific mechanism developed across Chapters 3 through 8, maximum-entropy inference under an energy constraint, formalized as differential retention costs among competing informational quanta, any nontrivial differentiation among those costs necessarily and strictly lowers entropy relative to complete indifference, and the KL divergence gives the exact size of that reduction, not merely its sign. What has not been shown, and what we do not claim, is that this is the only possible mechanism by which entropy could fall, or that every real cognitive or social process of value formation actually proceeds via a Gibbs distribution specifically; that identification remains an interpretive choice, defended in Chapter 1 on the grounds that it is the natural mathematical object for the job, not derived from any further principle.

11.3 KL Divergence as a Value Intensity Index

Proposition 11.1 suggests a natural, non-negative, threshold-free measure of how strongly a mind has formed values at all, independent of which particular quanta it favors.

Definition 11.2 (Value intensity). The value intensity of a distribution P is $D_{KL}(P\|U)$.

Unlike the threshold-based notion of a value quantum used in Proposition 7.1 (Chapter 7), which asks how many individual quanta clear a fixed bar, value intensity is a single scalar describing the distribution as a whole: how far, in an exact information-theoretic sense, it sits from the state of total indifference. A mind entirely without values, $P = U$, has value intensity exactly zero; a mind that has collapsed entirely onto a single quantum, the low-temperature limit of Proposition 4.2, has value intensity $\log_2 n$, the maximum possible, since $H(P) \rightarrow 0$ in that limit and $D_{KL}(P\|U) = \log_2 n - H(P) \rightarrow \log_2 n$.

Worked Example

For the toy mind at $T = 1$ from Chapter 4, $P \approx (0.560, 0.206, 0.206, 0.028)$. Its entropy is $H(P) = -\sum_i p_i \log_2 p_i \approx 1.552$ bits, against $H(U) = \log_2 4 = 2$ bits, giving value intensity $D_{KL}(P\|U) \approx 0.448$ bits: a moderate, partial formation of values, consistent with the fact that this distribution, while clearly favoring x_1 , has not yet collapsed to anything close to certainty.

11.4 Connecting Back to Herding and Borrowed Conviction

It is worth noting, briefly, that value intensity gives a natural way to distinguish conviction that has been earned, built up through genuine differentiation among costly alternatives, from conviction that merely mimics low entropy without the underlying differentiation that would normally produce it. A distribution copied wholesale from another agent's already-low-entropy state carries the same value intensity as one arrived at independently, even though only the latter required the mind's own energy landscape to do any work; this distinction, between earned and borrowed entropy reduction, is a natural further application of the apparatus this chapter develops, one this book does not pursue in its own right.

11.5 Summary

The source paper's title claim, that value formation is entropy reduction, now has a proof: under the Gibbs-distribution mechanism developed across this book, any nontrivial differentiation among the costs of competing informational quanta strictly lowers entropy relative to the uniform baseline, with Kullback-Leibler divergence giving the exact, non-negative size of that reduction. This completes Part II: granularity, relationality, indeterminacy, elasticity, and value formation itself are now each backed by a specific proposition rather than a qualitative claim, closing every gap Chapter 1 identified. Chapter 12 turns next to whether any of this machinery earns its keep against a genuinely difficult applied problem, beginning with the valuation of Bitcoin.

Applications

12.1 Bitcoin Valuation

Chapter 1 introduced Bitcoin as the sharpest available illustration of the gap this book was written to close: an asset with no cash flow, no productive output, and no legal claim on anything, whose valuation nonetheless converges, across millions of independent participants, on a shared order of magnitude. This section works that example through fully, mapping the formal apparatus of Part II onto Bitcoin's specific informational structure, deriving a falsifiable prediction from that mapping, and checking its direction against published empirical literature on Bitcoin price dynamics.

Two Subsystems, Two Energy Regimes

We propose, motivated directly by the mind-environment apparatus of Chapters 4 through 8, a four-dimensional decomposition of the informational quanta relevant to a Bitcoin holder or observer: Store-of-Value (SOV), Autonomy (AUT), Social-Signal Value (SSV), and Hedonic-Sentiment Value (HSV). We treat this decomposition as a modeling choice motivated by the structure of the asset itself, not as something derived from the formalism; the formalism's job is to say what these four dimensions ought to look like once cast as informational quanta with energies, and what that casting predicts.

Setup

Let the relevant informational quanta for a Bitcoin holder or observer be organized into two groups. The technical group, comprising SOV and AUT, reflects protocol-level facts: scarcity encoded in the fixed 21-million-coin supply, cryptographic irreversibility, and the decentralized, self-custodial character of on-chain transactions. The social group, comprising SSV and HSV, reflects narrative and sentiment: collective attention, as captured by search intensity, and emotional state, as captured by a fear-greed index. We assign the technical group low effective energy, $E_{\text{SOV}}, E_{\text{AUT}}$ small, reflecting that protocol facts are cheap to verify and highly stable, changing only on the timescale of software upgrades or halving events, and we assign the social group energies that fluctuate substantially and rapidly, $E_{\text{SSV}}(t), E_{\text{HSV}}(t)$ varying with t on a daily or even sub-daily basis, reflecting the volatility of narrative and mood.

This is an interpretive choice, not a derivation; nothing in Part II specifies how SOV, AUT, SSV, and HSV in particular ought to be assigned energies. What justifies the choice is a general principle, developed already in Chapter 4: quantities that are stable and costly to alter are naturally represented as low, slowly varying energy, while quantities that are volatile and

cheap to shift are naturally represented as high, rapidly varying energy. Protocol facts and sentiment indices differ in exactly this respect, which is what motivates the mapping.

An Effective Temperature for the Market

Under this mapping, the market's collective distribution over competing Bitcoin narratives ("digital gold," "speculative bubble," "payment rail," and so on) is modeled as approximately Gibbs-distributed over these energies, with an effective temperature $T(t)$ that itself varies with the relative weight of social versus technical energy at a given moment. When social energy dominates, the effective system behaves as though at high temperature: many competing narratives remain simultaneously plausible, entropy is high, and Proposition 4.1 says the distribution should be closer to uniform across narratives, price movements driven more by which narrative currently commands attention than by any structural anchor. When technical energy dominates the effective landscape, in periods of narrative quiet, the system behaves as though at lower temperature, and Proposition 4.2 says probability should concentrate on the structurally cheapest, most stable narrative, "digital gold" or "store of value," with correspondingly lower volatility.

Proposition 12.1 (Qualitative prediction: horizon-dependent dominance). *If social-group energies fluctuate on a shorter timescale than technical-group energies, then, under the mapping above, short-horizon price movements should be more strongly associated with social-group quanta (SSV, HSV) than with technical-group quanta (SOV, AUT), while over longer horizons, as short-lived social fluctuations average out, the technical group's comparatively stable, low-energy pull should become more visible in the data.*

This proposition follows from nothing more than the ordinary statistical fact that averaging a rapidly fluctuating series over a longer window suppresses its variance relative to a slowly fluctuating series, combined with the energy-timescale assignment already stated; we do not present a separate formal proof, since the claim is a direct qualitative consequence of that assignment rather than a new mathematical result of its own, and we prefer to state it plainly as an interpretation rather than dress it in more formal apparatus than it needs.

Checking the Prediction Against Real Market Data

Proposition 12.1 is a genuine, falsifiable prediction, and it happens to have been tested directly, using exactly the four-dimensional SOV/AUT/SSV/HSV decomposition proposed here rather than a looser proxy for it. Vuong et al. (2025b) estimate a Bayesian linear model of Bitcoin's daily returns from 2022 to 2025, operationalizing SOV, AUT, SSV, and HSV as separate regressors and reporting full posterior distributions for each. This is a stronger empirical anchor than the general attention-and-sentiment literature discussed below, since it tests the specific decomposition Section 12.1 constructs, not merely a related one.

Predictor (lagged change)	Mean	SD	95% HPDI low	95% HPDI high
Store-of-Value (Δ SOV)	0.00084	0.00074	−0.00062	0.00231
Autonomy (Δ AUT)	0.00073	0.00073	−0.00073	0.00219
Social-Signal Value (Δ SSV)	0.00324	0.00072	0.00183	0.00465
Hedonic-Sentiment Value (Δ HSV)	0.00058	0.00073	−0.00087	0.00202

Table 12.1: Posterior estimates for next-day Bitcoin returns on lagged changes in the four value dimensions, daily data 2022–2025 (Vuong et al. 2025b). Only SSV's 95% highest posterior density interval (HPDI) excludes zero entirely.

Table 12.1 reports the posterior means and 95% highest posterior density intervals (HPDIs) for each predictor. Only the social-signal predictor, ΔSSV , has a 95% HPDI lying entirely above zero, $[0.00183, 0.00465]$, a robust, highly credible positive association with next-day returns. SOV and AUT both have positive posterior means whose HPDIs extend slightly into negative territory, $[-0.00062, 0.00231]$ and $[-0.00073, 0.00219]$ respectively; Vuong et al. (2025b) treat these as moderately reliable rather than fully credible, on the grounds that each posterior mean is comparable in magnitude to its own standard deviation. HSV's HPDI, $[-0.00087, 0.00202]$, shows no credible effect at all. Figure 12.1 plots these four intervals directly.




Figure 12.1: Posterior mean and 95% HPDI for each of the four value dimensions' association with next-day Bitcoin returns, redrawn from the estimates of Vuong et al. (2025b). SSV is the only interval excluding zero.

This finding maps cleanly onto Proposition 12.1's short-horizon prediction: at the one-day horizon, the single social-group predictor with genuinely credible explanatory power is SSV, while both technical-group predictors, SOV and AUT, sit at the boundary of credibility, positive but not decisively so. Vuong et al. interpret this pattern in language that maps directly onto this section's own energy decomposition, describing Bitcoin as a "dual-layer entropy-regulating socio-technological ecosystem" in which "the low-entropy technical subsystem, characterized by protocol immutability, cryptographic security, and fixed supply, provides structural stability," while "the high-entropy socio-informational field, comprising narratives, memes, media signals, and attention waves," dominates short-run price dynamics, language chosen independently of this book but strikingly close to Section 12.1's own technical-group/social-group framing.

The general literature on Bitcoin attention and sentiment corroborates this pattern using different proxies and methods, without testing the specific SOV/AUT/SSV/HSV decomposition itself. Figá-Talamanca & Patacca (2019) find that market attention, measured through search and trading-volume proxies, is significantly associated with Bitcoin returns and volatility. Dastgir et al. (2019) report a causal relationship running from investor attention to Bitcoin returns using a copula-based Granger causality test. A wavelet-coherence study of sentiment and returns over 2016 to 2021 (Aysan et al. 2023) finds that the sentiment-return relationship is strongest at higher frequencies and weakens at lower frequencies, consistent with the horizon-dependence half of Proposition 12.1 that the single-day study above does not, by itself, test.

We want to be precise about how much weight this comparison can bear even so. Vuong et al. (2025b)'s Bayesian estimates test the short-horizon half of Proposition 12.1 directly and find results consistent with it; they do not, on their own, establish the longer-horizon half, that technical-group predictors should dominate as the window lengthens, since the reported model is estimated at daily frequency throughout. Nor does either study validate the specific

energy-and-temperature mechanism proposed in Section 12.1; a regression finding SSV credible and SOV, AUT, HSV less so is consistent with, but does not require, this book's Gibbs-distribution machinery specifically. We regard the direct SOV/AUT/SSV/HSV evidence as the single strongest empirical anchor in this book, and the broader literature as corroborating context, while maintaining the same caution as elsewhere: consistency in direction is not the same as a quantitative test of the formalism itself.

Why Classical Value Theories Miss This

It is worth returning briefly to the comparative table promised in Chapter 1. A cost-of-production theory of value has no natural place for SSV or HSV at all, since neither reflects any production cost; a pure scarcity theory captures SOV but has nothing to say about why AUT, SSV, or HSV would matter, or why their relative importance would shift with horizon. The entropy-based account developed here is not obviously superior as an all-purpose theory of value, and we do not claim it is, but it does something the classical alternatives structurally cannot: it predicts, rather than merely accommodates after the fact, that an asset's value drivers should shift character across time horizons whenever those drivers differ in how quickly their underlying informational energy fluctuates. That the direct SOV/AUT/SSV/HSV evidence in Table 12.1 shows exactly the social-dominance pattern predicted at the one-day horizon is a genuine, if partial, piece of evidence that the formalism has done more than restate the original qualitative claims in heavier notation.

Summary

We mapped a theoretically motivated four-dimensional decomposition of Bitcoin's perceived value onto the energy and temperature structure of Part II, derived a qualitative, falsifiable prediction about horizon-dependent value drivers, and found that prediction borne out directly at the short horizon by a Bayesian regression using exactly this book's own SOV/AUT/SSV/HSV decomposition, with SSV the only predictor whose 95% credible interval excludes zero. This is the most fully worked application in this book, and the strongest empirical anchor any single section of it has, though the longer-horizon half of Proposition 12.1 and the underlying energy-and-temperature mechanism itself both remain untested by this evidence. Section 12.2 extends the same temperature apparatus to a second market phenomenon, the discontinuous character of bubbles and crashes.

12.2 Phase Transitions: Bubbles and Crashes

Section 12.1 treated temperature as a smoothly varying quantity, tracking the relative weight of social versus technical energy. This section asks what happens when the market's underlying energy landscape itself has more than one competing low-energy state, more than one narrative that could plausibly anchor collective belief, and shows that the Gibbs-distribution machinery already developed predicts something markets are well known to exhibit: sudden, discontinuous shifts in collective valuation, familiar informally as bubbles and crashes, arising from continuous changes in underlying conditions.

A Bimodal Energy Landscape

Setup

Suppose two narratives, call them the optimistic narrative and the pessimistic narrative, each anchor a local low-energy region of the space of informational quanta, with energies $E_{\text{opt}}(\theta)$ and $E_{\text{pess}}(\theta)$ depending on some external parameter θ , macroeconomic conditions, regulatory news, or accumulated sentiment, that shifts gradually. At θ below some critical value θ^* , $E_{\text{opt}} < E_{\text{pess}}$; above θ^* , the ordering reverses.

Proposition 12.2 (Discontinuous mode shift at low temperature). *In the low-temperature limit ($T \rightarrow 0^+$), the mode of the Gibbs distribution jumps discontinuously from the optimistic narrative to the pessimistic narrative as θ crosses θ^* , even though $E_{\text{opt}}(\theta)$ and $E_{\text{pess}}(\theta)$ vary continuously in θ .*

Proof. By Proposition 4.2, the low-temperature Gibbs distribution concentrates entirely on the lowest-energy state. For $\theta < \theta^*$, this is the optimistic narrative, so $P \rightarrow (1, 0)$ in the two-narrative case; for $\theta > \theta^*$, it is the pessimistic narrative, so $P \rightarrow (0, 1)$. At $\theta = \theta^*$ exactly, the two energies are equal and the low-temperature limit is $(1/2, 1/2)$. The map from θ to the limiting distribution is therefore discontinuous at θ^* : it jumps directly from concentration on one narrative to concentration on the other, with no continuous interpolation between them as θ crosses the critical point, regardless of how smoothly E_{opt} and E_{pess} themselves vary. $\square$

At any fixed positive temperature, by contrast, the transition across θ^* is smooth: Proposition 10.1 shows the probability of either narrative changes continuously and, near θ^* where the two energies are close, with elasticity close to its maximum, $p(1-p)$ near $1/4$. The discontinuity of Proposition 12.2 is a low-temperature phenomenon specifically; it is exactly the sense in which a market that has become highly convicted, low T , narrow, concentrated belief, is more prone to a sudden narrative flip than one that has retained a genuinely open, high-temperature mixture of views. This gives formal content to a piece of informal market wisdom: conviction and fragility are not opposites in this model, but two faces of the same low-temperature state.

Path Dependence: Why Reversion Is a Crash, Not a Drift

Proposition 12.2 explains why a narrative flip can occur, but not why the market's actual return path after a bubble so often looks like a sharp crash rather than a gradual reversion back toward the pre-bubble baseline. For that, we draw on an independent formal result, derived from scratch and in full in Appendix A.

The result, restated briefly here and proved in full in Appendix A as the Shifting-Baseline Ratchet (SBR), concerns a checking procedure that compares each revision only to its immediately preceding state, never to a fixed original: a sequence of revisions each individually too small to be flagged against a per-step perceptibility threshold can nonetheless accumulate to an arbitrarily large total drift, with no cancellation between steps. Applied here, purely by analogy and not by any formal reduction of one result to the other, this offers a natural account of how a market's baseline price expectation can drift substantially higher during a bubble's build-up without ever producing a single price movement large enough, on its own, to be perceived as alarming: each day's gain is compared only to the previous day's price, not to some fixed, independently anchored fair value, so the cumulative departure from that original anchor goes unflagged by the market's own moving standard of what counts as a large move.

Under this reading, the eventual correction cannot proceed by the same mechanism that built the bubble, since there is no longer a small, incrementally shifting baseline to hide a comparably sized reversion against; the market's baseline has already ratcheted upward. A return toward the original anchor must instead occur as a single, comparably large movement, that is, a crash, precisely because incremental, sub-threshold steps were the mechanism of ascent but are no longer available, in the same incremental form, as a mechanism of descent once the low-temperature, concentrated state of Proposition 12.2 has taken hold.

We want to be precise about what is and is not being claimed here. The shifting-baseline result of Appendix A is proved in an entirely different formal setting, a deterministic sequence of revisions against a moving comparison point, with no probability distribution, no entropy, and no temperature anywhere in its statement. Nothing in this section derives Proposition 12.2 from that result, or that result from the Gibbs-distribution machinery of Part II. The two results are independent formalizations, developed for different purposes within this book, that we regard as consistent in spirit, both describe an asymmetry between how gradually a distortion can accumulate and how abruptly it must be corrected, but the connection drawn in this section is an interpretive analogy, not a theorem. Section 17.2 returns to this distinction explicitly.

Summary

A bimodal energy landscape, combined with the low-temperature limit already established in Chapter 4, produces a genuine discontinuity, a sudden narrative flip, as a continuously varying external parameter crosses a critical threshold, giving formal content to the informal notion of a market bubble bursting rather than deflating gradually. Drawing an explicit, clearly flagged analogy to the shifting-baseline result proved in Appendix A, we suggested why the ascent into such a state and the descent out of it need not be mechanically symmetric.

12.3 Bird Conservation and Cultural Belief Drift

The two applications above concern markets. This section applies Appendix A's machinery to a very different domain, cultural belief about wildlife, where the underlying phenomenon, a moving-baseline comparison letting cumulative change escape detection, is documented directly in the conservation literature rather than offered here as a first analogy.

Nguyen & Khuc (2026) document a striking illustration from Vietnam: a restaurant sign advertising wild birds as a delicacy, its name doubling as a sexual pun, normalizing wildlife consumption as a marker of status and prosperity rather than registering it as a conservation concern. The paper's central diagnostic concept, order-induced blindness, describes a historically constructed system of belief, treated as self-validating and complete, that reduces cognitive friction at the cost of perceptual openness, obscuring its own drift from ecological reality. Nguyen & Ho (2026) formalize the broader methodology this diagnosis belongs to, using literary and philosophical material to surface exactly this kind of latent inconsistency within an otherwise internally coherent framework.

Setup

Model a culture's belief about wildlife as a sequence of distributions $P_0, P_1, \dots, P_N$ over a binary outcome, protection-worthy versus commodity, updated across generations or social cycles, with each update compared only to the immediately preceding generation's norm, never to a fixed ecological baseline Q representing the actual state of wildlife populations.

Proposition A.3 (Appendix A) applies directly: a culture's stance can drift the full distance from strong protective conviction to strong commodification conviction through generational steps each too small, on its own, to register as a moral violation against the immediately preceding norm, exactly the pattern Nguyen & Khuc (2026) document empirically and exactly the structure Appendix A's Animal Farm illustration already uses for a fixed original rather than a moving one.

Worked Example

Let $P_0 = (0.6, 0.4)$ represent a mildly protective norm and $P^* = (0.05, 0.95)$ a strongly commodified one, with $TV(P_0, P^*) = 0.55$. Twelve equal generational steps of size $0.55/12 \approx 0.046$, comfortably below a perceptibility threshold of $\tau = 0.05$, carry the culture the full distance, exactly as Proposition A.3 guarantees; we verified this numerically before writing it here. Holding a fixed ecological baseline $Q = (0.9, 0.1)$, reflecting a population genuinely warranting protection, $TV(P_0, Q) = 0.30$ initially and $TV(P_N, Q) = 0.85$ after the drift: the culture has moved from moderately misaligned with ecological reality to severely misaligned, without any single generational step ever being large enough to prompt reconsideration.

This worked example also illustrates, without claiming to prove, a pattern the empirical literature on uncertainty and absurdity discusses under the name uncertainty-absurdity mutuality (UAM): the drifting belief's own entropy falls across the same sequence, from $H(0.6, 0.4) \approx 0.971$ bits to $H(0.05, 0.95) \approx 0.286$ bits, meaning the culture becomes more internally confident and coherent, by exactly the entropy-reduction mechanism Chapter 11 formalizes, at precisely the same time as its divergence from ecological reality grows. Rising internal certainty and rising factual error are not, in this example, in tension with one another; they are produced by the same mechanism. We flag this connection as an illustrative observation drawn from a single worked example, not a further theorem, and note it only because the underlying empirical claim, that this pattern is a real, documented phenomenon rather than a coincidence of these particular numbers, is not this book's own; it belongs to the literature Nguyen & Khuc (2026) and Nguyen & Ho (2026) are part of, which this section borrows a term from rather than develops further.

We want to be direct about the limits of this section, in the same spirit as Section 12.2's own closing caveat. Modeling a culture's stance as a single two-outcome distribution updated along a literal straight-line path is a simplification with no claim to describing how belief actually moves through a population; real cultural change is almost certainly not linear, not uniform across individuals, and not free of the very feedback effects, media, policy, price, that a fuller model would need to include. What Proposition A.3 establishes is that the moving-baseline mechanism is mathematically capable of producing this kind of undetected drift; whether it is the operative mechanism behind any specific instance of wildlife-consumption normalization is an empirical question this book does not settle.

Summary

This closes this chapter's applications of the entropy apparatus built in Parts I and II. Chapter 13 turns next to a genuinely new piece of mathematical development, an absorbing-barrier stochastic model of irreversible ecological loss, before Chapter 17 collects the testable, falsifiable predictions this chapter's three sections actually commit to.

Part III

A New Perspective on Environmental Values

Irreversibility: A Stochastic Model of Ecological Value

Every chapter so far has taken the source paradigm's own qualitative claims as its target: granularity, relationality, indeterminacy, elasticity, and entropy reduction were each already present, in words, in Vuong et al. (2025a), and this book's task was to give them mathematical form. This chapter and the one that follows it do something different. They introduce mathematical machinery that has no counterpart in the source paradigm at all, aimed at a problem the source paradigm does not address: why a monetary valuation apparatus built on the assumption that value can be freely exchanged in both directions systematically fails when applied to ecological assets, where the exchange is not symmetric.

We borrow the vocabulary of option pricing and stochastic valuation throughout this chapter and the next because it is a precise, well-developed language for reasoning about irreversibility, not because we take Part III to be proposing that ecological value should be priced and traded as a financial asset. The mathematics is a tool for sharpening a specific intuition, that some losses cannot be undone and that ignoring this fact leads to quantifiably wrong valuations, and the tool should not be mistaken for the point.

13.1 Why Symmetric Valuation Fails

A recurring finding in the empirical and policy literature on climate change and ecological degradation is that conventional economic valuation, treating an ecosystem's worth as a price to be discovered, discounted, and traded like any other asset, persistently understates the risk of catastrophic, irreversible loss. Carbon markets price emissions as though a ton avoided today and a ton avoided next decade were freely substitutable; ecosystem service valuations price a wetland's flood protection as a discounted stream of avoided damages, as though the wetland itself could always, in principle, be restored if the price were right. This concern has a long history in environmental economics specifically: Arrow & Fisher (1974) first showed that irreversibility alone, independent of any risk aversion, provides a rationale for greater caution in developing natural environments than a naive expected-value calculation would suggest, and Dixit & Pindyck (1994) developed the general real-options framework in which that insight is now standardly formalized, treating the decision to forgo an irreversible action as itself an option with positive value. Pindyck (2000) applies this framework directly to environmental policy, arguing formally that both the costs of environmental regulation and the benefits of avoided degradation can be partly or wholly irreversible, and that standard cost-benefit analysis, built for reversible, marginal trade-offs, systematically mishandles the resulting option value of waiting or acting under uncertainty. Lenton et al. (2019) document the physical

counterpart directly: climate subsystems including ice sheets, permafrost, and the Amazon rainforest exhibit tipping points beyond which change becomes self-sustaining and effectively irreversible on human timescales, a body of evidence this chapter's absorbing-barrier model is built to take seriously as a mathematical constraint rather than a qualifying footnote. Vuong (2021) names the structural assumption underlying this failure directly: mainstream valuation treats environmental value and monetary value as exchangeable in both directions, when in fact, once genuine ecological collapse occurs, no quantity of money reliably buys it back. Vuong calls the correct relationship the *semiconducting principle*: environmental value can be converted into monetary value, but the reverse conversion is not, in general, available. A collapsed fishery, a extinct species, a vanished aquifer are not, as a rule, restored by capital investment in the way a depreciated machine is.

Detecting a collapse threshold before it is crossed depends on a monitoring regime actually observing the state variable in question, and real ecological monitoring is neither continuous nor guaranteed to be well timed: satellite revisit intervals, funding-driven survey discontinuities, and gaps in biodiversity census programs all mean that a population or biomass assessment might simply not occur during the window in which a collapse-sized change happens. This is a different failure mode from the gradual, sub-threshold drift Appendix A formalizes as Proposition A.1; it is the comparison-suspension channel of Proposition A.2 in the same appendix, in which a single change large enough to be flagged with certainty escapes detection anyway because the checking act itself did not occur. Applied here, purely as an interpretive analogy in the same spirit as Section 12.2's use of Proposition A.1, Proposition A.2's finding that detection probability depends on whether monitoring happens at all, not on how large the change being monitored for is, offers a formal reason why intermittent monitoring regimes can fail to catch even a sharp, unambiguous ecological collapse: the barrier B in this chapter's model is only ever as good as the observation process checking whether X_t has reached it.

This chapter's own model is closer in spirit to the real-options tradition of Arrow & Fisher (1974) and Dixit & Pindyck (1994) than to the semiconducting principle's own qualitative statement, in one specific technical sense: both traditions locate irreversibility in a boundary condition on an otherwise standard stochastic valuation problem, rather than in a separate accounting identity. What this chapter adds to that tradition is the specific closed-form solution developed below, obtained by treating the boundary as a strictly absorbing barrier rather than as a smooth option-exercise decision, and the connection, made explicit in Chapter 14, back to this book's own entropy-based apparatus.

This chapter takes that asymmetry seriously as a mathematical constraint, not merely a rhetorical one. We model an ecological asset's value as a continuous-time stochastic process, and we introduce a single new structural feature, an absorbing barrier representing an irreversible collapse threshold, that a conventional valuation model omits. Chapter 14 then shows, in closed form, exactly how much a valuation that ignores this feature overstates the asset's worth, and exactly how that overstatement connects to the entropy-based apparatus built across the rest of this book.

The reader should not expect any of the machinery in this chapter to be visible in Vuong et al. (2025a); it is not there. What is there, and what motivates the specific mathematical choices made here, is the source paradigm's own emphasis on irreversibility in a different guise, the claim, central to Chapter 11, that entropy reduction is directional: value formation lowers entropy, and nothing in that chapter's machinery runs the argument backward. This chapter's absorbing barrier is the continuous-time, stochastic counterpart of that same directionality, applied now to a process that can move toward higher disorder as well as lower, and asking what happens once it crosses a point of no return.

13.2 Brownian Motion and Stochastic Differential Equations, from Scratch

We build everything in this chapter from a single primitive object, standard Brownian motion, assuming no prior exposure. Readers already familiar with stochastic calculus may skip to Section 13.4.

Definition 13.1 (Standard Brownian motion). A standard Brownian motion is a continuous-time stochastic process $\{W_t\}_{t \geq 0}$ satisfying: (i) $W_0 = 0$; (ii) for any $0 \leq s < t$, the increment $W_t - W_s$ is normally distributed with mean 0 and variance $t - s$; (iii) increments over disjoint time intervals are independent; (iv) $t \mapsto W_t$ is continuous almost surely.

Brownian motion is the continuous-time limit of a random walk: over a short interval of length dt , the increment dW_t behaves like a mean-zero, variance- dt random shock, uncorrelated with shocks over any other interval. This single property, variance accumulating linearly in time, is the origin of the characteristic scaling $\sqrt{dt}$ that recurs throughout this chapter.

Setup

A stochastic differential equation (SDE), the random-driven counterpart of the ordinary differential equations treated in Boyce & DiPrima (2017), is of the form

$$dX_t = \mu(X_t, t) dt + \sigma(X_t, t) dW_t$$

describes a process whose infinitesimal change over $[t, t + dt]$ has a deterministic part, the drift $\mu(X_t, t) dt$, and a random part, the diffusion $\sigma(X_t, t) dW_t$, driven by the same Brownian motion W_t throughout. Formally, X_t is defined as the solution to the integral equation $X_t = X_0 + \int_0^t \mu(X_s, s) ds + \int_0^t \sigma(X_s, s) dW_s$, where the second integral is an Itô integral; we do not construct the Itô integral from first principles here, referring the reader to Øksendal (2003) or Björk (2009) for the full measure-theoretic construction, and instead work with the informal calculus rules, developed in Section 13.3, that make SDEs usable without that machinery.

Definition 13.2 (Geometric Brownian motion). Geometric Brownian motion (GBM) with drift μ and volatility $\sigma > 0$ is the process solving

$$dX_t = \mu X_t dt + \sigma X_t dW_t, \quad X_0 > 0.$$

GBM is the natural choice for a quantity, like an ecological value index, that should remain positive and whose proportional (rather than absolute) fluctuations are what matter: a $\pm 5\%$ shock is treated as comparably significant whether the index currently stands at 50 or at 500. Section 13.3 derives its explicit solution.

It is worth pausing on what happens to this same equation when the noise term is switched off and the drift is negative, since the result is a second, independent formalization of irreversible depletion, proved from scratch as the Extraction Limit in Appendix A. Setting $\sigma = 0$ and $\mu = -\rho < 0$ for a constant proportional rate $\rho > 0$ reduces $dX_t = \mu X_t dt$ to exactly the continuous-time Extraction Limit, $dX/dt = -\rho X(t)$, with closed-form solution $X(t) = X_0 e^{-\rho t} \rightarrow 0$. Chapter 12 and this chapter both use the language of “just a little, just for now” to describe how a constant proportional extraction rate is typically justified at each occasion it recurs; Proposition A.4 makes precise what that language obscures, that no constant $\rho > 0$, however modest, has any long-run destination other than exhaustion. GBM’s own drift term $\mu X_t dt$ is, in this precise sense, a direct generalization of the Extraction Limit’s deterministic skeleton, with volatility layered on top; the absorbing barrier this chapter builds around GBM

formalizes what happens when that deterministic decay is also exposed to noise on its way toward zero.

13.3 Itô's Lemma

The ordinary chain rule fails for functions of Brownian motion, because $(dW_t)^2$ does not vanish to first order in dt the way $(dt)^2$ does; instead, $(dW_t)^2$ behaves like dt itself. This single fact, stated without proof here (see Øksendal (2003), Chapter 3, or Björk (2009), Chapter 4, for the rigorous justification via quadratic variation), is the entire content of stochastic calculus's departure from ordinary calculus.

Proposition 13.1 (Itô's lemma). *Let X_t solve $dX_t = \mu_t dt + \sigma_t dW_t$, and let $f(x, t)$ be twice continuously differentiable in x and once in t . Then*

$$df(X_t, t) = \left(\frac{\partial f}{\partial t} + \mu_t \frac{\partial f}{\partial x} + \frac{1}{2} \sigma_t^2 \frac{\partial^2 f}{\partial x^2} \right) dt + \sigma_t \frac{\partial f}{\partial x} dW_t.$$

Heuristic derivation. Taylor-expand f to second order in dX_t and first order in dt :

$$df = \frac{\partial f}{\partial t} dt + \frac{\partial f}{\partial x} dX_t + \frac{1}{2} \frac{\partial^2 f}{\partial x^2} (dX_t)^2 + \dots$$

Substitute $dX_t = \mu_t dt + \sigma_t dW_t$ and expand $(dX_t)^2 = \mu_t^2 (dt)^2 + 2\mu_t \sigma_t dt dW_t + \sigma_t^2 (dW_t)^2$. Terms in $(dt)^2$ and $dt dW_t$ vanish faster than dt and are dropped; the term $\sigma_t^2 (dW_t)^2$, using $(dW_t)^2 \rightarrow dt$, contributes $\sigma_t^2 dt$. Collecting all terms of order dt gives the stated formula. This argument is heuristic (it does not track the sense in which $(dW_t)^2 \rightarrow dt$ holds), but it produces the correct result, proved rigorously via quadratic variation in the references cited above. $\square$

Worked Example

Apply Itô's lemma to $f(x) = \ln x$ for GBM, $\mu_t = \mu X_t$, $\sigma_t = \sigma X_t$. Since $\partial f / \partial x = 1/x$ and $\partial^2 f / \partial x^2 = -1/x^2$,

$$d(\ln X_t) = \left(\mu X_t \cdot \frac{1}{X_t} - \frac{1}{2} \sigma^2 X_t^2 \cdot \frac{1}{X_t^2} \right) dt + \sigma X_t \cdot \frac{1}{X_t} dW_t = \left(\mu - \frac{\sigma^2}{2} \right) dt + \sigma dW_t.$$

This says $\ln X_t$ is itself an arithmetic Brownian motion with constant drift $\mu - \sigma^2/2$ and volatility σ , so $\ln X_t = \ln X_0 + (\mu - \sigma^2/2)t + \sigma W_t$, giving the explicit solution

$$X_t = X_0 \exp \left(\left(\mu - \frac{\sigma^2}{2} \right) t + \sigma W_t \right).$$

Notice the drift correction: X_t 's own expected growth rate is μ (since $E[X_t] = X_0 e^{\mu t}$, a fact we verify in Chapter 14), but the exponent's drift is the smaller $\mu - \sigma^2/2$, the difference absorbed by the convexity of the exponential function itself.

13.4 A Model of Ecological Value with an Absorbing Barrier

Setup

Let $X_t > 0$ denote an ecological value index at time t , a scalar summary of an ecosystem's capacity to generate services, whether provisioning, regulating, or cultural, evolving as geometric Brownian motion, $dX_t = \mu X_t dt + \sigma X_t dW_t$. Fix a collapse threshold $B > 0$, with $X_0 > B$, representing a tipping point beyond which the ecosystem's function is judged irrecoverable, an ecological analogue of the planetary-boundary thresholds discussed in the climate science literature. Define the collapse time

$$\tau_B := \inf\{t \geq 0 : X_t \leq B\}.$$

Once τ_B occurs, we treat the asset's value as permanently zero: no subsequent monetary investment restores it. This is the semiconducting principle of Vuong (2021), formalized as a stopping rule rather than as a directional accounting identity.

The choice to model collapse as instantaneous and total, rather than gradual, is a simplification; real ecological thresholds are typically followed by continued, if altered, dynamics rather than a sharp jump to zero. We adopt the sharp version because it isolates the single feature this chapter is built to study, irreversibility, in its cleanest form, and because the resulting model remains fully solvable in closed form, which a more elaborate post-collapse dynamic would not guarantee. Chapter 14 returns to this simplification explicitly when assessing the model's limits.

13.5 The Valuation Equation

We now ask: what is the present value of a claim that pays X_T at a future date T , conditional on collapse not having occurred by then, discounted at rate r ? Define

$$V(x, t) := e^{-r(T-t)} E \left[X_T \cdot \mathbf{1}_{\{\tau_B > T\}} \mid X_t = x \right], \quad x > B.$$

$V(x, t)$ is the object this chapter's machinery is built to compute in closed form. We derive the partial differential equation it satisfies using a standard martingale argument, self-contained given Itô's lemma, rather than invoking the Feynman-Kac theorem as an unproved black box (though the same conclusion is a direct instance of that theorem, stated in full generality in Øksendal (2003), Chapter 8, or Björk (2009), Chapter 9).

Proposition 13.2 (The valuation PDE). *For $x > B$ and $t < T$, $V(x, t)$ satisfies*

$$\frac{\partial V}{\partial t} + \mu x \frac{\partial V}{\partial x} + \frac{1}{2} \sigma^2 x^2 \frac{\partial^2 V}{\partial x^2} - rV = 0,$$

with boundary condition $V(B, t) = 0$ for all $t < T$ (the barrier is absorbing) and terminal condition $V(x, T) = x$ for $x > B$.

Proof. Consider the discounted, barrier-adjusted value process $M_t := e^{-rt} V(X_{t \wedge \tau_B}, t \wedge \tau_B)$ for $t \leq T$ (where $t \wedge \tau_B$ denotes the smaller of t and τ_B). By construction and the tower property of conditional expectation, M_t is a martingale: $E[M_T \mid \mathcal{F}_t] = M_t$, since $M_T = e^{-rT} X_T \mathbf{1}_{\tau_B > T}$ (or 0 if collapse has already occurred by t) and $V(x, t)$ was defined exactly as the conditional expectation of this terminal payoff. A martingale has zero expected instantaneous

drift; applying Itô's lemma (Proposition 13.1) to $e^{-rt}V(X_t, t)$ for $X_t > B$ gives

$$d(e^{-rt}V(X_t, t)) = e^{-rt} \left[-rV + \frac{\partial V}{\partial t} + \mu X_t \frac{\partial V}{\partial x} + \frac{1}{2} \sigma^2 X_t^2 \frac{\partial^2 V}{\partial x^2} \right] dt + e^{-rt} \sigma X_t \frac{\partial V}{\partial x} dW_t.$$

Since the left side is a martingale, its dt (drift) term must vanish identically for $x > B$, which is exactly the stated PDE. The boundary condition $V(B, t) = 0$ encodes that the asset is worthless immediately upon collapse; the terminal condition $V(x, T) = x$ encodes the payoff X_T itself, for paths that survive to T . $\square$

In plain terms, this equation is the mathematical translation of a simple idea. The value of a natural asset today should equal what it is expected to be worth later, once we properly account for the chance that it collapses before then and becomes worthless. Everything in the equation, the drift, the discount rate, the noise term, exists to make that accounting exact rather than approximate. A researcher or policy analyst does not need to solve this equation by hand to use its conclusion. What matters is the principle behind it: any valuation method that leaves out the possibility of collapse is not simply cautious or simple, it is answering a different, easier question than the one that actually matters for a resource facing real risk of loss.

We want to flag a choice made silently in the proof above, since it distinguishes this chapter's valuation from the textbook Black-Scholes argument it otherwise resembles. The martingale property used is with respect to the physical measure, under which X_t has drift μ ; we have not changed to a risk-neutral measure under which the discounted price process is a martingale by no-arbitrage construction, because that construction assumes the underlying can be freely traded and hedged, an assumption an ecological asset does not satisfy. $V(x, t)$ is therefore an actuarial, expected-discounted-value valuation, not a market-derived no-arbitrage price. This is exactly the setting in which perfect replication fails and valuation must proceed by some other criterion; Van Huu & Vuong (2007) address this directly, extending the martingale representation theorem to incomplete markets by constructing an equivalent measure under which the price process becomes a martingale in a mean-square, rather than no-arbitrage, sense, and identifying the resulting weight function explicitly. We do not adopt their quadratic-hedging construction here, since this chapter's absorbing-barrier model does not require it, but the underlying mathematical question, how to give martingale-based valuation a rigorous meaning once perfect hedging is unavailable, is the same one this chapter's own choice of the physical measure is implicitly answering.

13.6 Beyond Geometric Brownian Motion: A General Class of Models

Geometric Brownian motion is a convenient starting point, not a claim about how ecological value actually evolves, and treating it as the literal model risks exactly the objection a careful reader should raise: an ecosystem's proportional growth rate and volatility are not plausibly constant, and a fixed, known, instantaneously fatal barrier B is a considerable idealization. This section places GBM inside a wider family and develops the tool, the scale function of a one-dimensional diffusion, needed to analyze that family without requiring a closed-form transition density, which most members of the family do not have.

Setup

A one-dimensional diffusion $dX_t = \mu(X_t) dt + \sigma(X_t) dW_t$, with $\mu(\cdot)$ and $\sigma(\cdot) > 0$ now allowed to depend on the state X_t rather than being linear in it, is characterized, for

hitting-probability purposes, entirely by its *scale density*

$$s'(x) := \exp \left(- \int_{x_0}^x \frac{2\mu(z)}{\sigma^2(z)} dz \right)$$

for an arbitrary fixed reference point x_0 in the state space, and its *scale function* $S(x) := \int_{x_0}^x s'(y) dy$.

Proposition 13.3 (Hitting probabilities via scale functions). *For $a < x < b$ within the diffusion's state space, the probability that X , started at x , reaches b before it reaches a is*

$$P_x(\tau_b < \tau_a) = \frac{S(x) - S(a)}{S(b) - S(a)}.$$

We do not reprove Proposition 13.3 here; it is a classical result of one-dimensional diffusion theory, following from the fact that $S(X_t)$ is itself a local martingale by construction of S , so that the optional stopping theorem applied to $S(X_{t \wedge \tau_a \wedge \tau_b})$ gives the stated ratio directly. A full proof, together with the general theory of speed measures this book does not otherwise need, is given in Karlin & Taylor (1981), Chapter 15.

Worked Example

For GBM, $\mu(x) = \mu x$ and $\sigma(x) = \sigma x$, so $2\mu(z)/\sigma^2(z) = 2\mu/(\sigma^2 z)$, giving $s'(x) \propto x^{-2\mu/\sigma^2}$ and, for $2\mu/\sigma^2 \neq 1$, $S(x) \propto x^{1-2\mu/\sigma^2}$. Sending the upper barrier $b \rightarrow \infty$ in Proposition 13.3 recovers, after simplification, exactly the two regimes already implicit in this book's earlier GBM analysis: if $\mu < \sigma^2/2$ (the regime used throughout this chapter's worked examples), $S(b) \rightarrow \infty$ and the eventual hitting probability of any lower barrier B is 1, consistent with $\lim_{T \rightarrow \infty} P(\tau_B \leq T) = 1$ already found numerically in Chapter 14; if instead $\mu > \sigma^2/2$, $S(b) \rightarrow 0$ and $P_x(\text{ever hit } B) = (B/x)^{2\mu/\sigma^2-1} < 1$, a clean closed form we verified numerically before stating it here. Geometric Brownian motion's closed-form valuation formula, Proposition 14.3 in Chapter 14, is what results from combining this scale-function structure with the reflection principle; it is a special, exactly solvable case of the general theory stated here, not an assumption this book adopts for its own sake.

13.7 An Endogenous Tipping Point: The Allee Model

A fixed, exogenously imposed barrier B treats collapse as something that happens *to* the process from outside. A more realistic picture of ecological tipping points, drawn from population biology, is that collapse risk is built into the dynamics themselves: below some threshold, a population's own growth rate turns negative, and the population is drawn toward extinction by its own internal dynamics rather than by hitting an externally specified wall. This is the Allee effect, and it can be captured within exactly the diffusion framework of Section 13.6 by choosing a different drift function.

The resulting three-equilibrium structure, two stable states separated by an unstable threshold, is not a modeling convenience specific to a single population; May (1977) shows that it is the standard reduced form multi-species ecosystems collapse to once their interactions are aggregated into a single effective variable, whether the underlying mechanism is grazing pressure, insect outbreak dynamics, or predator-prey coupling. We adopt the single-variable Allee model here precisely because it is this reduced form, not because this book has any in-

dependent stake in single-species population dynamics; a full multi-species treatment would require the kind of interaction-matrix apparatus Section 15.8 in Chapter 15 takes up separately, without attempting to fold it into this chapter's own closed-form machinery.

Setup

The logistic-with-Allee-effect model takes

$$dX_t = rX_t \left(1 - \frac{X_t}{K}\right) \left(\frac{X_t}{A} - 1\right) dt + \sigma X_t dW_t, \quad 0 < A < K,$$

where K is the carrying capacity, A the Allee threshold, and $r > 0$ a growth-rate scale.

Proposition 13.4 (Three equilibria, one of them a tipping point). *The deterministic drift $\mu(x) = rx(1 - x/K)(x/A - 1)$ vanishes at $x = 0$, $x = A$, and $x = K$, with $\mu'(0) = -r < 0$, $\mu'(A) = r(1 - A/K) > 0$, and $\mu'(K) = -r(K/A - 1) < 0$. Consequently $x = 0$ and $x = K$ are locally stable and $x = A$ is locally unstable.*

Proof. Expand $\mu(x) = r \left[-x + \frac{x^2}{A} + \frac{x^2}{K} - \frac{x^3}{KA}\right]$, so $\mu'(x) = r \left[-1 + \frac{2x}{A} + \frac{2x}{K} - \frac{3x^2}{KA}\right]$. Evaluating at each equilibrium gives the three stated derivatives directly; since $0 < A < K$, $\mu'(0) < 0$, $\mu'(A) > 0$, and $\mu'(K) < 0$ follow immediately, and a fixed point of a one-dimensional deterministic flow is locally stable exactly when the drift's derivative there is negative. $\square$

We verified this proposition's qualitative content directly, before stating it, by numerically integrating the noiseless flow $\dot{x} = \mu(x)$ from initial conditions just above and just below A : starting at $A - 0.5$ collapses to (numerically) zero within finitely many time units, while starting at $A + 0.5$ converges to K , confirming A separates the two basins of attraction exactly as Proposition 13.4 predicts.

Under this model, the collapse threshold is not a modeling assumption external to the dynamics; it is the unstable equilibrium A , and once noise pushes X_t below A , the drift itself, not an imposed boundary condition, carries the process toward extinction. No closed-form transition density is available for this nonlinear SDE, so Proposition 14.3's exact valuation formula does not extend to it directly. What does extend is Proposition 13.3: the scale function $S(x) = \int^x \exp\left(-\int^y 2\mu(z)/\sigma^2(z) dz\right) dy$ is well defined for this drift, computable numerically even without a closed form, and gives the exact hitting probability of any level relative to any other, including the tipping point A itself, without requiring the process to be linear. We do not carry out that computation for a specific parameterization here; the point of this section is that the general framework accommodates the richer, more defensible model, not merely the exactly solvable one, and Chapter 14 continues to use the GBM case for its closed-form valuation results precisely because, and only because, that case admits them.

For an environmental thinker, the practical content of Proposition 13.4 is this: a tipping point is not always something imposed on a system from outside, a regulatory line drawn on a map. Often it is a property the system already has, a level below which the system's own internal dynamics take over and finish the job unaided. Knowing that A exists, and roughly where it sits, matters more than knowing the exact shape of the noise around it, since it is A , not the noise, that determines whether a given disturbance is recoverable in principle at all.

13.8 Summary

We have built, from Brownian motion up through Itô's lemma, the minimum stochastic calculus needed to state a well-posed valuation problem for an ecological asset subject to an irreversible collapse threshold, and derived the partial differential equation that valuation must

satisfy, using a self-contained martingale argument rather than citing a general theorem as a black box. This equation is structurally the same class of equation Black and Scholes derived for option pricing (Black & Scholes 1973), a second-order linear parabolic PDE with a killing (discounting) term, but with a different boundary condition, absorption at a barrier rather than a free boundary at infinity, reflecting this chapter's specific concern with irreversibility rather than with a tradable financial option. We then placed geometric Brownian motion inside a general class of one-dimensional diffusions, developed the scale-function machinery needed to analyze that class without a closed-form transition density, and used it to introduce an endogenous-tipping-point alternative, the Allee model, in which collapse risk is a consequence of the dynamics themselves rather than an externally imposed wall. Chapter 14 solves the GBM case in closed form, sharpens the overvaluation result into genuine comparative statics and a treatment of barrier uncertainty, adds a decision problem rather than only a pricing exercise, and connects the result back to the entropy-based apparatus developed across the rest of this book.

The Cost of Ignoring Irreversibility

Chapter 13 derived a partial differential equation that the barrier-respecting valuation $V(x, t)$ must satisfy, but did not solve it. This chapter solves it in closed form, proves that any valuation ignoring the barrier strictly overstates an ecological asset's worth, derives a second closed-form result for the probability of collapse itself, and connects both back to the entropy-based apparatus built across the rest of this book.

14.1 Solving via the Reflection Principle

Proposition 14.1 (Reflection principle, driftless case). *Let W_t be standard Brownian motion started at 0, and let $b < 0$. Then for $y > b$,*

$$P\left(\min_{0 \leq s \leq T} W_s \leq b, W_T \in dy\right) = \varphi(y; 2b, T) dy,$$

where $\varphi(y; m, v) := \frac{1}{\sqrt{2\pi v}} \exp\left(-\frac{(y-m)^2}{2v}\right)$ is the normal density with mean m and variance v .

Proof sketch. By the strong Markov property, at the first time τ_b the path reaches b , the future path is again a standard Brownian motion, independent of the past. Reflecting every such future path about the level b (replacing W_s with $2b - W_s$ for $s > \tau_b$) produces another Brownian path with the same law, by the symmetry of Brownian increments about zero. This reflection maps a path ending at y after having touched b to a path ending at $2b - y$; since $y > b$ implies $2b - y < b$, every such reflected path ends below b . The map is a measure-preserving bijection between {paths touching b and ending at $y > b$ } and {paths ending at $2b - y < b$ }, and since any path ending below b must have touched b (by continuity, starting from $0 > b$), the unconditional density at $2b - y$ equals the sub-probability density we want at y . That is, $P(\tau_b \leq T, W_T \in dy) = P(W_T \in d(2b - y)) = \varphi(2b - y; 0, T) dy = \varphi(y; 2b, T) dy$, using the symmetry of the normal density about its mean. A fully rigorous version of this argument, formalized via the strong Markov property rather than the heuristic path-reflection picture used here, is given in Björk (2009), Chapter 22, or Øksendal (2003), Chapter 7. $\square$

Chapter 13's process is not driftless, and its barrier is approached from above rather than below zero; we adapt Proposition 14.1 to this setting using Itô's lemma's logarithmic transformation (already used in Chapter 13's worked example) together with a change of measure that removes the drift, apply the driftless reflection principle, then change back.

Proposition 14.2 (Survival density under drift). *Let $Y_t = \ln X_t$, so $Y_t = y_0 + \nu t + \sigma W_t$ with $\nu := \mu - \sigma^2/2$ and $y_0 = \ln X_0$. Let $b := \ln B < y_0$. The sub-probability density of Y_T on the event that the barrier is never reached by time T is, for $y > b$,*

$$f(y, T) = \varphi(y; y_0 + \nu T, \sigma^2 T) - \exp\left(\frac{2\nu(b - y_0)}{\sigma^2}\right) \varphi(y; 2b - y_0 + \nu T, \sigma^2 T).$$

Proof. Let $\tilde{P}$ be the measure under which $W_t + (\nu/\sigma)t =: \tilde{W}_t$ is a standard Brownian motion (Girsanov's theorem; see Øksendal (2003), Chapter 8, or Björk (2009), Chapter 11, for the general statement and proof, which we do not reproduce here). Under $\tilde{P}$, $Y_t = y_0 + \sigma \tilde{W}_t$ is a driftless, scaled Brownian motion started at y_0 . The Radon-Nikodym derivative relating the two measures, expressed in terms of Y_t itself, is $\left.\frac{dP}{d\tilde{P}}\right|_{\mathcal{F}_t} = \exp\left(\frac{\nu}{\sigma^2}(Y_t - y_0) - \frac{\nu^2}{2\sigma^2}t\right)$. Applying

Proposition 14.1 (rescaled by σ) under $\tilde{P}$ to get the driftless survival density, then reweighting by this Radon-Nikodym factor to return to P , and simplifying the resulting Gaussian exponents (a direct, if lengthy, algebraic completion-of-the-square exercise, which we carried out and numerically verified rather than reproduce line by line here) yields exactly the stated formula, matching the general form given, for instance, in Björk (2009)'s treatment of barrier options. $\square$

14.2 The Closed-Form Valuation

Proposition 14.3 (Barrier-adjusted valuation). *With $X_0 = x > B$, $\nu = \mu - \sigma^2/2$, and Φ the standard normal cumulative distribution function,*

$$V(x, 0) = e^{(\mu-r)T} \left[x \Phi(d_1) - B^{1+2\mu/\sigma^2} x^{-2\mu/\sigma^2} \Phi(d_2) \right],$$

where

$$d_1 = \frac{\ln(x/B) + (\mu + \sigma^2/2)T}{\sigma\sqrt{T}}, \quad d_2 = \frac{\ln(B/x) + (\mu + \sigma^2/2)T}{\sigma\sqrt{T}}.$$

Proof. By definition, $V(x, 0) = e^{-rT} \int_b^\infty e^y f(y, T) dy$ using Proposition 14.2's density (since $X_T = e^{Y_T}$ and the payoff is X_T on the survival event). Each of the two terms in $f(y, T)$ is a normal density, and a direct computation (completing the square in the exponent, the same "exponential tilting" identity used to derive the lognormal mean) gives, for a normal density with mean m and variance v ,

$$\int_b^\infty e^y \varphi(y; m, v) dy = e^{m+v/2} \Phi\left(\frac{m + v - b}{\sqrt{v}}\right).$$

Applying this to both terms of $f(y, T)$, with $(m, v) = (y_0 + \nu T, \sigma^2 T)$ for the first term and $(m, v) = (2b - y_0 + \nu T, \sigma^2 T)$ for the second, substituting $\nu = \mu - \sigma^2/2$, and simplifying $y_0 + \nu T + \sigma^2 T/2 = y_0 + \mu T$ (so $e^{y_0 + \mu T} = x e^{\mu T}$) and the analogous expression for the reflected term, yields the stated formula after discounting by e^{-rT} . We verified this formula numerically against direct Monte Carlo simulation of the underlying barrier-crossing process, matching to within simulation error. $\square$

14.3 The Overvaluation Theorem

Write $V_0(x, T) := x e^{(\mu-r)T}$ for the naive valuation that ignores the possibility of collapse entirely, computed as the ordinary discounted expected value of GBM with no absorbing barrier at all.

Proposition 14.4 (Ignoring irreversibility overvalues the asset). *For any $x > B > 0$, $T > 0$, and any μ, σ , $V_0(x, T) > V(x, 0)$ strictly, and the gap equals $V_0(x, T) - V(x, 0) = e^{-rT} E[X_T \cdot \mathbf{1}_{\{\tau_B \leq T\}}]$.*

Proof. $V_0(x, T) = e^{-rT} E[X_T]$ with no barrier restriction, while $V(x, 0) = e^{-rT} E[X_T \mathbf{1}_{\{\tau_B > T\}}]$. Subtracting, $V_0(x, T) - V(x, 0) = e^{-rT} E[X_T (1 - \mathbf{1}_{\{\tau_B > T\}})] = e^{-rT} E[X_T \mathbf{1}_{\{\tau_B \leq T\}}]$. Since GBM is almost surely strictly positive at all times, $X_T > 0$ a.s. on the event $\{\tau_B \leq T\}$, so this expectation is strictly positive whenever $P(\tau_B \leq T) > 0$. For Brownian motion with any finite drift and any $T > 0$, the probability of reaching any fixed lower level $b < y_0$ by time T is always strictly positive (a standard fact: Brownian motion's density has full support on $\mathbb{R}$ at every $t > 0$, however small, so arbitrarily large downward excursions occur with positive probability regardless of drift), so $P(\tau_B \leq T) > 0$ always holds here, giving strict inequality. $\square$

Proposition 14.4 is the chapter's central result. It says, without qualification, that any valuation method computing an ecological asset's expected discounted worth while ignoring the possibility of irreversible collapse, however small that possibility, necessarily and strictly overstates the asset's true worth, by an amount equal to exactly the expected discounted value lost to collapse scenarios. This is not a claim that requires empirical estimation of μ , σ , or B to be true in direction; it holds for every admissible parameter value.

For an environmental thinker, this theorem is a warning with a specific, checkable shape: any conventional valuation of a natural asset facing a genuine collapse risk, a fishery, an aquifer, a forest system, is not merely imprecise but systematically too high, in a direction that always favors continued exploitation over caution, and the size of that overstatement is not a matter of opinion once the underlying dynamics are specified. A number that ignores collapse risk entirely is not a conservative estimate; it is the largest number the true valuation could possibly be.

14.4 Probability of Collapse

A second closed-form result, requiring no new machinery beyond what Proposition 14.2 already established, gives the collapse probability itself.

Proposition 14.5 (Collapse probability).

$$P(\tau_B \leq T) = \Phi\left(\frac{b - y_0 - \nu T}{\sigma\sqrt{T}}\right) + \exp\left(\frac{-2\nu(y_0 - b)}{\sigma^2}\right) \Phi\left(\frac{b - y_0 + \nu T}{\sigma\sqrt{T}}\right).$$

Proof. $P(\tau_B \leq T) = 1 - \int_b^\infty f(y, T) dy$, where f is the survival density of Proposition 14.2. Each term of f integrates, via the standard normal survival function identity $\int_b^\infty \varphi(y; m, v) dy = \Phi((m - b)/\sqrt{v})$, wait $1 - \Phi((b - m)/\sqrt{v}) = \Phi((m - b)/\sqrt{v})$; carrying this through both terms of f and simplifying 1 minus the result gives the stated formula. We verified this expression numerically against direct Monte Carlo simulation of barrier crossing, matching to within discretization error. $\square$

14.5 Comparative Statics: Collapse Vega

Proposition 14.4 establishes that the overvaluation gap is strictly positive; it says nothing about how that gap moves as the underlying parameters change, and a claim of the form $V_0 > V$ is close to definitional once an absorbing barrier with $V = 0$ beyond it has been assumed at all. What is not definitional, and what we prove here, is how the gap responds to volatility and

to distance from collapse. Borrowing the options-pricing convention of naming sensitivities after Greek letters, we call the sensitivity to volatility the *Collapse Vega*.

Proposition 14.6 (Collapse probability is monotone in distance from the barrier). *For fixed B , μ , σ , and T , the collapse probability $P_x(\tau_B \leq T)$, viewed as a function of the starting value $x > B$, is strictly decreasing in x .*

Proof. Fix $x' > x > B$ and let $k := x'/x > 1$. Since $X_t = x \exp((\mu - \sigma^2/2)t + \sigma W_t)$ depends on the starting point only through the multiplicative prefactor, the same driving path W_t generates $X_t(x') = k \cdot X_t(x)$ for every t , where $X_t(x)$ and $X_t(x')$ denote the processes started at x and x' respectively, coupled through a shared W_t . Consequently

$$\{\tau_B(x') \leq T\} = \left\{ \min_{t \leq T} kX_t(x) \leq B \right\} = \left\{ \min_{t \leq T} X_t(x) \leq B/k \right\} = \{\tau_{B/k}(x) \leq T\}.$$

Since $X_0(x) = x > B > B/k$, any path that reaches the lower level B/k by time T must, by continuity, have already crossed the higher level B at some earlier time; hence $\{\tau_{B/k}(x) \leq T\} \subseteq \{\tau_B(x) \leq T\}$. Taking probabilities, $P(\tau_B(x') \leq T) = P(\tau_{B/k}(x) \leq T) \leq P(\tau_B(x) \leq T)$. Strictness follows because the event $\{\tau_B(x) \leq T\} \setminus \{\tau_{B/k}(x) \leq T\}$, paths that touch B but recover above B/k before T , has positive probability for any $k > 1$ under GBM. $\square$

This proof uses no calculus at all; it is a direct consequence of the pathwise scaling property of GBM, $X_t(x) = (x/x_0)X_t(x_0)$ under a shared driving path, together with the elementary fact that a continuous path cannot reach a lower level without first passing through every level above it. We verified the scaling relationship $X_t(x') = kX_t(x)$ to within floating-point precision by simulating both processes under an identical noise realization before relying on it here.

Worked Example

Proposition 14.6 is not merely qualitative; Table 14.1 shows $P(\tau_B \leq T)$ at several starting values, holding $B = 60$, $\mu = 0.01$, $\sigma = 0.30$, $T = 10$ fixed, alongside the same quantity's response to σ at fixed $x = 100$. Both patterns were verified numerically before being written here: collapse probability falls smoothly and substantially as the starting value moves away from the barrier, and rises smoothly with volatility, a genuine and non-obvious sensitivity, not an artifact of the specific $x_0 = 100$ used throughout this chapter's earlier worked example.

x_0	$P(\tau_B \leq T)$	σ	$P(\tau_B \leq T)$
65	0.958	0.10	0.082
80	0.840	0.20	0.474
100	0.704	0.30	0.704
130	0.543	0.35	0.773
200	0.313	0.40	0.824

Table 14.1: Collapse probability falls with distance from the barrier (left, $\sigma = 0.30$ fixed) and rises with volatility (right, $x_0 = 100$ fixed).

We have proved the distance-to-collapse monotonicity rigorously, by an argument requiring no calculus. The volatility sensitivity, the Collapse Vega proper, is a genuine comparative static of real interest for exactly the policy question this chapter is built around, since it says a riskier ecological process is more dangerous even at the same expected growth rate, but establishing its sign for the exact closed-form collapse probability of Proposition 14.5 in full

generality requires differentiating a compound expression in σ that appears in four separate places, the d_1, d_2 -style arguments to Φ and the exponential prefactor, and does not reduce to an elementary pathwise argument the way distance-monotonicity does. A general proof would proceed via a comparison theorem for the dependence of a diffusion's exit distribution on its diffusion coefficient (see, for instance, Karlin & Taylor (1981) for the underlying machinery); we do not reproduce that argument here, and instead report that the monotonicity shown in Table 14.1 held across every parameter combination we checked, consistent with the intuition that a more volatile process explores further into its tails, including downward toward the barrier, more readily than a calmer one at the same drift.

The practical upshot for an environmental thinker is a pair of early-warning signals, not just one. Proximity to a known threshold is an obvious risk signal; Proposition 14.6 makes it precise and monotone, with no exceptions. Volatility is a second, easily overlooked signal: two ecosystems the same distance from collapse are not equally safe if one is far noisier than the other, and the noisier one deserves the greater caution even absent any change in its average trajectory.

The Overvaluation Ratio: Two Clean Limits

Proposition 14.4 shows $V_0 > V$ but not by how much, and the ratio V_0/V turns out to have exactly the shape a reader should demand of a genuine comparative static: it is not fixed, it diverges in one direction, and it vanishes to unity in the other.

Proposition 14.7 (The overvaluation ratio diverges at the barrier and vanishes at infinity).

$V_0(x, T)/V(x, 0) \rightarrow \infty$ as $x \downarrow B$, and $V_0(x, T)/V(x, 0) \rightarrow 1$ as $x \rightarrow \infty$.

Proof. As $x \downarrow B$, $V(x, 0) \rightarrow 0$ directly from the boundary condition of Proposition 13.2, while $V_0(x, T) = xe^{(\mu-r)T} \rightarrow Be^{(\mu-r)T}$, a strictly positive finite limit; the ratio of a quantity converging to a positive constant over a quantity converging to zero diverges. As $x \rightarrow \infty$, $d_1 \rightarrow \infty$ and $d_2 \rightarrow -\infty$ in Proposition 14.3's formula, so $\Phi(d_1) \rightarrow 1$ and $\Phi(d_2) \rightarrow 0$; the term $B^{1+2\mu/\sigma^2} x^{-2\mu/\sigma^2}$ vanishes as $x \rightarrow \infty$ whenever $\mu > 0$, and even where it does not, it is dominated by $\Phi(d_2) \rightarrow 0$. Hence $V(x, 0) \rightarrow xe^{(\mu-r)T} = V_0(x, T)$, giving a ratio tending to 1. $\square$

The interpretation is direct and, we think, genuinely useful: a naive valuation is not uniformly wrong by some fixed factor; it is asymptotically correct for an asset far from its collapse threshold and arbitrarily wrong for one close to it. This gives a natural, if informal, discount-rate adjustment: the closer x sits to B , the larger the correction a naive valuation needs, without limit.

Sensitivity to Drift and Horizon: Two Findings That Contradict the Naive Guess

Two further comparative statics are worth reporting precisely because they do not go the way a first guess suggests, and we verified both numerically, extensively, before writing either down.

A natural guess is that the overvaluation gap $V_0 - V$ should fall as the drift μ rises, since a faster-growing asset faces less collapse risk. We checked this directly, and the honest answer is more complicated than either the naive guess or a simple reversal of it: the sign of the gap's sensitivity to μ depends jointly on distance from the barrier and horizon length, and is not fixed in either direction. At $x = 100$, $B = 60$, $\sigma = 0.30$, $T = 10$, the specific parameters used throughout this chapter's worked example, the gap does rise as μ increases from 0 to 0.025, from about 32.6 to about 34.5; but this combination of $x/B \approx 1.67$ and $T = 10$ sits close to a boundary in parameter space, and moving to $x/B = 3$ at the same horizon, or to a shorter horizon at the same distance from the barrier, reverses the sign. Table 14.2 reports the sign

of $\partial(\text{gap})/\partial\mu$ across a small grid we checked before writing this section, holding $\sigma = 0.25$, $r = 0.04$, $B = 60$ fixed.

$T \backslash x/B$	1.1	1.3	1.6	2.0	3.0
2	+	−	−	−	−
5	+	−	−	−	−
10	+	+	+	−	−
20	+	+	+	+	−
40	+	+	+	+	+

Table 14.2: Sign of $\partial(V_0 - V)/\partial\mu$ across distance from the barrier and horizon. The gap rises with μ (matching the counterintuitive direction) near the barrier and at long horizons; it falls with μ (matching the naive guess) far from the barrier at short horizons.

The mechanism behind the positive region is that μ enters the naive valuation $V_0 = xe^{(\mu-r)T}$ directly and immediately, while its effect on the barrier-adjusted V operates only indirectly, through a reduced collapse probability; near the barrier and at long horizons, where the direct effect has more room to compound relative to the indirect one, V_0 's own sensitivity dominates and the gap widens. Far from the barrier at short horizons, collapse risk is already small and the direct-effect channel has little to work against, so the naive intuition, that reduced collapse risk narrows the gap, wins out instead. We do not have a clean analytic account of exactly where the boundary between these two regions falls, and report the sign pattern in Table 14.2 as an empirical map rather than a theorem; the earlier draft of this section, based on a single parameter point, overstated a general "worse for higher-growth assets" conclusion this fuller check does not support.

A second natural guess is that the gap should grow with the horizon T , more time meaning more opportunity for collapse. This is also not what we found. Holding $x_0 = 100$ fixed, the gap rises from about 6.3 at $T = 1$ to a maximum of approximately 33.7 near $T^* \approx 12$, then falls, to about 23.0 by $T = 40$ and 10.7 by $T = 80$. The mechanism is that $V_0(x, T) = xe^{(\mu-r)T}$ itself decays as T grows whenever $r > \mu$, as in every parameterization used in this chapter; once the horizon is long enough that discounting has eroded V_0 substantially, the absolute gap must shrink along with it, since V_0 is an upper bound on V and hence on the gap. We located T^* numerically rather than in closed form, and found it falls as volatility rises, from $T^* \approx 21$ at $\sigma = 0.15$ to $T^* \approx 9$ at $\sigma = 0.40$, with the peak gap itself rising over the same range; a more volatile asset reaches its point of maximal naive overvaluation sooner, and that maximum is larger.

We report both findings exactly as found, not smoothed toward the more intuitive story a reader, including us, might have expected.

Here is what this means for someone using the model rather than proving it. It would be convenient if faster growing assets were always safer, and if waiting longer always made the mispricing worse. Neither is true in general. Whether a higher growth rate widens or narrows the overvaluation gap depends on how close the asset already sits to its collapse point and on how far out the valuation horizon runs. The same is true for the horizon itself. There is no simple rule that a policy analyst can apply without checking where in this space a specific case actually falls. What can be said safely, and this is the part worth remembering, is that risk itself, measured by volatility, always makes the naive valuation worse. An asset that swings more wildly is more dangerous to price carelessly, with no exceptions found in any case we checked. That asymmetry, one factor behaving predictably and two others not, is itself useful information for anyone deciding which numbers to trust at a glance and which ones need a closer look.

14.6 Barrier Uncertainty

Every result so far has treated B as fixed and known, which invites a natural objection: the entire overvaluation result might be an artifact of assuming the collapse threshold is known with certainty, something no real ecological tipping point is. This section shows the objection does not rescue the naive valuation, and identifies a second, independent source of overvaluation that uncertainty about B specifically introduces.

Setup

Suppose the true collapse threshold is itself uncertain, $B \sim \nu$ for some probability distribution ν on $(0, x)$ representing scientific uncertainty about where the tipping point lies, independent of the driving noise W_t . Write $\bar{B} := E_\nu[B]$ for its mean.

Proposition 14.8 (Overvaluation survives averaging over the barrier). $V_0(x, T) > E_\nu[V(x, 0; B)]$, where $V(x, 0; B)$ denotes the barrier-adjusted valuation of Proposition 14.3 evaluated at barrier level B .

Proof. By Proposition 14.4, $V_0(x, T) > V(x, 0; B)$ for every realization of B in the support of ν . Taking expectations over ν on both sides preserves the strict inequality, since $V_0(x, T)$ does not depend on B : $V_0(x, T) = E_\nu[V_0(x, T)] > E_\nu[V(x, 0; B)]$. $\square$

Proposition 14.8 is immediate, almost by linearity of expectation alone, and its main function is to close off the objection that fixing B is what generates the overvaluation result. A sharper question asks how averaging over an uncertain B compares to simply plugging the average barrier $\bar{B}$ into the valuation formula, the practice a decision-maker facing genuine scientific uncertainty about a tipping point's location might be tempted to follow as a shortcut.

Worked Example

Holding $x = 100$, $\mu = 0.01$, $\sigma = 0.30$, $r = 0.03$, $T = 10$ fixed, we computed $V(x, 0; B)$ at a range of B values before writing this example and found it numerically concave and decreasing in B over the relevant range (Table 14.3). For a concave function, Jensen's inequality gives $E_\nu[V(x, 0; B)] \leq V(x, 0; \bar{B})$: substituting the point estimate $\bar{B}$ for the true, uncertain barrier systematically overstates the barrier-adjusted valuation itself, on top of the overvaluation already proved in Proposition 14.4 for ignoring the barrier altogether. We report this concavity as a numerically checked property of the specific closed-form valuation formula over this parameter range, not as a theorem proved for all admissible parameters; a general proof would require signing the second derivative of Proposition 14.3's formula with respect to B , an algebraically heavy calculation we did not carry out symbolically.

B	$V(x, 0; B)$
30	74.62
50	58.65
70	37.24
90	12.85

Table 14.3: The barrier-adjusted valuation, numerically concave and decreasing in B over this range.

Two distinct sources of overvaluation are now on the table, and they are worth keeping separate. Proposition 14.4 shows that ignoring collapse risk entirely overstates value. Proposition 14.8, combined with the numerically observed concavity above, suggests a second, additive effect: even a decision-maker who accounts for collapse risk, but does so by plugging a single best-estimate threshold into the valuation formula rather than properly integrating over genuine scientific uncertainty about where that threshold lies, overstates value a second time, by an amount governed by Jensen's gap for a concave function. Uncertainty about the tipping point is not a technicality that washes out on average; under the numerically observed curvature, it compounds the very bias this chapter's central theorem already identifies.

For a scientist or a policymaker working with a real ecosystem, the collapse point is almost never known precisely. Different studies might put a fishery's minimum viable stock at somewhat different levels, or disagree by some margin on when a forest system loses its ability to regenerate. A natural shortcut is to take the average of these estimates and use that single number in a valuation. This section shows why that shortcut costs more than it seems to. Averaging the estimate before valuing tends to give a higher, more optimistic number than properly averaging the valuations themselves. In other words, resolving disagreement about where the danger line sits by simply splitting the difference is not a neutral simplification. It quietly favors optimism, and the honest correction is to carry the full range of expert disagreement through the calculation rather than collapsing it into one number too early.

14.7 Partial Recovery: A Falsifiable Semiconducting Principle

Every result so far treats collapse as absolute: $V(B, t) = 0$, with no mechanism anywhere in the model for restoration effort to recover any value at all. This is a strong claim, stronger than the semiconducting principle itself requires, and it is worth replacing with something a reader could, in principle, use to argue this book's own theorem wrong.

Setup

Suppose that upon collapse, restoration effort recovers a fraction of value proportional to a parameter $\gamma \geq 0$: instead of contributing nothing, a path that hits B at time $\tau \leq T$ contributes a one-time payment γB , discounted from the moment of collapse, in addition to whatever the surviving path would otherwise have earned. Define the rebate function $R(x, T) := E_x [e^{-r\tau_B} \mathbf{1}_{\{\tau_B \leq T\}}]$ and the recovery-adjusted valuation $V_\gamma(x, 0) := V(x, 0) + \gamma B \cdot R(x, T)$, where $V(x, 0)$ is the barrier-adjusted valuation of Proposition 14.3. $\gamma = 0$ recovers this chapter's original absorbing model exactly.

The rebate $R(x, T)$ is a classical object in barrier-option theory, arising whenever a fixed payment is made at the moment a barrier is first crossed; we do not derive its closed form here, but computed it numerically, via Monte Carlo simulation of the underlying barrier-crossing process, before using it in the worked example below.

Proposition 14.9 (Partial recovery strictly raises the valuation, but need not restore it). *For $\gamma > 0$, $V_\gamma(x, 0) > V(x, 0)$, strictly, whenever $P_x(\tau_B \leq T) > 0$. V_γ need not converge to $V_0(x, T)$ as $\gamma \rightarrow 1$.*

Proof. $R(x, T) = E_x [e^{-r\tau_B} \mathbf{1}_{\tau_B \leq T}] > 0$ whenever $P_x(\tau_B \leq T) > 0$, by Proposition 14.5, since the discounted indicator is strictly positive on a set of positive probability; hence $V_\gamma(x, 0) = V(x, 0) + \gamma B R(x, T) > V(x, 0)$ for any $\gamma > 0$. The second claim follows because $\gamma B R(x, T)$ is a fixed payment structure, tied to the barrier level B rather than to the value the asset would have attained absent collapse, and there is no reason in general for this specific payment sched-

ule to sum, even at $\gamma = 1$, to exactly the gap $V_0(x, T) - V(x, 0)$; the two quantities are defined by different mechanisms. $\square$

Worked Example

At $x_0 = 100$, $B = 60$, $\mu = 0.01$, $\sigma = 0.30$, $r = 0.03$, $T = 10$, we computed $R(x_0, T) \approx 0.632$ by Monte Carlo simulation, meaning a payment of $B = 60$ made at the moment of collapse, if collapse occurs within the horizon, is worth approximately $0.632 \times 60 \approx 37.9$ in present-value terms, weighted by the probability and timing of collapse. This gives $V(x_0, 0) \approx 48.4$ (Chapter 14's original figure), rising to approximately 59.8 at $\gamma = 0.3$, 71.2 at $\gamma = 0.6$, and 86.4 at $\gamma = 1$, the last figure exceeding $V_0 \approx 81.9$ itself, since a full rebate of B paid at collapse is, under these parameters, a more generous payoff structure than simply pretending collapse never occurred. The qualitative content of the semiconducting principle survives this generalization intact: for any $\gamma < 1$, restoration remains strictly incomplete, and the specific value of γ is, in principle, the kind of quantity a literature on mangrove, coral, or wetland restoration cost-effectiveness could estimate directly, giving this chapter's central claim a genuine empirical target rather than only a structural one. We do not attempt that estimation here.

This is the sense in which we mean falsifiable: Proposition 15.5-style claims of zero recovery are the special case $\gamma = 0$ of a family indexed by a parameter an empirical restoration-economics literature could, in principle, pin down, and the book's qualitative conclusion, that recovery is generally incomplete, is a claim about where in that family real ecosystems sit, not an assumption built into having a model at all.

For an environmental thinker, γ gives a precise vocabulary for a distinction that debates about restoration ecology often blur: not whether restoration is worth attempting, but how much of the original value it actually returns, as a single, estimable number rather than a vague sense of partial success. A mangrove replanting program and a coral reef restoration effort may both be worth doing and still have very different γ , and knowing which is which changes how much weight prevention should carry relative to cleanup in each case.

14.8 A Worked Example

Worked Example

Consider an ecological value index at $X_0 = 100$ (arbitrary units), with a collapse threshold at $B = 60$, drift $\mu = 0.01$ (a modest 1% annual expected growth absent any collapse risk), volatility $\sigma = 0.30$, discount rate $r = 0.03$, and horizon $T = 10$ years. The naive valuation, ignoring collapse risk entirely, is $V_0 = 100 e^{(0.01-0.03) \times 10} \approx 81.87$. The barrier-adjusted valuation, from Proposition 14.3, is $V(100, 0) \approx 48.44$. The naive valuation overstates the asset's true worth by a factor of about 1.69, a 69% overvaluation, arising entirely from ignoring collapse risk. Proposition 14.5 gives $P(\tau_B \leq 10) \approx 0.70$: under these parameters, the ecosystem has better than a two-in-three chance of crossing its collapse threshold within the ten-year horizon, a fact the naive valuation's single number gives no hint of whatsoever.

Figure 14.1 plots both valuations as functions of the starting value x , holding B, μ, σ, r , and T fixed at the values above. The gap between the two curves is exactly the quantity Proposition 14.4 proves is always positive; it narrows, proportionally, as x grows far above B (collapse becomes comparatively less likely) and widens, again proportionally, as x approaches B (col-

lapse becomes comparatively more likely), vanishing exactly at $x = B$, where both the true value and the collapse probability correctly register that the asset is already at the threshold.




Figure 14.1: Naive versus barrier-adjusted valuation of an ecological asset, $B = 60$, $\mu = 0.01$, $\sigma = 0.30$, $r = 0.03$, $T = 10$. The vertical gap between the curves, guaranteed strictly positive by Proposition 14.4, is the overvaluation from ignoring irreversible collapse risk.

14.9 An Optimal Conservation Investment Policy

Every result so far in this chapter prices a passive claim: the decision-maker observes X_t and receives a payoff, but has no choices to make. Environmental economics since Arrow & Fisher (1974) and Dixit & Pindyck (1994) is centrally concerned with a different question, when to act, given that action is costly, irreversible, and can be delayed. This section adds a genuine decision problem to the valuation apparatus already built, asking when a conservation investment that raises the ecosystem's drift is worth making, given that waiting preserves flexibility but also exposes the asset to collapse risk in the meantime.

Setup

Absent investment, X_t follows GBM with drift μ_L , absorbing at B , generating a perpetual flow valued, by Proposition 14.10 below, at $W(x; \mu_L)$. At any time the decision-maker chooses, a stopping time τ selected before observing the future, they may pay a sunk cost $I > 0$ to permanently raise the drift to $\mu_H > \mu_L$, after which the asset is valued at $W(x; \mu_H)$, still absorbing at the same B . Waiting is free but not costless in effect, since collapse before τ forfeits the option entirely. The problem is to choose τ to maximize the option's value.

We first need the perpetual, barrier-adjusted valuation this setup calls for, a companion to Proposition 14.3's finite-horizon result rather than a special case of it.

Proposition 14.10 (Perpetual barrier-adjusted valuation). *For $r > \mu$, the perpetual value $W(x; \mu) := E_x \left[\int_0^{\tau_B} e^{-rt} X_t dt \right]$ satisfies*

$$W(x; \mu) = \frac{1}{r - \mu} (x - B^{1-p_2} x^{p_2}), \quad p_2 = \frac{-(\mu - \sigma^2/2) - \sqrt{(\mu - \sigma^2/2)^2 + 2\sigma^2 r}}{\sigma^2},$$

the negative root of $\frac{1}{2}\sigma^2 p^2 + (\mu - \sigma^2/2)p - r = 0$.

Proof. W satisfies the ODE $\frac{1}{2}\sigma^2 x^2 W'' + \mu x W' - rW + x = 0$ for $x > B$, by the same martingale argument as Proposition 13.2, now with a running source term x in place of a terminal condition, and boundary conditions $W(B) = 0$, $W(x)/x$ bounded as $x \rightarrow \infty$. Direct substitution confirms $W_p(x) = x/(r - \mu)$ solves the ODE exactly (the source term x is matched by the particular solution's own linear growth). The homogeneous equation $\frac{1}{2}\sigma^2 x^2 W'' + \mu x W' - rW = 0$ admits power solutions x^p for p solving the stated quadratic, with roots $p_1 > 0 > p_2$; boundedness of $W(x)/x$ as $x \rightarrow \infty$ rules out the x^{p_1} term (since p_1 exceeds 1 whenever $r > \mu$, making x^{p_1} grow faster than linearly), leaving $W(x) = x/(r - \mu) + Cx^{p_2}$. Imposing $W(B) = 0$ pins down $C = -B^{1-p_2}/(r - \mu)$, giving the stated formula. $\square$

We verified this formula two ways before relying on it: against an independently computed numerical solution of the same boundary value problem using a standard ODE solver, matching to eight significant figures, and, at coarser precision, against direct Monte Carlo simulation of the discounted integral, where residual disagreement at achievable simulation resolutions was consistent with the known discrete-monitoring bias inherent to simulating continuously monitored barriers, not with any error in the formula.

Proposition 14.11 (Optimal investment threshold). *The value of the option to invest, $F(x) := \sup_{\tau} E_x [e^{-r\tau} (W(X_{\tau}; \mu_H) - I) \mathbf{1}_{\{\tau < \tau_B\}}]$, is, in the continuation region $B < x < x^*$,*

$$F(x) = A(x^{p_1} - B^{p_1-p_2}x^{p_2})$$

for p_1, p_2 the roots (Proposition 14.10, evaluated at μ_L) of the low-drift characteristic equation, where the constant A and the optimal exercise threshold x^ are the unique solution to the value-matching and smooth-pasting conditions*

$$F(x^*) = W(x^*; \mu_H) - I, \quad F'(x^*) = W'(x^*; \mu_H).$$

For $x \geq x^$, immediate investment is optimal and $F(x) = W(x; \mu_H) - I$.*

Proof sketch. In the continuation region, no running payoff accrues to the option holder, so the same homogeneous ODE used in Proposition 14.10's proof governs F , now with μ_L in place of μ ; the boundary condition $F(B) = 0$ (the option is worthless once the ecosystem has already collapsed) rules out the same combination of homogeneous solutions as before, giving $F(x) = A(x^{p_1} - B^{p_1-p_2}x^{p_2})$ for a single free constant A . Two further conditions are needed to pin down A and the free boundary x^* : value matching requires that the option's value equal the exercise payoff exactly at the boundary, since the holder is indifferent between exercising and not exercising there by definition of an optimal stopping boundary; smooth pasting, the requirement that F and the exercise payoff meet tangentially rather than at a kink, is the standard optimality condition for this class of free-boundary problem, since a kink at x^* would imply a nearby, better exercise policy, contradicting optimality of τ^* . A full verification argument, confirming this construction actually solves the stochastic control problem rather than merely satisfying its formal optimality conditions, follows the standard approach for this problem class; see Dixit & Pindyck (1994) for the argument in full generality. $\square$

Worked Example

Take $\mu_L = 0.01$, $\mu_H = 0.04$, $\sigma = 0.30$, $r = 0.05$, $B = 60$, and a sunk investment cost $I = 50$. Solving the value-matching and smooth-pasting system numerically gives $A \approx 10.063$ and $x^* \approx 66.61$: the optimal policy is to invest as soon as the ecological value index reaches approximately 66.61, a modest premium above the collapse barrier itself, and to invest immediately if the index already exceeds this threshold on obser-

vation. Since this book's running worked example uses $x_0 = 100 > x^* \approx 66.61$, the model's own policy recommendation for that starting point is immediate investment: waiting is dominated by the collapse risk incurred during the delay, once the ecosystem is already comfortably above the barrier. We solved for (A, x^*) numerically rather than in closed form, since the transcendental system defined by value matching and smooth pasting does not reduce to elementary functions in general; the solution was confirmed by substituting back into both defining equations, matching to numerical precision.

This result formalizes Vuong's semiconducting principle as exactly the kind of irreversible-investment problem Arrow & Fisher (1974) and Dixit & Pindyck (1994) analyze, rather than only asserting the principle's directionality as Proposition 14.12 below does. Money spent on conservation is an option exercise, valuable precisely because it is irreversible and because delay carries collapse risk; nothing in this section's construction gives money any power to reverse a collapse that has already occurred, consistent with, and now embedded inside, a genuine optimization problem rather than stated as a separate postulate.

For an environmental thinker, the practical content of the threshold x^* is a direct answer to a question conservation funders and managers actually face: not "should we invest," but "have we waited too long to still benefit from waiting." Below x^* , delay still has option value, information may arrive, cheaper interventions may emerge, and holding off is not simply procrastination. Above x^* , the model says the option value of waiting has already been exhausted by collapse risk, and further delay is pure loss with no offsetting benefit.

We can now state, precisely, the sense in which this chapter's model realizes Vuong's semiconducting principle rather than merely gesturing toward it.

Proposition 14.12 (Directional exchange). *Monetary investment, modeled as an increase in the drift parameter μ , can raise $V(x, 0)$ and can lower $P(\tau_B \leq T)$ for $t < \tau_B$. No choice of μ , applied after τ_B has occurred, restores V to a positive value: $V(X_t, t) = 0$ for all $t \geq \tau_B$, identically in μ .*

Proof. The first claim follows from Proposition 14.3: $V(x, 0)$ is continuous and, for fixed $x > B$, strictly increasing in μ over the relevant range (a direct, if tedious, calculus verification using the explicit formula, which we do not reproduce symbol by symbol), reflecting that a higher expected growth rate raises both the expected terminal payoff and lowers collapse probability, both effects working in the same direction. The second claim is immediate from the model's own construction in Chapter 13: τ_B is defined as the first hitting time, after which the asset's value is set permanently to zero by fiat, a structural feature of the model rather than a computed consequence of any parameter, so no value of μ chosen for $t \geq \tau_B$ changes it. $\square$

This is, admittedly, a simple observation once the model is built, close to definitional. Its value is not in its difficulty but in what it makes explicit: within this model, exactly as Vuong (2021) argues informally, monetary intervention is only ever effective *before* collapse, working through the drift and volatility of an ongoing process, and has no formal mechanism, none at all, for reversing collapse once it occurs. The asymmetry is not an empirical regularity this model happens to reproduce; it is built into the model's absorbing boundary condition from the start, exactly as the semiconducting principle claims it should be.

Figure 14.2 draws the mechanism using the same visual language a semiconductor diode uses: a component that permits current in one direction and blocks it in the other. Environmental value converts to monetary value under forward bias, the direction the model's own drift and valuation machinery formalize throughout this chapter; the reverse conversion is

blocked, exactly as Proposition 14.12 proves, with the one partial exception already established in Section 14.7: a restoration effort can recover a fraction $\gamma < 1$ of what was lost, drawn here as a thin leakage past the barrier rather than a second forward channel.




Figure 14.2: The semiconducting principle of value exchange, drawn as a diode. Environmental value converts to monetary value under forward bias (top); the reverse conversion is blocked (bottom), with only a thin, partial leakage governed by the restoration parameter γ .

14.10 A Correction Factor for Conventional Valuation

Contingent valuation, ecosystem service accounting, and standard discounted cash flow methods, as ordinarily applied to natural capital, compute something structurally equivalent to $V_0(x, T) = xe^{(\mu-r)T}$: an expected discounted flow with no allowance for an absorbing collapse threshold. This section makes the correction such methods are missing explicit and reusable, rather than leaving Proposition 14.3's formula as a standalone result.

Setup

Define the correction factor $\kappa(x/B, \sigma, \mu, r, T) := V(x, 0)/V_0(x, T)$, the fraction of a conventional valuation that survives once the collapse barrier is properly priced.

x/B	$\sigma = 0.15$	$\sigma = 0.20$	$\sigma = 0.25$	$\sigma = 0.30$	$\sigma = 0.35$
1.10	0.243	0.191	0.161	0.143	0.130
1.25	0.499	0.404	0.346	0.309	0.284
1.50	0.745	0.632	0.555	0.502	0.464
2.00	0.930	0.847	0.774	0.716	0.672
3.00	0.993	0.967	0.927	0.885	0.848
5.00	1.000	0.997	0.987	0.970	0.950

Table 14.4: Correction factor $\kappa = V/V_0$, computed from Proposition 14.3 at $\mu = 0.01$, $r = 0.03$, $T = 10$. A conventional valuation should be multiplied by κ , not taken at face value, whenever the asset in question faces a collapse threshold at the stated proportional distance.

Table 14.4 is directly usable: given an estimate of how far a natural asset's current state sits above a literature-derived collapse threshold, expressed as the ratio x/B , and a volatility estimate, a practitioner already computing V_0 by conventional means can read off the multiplica-

tive correction this book's model implies, rather than treating V_0 as the answer. At $x/B = 1.5$ and $\sigma = 0.30$, for instance, the correction factor is approximately 0.50: a conventional valuation is roughly double what the barrier-adjusted model gives, under these parameters, for an asset only fifty percent above its own collapse threshold. We emphasize that this table is specific to the μ, r, T chosen here, not a universal constant; recomputing it for a different discount rate or horizon is a direct application of Proposition 14.3, not a new derivation.

This Is a Down-and-Out Barrier Claim

The valuation formula this chapter has built is, in the vocabulary of derivatives pricing, a down-and-out claim: a payoff that is cancelled entirely if the underlying falls to a barrier before maturity, priced by exactly the reflection-principle technique standard in that literature (Björk 2009). This is not a loose analogy; Proposition 14.3's formula has the same structural form, a Black-Scholes-type term minus a reflected correction term, as the classical down-and-out barrier option formula, specialized here to a linear rather than a convex payoff.

This equivalence gives Section 14.9's optimal conservation investment a second, complementary reading. Raising μ from μ_L to μ_H at cost I does not merely change a growth rate; in the language of this section, it is equivalent to purchasing barrier protection, converting a down-and-out claim on the low-drift process into a more valuable claim with a more favorable survival profile. Real-options practitioners already price exactly this kind of protection in other domains, catastrophe bonds and knock-out reinsurance being the closest financial analogues; this chapter's contribution is not the pricing technique itself, which is standard, but its application to an asset, ecological value, and a decision, conservation investment, that the existing literature on barrier claims does not itself address.

14.11 Connecting Back to Entropy

Nothing in this chapter's construction refers to entropy, Gibbs distributions, or the informational quanta apparatus of Chapters 2 through 11, and we do not claim a formal reduction of one to the other. The connection is thematic rather than technical: Chapter 11's value formation theorem established that entropy reduction, within this book's discrete apparatus, is directional, a mind's distribution can be driven away from uniformity, but nothing in that chapter's machinery runs in reverse to restore indifference once prioritization has occurred. This chapter's absorbing barrier is the same directionality, transplanted into continuous time and applied to an external, physical quantity, ecological value, rather than to an internal, cognitive one. Both chapters formalize a one-way street; only the address has changed.

14.12 Implications for the Climate and Ecological Crisis

The parameters used in this chapter's worked example were chosen for illustration, not fitted to any actual ecosystem, and we want to be direct about that limitation. What the model does establish, independent of any particular parameter choice, is Proposition 14.4's categorical claim: a valuation method that omits an absorbing collapse threshold will overstate an ecological asset's worth by an amount that grows precisely as collapse becomes more likely, which is to say, the overvaluation is largest exactly where the true stakes are highest. This is a mathematical mirror of a concern raised repeatedly in the ecological economics and climate policy literature, that market-based instruments calibrated to smooth, continuous price discovery are poorly suited to assets facing threshold effects and tipping points, and that this mismatch is not a minor calibration error but a structural feature of applying the wrong class of model. Nguyen et al. (2025b) document a concrete institutional version of this mechanism, arguing that costly, high-entropy technological interventions such as carbon capture

and storage, bioplastics, and glacier geoengineering are systematically favored by climate policy over cheaper, better-tested, nature-based alternatives precisely because a semiconducting-principle violation, treating environmental harm as monetarily offsettable, is built into how such interventions are evaluated and funded; our Overvaluation Theorem gives that same institutional pattern an exact, if idealized, quantitative form. We do not claim this chapter resolves that policy debate; we claim only that the specific, quantifiable overvaluation Proposition 14.4 proves gives that broader concern a precise mathematical form, in at least one tractable case.

14.13 Limitations

Three simplifications in this chapter's model are worth naming plainly, in the same spirit as the rest of this book's scope discussions. First, collapse here is instantaneous and total; real ecological thresholds more often involve degraded but nonzero continued function, which this model does not represent. Second, μ and σ are held constant, when actual ecological and climate dynamics plausibly involve time-varying, and possibly accelerating, degradation pressure. Third, the barrier B is fixed and known with certainty, when actual tipping points are, in practice, uncertain and only partially observable in advance, arguably the single most policy-relevant complication this chapter's model omits. Each of these is a natural direction for extending the model, and each would very likely require abandoning closed-form solvability in favor of numerical methods; we have chosen closed-form tractability deliberately, to make the overvaluation mechanism as transparent as possible, at the acknowledged cost of realism on all three fronts.

14.14 Summary

We solved the valuation PDE of Chapter 13 in closed form using the reflection principle, extended to drift via a change of measure; proved, without qualification, that any valuation ignoring an ecological asset's collapse risk strictly overstates its worth; and derived a companion closed-form formula for the collapse probability itself. We then went beyond that pricing result in three directions. We proved, by an elementary pathwise scaling argument requiring no calculus, that collapse probability falls strictly with distance from the barrier, and reported the companion volatility sensitivity, the Collapse Vega, as numerically verified rather than proved in general; we showed that overvaluation survives averaging over a genuinely uncertain barrier, and that the standard practice of substituting a point estimate for that uncertain barrier introduces a second, compounding source of overvaluation under the concavity we observed numerically; and we added a genuine decision problem, an optimal stopping formulation of when to undertake an irreversible conservation investment, solved via value matching and smooth pasting in the tradition of Arrow & Fisher (1974) and Dixit & Pindyck (1994), giving the semiconducting principle a policy-theorem form rather than only a directional postulate. We connected the model's absorbing-barrier mechanism explicitly to Vuong's semiconducting principle throughout, and to this book's own entropy-based apparatus, as a second, independent formalization of the same directional theme. Chapter 17 closes the book by returning to its full set of claims, old and new, and stating plainly what has been proved, what has been checked against evidence, and what remains speculative.

Biodiversity as an Irreplaceable Source of Cultural Information

Chapters 13 and 14 modeled irreversibility in continuous time, treating an ecological asset's monetizable worth as a stochastic process that can be permanently destroyed. This chapter models a different, complementary channel by which biodiversity loss destroys value, not the monetizable services an ecosystem provides, but its role as a source of informational material that human cultural creativity draws on and cannot manufacture from anything else. The formalization here is discrete and combinatorial rather than continuous and stochastic, and it connects back to this book's own apparatus directly, through Chapter 2's entropy ceiling and Chapter 6's combinatorics, rather than only thematically as Chapters 13 and 14 do.

15.1 Biodiversity as Informational Input

A body of work in anthropology, geography, and cultural ecosystem services research argues that the ecosystems people inhabit shape their cultural values, practices, and creative output, and that this relationship has grown more, not less, significant under digital mediation, where a single piece of nature-derived information, an image, a species, a landscape, can be copied, recombined, and redistributed globally at negligible cost. Within the granular interaction thinking apparatus this book has built, that claim has a natural formal home: if the mind is, as Chapter 9 modeled it, a filtered subset of a larger environment, then biodiversity is part of that environment, a source of candidate informational quanta, species, ecological relationships, landscapes, that human minds draw on, recombine, and retain as cultural material. Vuong & Nguyen (2025) develop this connection under the name Nature Quotient, the capacity to perceive, process, and organize information about the interconnections among humans, wildlife, ecosystems, and climate systems, and argue that this capacity is foundational to sustaining creative and cohesive cultures. Empirical work on how people actually perceive the consequences of biodiversity loss lends further support to treating this as a live policy concern rather than an abstraction: Nguyen & Jones (2022a) find that Vietnamese urban residents who perceive environmental degradation, lost economic opportunity, and lost knowledge as consequences of biodiversity decline are more likely to support restrictions on illegal wildlife consumption, and Nguyen & Jones (2022b) find that the same perceived losses, transmitted through an attitude of conservation endorsement, predict willingness to pay for conservation in protected areas, evidence that the informational and cultural stakes of biodiversity loss are ones ordinary people already register and act on, not only ones a formal model asserts.

What this literature establishes qualitatively, and what this chapter formalizes, is a specific claim: biodiversity is not merely one input to cultural creativity among many substitutable others, but a non-substitutable one, in the precise sense that a species once extinct cannot be

reconstituted as a source of new informational quanta by any amount of subsequent cultural or economic effort. This is a claim about the *environment* in Chapter 9's sense, not about the mind: it is not that people forget, but that the candidate set they could in principle draw from has permanently shrunk.

15.2 A Shrinking Environment

Setup

Let $N(t) \in \mathbb{N}$ denote the number of biodiversity-derived informational quanta available in the environment at time t , a count of the distinct species, ecological relationships, or landscape features that could, in principle, enter some mind's candidate set under the filtering process of Chapter 9. Model extinction as a pure death process: $N(t)$ is non-increasing, decreasing by one whenever a species or ecological feature is permanently lost, with no corresponding upward jump. Formally, $N(t) = N(0) - D(t)$, where $D(t)$ is a counting process (for instance, an inhomogeneous Poisson process with intensity $\lambda(t)$ reflecting the current rate of biodiversity loss) that only increases.

This is a deliberately minimal model, and simpler than Chapter 13's continuous-diffusion treatment of a single ecological asset's value; it says nothing about which species are lost, when, or why, only that the count of surviving candidate quanta moves in one direction. That minimality is the point: the results below hold for any intensity $\lambda(t)$ whatsoever, including one that varies with policy, climate feedback, or conservation effort, since nothing in what follows uses $\lambda(t)$'s specific form.

15.3 The Cultural Entropy Ceiling Theorem

Proposition 15.1 (Entropy ceiling shrinks with biodiversity). *Let $H_{\max}(t)$ denote the maximum entropy achievable by any distribution a mind could form over the environment's $N(t)$ biodiversity-derived quanta at time t . Then $H_{\max}(t) = \log_2 N(t)$, and $H_{\max}(t)$ is non-increasing in t .*

Proof. By Proposition 2.1 (Chapter 2), the maximum entropy of a distribution on n outcomes is $\log_2 n$, achieved by the uniform distribution. Applying this with $n = N(t)$ gives $H_{\max}(t) = \log_2 N(t)$. Since $\log_2$ is strictly increasing and $N(t)$ is non-increasing by construction, $H_{\max}(t)$ is non-increasing. $\square$

The interpretation is direct, given Chapter 11's own reading of entropy: $H_{\max}(t)$ is not the entropy any actual culture exhibits, but the ceiling on how much undifferentiated raw material, how large a space of distinguishable possibilities, is even available for a culture to prioritize among in the first place. A culture cannot form more distinct biodiversity-derived cultural values than its environment currently contains candidate quanta for, regardless of how sophisticated its own internal information processing is. Proposition 15.1 says this ceiling can only fall, never rise, as extinction proceeds, formalizing, in the discrete apparatus already built across Parts II and III, the claim that biodiversity loss constrains the raw space of future cultural possibility, not merely the stock of already-realized cultural products.

The result inverts as cleanly as it states, and the inversion is worth recording explicitly, since it converts a descriptive theorem into a usable policy target.

Corollary 15.1 (Minimum richness for a target entropy ceiling). *To guarantee a cultural entropy ceiling of at least H_{target} bits, richness must satisfy $N(t) \geq 2^{H_{\text{target}}}$.*

Proof. Immediate from Proposition 15.1: $H_{\text{max}}(t) = \log_2 N(t) \geq H_{\text{target}}$ if and only if $N(t) \geq 2^{H_{\text{target}}}$, since $\log_2$ is strictly increasing. $\square$

Worked Example

A policy target of preserving at least $H_{\text{target}} = 20$ bits of cultural entropy ceiling requires $N(t) \geq 2^{20} = 1,048,576$ candidate quanta; a more modest target of $H_{\text{target}} = 10$ bits requires only $N(t) \geq 2^{10} = 1,024$. Corollary 15.1 is exact given Proposition 15.1, but we repeat the caution already stated there: it bounds a ceiling on possibility, not a guarantee that any given culture achieves entropy anywhere near that ceiling, and the identification of $N(t)$ with any specific empirical count remains the same illustrative simplification flagged in Section 15.10 below.

For an environmental thinker, this reframes what is lost when a species disappears. The immediate loss is the organism and its ecological role; this proposition points at a second, easily overlooked loss, a permanent reduction in the raw informational material future generations could have drawn on, whether for science, art, or simply noticing something about the world that would otherwise never have been noticed at all. What information is being lost is not only about the present.

Proposition 15.1 is worth distinguishing carefully from a second, independent mechanism this book already proved, elsewhere, for a different purpose: Proposition 4.2's low-temperature limit (Chapter 4), sharpened into a general theorem by Proposition 11.1 (Chapter 11). That result concerns a mind's distribution collapsing toward a single option as effective temperature falls, entirely at fixed support: H_{max} does not move, since N itself is untouched, but $H(P)$ falls toward zero as probability concentrates on whichever option currently carries the lowest energy. Proposition 15.1 and this earlier result are not competing descriptions of the same failure; they describe two different things going wrong, one in the space of possibility itself, the other in how attention is distributed across a possibility space that remains fully intact.

The distinction has a direct illustration already cited in this book: Nguyen et al. (2025b)'s account of technological solutionism in climate policy, discussed in Chapter 14 for a different reason, is a case of the second mechanism rather than the first. Carbon capture, bioplastics, and geoengineering do not reduce the number of available response options, nature-based alternatives remain formally on the table, so H_{max} is unaffected; what collapses is $H(P)$, as policy attention and funding concentrate on the technological option to the near-exclusion of cheaper, better-tested alternatives. A reader applying Section 17.2's first question, what information is being lost, should ask it twice: once about the size of the possibility space itself, Proposition 15.1's question, and once about how narrowly attention is concentrated within whatever space remains, Proposition 11.1's question. A policy response can fail either test independently of the other, and biodiversity loss and technological solutionism are, in this precise sense, different failures rather than instances of the same one.

15.4 Combinatorial Diversity Collapse

Proposition 15.1 concerns the ceiling on undifferentiated possibility. A sharper question asks about differentiated possibility: given that a culture prioritizes some fixed number k of biodiversity-derived quanta as values, in the sense of Chapter 7's capacity bound, how many genuinely distinct ways are there to choose which k quanta those are?

Proposition 15.2 (Shrinking choice of cultural configurations). *Fix $k \geq 1$, and let $C(t) := \binom{N(t)}{k}$ denote the number of distinct k -element subsets of the environment's quanta available at time t (each such subset representing one possible configuration of k simultaneously prioritized biodiversity-derived values, in the sense of Proposition 7.1). Then $C(t)$ is non-increasing in t , and strictly decreasing whenever $N(t) \geq k$ falls.*

Proof. $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ is strictly increasing in n for fixed $k \geq 1$ and $n \geq k$, since $\binom{n+1}{k}/\binom{n}{k} = (n+1)/(n+1-k) > 1$ whenever $n+1-k > 0$. Since $N(t)$ is non-increasing, $C(t) = \binom{N(t)}{k}$ is non-increasing, strictly so whenever $N(t)$ strictly decreases while remaining at least k . $\square$

The combinatorial magnitude here is worth pausing on, since it is easy to underestimate. Table 15.1 shows $\binom{N}{5}$, the number of distinct five-value cultural configurations available, for several values of N spanning two orders of magnitude. A modest proportional decline in N produces a vastly larger proportional decline in the number of distinct configurations available, since $\binom{N}{k}$ grows polynomially of degree k in N : halving N roughly halves $H_{\max}$ by only one bit (Proposition 15.1), but can reduce $\binom{N}{k}$ by a factor close to 2^k .

N	$\binom{N}{5}$
1000	8,250,291,250,200
100	75,287,520
50	2,118,760
20	15,504
10	252

Table 15.1: The number of distinct five-quantum cultural configurations, $\binom{N}{5}$, collapses far faster than N itself as biodiversity declines.

Proposition 15.3 (Elasticity of configuration count with respect to richness). *For $k \ll N$, $\frac{d \log C}{d \log N} \approx k$: a small proportional decline in richness produces, to leading order, a k -fold larger proportional decline in the number of distinct achievable configurations.*

Proof. $\log C = \log N! - \log k! - \log(N-k)!$, so $\frac{d \log C}{dN} = \psi(N+1) - \psi(N-k+1)$ where ψ is the digamma function; for $k \ll N$, $\psi(N+1) - \psi(N-k+1) \approx k/N$ by a first-order expansion of ψ . Multiplying by N to convert to the elasticity $\frac{d \log C}{d \log N} = N \cdot \frac{d \log C}{dN} \approx N \cdot (k/N) = k$. $\square$

Worked Example

Proposition 15.3 converts directly into a headline number. For $k = 5$, a 50% decline in richness, $N(t)/N(0) = 0.5$, implies a configuration-space decline of approximately $(0.5)^5 = 0.03125$, a 96.9% loss, matching, to within rounding, the exact figures already computed for $N = 1000 \rightarrow 500$ in Table 15.1's underlying data: $\binom{500}{5}/\binom{1000}{5} \approx 0.0309$, a 96.9% decline. We verified the elasticity's convergence to exactly k numerically before stating Proposition 15.3: at $N = 1000$, the numerically computed elasticity is 5.01; at $N = 100$, still 5.10, confirming the approximation holds well outside the strict $k \ll N$ regime the proof formally requires.

For an environmental thinker, this is the sharpest single number this chapter produces, and the one most worth remembering: a linear-sounding loss, half of a region's species gone, does not translate into a linear-sounding loss of possibility. It compounds, and compounds fast, for

exactly the same combinatorial reason that losing half the letters of an alphabet destroys far more than half of the words that could be spelled with it. What future possibilities disappear with today's loss is not a metaphor; it has an exponent, and the exponent is the number of things being simultaneously valued at once.

15.5 Beyond Richness: Hill Numbers

Every result so far treats the $N(t)$ surviving quanta as interchangeable, tracking only their count. Real biodiversity is not like this: species differ enormously in abundance, and a community of 100 species in which one species accounts for 91% of all individuals is not, by any reasonable standard, as diverse as a community of 100 equally common species, even though both have identical raw richness $N = 100$. This section replaces the raw count with the standard family of diversity indices that takes abundance into account, and shows that this book's own entropy-ceiling result is one member of that family, not a competitor to it.

Setup

Let $p_1(t), \dots, p_{N(t)}(t)$ denote the relative abundances of the $N(t)$ surviving quanta at time t , $\sum_i p_i(t) = 1$. The Hill number of order $q \geq 0$ is

$${}^qD(t) := \left(\sum_i p_i(t)^q \right)^{1/(1-q)}, \quad q \neq 1; \quad {}^1D(t) := \exp \left(- \sum_i p_i(t) \ln p_i(t) \right),$$

the second expression recovered as the continuous limit of the first as $q \rightarrow 1$.

Proposition 15.4 (Hill numbers and the entropy ceiling). ${}^0D(t) = N(t)$ identically, regardless of the abundance distribution. For $q \in \{1, 2\}$, ${}^qD(t) \leq N(t)$, with equality if and only if the distribution is uniform.

Proof. For $q = 0$, $\sum_i p_i^0 = N(t)$ (a sum of $N(t)$ ones, since $p_i > 0$ for every surviving quantum), and $(N(t))^{1/(1-0)} = N(t)$, giving ${}^0D(t) = N(t)$ exactly, independent of the p_i 's. For $q = 1$: writing $H(t)$ (in nats) for the Shannon entropy of the abundance distribution, Proposition 2.1 (Chapter 2, there stated in bits but identical in any base) gives $H(t) \leq \ln N(t)$, with equality iff uniform; exponentiating both sides, since $\exp$ is strictly increasing, gives ${}^1D(t) = \exp(H(t)) \leq N(t)$, equality under the same condition. For $q = 2$: by the Cauchy-Schwarz inequality applied to the vectors $(p_1, \dots, p_{N(t)})$ and $(1, \dots, 1)$, $(\sum_i p_i \cdot 1)^2 \leq N(t) \sum_i p_i^2$, and since $\sum_i p_i = 1$, this gives $\sum_i p_i^2 \geq 1/N(t)$, so ${}^2D(t) = 1/\sum_i p_i^2 \leq N(t)$, with equality in Cauchy-Schwarz, and hence here, if and only if all p_i are equal. $\square$

The $q = 1$ case of Proposition 15.4 is, transparently, this book's own Proposition 2.1 wearing different notation, exponentiated rather than left in log form; we did not need a new proof, only a change of units. This is exactly the sense in which Chapter 1's scaffolding claim is meant: the apparatus built for an entirely different purpose in Part I is not merely analogous to what biodiversity science already uses, it is, after the substitution $q = 1$, the identical mathematical object.

Proposition 15.4 bounds ${}^qD(t)$ above by $N(t)$; it does not say ${}^qD(t)$ declines faster or slower than $N(t)$ as biodiversity erodes, since that depends on how the abundance distribution itself changes, not only on how many species remain. This turns out to depend critically on which species are lost.

Worked Example

Starting from a single skewed baseline abundance distribution over $N_0 = 100$ species (Shannon and Simpson diversity of 46.4 and 31.8 respectively at full richness, already well below the ceiling of 100, reflecting the baseline's own unevenness), we computed 1D and 2D under three erosion scenarios, verified numerically rather than asserted: losing the rarest species first (an extinction filter, the empirically common pattern in which specialist, low-abundance species are lost before generalists), losing the most abundant species first (as under selective harvest or hunting pressure targeting common or economically valuable species), and losing species at random independent of abundance. At 60% of original richness remaining, the rarest-first scenario left 1D at 89% of its original value, essentially unaffected, since the lost species were barely contributing to either diversity index in the first place; the random-loss scenario left 1D at only 53%; and the abundant-first scenario left 1D at 73%, an intermediate but still substantial decline. The same pattern holds for 2D . Raw richness, by construction, fell to exactly 60% of its baseline in every scenario, since all three remove the same number of species.

The qualitative lesson is that raw richness, $N(t)$, is not a reliable proxy for abundance-weighted diversity once a specific mechanism of loss is in view: the same fractional decline in $N(t)$ can leave 1D and 2D nearly untouched or can more than halve them, depending on whether the species being lost were already marginal contributors to diversity or were central to it. This is a genuine complication for the entropy-ceiling and combinatorial-collapse results of the two preceding sections, both of which are stated in terms of $N(t)$ alone; a full account would track the abundance distribution's own evolution, not only its support size, a natural extension we do not carry out here.

For researchers and policymakers, the practical lesson is that a species count on its own can be a misleading headline number. Two regions could lose the same number of species and end up in very different states of real diversity, depending on whether the species lost were rare specialists that barely registered in the ecosystem's overall balance, or common, central species that many other parts of the system depended on. A monitoring program that only tracks how many species remain, without asking which ones, can miss the more important story. This does not mean richness counts are useless. It means they should be read alongside some sense of which species are actually disappearing, not treated as a complete picture on their own.

15.6 The Semiconducting Principle, Informational Version

Chapter 14 formalized Vuong (2021)'s semiconducting principle for monetary and environmental value directly: money converts to environmental value only in one direction. The model of this chapter makes the same asymmetry visible in a different currency, informational rather than monetary.

Proposition 15.5 (No informational buy-back). *No expenditure of money, labor, or cultural effort at any time $t' > t$ increases $N(t)$ above $N(t)$ once a decrease has occurred: for all $t' > t$, $N(t') \leq N(t)$.*

Proof. Immediate from the model's construction: $N(t) = N(0) - D(t)$ with $D(t)$ a non-decreasing counting process, by definition of a pure death process. $\square$

As with Proposition 14.12 in Chapter 14, this is close to definitional once the model is stated, and its value lies in what it makes explicit rather than in its difficulty. Money, applied

before a species is lost, can lower the intensity $\lambda(t)$ governing further loss, exactly as it could raise the drift μ in Chapter 14's continuous model. No expenditure, applied after the fact, restores what has already been lost, because the model gives it no mechanism to do so. Cultural memory, restoration ecology, and de-extinction technology may partially soften this in practice, restoring some *use* of information about an extinct species even if not the species itself as an ongoing source of new information. This section makes that softening a parameter rather than an unmodeled caveat.

Setup

Let $\delta \in [0, 1]$ denote the fraction of an extinct quantum's informational content that survives in archived form, museum specimens, genomic libraries, ethnographic and artistic record, recoverable for cultural use though not as an ongoing source of new information. Define the effective count $N_{\text{eff}}(t) := N(t) + \delta \cdot D(t)$, where $D(t) = N(0) - N(t)$ is the number of extinctions by time t . $\delta = 0$ recovers this chapter's original model exactly; $\delta = 1$ treats archived information as fully equivalent, for entropy-ceiling purposes, to a living quantum.

Proposition 15.6 (Archival memory raises the effective ceiling but cannot restore the original). *For $\delta \in [0, 1)$, $N_{\text{eff}}(t) < N(0)$ whenever $D(t) > 0$, so $\log_2 N_{\text{eff}}(t) < \log_2 N(0) = H_{\text{max}}(0)$ strictly. As $\delta \rightarrow 1$, $N_{\text{eff}}(t) \rightarrow N(0)$ exactly.*

Proof. $N_{\text{eff}}(t) = N(t) + \delta D(t) = N(0) - D(t) + \delta D(t) = N(0) - (1 - \delta)D(t)$. For $\delta < 1$ and $D(t) > 0$, $(1 - \delta)D(t) > 0$, so $N_{\text{eff}}(t) < N(0)$ strictly. As $\delta \rightarrow 1$, $(1 - \delta)D(t) \rightarrow 0$, so $N_{\text{eff}}(t) \rightarrow N(0)$. $\square$

The interpretation parallels Chapter 14's γ : archiving is a genuinely recoverable fraction of what was lost, not nothing, but Proposition 15.6 shows that unless archival fidelity is perfect, $\delta = 1$, the entropy ceiling never fully returns to its pre-extinction level, since archived information, by construction, no longer generates new informational quanta the way a living, evolving population does. Whether real archival practice sits closer to $\delta = 0$ or $\delta = 1$ for any given form of biodiversity-derived cultural information, a museum specimen's information content plausibly retains more of its value than an extinct species' unrecorded ecological relationships did, is an empirical question this book does not resolve; the parameter is stated so that resolving it, in either direction, would tell us something Proposition 15.5 alone cannot.

The point worth carrying forward for policy is a simple one, even though the model behind it is not. Museums, seed banks, and genetic archives are genuinely valuable, and this section does not argue otherwise. What it argues is that archiving is not the same thing as saving. A specimen in a collection can preserve information about a species long after it is gone, but it cannot put that species back into the living, changing process that once generated new cultural and ecological meaning from it. Funding for archives and funding for preventing extinction in the first place are addressing two different problems, and treating a well-stocked archive as a substitute for prevention overstates what archiving can actually deliver.

15.7 Coupling to Ecological Value

Chapter 15's pure death process for $N(t)$ and Chapters 13 and 14's stochastic model of ecological asset value X_t have, until now, been presented as two separate models sharing a theme rather than a mechanism. This section couples them directly: the rate at which species are lost is allowed to depend on the same process, X_t , whose collapse Chapters 13 and 14 price, giving one joint model in place of two parallel ones.

Setup

Let $D(t)$, the extinction counting process introduced in Chapter 15, be a doubly stochastic Poisson (Cox) process with intensity $\lambda(t) = \lambda_0 + \alpha/X_t$, $\lambda_0, \alpha \geq 0$, where X_t is the same absorbing GBM of Chapter 13: intensity rises as the ecological asset's value falls, reflecting that ecological stress on the specific asset X_t prices plausibly co-occurs with elevated extinction pressure on the broader biodiversity it is embedded in. As before, $N(t) = N(0) - D(t)$.

Proposition 15.7 (Expected extinction count under coupling).

$$E[N(0) - N(T \wedge \tau_B)] = E \left[\int_0^{T \wedge \tau_B} \lambda(t) dt \right] = \lambda_0 E[T \wedge \tau_B] + \alpha E \left[\int_0^{T \wedge \tau_B} \frac{dt}{X_t} \right].$$

Proof. For a Cox process with intensity $\lambda(t)$, the expected number of events by time t^* , conditional on the intensity path, is the compensator

$$\int_0^{t^*} \lambda(s) ds;$$

taking a further expectation over the randomness in X_t itself, by Campbell's theorem for doubly stochastic Poisson processes, gives the unconditional expectation as the expectation of the integrated intensity. Truncating at $T \wedge \tau_B$, since the model's dependence of λ on X_t is defined only while X_t remains above the collapse barrier, and substituting $\lambda(t) = \lambda_0 + \alpha/X_t$ and splitting the integral by linearity gives the stated decomposition. $\square$

Proposition 15.7 is exact but not, in general, closed-form: the second term involves $E \left[\int_0^{T \wedge \tau_B} X_t^{-1} dt \right]$, an expectation of a functional of GBM's occupation time weighted by $1/X_t$, which does not reduce to elementary functions for finite T . We evaluated it numerically rather than pursue a further closed-form reduction.

Worked Example

Taking $\mu = 0.01$, $\sigma = 0.30$, $x_0 = 100$, $B = 60$, $T = 10$ as in this book's running example, with $\lambda_0 = 0.5$ and $\alpha = 20$, Monte Carlo simulation gives $E[N(0) - N(T \wedge \tau_B)] \approx 3.64$ expected extinction events. This is, at first glance, counterintuitive: holding X_t fixed at its starting value $x_0 = 100$ (no stochasticity, no collapse risk) gives a larger expected count, $(\lambda_0 + \alpha/x_0)T = 7.0$, over the same horizon. The coupling does not straightforwardly increase extinction risk; it introduces two offsetting effects. Stress elevates the instantaneous rate, α/X_t grows as X_t falls, which by itself would push the count up. But truncation at τ_B cuts the integration window short whenever the asset collapses before T , an event with probability 70% under these parameters by Proposition 14.5, removing exposure to further extinction events after that point under this model's own bookkeeping convention. In this parameterization the second effect dominates the first, giving a lower expected count than the naive, uncoupled calculation; we report this as a genuine, non-obvious feature of the coupled model rather than smooth it into a simpler but less accurate directional story; a reader who prefers the intuition that collapse should trigger, not truncate, further biodiversity loss should replace the truncation convention with one in which $N(t)$ jumps to some residual floor at τ_B , a natural variant this section does not pursue.

The expectation alone does not say how reliable a guide it is; a mean of 3.64 could reflect a tightly concentrated count or one spread widely around it. The Cox process structure already in place gives the variance for free, without any new modeling assumptions.

Proposition 15.8 (Variance of the coupled extinction count). *Write $\Lambda := \int_0^{T \wedge \tau_B} \lambda(t) dt$ for the integrated intensity along a given path. Then*

$$\text{Var}[N(0) - N(T \wedge \tau_B)] = E[\Lambda] + \text{Var}[\Lambda].$$

Proof. Conditional on the full path of X_t (equivalently, conditional on Λ), $D(T \wedge \tau_B) := N(0) - N(T \wedge \tau_B)$ is Poisson distributed with parameter Λ , since a Cox process is, by definition, an ordinary Poisson process once its intensity path is fixed. A Poisson distribution has equal mean and variance, so $E[D | \Lambda] = \Lambda$ and $\text{Var}[D | \Lambda] = \Lambda$. The law of total variance gives

$$\text{Var}[D] = E[\text{Var}(D | \Lambda)] + \text{Var}[E(D | \Lambda)] = E[\Lambda] + \text{Var}[\Lambda],$$

using $E[\text{Var}(D | \Lambda)] = E[\Lambda]$ and $\text{Var}[E(D | \Lambda)] = \text{Var}[\Lambda]$ directly. $\square$

Worked Example

At the same parameters as above, Monte Carlo simulation gives $E[\Lambda] \approx 3.65$ (matching Proposition 15.7) and $\text{Var}[\Lambda] \approx 5.45$, so Proposition 15.8 predicts $\text{Var}[D] \approx 9.10$, a standard deviation of approximately 3.02 extinction events against a mean of 3.64. We verified this prediction directly against the variance of simulated extinction counts themselves, rather than only the formula's two components separately, and found agreement to within simulation error. A standard deviation comparable to the mean signals a genuinely dispersed count, not a tightly predictable one: this model does not support treating 3.64 as a reliable point forecast, only as the center of a fairly wide distribution.

The reason this matters beyond the mathematics is that a single expected number can create a false sense of precision. If a report says an ecosystem is expected to lose about four species over the next decade, it is tempting to plan around that number as though it were close to certain. This result shows why that would be a mistake in a case like this one. The uncertainty around that expected number can be as large as the number itself, meaning outcomes noticeably higher or lower than the average are genuinely plausible, not rare exceptions. Good planning should budget for that range, not just the midpoint.

Proposition 15.8 is exact for any horizon T , but it does not, by itself, say what shape that distribution takes, only its first two moments. Appendix B takes up that further question, asking when the count itself is approximately normally distributed, and shows that answering it well requires a different, complementary model from the one used here.

15.8 The Missing Direction: Does Richness Affect Volatility?

Section 15.7 makes $N(t)$ depend on X_t ; nothing so far makes X_t depend on $N(t)$ in return, despite this section's own title having called the result a joint model. This is a real gap, and closing it fully would require the kind of multi-species interaction model, a community matrix linking every pair of species' population dynamics, that this book's single-variable diffusion for X_t was deliberately built to avoid. What we can do without abandoning that tractability is bring in what community ecology already knows about how the number of interacting species relates to aggregate volatility, and let that knowledge motivate, rather than derive, a coupling in the missing direction.

Setup

May (1972) shows that a randomly interacting community of S species, with connectance C (the probability any two species interact) and interaction strengths of standard deviation α , is stable, in the local, linearized sense, only if $\alpha\sqrt{SC} < 1$; above this threshold, instability becomes near certain as S grows. Taken at face value, this says that *more* species, holding interaction strength and connectance fixed, pushes a community toward instability, not away from it. Tilman (1996), using two decades of grassland data, finds the reverse at the level of aggregate ecosystem properties: total community biomass becomes *more* stable as plant diversity rises, even as May's prediction of reduced population-level stability for individual species is simultaneously borne out in the same data. Tilman's own reading is that both are correct, at different levels of aggregation: many weakly correlated species buffer aggregate biomass through statistical averaging, the insurance effect, even while the individual populations themselves fluctuate more as the community's interaction structure grows more complex.

This book's process X_t is an aggregate, closer in spirit to Tilman's community biomass than to any single species' population, which is some reason to favor the insurance-effect direction; but X_t was also introduced, in Chapter 13, as a general ecological value index, not specifically as aggregate biomass, so we do not think the choice is settled by definition. We present both directions explicitly rather than picking one.

Worked Example

Suppose $\sigma(N) = \sigma_0\sqrt{N_0/N}$ (Tilman-type: richness dampens aggregate volatility through statistical averaging) against $\sigma(N) = \sigma_0\sqrt{N/N_0}$ (May-type: richness amplifies aggregate volatility through interaction complexity), calibrated to agree at $\sigma_0 = 0.30$ when $N = N_0 = 1000$. A 50% richness decline, $N(t)/N_0 = 0.5$, gives $\sigma = 0.30\sqrt{2} \approx 0.424$ under the Tilman-type coupling and $\sigma = 0.30/\sqrt{2} \approx 0.212$ under the May-type coupling. Feeding each into Proposition 14.5 (Chapter 14) at $x_0 = 100$, $B = 60$, $\mu = 0.01$, $T = 10$, collapse probability rises to approximately 0.84 under the Tilman-type coupling and falls to approximately 0.51 under the May-type coupling, against an uncoupled baseline of 0.70 at fixed $\sigma_0 = 0.30$. The same biodiversity decline implies a substantially more dangerous ecological asset under one hypothesis and a substantially safer one under the other, using nothing more than functional forms and Chapter 14's own, already-established Collapse Vega.

We regard this divergence as the point, not a defect to be resolved. The stability-complexity debate this section inherits from May and Tilman is a genuine, still-active one in ecology, not something this book's own apparatus can settle, and a model that silently picked one direction would be quietly taking a side in a live empirical dispute while presenting the choice as a technical detail. Which functional form, if either, actually describes a given ecosystem's aggregate volatility as a function of its richness is an empirical question a real dataset, tracking both $N(t)$ and a genuine proxy for X_t 's volatility over time, could in principle answer; this book does not have that dataset, and states the coupling as an open, falsifiable choice rather than resolving it by assumption.

15.9 A Worked Example

Worked Example

Global vertebrate population sizes fell by an average of approximately 69% between 1970 and 2018 (World Wide Fund for Nature 2022). Taking this proportional decline as an illustrative, order-of-magnitude proxy for the decline in $N(t)$, an environment with $N(0) = 8,000,000$ candidate quanta (roughly the estimated total number of species on Earth) falls to approximately $N(t) \approx 2,480,000$. By Proposition 15.1, the entropy ceiling falls from $\log_2(8,000,000) \approx 22.93$ bits to $\log_2(2,480,000) \approx 21.24$ bits, a drop of about 1.69 bits. This is a modest-looking number, and deliberately so: entropy is a logarithmic measure, and Proposition 15.1 alone understates the practical severity of the decline. Proposition 15.2's combinatorial count does not: for $k = 5$, the number of distinct five-quantum configurations falls by a factor of $\binom{8,000,000}{5} / \binom{2,480,000}{5} \approx (8,000,000/2,480,000)^5 \approx 349$, roughly two and a half orders of magnitude, far outpacing the proportional decline in N itself. We verified this figure directly, both via the power-law approximation and via an exact log-gamma computation of the two binomial coefficients, which agree to three significant figures; the ratio $\binom{N_1}{k} / \binom{N_2}{k} \approx (N_1/N_2)^k$ holds this closely because $k \ll N_1, N_2$ here.

We repeat the caution already familiar from this book's other applications: the mapping from a 69% population decline to a proportional decline in $N(t)$ specifically is an illustrative simplification, not a measurement. Population decline and species-level extinction are related but distinct quantities, and $N(t)$ in this chapter's model counts something closer to the latter. The qualitative conclusion, that combinatorial diversity collapses faster than any linear or logarithmic measure of decline would suggest, does not depend on this simplification; the specific number 349 does.

15.10 Limitations

This chapter's model shares Chapters 13 and 14's commitment to tractability over realism, and the resulting simplifications are worth naming plainly. The archival parameter δ introduced in Section 15.6 softens, but does not remove, the model's core commitment to one-directional loss: rewilding, captive breeding and reintroduction, and seed and genome banking represent real attempts to restore living populations rather than merely archive information about them, and none of these is what δ represents. Treating all quanta as informationally interchangeable within the count $N(t)$ ignores that species differ enormously in their cultural salience and creative influence, a concern Section 15.5's three-scenario comparison addresses only partially, showing that the direction of the effect depends on which species are lost without fully resolving how to weight cultural salience specifically. And the identification of $N(t)$ with any single empirical proxy, population counts, species counts, or otherwise, is, as the worked example above states directly, an illustrative simplification rather than a measurement this chapter defends.

15.11 Summary

Biodiversity loss constrains cultural creativity through a channel distinct from, and formalized differently than, the monetary asset destruction of Chapters 13 and 14: a shrinking pool of candidate informational quanta lowers the entropy ceiling on cultural possibility (Proposition 15.1) and collapses the combinatorial space of distinct achievable cultural configurations far faster than the underlying decline itself (Proposition 15.2), an effect Proposition 15.3 con-

verts into an exact elasticity, approximately k -fold for k simultaneously prioritized values, so that a 50% richness decline implies roughly a 97% configuration-space decline at $k = 5$. Archival memory raises the effective ceiling but, unless archival fidelity is perfect, cannot restore it fully (Proposition 15.6), a falsifiable weakening of the model's original all-or-nothing claim in the same spirit as Chapter 14's γ . We then went beyond raw species counts in two further directions. Hill numbers generalize the entropy ceiling to abundance-weighted diversity, recovering this chapter's own $q = 1$ result as a special case of the same maximum-entropy theorem used throughout Part I, and a numerically verified comparison across three loss scenarios showed that whether abundance-weighted diversity falls faster or slower than raw richness depends critically on which species are lost. And coupling the extinction intensity directly to the ecological value process X_t of Chapters 13 and 14 turned two thematically related models into one joint model, with a Campbell's-theorem expectation formula and a numerically verified example whose direction, a coupled model producing fewer expected extinctions than a naive fixed-value baseline under this section's own truncation convention, we reported honestly rather than smoothing into a simpler story. We closed the missing direction of that coupling explicitly rather than silently: whether richness dampens or amplifies X_t 's own volatility is a genuine, unresolved question in ecology, May (1972)'s complexity-instability result and Tilman (1996)'s insurance-effect finding pointing opposite ways, and we showed, using Chapter 14's own Collapse Vega, that the same biodiversity decline implies sharply different collapse risk depending on which direction is correct, a divergence we present as an open empirical question rather than resolve by assumption. Unlike Chapters 13 and 14, this chapter's connection to the rest of the book is technical rather than merely thematic: its central propositions follow directly from apparatus, the maximum entropy principle and elementary combinatorics, built in Parts I and II for an entirely different purpose. Chapter 17 closes the book by returning to its full set of claims, old and new, across the book as a whole.

Sufficiency as an Ecological Strategy

Chapters 13 through 15 ask what happens once collapse is already underway, or already priced. This chapter asks a question that belongs earlier in that sequence: how much restraint, and what kind, is enough to avoid needing those results at all. The question turns out to have a precise formal shape once posed in this book's own vocabulary, requiring no apparatus beyond what Chapters 2 and 13 already built.

16.1 The Earth Sufficiency Paradigm

The sustainability literature's own concept of sufficiency, meeting human needs within planetary limits by moderating consumption, has been criticized from within that literature as still anthropocentric: it treats ecological integrity as a boundary condition for human consumption to respect, rather than as a requirement with standing of its own. Vuong & Nguyen (2026) propose instead an *earth sufficiency* paradigm, in which the sufficiency of non-human species and ecological functions is a primary objective, not a byproduct of human restraint, secured through what they develop as an operational, three-tier hierarchy for the semiconducting principle already central to Chapter 14: preventing irreversible loss and restoring regenerative capacity takes strict priority; reducing non-essential throughput comes second; and technological intervention in genuinely hard-to-abate sectors comes third, explicitly barred from legitimizing continued expansion of the activities it is meant to offset. This is the same directional asymmetry Proposition 14.12 formalizes for a single valuation problem, generalized here into a standing policy over an entire portfolio of activities, and the same boundary the Innovation Curse discussion in Section 15.7 already illustrates for one specific technology tier.

Vuong & Nguyen (2026) tie the paradigm's practical success to cultivating Nature Quotient (NQ), the capacity to perceive ecological interdependence, and warn that its absence produces what they call *order-induced blindness*, self-validating cultural and institutional structures that obscure their own drift from ecological reality, precisely the phenomenon Section 12.3's belief-drift application already formalizes using Appendix A's machinery.

The intuition that non-extractive standing can itself be a form of flourishing, rather than a waste of extractable value, is considerably older than the sustainability literature. In *Free and Easy Wandering*, the opening chapter of the Zhuangzi (2013), Huizi complains to Zhuangzi that his enormous tree is worthless, too gnarled for a carpenter's plumb line, too twisted for compass or square. Zhuangzi's answer is that the tree's very uselessness to instrumental purposes is what has let it grow unmolested for centuries, standing where an axe will never find a reason to fall. We do not claim the classical text anticipates this chapter's mathematics, only that the reordering Vuong & Nguyen (2026) ask readers to make, treating an ecosystem's

standing value as prior to its extractable value rather than secondary to it, has a very old lineage.

This chapter takes the paradigm's remaining two claims, that ecological priorities are bounded in number and that the transition between an eco-deficit and an eco-surplus posture is not smoothly reversible, and gives each a precise, provable form using apparatus this book has already built.

16.2 A Capacity Bound for Conservation Priorities

The first claim requires no new machinery at all; it is a direct corollary of Proposition 2.1, restated here because the restatement is the point.

Corollary 16.1 (Capacity bound for conservation priorities). *Suppose an institution commits to treating m conservation priorities, among a wider candidate set of n , as genuinely held, meaning each receives probability weight at least θ in the institution's own resource-allocation distribution P . Then necessarily $\theta \leq 1/m$.*

Proof. If each of m priorities satisfies $p_i \geq \theta$, then $1 = \sum_i p_i \geq m\theta$, since P is a probability distribution and the m held priorities are a subset of its support; dividing gives $\theta \leq 1/m$ directly. $\square$

The reading is exact rather than merely suggestive: an institution cannot widen how many things it treats as genuine conservation priorities without lowering the guaranteed floor of attention or resourcing any one of them can claim. Committing to two priorities can guarantee each at most half of the relevant budget; committing to ten caps the guarantee at a tenth, a fivefold drop in worst-case robustness for a fivefold increase in breadth. Corollary 16.1 does not say which number of priorities is correct; it says only that the trade-off between how many and how robustly is not avoidable by better management, since it follows from the probability simplex alone, the same fact Chapter 7 already uses for a different purpose.

For an agency or a funding body, this speaks to a temptation that comes up constantly. It always sounds responsible to say that an organization cares about many issues at once, biodiversity, water, soil, climate, community livelihoods, all treated as top priorities together. This result shows that treating everything as a top priority is not free. Every additional item added to the list of guaranteed priorities lowers the minimum commitment any single one of them can be promised. This is not a criticism of breadth for its own sake. It is a reminder that breadth and depth trade off against each other by simple arithmetic, not by poor planning, and an honest strategy should say which of the two it is choosing rather than claiming both without limit.

16.3 Diminishing Returns to Ecological Prioritization

Corollary 16.1 bounds the floor; it says nothing about how much differentiation an institution actually achieves by holding m priorities rather than treating all n candidates as equally worth pursuing. The next result addresses this, and requires slightly more than a restatement, though still nothing beyond Chapter 2's own apparatus.

Proposition 16.1 (Diminishing returns to ecological prioritization). *Consider an institution that holds exactly m priorities out of n candidates, weighting the m held priorities uniformly at $\theta = 1/m$, the maximal floor Corollary 16.1 permits. Define the entropy reduction this achieves, relative to spreading attention uniformly across all n candidates, as $C(m) := \log_2 n - \log_2 m$,*

and the average reduction per priority as $c(m) := C(m)/m$. Then $c(m)$ is strictly decreasing on $(0, n]$.

We prove Proposition 16.1 in Appendix C, since the argument, while elementary, is a genuine calculus derivation rather than a restatement, and belongs with this book's other appendix-length supporting results rather than interrupting this chapter's main line.

Worked Example

For $n = 50$ candidate conservation targets, $c(1) = \log_2 50 \approx 5.64$ bits per priority; $c(5) \approx 0.66$; $c(10) \approx 0.23$; $c(25) = 0.04$. Moving from 5 to 10 held priorities does not merely halve the average return per priority, it cuts it to roughly a third, consistent with Proposition 16.1's strict-decrease claim. We verified this numerically, computing $c(m)$ directly across the full range $1 \leq m \leq n$, before stating the general result.

Read against the semiconducting principle's own three-tier structure, Proposition 16.1 gives a formal reason to expect the first tier, preventing irreversible loss, to concentrate on comparatively few, well-chosen priorities rather than spreading equally across every candidate threat: the marginal entropy-reduction return on the fifth or tenth simultaneously held priority is a fraction of the return on the first. This is a claim about a stylized boundary configuration, priorities weighted at the maximal permitted floor, not about how any actual institution's attention happens to be distributed; Vuong & Nguyen (2026)'s own discussion of a redistribution bottleneck under rigid consumption ceilings is, on this reading, a description of exactly the tension this proposition makes precise.

In everyday terms, adding a new priority to an already long list gives less benefit than adding it to a short one. The first priority an organization commits to concentrates its attention enormously, since it picks one focus out of many possible ones. The tenth priority adds much less, since attention is already spread thin by that point. This does not mean a tenth priority is worthless. It means the honest expectation for what that tenth priority will accomplish should be modest, and an organization that keeps adding priorities without acknowledging this pattern risks promising more than the arithmetic of attention allows it to deliver.

16.4 Hysteresis Between Eco-Deficit and Eco-Surplus

The two results above describe an institution at a fixed moment. The final result of this chapter describes how an institution moves, over time, between what Vuong & Nguyen (2026) call an eco-deficit posture, accumulating ecological debt in exchange for growth, and an eco-surplus posture, in which activity contributes to regenerative capacity rather than depleting it, and asks whether that movement is symmetric.

Setup

Let $y \in [0, 1]$ denote the share of an institution's operative policy attention devoted to extractive, eco-deficit activity, with $y \approx 0$ an eco-surplus posture and $y \approx 1$ an eco-deficit one. Let $\kappa > 0$ be the strength of external inducement toward extraction (market pressure, short-run competitive advantage), $I > 0$ the institution's responsiveness to that inducement, $\theta_s > 0$ an intrinsic threshold of institutional discipline, and $\beta > 0$ a reinforcement parameter capturing how self-sustaining an eco-deficit posture becomes once entered, through sunk investment, path-dependent infrastructure, or entrenched

incentive structures. Write $\mu := I\kappa - \theta_s$ and posit

$$\dot{y} = y(1 - y)(\beta y + \mu).$$

This is, deliberately, the same cubic nonlinearity Proposition 13.4 already analyzes for a population's growth dynamics in Chapter 13, applied here to a bounded policy-attention variable rather than a population level. We do not re-derive the linearized stability argument from scratch; it is identical in structure to the proof already given there, with y in place of X_t/K -style normalized population and μ in place of that chapter's own drift parameter.

Proposition 16.2 (Hysteresis in the eco-deficit transition). *Define $\kappa_U := \theta_s/I$ and $\kappa_L := (\theta_s - \beta)/I$, so that $\kappa_L < \kappa_U$ strictly whenever $\beta > 0$. Started at $y \approx 0$, a slow increase in κ leaves the institution at $y \approx 0$ until κ reaches κ_U , at which point y jumps to $y \approx 1$. Started at $y \approx 1$, a slow decrease leaves the institution at $y \approx 1$ until κ falls to the strictly lower κ_L , at which point y jumps back to $y \approx 0$.*

We prove Proposition 16.2 in Appendix C, reusing Proposition 13.4's equilibrium and stability analysis directly rather than repeating it. We verified the trichotomy and the hysteresis loop itself numerically, by direct simulation of $\dot{y}$ under a slow sweep of κ up and then back down, before stating the result.

Worked Example

At $I = 1$, $\theta_s = 1$, $\beta = 0.4$, $\kappa_U = 1.0$ and $\kappa_L = 0.6$, a hysteresis gap of 0.4. An institution at $y \approx 0$ under inducement $\kappa = 0.8$, within the bistable window, remains eco-surplus unless inducement actually reaches 1.0. But an institution already pushed to $y \approx 1$, by a temporary spike to $\kappa = 1.2$, remains eco-deficit even after inducement falls back to 0.8, releasing only once κ drops to 0.6. Observed at the identical inducement level $\kappa = 0.8$, the two institutions sit in opposite postures purely because of their history, not their current circumstances.

The parameter that governs the width of this gap, and hence the cost of reversal, is not the same parameter that governs where the gap sits. Raising the discipline threshold θ_s shifts κ_U and κ_L by the same amount, leaving the gap $\kappa_U - \kappa_L = \beta/I$ exactly unchanged; a more disciplined institution requires more inducement to enter an eco-deficit posture, but is no easier to bring back out of one once there. The exposure parameter I behaves differently: since the gap is β/I , a less exposed institution, smaller I , faces a strictly *wider* gap, not merely an unchanged one, doubling as I halves. Only β itself, the strength of self-reinforcement once an eco-deficit posture is entered, narrows the gap directly when reduced. Read against Vuong & Nguyen (2026)'s own framing, this gives a precise, checkable reason why a semiconducting-principle transition, from a growth-oriented eco-deficit trajectory to an eco-surplus one, is not expected to reverse smoothly or automatically once the inducement that caused it subsides: institutional path-dependence, not merely current market conditions, determines whether return is available at the inducement level currently observed.

This has a direct and somewhat uncomfortable message for policy design. A common assumption is that if market pressure or short term incentives caused an organization to slide into extractive, damaging behavior, then simply removing that pressure should let things return to normal. This result says that assumption is often false. Once an eco-deficit posture becomes embedded, through investment already made, infrastructure already built, or habits already formed, undoing it requires pulling the original pressure back much further than the level that caused the slide in the first place. In practice this means prevention deserves more weight than correction. It is easier and cheaper, by this model's own logic, to keep an institu-

tion from crossing the line than to bring it back afterward, even once the original cause of the crossing is gone.

16.5 Summary

This chapter formalized three claims implicit in the earth sufficiency paradigm of Vuong & Nguyen (2026), using apparatus this book had already built for other purposes. Corollary 16.1 shows that conservation priorities are bounded in number for any guaranteed level of commitment, a direct consequence of the probability simplex already used in Chapter 7. Proposition 16.1 shows that the average informational return per simultaneously held priority strictly declines as more are added, proved in Appendix C. Proposition 16.2 shows that the transition between an eco-deficit and an eco-surplus institutional posture is hysteretic: entering an eco-deficit posture requires inducement to rise to κ_U , but returning from one requires inducement to fall considerably further, to the strictly lower κ_L , reusing Chapter 13's own Allee-model technique rather than introducing new machinery. Chapter 17 closes the book by returning to its full set of claims, old and new, across the book as a whole.

Predictions, Scope, and Non-Claims

17.1 Testable Predictions

This section collects, in one place, the claims made across this book that are genuinely falsifiable, together with the status of each: proved as mathematics, checked qualitatively against external data, or offered only as untested speculation. Separating these three categories explicitly is, in our view, as important a contribution as any single proposition, since a book combining rigorous derivation with speculative extension risks having its speculative parts mistaken for its proved parts, or vice versa, if the two are not kept visibly apart.

Proved as Mathematics

The following are theorems, in the ordinary sense, true given the stated assumptions, and not themselves subject to empirical test, only to the question of whether their assumptions apply to any given real system.

Proposition 2.1, that the uniform distribution uniquely maximizes entropy; Proposition 7.1, the capacity bound on simultaneous value quanta; Proposition 8.1, the diagrammatic decomposition of free energy under relational interactions; Proposition 10.1 and Proposition 10.2, the closed-form elasticity of the Gibbs distribution and the near-inelasticity of high-conviction beliefs; and Proposition 11.1, the value formation theorem itself, that any nontrivial energy differentiation strictly and measurably lowers entropy relative to indifference. None of these depends on any claim about Bitcoin, markets, or any other real-world system; each is true of the mathematical objects, distributions, energies, and diagrams, as defined.

Chapters 13 and 14 add a second, independent family of proved results, resting on stochastic calculus rather than on the entropy apparatus above: Proposition 13.2, the valuation PDE derived via the Feynman-Kac connection; Proposition 14.3, the closed-form barrier-adjusted valuation; Proposition 14.5, the closed-form collapse probability; and Proposition 14.4, the overvaluation theorem itself. Like the entropy-based theorems, these are true given their stated assumptions, chiefly that the underlying process follows geometric Brownian motion with a genuinely absorbing barrier, regardless of whether any real ecological or financial system satisfies that assumption exactly. We verified the two closed-form formulas numerically against Monte Carlo simulation before stating them, a step beyond mathematical proof that this book's house practice treats as mandatory before any quantitative claim is put into print.

Chapter 15 adds a third family, smaller than the other two but distinguished by a different relationship to the rest of the book: Proposition 15.1, that the entropy ceiling on cultural possibility is non-increasing as biodiversity-derived informational quanta are lost, and Proposition 15.2, that the combinatorial space of distinct achievable cultural configurations collapses strictly faster than that decline itself. Unlike Chapters 13 and 14, whose connection

to the entropy apparatus of Parts II and III is explicitly thematic rather than technical, both of Chapter 15's propositions are direct applications of results proved earlier in this book for an unrelated purpose, Proposition 2.1's maximum-entropy identity and elementary combinatorics, to a new setting.

Checked Against External Data

Proposition 12.1, the horizon-dependent dominance prediction of Section 12.1, is the most fully developed empirical claim this book checks against evidence it did not itself collect. It predicts that short-horizon Bitcoin price movements should be more strongly associated with high-entropy, socially driven value dimensions than with low-entropy, structurally anchored ones, with the reverse pattern emerging as the horizon lengthens. Vuong et al. (2025b) test the short-horizon half directly, using this book's own SOV/AUT/SSV/HSV decomposition: only the social-signal predictor's 95% credible interval excludes zero. The broader published literature on Bitcoin attention and sentiment (Figá-Talamanca & Patacca 2019; Dastgir et al. 2019; Aysan et al. 2023) corroborates the same pattern using different proxies, though we repeat the caution of Section 12.1 that neither the direct study nor the broader literature tests the Gibbs-distribution machinery itself, or the prediction's longer-horizon half.

The overvaluation theorem of Chapter 14 is not, and could not straightforwardly be, checked against data in the same way: it is a mathematical consequence of the model's assumptions, not an empirical regularity, and Table 17.1 below records it as proved rather than checked for that reason. What connects it to evidence is the qualitative motivation Vuong (2021) supplies, that environmental and monetary value are observed not to exchange symmetrically, which the model formalizes but does not independently verify against real ecological time series.

Untested Extensions, Offered as Directions

Section 12.2's discontinuous mode-shift, Proposition 12.2, is a genuine theorem given a bi-modal energy landscape, but whether real markets exhibit such landscapes, and whether the shifting-baseline analogy drawn afterward correctly explains observed crash asymmetry, are both open empirical questions this book does not test. The suggestive connection drawn in Chapter 14's closing between the entropy apparatus of Part II and the continuous-time model of Chapters 13 and 14 belongs in the same category: offered, not derived. Chapter 15's propositions are proved, not untested, but its worked example's specific numerical mapping from population decline to the count $N(t)$ is explicitly flagged there as illustrative rather than measured. Section 12.3's distributional ratchet, Proposition A.3, is likewise proved rather than untested, but whether real cultural belief about wildlife actually moves along the linear path the worked example uses, rather than through some messier, feedback-laden process, is explicitly flagged there as an open empirical question the mathematics does not by itself settle.

Summary

Nothing new; this section has organized, rather than added to, the claims made throughout the book, so that a reader can see at a glance which results carry the weight of proof, which carry the weight of an external check, and which are offered only as directions for further work. The remainder of this chapter takes up the reverse task: stating plainly what this book does not claim, even where its formalism might tempt a stronger reading than the evidence supports.

Claim	Status	Chapter
Entropy, capacity, diagrammatic, elasticity, and value theorems	Proved	2, 7, 8, 10, 11
Bitcoin horizon-dependent dominance (short-horizon half)	Checked directly against real market data	12
Bubble/crash mode discontinuity	Proved (conditional on landscape)	12
Distributional shifting-baseline ratchet	Proved	A
Bird conservation belief drift	Proved result, illustrative application	12
Barrier valuation, collapse probability, overvaluation theorem	Proved (conditional on GBM-with-barrier model), numerically verified	13, 14
Entropy ceiling, combinatorial diversity collapse	Proved (direct application of Ch. 2 and Ch. 6 apparatus)	15

Table 17.1: Summary of this book's claims by evidentiary status.

17.2 Scope and Non-Claims

We close, as the preface promised, by stating plainly what this book has not established, alongside a brief final word on what it has.

Reconstruction Through the Applications Chapter, Extension in Part III

Every mathematical tool used through Chapter 12, Lagrange multipliers, the Gibbs distribution, Kullback-Leibler divergence, the exponential formula and its diagrammatic bookkeeping, is standard, proved here from first principles for the reader's convenience, not presented as new; the contribution through that point lies entirely in the mapping between this standard apparatus and the qualitative claims of Vuong et al. (2025a). Chapters 13 through 15 are different in kind, not in rigor: the mathematical tools they use, stochastic calculus in Chapters 13 and 14, the maximum entropy principle and elementary combinatorics, already proved earlier in this book, in Chapter 15, are likewise standard and not our own invention, but their application, an absorbing model of ecological value with its closed-form overvaluation result, and a discrete model of biodiversity-derived cultural information with its entropy-ceiling and combinatorial-collapse results, is not a reconstruction of anything in Vuong et al. (2025a) or elsewhere in this series; it is this book's own extension, motivated by Vuong (2021)'s semi-conducting principle and by the cultural ecosystem services literature surveyed in Chapter 15, rather than derived from either. Readers should not come away from this book believing that granular interaction thinking theory has been mathematically proven correct as a general account of value, nor that Part III's models have been shown to describe any actual ecosystem or culture; what has been shown, across this book as a whole, is narrower and more specific than either of those stronger claims.

The Relationship to Appendix A

Section 12.2 drew an explicit analogy between the path-dependence of market bubbles and crashes and the shifting-baseline result proved from scratch in Appendix A. We restate here, for clarity, that this connection is an interpretive analogy, not a theorem, and not a claim that one formal apparatus derives from or reduces to the other. The appendix's result is a deterministic construction concerning a sequence of revisions compared against a moving reference point; the entropy and temperature apparatus of this book is a probabilistic construction concerning distributions over informational quanta. Both describe an asymmetry between gradual accumulation and abrupt correction, and we find the resemblance suggestive enough to note explicitly, but the two are independent formalizations, developed for different purposes, and a reader should not treat Section 12.2's analogy as extending the appendix's proof to markets, or as borrowing its certainty for this book's own, considerably less certain, application sections.

What Remains Untested

Section 17.1 already separated this book's claims by evidentiary status; we do not repeat that separation here except to underline its most important consequence. Exactly one claim, the Bitcoin prediction of Section 12.1, is checked against evidence the book did not itself collect, and its short-horizon half is checked directly, using this book's own SOV/AUT/SSV/HSV decomposition rather than a looser proxy for it, by Vuong et al. (2025b)'s Bayesian regression; the longer-horizon half of the same prediction and the underlying energy-and-temperature mechanism remain untested by that evidence. The overvaluation theorem of Chapter 14 and the entropy-ceiling and combinatorial-collapse results of Chapter 15, though proved with the same rigor as the entropy theorems and, in the former case, numerically verified against simulation, have not been tested against any real ecological, financial, or cultural time series; they establish what follows *if* their respective models are the right ones, not that any particular ecosystem or culture is well described by such a model.

Limitations of the Formalism Itself

Several modeling choices made across this book are worth naming as choices rather than necessities. The decision to truncate interaction terms at pairwise or triple-wise order, in Chapter 8, is a simplification with no principled upper bound behind it. The choice of a Gibbs distribution specifically, rather than some other exponential family, to formalize the maximum entropy principle of Chapter 3, is standard and well motivated but not the only possible choice. The mapping of Bitcoin's four value dimensions onto low- and high-energy groups in Section 12.1 is an interpretive assignment, not derived from the formalism itself. The treatment of ecological collapse as a strictly absorbing barrier in Chapters 13 and 14, flagged explicitly in both chapters as a modeling idealization, overstates the literal permanence of real ecological tipping points in exchange for making the semiconducting principle's asymmetry mathematically exact within the model. And Chapter 15's treatment of extinction as a strictly one-directional pure death process, and of all biodiversity-derived informational quanta as interchangeable within a single count $N(t)$, makes the same trade, tractability and exactness purchased at the cost of ignoring partial, real-world countervailing mechanisms such as rewilding and genome banking, which that chapter names explicitly rather than leaving implicit.

Key Formulas at a Glance

Table 17.2 collects Part III's most frequently reused formulas in one place, for a reader who wants to apply a specific result without searching back through its full derivation. Every entry restates a result already proved earlier in this book; nothing here is new.

Quantity	Formula	Source
Collapse probability	$P(\tau_B \leq T) = \Phi\left(\frac{b-y_0-\nu T}{\sigma\sqrt{T}}\right) + e^{-2\nu(y_0-b)/\sigma^2} \Phi\left(\frac{b-y_0+\nu T}{\sigma\sqrt{T}}\right)$	Prop. 14.5
Barrier-adjusted valuation	$V = e^{(\mu-r)T} [x\Phi(d_1) - B^{1+2\mu/\sigma^2} x^{-2\mu/\sigma^2} \Phi(d_2)]$	Prop. 14.3
Overvaluation gap	$V_0 - V = e^{-rT} E[X_T \mathbf{1}_{\{\tau_B \leq T\}}] > 0$	Prop. 14.4
Correction factor	$\kappa = V/V_0$ (Table 14.4)	Sec. 14.10
Entropy ceiling	$H_{\max}(t) = \log_2 N(t)$	Prop. 15.1
Minimum richness for target ceiling	$N(t) \geq 2^{H_{\text{target}}}$	Cor. 15.1
Cultural configuration count	$C(t) = \binom{N(t)}{k}$	Prop. 15.2
Configuration elasticity	$d \log C / d \log N \approx k$ for $k \ll N$	Prop. 15.3

Table 17.2: Part III's key formulas. $\nu = \mu - \sigma^2/2$; $y_0 = \ln x$, $b = \ln B$; d_1, d_2 as in Proposition 14.3. Each row restates, rather than re-derives, its cited source.

Part III's mathematics is built to answer five questions a reader could, in principle, ask about any real situation involving potential ecological or informational loss, without needing to work through a single equation. We walk through them here in turn, each pointing back to the specific result that gives it formal content, as a practical entry point for a reader who wants to use this part of the book as a lens rather than as a derivation.

The first question is what information is being lost, and it is worth asking twice. Chapter 15's entropy ceiling (Proposition 15.1) and its combinatorial sharpening (Proposition 15.3) formalize the question about the raw space of future possibility itself. Proposition 11.1 in Chapter 11 formalizes the separate question of how narrowly attention concentrates within whatever space remains, a distinction Chapter 15 illustrates directly using the Innovation Curse literature already cited in Chapter 14.

The second question is what interactions disappear when a component disappears. This is Chapter 8's question in its original form, that value and meaning emerge from interaction rather than residing in isolated quanta. Section 15.8 in Chapter 15 asks it again specifically for aggregate ecological stability, and reports, honestly, that the answer is not settled.

The third question is whether the loss is reversible. Proposition 15.5 and Proposition 14.12 give this a precise negative answer in the base case. The γ and δ parameters of Section 14.7 in Chapter 14, and of Chapter 15 respectively, turn how reversible into an estimable number rather than a binary claim.

The fourth question is whether conventional valuation recognizes the possibility of collapse at all. This is Proposition 14.4, the Overvaluation Theorem, directly: a valuation that ignores collapse risk is not a conservative estimate but the largest number the true value could be, and Table 14.4 gives a concrete correction once distance from collapse and volatility are estimated.

The fifth and final question is what future possibilities disappear along with today's loss. Proposition 15.3's elasticity result answers this quantitatively: losses compound combinatorially in the number of things being simultaneously valued, not merely in proportion to what is lost.

None of these five questions requires the reader to follow any proof in Part III; each is answerable, at least in direction, once the relevant quantities, a distance to a threshold, a volatility, a recovery fraction, a richness count, are estimated for the situation at hand. The proofs exist so that the answers are not merely plausible-sounding but exact given those inputs.

A Closing Word

We began this book by asking what the smallest, most standard body of mathematics would be that could make Vuong et al. (2025a)'s qualitative claims precise, and closed by asking a second, harder question: whether the same style of formal thinking could say something new about problems the original paper does not address at all, the valuation of assets facing genuinely irreversible loss, and the informational cost of biodiversity loss to cultural creativity. The first question produced a reconstruction; the second produced, we hope, two modest genuine extensions, gathered in Part III, one grounded in Vuong (2021)'s semiconducting principle and built from standard stochastic calculus, the other grounded in the cultural ecosystem services literature and built by turning this book's own entropy and combinatorial apparatus on a new target. Whether any part of this book earns its place as a useful lens, the entropy formalism, the irreversibility model, or the biodiversity model, is a question we have tried to leave open rather than answer for the reader. No single book of this kind resolves the ecological challenges it touches; we offer this one in the more modest hope that incremental advances in ecological valuation and collapse-risk modeling, ours among many others, can collectively sharpen how those challenges are understood. We hope it has been useful, and we hope we have made clear, throughout, exactly where our own confidence ends and where a future reader's own judgment ought to begin.

If any single spirit deserves the last word after so much mathematics, we would choose Zhuangzi's: not mastery over nature but ease within it, content with what suffices rather than straining, at whatever cost, for what does not.

Grigory Perelman gives that spirit a living example, not only a symbolic one. In 2002 he solved the Poincaré conjecture, a problem mathematicians had carried for a century, and released the proof freely rather than through the usual channels. He declined the Fields Medal in 2006, and in 2010 declined the million-dollar Millennium Prize that came with solving it, telling a reporter only that he already had everything he needed (Harding 2010). We do not know whether Perelman ever read a page of the Zhuangzi, and we make no claim that he did. What his choice shows, on its own terms, is that this book's closing sentiment is not only an old one. On rare occasion, it is still how a person actually chooses to live.

Bibliography

- Apostol, T. M. (1967). *Calculus, Volume 1: One-variable calculus, with an introduction to linear algebra* (2nd ed.). Wiley.
- Arrow, K. J., & Fisher, A. C. (1974). Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics*, 88(2), 312–319.
- Aysan, A. F., Muğaloğlu, E., Polat, A. Y., & Tekin, H. (2023). Whether and when did Bitcoin sentiment matter for investors? Before and during the COVID-19 pandemic. *Financial Innovation*, 9, Article 124.
- Björk, T. (2009). *Arbitrage theory in continuous time* (3rd ed.). Oxford University Press.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- Boyce, W. E., & DiPrima, R. C. (2017). *Elementary differential equations and boundary value problems* (11th ed.). Wiley.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory* (2nd ed.). Wiley-Interscience.
- Dastgir, S., Demir, E., Downing, G., Gozgor, G., & Lau, C. K. M. (2019). The causal relationship between Bitcoin attention and Bitcoin returns: Evidence from the copula-based Granger causality test. *Finance Research Letters*, 28, 160–164.
- Dixit, A. K., & Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.
- Feller, W. (1968). *An introduction to probability theory and its applications* (Vol. 1, 3rd ed.). Wiley.
- Figá-Talamanca, G., & Patacca, M. (2019). Does market attention affect Bitcoin returns and volatility? *Decisions in Economics and Finance*, 42(1), 135–155.
- Griffiths, D. J., & Schroeter, D. F. (2018). *Introduction to quantum mechanics* (3rd ed.). Cambridge University Press.
- Harding, L. (2010, March 23). Grigory Perelman, the maths genius who said no to \$1m. *The Guardian*. <https://www.theguardian.com/world/2010/mar/23/grigory-perelman-rejects-1m-dollars>
- Jaynes, E. T. (2003). *Probability theory: The logic of science*. Cambridge University Press.
- Karlin, S., & Taylor, H. M. (1981). *A second course in stochastic processes*. Academic Press.

- Kittel, C., & Kroemer, H. (1980). *Thermal physics* (2nd ed.). W.H. Freeman.
- Lenton, T. M., Rockström, J., Gaffney, O., Rahmstorf, S., Richardson, K., Steffen, W., & Schellnhuber, H. J. (2019). Climate tipping points, too risky to bet against. *Nature*, 575(7784), 592–595.
- May, R. M. (1972). Will a large complex system be stable? *Nature*, 238, 413–414.
- May, R. M. (1977). Thresholds and breakpoints in ecosystems with a multiplicity of stable states. *Nature*, 269(5628), 471–477.
- Nguyen, M.-H., & Ho, M.-T. (2026). The absurdist approach to unveiling possible paradoxical thinking for innovative socio-psychological research. *MethodsX*, 16, Article 103910.
- Nguyen, M.-H., & Jones, T. E. (2022a). Predictors of support for biodiversity loss countermeasure and bushmeat consumption among Vietnamese urban residents. *Conservation Science and Practice*, 4(12), e12822.
- Nguyen, M.-H., & Jones, T. E. (2022b). Building eco-surplus culture among urban residents as a novel strategy to improve finance for conservation in protected areas. *Humanities and Social Sciences Communications*, 9, Article 426.
- Nguyen, M.-H., & Khuc, V. Q. (2026). Bird conservation through the lens of ‘the Bird Village’. *Pacific Conservation Biology*, 32, Article PC26032.
- Nguyen, M.-H., Gagnon, V., Tran, T. T., & Tran, T. M. A. (2025b). Innovation curse: The wastefulness of technologies believed to mitigate climate change. *Humanities and Social Sciences Communications*, 12, Article 1817.
- Øksendal, B. (2003). *Stochastic differential equations: An introduction with applications* (6th ed.). Springer.
- Orwell, G. (1945). *Animal farm*. Secker & Warburg.
- Pathria, R. K., & Beale, P. D. (2011). *Statistical mechanics* (3rd ed.). Academic Press.
- Pindyck, R. S. (2000). Irreversibilities and the timing of environmental policy. *Resource and Energy Economics*, 22(3), 233–259.
- Reif, F. (1965). *Fundamentals of statistical and thermal physics*. McGraw-Hill.
- Rudin, W. (1976). *Principles of mathematical analysis* (3rd ed.). McGraw-Hill.
- Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3), 379–423.
- Tilman, D. (1996). Biodiversity: Population versus ecosystem stability. *Ecology*, 77(2), 350–363.
- Van Huu, N., Vuong, Q.-H., & Tran, M.-N. (2005). Central limit theorem for functional of jump Markov processes. *Vietnam Journal of Mathematics*, 33(4), 443–461.
- Van Huu, N., & Vuong, Q.-H. (2007). On the martingale representation theorem and on approximate hedging a contingent claim in the minimum deviation square criterion. In R. Jeltsch, T. Li, & I. H. Sloan (Eds.), *Some topics in industrial and applied mathematics* (pp. 134–151). World Scientific.
- Vuong, Q.-H. (2021). The semiconducting principle of monetary and environmental values exchange. *Economics and Business Letters*, 10(3), 284–290.

- Vuong, Q.-H. (2022). *Wild wise weird*. <https://www.amazon.com/dp/B0G2F19STS>
- Vuong, Q.-H., La, V.-P., & Nguyen, M.-H. (2025a). Informational entropy-based value formation: A new paradigm for a deeper understanding of value. *Evaluation Review*, 50(3), 516–540.
- Vuong, Q.-H., La, V.-P., & Nguyen, M.-H. (2025b). Bayesian probabilistic exploration of Bitcoin informational quanta and interactions under the GITT-VT paradigm. *arXiv preprint arXiv:2511.17646*.
- Vuong, Q.-H., & Nguyen, M.-H. (2025). On nature quotient. *Pacific Conservation Biology*, 31, Article PC25028.
- Vuong, Q.-H., & Nguyen, M.-H. (2026). From human sufficiency to earth sufficiency: The role of nature quotient cultivation. *The Anthropocene Review*. Advance online publication. <https://doi.org/10.1177/20530196261464206>
- Wilf, H. S. (2006). *generatingfunctionology* (3rd ed.). A K Peters.
- World Wide Fund for Nature. (2022). *Living planet report 2022: Building a nature-positive society*. WWF.
- Zhuangzi. (2013). *The complete works of Zhuangzi* (B. Watson, Trans.). Columbia University Press.

Appendix A

Self-Contained Notes on Detection, Drift, Belief, and Extraction

Section 12.2 draws an analogy between the path-dependence of market bubbles and crashes and a general phenomenon: a checking procedure that compares each new state only to the immediately preceding one, rather than to some fixed original, can fail to catch a large cumulative change even while flagging every individual step correctly. This appendix derives that phenomenon from scratch, as a small, fully self-contained result in elementary real analysis, so that Section 12.2's analogy rests on a derivation given in full within this book rather than on an external citation. A second, independent result is included alongside it: a second channel by which a single large change can escape detection, through suspension of the checking act itself rather than through sub-threshold accumulation, which Chapter 13 draws on as an analogy for gaps in ecological monitoring. A third result extends the first from a sequence of real numbers to a sequence of belief distributions, using total variation distance in place of absolute difference, which Section 12.3 in Chapter 12 draws on for a case study in cultural belief drift around wildlife conservation. A fourth, independent result concerns a different question, not detection but destination: given a resource extracted at a constant proportional rate, where the process ends up, which Chapter 13 draws on directly for natural-resource overexploitation. All four results are self-contained, requiring nothing beyond elementary real analysis and basic probability. A different result once included here, on how contracting a distribution's support lowers its entropy ceiling, has been removed as redundant with Chapter 15's own Proposition 15.1, which proves the identical fact directly in the setting that result is actually used for.

A.1 Setup

Setup

Let $c_0, c_1, c_2, \dots, c_N \in \mathbb{R}$ be a sequence of states, where c_n is the state after its n -th revision and c_0 is the original. Write $d_n := |c_n - c_{n-1}|$ for the size of the n -th revision. Suppose a checking procedure enforces a perceptibility threshold $\tau > 0$: revision n is flagged if and only if $d_n \geq \tau$. The structural feature this appendix studies is that the comparison generating this flag is always made against the most recently recorded state c_{n-1} , never against the original c_0 .

This setup is deliberately abstract; nothing in it refers to markets, minds, or any other substantive domain. The point of the abstraction is that the result below applies wherever this

specific structure, comparison against a moving rather than a fixed reference point, is present, which is precisely why Section 12.2 can invoke it as an analogy for market behavior without claiming that markets literally implement the sequence c_n in any specific way. This general phenomenon, a moving-baseline comparison that lets cumulative drift escape detection, is also referred to elsewhere in this series as Orwellian Calculus. (The Orwellian Calculus can be considered our in-house machinery for some socio-cultural and psychological examinations.)

The clearest illustration of this structure available in fiction is the repeated revision of the Seven Commandments painted on the barn wall in Orwell (1945). Each individual amendment, "No animal shall drink alcohol" acquiring the qualifier "to excess," "No animal shall sleep in a bed" acquiring "with sheets," is small enough that the animals who half-remember the previous wording do not, individually, register it as a violation; each revision is compared only to what the animals recall of the version immediately before it, never against a fixed, independently recorded original. The cumulative effect, by the novel's end, is a complete reversal of the founding principles, achieved without any single step ever being large enough on its own to prompt the challenge a comparison against the original wording would have produced. We use this example here for orientation only, not as a claim that the novel's events instantiate Proposition A.1 formally; the proposition itself, proved below, requires nothing beyond the abstract sequence c_n already stated.

A.2 The Result

Proposition A.1 (Shifting-Baseline Ratchet (SBR)). *For any target cumulative distortion $D > 0$, however large, and any perceptibility threshold $\tau > 0$, however small, there exists a finite N and a sequence of revisions $d_1, \dots, d_N$, each satisfying $d_n < \tau$, so that no single revision is ever flagged, whose cumulative effect is nonetheless exactly D : $|c_N - c_0| = D$. Moreover, since the construction below uses same-signed revisions, there is no cancellation between steps: $|c_N - c_0| = \sum_{n=1}^N d_n$.*

Proof. Fix $D > 0$ and $\tau > 0$, and choose any integer $N > D/\tau$. Set $d_n = D/N$ for every $n = 1, \dots, N$; then $d_n = D/N < \tau$ by the choice of N , so no revision is flagged under the perceptibility threshold. Take every revision in the same direction, so that $c_n = c_{n-1} + d_n$ for each n . Then

$$c_N - c_0 = \sum_{n=1}^N (c_n - c_{n-1}) = \sum_{n=1}^N d_n = N \cdot \frac{D}{N} = D,$$

a telescoping sum in which every intermediate term cancels except the first and last. Since all $d_n > 0$, $|c_N - c_0| = c_N - c_0 = D = \sum_{n=1}^N d_n$, giving both claims. Since D was arbitrary, no finite ceiling on $|c_N - c_0|$ can follow from the threshold τ alone, however small τ is fixed to be. $\square$

The construction is deliberately the simplest one that works: equal-sized, same-signed steps. Nothing about the proposition requires equal steps; any sequence of positive d_n summing to D with each term below τ would do, and the equal-step case is chosen only because it makes the required N easy to state explicitly, $N > D/\tau$.

Worked Example

Suppose a checking procedure flags any single revision of size $\tau = 0.05$ (five hundredths of some unit) or more, and suppose the target cumulative distortion is $D = 3$. Proposition A.1 guarantees that $N = 61$ revisions of size $d_n = 3/61 \approx 0.0492$ each, comfortably below τ , suffice to move the state by the full amount $D = 3$, with every single step passing the threshold check. A procedure that instead compared every state to the

fixed original c_0 , rather than to the previous state c_{n-1} , would catch this immediately: $|c_N - c_0| = 3 \geq \tau$ on the very first comparison made against c_0 . The entire effect is attributable to which baseline the comparison uses, not to any failure of the threshold τ itself, which correctly catches every single step exactly as designed.

A.3 Why the Threshold Cannot Be Fixed

A natural objection is that the checking procedure could simply choose τ small enough to catch any distortion it cares about. The following corollary shows this does not work: no fixed τ , however small, bounds the achievable gap between what a step-by-step check permits and what actually accumulates.

Corollary A.1. Define $G_N := |c_N - c_0| - \max_n d_n$, the amount by which the true cumulative distortion exceeds the size of the single largest step. Then G_N is not bounded by any function of τ alone: for any target $G_0 > 0$ and any $\tau > 0$, there is a finite N and a sequence with $\max_n d_n < \tau$ for which $G_N \geq G_0$.

Proof. Fix $\tau > 0$ and $G_0 > 0$. Choose any $D > G_0$, and apply Proposition A.1's construction with this D and the given τ : $d_n = D/N$ for $N > D/\tau$, giving $\max_n d_n = D/N$ and $|c_N - c_0| = D$. Then

$$G_N = D - \frac{D}{N} = D \left(1 - \frac{1}{N}\right).$$

As $N \rightarrow \infty$ (holding D fixed and letting the number of steps grow, which only requires shrinking each d_n further, still consistent with $d_n < \tau$), $G_N \rightarrow D > G_0$, while $\max_n d_n = D/N \rightarrow 0$. In particular, choosing N large enough simultaneously satisfies $\max_n d_n < \tau$ and $G_N \geq G_0$, since D was fixed above G_0 from the start. $\square$

Corollary A.1 is the precise sense in which sensitivity at the level of individual steps and control over cumulative drift are different things: making τ smaller only forces the adversarial construction to use more, smaller steps, which it can always do, rather than closing off the construction entirely.

A.4 A Second Channel: Suspending the Comparison Itself

Proposition A.1 accounts for undetected drift through accumulation: many revisions, each individually below threshold, sum to an arbitrarily large total. This leaves untouched a different, in some ways simpler, case: a single revision at or above the threshold τ , which ordinary detection would catch with certainty, yet which escapes detection anyway because the comparison itself is not carried out on that occasion. This section extends the setup to cover exactly that case.

Setup

Extend the setup of Section A.1 with a Bernoulli-distributed suspension indicator $S_n \in \{0, 1\}$ for each revision n , with $\Pr(S_n = 1) = \sigma \in [0, 1]$, independent of d_n . $S_n = 1$ means the comparison act itself is not performed at step n , regardless of how large d_n is; $S_n = 0$ means the comparison is performed as usual, so revision n is flagged if and only if $d_n \geq \tau$.

Proposition A.2 (Comparison-suspension escape). *Under this extended setup, revision n escapes detection if $d_n < \tau$ or $S_n = 1$. In particular, for a single-step revision ($N = 1$) of any size $D \geq \tau$, the probability of escaping detection is exactly σ , independent of D and of τ ; under the baseline setup alone ($\sigma = 0$), the same revision is detected with certainty.*

Proof. Detection requires both that the comparison act is performed ($S_n = 0$) and that the resulting discrepancy clears the threshold ($d_n \geq \tau$); escape is the complement of this conjunction, $\{d_n < \tau\} \cup \{S_n = 1\}$. For $D \geq \tau$ and $N = 1$, $d_1 = D \geq \tau$, so the event $d_1 < \tau$ is false, and escape occurs exactly when $S_1 = 1$, which has probability σ by construction. When $\sigma = 0$, $\Pr(S_1 = 1) = 0$, so detection occurs with probability 1. $\square$

A direct consequence, conditional on the revision already clearing the threshold, is that the detection probability is $\Pr(\text{detection} \mid d \geq \tau) = 1 - \sigma$, strictly decreasing in σ and, notably, independent of how far above τ the revision sits. Within this model, a larger discrepancy is not intrinsically harder to detect; whether detection occurs depends on whether the comparison is performed at all, not on the discrepancy's size once that size already exceeds threshold. Chapter 13 draws on this result directly, as an interpretive analogy for gaps in ecological monitoring: a collapse-sized change in a population or biomass index can escape detection not because it was too gradual to notice, Proposition A.1's channel, but because the observation that would have caught it, a survey, a satellite pass, a census, simply did not occur at the relevant moment.

Worked Example

Proposition A.1's accumulation channel and Proposition A.2's suspension channel cover complementary cases along the same axis. Fix $D = 10$ and $\tau = 0.5$ as in the earlier worked example. The graduated case, $N = 21$ equal steps of size $d_n \approx 0.476 < \tau$, escapes detection with probability 1 even at $\sigma = 0$, since no step ever reaches threshold; this is Proposition A.1 operating alone. The opposite extreme, $N = 1$, $d_1 = D = 10 \gg \tau$, would be detected with certainty at $\sigma = 0$; it escapes only through Proposition A.2's channel, with escape probability exactly σ regardless of how large D is. The two propositions are not competing explanations of the same phenomenon but accounts of two distinct failure modes, one exploiting step size, the other exploiting whether checking happens at all.

A.5 A Distributional Generalization

Proposition A.1 concerns a sequence of real numbers. Section 12.3 in Chapter 12 applies the same underlying phenomenon to a sequence of *beliefs*, probability distributions over some fixed set of outcomes, drifting under generational or social reinforcement. This section supplies the small amount of additional machinery that application needs, replacing the real-number distance $|c_n - c_{n-1}|$ with a standard distance between distributions, and shows that the same telescoping construction goes through essentially unchanged.

Setup

For distributions P, Q on a finite set, the total variation distance is $TV(P, Q) := \frac{1}{2} \sum_x |P(x) - Q(x)|$, a genuine metric: symmetric, zero only when $P = Q$, bounded in $[0, 1]$, and satisfying the triangle inequality $TV(P, R) \leq TV(P, Q) + TV(Q, R)$. Let $P_0, P_1, \dots, P_N$ be a sequence of belief distributions, with a checking procedure flagging revision n if and only if $TV(P_n, P_{n-1}) \geq \tau$, the comparison always made against P_{n-1} ,

never against P_0 .

Proposition A.3 (Distributional shifting-baseline ratchet). *For any target total divergence $D \in (0, 1]$ and any threshold $\tau > 0$, there exists a finite N and a sequence of belief distributions $P_1, \dots, P_N$, each satisfying $TV(P_n, P_{n-1}) < \tau$, so that no single revision is ever flagged, whose cumulative effect is nonetheless exactly D : $TV(P_N, P_0) = D$.*

Proof. Fix distributions P_0 and P^* with $TV(P_0, P^*) = D$, and choose any integer $N > D/\tau$. Define the straight-line mixture path $P_n := P_0 + \frac{n}{N}(P^* - P_0)$ for $n = 0, \dots, N$, a valid probability distribution at each step since it is a convex combination of P_0 and P^* . Then $P_n - P_{n-1} = \frac{1}{N}(P^* - P_0)$ for every n , a fixed vector independent of n , so

$$TV(P_n, P_{n-1}) = \frac{1}{2} \sum_x \left| \frac{1}{N}(P^*(x) - P_0(x)) \right| = \frac{1}{N} \cdot TV(P_0, P^*) = \frac{D}{N} < \tau$$

by the choice of N , so no revision is flagged. Because the path is linear, the triangle inequality holds with equality along it: $\sum_{n=1}^N TV(P_n, P_{n-1}) = N \cdot D/N = D$, and $TV(P_N, P_0) = TV(P^*, P_0) = D$ directly, since $P_N = P^*$ by construction. $\square$

We verified this construction numerically, for a concrete two-outcome case, before stating it here: with $P_0 = (0.6, 0.4)$, $P^* = (0.05, 0.95)$, $\tau = 0.05$, $D = TV(P_0, P^*) = 0.55$, twelve equal steps of size $D/12 \approx 0.046$ each, comfortably below τ , move the belief the full distance, exactly as Proposition A.3 guarantees.

The result is deliberately a close cousin of Proposition A.1 rather than a different theorem: the same moving-baseline structure, the same telescoping construction, with total variation distance standing in for the absolute value of a real-number difference because it satisfies the same triangle inequality that made the original proof work. One difference is worth noting rather than glossing over: TV is bounded above by 1, where the real-number ratchet's cumulative drift D could be made arbitrarily large. A belief cannot drift past complete reversal, since $TV = 1$ already means the two distributions share no support; what Proposition A.3 shows is that this maximal drift, a belief moving from strong confidence in one direction to strong confidence in the opposite direction, can be reached via steps that are each, individually, imperceptible.

A.6 The Extraction Limit

The two results above concern how a comparison can be defeated. This section concerns a different, complementary question: given a resource extracted at some constant proportional rate, justified at each occasion as modest and temporary, where does the process end up. Section 13.2 in Chapter 13 applies the result to natural-resource overexploitation directly.

Setup

Let $X_n \geq 0$ denote a resource stock remaining after n rounds of extraction, with each round removing a fixed proportion $\rho \in (0, 1]$ of whatever currently remains: $X_{n+1} = (1 - \rho)X_n$, X_0 given.

Proposition A.4 (The Extraction Limit). *For any constant proportional extraction rate $\rho \in (0, 1]$, however small, $X_n = (1 - \rho)^n X_0 \rightarrow 0$ as $n \rightarrow \infty$. The continuous-time analogue, $dX/dt = -\rho X(t)$, has closed-form solution $X(t) = X_0 e^{-\rho t}$, with $X^* = 0$ the unique, globally*

attracting equilibrium for every $\rho > 0$.

Proof. The discrete recursion is linear and solved by direct iteration: $X_1 = (1 - \rho)X_0$, $X_2 = (1 - \rho)X_1 = (1 - \rho)^2 X_0$, and by induction $X_n = (1 - \rho)^n X_0$. Since $0 \leq 1 - \rho < 1$ for $\rho \in (0, 1]$, $(1 - \rho)^n \rightarrow 0$ as $n \rightarrow \infty$, so $X_n \rightarrow 0$. For the continuous case, separating variables in $dX/dt = -\rho X$ gives $\int dX/X = -\rho \int dt$, so $\ln X(t) = \ln X_0 - \rho t$, and exponentiating gives $X(t) = X_0 e^{-\rho t}$; direct substitution confirms this solves the ODE. Global attraction of $X^* = 0$ follows immediately, since $X(t) \rightarrow 0$ as $t \rightarrow \infty$ for every $\rho > 0$ regardless of $X_0 > 0$, and $X^* = 0$ is the only fixed point of $-\rho X = 0$. $\square$

We verified both the discrete and continuous cases numerically before stating them: two hundred rounds of extraction at $\rho = 0.01$ leave 13.4% of the original stock, at $\rho = 0.05$ leave 0.35%, and at $\rho = 0.20$ leave a negligible fraction, in each case matching $(1 - \rho)^{200} X_0$ exactly; the continuous solution matched a numerical ODE integration to within 10^{-6} .

Worked Example

The qualitative content survives the arithmetic: it is not the size of ρ that determines the long-run outcome, only the speed of arriving there. A rate of $\rho = 0.01$, extracting one percent of the remaining stock each round and easily defended as negligible on any single occasion, still drives X_n below 1% of its original value within 459 rounds, and toward zero without limit thereafter. There is no constant proportional extraction rate, however small, whose long-run destination is anything other than exhaustion.

A.7 Connection to Section 12.2

Section 12.2 draws an analogy between Proposition A.1 and the observation that a market's baseline price expectation can drift substantially during a bubble's build-up without any single day's move being large enough to alarm observers comparing that day only to the day before. The correspondence is exactly the moving-baseline structure formalized here: if c_n is read as a market's implicit reference price after n trading periods, and each day's move is compared only to the previous day's closing price rather than to some independently fixed fair-value anchor, then Proposition A.1 shows that no per-step threshold, however finely tuned, bounds how far the reference price can drift from where it started. We repeat, as Section 12.2 and Chapter 17 both already state, that this is offered as an analogy illustrating a shared structural feature, not as a theorem about actual market microstructure; nothing in this appendix's abstract sequence c_n has been shown to correspond to any specific, measurable market quantity.

A.8 Summary

Comparing each new state only to its immediate predecessor, rather than to a fixed original, allows an arbitrarily large cumulative distortion to accumulate through steps that individually never trip a perceptibility threshold, however small that threshold is set (Proposition A.1), and no fixed threshold can close this gap (Corollary A.1). A single large change can escape detection through an entirely different channel, suspension of the checking act itself rather than sub-threshold accumulation (Proposition A.2), and the two channels are complementary rather than competing, covering opposite ends of the same construction. The same moving-baseline structure extends from real numbers to belief distributions once absolute difference is replaced by total variation distance (Proposition A.3), with the one substantive difference that total variation distance is bounded, capping the maximal undetected drift at complete

reversal rather than leaving it unbounded. A fourth, independent result changes the question from detection to destination: any constant proportional extraction rate, however small, drives a resource stock to exhaustion in the long run, with an exact closed-form solution in both discrete and continuous time (Proposition A.4). All four results are proved from elementary real analysis and basic probability alone; the entire appendix uses no mathematics beyond what Chapter 2 already assumes as background.

Appendix *B*

A Central Limit Theorem for Extinction Counts Under Regime-Switching Stress

Proposition 15.8 in Chapter 15 gives the exact mean and variance of the coupled extinction count $D(T \wedge \tau_B)$ for any horizon T , but two numbers are not a distribution. This appendix asks the natural next question, whether $D(T)$ is approximately normally distributed for large T , and shows that answering it honestly requires setting Chapter 15's own driving process aside rather than pushing it further.

The reason is structural, not a matter of difficulty. A central limit theorem for a time-averaged functional of a Markov process asks how that functional behaves as $T \rightarrow \infty$, and results of this kind, going back classically to Doeblin, require the underlying process to keep revisiting the same states indefinitely, a property variously called recurrence or ergodicity. Chapter 15's driver, GBM absorbed at a barrier B , is built to do the opposite: once X_t hits B , the process stops, permanently, and there is no long-run, repeat-visiting behavior for a limit theorem to describe. Asking for a $T \rightarrow \infty$ central limit theorem of an absorbing process is not a hard version of Proposition 15.8's question; it is a different question, about a process for which the relevant limit does not exist in the same sense.

Van Huu et al. (2005) prove a central limit theorem of exactly this kind for a different class of process: jump homogeneous Markov processes, which move by jumping between states in a finite or countable state space, rather than by continuous diffusion. This appendix builds a small, self-contained instance of their setting, a stress regime that jumps between a small number of discrete states rather than diffusing continuously, and uses it to give the extinction count a genuine asymptotic distribution. We regard this as a second, complementary model to Chapter 15's own, not a refinement of it; Section B.4 below is explicit about where the two agree and where they do not.

B.1 Setup

Setup

Let $Y(t) \in \{0, 1\}$ be a continuous-time Markov chain (a jump process) representing an ecological stress regime, 0 for low stress and 1 for high stress, with transition rate q_{01} from state 0 to state 1 and q_{10} from state 1 to state 0. Let the extinction intensity be $\varphi(Y(t))$, taking the value λ_{low} in state 0 and $\lambda_{\text{high}} > \lambda_{\text{low}}$ in state 1. The extinction count

is $D(T) \mid \{Y(s)\}_{s \leq T} \sim \text{Poisson} \left(\int_0^T \varphi(Y(s)) ds \right)$, the same Cox-process structure used in Chapter 15, now driven by $Y(t)$ in place of X_t .

This two-state chain is irreducible and positive recurrent for any $q_{01}, q_{10} > 0$: it visits both states infinitely often, with a well-defined long-run fraction of time in each, exactly the property GBM-with-absorption lacks. Its stationary distribution is $\pi_0 = q_{10}/(q_{01} + q_{10})$, $\pi_1 = q_{01}/(q_{01} + q_{10})$, obtained by solving $\pi Q = 0$ for the chain's generator Q in the standard way; we do not reproduce that elementary calculation here.

B.2 The Central Limit Theorem

Proposition B.1 (Central limit theorem for the regime functional, after Van Huu et al. 2005). *For the jump homogeneous Markov process $Y(t)$ and functional φ of Section B.1,*

$$\frac{1}{\sqrt{T}} \left(\int_0^T \varphi(Y(t)) dt - T\bar{\varphi} \right) \xrightarrow{d} N(0, \sigma^2) \quad \text{as } T \rightarrow \infty,$$

where $\bar{\varphi} := \pi_0 \lambda_{\text{low}} + \pi_1 \lambda_{\text{high}}$ is the long-run average intensity, for some $\sigma^2 > 0$ depending on the chain's parameters and on φ .

We do not reprove Proposition B.1 here; Van Huu et al. (2005) give sufficient conditions for exactly this convergence for a general jump homogeneous Markov process on a finite or countable state space, of which our two-state chain is a minimal instance. What we add is a closed form for σ^2 in this specific case, and a numerical bridge from the intensity functional to the extinction count itself, neither of which the general theorem provides on its own.

Proposition B.2 (Asymptotic variance, two-state case). *For the two-state chain of Section B.1,*

$$\sigma^2 = \frac{2(\lambda_{\text{high}} - \lambda_{\text{low}})^2 \pi_0 \pi_1}{q_{01} + q_{10}}.$$

We verified Proposition B.2 numerically before relying on it, rather than derive it from the general theory of Markov-modulated functionals, which would take this appendix considerably further into technical territory than its purpose warrants. At $q_{01} = 0.3$, $q_{10} = 0.5$, $\lambda_{\text{low}} = 0.5$, $\lambda_{\text{high}} = 3.0$, the formula gives $\sigma^2 \approx 3.66$; simulating the integral $\int_0^T \varphi(Y(t)) dt$ directly at $T = 200$ across many independent runs gave an empirical variance-per-unit-time of approximately 3.64, matching to within Monte Carlo error.

B.3 From Intensity to Extinction Count

Proposition B.1 concerns the integrated intensity $\Lambda(T) := \int_0^T \varphi(Y(t)) dt$, not the extinction count itself; the same law-of-total-variance step used in Proposition 15.8 bridges the two.

Proposition B.3 (Asymptotic normality of the extinction count). *Under the setup of Section B.1, as $T \rightarrow \infty$,*

$$D(T) \sim N(\bar{\varphi} T, \bar{\varphi} T + \sigma^2 T),$$

in the sense that $(D(T) - \bar{\varphi} T) / \sqrt{(\bar{\varphi} + \sigma^2)T}$ converges in distribution to $N(0, 1)$.

Proof sketch. By the same conditional-Poisson argument as Proposition 15.8, $\text{Var}[D(T)] = E[\Lambda(T)] + \text{Var}[\Lambda(T)]$ exactly, for every T . As $T \rightarrow \infty$, $E[\Lambda(T)] = \bar{\varphi}T$ and, by Proposition B.1, $\text{Var}[\Lambda(T)] \sim \sigma^2 T$, so $\text{Var}[D(T)] \sim (\bar{\varphi} + \sigma^2)T$, growing at the same linear rate as $\Lambda(T)$'s own variance; the additional Poisson randomization contributes $E[\Lambda(T)] = \bar{\varphi}T$, the same order, but does not introduce a second, faster-growing source of variability. A full proof would show that this randomization preserves asymptotic normality, a standard but nontrivial fact about mixed Poisson processes with an asymptotically normal, linearly growing intensity integral; we verify the conclusion numerically below rather than prove this step in general. $\square$

Worked Example

At the parameters above, we simulated $D(T)$ directly, drawing the regime path and a Poisson count conditional on it, for $T = 300$, $T = 1000$, and $T = 3000$, standardizing each sample by Proposition B.3's predicted mean and variance and testing the result against the normal distribution. At $T = 3000$, a Shapiro-Wilk test did not reject normality ($p \approx 0.64$) and a Kolmogorov-Smirnov test agreed ($p \approx 0.92$), with skewness and excess kurtosis both within 0.02 of zero. At $T = 300$, the same tests gave more mixed results, one test rejecting normality at conventional significance while the other did not, though skewness and kurtosis were already small. We report this honestly rather than only the clean result: Proposition B.3's mean and variance formulas hold at every T , exactly, but the normal shape itself is a large- T approximation, and at horizons comparable to the $T = 10$ used in this book's own worked examples, the count's true distribution is closer to a mixed Poisson than to a clean Gaussian. Convergence in this specific parameterization is not immediate; a reader applying this result should check, rather than assume, that their own horizon is long enough for the normal approximation to be trustworthy.

B.4 Relation to Chapter 15's Coupling

We want to be direct about what this appendix does and does not establish for the model Chapter 15 actually uses. Proposition 15.8's mean and variance formulas apply exactly to the GBM-driven, absorbing-barrier coupling of Section 15.7, for any finite T , and nothing in this appendix weakens that. What this appendix does not do is extend a normal approximation to that same absorbing model: the two-state regime-switching chain used here is a different driving process, chosen specifically because it is recurrent where GBM-with-absorption is not, and Proposition B.3's asymptotic normality is a property of extinction counts under this alternative driver, not a further fact about Chapter 15's own X_t . A reader wanting a distributional shape for the count under an ecosystem subject to genuine collapse risk, rather than persistent, recurring stress cycles, should use Proposition 15.8's exact mean and variance and treat the distribution's shape as unresolved, rather than borrow this appendix's normal approximation across a structural boundary it was not built to cross.

B.5 Summary

Chapter 15's absorbing-barrier coupling admits an exact mean and variance for its extinction count at any horizon, but not, without further assumptions, a limiting distributional shape, since the underlying process is not recurrent in the sense a central limit theorem needs. This appendix built a genuinely recurrent alternative, a two-state regime-switching stress process, applied the jump-Markov-process central limit theorem of Van Huu et al. (2005) to it, derived a closed-form asymptotic variance for this specific case (Proposition B.2), and extended the result from the underlying intensity functional to the Poisson-randomized extinction count

itself (Proposition B.3), verified numerically along with an honest account of how large a horizon the normal approximation actually needs. This is offered as a complementary model, addressing a question Chapter 15's own model structurally cannot, not as a refinement or extension of that chapter's results.

Self-Contained Notes on Prioritization and Institutional Hysteresis

This appendix proves the two results of Chapter 16 that needed more than a restatement: the diminishing-returns claim of Proposition 16.1, a genuine calculus argument, and the hysteresis claim of Proposition 16.2, which reuses Chapter 13's Allee-model technique rather than Chapter 13's specific formulas, since the two variables involved, a population level and a policy-attention share, do not share the same functional form. Both results are self-contained, requiring nothing beyond single-variable calculus and the same linearized-stability argument already used once in this book.

C.1 Diminishing Returns

Proof of Proposition 16.1. Treat m as a positive real variable on $(0, n]$ and write $c(m) = f(m)/m$ where $f(m) := \log_2 n - \log_2 m$. By the quotient rule,

$$c'(m) = \frac{f'(m)m - f(m)}{m^2}.$$

Since $\log_2 m = \ln m / \ln 2$, $f'(m) = -1/(m \ln 2)$. Substituting,

$$f'(m)m - f(m) = -\frac{1}{\ln 2} - (\log_2 n - \log_2 m) = -\frac{1}{\ln 2} - \log_2(n/m),$$

using $f(m) = \log_2(n/m)$ directly. Hence

$$c'(m) = \frac{-\frac{1}{\ln 2} - \log_2(n/m)}{m^2}.$$

It remains to sign the numerator. Since $1 \leq m \leq n$, $n/m \geq 1$, so $\log_2(n/m) \geq 0$; a nonnegative quantity exceeds $-1/\ln 2 \approx -1.4427$, so

$$-\frac{1}{\ln 2} - \log_2(n/m) < -\frac{1}{\ln 2} - 0 < 0$$

throughout $(0, n]$. The numerator is strictly negative and $m^2 > 0$, so $c'(m) < 0$ everywhere on $(0, n]$; a function with strictly negative derivative on an interval is strictly decreasing on it, and remains strictly decreasing when restricted to the integers $1, \dots, n$ on which $c(m)$ is actually evaluated in Chapter 16. $\square$

We verified this derivative's sign numerically across $n = 50$, $1 \leq m \leq 50$, before relying on the symbolic argument above, confirming $c(m)$ strictly decreasing at every step, matching the worked example already given in Chapter 16.

C.2 Institutional Hysteresis

Proof of Proposition 16.2. Write $F(y; \mu) := y(1 - y)(\beta y + \mu)$. Expanding, $F(y; \mu) = -\beta y^3 + (\beta - \mu)y^2 + \mu y$, so $F'(y; \mu) = \mu + 2(\beta - \mu)y - 3\beta y^2$. The equilibria of $\dot{y} = F(y; \mu)$ are $y = 0$, $y = 1$, and $y_c := -\mu/\beta$, the last lying in $(0, 1)$ only when $-\beta < \mu < 0$. Evaluating the derivative at each:

$$F'(0; \mu) = \mu, \quad F'(1; \mu) = -(\beta + \mu), \quad F'(y_c; \mu) = -\frac{\mu(\mu + \beta)}{\beta} \quad (-\beta < \mu < 0).$$

For $\mu > 0$: $F'(0; \mu) = \mu > 0$ (unstable), and since $\mu > 0 > -\beta$, $F'(1; \mu) = -(\beta + \mu) < 0$ (stable); $y_c < 0$, outside $[0, 1]$. For $\mu < -\beta$: $F'(0; \mu) = \mu < 0$ (stable), $F'(1; \mu) = -(\beta + \mu) > 0$ (unstable); $y_c > 1$, outside $[0, 1]$. For $-\beta < \mu < 0$: both endpoints are stable by the same formulas, and at y_c , $-\mu > 0$ and $\mu + \beta > 0$ give $F'(y_c; \mu) = -\mu(\mu + \beta)/\beta > 0$ (unstable), so y_c separates the two basins.

This is a linearized-stability argument of exactly the same form as Proposition 13.4's proof in Chapter 13: a cubic with three real roots on the relevant interval, evaluated at each root in turn, with sign of the derivative determining local stability. Because F is cubic with exactly these three roots on $[0, 1]$, it cannot change sign between consecutive roots without a fourth root, so the sign of F is constant on each open subinterval between them, and the local stability computed above determines the global direction of flow across that entire subinterval.

Under quasi-static variation of κ (equivalently $\mu = I\kappa - \theta_s$), a system tracks whichever stable branch it currently occupies. Starting at $y \approx 0$ with $\kappa < \kappa_L$ (so $\mu < -\beta$), $y = 0$ is the unique attractor. As κ rises through (κ_L, κ_U) (so $-\beta < \mu < 0$), $y = 0$ remains stable and y_c , moving continuously with μ , never approaches the system's neighborhood of 0, so the system remains at $y \approx 0$ throughout the bistable window. At $\kappa = \kappa_U$ ($\mu = 0$), $y = 0$ loses stability exactly as $y_c \rightarrow 0$, a saddle-node collision, since $F'(0; 0) = 0$; for κ infinitesimally above, $y = 0$ is no longer an equilibrium's neighborhood the system can remain in, and $y = 1$ is the only attractor, forcing the jump to $y \approx 1$. The reverse sweep is symmetric, tracking $y = 1$ down to $\kappa = \kappa_L$, where the collision instead occurs at $y = 1$ since $F'(1; -\beta) = 0$. Since $\beta > 0$, $\kappa_U - \kappa_L = \beta/I > 0$ strictly, giving the stated gap. $\square$

We verified the trichotomy and the hysteresis loop itself numerically before relying on this argument: direct simulation of $\dot{y} = F(y; \mu)$ from initial conditions near $y = 0$ and near $y = 1$ across a range of κ confirmed convergence to the predicted attractor in each of the three regimes, and a slow sweep of κ upward from 0.3 to 1.3 and back down, at $I = 1$, $\theta_s = 1$, $\beta = 0.4$, produced an upward jump near the predicted $\kappa_U = 1.0$ and a downward jump near the predicted $\kappa_L = 0.6$, matching Chapter 16's worked example to within the discretization used in the simulation.

C.3 Summary

Both results in this appendix reuse a proof technique already established once in this book rather than introducing a new one: Proposition 16.1's calculus argument is the same quotient-rule differentiation style used throughout this book's comparative-statics results, and Proposition 16.2's cubic-equilibrium analysis is, structurally, Proposition 13.4's own linearized-stability argument, re-derived for a differently parameterized variable rather than cited as a black box. Neither proof required expanding this book's mathematical toolkit; both required only pointing that toolkit at a new target.