The Economics of Viable Continuation

Production, Obligation, Institutions,
and the Organization of Feasible Futures

Flyxion
Independent Researcher

First full draft, October 2026

Contents

Preface **ix**

Reading Paths and Dependencies **x**

Common Model and Notation **xii**

The Housing Case **xiii**

The Proof Programme **xvi**

I Economic Activity and the Establishment of Futures **1**

1 What an Economy Maintains **3**
1.1 A transaction that produces income and uses up capacity 4
1.2 Relation to established accounting and capital theory 5
1.3 What the owner now decides 6

2 States, Histories, and Continuations **7**
2.1 States and histories 7
2.2 Cash and arrears together 9

3 Feasibility, Accessibility, and Admissibility **10**
3.1 Four sets of actions 10
3.2 Examples across domains and a boundary case 11

4 Establishing the Domain of Choice **13**
4.1 A two-stage model 13

5 Welfare Without Counting Options **16**
5.1 Why counting options fails 16
5.2 Conflict between continuations 17

II Production, Maintenance, and Irreversibility 19

6 Production as Transformation of Conditions 21
6.1 The housing transition . . . 21
6.2 Outputs, inputs, residuals, obligations, and capacity . . . 22
6.3 What each representation can answer . . . 23

7 Capital as Maintained Capability 24
7.1 Identical assets, different complements . . . 24

8 Maintenance and the Reproduction of Capacity 27
8.1 Deterioration, thresholds, and the viability boundary . . . 27
8.2 Deferral and the loss of viability . . . 28
8.3 Preventive maintenance, restoration, and backlog . . . 28

9 Irreversible Decisions and the Value of Waiting 30
9.1 The two-period comparison . . . 30
9.2 Symmetry of action and inaction . . . 31

10 Residuals, Waste, and the Afterlives of Production 33
10.1 What persists . . . 33
10.2 Who bears the residual . . . 34

III Obligations, Degradation, and Recovery 36

11 Obligations Through Time 38
11.1 Obligations as objects with structure . . . 38

12 Required Functions and Optional Capabilities 41
12.1 Classification as an institutional decision . . . 41
12.2 Declaring a function optional . . . 42

13 Viability Under Scarcity 44
13.1 The viability kernel . . . 44
13.2 Comparative statics . . . 44
13.3 Local acceptability and almost-sure viability . . . 45

14 Compensating Continuation 47
14.1 A rank certificate . . . 47
14.2 Credible and non-credible compensation . . . 48

15 Corrective Capacity and Recoverability 49
15.1 Lookahead depth and the point of no return . . . 49

15.2 The option value of corrective capacity 50

IV Value, Prices, Accounting, and Debt 52

16 Valuing Feasible Futures 54
16.1 Incomparable outcomes 54
16.2 Money as one representation 55

17 Prices as Compressed Coordination Signals 57
17.1 Scarcity prices as dual variables 57
17.2 What the price hides 58

18 Profit and the Location of Costs 60
18.1 Decomposing profit 60
18.2 Profit supported by reproduction, and profit supported by deferral 61

19 Monetary Debt and Restoration Obligations 62
19.1 Why money does not suffice 62
19.2 Monetization and substitution 63

20 Accounting for Continuation 64
20.1 The factorization problem 64
20.2 Supplementary accounts 64

V Knowledge, Attention, and Productive Competence 66

21 Economic Representations and Semantic Compression 68
21.1 Task-relative adequacy 68
21.2 Compression in forecasting and classification 69

22 Memory as Organizational Capital 70
22.1 Retained, retrievable, and usable 70
22.2 Forgetting without erasing obligations 71

23 Aspect Relegation and the Economics of Skill 72
23.1 A margin for detection and return 72
23.2 The practice that preserves the margin 73

24 Automation and the Preservation of Intervention 75
24.1 The cost of a delayed return 75
24.2 Comparing policies over the horizon 75

25 Verification, Proxies, and Institutional Confidence **77**
25.1 What a passed check establishes 77
25.2 Proxies and institutional confidence 78

VI Institutions, Power, and Distribution **79**

26 Distinctions That Allocate Futures **81**
26.1 Labels as state variables 81
26.2 Relation to institutional economics 82

27 Property as Structured Authority **83**
27.1 Permissions and capacity 83

28 Contracts and the Evolution of Dependence **85**
28.1 Switching costs and the bargaining position 85
28.2 Interventions that accumulate 86

29 Distribution of Exit and Refusal **87**
29.1 Unequal capacities to leave 87

30 Public Authority and the Admissibility of Constraints **89**
30.1 Procedure and evidence 89
30.2 Authority, contestability, and legitimacy 90

VII Markets, Firms, Commons, and Coordination **91**

31 Exchange and Compatible Commitments **93**
31.1 Exchange as transformation of resources, permissions, and obligations . . . 93
31.2 Reservations and concurrent requests 94

32 Firms as Organizations of Complementary Continuation **95**
32.1 Coordination, authority, and the boundary 95
32.2 Assume-guarantee composition 95

33 Platforms and Infrastructural Dependence **97**
33.1 Quality, participation, and revenue 97
33.2 Appropriation of complementary investment 97

34 Commons and Shared Maintenance **99**
34.1 Contribution, allocation, and replenishment 99

35 Contribution, Reward, and Measurement **101**

35.1 Proxy rewards and what they miss . . . 101
35.2 Rewards change what they measure . . . 102

VIII Ecology, Care, Technology, and Development 103

36 Ecological Reproduction and Economic Time 105
36.1 Regeneration, thresholds, and horizons . . . 105
36.2 Damp and the dwelling . . . 106

37 Care and the Maintenance of Persons 107
37.1 The caregiver's viability . . . 107
37.2 Dependency and the danger of valuing persons by contribution . . . 108

38 Local AI and the Ownership of Productive Infrastructure 109
38.1 Total costs and dependence . . . 109
38.2 Distribution of control . . . 110

39 Development as the Establishment of Capabilities 111
39.1 Establishing domains and who can reach them . . . 111
39.2 Output and authority . . . 112

40 Crisis, Transition, and Institutional Repair 113
40.1 A viable destination and a viable path . . . 113
40.2 Waiting for a subsidy . . . 114
40.3 Mechanisms for revising a plan . . . 114

IX Mathematical Foundations and Economic Results 115

41 Transition Systems and History Compression 117
41.1 Setting and the factorization lemma . . . 117
41.2 What a representation forgets . . . 118

42 Orders, Dominance, and Incomparability 119
42.1 Dominance orders . . . 119

43 Reachability, Robustness, and Recovery 121
43.1 The kernel and its algorithms . . . 121
43.2 Worst-case, almost-sure, and risk-budget guarantees . . . 121
43.3 The waiting theorem . . . 122

44 Composition and Systemic Failure 123
44.1 The obstruction and the positive results . . . 123

44.2 Accounting consistency and operational progress . . . 123

45 Measurement and Empirical Identification 125
45.1 From quantities to indicators . . . 125
45.2 Which causal claims need which designs . . . 125

X Categorical Interpretations 127

46 Economic Processes as Morphisms 129
46.1 The syntactic category . . . 129
46.2 What the construction adds . . . 130

47 Parallel Production and Monoidal Structure 131
47.1 Independence fails under a shared sink . . . 131

48 Resources, Obligations, and Enriched Categories 133
48.1 Costs that compose . . . 133

49 Accounting as a Structure-Preserving Map 135
49.1 The ledger functor . . . 135

50 Open Systems, Institutions, and Compositional Guarantees 137
50.1 Interfaces and the assume-guarantee form . . . 137

XI Type Theoretic Interpretations 139

51 Economic States and Admissible Actions as Types 141
51.1 Actions indexed by state . . . 141
51.2 Inhabitation is not executability . . . 141

52 Linear and Affine Resources 143
52.1 Which resources may be copied . . . 143

53 Obligations, Deadlines, and Protocols 145
53.1 Escrow as a protocol with failure branches . . . 145

54 Proof-Carrying Commitments and Verified Degradation 147
54.1 A certificate that can be checked . . . 147
54.2 Replay integrity and historical authority . . . 148

55 Refinement, Interpretation, and the Limits of Formal Assurance 149
55.1 A verified scheduler that fails its purpose . . . 149

XII Sheaf Theoretic Interpretations **151**

56 Local Economic Knowledge and Overlapping Contexts **153**
56.1 A base of contexts . . . 153

57 Presheaves of Plans, Accounts, and Commitments **155**
57.1 Restriction and what it forgets . . . 155

58 Gluing and Global Feasibility **157**
58.1 Matching plans that cannot be carried out . . . 157

59 Obstructions, Inconsistency, and Institutional Boundaries **159**
59.1 An exchange graph . . . 159
59.2 A twisted cover . . . 160

60 Descent, Federation, and the Organization of Collective Futures **161**
60.1 Federated records . . . 161
60.2 The three interpretations together . . . 161

A Prerequisites **163**
A.1 Optimization . . . 163
A.2 Dynamical Systems . . . 163
A.3 Probability . . . 163
A.4 Order Theory . . . 163
A.5 Category Theory . . . 163
A.6 Dependent and Linear Types . . . 164
A.7 Sheaves . . . 164

B The Computational Companion **165**

C Framework Concordance **167**

Preface

This book develops an economic theory in which the central object is the set of futures that an economy keeps open. Production, maintenance, obligation, exchange, and institutional design are studied by asking what continuations each arrangement makes possible, to whom, and at what cost to the conditions that later activity requires. The same transaction can produce income in the present while consuming capacities elsewhere, and the theory is built to make that difference visible and testable.

The order of presentation is deliberate. Parts I through VIII develop the economic concepts in their own terms, with motivating situations, worked examples, and identifiable consequences. Part IX consolidates the formal machinery and states the book's principal results with their hypotheses. Parts X, XI, and XII then give categorical, type theoretic, and sheaf theoretic interpretations. Each interpretation recovers structure already established in the economic theory and states what it preserves and what it leaves out. No chapter treats a mathematical analogy as a result.

Every chapter follows a common pattern: a motivating economic situation, definitions, a developed argument, a worked example, propositions with proofs where warranted, empirical questions, and exercises. Formal results are consequences of stipulated models. Claims about actual economies are labelled as such and are the subject of the empirical questions. The book's central research commitment is to determine when continuation-based descriptions yield explanations or decisions that ordinary accounting, optimization, and institutional analysis would miss, and to say plainly when they do not.

Draft note. Scope. The outline does not yet treat money and banking, labour markets, international trade, long-run growth, or mechanism design as separate topics. Either state this scope explicitly here, or fold labour into the chapters on exit and care, and mechanism design into the chapters on exchange and protocols. A possible additional chapter would recover Arrow–Debreu equilibrium, Solow growth, and dynamic programming as limiting cases (history compresses, viability constraints are slack, valuation is scalar), giving graduate readers a benchmark for later claims.

Draft note. Length. Sixty chapters is on the order of a thousand pages. Reading paths are given in the next section. Parts X to XII could be issued as a second volume if the length becomes unwieldy.

Reading Paths and Dependencies

The twelve parts are not meant to be read in a single sitting. Three routes through the book are offered. The economics route covers Parts I through VIII and needs no category theory, type theory, or sheaf theory. The formal-core route takes Parts I to III for the economic foundation, then Part IX for the principal results, and then one of the interpretation parts, X, XI, or XII. The full route reads the parts in order. Table 1 marks the parts each route includes, and Figure 1 shows which parts each part presupposes.

Route	I	II	III	IV	V	VI	VII	VIII	IX	X	XI	XII
Economics	■	■	■	■	■	■	■	■				
Formal core	■	■	■						■	one of X, XI, XII		
Full	■	■	■	■	■	■	■	■	■	■	■	■

Table 1: Reading routes. Shaded cells are read.

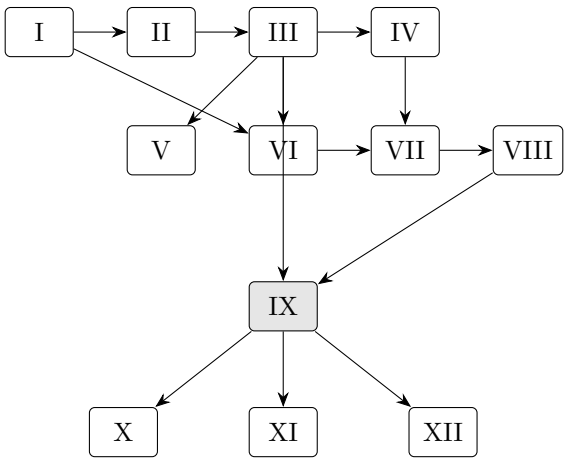


Figure 1: Dependencies among parts. An arrow from A to B means B presupposes A. Part IX is formally self-contained given Parts I to III, and the arrow from Part VIII records that its results are motivated by the whole economic text.

Housing as the running case

A housing-maintenance system recurs throughout. It begins with physical deterioration and tenant obligations, and then carries the argument through accounting, exit capacity, contractor composition, typed repair commitments, and the compatibility of local records. Each part opens with a short paragraph marked *Housing thread* stating what the case contributes there. Reading these paragraphs in order shows what each additional mathematical language adds to the same example.

Common Model and Notation

A single model is used throughout, and each chapter states any departure from it. Let $s \in S$ be an augmented economic state containing material resources, relevant institutional permissions, and outstanding obligations. In a state s the available actions are $a \in A(s)$, and the allowed disturbances are $w \in W(s, a)$. Successor states are given by

$$s' = T(s, a, w).$$

Both $A(s)$ and $W(s, a)$ are assumed nonempty unless a result says otherwise. Let $G \subseteq S$ be the set of acceptable states.

A *policy* is nonanticipating: it may use the observed history but cannot depend on disturbances that have not yet been observed. Unless stated, observation is complete, so the state and the past disturbances are observed. A policy keeps the system *in G forever* against a class of disturbances if every trajectory it generates under an allowed disturbance sequence stays in G.

Three structures are kept separate throughout: *admissibility* (which trajectories satisfy the stated obligations), *welfare* (which trajectories are better for whom), and *financial value* (how trajectories are priced or accounted). A result about one is never silently a result about another.

Symbol	Meaning
$S,\ A(s),\ W(s,a),\ T$	states, actions, disturbances, transition map
G	acceptable states
Φ	robust predecessor operator, $\Phi(X) = \{s \in G : \exists a\ \forall w,\ T(s,a,w) \in X\}$
$K = K(G)$	viability kernel: greatest fixed point of Φ
$L(h)$	set of obligation-satisfying future sequences after history h
$h \sim h'$	continuation equivalence, $L(h) = L(h')$
$F : S \to Y$	an accounting or other representation
$\preceq$	partial order on continuation profiles

The Housing Case

A single small housing economy runs through the book. It is introduced here in a deterministic version that can be calculated by hand, and later chapters add disturbances, shared resources, information problems, and institutional change to the same case with the same units, timing, and signs. Amounts are in currency units, time is in months, and nothing is discounted until a chapter says otherwise.

Participants and what is observed

A *tenant* holds liquid resources $x \geq 0$ and pays rent r per month. After rent and living costs the tenant's net monthly surplus on a monthly contract is $e \geq 0$. A *landlord* owns the dwelling and may offer contracts. A *contractor* and a *public authority* enter in later chapters, as do financial accounts and outside housing opportunities. Here the tenant's x, e, and the exit cost below are observed by the tenant. The landlord observes the contract terms and whether rent is paid.

The decision

Suppose a landlord offers a prepaid lease. The tenant pays $P = (1 - \delta)\, rT$ at the start for a lease of T months at a discount $\delta \in [0, 1)$. The prepayment is non-refundable, and the tenant pays no further rent during the lease. The landlord saves a collection and arrears cost $\kappa > 0$ per prepaid contract. The tenant must hold $x \geq P$ to be eligible.

The offer looks attractive on its financial face. The discount δrT is a transfer from landlord to tenant, so aggregate surplus rises by κ, and if $\delta rT \leq \kappa$ the landlord's profit does not fall. The difficulty is that a tenant who prepays has less cash on hand for the next H months, which is the window in which a job loss, a repair failure, or a safety problem might force a move. A move at month $t \leq H$ needs resources of at least c at that date, where c covers moving costs, a new deposit, and a first month's rent.

Definition 0.1 (Executable exit)**.** An exit within horizon H is *executable* for a tenant if there is a date $t \in \{0, \ldots, H\}$ at which the tenant's resources, with every earlier balance nonnegative, are at least c. This is the constrained reachability of P08 in the present model.

On a monthly contract with $e \geq 0$, resources at month t are $x + te$, nondecreasing, so the exit is executable exactly when $x + He \geq c$. On the prepaid lease, resources at month t are $x - P + t(e + r)$ because rent is no longer paid, and the exit is executable exactly when $x - P + H(e + r) \geq c$.

Theorem H1 (P28 in the housing model)**.** *Fix r, T, δ, e, c, H with $e \geq 0$ and $T > H$. Among tenants eligible for the prepaid lease:*

(a) Aggregate surplus, counted as the sum of landlord profit and tenant flow payoff with the tenant's stay decision held fixed, rises by κ. Landlord profit changes by $\kappa - \delta rT$ and the tenant's by δrT.

(b) The tenants who had an executable exit on the monthly contract and lose it on the prepaid lease are exactly those with

$$\max\{P,\ c - He\} \ \leq \ x \ < \ c - He + P - rH.$$

(c) This set is nonempty if and only if $H < (1 - \delta)T$ and $c > H(e + r)$.

Proof. (a) The discount is paid by the landlord and received by the tenant, so it cancels in the sum, leaving the cost saving κ. (b) A tenant is eligible if $x \geq P$. The exit was executable before if $x \geq c - He$. It is not executable after if $x - P + H(e + r) < c$, that is $x < c - He + P - rH$. Intersecting the three conditions gives the interval. (c) The interval is nonempty if and only if its upper end exceeds both lower ends. Exceeding $c - He$ requires $P - rH > 0$, that is $H < (1 - \delta)T$. Exceeding P requires $c - He - rH > 0$, that is $c > H(e + r)$. □

Hypotheses. The conclusion is carried by two assumptions: the prepayment is nonrefundable, and the exit cost c must be met from liquid resources within the window. Nonnegativity of e only simplifies the safety check, since otherwise the balance path must be tracked. The assumption most likely to fail in applications is that the surplus measure ignores the value of the exit option. If it were priced, the sign of the aggregate change would depend on the probability of needing to exit.

Worked numbers. Take $r = 1000$, $T = 12$, $\delta = 0.05$, $e = 200$, $c = 4000$, $H = 3$, $\kappa = 800$. Then $P = 11{,}400$ and $\delta rT = 600$. Aggregate surplus rises by 800, the landlord gains 200, and a tenant who never needs to move gains 600. On the monthly contract the exit is executable for $x \geq 4000 - 600 = 3400$. On the prepaid lease it is executable for $x \geq 4000 - 11{,}400 + 3 \cdot 1200 = 11{,}800$ among eligible tenants. Tenants with $11{,}400 \leq x < 11{,}800$ are eligible, would have been able to exit, and can no longer do so. Both conditions of (c) hold, since $3 < 11.4$ and $4000 > 3 \cdot 1200 = 3600$.

Interpretation. The financial improvement and the loss of exit capacity are properties of different structures, as the book's separation of welfare, admissibility, and financial value requires. The result does not say prepaid leases harm tenants in general. It gives the exact condition under which a gain in surplus coexists with a smaller safe exit set, and the exact set of tenants affected.

When compensation restores the lost capability

Suppose the landlord compensates an affected tenant with a transfer m delivered at month τ. An outside unit can be taken only within a window of months $[a, E]$, so the exit date t must satisfy $a \le t \le \min\{H, E\} =: \bar{t}$. The transfer counts only at dates $t \ge \tau$.

Theorem H2 (P29 in the housing model)**.** *Let a tenant be eligible with a positive shortfall,*

$$m^*(x) \;=\; c - (x - P) - \bar{t}\,(e + r) \;>\; 0.$$

The transfer m delivered at τ restores an executable exit if and only if all three hold: (i) availability, $a \le \bar{t}$; (ii) timing, $\tau \le \bar{t}$; (iii) sufficiency, $m \ge m^*(x)$. If (i) or (ii) fails, no amount of money restores the capability.

Proof. Resources at date t are $x - P + t(e + r) + m\,\mathbf{1}[\tau \le t]$, which is nondecreasing in t when $e + r > 0$. If $\tau \le \bar{t}$ the best date is $\bar{t}$ and the condition is $m \ge m^*(x)$. If $\tau > \bar{t}$ the transfer never counts and the shortfall stays positive. If $a > \bar{t}$ there is no admissible date. □

Hypotheses and interpretation. The three conditions correspond to the three ways money can fail: the thing to be bought may not exist in the window, the money may arrive too late, or it may be too small. A transfer that increases holdings without satisfying (i) to (iii) has not restored the capability. In the worked numbers, a tenant with $x = 11{,}500$ has $m^* = 4000 - 100 - 3 \cdot 1200 = 300$. A transfer of 300 delivered by month 3 restores the exit. A transfer of 10^6 delivered at month 4 does not. The largest shortfall among affected tenants is $m^*(P) = c - \bar{t}(e + r) = 400$, which exceeds the landlord's net gain $\kappa - \delta r T = 200$, so in this example the surplus gain does not by itself finance full restoration. That comparison is a property of these numbers, not a theorem.

Scope. Both results are proved only in this deterministic model. The general statements remain specifications (P28, P29). The companion reproduces the closed forms by brute-force enumeration over integer parameter sets, as a reported check.

Next. Part I introduces the states and histories that let this case be stated without the special structure used here. The chapter on exit (Chapter 29) generalizes Theorem H1, and the chapter on restoration obligations (Chapter 19) generalizes Theorem H2.

The Proof Programme

The results below are theorem *specifications* for the writing programme. Several elementary results already admit the short proofs sketched here. The broader economic interpretations depend on assumptions that the chapters must make explicit, and some entries are statements whose exact form is still to be fixed. Each entry carries a family tag and a status:

Elementary
The statement is fixed and a short proof is indicated.

Hypotheses to fix
The statement is fixed in outline but the hypotheses must be pinned down in the chapter that proves it.

Specification
The result is described but the precise statement is not yet settled.

The six families are REP (representation and economic knowledge), VIA (viability and comparative statics), PRD (production, deterioration, irreversible investment), INS (institutions, distribution, dependence), CPT (competence, automation, corrective control), and CMP (composition and mathematical interpretation). Identifiers P01 to P29 are stable and are cited by the chapters; the family tag organizes Part IX.

Dependency structure

P01 to P03 are the foundation. P04 to P10, together with P23 and P25, form the theory of sustainable commitments. P11 to P17, together with P24, form the theory of economic representation and control. P26 to P29 add the economic substance for investment, institutions, and compensation. P18 to P22 are the mathematical reinterpretations, and P27 organizes composition across all three. The housing-maintenance example recurs across every group.

Representation and economic knowledge

P01. History sufficiency and continuation equivalence. [REP]

Hypotheses. For an observed history h let $L(h)$ be the set of future action-and-disturbance sequences that satisfy the specified obligations. A summary q is *exact for admissibility* if, for every history h and every future sequence u, whether $u \in L(h)$ is determined by $q(h)$ and u.

Statement. Define $h \sim h'$ if and only if $L(h) = L(h')$. Any exact summary satisfies $q(h) = q(h') \Rightarrow h \sim h'$, so $\sim$ is the coarsest exact summary. If the continuation language has the residual property $L(he) = e^{-1}L(h)$ for each event e, then $\sim$ is a right congruence: $h \sim h' \Rightarrow he \sim h'e$, and the classes support well-defined updates.

Consequence. Debt, deferred maintenance, accumulated exposure, and contractual commitments cannot be dropped merely because current physical holdings match. The result concerns exact admissibility recognition. Predictive sufficiency for probabilities or welfare needs further conditions.

Status. Elementary. Congruence requires the residual property.

P11. Accounting sufficiency requires factorization. [REP]

Statement. Let $F : S \to Y$ be an accounting representation and $P : S \to \{0, 1\}$ an economic property. The property is determined exactly by the accounts if and only if $F(s) = F(s') \Rightarrow P(s) = P(s')$, equivalently $P = \bar{P} \circ F$ on the image of F for some $\bar{P} : F(S) \to \{0, 1\}$.

Relation to P01. P01 and P11 are one factorization lemma. P01 is its dynamic version, where the representation must in addition be a congruence for event updates. Minimal sufficient representation is the quotient by the kernel of the property.

Consequence. Equal reported profit can coexist with different answers about maintenance adequacy, restoration exposure, or viable operation, whenever F identifies states that differ on those properties.

Status. Elementary.

P13. Scalar representation fails for genuine incomparability. [REP]

Statement. (a) If a partial order $\preceq$ on continuation profiles contains incomparable x, y, no real-valued v satisfies $x \preceq y \Leftrightarrow v(x) \leq v(y)$ for all pairs, since $\mathbb{R}$ is totally ordered. (b) Positive companion: every partial order admits a multi-utility representation, a family (v_i) of monotone real functions with $x \preceq y \Leftrightarrow v_i(x) \leq v_i(y)$ for all i. For a finite order the minimal number of functions equals the order dimension (Dushnik and Miller [9]).

Consequence. Useful prices, utilities, and monotone indicators are not ruled out. What is ruled out is an exact, order-reflecting scalar representation of all distinctions.

Status. (a) elementary; (b) hypotheses to fix, in particular the infinite case.

P15. Proxy rewards underdetermine substantive contribution. [REP]

Statement. Let $m(a)$ measure rewarded activity and $V(a)$ the substantive contribution sought. If $m(a) = m(b)$ and $V(a) \neq V(b)$, a reward depending only on m cannot distinguish a from b. If $m(a) > m(b)$ and $V(a) < V(b)$, a strictly increasing function of m ranks them opposite to V.
Consequence. This is a limit on information, not a theory of behaviour. Behaviour requires the participants' costs, available actions, and preferences to be specified separately.
Status. Elementary.

P17. Verification transfers only through an adequate interpretation. [REP]

Hypotheses. Let M be a formal model, E the empirical system, and R a relation connecting their states. Assume soundness of overapproximation: every relevant empirical execution is R-related to a model execution, and R-related states satisfy that model safety implies empirical safety.
Statement. A verified universal safety property of M transfers to E under these hypotheses.
Consequence. Passing a check establishes what its specification and interpretation support. The connection to the economic situation is an additional proof obligation. Liveness and probabilistic claims need stronger relations.
Status. Hypotheses to fix.

P24. Scarcity prices as dual variables. [REP]

Hypotheses. A convex allocation problem with a capacity constraint $\sum_j d_j \leq c$, Slater's condition, and attainment, so that strong duality holds and a multiplier exists.
Statement. There is a multiplier $\lambda^* \geq 0$ supporting the optimal allocation, with complementary slackness $\lambda^*(c - \sum_j d_j) = 0$. If in addition the objective and constraints decompose across participants (separability) and each participant's choice rule is the individual Lagrangian problem at λ^*, then the multiplier is a decentralized price. A multiplier alone does not establish a market implementation.
Consequence. Supporting prices, equilibrium prices, sensitivity derivatives, and sufficient statistics are distinct notions, and uniqueness, decentralizability, and informational sufficiency are separate claims. By P11 a price is a sufficient statistic for a given decision only if that decision factors through it.
Status. Hypotheses to fix.

Viability and comparative statics

P03. Robust viability as a greatest fixed point. [VIA]

Statement. Define $\Phi(X) = \{s \in G : \exists a \in A(s)\ \forall w \in W(s,a),\ T(s,a,w) \in X\}$. The operator Φ is monotone on subsets of S, so by Knaster and Tarski [21] it has a greatest

fixed point $K = \bigcup\{X \subseteq G : X \subseteq \Phi(X)\}$. With $A(s)$ and $W(s,a)$ nonempty and complete observation, K is exactly the set of states from which a policy keeps the system in G forever against all allowed disturbances, and a state-feedback policy suffices.
Convergence. The iteration $K_0 = G$, $K_{n+1} = \Phi(K_n)$ is decreasing. It stabilizes at K in at most $|S|$ steps when S is finite. In general the iteration may need transfinite stages, and the companion code exhibits an infinite chain where each K_n is nonempty but $K = \varnothing$. Termination claims are therefore stated only under separately stated hypotheses. See Aubin [4] for continuous-time viability theory.
Consequence. Present acceptability and sustainable acceptability are different properties.
Status. Elementary for the fixed point and the witnessing action. Convergence needs finiteness.

P23. Comparative statics of viability. [VIA]

Hypotheses. Systems compared on a common state space.
Statement. The kernel K is monotone increasing in the acceptable set G, increasing in each action set $A(s)$, and decreasing in each disturbance set $W(s,a)$, all else equal. The proof is monotonicity of Φ in each parameter, hence of its greatest fixed point.
Consequence. Added duty, added capacity, and added exposure are treated in one statement. P04 and P14 are corollaries.
Status. Elementary.

P04. Obligation tightening cannot enlarge viability. [VIA]

Statement. On a common state space with identical dynamics, actions, and disturbances, $G_2 \subseteq G_1$ implies $K(G_2) \subseteq K(G_1)$. If a new obligation enlarges the state, compare through an explicit simulation relation $\rho \subseteq S_2 \times S_1$ under which ρ-related states respect acceptability and every action and disturbance of the tighter system is matched in the weaker one.
Consequence. Tightening can leave the kernel unchanged, so an added duty does not by itself require extra resources or changed technology. The requirement arises only when the new constraint removes previously viable states or policies, that is, when $K(G_2) \subsetneq K(G_1)$. Preserving the old kernel then requires altering resources, permissions, technology, or disturbance exposure.
Status. Elementary on a common state space. Hypotheses to fix for the simulation version.

P05. Local acceptability does not imply viable continuation. [VIA]

Statement. There exist systems with $s_0 \in K$, $s_1 \in G$, $s_1 \notin K$, and an available action from s_0 to s_1. Three states suffice: from s_0 one action stays at s_0 and another enters s_1, and from s_1 every action enters an unacceptable state $b \notin G$.
Consequence. A current-state check can pass while continued feasibility is destroyed. This supports the trajectory-level analysis of maintenance cuts, cumulative dependence, and repeated interventions.

Status. Elementary. Reported check in the companion (Appendix B).

P06. Safe degradation requires successor viability. [VIA]

Statement. Let $A_K(s) = \{a \in A(s) : T(s, a, w) \in K \text{ for every allowed } w\}$. For $s \in K$, a degradation action preserves the possibility of satisfying the ongoing regime exactly when it lies in $A_K(s)$, and $K = \{s \in G : A_K(s) \neq \varnothing\}$. If degradation creates a restoration duty, that duty must be included in the successor state before A_K is evaluated.
Consequence. Temporary suspension cannot be evaluated apart from the obligation it leaves behind.
Status. Elementary, from P03.

P07. A constructive certificate for compensating continuation. [VIA]

Hypotheses. Let $R \subseteq G$ be the states in which a restoration obligation is discharged. Let $D \subseteq G$ carry a rank $r : D \to \mathbb{N}$ with $r(s) = 0$ exactly on R, and let a policy be such that every allowed successor remains in D and strictly decreases r whenever $r > 0$.
Statement. From $s_0 \in D$ restoration occurs within $r(s_0)$ steps while remaining acceptable. Induction on r.
Consequence. Reserved resources and permissions must support a policy whose progress survives the modelled disturbances.
Status. Elementary.

P14. Corrective control has weak option value. [VIA]

Statement. If $\Pi_0 \subseteq \Pi_1$ are policy sets for identical environments and criteria, $\sup_{\Pi_1} J \geq \sup_{\Pi_0} J$. Strict improvement requires a circumstance in which the added action helps. For viability the same statement is the action-set case of P23.
Consequence. Corrective capacity has potential value, but its acquisition is a substantive investment problem. Training, monitoring, reserves, and intervention infrastructure must be priced explicitly.
Status. Elementary.

P25. Almost-sure and risk-constrained viability. [VIA]

Settled statement (countable case). Robust viability over disturbance supports is not in general equivalent to almost-sure viability. In a finite or countable model, requiring every positive-probability successor to stay in X gives the almost-sure predecessor condition. For general kernels, the operator to be studied is $\Phi_{\text{as}}(X) = G \cap \{s : \exists a \in A(s),\ \mathbb{P}(X \mid s, a) = 1\}$.
Hypotheses to fix. Identifying the greatest fixed point of Φ_{as} with an implementable policy needs measurability and selection assumptions. Moving to universally measurable or analytic sets does not by itself solve the problem, since those classes are not closed under all the operations and transfinite iterations involved (Bertsekas and Shreve [5]). The measurable domain and the policy-selection theorem must be fixed first. A support formulation holds

for closed X, where full measure implies containment of the support. Risk-constrained viability needs an augmented state carrying a remaining risk budget and a justified rule for allocating conditional risk across successors.
Status. Countable case: elementary. General-state version: research specification until the measurable domain and selection theorem are fixed.

Production, deterioration, and irreversible investment

P12. Financial improvement can accompany capability depletion. [PRD]
Statement. Let surplus be $\pi = R - C - M$ with maintenance expenditure M, and capacity evolve by $k_{t+1} = k_t - d(k_t) + g(M_t)$. (a) If a reduction $\delta > 0$ in maintenance leaves revenue and other costs unchanged, current surplus rises by δ, and if g is strictly increasing over the relevant interval, successor capacity falls relative to the maintained baseline. (b) Threshold: derive the viability boundary $k^*(\bar{M})$ for a maintenance budget $\bar{M}$, and state conditions under which it jumps or has a kink.
Consequence. Nonlinear deterioration alone does not guarantee a discontinuity or a kink. The threshold result must say what does.
Status. (a) elementary. (b) specification.

P26. When waiting improves attainable value. [PRD]
Elementary two-period form. Investment of cost I yields a payoff X with mean μ, revealed after one period. Acting now yields $V_{\text{act}} = (\mu - I) + r$, where r is the first-period flow return forgone by waiting. Waiting yields $V_{\text{wait}} = \beta\,\mathbb{E}[\mathbf{1}_{\text{surv}}(X - I)^+]$ with discount β, where $\mathbf{1}_{\text{surv}}$ indicates that the opportunity survives. If survival is independent of the revealed payoff with probability q, this equals $\beta q\,\mathbb{E}[(X - I)^+]$. Otherwise it does not, and the dependence must be modelled. Waiting is better exactly when $V_{\text{wait}} > V_{\text{act}}$. By Jensen, $\mathbb{E}[(X - I)^+] \geq (\mu - I)^+$, so the option value of information is nonnegative before discounting, deterioration, and expiry. See Dixit and Pindyck [8].
General statement. Compare the value of information acquired through delay with deterioration, forgone returns, and expiring opportunities, treating action and inaction symmetrically as potentially irreversible.
Status. Two-period form elementary. General form specification.

Institutions, distribution, and dependence

P02. Establishment precedes selection. [INS]
Statement. Let D be an opportunity domain, $E(e, D)$ the domain established by an acquisition action e, and $C(D)$ the outcomes obtainable by selection inside D. If $x \notin C(D)$ and

$x \in C(E(e, D))$, selection within D cannot implement x whereas establishment followed by selection can. Budgeted version: assume a zero-cost identity establishment action (or an explicit permission to skip establishment). Then with budget b, outcomes under establish-then-select policies contain those under selection only, and the inclusion is strict only if an affordable establishment action makes an otherwise inaccessible outcome reachable.
Consequence. Expenditure on access, infrastructure, learning, or recognition can change the choice problem itself.
Status. Elementary.

P08. Nominal exit and executable exit are distinct. [INS]

Statement. Let X be an exit target and B the states satisfying the safety and obligation constraints during departure. Define $\text{Reach}_0 = X$ and $\text{Reach}_{k+1} = \text{Reach}_k \cup \{s \in B : \exists a\ \forall w,\ T(s, a, w) \in \text{Reach}_k\}$. An exit is executable within horizon H exactly when $s \in \text{Reach}_H$. A legal permission to exit does not imply membership. A permitted departure requiring a fee beyond the person's accessible resources is a counterexample.
Consequence. Exit capacity is an operational variable and its distribution can be studied independently of whether contracts contain an exit clause.
Status. Elementary.

P28. Surplus gain with a shrunken safe exit set. [INS]

Statement. Characterize when an intervention raises aggregate surplus while reducing a specified participant's executable exit set, building on P08, and relate this to bargaining, cumulative contractual dependence, and portability.
Status. Proved in the hand-calculable housing model (Theorem H1 in the housing case). General statement: specification.

P29. When compensation restores capability. [INS]

Statement. Identify conditions under which a transfer restores the lost capability rather than merely increasing monetary holdings: the compensating bundle must lie in the reachability preimage of the lost capability under prices and access actually available, using P08 and P07.
Status. Proved in the housing model (Theorem H2 in the housing case): restoration requires availability, timing, and sufficiency. General statement: specification.

Competence, automation, and corrective control

P16. Safe aspect relegation requires a detection-and-return margin. [CPT]

Statement. Let an automated operating region lie at distance d from a failure boundary, with adverse motion at rate at most v, detection and return to explicit control within

time τ, and intervention able to arrest deterioration before the boundary. Then $d > v\tau$ is sufficient, plus allowances for observation error and stopping distance. Endogenous version: response time depends on retained competence, $\tau_t = \tau(c_t)$, and $c_{t+1} = f(c_t, u_t, \ell_t)$ with practice u_t and reliance on relegated operation ℓ_t. Under monotonicity assumptions (τ decreasing in c, f increasing in u and decreasing in ℓ), the condition is $d_t > v_t\,\tau(c_t) +$ allowances.
Economic result. The minimal practice expenditure u_t^* that preserves the margin, and the circumstances in which apparent automation savings are consumed by it.
Status. Fixed-τ form elementary. Endogenous form hypotheses to fix.

Composition and mathematical interpretation

P09. Local feasibility is insufficient under shared scarcity. [CMP]

Statement. Two commitments each needing one unit of a resource at the same time, with one unit available, are each feasible alone and jointly infeasible. For divisible simultaneous demands d_j on capacity c, $d_j \le c$ for every j does not imply $\sum_j d_j \le c$, and $\sum_j d_j \le c$ is necessary and sufficient for joint feasibility in the static divisible case.
Consequence. This is the obstruction every later compositional guarantee must address. Departments, firms, and agencies can each hold feasible plans that cannot be executed together.
Status. Elementary. Reported check in the companion (Appendix B).

P10. Independent viability composes under explicit independence. [CMP]

Hypotheses. Two systems with product state space, independent actions, product disturbance sets, factorized dynamics, nonempty action and disturbance sets, and acceptable set $G_1 \times G_2$.
Statement. The robust kernel is $K_{12} = K_1 \times K_2$. For $\supseteq$, combine the witnessing policies. For $\subseteq$, a joint action may depend on both components, so the projection is not immediate. Fix an allowed disturbance selection for system 2 and simulate system 2 inside a policy for system 1. The induced policy depends only on system 1's history and keeps it in G_1 against all of its disturbances. This is where nonemptiness and the policy class matter.
Consequence. Shared budgets, coupled timing, or cross-system effects violate the premises. The theorem states exactly what must be established before local guarantees assemble without further analysis.
Status. Hypotheses to fix. Reported check in the companion, with and without coupling.

P27. Circular assume-guarantee composition. [CMP]

Statement. Components carry assumption-guarantee pairs (A_i, G_i) on safety properties, with shared capacity appearing in the interface as a time-indexed assumption on the

combined draw. If each guarantee at step $n+1$ depends only on assumptions through step n, the circular argument is well founded by induction on n, and the conjunction of guarantees holds for the composition (Abadi and Lamport [1]).
Consequence. The temporal discipline is what prevents circular assumptions from justifying themselves. It restricts the argument to safety properties.
Status. Specification. The inductive form is the target statement.

P18. Economic processes form a category when interfaces compose. [CMP]
Hypotheses. Let $\mathcal{P}$ be the full path category: objects are interfaces, morphisms are all histories between them (admissible or not), composition is concatenation. Let Adm be a set of morphisms of $\mathcal{P}$, the admissible histories.
Statement. For parallel morphisms $f, g : A \to B$ define $f \approx g$ if for every pair of composable contexts $u : C \to A$ and $v : B \to D$, $v \circ f \circ u \in \mathrm{Adm}$ if and only if $v \circ g \circ u \in \mathrm{Adm}$. Only parallel morphisms are compared, so composition preserves sources and targets. Then $\approx$ is a congruence for composition, $\mathcal{P}/\approx$ is a category (the syntactic category of Adm) with identities the classes of empty histories, and Adm is a union of $\approx$-classes, so it is recognized within the quotient. The admissible histories themselves need not form a subcategory: they need not be closed under composition, as P05 and P09 show, and constructing the quotient of the full path category avoids reintroducing that closure problem. A monoidal structure needs a parallel composition satisfying the interchange law, which fails under shared resources unless the interface records them (P09, P10). See Mac Lane [14].
Consequence. One-sided continuation equivalence from P01 gives only a right congruence. Quotienting morphisms needs stability under both preceding and following composition, so the categorical chapter begins with this syntactic congruence, or with the transition category of the minimized continuation system.
Status. Hypotheses to fix.

P19. Categorical accounting preserves specified composition laws. [CMP]
Statement. An accounting interpretation is a functor when $F(g \circ f) = F(g) \circ F(f)$ and $F(\mathrm{id}_s) = \mathrm{id}_{F(s)}$. For a chosen accounting system these equations follow from its posting rules. Its limitations are then analyzed through P11: a composition-preserving representation can still forget distinctions needed to assess continuation.
Consequence. Coherent bookkeeping and sufficient economic description are different properties.
Status. Elementary for a given posting system.

P20. Typed execution: preservation and progress. [CMP]
Statement. Let $\Gamma \vdash a : s \Rightarrow s'$ require the stated permissions, account for resource use, and update obligations. If every primitive rule is sound for the transition semantics, sequential composition yields sound judgments by induction on derivations, and well-typed execution maintains a specified invariant (*preservation*). *Progress*: a well-typed commitment

is executable, conditional on available resources, enabled counterparties, and any environmental assumptions needed for execution. Linear or affine disciplines prevent duplication of exclusive resources when the semantics and trusted boundary enforce them.
Consequence. The gap between formal inhabitation and executability is exactly where progress can fail. The theorem does not establish the truth of sensor readings, the legitimacy of obligations, or the completeness of the model.
Status. Preservation elementary. Progress hypotheses to fix.

P21. Compatible local records glue exactly under the sheaf condition. [CMP]
Statement. Let economic contexts form a site and let $\mathcal{F}$ assign records with coherent restriction maps. If $\mathcal{F}$ is a sheaf, every matching family on a cover has a unique global section.
Consequence. The substantive work is proving that a particular record system satisfies the condition. Feasible allocations can fail to form a sheaf when local contexts omit globally shared capacity, as in P09. Gluing records and establishing executable coordination are separate tasks unless the construction links them.
Status. Hypotheses to fix, per record system.

P22. A cohomological obstruction in an affine coordination model. [CMP]
Statement. If compatible configurations form a torsor under an abelian sheaf $\mathcal{A}$, local choices give transition differences on overlaps, forming a degree-one cocycle. A global configuration exists exactly when the torsor class vanishes in $H^1(X, \mathcal{A})$, under the standard classification hypotheses (Mac Lane and Moerdijk [15]).
Worked case, frictionless. On a graph give each oriented edge a log exchange rate c_{ij} with $c_{ji} = -c_{ij}$. A global log-price potential p with $c_{ij} = p_j - p_i$ exists exactly when the sum around every cycle is zero, that is, when the class $[c] \in H^1(\Gamma; \mathbb{R})$ vanishes. A nonzero cycle sum gives an arbitrage direction.
Worked case, frictions. With bid–ask spreads, let u_{ij} be an upper bound on the log rate from i to j. A potential with $p_j - p_i \le u_{ij}$ exists if and only if there is no negative cycle (shortest-path duality, Bellman–Ford), and the vanishing of $[c]$ is no longer required. The frictionless case is the limit $u_{ji} = -u_{ij}$, where the two inequalities force equality.
Consequence. The example gives a genuine obstruction theory for the model. Broader institutional conflicts need separate constructions and must not be described as cohomological without one.
Status. Graph case elementary, with a reported check in the companion. General sheaf case specification.

Where each result appears

Chapter assignments are provisional. For each result, the chapters are grouped by where it is stated early (Parts I to VIII), where it is stated rigorously (Part IX), and where it is

reinterpreted (Parts X to XII). A dash means the band has no assigned chapter.

Result	Family	Early (I–VIII)	Rigorous (IX)	Reinterpreted (X–XII)
P01	REP	2	41	46
P02	INS	4, 39	41	51
P03	VIA	13	43	54
P04	VIA	13	43	50
P05	VIA	15	43	46
P06	VIA	12	43	54
P07	VIA	14	43	53
P08	INS	3, 29	43	53
P09	CMP	31	44	47, 58
P10	CMP	31	44	47
P11	REP	20, 21	41	49
P12	PRD	8, 18	45	49
P13	REP	16	42	48
P14	VIA	15	43	54
P15	REP	35	45	55
P16	CPT	23, 24	45	54
P17	REP	25	45	55
P18	CMP	31	44	46, 47, 60
P19	CMP	20	41	49
P20	CMP	12	–	51, 52, 53, 54, 60
P21	CMP	22	–	56, 57, 58, 60
P22	CMP	17	–	59, 60
P23	VIA	13	43	50
P24	REP	17	44	48
P25	VIA	13	43	–
P26	PRD	9, 40	43	–
P27	CMP	32	44	50
P28	INS	29	45	53
P29	INS	19	45	49

Running example

A housing-maintenance system carries all six families. Deterioration establishes the dynamics. Tenancy and repair duties establish obligations. Contractor availability deter-

mines progress. Budgets introduce scarcity prices. Records test representation sufficiency. Shared crews expose the failure of local composition. The finite-state companion reports checks of the viability and independence results and deliberately reproduces counterexamples when their hypotheses are removed, so that the proof programme is a working specification and not a parallel collection of formal propositions.

Part I

Economic Activity and the Establishment of Futures

Part I fixes what an economy maintains and introduces the objects the rest of the book relies on: states, histories, continuations, and the distinction between what is possible, accessible, permitted, and acceptable.

HOUSING THREAD. Define the building, its tenants, and the repair crew, and list the states and histories that matter: rent arrears, deferred work, and notices served.

The housing thread of this part is carried by a four-unit building. Each unit is let at a monthly rent of 1000 and costs 300 per month to run before repairs. Its physical condition k is measured on a scale of 0 to 10 in steps of one tenth, a unit with condition below 4 cannot lawfully be let, and a single repair crew serves the whole building up to a monthly spending ceiling of 400. The tenants are described by their liquid resources, their monthly surplus, and the costs of moving. The states that matter are the condition of each unit, the cash of the owner and of each tenant, the number of consecutive missed rents, the notices that have been served, and the repair work that has been promised but not yet done. The remaining chapters of the part show why each of these is needed and what is lost when it is omitted.

Chapter 1

What an Economy Maintains

An owner of the four-unit building must fix next year's maintenance budget before the roof, the plumbing, and the tenants' circumstances are known. The owner's accountant reports the building's result as rent less running cost less maintenance spending, so that a unit let at 1000 with running cost 300 and maintenance 100 yields a monthly surplus of 600. If the maintenance line is cut to zero the report shows 700 and the next year's statement looks better by 1200 per unit. Nothing in that statement says whether the building will still be lettable in three years. The decision is hard because the object the owner cares about, a building that can go on producing lettings, is not an entry in any ledger the owner keeps, and because the cost of the cut arrives later than the gain from it, in a form that is not proportional to the cut. The participants are the owner, who knows the budget and the rent roll, the tenants, who know how they use the unit but cannot see the roof, and a repair crew whose capacity is limited. The owner's objective is a long stream of surplus and the constraint is that each unit remain lettable.

This chapter argues that the useful description of an economic activity records two things, the output obtained now and the conditions that the activity leaves behind, and that the second is a state in its own right and not a footnote to the first. Four activities show the range. A household cooks, cleans, and raises children, and what it maintains is health, skills, a dwelling, and the household's own capacity to continue the same routines tomorrow. A construction firm completes a building, and what it maintains is its equipment, its crews, its bonding capacity, and a reputation that allows it to win the next contract. A hospital treats patients, and what it maintains is trained staff, sterile theatres, supplies, and a referral network that fails slowly and recovers slowly. A watershed supplies water, fish, and timber, and what it maintains is soil, forest cover, and groundwater, which a harvest can draw down without any entry on any account. In every case the quantity sold or consumed in a period, which the national accounts call output, is the visible part of a larger process, and the larger process includes the reproduction of the conditions under which output can be repeated.

1.1 A transaction that produces income and uses up capacity

The housing case gives the simplest form of the point. Let the condition of a unit be k_t at the start of month t, measured in tenths of a condition point so that the starting value is 90 and the lettable threshold is 40. Deterioration per month is 1 tenth while $k > 80$, 2 tenths while $50 < k \leq 80$, and 4 tenths while $k \leq 50$. Spending M on maintenance restores $M/100$ tenths per month, so $M = 100$ exactly offsets deterioration at the top tier. These numbers are chosen so that every statement below can be computed by hand and are reproduced by the companion (check H3). With $M = 100$ each month the condition never moves from 90 and the monthly surplus is $\pi = 1000 - 300 - 100 = 600$. With $M = 0$ the monthly surplus is 700 and the condition follows the sequence $90, 89, \ldots, 81, 80$ over the first ten months, then falls by two each month to 50 at month 25, then by four to 46, 42, and 38 at months 26, 27, and 28. The unit is lettable in months 0 through 27, which is 28 rent months, and the cumulative surplus over those months is $28 \times 700 = 19{,}600$. At month 28 the condition 38 is below the threshold and the unit earns nothing until it is repaired.

Restoring the unit to 90 costs 100 per tenth, which is 5200 from a condition of 38. The owner saved 100 per month for 28 months, which is 2800, and so the cut is paid for twice over by the restoration bill, and the loss is 2400 before counting the rent forgone while the unit is unlettable. The pattern is not an accident of the numbers. In the first tier, deterioration equals the maintenance saved tenth for tenth, so deferral for up to ten months costs nothing in net terms; the net value of deferring n months is 0 at $n = 10$, -1000 at $n = 20$, -1500 at $n = 25$, and -2400 at $n = 28$. A manager who observes only the first ten months of the experiment will see a free gain, and a manager who observes the first twenty-eight will see that the free gain was a loan whose terms worsened with every month the loan was held.

Definition 1.1 (Income and capacity)**.** An *activity* in period t consists of a flow vector of receipts and costs, whose sum is the period surplus π_t, and a change in a vector k_t of *capacity variables*, the stocks of condition, skill, access, and entitlement on which later flows depend. Two activities are *financially equivalent* if they have the same π_t and *capacity equivalent* if they leave the same k_{t+1}.

The definition is followed by an example and a boundary case. The example is the pair of maintenance policies $M = 100$ and $M = 0$ above: they differ in π_t by 100 and in k_{t+1} by 1 tenth at the top tier, and they are therefore neither financially nor capacity equivalent. The boundary case is a pure financial claim, such as a government bond held for a period, whose purchase and sale change π_t and cash but leave the conditions of production untouched. Here the capacity variable is trivial and the two notions of equivalence coincide. The point of the definition is that the two coincide rarely, and that when they diverge the divergence is the economic content.

Proposition 1.2 (Wedge between surplus and sustainable surplus)**.** *Consider an activity with revenue $R(k)$ that is nondecreasing in capacity, maintenance cost M, and capacity dynamics $k_{t+1} = k_t - d(k_t) + g(M_t)$ with g strictly increasing. For any baseline M^b and any reduction $\delta > 0$ with $g(M^b - \delta) < g(M^b)$, surplus rises by δ in the period of the reduction and capacity next period is strictly lower than under the baseline. If in addition R is strictly increasing at the baseline capacity, revenue in the following period is also lower or equal, and strictly lower whenever the reduced capacity falls in an interval where R is strictly increasing.*

Proof. Surplus is $R - C - M$ with R and C unchanged in the period of the reduction, so it rises by δ. Next period capacity is $k - d(k) + g(M^b - \delta)$, which is strictly lower than $k - d(k) + g(M^b)$ because g is strictly increasing. Revenue is a nondecreasing function of capacity, so it is not higher, and it is strictly lower if R is strictly increasing between the two values. □

The proposition is elementary and it is stated in the proof programme as P12(a). It is recorded here because it licenses a specific criticism of income accounting and not a general suspicion of it: a period surplus is a correct measure of what was received less what was spent, and it is a wrong guide to action whenever the action changes the capacity that produces the next period's receipts. What the proposition leaves open, and what Chapter 8 takes up, is how large the later loss is, because that depends on how deterioration behaves, and in the housing numbers the loss is zero for ten months and then grows. An assumption doing real work is that g is strictly increasing over the relevant interval; if some of the maintenance budget bought nothing, a cut of that part would be a pure saving, and a good deal of ordinary budget-trimming is exactly that.

1.2 Relation to established accounting and capital theory

The distinction between income and the preservation of capital is old. Hicks defined income as the amount a person can consume in a period while leaving themselves as well off at the end as at the start [12], and a long line of work on capital maintenance, environmental accounting, and the measurement of net product follows from asking what the phrase "as well off" should mean. The present book reproduces the Hicksian point in the proposition above, in a notation that keeps the capacity variable explicit. It does not follow Hicks in condensing capacity into a single wealth number, because Chapter 16 shows that capacities of different kinds are not always comparable on a common scale. The capability approach of Sen [19] shares with the present account the refusal to read a person's situation off their holdings of commodities, and Chapter 5 draws on it. The book differs from that approach in asking what it takes to keep a capability available through time against specified disturbances, a question about trajectories and not only about a snapshot of opportunity.

1.3 What the owner now decides

Return to the owner. The owner who reads only the accountant's statement sees a surplus of 700 against 600 and cuts maintenance. The owner who adds the capacity variable sees that the cut is free for ten months and costly afterward, that the building will be unlettable in month 28 unless the work is restored, and that restoring costs 5200 against savings of 2800. The decision is now a comparison of trajectories, and the question is how long the owner can defer without crossing a point at which the restoration is no longer possible. That question needs a model of the crew's capacity, of the owner's cash, and of the tenants' obligations, and the next four chapters construct the pieces from which it is answered.

What has been established. Income and capacity are distinct, an activity can raise one while lowering the other, and the housing numbers show a case where the apparent gain is a loan on increasingly unfavourable terms. The decision now enabled is a trajectory comparison for the maintenance budget. The assumptions still required are the form of the deterioration function and the effect of spending on condition, both of which are empirical. The question passed on to Chapter 2 is what must be recorded about the past for the future questions to be answerable, since the accountant's statement plainly does not suffice.

Exercise 1.3. For a unit starting at condition 90 with the tiers above, compute the net value of deferring maintenance for $n = 15$ months and then restoring. (The answer is the saved spending 1500 less the restoration cost to condition 90.)

Exercise 1.4. Prove that the net value of deferral $V(n) = 100n - 100(90 - k_n)$ is nonincreasing in n for the stated tiers, and identify the first month at which it becomes strictly decreasing.

Exercise 1.5. Choose one of a clinic, a school, or a small farm. List the capacity variables its annual accounts omit, state for each how it deteriorates and what restores it, and say which of the three omitted variables you would measure first and why.

Chapter 2

States, Histories, and Continuations

A tenant has missed one rent. The landlord's lease says that a notice to end the tenancy is served after two consecutive missed rents, and that paying the arrears within a month of notice cancels it. A second tenant in the building also holds 500 in cash and has never missed a rent. At a glance the two are alike, and a lender or a statistical model that sees only current holdings would treat them alike. The decision for the landlord is whether to extend a credit, to schedule a repair visit that will cause the tenant inconvenience, or to start planning for vacancy. The difference between the two tenants is invisible in their holdings and decisive for what can happen next, and the chapter makes that difference into a definition.

2.1 States and histories

An economic *state* is a description of the situation sufficient for the questions at issue. An *event* is something that happens, an action taken or a disturbance realized, and a *history* is a finite sequence of events. A model specifies which histories are possible, which futures are open after each, and which of those futures meet the stated obligations. The basic object in the book is the set of possible trajectories, in contrast to a single forecast. In the common model of the notation chapter, $x_{t+1} \in T(x_t, h_t, a_t)$ and $h_{t+1} = U(h_t, x_t, a_t, x_{t+1})$ so that a history summary h travels with the physical state and is updated by an explicit rule.

Take the missed-rent rule. Let the events be p (rent paid) and m (rent missed), and let the *continuation language* $L(h)$ be the set of future event sequences that do not lead to notice. For a tenant with no recent misses, L contains every sequence with no two consecutive m's. For a tenant whose last event was m, L contains every sequence that begins with p and then continues as for the first tenant. The two languages differ, so the two tenants differ in a precise sense even when their cash is the same, and the difference

is a count of consecutive misses that is 0 for one and 1 for the other. A third class, with count 2, is the set of histories after which notice has been served, and its continuation language is empty for the purpose of avoiding notice.

Definition 2.1 (Exact summary and continuation equivalence)**.** Let H be the set of histories and $L(h)$ the set of future sequences u such that hu meets the specified obligations. A function $q : H \to Q$ is *exact for admissibility* if, for all h and u, whether $u \in L(h)$ is determined by $q(h)$ and u. Write $h \sim h'$ if $L(h) = L(h')$.

An example is $q(h)$ = the number of consecutive misses at the end of h, capped at 2. It is exact for the notice rule and it has three values. A boundary case is $q(h)$ = the tenant's cash balance, which is not exact because two histories with equal cash and different counts have different languages. A second boundary case is the full history itself, which is trivially exact and useless as a summary because it grows without bound. The question the next theorem answers is how much can be thrown away.

Theorem 2.2 (P01: coarsest exact summary)**.** *(i) If q is exact for admissibility then $q(h) = q(h')$ implies $h \sim h'$. Consequently $\sim$ is the coarsest exact summary, and $q_\sim(h) = [h]_\sim$ is itself exact. (ii) If the language has the residual property $L(he) = e^{-1}L(h)$ for every event e, where $e^{-1}L = \{u : eu \in L\}$, then $\sim$ is a right congruence: $h \sim h'$ implies $he \sim h'e$ for every e, so updates are well defined on classes.*

Proof. (i) If $q(h) = q(h')$ then for every u the statement $u \in L(h)$ is determined by $q(h)$ and u, which is also $q(h')$ and u, so $L(h) = L(h')$. The class map $[\cdot]_\sim$ is exact by definition of $\sim$. Any exact q therefore refines $\sim$, so $\sim$ is the coarsest. (ii) If $h \sim h'$ then $L(he) = e^{-1}L(h) = e^{-1}L(h') = L(h'e)$, hence $he \sim h'e$. □

The theorem is the dynamic form of a factorization property that Chapter 18 states for accounts in general as P11. Its hypotheses are light. The residual property is what makes a history summary behave like a state; it holds whenever $L(h) = \{u : hu \in \Lambda\}$ for a fixed language Λ of acceptable full histories, which is the natural case when the obligations are fixed in advance. It fails when the rules themselves change in a way not recorded as an event, for example when a legislature raises the number of misses that triggers notice without that change being part of any history, and in that case the classes are not closed under updates and the state has to be enlarged to include the rule in force. When the number of classes is finite the minimal summary is a finite automaton, and this is the content of the Myhill–Nerode construction [17]; when it is infinite, as with an obligation that depends on an unbounded running total, an exact finite summary does not exist and the model has to carry a numerical state.

The economic reading is separate from the theorem. Debt, deferred maintenance, accumulated exposure, and contractual commitments can be dropped from a state description only if they do not change which futures are acceptable. Two households with the same cash and different arrears are different states for any question about whether they can

continue, and two buildings with the same present condition and different promised-but-undone repairs are different in the same way. The theorem concerns exact recognition of acceptable continuations. Predictive sufficiency for probabilities or for welfare needs further conditions, and no claim is made here that the coarsest exact summary for admissibility is also sufficient for forecasting.

2.2 Cash and arrears together

Now add cash. Let the rent be 1000, let the tenant's net monthly income after living costs but before rent be 1200, and let the tenant start with a balance x. Paying is possible when the balance after receiving income is at least 1000. A tenant at $x = 500$ can pay this month, ending with $x + 1200 - 1000 = 700$. A tenant at $x = -300$ cannot pay without borrowing: the balance after income is 900. The relevant summary is the pair (c, x) of consecutive misses and balance, and Theorem 2.2 says the pair can be collapsed further if two pairs have identical languages. For instance, every pair with $x \geq 0$ and $c = 0$ is equivalent in the first month, since the tenant can pay and does not need to miss, but they are not equivalent over longer horizons if the income is shocked, a point developed in Part III. The state is therefore a function of the question as well as of the system.

What has been established. A state description for economic continuation must include the features of the past that change what can happen next, and the minimal such description is the quotient by continuation equivalence. The decision now enabled is the choice of which variables to carry: arrears count, cash, outstanding repairs, and notices, but not features that leave every continuation language unchanged. Assumptions still required are the residual property and a finite number of classes. The question passed on to Chapter 3 is that a continuation language only says which futures meet the obligations; it does not say which actions the person can actually take, and the two are often confused.

Exercise 2.3. Write the continuation language and the minimal summary for a rule under which notice follows three missed rents in any window of six months. Is the summary finite?

Exercise 2.4. Show by an explicit pair of histories that the cash balance alone is not exact for the two-consecutive-misses rule, and give a rule change under which the residual property fails.

Exercise 2.5. Identify a regulated service, such as a credit file or a medical record, in which the past is summarized by a score. State what continuation questions the score is exact for and a pair of histories with equal scores and different continuations.

Chapter 3

Feasibility, Accessibility, and Admissibility

A tenant in the building wishes to leave. The lease permits departure on one month's notice, the tenant knows of a cheaper unit across town, and nothing in the law forbids moving. The tenant has 3000 in savings and a monthly surplus of 200. Moving costs, a new deposit, and the first month's rent come to 4000. The tenant is permitted to move, physically able to move, and aware of a place to go, and still cannot do it for another five months. Four things have been run together in ordinary speech under the word "can": what is physically possible, what is practically accessible, what is institutionally permitted, and what is consistent with the obligations one has taken on. The chapter separates them because each fails independently and the failures have different remedies.

3.1 Four sets of actions

At state s, let $A_{\text{phys}}(s)$ be the actions not excluded by physical or technical laws. Let $A_{\text{acc}}(s) \subseteq A_{\text{phys}}(s)$ be the actions the person can in fact take given resources, knowledge, time, and standing, so that a fee that cannot be paid removes the action. Let $A_{\text{perm}}(s)$ be the actions permitted by law, contract, and custom. Let $A_{\text{eff}}(s) = A_{\text{acc}}(s) \cap A_{\text{perm}}(s)$ be the effective set, the actions that are both takeable and allowed. Finally let $A^{G}_{\text{acc}}(s)$ be those effective actions whose successors remain in the acceptable set G, the actions consistent with stated obligations. Acceptability is itself a judgement about continuation, and Chapter 13 shows that the set of actions with acceptable successors is still too large, because an action can lead to a state that is acceptable but not viable.

In the housing example the monthly lease gives a permitted exit with fee zero, so A_{perm} contains moving, and A_{phys} contains it too. The accessibility condition is the resource threshold $x + te \geq c$ at the exit date t. For $x = 3000$, $e = 200$, and $c = 4000$ the smallest t with $3000 + 200t \geq 4000$ is $t = 5$, so exit is effective only from month 5. A tenant whose resources are 3400 can exit by month 3. This is the same calculation as the front chapter,

and it is the first appearance of what the proof programme calls P08.

Definition 3.1 (Executable exit)**.** Let X be an exit target, B the states satisfying the safety and obligation constraints during departure, and define $\text{Reach}_0 = X$ and $\text{Reach}_{k+1} = \text{Reach}_k \cup \{s \in B : \exists a \in A_{\text{eff}}(s)\ \forall w,\ T(s,a,w) \in \text{Reach}_k\}$. An exit is *executable within horizon* H from s if $s \in \text{Reach}_H$.

Theorem 3.2 (P08: nominal and executable exit differ)**.** *(i) A state belongs to* Reach_H *if and only if there is a policy using only effective actions that reaches* X *within* H *steps while remaining in* B *against every allowed disturbance. (ii) Legal permission to exit does not imply membership: there is a model in which the exit action is in* $A_{\text{perm}}(s)$ *at every* s *and* $\text{Reach}_H = X$*, so that no state outside* X *can exit.*

Proof. (i) By induction on k: $s \in \text{Reach}_{k+1}$ means that $s \in X$, or that $s \in B$ and some effective action sends every allowed successor into Reach_k, which by hypothesis admits a policy reaching X within k further steps. (ii) Let the exit fee be c, let resources never exceed $c-1$, and let X be the set of states in which the fee has been paid. The exit action is permitted everywhere and accessible nowhere outside X. □

The hypothesis carrying the conclusion is that accessibility is determined by resources. The simplification is a deterministic income stream; the assumption most likely to fail in applications is that the fee c is known, because exit costs frequently include search time, references, and deposit rules that the tenant discovers at the point of exit. The economic interpretation is that exit capacity is an operational variable, whose distribution across people can be studied separately from whether contracts contain an exit clause. A lease that says exit is free and a lease that says exit costs 4000 differ in the second case by a charge, but a lease that says exit is free for a tenant who cannot cover the next deposit has only a nominal exit clause.

3.2 Examples across domains and a boundary case

The same four sets separate in employment, in medical treatment, and in credit. A worker is permitted to take a job in another city, physically able, and not able to pay the first month's housing there before the first pay arrives. A patient is entitled to a treatment, which exists, and cannot reach the clinic. A borrower is eligible for a loan product and has no collateral to meet its terms. In each of these the formal list of options is the same as for a wealthier person and the effective set is smaller. A boundary case is the opposite error, treating whatever is accessible as acceptable: a tenant may be able to pay rent by drawing the moving fund, which is accessible and permitted, and which satisfies this month's obligation while leaving the tenant's exit unexecutable. This is the sense in which an available choice need not be an acceptable continuation, and it is the first sign of the gap formalized in P05 in Chapter 13.

What has been established. Effective actions are the intersection of the accessible and the permitted, nominal exit and executable exit differ, and the difference can be calculated from resources. The decision now enabled is to ask, for any institution that advertises a right, what must be held to use it. Assumptions still required are the resource model of access and a known fee. The question passed on to Chapter 4 is who bears the cost of making an alternative accessible in the first place.

Exercise 3.3. For $e = 200$, $c = 4000$ and $H = 3$, find the minimum initial resources for an executable exit on the monthly contract, then recompute when a once-only cost of 500 is due at the end of month 1.

Exercise 3.4. Prove that if $A_{\text{eff}}(s) \subseteq A'_{\text{eff}}(s)$ for all s then $\text{Reach}_H \subseteq \text{Reach}'_H$, and state the economic meaning.

Exercise 3.5. For a public service you know, such as a housing voucher or a vaccine appointment, list the four sets and estimate which of them excludes the largest share of eligible people.

Chapter 4

Establishing the Domain of Choice

A tenant in the four-unit building pays 1000 and has a budget of 500 set aside. A rental listing service charges 150 for an introduction that reveals three alternatives, with monthly rents of 800, 950, and 1200. Without the introduction the tenant knows of no alternative to the present unit. Moving to any alternative costs 400. The decision is whether to buy the introduction, and what makes it hard is that the benefit of selecting among the alternatives is available only after the cost of establishing them has been paid, and that selecting is itself expensive. A standard model of choice begins with a set of options and asks which is best. This chapter asks where the set came from.

4.1 A two-stage model

Let D be an opportunity domain, the set of alternatives available to choose among. An *acquisition action* e has a cost $\kappa(e)$ and produces a larger domain $E(e, D) \supseteq D$. Selection inside a domain D' has a cost $\sigma(x)$ for each chosen outcome x, and $C(D')$ denotes the outcomes obtainable by selection, with their costs. The tenant's problem is to pick e and then $x \in C(E(e, D))$, with total cost not exceeding the budget b.

For the tenant, $D = \{\text{stay}\}$ with selection cost 0. The introduction is an acquisition action with $\kappa = 150$ and $E(e, D) = \{\text{stay}, 800, 950, 1200\}$. Choosing the unit at 800 costs $\sigma = 400$ and saves 200 per month. The total up-front outlay is 550, so the smallest budget at which the outcome "live at 800" is reachable is $b^* = 550$, and the payback period is $550/200 = 2.75$ months, so the move is worthwhile if the tenant expects to stay at least three months in the new unit. For a tenant with the stated budget of 500 the introduction is affordable and the move is not: after paying 150 the tenant has $350 < 400$. Buying the introduction in that case wastes 150, and a good adviser would say that the first stage must be evaluated with the second in view.

Theorem 4.1 (P02: establishment precedes selection)**.** *(i) If $x \notin C(D)$ and $x \in C(E(e, D))$, selection within D cannot implement x whereas establishment followed by selection can. (ii) Suppose the action set includes a zero-cost identity establishment action e_0 with $E(e_0, D) =$*

D. Then for any budget b the outcomes implementable by establish-then-select policies contain those implementable by selection only, and the inclusion is strict if and only if some establishment action with $\kappa(e) \leq b$ makes an outcome reachable with total cost within b that was not reachable within D.

Proof. (i) is the definition. (ii) A selection-only policy is the establish-then-select policy that uses e_0, so the containment holds. If the inclusion is strict, some outcome x is implemented by a policy with some $e \neq e_0$, whose total cost is within b and which makes x available, and which is not reachable within D. Conversely any such e and x give a strictly larger set. □

The hypothesis carrying the conclusion is the zero-cost identity action. If establishment is compulsory, as when a licence or registration must be bought before any selection is allowed, the containment can fail, and the counterexample is a market in which a mandatory fee exceeds the budget of a participant who could otherwise have chosen among the free alternatives. For the tenant of the example the budget analysis is as follows. At $b = 400$ the tenant cannot move to 800 at all and the identity action is the only sensible policy. At $b = 500$ the introduction can be bought and the move cannot be made. At $b = 550$ the move is reachable, and the strictness of the inclusion begins exactly there.

The interpretation is economic and not logical. Expenditure on search, on education, on infrastructure, on legal recognition, and on interface design can change the choice problem itself. These are costs of making alternatives available, and they are paid by someone. When a city builds a transit line it establishes a domain of commuting options for many households at once, at a cost that no household would pay alone. When a platform lists apartments it establishes a domain, and the fee structure determines who can afford to see it. A familiar analysis would treat the cost of search as a transaction cost within a given model of choice. The point here is slightly different: the domain is itself a variable of the model, and some populations differ from others in the domains that have been established for them, which is a form of inequality that a model beginning with a fixed choice set cannot represent.

What has been established. Establishment and selection are distinct expenditures, the break-even budget for the example is 550, and the budgeted inclusion has a precise condition for strictness. The decision now enabled is to evaluate any proposed access intervention by what it adds to the established domain of the people it targets and not only by what it costs per head. Assumptions still required are that establishment is optional, that its cost is observed, and that selection costs can be computed. The question passed on to Chapter 5 is how to judge whether a larger established domain makes a life better.

Exercise 4.2. Recompute the break-even budget and payback period if the introduction fee is 250 and the unit at 800 is replaced by one at 850.

Exercise 4.3. Give an example in which a mandatory establishment fee makes the set of

implementable outcomes smaller than under selection only, and state which hypothesis of Theorem 4.1(ii) fails.

Exercise 4.4. Choose a public programme such as a library, a language class, or a bus line. State what domain it establishes, who pays, and which population's domain it enlarges most.

Chapter 5

Welfare Without Counting Options

A city council is deciding whether to permit landlords to offer prepaid leases of the kind in the front chapter. The proposal adds an option for tenants who can afford it and, according to supporters, adds choices for everyone. Opponents say that more options are not a gain for a person whose exit is lost. The council needs a way to compare outcomes that does not reduce to counting choices, does not hide conflict among the people affected, and does not pretend that a single number measures a life. This chapter sets out the normative foundations that the rest of the book assumes and says where they stop.

5.1 Why counting options fails

Suppose welfare is measured by the number of trajectories open to a person. The count is not invariant to how options are individuated: listing five hundred near-identical apartments multiplies the count without adding any real alternative. The count also ignores that some options are worth less to a person who bears a burden in using them, and it ignores that some constraints protect. A tenant protected against eviction for the first year has fewer open trajectories than one without that protection, and many people would say the protected tenant is better placed. Equally, a worker bound by a non-compete clause has fewer options and is not thereby better off. Therefore neither the count of options nor its absence tells us what is good.

The alternative developed here is a *plural assessment*. A continuation profile for a person has coordinates for substantive achievement (what the person is able to do and be, including housing of adequate condition), security (the likelihood that this persists against specified disturbances), burden (the effort and cost of maintaining it), access (what the person can reach when circumstances change, which is the measure of Chapters 3 and 4), and corrective freedom (the capacity to change course without catastrophic loss, taken up in Chapter 15). The coordinates are not claimed to be exhaustive and the choice of them

is a normative commitment that the reader may reject; the book asks only that any rival list be stated as explicitly.

Theorem 5.1 (P13: scalar representation fails under incomparability)**.** *(a) If a partial order $\preceq$ on profiles contains incomparable elements x, y, there is no real-valued v such that $x \preceq y \Leftrightarrow v(x) \le v(y)$ for all pairs. (b) Every partial order on a finite set admits a multi-utility representation: a finite family (v_i) of real functions such that $x \preceq y$ if and only if $v_i(x) \le v_i(y)$ for every i. The minimal size of such a family is the order dimension of Dushnik and Miller [9].*

Proof. (a) Since x and y are incomparable, neither $x \preceq y$ nor $y \preceq x$, but for reals either $v(x) \le v(y)$ or $v(y) \le v(x)$, so the equivalence fails for one of the pairs. (b) For a finite partial order the intersection of its linear extensions is the order itself, because for any incomparable pair there is an extension that places each above the other. Each linear extension ℓ_i gives a real function v_i by position, and the family has the stated property. The minimal number of extensions is the dimension. □

The consequence is moderate. Prices, utilities, and monotone indicators remain useful and are not ruled out. What is ruled out is an exact scalar representation of all the distinctions the assessment records. In the infinite case, the existence and size of a multi-utility representation need further hypotheses, which Chapter 16 states. For the council, the consequence is that a recommendation must say which coordinates it favours and must show the profiles of the groups affected rather than a single total.

5.2 Conflict between continuations

Suppose the building's owner holds 1500 in cash and pays a mortgage of 700 per month. A tenant stops paying. The tenant's protection requires two consecutive missed rents before notice, and notice takes a further month to run, so that the owner's exposure is three months. Three months of the mortgage is 2100, which exceeds the owner's 1500. The tenant's security and the owner's solvency are both coherent obligations, and for this owner they cannot hold together: the set of trajectories that satisfy both is empty from the state of the example. In the language of Chapter 13 the owner is outside the viability kernel for the joint regime. Protecting the tenant restricts the owner, and relaxing the protection restricts the tenant. The conflict is not a mistake in the model.

Three responses are available and each distributes burdens differently: the owner may be required to hold reserves of 2100, which burdens small owners disproportionately; a public arrears fund may cover the gap, which moves the burden to taxpayers; or the tenant protection may be shortened, which moves the burden to tenants. The plural assessment lets each of these be reported as a change in the profiles of identified people, and it makes explicit that the choice among them depends on how the coordinates are weighed. The disagreement is thereby displayed as a disagreement about weighing, instead of being concealed in the choice of a single metric.

What has been established. Counting options does not measure welfare, a plural profile does not reduce to a scalar when its coordinates are incomparable, and protections for one party can make the position of another unviable. The decision now enabled is to evaluate policies by the profiles of the people they change and to name the weights used when a ranking is made. Assumptions still required are the choice of coordinates and their measurement. The question passed on to Part II is how production, the activity that supplies the housing, itself changes the conditions on which every profile depends.

Exercise 5.2. For the owner with 1500 in cash and a 700 mortgage, find the smallest reserve under which three months of exposure can be met and the longest protection period a reserve of 1500 can support.

Exercise 5.3. Give a four-element partial order with order dimension two and exhibit the two linear extensions that represent it.

Exercise 5.4. Select a policy that creates an additional option for some households. State which coordinates of the profile it raises, which it lowers, and for whom.

Part II

Production, Maintenance, and Irreversibility

Part II treats production as the transformation of conditions, capital as maintained capability, and irreversibility as a symmetric property of action and inaction.

HOUSING THREAD. Roof and plumbing deteriorate; producing a habitable unit transforms the building's conditions; the cost of deferred maintenance is traced forward.

The roof and plumbing of the four-unit building deteriorate on the schedule fixed in Part I, a renovation changes a unit's condition by an amount that can be bought at a known price, and the question of this part is what it means to produce something when production itself changes the conditions of later production.

Chapter 6

Production as Transformation of Conditions

An owner has a unit at condition 50 and two bids to renovate it to condition 80. The first contractor charges 3000, which is 100 per tenth for 30 tenths, and offers a six-month warranty. The second charges 2400 and offers none. Both promise a renovated unit, and a record of the work would show the same output, one unit raised to 80. The decision is which bid to accept, and it is hard because the difference between them cannot be seen in the output and appears only in how the unit behaves afterward. A production function that maps inputs to an output quantity records that two renovations were produced; the question the owner faces is about what each leaves behind.

6.1 The housing transition

Every later calculation in the book uses one transition rule for a unit, fixed in the shared specification `spec/housing_v2.json` and implemented unchanged in the companion and in the simulator. The state is the condition k_t in tenths of a condition point, with $k_{\max} = 100$ and a lettable minimum of 40. Within month t the order of events is fixed. The manager observes k_t and chooses a maintenance spend M_t, a multiple of 100 no larger than the crew ceiling, which is 400 for the building as a whole. If $k_t \geq 40$ the unit is let, so that rent 1000 is collected and running cost 300 is paid, and if $k_t < 40$ neither occurs. The spend M_t is paid. Finally $k_{t+1} = \min\{100,\ k_t - d(k_t) + M_t/100\}$, where $d(k) = 1$ for $k > 80$, $d(k) = 2$ for $50 < k \leq 80$, and $d(k) = 4$ for $k \leq 50$. The monthly surplus is $\pi_t = 700 \cdot \mathbf{1}[k_t \geq 40] - M_t$. Deterioration and repair are simultaneous within the month and condition is observed before the choice, so a repair bought in month t affects the condition seen in month $t + 1$ and the rent of month t depends on the condition at its start.

Definition 6.1 (Housing transition)**.** A *housing unit* is the system with state $k \in \{0, \ldots, 100\}$, action $M \in \{0, 100, \ldots, 400\}$, acceptable set $G = \{k \geq 40\}$, deterministic successor

$T(k, M) = \min\{100, k - d(k) + M/100\}$, and surplus $\pi(k, M) = 700 \cdot \mathbf{1}[k \geq 40] - M$. A building of several units shares the ceiling: the sum of the spends on all units in a month is at most 400.

An example is the unit at 90 under $M = 100$: it earns $\pi = 600$ in month 0 and stays at 90 because $d(90) = 1$ is exactly offset by the repair. Under $M = 0$ the same unit earns 700 and moves $90, 89, 88, 87, \ldots$. A boundary case is a unit at 38: it earns nothing whatever the spend, so $\pi = -M$, and a spend is a pure restoration outlay. Because the production of a month of housing is exactly this transition, the owner's output (a let unit-month) and the change in capacity (the next condition) are produced by the same act, and the two sides of the accountant's statement come from one rule.

6.2 Outputs, inputs, residuals, obligations, and capacity

Represent production in a period by five parts. There are the *inputs* consumed, labour, materials, energy, and the use of existing capacity. There are the *outputs* delivered, goods and services that others receive. There are *residuals*, material and informational leftovers whose disposition is a further task (Chapter 10). There are *new obligations*, such as warranties, safety duties, and regulatory filings that outlast the transaction. And there is the *change in productive capacity*, the difference between the stock of conditions that supports production after the period and before it. The conventional production function $Q = f(K, L)$ records inputs and the output quantity and leaves the other three parts implicit or absent. It answers questions such as how much can be produced from a given endowment in one period, and it is silent on whether the same endowment will be there next period.

Suppose the second contractor leaves a hidden defect, a poorly sealed pipe joint, that raises the deterioration in the middle tier ($50 < k \leq 80$) from 2 to 3 tenths per month. Under no further maintenance, the unit from condition 80 with sound work lasts 18 rent months before falling below 40, since the path reaches 50 after 15 months and then loses 4 per month for three more, and the unit with the defect lasts 13 months, since it reaches 50 after ten months and then follows the same last three steps. The defect costs five rent months. At the no-maintenance surplus of 700 per month this is a loss of 3500 against a saving on the bid of 600. The production record shows two identical outputs and a saving; the transition description shows a loss nearly six times the saving. The calculation is reproduced in the companion (check BOOK, Chapter 6).

Definition 6.2 (Transition description of production)**.** A production activity from state s is a tuple (u, y, ρ, ω, s') of inputs u, outputs y, residuals ρ, new obligations ω, and successor state s', subject to a transition relation $s' \in T(s, u)$. A *production function* is the restriction to the map $u \mapsto y$ when the other components are constant or ignored.

An example is the renovation above, in which ω is the warranty (present in the first bid, absent in the second) and the successor state includes the defect. A boundary case is

the production of a standardized good, such as a bolt of cloth in a factory whose machines are in steady state, where the residual, obligation, and capacity components are constant across alternatives and the production function loses nothing. The boundary shows that the richer description is a generalization, and that it matters when the omitted components vary across the alternatives being compared.

6.3 What each representation can answer

The comparison with the production function can be made precise. A production function answers questions that depend on (u, y) only: the efficient mix of inputs for a given output, the returns to scale, the marginal product of labour. It cannot answer a question whose answer depends on ρ, ω, or s', such as which of two equal-output methods leaves the unit lettable for longer, or which leaves the lower exposure to a repair claim. By the factorization of Chapter 2 the production function is a sufficient description for exactly those questions that factor through (u, y); for the owner's decision it is not sufficient because the bids agree on y and differ on ω and s'.

The transition description is not free. It needs a model of how conditions evolve, and a poor model of deterioration is worse than none for decisions that hinge on it. The standard literature on production correctly supplies, for questions about prices and allocation, an economy of tractable models; this book does not displace those models and supplies the extension needed when the questions concern continuation. Where the two give the same answer the production function is preferable for its simplicity, and where they differ the difference is exactly the object of study.

What has been established. Production is better described by inputs, outputs, residuals, obligations, and capacity change than by the output quantity alone, and the housing example shows a hidden defect costing nearly six times the saving that hid it. The decision now enabled is to compare bids by their effect on the successor state. Assumptions still required are a deterioration model and a way to observe defects. The question passed on to Chapter 7 is how to value the capability that production maintains or consumes.

Exercise 6.3. For a unit at condition 80 and a defect that raises middle-tier deterioration from 2 to 4, compute the number of rent months with no maintenance and the loss relative to the sound renovation.

Exercise 6.4. Show that the production function is sufficient for the question "which bid maximizes output" and not for "which bid leaves the larger lettable horizon", by giving two states with the same (u, y) and different horizons.

Exercise 6.5. Choose an industry, such as software or food service, and list its residual and obligation components. State which are most often omitted from its published productivity statistics.

Chapter 7

Capital as Maintained Capability

Two units in the building have the same condition, 60, the same walls and fittings, and the same lease. One is served by a crew with the full ceiling of 400 per month. The other can call on a repairer who will spend at most 100 per month on it, because the rest of the crew's time is committed. A buyer who inspects both finds the same asset. The decision for an owner choosing which to acquire or which to fund is what each is worth, and the difficulty is that the physical asset is identical while the capability it provides over the next three years is not.

7.1 Identical assets, different complements

Capital goods, in the sense used here, are machines, buildings, knowledge, networks, and institutional arrangements that provide a capability on the condition that complementary things continue to be provided. The value of a capital good to its owner is the value of a flow of services conditioned on a stream of complements, and it depends on repairability, on the skills available, and on access to the complements in the market. Possession is a necessary part of this and not the whole.

Apply the transition of Chapter 6 to the two units at $k = 60$ over a horizon of 36 months. With the full ceiling the manager spends $d(60) \times 100 = 200$ per month, which holds the unit at 60 indefinitely, so the unit is let in all 36 months at a surplus of $700 - 200 = 500$ per month, a total of 18,000. With the ceiling of 100 the best the manager can do is spend 100 every month, giving a net change of -1 tenth per month in the middle tier. The unit reaches 50 after ten months and then falls by 3 per month, to 47, 44, 41, and 38, and it is let in 14 months at 600 per month, a total of 8400, with no rent afterward. The same asset is let for the whole horizon in one case and for 14 months in the other.

A third case shows an exposed dependence. Suppose the full ceiling is available for twelve months and then withdrawn, for example because a vendor ends its parts support, so that no maintenance can be bought afterward. The manager spends 200 per month for twelve months, the unit stays at 60, and then with $M = 0$ it falls 2 per month to 50 after

five months and by 4 per month to 46, 42, 38 in three more. The unit is let for $12+8=20$ months and earns $12\times 500+8\times 700=11{,}600$. The three totals 18,000, 8400, and 11,600 are produced by identical assets under three different supplies of the complement, and they are reproduced in the companion (check BOOK, Chapter 7).

Definition 7.1 (Usable capability)**.** Let $y(k,z)$ be the service flow, here the number of let months over a stated horizon, that the owner can obtain from an asset in condition k when the complement z is available, where z is the profile of maintenance spending the owner can in fact buy (a ceiling, possibly changing over time). Two assets are *physically identical* if their conditions coincide and *capability equivalent* if $y(k,z)$ coincides for every complement profile z in the relevant set. A complement is *exposed* if one supplier controls its availability and *controllable* otherwise.

The three runs above show physically identical assets that are not capability equivalent. A boundary case is an asset with no maintenance requirement, such as bare land, for which y does not depend on z and the physical and capability descriptions coincide.

Proposition 7.2 (Usable capability is monotone in the complement)**.** *For a unit with nonincreasing deterioration d and a repair function increasing in spend, if a ceiling profile z' is pointwise at least z then $y(k,z')\geq y(k,z)$ for every starting condition k, and the inequality is strict when the unit under z eventually becomes unlettable while under z' it does not.*

Proof. Any spending plan feasible under z is feasible under z'. Under the monotone dynamics a plan that spends at least as much at each date produces a condition path at least as high, hence at least as many lettable months. The strict case is the definition. □

The hypothesis carrying the conclusion is monotonicity of the dynamics in spend; the simplification is the deterministic ceiling; the assumption most likely to fail is that the ceiling is known in advance, because exposed complements are usually withdrawn without warning. The comparison with standard capital theory is as follows. Neoclassical theory measures the capital stock by cumulating investment less depreciation, so that identical assets are worth the same. This book adds the complement as a variable. It does not deny the value of ownership; it observes that two items with identical ownership, age, and condition differ by more than half in earnings over three years because of what each is paired with, which is a statement about their position in the viability analysis of Chapter 13: under the ceiling of 100 a unit at 60 is outside the kernel, and under the ceiling of 400 it is inside.

What has been established. Physically identical assets can yield very different usable capabilities, the difference is computable from the complement profile, and exposed dependence can reduce a capability sharply. The decision now enabled is to price the complement along with the asset. Assumptions still required are the ceiling profile and a

rule for what the owner can do when a complement is withdrawn. The question passed on to Chapter 8 is how to schedule maintenance when its price is paid now and its benefit later.

Exercise 7.3. Compute the let months and total surplus over 36 months for the unit at 60 if the full ceiling is available for twenty-four months and then withdrawn.

Exercise 7.4. Prove Proposition 7.2 formally for the housing transition, and give a non-monotone deterioration under which it fails.

Exercise 7.5. Name an exposed complement in a domain of your choice, such as a cloud service or a medical device, and write down what the dependent operator would have to hold to be indifferent to its withdrawal.

Chapter 8

Maintenance and the Reproduction of Capacity

The crew serving the building can spend at most 400 per month across all units. A manager is told to set the ceiling for the next year and wants to know how low it can safely go. The answer depends on how fast units deteriorate in their current condition and on what a given spend restores, and it is not smooth. If the ceiling is 100, only units above condition 80 can be held indefinitely. If it is 200, units above condition 50 can be held. Only at 400 can every lettable unit be held. The manager's difficulty is that the cut from 400 to 300 looks like a quarter reduction and removes the ability to hold every unit at condition 50 or below, which is exactly the units in greatest need.

8.1 Deterioration, thresholds, and the viability boundary

Take the model of Chapter 1: $k_{t+1} = \min\{k_{\max},\, k_t - d(k_t) + g(M_t)\}$ with $g(M) = M/100$ tenths, lettable if $k \geq 40$, and d equal to 1, 2, 4 on the three tiers. A budget ceiling $\bar{M}$ restricts $M_t \leq \bar{M}$. The *viability kernel* of the unit under ceiling $\bar{M}$ is the set of conditions from which some admissible policy keeps $k_t \geq 40$ for all t.

Theorem 8.1 (P12(b): threshold under a budget ceiling)**.** *Suppose d is nonincreasing in k, g is nondecreasing, and $d(k) > 0$ for all k in the lettable range. Then the kernel under ceiling $\bar{M}$ is $K(\bar{M}) = \{k \geq 40 : d(k) \leq g(\bar{M})\}$, so the viability threshold is $k^*(\bar{M}) = \min\{k \geq 40 : d(k) \leq g(\bar{M})\}$ (or no state is viable if the set is empty).*

Proof. If $d(k) \leq g(\bar{M})$ then spending $\bar{M}$ gives a net change $g(\bar{M}) - d(k) \geq 0$, and since d is nonincreasing in k the next condition $k' \geq k$ also satisfies $d(k') \leq d(k) \leq g(\bar{M})$, so the argument repeats and the unit never falls. If $d(k) > g(\bar{M})$ then under any policy with $M_t \leq \bar{M}$ the condition falls by at least $d(k) - g(\bar{M}) > 0$ in the first month, and in every later month by at least that amount because d can only rise as k falls, so k_t drops below 40 in finitely many steps. □

The numbers are as follows. For $\bar{M} < 100$ no unit is viable since $g(\bar{M}) < 1 \leq d(k)$. For $100 \leq \bar{M} < 200$ the threshold is 81. For $200 \leq \bar{M} < 400$ the threshold is 51. For $\bar{M} \geq 400$ the threshold is 40. The companion confirms these against exhaustive computation of the kernel (check BOOK, Chapter 8). Hypotheses doing work are monotonicity of d and the fact that the ceiling is binding in each month; a model in which d had a hump would have disconnected kernels. The threshold jumps here because d is a step function of k. If d were continuous and strictly decreasing, the threshold $k^*(\bar{M})$ would be continuous in $\bar{M}$, and it would have a kink wherever d does. Nonlinear deterioration alone, therefore, does not guarantee discontinuity; the discontinuity is created by jumps in the rate. The economic reading is that a modest reduction in the ceiling can remove the ability to hold an entire class of units, and the units removed are those in the worst condition.

8.2 Deferral and the loss of viability

Chapter 1 computed the financial cost of deferral. The viability question is different: how long can maintenance be deferred before the unit has left the set from which the ceiling can still hold it. Take a unit at 90 and defer all maintenance, $M = 0$. Its condition follows $90, 89, \ldots, 80$ at month ten, then $78, 76, \ldots, 50$ at month twenty-five, then $46, 42, 38$. Under a ceiling of 100 the kernel is the conditions from 81 upward, so the unit is inside the kernel through month nine, where $k = 81$, and outside it from month ten, where $k = 80$. Under the ceilings 200 and 300 the kernel starts at 51, and the unit is inside through month twenty-four, where $k = 52$, and outside from month twenty-five. Under the ceiling of 400 the kernel is every lettable condition, and the unit is inside through month twenty-seven, where $k = 42$. The unit is therefore deferrable for at most nine, twenty-four, twenty-four, and twenty-seven months under the four ceilings. The numbers have a plain reading: with the ceiling at 100 an owner who has deferred for ten months has permanently lost the ability to hold the unit and now needs a higher ceiling to prevent decline. The amount of deferral that can be undone by the available crew depends on the crew, and a manager who watches only condition, which stays acceptable until month twenty-seven in all four cases, sees nothing. These values are reproduced in the companion (check BOOK, Chapter 8).

8.3 Preventive maintenance, restoration, and backlog

Preventive maintenance holds k where it is at a cost of $d(k) \times 100$ per month. Restoration after failure raises k at 100 per tenth. The cost of preventing a fall of n tenths is $100n$ and so is the cost of repairing it, at the tier where each month's decay is one tenth. In lower tiers decay is 2 or 4 per month and the cost of holding is a multiple of the cost at the top. The sum of the costs of holding a unit at $k = 50$ forever is 400 per month, which is 4800 per year, against a restoration cost from 38 to 90 of 5200 once. The two are the same order of magnitude, and the owner who faces a unit at 50 is comparing a perpetual flow to a lump sum and an uncertain period of lettability in between.

Maintenance schedules also create backlogs. Let $\rho_M(t)$ be the rate at which maintenance need accrues and $\rho_R(t)$ the rate at which crews retire it, and let $D(t) = \int_0^t (\rho_M - \rho_R)\,du$ be the cumulative excess. One might hope that a nonpositive long-run average net accumulation rate guarantees a bounded backlog, but it does not. Take $\rho_M - \rho_R = 1/(2\sqrt{1+t})$, so that $D(t) = \sqrt{1+t} - 1$. The average rate $D(t)/t$ tends to zero, being 0.0909 at $t = 99$ and about 0.0099 at $t = 9999$, while $D(99) = 9$ and $D(9999) = 99$ are unbounded in t. A bounded-backlog result therefore needs the stronger condition $\sup_{t \geq 0} \int_0^t (\rho_M - \rho_R)\,du < \infty$, or a reflected formulation in which backlog cannot be negative, $B(t) = \sup_{0 \leq s \leq t} \int_s^t (\rho_M - \rho_R)\,du$, bounded under the corresponding condition on all intervals. The correction was carried over from earlier work and is stated here as a hypothesis of the backlog results of Part IV.

What has been established. Under monotone deterioration the viable set under a spending ceiling is an upper set of conditions with a threshold that jumps where the deterioration rate does, and a zero average accumulation rate does not bound a backlog. The decision now enabled is to set the ceiling with reference to the conditions of the worst units. Assumptions still required are monotone deterioration and a spending-restoration function, which are empirical. The question passed on to Chapter 9 is when to commit resources to an irreversible repair and when to wait for information.

Exercise 8.2. For the three tiers, find the annual cost of holding a unit at conditions 90, 70, and 50 and the ceiling needed in each case.

Exercise 8.3. Construct a deterioration function that is not monotone and give a ceiling for which the kernel is not an upper set.

Exercise 8.4. Using an official maintenance-backlog figure for a transit system or school district, say whether the reported figure supports a boundedness claim and what additional series would be needed.

Chapter 9

Irreversible Decisions and the Value of Waiting

The owner has learned that the roof of one building may need replacing. Replacing it costs $I = 3500$. If the roof is replaced, the payoff is an increase in the building's long-run value that depends on a revealed fact, namely whether a planned neighbourhood development proceeds. If it proceeds the payoff is $X = 6000$ and if not $X = 2000$, each equally likely, so $\mu = 4000$. The fact is revealed in one month. Acting now also earns a first-month flow return of $r = 100$ that waiting forgoes. The roof belongs to a unit now at condition 55, so that the investment of $I = 3500$ is exactly the restoration to 90 at 100 per tenth. Waiting one month is not free of physical cost: the unit deteriorates by $d(55) = 2$ tenths to 53, and the same restoration costs 3700 after the wait. The owner discounts the next month at $\beta = 0.9$ and may decline to invest after the revelation. The opportunity might not survive: with probability 0.8, independent of the payoff, the contractor's offer is still available after a month. The assumptions are therefore stated in full: the development decision is revealed perfectly at the end of month one; the owner may decline after the revelation; the contractor's offer expires at the end of the month with survival probability q, taken independent of the revelation in the first calculation and dependent in the second; the discount factor for the month is $\beta = 0.9$; and the cost of the investment after waiting is $I' = 3700$ because of the deterioration. Waiting is also irreversible in a sense, because the month cannot be recovered, and the problem is to compare the two.

9.1 The two-period comparison

Acting now yields $V_{\text{act}} = (\mu - I) + r = (4000 - 3500) + 100 = 600$. Waiting yields $V_{\text{wait}} = \beta\,\mathbb{E}[\mathbf{1}_{\text{surv}}(X - I)^{+}]$. After waiting the cost is $I' = 3700$. In the low state $X - I' = -1700 < 0$ and the owner declines; in the high state $X - I' = 2300$. So $\mathbb{E}[(X - I')^{+}] = \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 2300 = 1150$. With independent survival of probability $q = 0.8$, $V_{\text{wait}} = 0.9 \times 0.8 \times 1150 = 828 > 600$ and the owner should wait. The comparison reverses when survival is not independent

of the payoff. Suppose the contractor's offer survives with probability 0.4 in the high state (because the development also raises demand for contractors) and with probability 1 in the low state. Then $\mathbb{E}[\mathbf{1}_{\text{surv}}(X - I')^+] = \frac{1}{2} \cdot 0.4 \cdot 2300 = 460$ and $V_{\text{wait}} = 0.9 \times 460 = 414 < 600$, so acting now is better. The calculation is in the companion (check BOOK, Chapter 9).

Proposition 9.1 (P26: two-period value of waiting)**.** *Let the investment cost be I now and I' after one period, and the payoff X with mean μ, revealed after one period, and let $\mathbf{1}_{\text{surv}}$ indicate that the opportunity survives. Then $V_{\text{act}} = (\mu - I) + r$ and $V_{\text{wait}} = \beta\,\mathbb{E}[\mathbf{1}_{\text{surv}}(X - I')^+]$. If survival is independent of X with probability q then $V_{\text{wait}} = \beta q\,\mathbb{E}[(X - I')^+]$, and waiting is better exactly when $q > q^* = V_{\text{act}}/(\beta\,\mathbb{E}[(X - I')^+])$. In every case $\mathbb{E}[(X - I')^+] \geq (\mu - I')^+$.*

Proof. The expressions are the definitions. Independence factorizes the expectation. The inequality is Jensen's for the convex function $z \mapsto z^+$. □

In the example $q^* = 600/(0.9 \times 1150) = 40/69 \approx 0.580$, so waiting is better for any survival probability above that value. The hypotheses doing work are independence of survival and the option to decline: if the owner were obliged to invest after the revelation, the value of waiting would be $\beta q(\mu - I') = 0.9 \times 0.8 \times 300 = 216$, which is far less than acting now. The deterioration of 200 in the cost is also material: with no deterioration the option value would be 1250 and the value of waiting 900. The assumption that the revelation is costless also matters; if learning the fact required a payment, the comparison would change by that amount. The counterexample with dependent survival shows that independence is not innocuous.

9.2 Symmetry of action and inaction

The classical treatment of irreversible investment, from Arrow and Fisher [3] to the real-options literature of Dixit and Pindyck [8], emphasizes that committing resources forecloses the option to wait. The point made here is that waiting is also a choice that can foreclose things: the contractor's offer may lapse, the roof may deteriorate, and a resource that deteriorates while unused costs the owner something. Both action and inaction are potentially irreversible, and a complete treatment compares them symmetrically by listing, for each, which futures it closes. The two-period form above handles this by the survival indicator and the forgone return r. The deterioration of the roof is included above through I'. A richer form would make the cost $I(t)$ a function of the delay, and it would model the owner's remaining budget as part of the state; Part IX states the general version as a specification.

What has been established. Waiting is better than acting in the example exactly when the survival probability exceeds 40/69 if survival is independent of the payoff, and acting is better when survival is negatively related to the payoff. The decision now enabled

is a transparent comparison that lists the survival assumption. Assumptions still required are the probability model for survival and the discount factor. The question passed on to Chapter 10 is what remains after an activity ends, which is one of the things that make an action irreversible.

Exercise 9.2. Recompute q^* for a first-period return of $r = 300$, and for the variant in which waiting does not raise the cost, $I' = I = 3500$.

Exercise 9.3. Prove that the option value $\mathbb{E}[(X - I)^+] - (\mu - I)^+$ is nonnegative, and find a distribution of X with zero option value.

Exercise 9.4. Find a public decision, such as whether to repair a bridge now or study it further, and list the ways that waiting forecloses options.

Chapter 10

Residuals, Waste, and the Afterlives of Production

The unit that Chapter 1 left to deteriorate for twenty-eight months is at condition 38 and earns nothing. Restoring it to 90 costs 5200. The owner holds the unit in a company whose only asset it is, valued at 500 in a forced sale. The owner can restore the unit, replace its boiler and continue, or abandon the unit. Abandonment leaves a vacant shell with a hazardous old boiler, and someone will have to pay to secure and remove it. The decision is made by the owner, the cost of abandoning falls elsewhere, and the aim of this chapter is to account for that cost before the decision is made.

10.1 What persists

Production leaves things behind. A demolished building leaves rubble and a site. A factory leaves contaminated soil. A retired software system leaves data in formats no one reads and interfaces others still depend on. An agency leaves forms and rules that continue to bind after their purpose ends. Each is a residual in the sense of Chapter 6, and each has a disposition duty that is assigned to someone or unassigned. In the housing case there are two events that generate residuals, and the accounts of each can be set out in the same units as before.

Replacement. If the owner replaces the unit's old boiler, the old boiler contains insulation that is hazardous when disturbed. Handling costs 400 and disposal 150, a residual obligation of 550 that arises at the moment of replacement and is discharged only when the work is done. An owner who leaves the old boiler in the basement has not removed the obligation; the obligation has been transferred silently to whoever holds the building next.

Abandonment. If the owner abandons the unit at condition 38, the shell remains. The residual obligations are the demolition, which costs 2000, the removal of the hazardous boiler, which costs 550, and the securing of the vacant shell at 50 per month for the six

months that the municipal process takes before demolition, which costs 300. The total disposition cost is $L = 2000 + 550 + 300 = 2850$. The accounts after abandonment show a balance sheet with the unit written off and, if the obligation is recorded, a liability of 2850. If it is not recorded, the balance sheet shows nothing and the obligation falls on the municipality or the neighbours.

10.2 Who bears the residual

Compare the options for the owner whose company holds assets $a = 500$ and has limited liability. Restoring costs 5200 and returns a unit that earns 700 per month while it lasts, so on any horizon beyond eight months the restoration is financially worth doing, but it needs 5200 the owner may not have. Abandonment costs the owner $\min\{L, a\} = \min\{2850, 500\} = 500$, the whole of the company's assets, and costs the community the remaining $2850 - 500 = 2350$. If the owner were liable for the full L the cost of abandoning would be 2850 and would still compare favourably with restoration at 5200 for an owner who expects to hold the unit only briefly, but it would not be a free exit. Limited liability makes the private comparison 500 against 5200 and the social comparison 2850 against 5200, and the two rank the options differently when the unit's continued earnings do not repay the restoration.

Proposition 10.1 (Residual liability and internalization)**.** *Let a decision yield private benefit π and generate a residual with disposition cost L, with liability limited to assets a. The decision-maker takes it if and only if $\pi \geq \min\{L, a\}$, while the socially efficient rule takes it if and only if $\pi \geq L$. The two differ exactly when $\min\{L, a\} \leq \pi < L$, and in that case a bond of at least $L - a$ posted at the start restores the efficient rule.*

Proof. Private payoff is $\pi - \min\{L, a\}$ and social payoff is $\pi - L$. They differ in sign exactly when $\min\{L, a\} \leq \pi < L$, which requires $a < L$. A bond b raises assets to $a + b$, so with $b \geq L - a$ the liability is L. □

In the abandonment example, the benefit of abandoning is the restoration cost avoided, $\pi = 5200$, so the comparison is between $\pi - \min\{L, a\} = 4700$ and $\pi - L = 2350$. Both are positive, so abandonment is efficient here too, once the owner reports the liability of 2850 in full. The proposition is strictly about the sign, and in the illustration with a smaller benefit, for example a project with private profit 1000 and residual cost 3000 for a firm with assets 500, the private payoff is $500 > 0$ and the social payoff is $-2000 < 0$, and a bond of 2500 is needed to align them. This is the judgment-proof problem of Shavell [20] in the notation of this book. The new element is the state: the residual is a record of an obligation that outlives the transaction, and the choice of whom it is assigned to determines viability for those who inherit it, which in the abandonment case is the municipality's budget. The numbers are reproduced in the companion (check BOOK, Chapter 10).

The hypothesis carrying the result is limited liability, and the simplification is a known and certain L. The assumption most likely to fail is that the owner and the authority agree on L. The cost of a residual is typically uncertain and contested, and a bond sized to a low estimate leaves a gap. The interpretation is that making disposition duties explicit, attached to the producer and secured at the outset, is a different kind of regulation from one that inspects output, since it takes the continuation of the site as the thing to be secured.

What has been established. Replacement and abandonment each generate residual obligations that can be accounted in the same units as the rest of the case, a total of 550 for replacement and 2850 for abandonment, and with limited liability the private and social costs of abandoning differ by 2350. The decision now enabled is to attach a bond to each project in proportion to its residual cost. Assumptions still required are a measure of L and a mechanism of enforcement. The question passed on to Part III is how to represent such duties, and obligations of every kind, through time.

Exercise 10.2. For $\pi = 1000$, $L = 3000$, and $a = 500$, compute the smallest bond that makes the firm decline the project, and the range of π for which the firm undertakes it with the bond in place.

Exercise 10.3. Show that if liability is unlimited the private and social rules coincide, and discuss what practical features make liability effectively limited.

Exercise 10.4. Select a residual from the chapter's list and describe who bears it today, who would bear it under an assigned-bearer rule, and what record would be needed to enforce the assignment.

Part III

Obligations, Degradation, and Recovery

Part III represents obligations through time and introduces viability kernels, safe degradation, compensating continuation, and corrective capacity.

Housing thread. Tenancy and repair duties become obligations, an unheated unit may not be let, and viability is studied under a maintenance budget.

In Part III the crew's ceiling of 400 per month and the lettable threshold of 40 become constraints on promises. The owner owes the tenants a heated, habitable unit, and the question is which states of the building and of the owner's accounts are ones from which that promise can still be kept.

Chapter 11

Obligations Through Time

The owner of the building holds 2500 in cash and expects 1000 of rent in thirty days. Three repair duties are outstanding. A heating failure must be fixed within 3 days at a cost of 600; a leak within 14 days at a cost of 900; and a roof section within 30 days at a cost of 2000. A second owner holds the same 2500 and has no outstanding duties. A lender looking at the two balance sheets sees the same cash, the same building, and the same rent roll. The decision for the first owner is whether the three promises can all be kept, and for the lender whether to extend credit. The difficulty is that the information needed to answer is held in the list of duties and not in the accounts.

11.1 Obligations as objects with structure

An obligation is represented by a *bearer* who owes, a *beneficiary* to whom it is owed, a *trigger* that brings it into force, a *deadline* by which it must be discharged, and a *discharge condition* that says what counts as discharge. A lease repair duty has the landlord as bearer, the tenant as beneficiary, a report of failure as trigger, a stated number of days as deadline, and restoration of function as the discharge condition. Contracts, promises, legal duties, maintenance commitments, and informal responsibilities all fit the same pattern, with the difference that informal ones often lack a stated deadline and a discharge condition, and a good deal of the difficulty in them lies there. Because an obligation constrains future action, it belongs in the state description, and Chapter 2 showed that two situations with the same holdings and different obligations are different states.

Definition 11.1 (Obligation)**.** An *obligation* is a tuple $(b, \beta, \tau, T, \Delta)$ of bearer b, beneficiary β, trigger τ, deadline T, and discharge condition Δ. It is *outstanding* at a history after its trigger has occurred and before Δ holds, and *breached* if T passes while it is outstanding.

The three repairs are obligations of the same bearer to the same beneficiary class, with deadlines 3, 14, and 30 and costs 600, 900, and 2000. A boundary case is a hope, such

as an owner's intention to repaint the hallway: it has no beneficiary able to enforce it, no deadline, and no discharge condition, and it does not constrain the state.

Proposition 11.2 (Joint feasibility of dated obligations)**.** *Suppose obligations $j = 1, \dots, n$ with deadlines $T_1 \le \cdots \le T_n$ and costs c_j can be paid at any date up to their deadlines from a single fungible resource, and resources arrive in amounts a_s at dates s with a_0 the initial cash. The obligations are jointly feasible without breach if and only if, for each j, $\sum_{i \le j} c_i \le \sum_{s \le T_j} a_s$.*

Proof. Necessity: everything due by T_j must be paid by T_j from what has arrived by T_j. Sufficiency: pay each obligation at its deadline in order. At T_j the cash on hand is $\sum_{s \le T_j} a_s - \sum_{i<j} c_i \ge c_j$ by hypothesis. □

In the example the cumulative amounts due are 600 by day 3, 1500 by day 14, and 3500 by day 30. The cumulative supply is 2500 by day 3, 2500 by day 14, and 3500 by day 30, since the 1000 of rent arrives on day 30. All three conditions hold, with equality at day 30. An owner with 2400 in cash fails the last condition by 100, and the failure is caused by the dated roof duty, which does not appear in the balance sheet. The second owner, with no duties, is not even at risk. The calculation is reproduced in the companion (check BOOK, Chapter 11).

The hypotheses doing work are fungibility of the resource and the ability to pay at any time before the deadline. If the contractor must be paid in advance or the roof section requires materials ordered weeks earlier, the effective deadline is earlier than the stated one, and the proposition applies to the effective deadlines. If rent may fail to arrive, the proposition holds with a_s replaced by guaranteed amounts, and Part III develops the robust version. Economically, the result says that a balance sheet is not a sufficient statistic for the ability to meet obligations, in the precise sense of P11: the accounts do not determine feasibility, because the same accounts are compatible with feasible and infeasible duty lists.

What has been established. Obligations have a structure that a balance sheet omits, and joint feasibility of dated obligations has a cumulative characterization. The decision now enabled is to test a set of commitments against dated supply before taking on another. Assumptions still required are fungibility of the resource and the date at which each cost is actually incurred. The question passed on to Chapter 12 is which obligations are really obligations, as against capabilities that may be allowed to lapse.

Exercise 11.3. For duties of 600 by day 3, 900 by day 14, and 2000 by day 30 and cash of 2500, find the latest day on which a new 400 duty can be added with deadline 20 while preserving feasibility, if rent of 1000 arrives on day 30.

Exercise 11.4. Show by an example that Proposition 11.2 fails if payment must be made a week before the deadline for some duties, and state the correct condition.

Exercise 11.5. Find an informal responsibility in a household or a workplace and write it as an obligation tuple, saying which components are left unspecified in practice and what disputes result.

Chapter 12

Required Functions and Optional Capabilities

The owner faces a heating failure in one unit in winter. Heating is a function the lease requires. The elevator in the same building is a convenience for tenants on the upper floors and is out of order for the same month. A cost-conscious owner would postpone both, and a city inspector will accept only one of these decisions. The question is what makes a function required and what makes a capability optional, and the answer is that the classification is a decision made by someone, with consequences, and not a fact visible in the object.

12.1 Classification as an institutional decision

A function is *required* if the regime in force demands that it continue, so that its failure is a breach and the state in which it fails is outside the acceptable set G. A capability is *optional* if it may be reduced or suspended without breach. The classification is established by statute, contract, regulation, or practice, and it can be contested and revised. Heating in cold months is typically required by housing codes. An elevator may be required for buildings above a given height or for tenants with certain disabilities, and optional otherwise. A public library's branch hours are optional under most budgets, and the same hours become required once a court rules that access to legal materials is a right. Calling a function optional is therefore a consequential act: it moves its failure from the set of breaches to the set of acceptable degradations, and what that costs depends on who relies on it.

Once a function is required the question is how it may be degraded safely. Let K be the viability kernel of the acceptable regime. The set of *safe degradation actions* at $s \in K$ is $A_K(s) = \{a \in A(s) : T(s, a, w) \in K \text{ for every allowed } w\}$.

Theorem 12.1 (P06: safe degradation requires successor viability)**.** *For $s \in K$, an action preserves the possibility of satisfying the ongoing regime exactly when it lies in $A_K(s)$, and*

$K = \{s \in G : A_K(s) \neq \varnothing\}$. If degradation creates a restoration duty, the duty must be included in the successor state before A_K is evaluated.

Proof. By P03, $K = \Phi(K)$, and $\Phi(K)$ is the set of $s \in G$ with an action all of whose successors lie in K, which is the second statement. An action not in $A_K(s)$ has some allowed successor outside K, from which by definition the regime cannot be satisfied forever against all disturbances. If a restoration duty is created, the extended state records it and the extended kernel K^{ext} replaces K in the definition of A_K. □

The result is elementary and the housing numbers show why it matters. A unit has condition 52, the crew ceiling is 400, and the owner is considering skipping a month of maintenance on it. Without a restoration duty the plain kernel under the ceiling is every condition from 40 upward, and a skipped month lands the unit at 50, which is in the kernel: the unit can be held at 50 at the full ceiling forever. Suppose, however, that the lease imposes a restoration duty after any degradation, that the unit must return to condition 60. At the ceiling of 400 the net change per month is $4 - d(k)$, which is $+2$ in the middle tier and 0 at conditions 50 and below. A unit at 51 rises to $53, 55, 57, 59, 61$ in five months, and a unit at 50 never rises. The extended kernel is therefore conditions from 51 upward, and the skipped month, which lands at 50, is not a safe degradation. The safe actions at 52 are the spending levels of 100 or more, which land at 51 or higher. The calculation is reproduced in the companion (check BOOK, Chapter 12).

The economic point is the one in the theorem's last sentence. A temporary suspension cannot be judged apart from the obligation it leaves behind, and an owner who evaluates the skipped month by its immediate effect on a plain viability test will make a degradation that the extended regime does not allow. The typed treatment of the classification of required functions, and of what an execution may and may not do, is taken up in Part XI under P20. A boundary case in the opposite direction is a function with no restoration cost, such as a service whose suspension leaves no trace once resumed; for it the extended and plain kernels coincide and safe degradation is just viable degradation.

12.2 Declaring a function optional

A public-service example shows the institutional side. A transit agency under budget pressure may classify night service as optional. If the agency's regime treats night service as optional, suspending it breaches nothing and is accounted as a saving. The workers whose shifts end after midnight face an expensive commute or the loss of their jobs, and their continuation set shrinks. The classification therefore allocates burdens, and a good deal of the politics of budgeting is a fight about the list of required functions. The book does not claim to know the right classification in any particular case. It claims that the classification should be explicit, contestable, and accompanied by a record of who relies on each function, so that reclassifying a function is a visible act with identified consequences.

What has been established. A degradation is safe only if its successor is in the kernel of the regime including any restoration duty it creates, and the classification of functions as required or optional is itself a consequential institutional decision. The decision now enabled is to evaluate suspension against the extended state. Assumptions still required are an explicit list of required functions and a model of the restoration duty. The question passed on to Chapter 13 is how to compute the kernel when resources are scarce and disturbances are possible.

Exercise 12.2. For the restoration duty to return to 70, compute the extended kernel under the ceiling of 400 and the safe actions at condition 60.

Exercise 12.3. Prove that if a degradation creates no restoration duty then $A_K(s)$ computed with the extended state equals the plain $A_K(s)$.

Exercise 12.4. Select a public function that is currently optional. Identify who relies on it, and write the sentence that a reclassification would have to include to be explicit about consequences.

Chapter 13

Viability Under Scarcity

The owner can now ask the question the earlier chapters were built to answer: from which conditions of the building can the owner keep it lettable, given that the crew can spend at most a stated amount each month. A target such as "every unit at condition 90" is desirable, but it is not the same as the set of conditions from which lettability can be preserved. The first is a goal. The second is a fact about what is reachable under constraint, and the difference decides whether a manager who is behind on a unit should push for the goal or accept a smaller aim.

13.1 The viability kernel

Take a state space S, an acceptable set $G \subseteq S$, actions $A(s)$, allowed disturbances $W(s,a)$, and transition T. The *viability kernel* is the set of states from which some policy keeps the system in G forever against every allowed disturbance,

$$K = \{x : \exists \pi \text{ such that all trajectories permitted by the model under } \pi \text{ satisfy the continuation constraints}\}.$$

With $\Phi(X) = \{s \in G : \exists a \ \forall w,\ T(s,a,w) \in X\}$, P03 identifies K as the greatest fixed point of Φ, which exists by monotonicity and the theorem of Knaster and Tarski [21]. The iteration $K_0 = G$, $K_{n+1} = \Phi(K_n)$ decreases and stabilizes in at most $|S|$ steps for a finite system, but the companion exhibits an infinite chain in which each K_n is nonempty and $K = \varnothing$, so termination claims need finiteness. For the continuous-time analogue see Aubin [4].

In the housing unit under a ceiling $\bar{M}$ the kernel was computed in Chapter 8: it is every condition from 81 upward at $\bar{M} \in [100, 200)$, from 51 at $[200, 400)$, and from 40 at $\bar{M} \geq 400$.

13.2 Comparative statics

Theorem 13.1 (P23 and P04: comparative statics of viability)**.** *On a common state space, K is monotone increasing in the acceptable set G, increasing in each action set*

$A(s)$, and decreasing in each disturbance set $W(s,a)$. In particular, if $G_2 \subseteq G_1$ with identical dynamics, actions, and disturbances then $K(G_2) \subseteq K(G_1)$.

Proof. Φ is monotone in X and, for fixed X, increasing in G and in each $A(s)$ and decreasing in each $W(s,a)$, since a larger G or A makes the defining condition easier to satisfy and a larger W makes it harder. The greatest fixed point inherits the same monotonicity, because $K = \bigcup\{X : X \subseteq \Phi(X)\}$ and enlarging Φ pointwise enlarges the union. □

The corollary on added duties needs care. Tightening the acceptable set cannot enlarge the kernel, but it need not shrink it. With a ceiling of 100 the kernel under the lettable threshold 40 is the conditions from 81 up. Raising the threshold to 60 leaves the kernel unchanged at 81 and above, since no state below 81 was viable and the new duty removes only states that were already outside. Raising it to 85 shrinks the kernel to conditions from 85 up. An added duty therefore requires extra resources or a changed technology only when it removes previously viable states, that is, when $K(G_2) \subsetneq K(G_1)$, and the owner who is told that a higher standard is coming can check by computing the kernel. The three ceilings give nested kernels $\{81, \dots\} \subset \{51, \dots\} \subset \{40, \dots\}$ in agreement with the action-set statement of the theorem. If an adversary may add one tenth of extra deterioration each month, a disturbance, the ceiling of 400 cannot hold units at or below 50, because the deterioration there is $4 + 1 = 5$ against a repair of 4, and the robust kernel is the conditions from 51 upward. The results are in the companion (check BOOK, Chapter 13).

13.3 Local acceptability and almost-sure viability

Proposition 13.2 (P05: local acceptability does not imply viable continuation)**.** *There are systems with $s_0 \in K$, $s_1 \in G$, $s_1 \notin K$, and an action from s_0 to s_1.*

Proof. Three states suffice. From s_0 one action stays at s_0 and another goes to s_1. From s_1 every action goes to a state $b \notin G$. Then $s_0 \in K$ by always staying, and $s_1 \in G \setminus K$. □

In the housing numbers, take the ceiling 100 and the unit at 90. The condition 70 is acceptable, since it is above 40, and it is reachable from 90 by declining to maintain for twenty months: the unit passes through 78 at month 11 and 60 at month 20. It is not viable, because it is below 81. A test at each date that checks only acceptability passes throughout, and the unit is nonetheless doomed once it has crossed 81. The failure of the test is not visible until the unit actually falls below 40 some years later.

For stochastic disturbances, P25 separates three formulations. Robust viability requires every allowed disturbance to be survived. In a finite or countable model, almost-sure viability requires only that every positive-probability successor remain in the set, and its predecessor operator is $\Phi_{\mathrm{as}}(X) = G \cap \{s : \exists a,\ \mathbb{P}(X \mid s,a) = 1\}$. Suppose the one-tenth shock occurs each month with probability 0.1. A unit at 50 under a ceiling of 400 falls by one tenth in each shock month and does not recover, so it needs eleven shocks in a year to

fall below 40 within the year, and the probability of that is $109/10^{12}$, about 1.1×10^{-10}. Yet over an unbounded horizon eleven shocks occur with probability one, so the unit at 50 is not almost-surely viable, and the almost-sure kernel agrees with the robust kernel, the conditions from 51. A risk-constrained notion, with a stated horizon and a tolerated failure probability, would admit the unit for one year and nothing more. For general state spaces the statement of the greatest fixed point of Φ_{as} as an implementable policy requires fixing a measurable domain and a selection theorem [5]; this book treats that case as a research specification.

What has been established. The viability kernel is the set of conditions from which obligations can still be met, it responds monotonically to duties, capacities, and exposures, and acceptable-now does not imply viable. The decision now enabled is to compute the kernel before adopting a standard. Assumptions still required are finiteness for termination of the iteration and a specified disturbance set. The question passed on to Chapter 14 is how a suspended obligation can be credibly restored.

Exercise 13.3. For a ceiling of 200 compute the kernel with adversarial shocks of one and two tenths per month.

Exercise 13.4. Prove that if G is forward invariant under some policy for all of S then $K = G$, and show that the converse holds in a finite system.

Exercise 13.5. Using a municipal reserve fund or a hospital bed capacity, describe the acceptable set, the actions, and the disturbances, and say whether tightening a stated standard would change the kernel.

Chapter 14

Compensating Continuation

A unit in the building has been run down to condition 51 because an earlier owner deferred work, and the current owner has promised the tenants that it will be restored to 60. The promise is a restoration obligation. The question is whether it is credible, and the answer depends on whether the resources needed are reserved, whether they are in fact available when needed, and whether other duties draw on the same resources. A promise that rests on a crew already committed elsewhere is a hope.

14.1 A rank certificate

Let $R \subseteq G$ be the set of states in which the restoration is discharged, let $D \subseteq G$ carry a rank $r : D \to \mathbb{N}$ with $r(s) = 0$ exactly on R, and let a policy be such that every allowed successor remains in D and strictly lowers the rank whenever $r > 0$.

Proposition 14.1 (P07: constructive certificate)**.** *From $s_0 \in D$ the restoration occurs within $r(s_0)$ steps while remaining acceptable.*

Proof. By induction on $r(s_0)$: if $r = 0$ the state is in R. Otherwise the policy leads to a successor in D of rank at most $r(s_0) - 1$, which by hypothesis reaches R within that many further steps, and all intermediate states are in $D \subseteq G$. □

For the unit, with the ceiling at 400 the net rise per month in the middle tier is $4-2 = 2$ tenths, so the rank to reach 60 is 5 months from 51, 4 from 52, 3 from 55, 1 from 58, and the state 50 has no finite rank because it cannot rise. The certificate is the table of ranks together with the spend that realizes it: full ceiling every month for five months, a reserve of $5 \times 400 = 2000$ in crew time. This number is what makes the promise concrete, and a promise without it is a statement of intention. The certificate does not prove the unit will be restored; it proves restoration follows if the reserved resources and permissions are actually supplied and the disturbances do not exceed those modelled.

14.2 Credible and non-credible compensation

Now give the owner two units, both at condition 51, and the same single crew with a ceiling of 400 for both. Each unit's restoration needs the whole ceiling, and holding the other unit at 51 needs 200. A joint schedule that restores both to 60 does not exist: the exhaustive search of the companion finds no sequence of allocations under which both units are at or above 60 while remaining lettable (check BOOK, Chapter 14). The two promises, made separately, each look credible and jointly are not. This is the shape of P09 applied to restoration, and it matches the case in which a promise of compensation depends on a capability that other promises have already consumed. Three worked cases show the same structure. Deferred building repairs are credible only with a reserve of crew time and funds that is not also pledged to another project. Reduced hospital services are credible only if the staff who would restore them are retained and not redeployed. Temporary infrastructure closures are credible only with a funded schedule and a contractor under contract.

What has been established. A restoration obligation is credible when a rank function and reserved resources support a policy whose progress survives the modelled disturbances, and two individually credible promises can be jointly incredible under shared scarcity. The decision now enabled is to require a reserve and a rank for every restoration promise. Assumptions still required are the disturbance set and an accurate record of other claims on the same resources. The question passed on to Chapter 15 is how an operation that appears to function can lose the capacity to recover.

Exercise 14.2. For the ceiling of 400 compute the rank to reach 70 from 60 and the crew reserve it requires.

Exercise 14.3. Show that for the two-unit case at $(51, 51)$ no policy reaches $(60, 60)$, and find the best achievable (a, b) after ten months.

Exercise 14.4. Take a public promise to restore a service after a closure. Identify the rank function, the reserve, and any competing claim on the reserve.

Chapter 15

Corrective Capacity and Recoverability

A manager reviews the building every year. At each review the manager checks that every unit is lettable and that the next year's budget covers what is needed to keep it so. In every review for ten years the building passes. In the eleventh year a tenant reports that a unit has failed, and the manager discovers that it has been past the point of no return for several years. The review was a good test of whether the building was acceptable and a poor test of whether it could still be saved. Corrective capacity is the resources, permissions, knowledge, and time needed to correct a course of action, and a system can remain operational while it loses the ability to recover from foreseeable disturbances.

15.1 Lookahead depth and the point of no return

Define the survival chain for a finite system by $K_0 = G$ and $K_{n+1} = \Phi(K_n)$, so that K_n is the set of states from which the acceptable set can be preserved for n further steps. A review that checks viability over a horizon of n steps certifies exactly K_n.

Theorem 15.1 (Depth of certification)**.** *Let the iteration stabilize at $K = K_m$ with $K_{m-1} \neq K_m$. Then any review of depth $n < m$ certifies a set $K_n \supsetneq K$, and a state in $K_n \setminus K$ passes the review while being unrecoverable: it will leave G in finite time under every policy.*

Proof. The chain is decreasing, and it is strictly decreasing up to its stabilization, since if $K_{j+1} = K_j$ for some j then every later set equals K_j and $m \leq j$. So $K_n \supsetneq K_m = K$ for $n < m$. A state outside K has no policy that keeps the system in G forever, and in a finite system any state outside K_N fails within N steps for $N = |S|$. □

The hypotheses doing work are finiteness, which gives stabilization, and the fact that the review certifies only a finite depth. In the housing unit under the ceiling 100 the chain stabilizes at the thirty-fourth step, so $m = 34$ months. A one-step review certifies

conditions from 43 upward. A twelve-month review certifies conditions from 59 upward. The kernel is the conditions from 81 upward. A unit at 60 passes every review of depth up to 13 months and will fail eventually, and the manager who reviews annually on a twelve-month depth sees only the first signs of trouble in the year the unit actually fails. The calculation is in the companion (check BOOK, Chapter 15).

The lesson is that corrective capacity has a time scale. A review is useful if its depth is at least the lead time needed to correct, and it is useless if the correction takes longer than the depth it can see. In the example, a unit below 81 at the ceiling of 100 needs the ceiling raised before it can be saved, and raising the ceiling takes the time of procurement and hiring, so a review with depth of twelve months, applied to a process whose correction takes eighteen months of lead time, certifies states that are already unrecoverable. This is the general statement of the opening of the chapter.

15.2 The option value of corrective capacity

Proposition 15.2 (P14: weak option value)**.** *If* $\Pi_0 \subseteq \Pi_1$ *are policy sets for identical environments and criteria then* $\sup_{\Pi_1} J \geq \sup_{\Pi_0} J$*. Strict improvement requires a circumstance in which the added action helps. For viability the same statement is the action-set case of P23.*

The proof is that a supremum over a larger set is at least as large. Raising the ceiling from 300 to 400 enlarges the action set, and it changes the kernel from conditions 51 and above to 40 and above, which is strict because units at 50 and below can then be held. The statement says only that corrective capacity has potential value. Whether it has net value depends on the cost of acquiring it: the standing cost of the extra 100 per month, the training of the crew, the monitoring that tells the manager when to use it, and the reserves it requires. Part V treats these as substantive investment problems with their own models.

What has been established. A review that certifies acceptability over a finite depth passes states that cannot be saved whenever the stabilization step exceeds the depth, and corrective capacity has weak option value whose net value depends on acquisition cost. The decision now enabled is to set the review depth to at least the lead time needed to correct. Assumptions still required are finiteness and a model of lead time. The question passed on to Part IV is how to compare the profiles that arise, as the many obligations and capacities that the building carries cannot be put on a single scale.

Exercise 15.3. For the ceiling of 200 compute the survival chain length m, and the smallest condition certified by a twelve-month review.

Exercise 15.4. Prove that in a finite system the chain stabilizes in at most $|S|$ steps, and give a system for which $m = |S| - 1$.

Exercise 15.5. For a financial-regulatory stress test of horizon one year, say what the horizon certifies, what lead time the correction of a failing institution requires, and whether the horizon suffices.

Part IV

Value, Prices, Accounting, and Debt

Part IV compares valuations, prices, profit, debt, and accounts as representations, and asks what each preserves of the underlying continuation structure.

HOUSING THREAD. Rent, reserves, and the landlord's reported profit are set beside the maintenance backlog, which no conventional account records as a liability.

From Part IV the building has four units at conditions $(90, 70, 55, 50)$, the starting state recorded in the shared specification. At the full ceiling the holding cost of the units is 100, 200, 200, and 400 per month, a total of 900 against a ceiling of 400, and restoring all four to 90 would cost 9500. None of these quantities appears in a conventional account, and this part is about how they can be made visible without pretending that they reduce to one number.

Chapter 16

Valuing Feasible Futures

The owner has twelve months of planning to do for the four-unit building and a crew ceiling of 400. The allocations differ in what they leave. Over a horizon of 24 months, sending nothing to any unit leaves one unit lettable, yields a cumulative surplus of 32,200, and leaves a restoration backlog of 30,800. Other allocations give very different results. The decision is which to choose, and the difficulty is that the outcomes are lists of different kinds of thing: money earned, units still lettable, and work owed.

16.1 Incomparable outcomes

Simulate the building under all constant allocations of the ceiling in multiples of 100 across the four units. There are 70 of them. For each, record the cumulative surplus over 24 months, the number of units lettable at the end, and the backlog, the cost of restoring every unit to 90 ignoring decay during repair. Three of the 70 are not dominated on all three coordinates by any other: spending $(0, 100, 200, 0)$ gives a surplus of 45,300 with two units lettable and a backlog of 21,500; spending $(0, 200, 200, 0)$ gives 42,900, three units lettable and a backlog of 18,300; and spending $(100, 100, 200, 0)$ gives 42,900, two units lettable and a backlog of 17,700. Each of the three beats the others on some coordinate. In all three the worst unit, at condition 50, is allowed to fall to zero, which is what the ceiling forces once unit 4 would absorb the whole of it. The calculations are reproduced in the companion (check BOOK, Chapter 16).

Anyone who ranks these plans with a single number has chosen weights for the coordinates. If surplus alone is weighed the first plan wins. If lettable units at the end are weighed the second wins, and if the backlog alone matters the third. None of the three is wrong, but each choice is a normative decision that a single total conceals. The point of the chapter is that the structure of the comparison, an order that leaves some pairs unranked, is not a defect of the data. It is what remains when the distinctions are kept.

Definition 16.1 (Continuation profile and dominance)**.** A *continuation profile* is a vector $(v_1, \dots, v_m)$ of coordinates, each oriented so that larger is better (the backlog enters with

a negative sign). Profile x *dominates* y, written $y \preceq x$, if $v_i(y) \le v_i(x)$ for every i. The order is partial: profiles with some coordinates larger and some smaller are incomparable.

The three plans form an antichain in this order, which is the example, and a boundary case is a single-coordinate profile, for which the order is total and a scalar is exact. Theorem 5.1 of Chapter 5 says that no real-valued function reflects the order on the antichain, and that a multi-utility family does. For the three plans, which form an antichain, two functions suffice: one that ranks them A, B, C from highest to lowest and one that ranks them C, B, A, so that every pair is ranked oppositely by the two and no pair is dominated. The existence of such families is not in doubt for a finite order, and what is open for infinite sets is stated as a hypothesis to be fixed in Part IX.

16.2 Money as one representation

Monetary valuation is not excluded. A conversion rate from lettable units and from backlog into currency yields a scalar, and with a stated conversion every comparison can be made. What the chapter claims is narrower: the conversion is a choice, it is made by someone, and two people with different rates rank the same plans differently. A rate of \$20,000 per lettable unit at the end of the horizon and a charge of the full backlog values plan A at $45{,}300 + 2 \times 20{,}000 - 21{,}500 = 63{,}800$, plan B at $42{,}900 + 60{,}000 - 18{,}300 = 84{,}600$, and plan C at $42{,}900 + 40{,}000 - 17{,}700 = 65{,}200$, so that B is chosen. Setting the rate at zero returns A. Circumstances in which a scalar is justified, such as a single owner with a known discount rate and no concern for tenants, are listed in the exercises.

A related quantity from information theory appears in the literature on constraints as costs. For a probability measure P and an event B with $0 < P(B) \le 1$, the conditional distribution $P(\cdot \mid B)$ satisfies $D_{\mathrm{KL}}(P(\cdot \mid B) \,\|\, P) = -\log P(B)$, because $P(x \mid B)/P(x) = 1/P(B)$ on B. Requiring that the system stay within B, for example that all units remain lettable, restricts the distribution by that amount. The identity measures the restriction in nats. It does not give a monetary cost, a physical dissipation, or a conservation law, and the book uses it only as a measure of how much a constraint narrows what can happen.

What has been established. The outcomes of the building's plans form an antichain on three coordinates, a scalar ranking requires conversion rates that someone must choose, and the order can still be represented by a finite family of monotone functions. The decision now enabled is to report plans as profiles with their conversion assumptions. Assumptions still required are the choice of coordinates and, for the infinite case, the existence of a multi-utility representation. The question passed on to Chapter 17 is when a price can serve to coordinate decisions without a shared ranking.

Exercise 16.2. Compute the profile of the plan $(0, 0, 0, 400)$ over 24 months and say which of the three plans in the chapter dominates it.

Exercise 16.3. Prove that if a finite partial order has dimension 2 then two functions represent it, and exhibit them for the three plans.

Exercise 16.4. List three conditions under which a single monetary conversion of the three coordinates would be justified for a public agency, and say which fails first in practice.

Chapter 17

Prices as Compressed Coordination Signals

The crew ceiling is 400 per month, and the manager wishes to let the four units compete for it without examining each. A price per unit of crew spending would do this: each unit's caretaker would request spending until the marginal benefit equals the price, and the price would be set so that requests sum to the ceiling. The manager's question is whether such a price exists, whether it is a good summary of what the units need, and what it hides.

17.1 Scarcity prices as dual variables

Take the stylized allocation problem $\max \sum_i a_i \sqrt{r_i}$ subject to $\sum_i r_i \le B$ and $r_i \ge 0$, where r_i is the spend on unit i and a_i weights the benefit. Let $a = (1, 2, 2, 4)$ and $B = 400$. The problem is convex, Slater's condition holds, and the maximum is attained. The first-order conditions give $r_i \propto a_i^2$: since $\sum a_i^2 = 25$, $r_i = 400\, a_i^2/25 = (16, 64, 64, 256)$, which sums to 400. The multiplier is $\lambda^* = a_i/(2\sqrt{r_i}) = 1/8$ for every i. Each unit, facing the price $1/8$ per currency unit and solving $\max_r a_i\sqrt{r} - \lambda^* r$ independently, chooses $r = (a_i/(2\lambda^*))^2 = 16a_i^2$, which is the same allocation. The decentralized choices clear the budget at the multiplier.

Theorem 17.1 (P24: scarcity prices as dual variables)**.** *For a convex allocation problem with capacity constraint $\sum_j d_j \le c$, Slater's condition, and attainment, there is a multiplier $\lambda^* \ge 0$ supporting the optimal allocation, with complementary slackness $\lambda^*(c - \sum_j d_j) = 0$. If, in addition, the objective and constraints decompose across participants and each participant's choice is the individual Lagrangian problem at λ^*, then λ^* is a decentralized price. A multiplier alone does not establish a market implementation.*

Proof. Strong duality under Slater's condition gives a multiplier and a saddle point of the Lagrangian; complementary slackness follows from optimality of the saddle. If the Lagrangian separates across participants, minimizing over allocations at λ^* separates into

individual problems whose solutions together form an optimal allocation when the individual solutions are unique, as here. □

The hypotheses are strong and each does work. Convexity gives the multiplier. Separability lets the price decentralize, and if two units shared a repair that benefited both, their problems would not separate. Uniqueness of each individual solution matters because with several maximizers the price clears the market only after an assignment among them. The economic reading is that a scarcity price is a sufficient statistic for the units' requests only when their choices factor through it, as P11 requires.

17.2 What the price hides

In the housing model the price of 1/8 is the same for all four units and says nothing about the one fact that matters most: unit 4 at condition 50 needs 400 per month to hold and cannot be raised by anything the crew can supply, and the stylized concave benefit assigns it only 256. The price clears the budget and the unit crosses its threshold. The sufficient statistic fails because the real problem has thresholds and tiers, and a price derived from a smooth surrogate does not carry them. Supplementary records, here the condition and hold cost of each unit, are needed alongside the price. This is an instance of the general point that a price compresses the causal structure underlying a decision, and that the omission is tolerable only for decisions that factor through it.

An exchange-rate graph shows an obstruction of another kind. Suppose three internal units of account exist: money, crew hours, and materials credits. One crew hour exchanges for 40 in money, one materials credit for 20 in money, and one crew hour for 2.5 credits. Going from money to hours to credits and back to money multiplies a stake by $(1/40) \times 2.5 \times 20 = 5/4$, so the quoted rates admit an arbitrage of 25% on each round trip. In the log notation of P22, the sum of the log rates around the cycle is $\ln(5/4) \neq 0$, so no potential function (no single price level) is consistent with all three quotes. If every exchange carries a proportional fee f the round trip multiplies the stake by $(5/4)(1-f)^3$, and arbitrage disappears once $(1-f)^3 \leq 4/5$, that is $f \geq 1 - (4/5)^{1/3} \approx 0.0717$. A fee of 8% gives $(5/4)(0.92)^3 \approx 0.973 < 1$ and a fee of 5% gives $\approx 1.072 > 1$. With frictions a potential need not exist as an equality but a no-arbitrage condition replaces it (shortest-path duality), and the frictionless case is the limit. This is the first obstruction example of the book; it is a statement about an exchange graph and is not extended to institutional conflict in general.

What has been established. A scarcity multiplier exists under convexity and Slater, decentralizes under separability, and is a poor statement of a threshold problem, and the quoted exchange rates of three internal units can be inconsistent in a way a cycle product reveals. The decision now enabled is to use prices only for decisions that factor through them. Assumptions still required are convexity, separability, and an account of what the

price omits. The question passed on to Chapter 18 is how profit, the commonest summary of all, compares with what it omits.

Exercise 17.2. Recompute the allocation and multiplier for $a = (1, 1, 2, 2)$ and $B = 400$.

Exercise 17.3. Give an example with two participants whose objectives do not separate, so that the multiplier exists but does not decentralize.

Exercise 17.4. Choose a market price that a public agency uses to allocate a scarce resource and list two features of the situation that the price does not carry.

Chapter 18

Profit and the Location of Costs

The owner of the building spends the whole ceiling of 400 per month on unit 4, which holds it at condition 50, and nothing on the other three. At the end of twelve months the books show revenue of 42,000 (forty-two unit-months let), running costs of 12,600, and maintenance of 4800, a profit of 24,600. A lender sees a profitable enterprise. The decision is whether to distribute the profit, and the difficulty is that the profit was earned partly by letting three units that were left to decline.

18.1 Decomposing profit

Write reported profit as revenue less recognized costs, $\pi = R - C - M$, and separate the effect on capacity. Over the twelve months the conditions move from $(90, 70, 55, 50)$ to $(76, 42, 13, 50)$. Unit 3 has been unlettable for six of the twelve months, and unit 2 sits at 42, two tenths above the threshold. The cost of restoring the first three units to their starting conditions, at 100 per tenth ignoring decay during repair, is $1400 + 2800 + 4200 = 8400$. Define the *capacity-adjusted profit* as reported profit less this restoration cost: $24{,}600 - 8400 = 16{,}200$. More than a third of the reported profit was a transfer from the future to the present.

Definition 18.1 (Capacity-adjusted profit)**.** For a period with reported profit π, opening capacity k_0 and closing capacity k_1, and a restoration price $\rho(k_0, k_1)$ for returning capacity from k_1 to k_0, the *capacity-adjusted profit* is $\pi - \rho(k_0, k_1)$ when ρ is nonnegative, and π when capacity has grown. The *location of costs* identifies who bears ρ: the owner at a later date, the tenants through reduced habitability, or third parties through residuals.

The example has an owner bearing most of the cost at a later date, the tenants bearing the habitability loss in the months unit 3 was unlettable, and no third party. Two other allocations of the same crew time make the point. Spending $(100, 200, 100, 0)$ for twelve months gives a reported profit of 20,400 and closing conditions $(90, 70, 29, 2)$, so the adjusted profit is $20{,}400 - 7400 = 13{,}000$. The first allocation therefore leads on reported

profit by 4200 and by 3200 on adjusted profit, and neither number says that one unit has been lost to lettability in the first case and two in the second.

Remark 18.2 (Profit and capacity move independently). In the housing model the maintenance spend changes reported profit by one amount and the restoration charge by another, and by Chapter 8 the two change at different rates in different tiers. The comparison of allocations above exhibits activities with higher reported profit and higher restoration charge, and the opposite. The remark is not a law about firms; it concerns a model in which maintenance is the only capacity variable.

18.2 Profit supported by reproduction, and profit supported by deferral

A worked enterprise model separates the two. In the first year of the example 16,200 of the profit is supported by reproducible productivity, in the sense that it remains after charging the restoration; the remaining 8400 is supported by deferral. In a second year at the same allocation the units that have become unlettable earn nothing, the deferral has little further to give, and reported profit falls from 24,600 to 12,700. The pattern is general: profit supported by deferral is not recurring, and the lender who capitalizes the first year's profit at a multiple has capitalized the deferral. The numbers are in the companion (check BOOK, Chapter 18). The same decomposition applies where damage is unpriced rather than deferred, for instance to a commons, and Chapter 36 returns to this.

What has been established. Reported profit and capacity change can move independently, and the capacity-adjusted profit exposes what part of the reported figure rests on deferral. The decision now enabled is to report the restoration charge beside profit. Assumptions still required are a restoration price and a closing state of capacity. The question passed on to Chapter 19 is when a monetary obligation and a restoration obligation can be exchanged.

Exercise 18.3. Compute reported and capacity-adjusted profit for the allocation $(100, 100, 100, 100)$ over twelve months.

Exercise 18.4. Prove that capacity-adjusted profit equals reported profit exactly when $k_1 \geq k_0$ componentwise.

Exercise 18.5. Select a listed company with large deferred maintenance or environmental liabilities and describe how its reported profit would change under a capacity-adjusted definition.

Chapter 19

Monetary Debt and Restoration Obligations

The unit that Chapter 10 left at condition 38 is owed a restoration. Chapter 1 priced the restoration at 5200 and the owner now has a court order and a sum of 5200 to pay for it. The decision is whether the sum discharges the obligation. A mortgage is discharged by paying its balance, from any source, to any holder. A repair is discharged only by a unit that is in fact restored, and money is a means that works only if the restoration can be performed.

19.1 Why money does not suffice

The crew ceiling for the building is 400 per month. For a unit at 38 the deterioration is $d = 4$ per month, so spending the full ceiling gives a net change of 0: the unit stays at 38 whatever sum is paid, because the crew cannot spend more than 400 in a month. No amount of money restores the unit under the present ceiling. Raising the ceiling to 500 gives a net rise of 1 tenth per month in the lowest tier, so the unit climbs from 38 to $39, 40, \ldots, 50$ and then, once above 50, faster; it reaches 90 after 26 months, and total spending at the ceiling is $26 \times 500 = 13{,}000$. At 600 the time is 17 months and the spend 10,200; at 800, 10 months and 8000; at 1200, six months and 7200. The sum of 5200 in Chapter 1 was the cost of restoration if the repair could be done at once, ignoring the decay during repair. The sum needed depends on how quickly the work can be done, and the speed is set by capacity that money alone does not provide.

Theorem 19.1 (P29 extended: compensation restores capability only with availability, timing, and sufficiency)**.** *Let a restoration obligation require reaching a target set R by a deadline T. A monetary transfer m restores the capability if and only if (i) the actions needed are in the effective set under the available ceiling (availability); (ii) the minimal time to reach R under that ceiling is at most T, and the transfer is available by the dates at which spending is needed (timing); and (iii) the minimal total spend C^* to reach R within*

T satisfies $m \geq C^*$ (sufficiency).

Proof. If the effective set contains no action raising the state toward R then no spending reaches R whatever m is. Otherwise the fastest path under the ceiling has a minimal time t^* and a minimal spend $C^*(T)$ among paths reaching R within T. Restoration within T is possible exactly when $t^* \leq T$, spending only at dates when funds are available, and $m \geq C^*(T)$. □

The theorem is the generalization of Theorem H2 of the housing case, where the three conditions were derived for exit. Availability, here, is the ceiling being at least 500 and not the sum of money; the transfer of 5200 satisfies neither the availability condition nor the sufficiency condition at ≥ 500, since the minimal spend at that ceiling is 13,000. A transfer of 13,000 paid at the start, with a ceiling of 500, restores the unit in 26 months; the same sum with the ceiling left at 400 restores nothing. The calculations are reproduced in the companion (check BOOK, Chapter 19).

19.2 Monetization and substitution

Financial liabilities can usually be transferred and monetized because they are claims to fungible resources. The obligation to restore a unit to a tenant, a stream to repair ecological damage, or a duty to a particular community cannot always be transferred in this way, because the beneficiary, the place, or the capability that was lost is part of the obligation. Substitution fails when the relevant beneficiary or place matters, as when a payment to a regulator does not restore a wetland, or when the capability lost cannot be bought at any price in the time available. A boundary case where monetization works is a damage that has a fungible repair available in the same place at a known price. For such cases a bond or insurance contract is a good instrument, and Chapter 20 proposes how to record the distinction.

What has been established. A restoration obligation is discharged by the restored state and not by a sum, money restores only given availability, timing, and sufficiency, and the minimal spend rises as the ceiling falls. The decision now enabled is to ask, of every compensation offer, what capacity it activates. Assumptions still required are the ceiling, the deadline, and the dynamics of decay. The question passed on to Chapter 20 is how the books should show what is owed in each of these senses.

Exercise 19.2. Compute the months and total spend to restore a unit from 38 to 90 at a ceiling of 700, and compare with the figures for 600 and 800 in the chapter.

Exercise 19.3. Prove that for a ceiling at or below 400 no spending raises a unit at 50 or below, and identify the smallest ceiling that raises a unit at 38.

Exercise 19.4. Name an obligation that cannot be discharged by payment to a third party and say which of availability, timing, or sufficiency fails if a payment is offered.

Chapter 20

Accounting for Continuation

Two owners each hold a four-unit building, spend 400 per month on maintenance, and let all four units. Each reports a monthly profit of 2400, a balance sheet with no liabilities, and an identical rent roll. Building X has units at $(90, 70, 55, 50)$. Building Y has all four at 90. The lender has to decide on which to lend for ten years. The accounts give no reason to distinguish them, and the physical state gives every reason to.

20.1 The factorization problem

Let $F : S \to Y$ be the accounting representation that sends a state to the reported figures, and $P : S \to \{0, 1\}$ the property "the building can keep all units lettable for the foreseeable future". By P11 the accounts determine the property if and only if $F(s) = F(s')$ implies $P(s) = P(s')$. Here $F(X) = F(Y)$ and the property differs. For Y spending 100 on each unit holds all four at 90 forever, so the building is in the viability kernel. For X it is not: unit 4 at condition 50 cannot be raised by anything the crew supplies and holding it needs the whole 400 every month, except for a total shortfall of at most 1000 over all time before it falls below 40; units 1 to 3 would then receive at most 1000 in total, which cannot keep them lettable indefinitely since each loses at least one tenth per month. The accounts do not factor through the property, and no refinement of the profit figure alone can repair this.

20.2 Supplementary accounts

The proposal is a set of accounts that accompany the conventional ones. A *maintenance backlog account* records the restoration cost to a stated standard: 9500 for X and 0 for Y. A *recovery reserve account* records resources and permissions committed to corrective action. A *critical dependency register* lists complements controlled by a single supplier (Chapter 7). An *outstanding disposition duties account* lists residual obligations of the kind in Chapter 10. None of these is aggregated with the others, and none is converted

into a liability on the conventional balance sheet unless a rule of the accounting regime says so. They add auditable visibility without presuming a single number. A fifth item, the hold-cost to ceiling ratio, is 900/400 for X and 400/400 for Y, and it gives the lender the simplest available warning.

Coherent bookkeeping and sufficient description are different properties. A ledger that records postings as additions is a functor from histories to the additive monoid of balances, in the sense of P19: $F(g \circ f) = F(g) + F(f)$ and $F(\mathrm{id}) = 0$, and these equations follow from the posting rules. That X and Y have equal ledger entries month by month is consistent with F being a functor, and it shows that functoriality does not give sufficiency.

A further design distinction concerns obligations and their discharge. An obligation, a reported action, evidence of performance, and the discharge of the obligation are four different institutional events. A two-ledger arrangement keeps official spending commitments separate from community observations and reports their divergence without averaging them; this is documented in a design review and has not been tested here, so it is used only to illustrate the separation. The separation matters because an entry in the first ledger records an intention, an entry in the second records what was seen, and the gap between them is the information.

What has been established. Two buildings with identical reported figures can differ on viability, so accounts alone do not determine continuation, and a set of supplementary accounts records what the conventional ones omit. The decision now enabled is to require the supplementary accounts as conditions of long-term lending. Assumptions still required are agreed standards for the backlog and the register. The question passed on to Part V is how the knowledge needed to maintain the building is itself maintained.

Exercise 20.1. Compute the backlog and hold-cost ratio for a building with units $(90, 90, 60, 60)$ and say whether the ratio alone settles membership in the joint kernel for a ceiling of 400.

Exercise 20.2. Prove that a ledger whose postings are additive defines a functor from the path category to $(\mathbb{Z}, +)$ and give a property of the building that no such functor determines.

Exercise 20.3. Draft the four events (obligation, report, evidence, discharge) for a city's sidewalk-repair programme and say which records would be needed from each party.

Part V

Knowledge, Attention, and Productive Competence

Part V studies representations, memory, skill, automation, and verification as supports for, or threats to, competent continued action.

Housing thread. Inspection records, contractor knowledge, and automated scheduling are examined, along with what a passed inspection establishes.

In Part V the building is managed by people who must know its condition, remember what has been promised, keep the skill to act, and decide how far to rely on automatic scheduling. The shared parameters for automation are in the specification: automation saves 300 per month of supervision, practice costs 200 per unit of effort per month, competence erodes by $3/50$ per month under full reliance and gains $3/20$ per unit of practice, and the time to return to explicit control is $\lceil 1 + 7(1 - c) \rceil$ months at competence $c \in [0, 1]$.

Chapter 21

Economic Representations and Semantic Compression

An inspector grades each unit on a scale of 0 to 10 by dividing its condition by ten and rounding. A manager reads the grades and decides which units to hold with the crew's limited spending. The grades are compact, they can be passed between agencies, and they are easy to audit. The decision is whether the grades are good enough, and the difficulty is that the answer depends on what the manager is going to do with them. Two units can have the same grade and require different spending.

21.1 Task-relative adequacy

A representation $\rho : S \to Y$ compresses a state to a description, and it preserves some distinctions and discards others. It is not good or bad in itself. It is adequate for a task if the task's answer depends only on the description. Let $\tau : S \to Z$ be the task, for example the monthly holding cost of a unit, $z = 100\,d(k)$. By P11 the representation is *adequate for* τ if $\rho(s) = \rho(s')$ implies $\tau(s) = \tau(s')$, equivalently if $\tau = \bar{\tau} \circ \rho$ for some $\bar{\tau}$.

The score $\rho(k) = \text{round}(k/10)$ is not adequate for holding cost. Units at conditions 80 and 81 both score 8, yet their holding costs are 200 and 100. Units at 50 and 51 both score 5, with holding costs 400 and 200. The tier boundaries lie inside score classes, and a manager who uses the score to allocate a ceiling of 400 over four units will misjudge the units at the boundaries by factors of two. A grade defined directly by the tiers, A for conditions above 80, B for 51 to 80, C for 40 to 50, and D below 40, is adequate for holding cost, since each lettable grade has a single cost (100, 200, 400). It is not adequate for the task of predicting how many months the unit stays lettable if left alone, because within grade B a unit at 80 lasts 18 months and one at 52 lasts 4. The same compression is adequate for one task and inadequate for another, and the check is a calculation, reproduced in the companion (check BOOK, Chapter 21).

Definition 21.1 (Task-relative adequacy)**.** A representation $\rho : S \to Y$ is *adequate for*

the task $\tau : S \to Z$ if τ is constant on each fibre of ρ. It is *minimally adequate* if no strictly coarser representation is adequate. The minimal adequate representation for τ is the quotient of S by the kernel of τ.

An example is the tier grade for the holding-cost task. A boundary case is the identity representation, adequate for every task and costing the most to keep and transmit. A useful representation is one that is minimally adequate for the tasks it serves and whose omissions are known.

21.2 Compression in forecasting and classification

Forecasts and classifications compress in the same way. A forecast of the monthly rent roll summarizes the number of let units and discards their conditions, and is adequate for the task of next month's revenue but not for the revenue in two years. A classification of tenants as low or high risk discards the information that separates tenants whose risk comes from irregular income from those whose risk comes from large expenses, and is adequate for a single decision on the credit limit but not for the decision whether to offer a payment plan. The general lesson is that the question to ask of any category system is which decisions it was built for, and whether its use has drifted to others. Chapter 26 treats classification as an allocator of futures and takes this up.

What has been established. A representation is adequate relative to a task, the score used in the example is inadequate for the holding-cost task because tier boundaries fall inside score classes, and the tier grade is adequate for holding cost and inadequate for time to failure. The decision now enabled is to state, for each record, the tasks for which it is adequate. Assumptions still required are an explicit list of tasks. The question passed on to Chapter 22 is how adequate representations are retained, retrieved, and kept usable over time.

Exercise 21.2. Find the smallest number of grade classes that is adequate for the task "months of lettability remaining with no maintenance" on conditions 40 to 100.

Exercise 21.3. Prove that the minimal adequate representation for τ is unique up to relabelling.

Exercise 21.4. Take a credit score, a hospital triage category, or a school grade, and name one task for which it is adequate and one for which it is not.

Chapter 22

Memory as Organizational Capital

The manager of the building keeps a log of the condition of each unit. A flood in the office destroys it. The leases and the rent roll are intact, but the conditions are lost. The decision is whether the log was worth keeping, and the difficulty is that the cost of keeping it was small and visible and the cost of losing it appears only when the manager next tries to allocate a crew.

22.1 Retained, retrievable, and usable

Information held somewhere is not thereby available for action. It is *retained* if it exists on some medium, *retrievable* if it can be found at the moment needed at acceptable cost, and *usable* if the person finding it can interpret it for the decision. A log in a format that no one can read is retained and not usable. The institutional memory of a firm is the part of its records, routines, and understanding that is usable, and it is a form of capital in the sense of Chapter 7: it provides a capability on the condition that complements, such as readers, indexes, and formats, continue to be provided.

The value of the log can be calculated. With the log the manager knows the conditions $(90, 70, 55, 50)$ and the holding costs $(100, 200, 200, 400)$, a total of 900. Without it the manager knows only that every unit is lettable, that is, that each is at condition 40 or above. A prudent allocation assumes the worst tier for every unit, 400 each, a total of 1600. The difference of 700 per month is the value of the records to the holding decision. If keeping the log costs 20 per unit per month, or 80 per month, the log is worth keeping at any horizon. Reinspecting after the loss costs 150 per unit, a total of 600, once, against an excess holding cost of 700 per month until the conditions are known again. The calculation is reproduced in the companion (check BOOK, Chapter 22).

Proposition 22.1 (Value of a record)**.** *Let ρ be a record of the state and τ a decision cost function. The value of the record to a decision is the difference between the worst-case cost over states consistent with the information without the record and the cost with the record. It is nonnegative and zero exactly when the decision is the same at every state consistent*

with the information.

The proof is the definition: removing information enlarges the set of states over which the worst case is taken. The hypothesis doing work is that the manager guards against the worst case; a manager who assumed the average would compare differently, and the average needs a probability model that this book does not supply.

22.2 Forgetting without erasing obligations

Forgetting can reduce costs. Purging old tenant records, closing a database, or retiring a form reduces storage and handling. It does not erase obligations: a repair duty that a tenant reported and that the log recorded is still outstanding after the log is lost, and the owner is in breach when the deadline passes whether or not any record remains. The danger is that forgetting is treated as discharge. A record system should distinguish between archiving, which removes a record from active use while keeping it retrievable, and deleting, which does not, and should not allow deletion of records of outstanding obligations.

A distinction from the literature on verifiable records separates replay integrity from historical authority. A log that can be replayed to produce the same state, or that carries a hash chain, has integrity in the sense that it is self-consistent. It does not follow that the inputs it recorded were truthful, or that the commitments it records were legitimate, or that the branch an institution has selected is the correct one. An institution can keep reproducible accounts without establishing that the rent was in fact paid. Chapter 25 develops this, and the sheaf-theoretic treatment of when local records glue into a global one is deferred to Part XII (P21).

What has been established. Records are capital whose value is the reduction in worst-case cost of decisions, in the example 700 per month against 80 in upkeep, forgetting reduces cost but does not discharge obligations, and integrity is distinct from authority. The decision now enabled is to keep records of states and obligations that feed decisions with large worst-case costs. Assumptions still required are the decision cost function and a rule for dealing with worst cases. The question passed on to Chapter 23 is how routine skill, which also saves attention, is kept ready.

Exercise 22.2. Recompute the value of the log if the units are $(90, 90, 60, 60)$.

Exercise 22.3. Show that a record has zero value to a decision that is the same at every state consistent with the remaining information, and give an example of a nonzero record that nonetheless has zero value to a stated decision.

Exercise 22.4. Select an organization's archive and write the retention rule that would keep records of outstanding obligations and release the rest.

Chapter 23

Aspect Relegation and the Economics of Skill

The manager has set up an automatic scheduler that decides each month how to allocate the crew's spending. The scheduler works well for as long as the conditions stay within the range it was built for, and the manager has stopped checking it. The scheduler saves 300 per month in supervision time. The decision is how much to keep practising the manual allocation, and what makes it hard is that the skill erodes under reliance, the erosion is invisible until the scheduler fails, and by then the time needed to take back control is longer.

23.1 A margin for detection and return

Practice and routine change the attention a task requires. A relegated aspect of a task is performed without attention, which frees resources, and the freeing creates dependence on conditions remaining stable. The economic issue is recognizing when a task must return to explicit supervision.

Fix an operating region at condition 60 with the scheduler holding units at that level. The failure boundary is condition 40, so the distance is $d = 20$ tenths. If the scheduler silently stops, the worst rate of decline is $v = 4$ tenths per month, which is the rate in the lowest tier. If detection and return to explicit control takes τ months, the safe condition is $d > v\tau$, that is, $\tau < 5$. The exact tolerance is larger: a unit at 60 left alone is lettable for 8 months, with $k = 42$ after seven and 38 after eight, so any return within seven months is safe, and the bound $\tau < 5$ is conservative because it uses the worst rate for the whole period. The planner who needs a margin that does not depend on a particular path uses the bound.

Theorem 23.1 (P16: safe relegation requires a detection-and-return margin)**.** *Let an automated operating region lie at distance d from a failure boundary, adverse motion at rate at most v, detection and return to explicit control within time τ, and intervention*

able to arrest deterioration before the boundary. Then $d > v\tau$ is sufficient, together with allowances for observation error and stopping distance. If response time depends on retained competence, $\tau_t = \tau(c_t)$, with τ decreasing in c, the condition at date t is $d_t > v_t\,\tau(c_t) + $ allowances.

Proof. In time τ the state moves at most $v\tau < d$ toward the boundary, so it has not reached it when explicit control returns, and the intervention then arrests deterioration. The endogenous statement replaces τ by $\tau(c_t)$ at each date. □

With $\tau(c) = \lceil 1 + 7(1-c) \rceil$ months, the margin $\tau \le 4$ requires $1 + 7(1-c) \le 4$, that is $c \ge 4/7$. Competence therefore has to stay at or above $4/7$.

23.2 The practice that preserves the margin

Let competence evolve by $c_{t+1} = \min\{1,\ c_t - \frac{3}{50}\ell_t + \frac{3}{20}u_t\}$ with reliance $\ell_t = 1$ under full automation and practice $u_t \in [0,1]$. With $u = 0$ competence falls by $3/50$ each month, so from $c = 1$ it is below $4/7$ at month 8, where $c = 0.52$. Constant practice holds competence if $\frac{3}{20}u = \frac{3}{50}$, that is $u^* = 2/5$. Practice costs 200 per unit, so holding competence costs 80 per month, leaving a net saving of $300 - 80 = 220$ per month. The apparent saving of 300 is consumed entirely if practice costs 750 per unit, since $u^* \times 750 = 300$. The numbers are reproduced in the companion (check BOOK, Chapter 23).

The result is conditional. If the failure is rare and costly, practice pays, and Chapter 24 computes how rare. If practice is cheap relative to the saving, the answer is to practise. If the scheduler's savings are small and practice is expensive, the margin may not be affordable, and the correct inference is that the automation should not be adopted at that operating point, not that the margin may be neglected. The distinction between expenditure and damage to regenerative capacity is the same one: the scheduler's saving is an expenditure avoided, and the erosion of skill is a damage to the capacity that regenerates readiness.

What has been established. A margin $d > v\tau(c)$ turns the question of how much skill to keep into a calculation, in the example $c \ge 4/7$ requiring practice $u^* = 2/5$ at a cost of 80 per month, and the saving from automation is consumed if practice costs 750 per unit. The decision now enabled is to budget for practice as part of the cost of automation. Assumptions still required are the erosion and gain rates, which are empirical. The question passed on to Chapter 24 is how to compare policies over a horizon long enough for erosion to matter.

Exercise 23.2. Recompute u^* and the break-even practice price if erosion is $1/20$ per month and the gain is $1/10$ per unit of practice.

Exercise 23.3. Prove that if τ is decreasing in c and the dynamics are monotone, then a policy that maintains $c_t \ge c^\dagger$ at all t satisfies the margin whenever $c^\dagger$ does.

Exercise 23.4. Choose an occupation in which routine skill is automated, such as pilots or radiologists, and identify the failure boundary, the adverse rate, and the return time.

Chapter 24

Automation and the Preservation of Intervention

The scheduler of Chapter 23 has run for thirty months with no practice. It fails silently. The manager notices after a delay that depends on how much skill remains, and the decision, taken thirty months earlier, was whether to spend 80 per month on practice. The question is how large the expected cost of failure must be before the practice was worth buying.

24.1 The cost of a delayed return

Suppose the scheduler has been holding a unit at 60 and stops. If control returns after τ months, the unit has fallen to the condition reached after τ months of no maintenance, and the loss is the cost of restoring it to 60 at 100 per tenth, plus 700 for the month lost if the condition is below 40. With the return times of Chapter 23 the losses are 200 for $\tau = 1$, 800 for $\tau = 4$ (condition 52), 1400 for $\tau = 6$ (condition 46), and 2900 for $\tau = 8$ (condition 38, below the lettable minimum). Under full reliance competence falls to zero at about the seventeenth month, where the return time is eight months. A failure at month 30 therefore costs 2900 without practice, and 200 with practice at $u^{*} = 2/5$ that has kept competence at 1.

24.2 Comparing policies over the horizon

Over 36 months, full reliance saves $300 \times 36 = 10{,}800$ and reliance with practice saves $(300 - 80) \times 36 = 7920$, a difference of 2880. If failures occur n times over the horizon and each occurs late, the practice policy saves $n(2900 - 200) = 2700n$. Practice is the better policy exactly when $2700n > 2880$, that is $n > 16/15 \approx 1.07$. If the expected number of failures over three years is greater than about one, buy the practice; if less, the reliance policy dominates. The threshold is not a probability estimate. It is the break-even that

the manager can compare with an estimate of the scheduler's failure rate made from the manager's own records, and it exposes how sensitive the conclusion is to the return-time function, which here is an assumption. The calculation is reproduced in the companion (check BOOK, Chapter 24).

The qualitative point is that full automation lowers the immediate cost and raises the cost of the rare failure, because the failure arrives when the people who could respond have lost the ability to do so. An automation policy should therefore be assessed on output, error, oversight burden, retained competence, and recovery time together. A tool that offers isolated working copies, a command filter, or a stop control has limited effect here if the effects of an action cannot be recovered after the fact. The economic question about any such tool is how much review, intervention time, and recoverability it actually preserves, and no claim of zero risk is admissible.

What has been established. The cost of a delayed return rises with eroded competence, full reliance and reliance with practice are compared by a break-even number of failures, here 16/15 per three years, and the conclusion depends on the return-time function. The decision now enabled is to count expected failures before choosing a policy. Assumptions still required are the return-time function and the failure process. The question passed on to Chapter 25 is how to know that a check of the automation checks what matters.

Exercise 24.1. Recompute the break-even number of failures if practice costs 400 per unit.

Exercise 24.2. Show that the break-even n^* is increasing in the practice price and decreasing in the differential loss, and interpret.

Exercise 24.3. Select a system with automated operation, such as an elevator maintenance contract or an airline autopilot, and list the quantities needed to apply the break-even calculation.

Chapter 25

Verification, Proxies, and Institutional Confidence

An inspector checks each unit in the building against the rule that its holding cost must not exceed the crew ceiling of 400. Every unit passes: the holding costs are 100, 200, 200, and 400. The inspection report says the building meets the maintenance standard, and a lender takes comfort from it. The decision is how much confidence the report warrants. The units together require 900 against a ceiling of 400, and the check never added them.

25.1 What a passed check establishes

Theorem 25.1 (P17: verification transfers only through an adequate interpretation)**.** *Let M be a formal model, E the empirical system, and R a relation connecting their states. Assume soundness of overapproximation: every relevant empirical execution is R-related to a model execution, and R-related states satisfy that model safety implies empirical safety. Then a verified universal safety property of M transfers to E.*

Proof. Take an empirical execution e. By assumption there is a model execution m related to it. The property holds at every state of m, and the states of e are related to those of m, so by the safety-transfer assumption each state of e is safe. □

The hypotheses are the content. The inspection check is a verified property of a model in which the building is a collection of independent units, and the empirical system is a building with a shared crew. The relation R fails to preserve the safety property because the shared ceiling is not in the model, and by P09 each unit is feasible alone and the whole is not. Passing a check establishes what its specification and interpretation support, and the connection to the economic situation is an additional proof obligation. Liveness and probabilistic claims need stronger relations than the safety relation used here.

An executable version makes the mismatch explicit. In the companion the function that tests each unit's holding cost against the ceiling returns true for the building $(90, 70, 55, 50)$,

and the function that sums the holding costs and compares with the ceiling returns false (check BOOK, Chapter 25). The repair is not to add more checks but to restate the specification: the property of interest is the joint one, and the check must be written for it. A second example from the literature on ledgers has the same shape. A verifier checks seals against stored roots without recomputing them from the contents, and checks debits one at a time without aggregating repeated postings. Two debits of 600 against a stated restriction of 1000 each pass, and together they exceed it. A tampered ledger whose contents have changed and whose stored root has not passes the seal check. In both cases the check succeeds while omitting the relation a reader assumes it verifies.

25.2 Proxies and institutional confidence

Institutions rely on proxies: a certification stands in for competence, a benchmark for performance, an audit for honesty. Each carries the risk that the passed check is read more broadly than its specification warrants, and that behaviour is then organized around passing the check. By P15 a proxy that depends only on the rewarded measure cannot distinguish two activities with the same measure and different substantive contribution. The effect on confidence is that the institution may hold a belief about the building that its evidence does not support, and any financing decision that rests on the belief inherits the gap. The response proposed here is to record, alongside every certificate, the specification and the interpretation relation it relies on, so that a user can see which of the hypotheses of Theorem 25.1 are being assumed.

What has been established. A verified property transfers only through an interpretation that preserves safety, the unit-by-unit inspection passes while the joint requirement of 900 against 400 fails, and a check can succeed while omitting the relation a reader assumes. The decision now enabled is to attach the specification and interpretation to every certificate. Assumptions still required are an explicit interpretation relation and a joint property. The question passed on to Part VI is how the authority to act and to refuse, the rights that structure these dependencies, is itself organized.

Exercise 25.2. For the building $(90, 90, 60, 60)$ write the unit-by-unit check and the joint check and say which one the inspector should report.

Exercise 25.3. Give a model in which every safety property transfers under the relation of Theorem 25.1 and a liveness property does not.

Exercise 25.4. Take a certification such as a food-safety or financial-audit certificate, and state the model it verifies, the empirical system, and one way the relation between them could fail.

Part VI

Institutions, Power, and Distribution

Part VI analyzes classifications, property, contracts, exit, and public authority as distinctions that alter which continuations are accessible and to whom.

HOUSING THREAD. Tenure, deposits, notice periods, and a tenant's executable capacity to leave are placed in one model.

In Part VI the participants of the housing case are placed in a system of categories, rights, contracts, and public rules. The state of a tenant now includes a label assigned by someone else, the exit cost c includes the deposit that the label implies, and the owner's authority over the building is a bundle of permissions with different practical reach.

Chapter 26

Distinctions That Allocate Futures

A landlord asks two applicants for the same unit to pay a deposit. Applicant A is classified as low risk and pays one month's rent. Applicant B has a record of one late payment, is classified as higher risk, and pays two months'. Each has 3000 in savings and a monthly surplus of 200, and each will at some point want to move. The decision for the landlord is the size of the deposit; for the tenant it is whether the unit is worth taking. The difficulty is that the classification looks like a description of the applicant and works as a rule that allocates access. What the label says about the person is less important to the outcome than what the label does.

26.1 Labels as state variables

A classification $\kappa : H \to \Lambda$ assigns a label to each history, for instance the credit status derived from payment events. An *access rule* $\alpha : \Lambda \to \mathcal{P}(A)$ assigns to each label a set of permitted or priced actions. Combined, they determine the effective set of Chapter 3 as a function of the label. In the example the exit cost is $c_\lambda = 4000 + 1000 \cdot \mathbf{1}[\lambda = B]$, since the second deposit at the next unit is a further 1000. With savings of 3000 and a surplus of 200 per month, applicant A can exit at month 5 and applicant B at month 10, since $3000 + 200t \geq 5000$ first holds at $t = 10$. Two persons with identical resources and different labels therefore have different continuation languages, and by Chapter 2 the label must be part of the state. The calculation is reproduced in the companion (check BOOK, Chapter 26).

Definition 26.1 (Allocating distinction)**.** A classification κ with access rule α is *allocating* if there exist histories h, h' with identical physical state and resources and $\kappa(h) \neq \kappa(h')$ such that the continuation languages differ, $L(h) \neq L(h')$. It is *persistent* if the label changes only after a stated waiting condition.

The tenant classification is allocating. Suppose the label reverts to A only after twelve consecutive on-time payments. Then the tenant who paid late once at month 0 is labelled B throughout the year, and any move in that year carries the extra deposit of 1000. A

single late payment costs a tenant who moves in the year at least 1000, which is five months of the monthly surplus. A boundary case is a label with no access rule, such as a nickname, for which the continuation languages coincide.

26.2 Relation to institutional economics

The classification literature in institutional economics treats categories as devices that reduce transaction costs, and this is right as far as it goes: a deposit rule keyed to a label is cheaper than individual assessment. The additional point here is that the label is a state variable, so that its errors are not noise but allocations. A misclassified applicant, in the sense that the label does not reflect the facts the rule is meant to track, bears a cost of 1000 on exit in this model, and the cost is incurred precisely when the person is least able to bear it, at the moment of a forced move. The construction connects to the theory of distinction and continuation: drawing a line is an act that fixes which futures are open on each side. Which lines are drawn, who draws them, and how they may be revised are the subjects of the rest of the part.

What has been established. Labels derived from history allocate futures, the exit date of a tenant with fixed resources moves from month 5 to month 10 with a deposit label, and a persistent label imposes a cost on a single past event. The decision now enabled is to list, for every classification an institution uses, the access rule attached to it. Assumptions still required are the access rule and the revision rule. The question passed on to Chapter 27 is what authority a property holder has over the things the labels apply to.

Exercise 26.2. Compute the exit date for a tenant with 4000 in savings and surplus 300 per month under each label.

Exercise 26.3. Prove that a classification is allocating exactly when its label is not redundant in the minimal exact summary of the continuation language.

Exercise 26.4. Select a credential, such as a licence or a visa category, and write its access rule and its revision rule.

Chapter 27

Property as Structured Authority

The owner of a unit can use it, exclude others, modify it, sell it, borrow against it, and abandon it. Each of these is a permission, and they come in different combinations under different regimes. The owner of a mortgaged unit may not sell without the lender's consent; the owner of a listed building may not modify it; the owner of a leased unit may not exclude the tenant. The decision for a buyer is what is being bought, and the difficulty is that the word ownership names the same thing in all of these cases and the practical powers differ.

27.1 Permissions and capacity

Represent property in a thing as a set $P \subseteq \{\text{use}, \text{exclude}, \text{modify}, \text{transfer}, \text{encumber}, \text{abandon}\}$ of permissions held by a party. The structure is a lattice under inclusion, and holdings can be compared when one set contains the other. Permission is not capacity. The party can exercise a permission only if the action is also in the accessible set of Chapter 3. Two owners with the same permission set and different crew ceilings are not equally able to use their unit: the unit at condition 60 is let for 36 months at the full ceiling and for 14 at a ceiling of 100 (Chapter 7), though both own the unit in full. The effective authority of a party is the intersection of the permission set with what is accessible, and distribution of property in the sense that matters for continuation is the distribution of effective authority.

Encumbrance shows the interaction. Suppose a unit can be sold for 30,000 and carries a lien of 28,000. The owner's transfer permission is conditional on discharging the lien, and the net proceeds of transfer are 2000. If the owner holds only 500 in other assets, the option to abandon is differently priced: the lender's claim of 28,000 ranks ahead, so the private cost of abandoning is bounded by the owner's equity in the company, as in Chapter 10, and the burden of the residual of 2850 passes to others. The permission to abandon exists, and its exercise places burdens on persons who have no say in it, which is why the sixth permission is the one that regimes most often qualify.

Proposition 27.1 (Effective authority is monotone in capacity)**.** *If party i holds permis-*

sion set P_i and capacity profile z_i, and $P_i \subseteq P_j$ and $z_i \leq z_j$ pointwise, then the set of continuations available to i is contained in that available to j, under monotone dynamics.

The proof is that every plan available to i is available to j. The proposition has a converse failure: the party with more permissions and less capacity is not comparable with one holding fewer permissions and more capacity, and no ranking by permissions alone, or by wealth alone, can capture the order.

What has been established. Property is a bundle of permissions whose practical force depends on capacity, encumbrance conditions the permission to transfer, and the permission to abandon imposes burdens on others. The decision now enabled is to describe a holding by permissions and capacity together. Assumptions still required are the list of permissions in the regime and a model of capacity. The question passed on to Chapter 28 is how contracts, which reallocate permissions through time, change the position of the parties.

Exercise 27.2. For a unit worth 30,000 with a lien of 28,000 and selling costs of 1500, compute the net proceeds of transfer.

Exercise 27.3. Give two permission-capacity pairs that are incomparable under the order of the proposition.

Exercise 27.4. Select a property regime, such as a condominium or a tenancy in common, and list the permissions it withholds from each party.

Chapter 28

Contracts and the Evolution of Dependence

The owner of the unit at condition 50 needs 400 of maintenance every month to hold it, and a single vendor supplies it. The competitive price of maintenance is 100 per unit of 100 spending, so the vendor's price per unit is $q = 100$ at cost. The vendor proposes a multi-year agreement. The decision for the owner is the price at which to accept, and what makes it hard is that the agreement changes the owner's alternatives for later periods.

28.1 Switching costs and the bargaining position

The owner holds the unit for a horizon of $N = 12$ months and needs $u = 4$ units of maintenance each month. The unit earns 700 per month while held, so the owner will pay a price per unit of at most $q \le 700/u = 175$, above which holding earns less than abandoning. The owner has an alternative supplier at the competitive price 100, available at a one-time switching cost S. The incumbent can charge at most the price that leaves the owner indifferent between paying and switching: $N\,u\,(q-100) \le S$, so $q^* = 100 + S/(Nu) = 100 + S/48$. If $S = 1200$ then $q^* = 125$, a markup of 25 per unit, or 100 per month. The incumbent's price is capped from above by the owner's fallback of abandoning, $q \le 175$, so the extractable amount is at most $4 \times 75 = 300$ per month, which is the entire surplus of 700 less the competitive holding cost of 400.

Dependence accumulates when switching costs grow with incumbency. Suppose the vendor accumulates knowledge of the unit, so that each year the switching cost rises by 600: $S_n = 1200 + 600n$ after n years. Then $q_n^* = 125,\ 137.5,\ 150,\ 162.5,\ 175$ for $n = 0, \dots, 4$, and from $n = 4$ the cap of 175 binds. After four years the incumbent can extract the entire surplus above competitive holding cost, 300 per month, whereas in the first year the extractable amount was 100. The sequence of contracts, each of which was acceptable at its signing since it was no worse than switching at that date, moves the owner from a position of bargaining strength to one of near-total dependence. The numbers are

reproduced in the companion (check BOOK, Chapter 28).

Definition 28.1 (Cumulative dependence)**.** A sequence of agreements exhibits *cumulative dependence* if the maximal price extractable at date n under the participation constraint is increasing in n because the switching cost S_n is increasing in the history of agreements.

The example is the vendor above. A boundary case is a commodity input with a spot market, for which $S_n = 0$ at every date and $q^* = 100$ throughout. The result is conditional on the assumed growth of S_n; if the owner can retain the know-how, or the vendor's knowledge is transferable by documentation, the growth stops. This is an assumption about the technology of the relationship, and it is stated rather than derived.

28.2 Interventions that accumulate

The same structure applies to continuing intervention in a physical system. A unit at 50 that is held by monthly maintenance is dependent on that maintenance; the longer it is held without being raised, the longer the dependence lasts, and the position of the supplier is the position of whoever controls the maintenance. A development agency that supports a region with continuing transfers, a platform that subsidizes complementors, and a lender that rolls over a loan are in the same logical position. The analysis here has a limited scope: it identifies the quantity, the switching cost over the horizon, that bounds extraction, and it points to the places where regulation can act, namely on the growth of the switching cost through portability, documentation requirements, and limits on the length of exclusive agreements.

What has been established. The price an incumbent can extract is bounded by the switching cost over the horizon, a sequence of individually acceptable agreements can move that bound from 125 to 175, and the numbers depend on an assumed growth in switching cost. The decision now enabled is to examine switching-cost growth before signing a long agreement. Assumptions still required are the growth of S_n and the fallback of the dependent party. The question passed on to Chapter 29 is how unequal capacities to exit shape what contracts can be extracted.

Exercise 28.2. Compute q^*_n for $n = 0, \dots, 5$ with $N = 24$ and the same switching-cost sequence.

Exercise 28.3. Show that q^* is nonincreasing in the horizon N for fixed S and interpret the result for short and long contracts.

Exercise 28.4. Select a long-term relationship, such as a software licence or a franchise, and describe how the switching cost grows with incumbency.

Chapter 29

Distribution of Exit and Refusal

A housing authority is considering whether to allow landlords to offer prepaid leases to its ten applicants. The tenants' liquid resources are 3400, 5000, 9000, 11,400, 11,500, 11,700, 11,800, 12,000, 15,000, and 20,000, and the parameters are those of the front chapter. Supporters point out that the lease raises aggregate surplus and that the landlord's discount helps tenants. The decision is whether to permit it, and the difficulty is that the effects fall on different tenants differently, and some of them are invisible in the aggregate.

29.1 Unequal capacities to leave

Exit rights and exit capacities differ. A tenant has a right to leave a monthly lease on notice, and a capacity to leave that depends on resources. The distribution of the capacity, over employment, housing, platforms, financial relationships, and harmful environments, is a dimension of inequality that income alone does not capture. The ten applicants all have executable exit before the proposal: the smallest, with 3400, satisfies $x + 3e \geq 4000$ with $e = 200$. After the proposal, the seven with at least $P = 11{,}400$ are eligible. Of these, the ones with x in $[11{,}400, 11{,}800)$, that is 11,400, 11,500, and 11,700, lose the executable exit. Three of the seven eligible tenants, and three of the ten, are worse off in exit capacity, while aggregate surplus rises by $7 \times 800 = 5600$ if all eligible tenants sign. The three who lose are those just above the eligibility line, and the tenants at the bottom, who cannot sign, are unaffected.

Proposition 29.1 (P28: surplus gain with a shrunken safe exit set, lump sum for flow)**.** *Let an intervention exchange a lump sum P paid at the start for the removal of a flow payment ρ per month for $T > H$ months, with the lump sum not refundable, and let the exit cost be c and the baseline monthly surplus $e \geq 0$. Among eligible tenants ($x \geq P$), those who had an executable exit within H months and lose it are exactly those with $\max\{P, c - He\} \leq x < c - He + P - \rho H$. The set is nonempty if and only if $P > \rho H$ and $c > H(e + \rho)$. Aggregate surplus is unchanged by the transfer itself and rises by whatever cost saving the intervention produces.*

The proof is that of Theorem H1 with P and ρ in place of $(1-\delta)rT$ and r: resources on the lease at month t are $x - P + t(e+\rho)$ and on the monthly contract $x + te$, and the intersection of the three conditions gives the interval. The hypotheses carrying the result are non-refundability and the liquidity requirement of the exit cost. The condition $P > \rho H$ says that the lump sum exceeds what the flow would have supplied within the exit window, which is why a prepayment that is small relative to the window does not shrink the exit set. The simplification is the deterministic monthly surplus; the assumption most likely to fail is that exit is a cash threshold, since people sometimes exit by borrowing from relatives or by sharing, and the model would need those channels as explicit actions.

The consequence for policy is that a nominal exit right and an executable exit path are different objects with different distributions. A proposal that is Pareto improving in expected surplus can shrink the safe exit set of identified tenants, and a regulator who looks only at the surplus will not see it. The remedies available are those of Chapter 14: making the exit cost smaller, making the lump sum refundable in part, or compensating the tenants in the interval with a transfer that is available when it is needed. The latter is the subject of Theorem H2 of the housing case, and the numbers for the ten applicants are reproduced in the companion (check BOOK, Chapter 29).

What has been established. Exit capacity is distributed unequally and can be shrunk by a proposal that raises aggregate surplus, three of ten applicants lose executable exit in the example, and the condition for the shrinkage is general for lump-sum-for-flow trades. The decision now enabled is to evaluate any contract innovation by its effect on the distribution of executable exit. Assumptions still required are the cash threshold model of exit and non-refundability. The question passed on to Chapter 30 is who has authority to impose a remedy.

Exercise 29.2. Recompute the lost-exit set for $\delta = 0.10$ and the other parameters of the front chapter.

Exercise 29.3. Prove that if the lump sum is refundable on exit the lost-exit set is empty.

Exercise 29.4. Name an exit that is nominally free and executably costly in employment, banking, or platform membership, and describe the resource threshold.

Chapter 30

Public Authority and the Admissibility of Constraints

The city council wants to require owners to repair reported defects, and must decide when a repair obligation may be treated as discharged. One proposal is that the obligation lapses after a waiting period of three months from the report. Another is that it is discharged only when an inspector certifies the unit is restored. The council also holds citizen reports, which it intends to publish grouped by neighbourhood. The decision is the rule, and the difficulty is that a procedural condition and a substantive fact can diverge. The council must also decide on what authority it imposes the constraint and how the people constrained can contest it.

30.1 Procedure and evidence

Take a unit at condition 51 that the owner has been ordered to restore to 60. The crew ceiling is 400, and the rank to reach 60 is five months (Chapter 14). Under the waiting rule the obligation lapses after three months. An owner who spends the full ceiling for three months leaves the unit at 57, which is below 60. An owner who spends nothing leaves it at 41, and both owners are discharged. The rule does not distinguish them, and the first owner, who made a full effort, is also discharged short of restoration. Lengthening the wait to five months gets the honest owner to 61 and still discharges the owner who does nothing, whose unit is at 33, below the lettable minimum. An elapsed period is an enforceable procedural condition. It cannot establish successful repair, because what it measures, time, is independent of the quantity of interest. The evidence rule discharges only at a restored condition, but it must give the owner at least the rank in time, which is five months, or the honest owner cannot satisfy it. A rule that closes only on procedural authorization and specified evidence, with a route for reopening a repair that later fails, is the design this chapter proposes. It is a proposed institutional design and not evidence of deployment. The condition values are reproduced in the companion (check BOOK,

Chapter 30).

A second feature of the design is the treatment of reports. Grouping reports by location suppresses duplicate influence from a single source but does not establish that the reports come from independent witnesses, and a minimum cluster size does not protect against all differencing attacks. Suppose two published groups $A = \{1, 2, 3\}$ and $B = \{2, 3, 4\}$ each satisfy a minimum size of three, with counts $n(A) = 2$ and $n(B) = 1$ of residents reporting a leak, each resident reporting 0 or 1. Then $x_1 + x_2 + x_3 = 2$ and $x_2 + x_3 + x_4 = 1$, so $x_1 - x_4 = 1$, which forces $x_1 = 1$ and $x_4 = 0$. Resident 1 has reported and resident 4 has not, though no group of size less than three was published.

30.2 Authority, contestability, and legitimacy

The previous chapters established that constraints can protect and can dominate. The viability formalism can say whether a constraint is satisfiable; it cannot say whether it is legitimate. Legitimacy depends on who establishes the continuation constraints, by what procedure, whether those subject to them can contest them, and how they are revised. A constraint is *contestable* if there is a procedure by which an affected party can have it reviewed and, if the review succeeds, changed. Preservation language, that a regime exists to protect the future, can legitimate domination if it is exempt from contest; the test is whether the constraint can be questioned by those who bear its costs. An institution that removes the route for reopening a failed repair, or that excludes tenants from the review of the standard, has made the constraint non-contestable in the sense used here, whatever its stated aim.

What has been established. A waiting period cannot establish a repair, an evidence rule needs at least the rank in time to be satisfiable by an honest owner, grouping and a minimum cluster size do not prevent differencing, and the legitimacy of a constraint depends on contestability that the formalism does not supply. The decision now enabled is to require evidence and reopening in discharge rules. Assumptions still required are an agreed evidentiary standard and a body with authority. The question passed on to Part VII is how commitments held by different parties compose into a feasible collective arrangement.

Exercise 30.1. For a unit at 53 and a target of 60 compute the rank at a ceiling of 400 and the shortest evidence-rule period that an honest owner can satisfy.

Exercise 30.2. Construct a differencing attack against groups of size four and show how adding noise changes it.

Exercise 30.3. Choose a public constraint and describe who can contest it, by what procedure, and with what effect.

Part VII

Markets, Firms, Commons, and Coordination

Part VII asks when individually acceptable commitments compose into a feasible collective arrangement, in markets, firms, platforms, and commons.

Housing thread. A shared crew serves several buildings, listings run through a platform, and a cooperative maintenance fund pools contributions.

In Part VII the crew, the listings, and the maintenance fund are shared among several buildings and owners. The question that runs through the five chapters is the one P09 posed in its simplest form: commitments that are individually acceptable need not be jointly feasible, and the institutions of exchange, the firm, the platform, and the commons are different ways of repairing that failure or failing to.

Chapter 31

Exchange and Compatible Commitments

Two building owners each contract with the repair crew for 300 of work per month. Each contract is within the crew's ceiling of 400, and the crew's manager signs both. The decision for the crew is whether the contracts can be honoured, and for a regulator whether bilateral exchange, each exchange being acceptable to both parties, produces a feasible arrangement. The difficulty is that feasibility is a property of the whole set of commitments and each exchange is judged one at a time.

31.1 Exchange as transformation of resources, permissions, and obligations

An exchange transforms the holdings of two parties: resources move, permissions are granted, and obligations are created. It is acceptable to each party if the successor state is in that party's acceptable set. Chapter 13 showed that acceptability does not suffice for viability, and the present chapter adds the shared-resource obstruction. For divisible simultaneous demands d_j on a capacity c, P09 states that $d_j \le c$ for every j does not imply $\sum_j d_j \le c$, and that $\sum_j d_j \le c$ is necessary and sufficient for joint feasibility in the static divisible case. With demands $(300, 300)$ and $c = 400$, each is feasible alone and the sum is 600. The physical version in the housing model is the pair of units at $(50, 50)$: each is viable alone under the ceiling, since holding costs 400, and the pair is not, since together they need 800. At $(51, 51)$ the pair is viable, because each needs 200 and the sum is 400. The best split at $(50, 50)$ keeps both lettable for six months (check H4 and BOOK, Chapter 31).

Theorem 31.1 (P10: composition under independence, and its failure)**.** *Two systems with product state space, independent actions, product disturbance sets, factorized dynamics, nonempty action and disturbance sets, and acceptable set $G_1 \times G_2$ have robust kernel $K_{12} = K_1 \times K_2$. With a shared resource the inclusion $K_{12} \subseteq K_1 \times K_2$ can be strict.*

Proof. For $\supseteq$, combine the witnessing policies. For $\subseteq$, fix an allowed disturbance selection for system 2 and simulate system 2 inside a joint policy; the induced policy depends only on system 1's history and keeps it in G_1 against all of its disturbances. The strict case is the pair $(50, 50)$ under a shared ceiling of 400, where each coordinate is in its own kernel and the pair is not in the joint kernel. □

The hypotheses carrying the result are independence of actions and the absence of a shared budget. A shared ceiling, coupled timing, or cross-system effects violate the premises, and the theorem states what must be established before local guarantees assemble without further analysis. The statement is reported, with and without the coupling, in the companion (Appendix B).

31.2 Reservations and concurrent requests

Reservations introduce time into exchange. A resource is withheld for a period, may expire, and returns. Authorization and commitment must stay consistent under concurrent requests. Suppose two requests for 300 arrive when 400 remains. If each is authorized by reading the remaining capacity and then committing, both read 400, both pass the check $300 \leq 400$, and the total committed is 600. If the check and the commitment are made together, the first commits and leaves 100, and the second fails. The failure of the first procedure is the failure of P09 at the level of the protocol, and a design that keeps authorization and commitment together restores consistency at the price of serializing requests. A documented case of this kind, two customers each offered the last item, has the same structure, and the present book treats it as an illustration of the obstruction and makes no claim about any particular system.

What has been established. Individually acceptable exchanges compose only under independence, the pair $(50, 50)$ shows strict failure under a shared ceiling, and concurrent authorization must be tied to commitment. The decision now enabled is to require a joint feasibility check for every set of commitments on a shared resource. Assumptions still required are an explicit list of shared resources and the timing of draws. The question passed on to Chapter 32 is how a firm resolves such conflicts by authority and what it costs.

Exercise 31.2. For three owners each requesting 200 against a ceiling of 400, find the largest number of requests that can be honoured and the allocation rule that selects them.

Exercise 31.3. Prove that the sum condition is necessary and sufficient for joint feasibility of divisible simultaneous demands and give a counterexample for indivisible demands.

Exercise 31.4. Describe a reservation system you have used and the point at which authorization and commitment could diverge.

Chapter 32

Firms as Organizations of Complementary Continuation

The owners of the four units could contract with the crew separately, or they could form a firm that employs the crew and allocates its time. The decision is which to do, and the difficulty is that the choice changes the way conflicts over the shared ceiling are settled: by negotiation among owners in the first case, and by authority in the second.

32.1 Coordination, authority, and the boundary

A firm coordinates tasks, assets, knowledge, and promises that cannot be maintained through isolated transactions alone. The explanation of its boundary in the literature, from Coase [6], is in terms of the costs of using the market; the present account adds the feasibility of the joint arrangement. The four owners together request 900 of crew spending against a ceiling of 400. Under separate contracts the conflict is settled by whoever contracts first, or by price, and the allocation can be one of the dominated plans of Chapter 16. Under a firm an authority chooses among the nondominated plans, such as $(0, 200, 200, 0)$ with three lettable units after two years, and bears responsibility for the choice. Authority also concentrates dependence: the firm's decisions are not contested by the units the way a contract can be, and the distribution of the residual responsibilities, the obligations and burdens that no contract assigns, is shifted to the firm.

32.2 Assume-guarantee composition

When components are developed separately, the composition of their guarantees requires care. Suppose the crew guarantees that it will supply the spend requested by each building each month, on the assumption that the total request does not exceed 400, and each building guarantees that its unit stays at condition 90 on the assumption that the crew supplies the request. Taken at face value the two statements are circular: the crew's

guarantee needs the buildings' behaviour, and the buildings' needs the crew's. A circular argument can prove too much. A model in which both guarantees fail satisfies the two implications "if the other holds, mine holds", so the pair of implications does not entail that both hold.

Theorem 32.1 (P27: well-founded circular assume-guarantee composition)**.** *Let components carry assumption-guarantee pairs (A_i, G_i) on safety properties, with shared capacity appearing in the interface as a time-indexed assumption on the combined draw. If each guarantee at step $n+1$ depends only on assumptions through step n, then the conjunction of the guarantees holds for the composition at every step, by induction on n.*

Proof. At step 0 the guarantees hold by initialization. If the guarantees hold through step n, then the assumptions through step n hold, because each assumption is a guarantee of another component, and so the guarantees at step $n+1$ hold by hypothesis. The argument is that of Abadi and Lamport [1]. □

In the example the temporal discipline is satisfied: the request at month $n+1$ is a function of the condition at month $n+1$, which depends only on supply through month n. With four units at 90 and a ceiling of 400, the requests are 100 each, the supply equals the request, and the units stay at 90 for every month, which the companion confirms over 24 months (check BOOK, Chapter 32). The theorem is restricted to safety properties and to the stated lag; if the crew's guarantee at month n depended on the request at month n and the request depended on the supply at month n, the induction would not start.

What has been established. A firm settles conflicts over a shared ceiling by authority and concentrates residual responsibility, and separately developed components compose only under a time-indexed assume-guarantee discipline that excludes same-step circularity. The decision now enabled is to specify interfaces with explicit lags. Assumptions still required are safety properties and the lag structure. The question passed on to Chapter 33 is what happens when the shared institution controls access to the opportunity domain.

Exercise 32.2. Show by an explicit two-state model that the circular implications are satisfied by the all-fail behaviour.

Exercise 32.3. State and prove the induction for three components in a ring, each guaranteeing at step $n+1$ on the assumption of its predecessor through step n.

Exercise 32.4. Describe a firm boundary decision, such as insourcing a service, and say which dependencies and residual responsibilities change.

Chapter 33

Platforms and Infrastructural Dependence

A listing platform lets owners advertise units to tenants. The platform charges a monthly fee and also decides how much verification of each listing's condition to require. An owner considers joining, which needs a one-time investment of 3000 in photographs and inspection documents that are of no use elsewhere. The decision for the owner is whether to join, and for the platform it is the fee and the verification standard. The difficulty is that the platform controls access to the domain of choice that it establishes, and can change its rules after the owner has invested.

33.1 Quality, participation, and revenue

A platform's revenue depends on how much it charges and how many participate, and participation depends on the quality it requires. Let $q \in [0,1]$ be the verification standard, $M(q)$ the number of participating listings, and revenue $R(q) = (\alpha q + \beta) M(q)$ where the fee rises with the standard. Then $R'(q) = \alpha M(q) + (\alpha q + \beta) M'(q)$, and revenue can fall as quality rises if participation falls strongly enough. Take $\alpha = 100$, $\beta = 200$, and $M(q) = 1000 - 400q$. Then $R(q) = (100q + 200)(1000 - 400q)$, $R'(q) = 20{,}000 - 80{,}000q$, positive for $q < 1/4$ and negative above. Revenue is 200,000 at $q = 0$, 202,500 at $q = 1/4$, 200,000 at $q = 1/2$, and 180,000 at $q = 1$. A platform maximizing revenue sets a verification standard of $1/4$ although tenants would gain from a higher one. This is a conditional incentive result, stated for the given participation function; the book does not claim that platforms in general under-provide quality.

33.2 Appropriation of complementary investment

The investment of 3000 is specific to the platform. Suppose the owner's outside option is to let by another channel with a surplus of 600 per month, and on the platform the

surplus is 700 less a fee. If the platform's fee rate is φ of the rent of 1000, the owner stays once joined if $700 - 1000\varphi \geq 600$, that is $\varphi \leq 10\%$. The sunk investment does not enter the stay decision. A platform that can change its rules unilaterally can therefore raise the fee to 10% after the owner has invested, and a forward-looking owner anticipates this. With $\varphi = 10\%$ the owner earns 600 on the platform and also paid 3000, so joining is dominated at every horizon N since $600N - 3000 < 600N$. With a committed fee of 5% the owner earns 650 per month, the gain over the outside option is 50 per month, and joining pays only if $50N \geq 3000$, that is $N \geq 60$ months. The platform's inability to commit prevents the investment altogether, and commitment, whether through contract, regulation, or portability, enables it. Interoperability and portability matter for the same reason: they improve the outside option and so lower the fee the platform can extract after the investment is sunk. The numbers are reproduced in the companion (check BOOK, Chapter 33).

What has been established. A revenue-maximizing verification standard can fall short of the standard that tenants prefer, and unilateral rule changes after sunk investment defeat the investment unless the platform can commit. The decision now enabled is to assess a platform by its commitment structure as well as its fees. Assumptions still required are the participation function and the outside option. The question passed on to Chapter 34 is how a community maintains a shared resource without a central owner.

Exercise 33.1. Maximize $R(q)$ for $M(q) = 1000 - 300q$ with the same α and β.

Exercise 33.2. Show that the maximum fee the platform can extract after the investment is independent of the size of the sunk investment, and state the exception.

Exercise 33.3. Select a platform that you know and identify the complementary investments made by participants and the rules the platform can change unilaterally.

Chapter 34

Commons and Shared Maintenance

The four owners decide to fund the crew jointly, each contributing to a common pool, and the ceiling is the pool. The decision is how much each contributes and how the pool is allocated. The difficulty is that contributions are voluntary, benefits are shared, and monitoring is costly. Common ownership of the pool does not by itself tell which units get repaired.

34.1 Contribution, allocation, and replenishment

Commons arrangements are defined by who may use the resource, who must contribute, how use and contribution are monitored, what sanctions apply, and how conflicts are resolved [18]. Consider three allocation rules for a pool of 400 over 24 months for the units $(90, 70, 55, 50)$. Equal shares of 100 each give a cumulative surplus of 33,100, one lettable unit at the end, and a backlog of 23,200. A need-based rule that sends 200 each to the second and third unit and nothing to the first and fourth gives 42,900, three lettable units, and a backlog of 18,300. If the first owner stops contributing, the pool falls to 300 and the allocation $(0, 100, 100, 100)$ gives 35,500, one lettable unit, and a backlog of 27,000. The same pool and the same ownership yield very different outcomes depending on the allocation rule. The pooled arrangement is common ownership; the plan chosen is the organization of maintenance.

Free riding appears when an owner withholds the contribution of 100 and the pool shrinks. The first owner gains 100 per month and, in this example, loses nothing immediately because unit 1 is the slowest to deteriorate; the others bear the loss. A sanction must exceed the gain from free riding at the probability of detection, and monitoring has a cost that must be paid out of the same pool. These are the standard design questions, and the housing model contributes a specification of what the pool is for, namely maintaining a set of conditions against thresholds, so that the design can be judged by the

profile it produces rather than by contribution levels. The numbers are reproduced in the companion (check BOOK, Chapter 34).

What has been established. Common ownership and the organization of maintenance are distinct, three allocations of the same pool yield different profiles, and free riding shifts losses to others. The decision now enabled is to evaluate commons rules by the profile of conditions they maintain. Assumptions still required are monitoring and sanction technology. The question passed on to Chapter 35 is how contributions are measured and rewarded.

Exercise 34.1. Compute the 24-month profile of the allocation $(0, 200, 100, 100)$ and compare with the three in the chapter.

Exercise 34.2. Show that the equal-share rule is not on the nondominated frontier of Chapter 16.

Exercise 34.3. Select a commons you know and list its contribution rule, monitoring, sanction, replenishment, and conflict-resolution mechanism.

Chapter 35

Contribution, Reward, and Measurement

A cooperative rewards members for maintenance work using a currency backed by recorded work. Each completed item of 100 spending earns one token. The decision is whether the token accurately rewards contribution to the cooperative's aim, which is that the units stay lettable, and the difficulty is that tokens count items while the aim concerns conditions.

35.1 Proxy rewards and what they miss

By P15, a reward that depends only on the measured activity m cannot distinguish two actions with the same m and different substantive contribution V, and if $m(a) > m(b)$ while $V(a) < V(b)$ a strictly increasing reward ranks them opposite to V. The housing model shows this directly. Take four units at 90 and a horizon of 28 months. Preventive maintenance, 100 per unit per month, generates $4 \times 28 = 112$ items, and keeps all four units at 90. Deferring all maintenance for 28 months and then restoring each unit from 38 to 90 requires 52 items per unit, 208 items in all. The deferral path earns almost twice as many tokens as the preventive path, and it is the path that has left all four units at the edge of being unlettable. The measure is increasing in the damage. Maintenance and preventive work are particularly hard to recognize through simple activity counts because success consists of an absence, units that did not fail. The numbers are reproduced in the companion (check BOOK, Chapter 35).

Proposition 35.1 (Activity counts can invert the ranking)**.** *There are two maintenance policies a, b with item counts $m(a) > m(b)$ and with the number of units lettable at the end of the horizon $V(a) < V(b)$.*

The example above is the proof: the deferral path is a and the preventive path is b. The point is a limit on information and not a theory of behaviour. Behaviour requires the participants' costs, available actions, and preferences to be specified, and a reward system

can nonetheless be designed to depend on a measure of V, such as the condition record of Chapter 22, or on a combination of measures audited by an independent party. Such a design inherits the problems of verification (Chapter 25): the condition record must be reliable and the interpretation relation must be stated.

35.2 Rewards change what they measure

When a reward depends on a measure, the measure becomes a target. A cooperative that pays per item may find that items are divided into smaller items, that work is done where it is countable and not where it is needed, and that restoration is preferred to prevention. A worked case in the literature is the work-backed currency, and the lesson drawn here is the structural one: any activity count is subject to inversion, and a reward should be tied to the state the activity is meant to produce.

What has been established. A per-item reward ranks deferral and restoration above prevention in the housing model, 208 items against 112, and activity counts can invert the ranking of substantive contribution. The decision now enabled is to base reward on measured conditions, with their verification stated. Assumptions still required are a reliable condition record and a verification chain. The question passed on to Part VIII is how these structures work when the resource is a living system, a person, or a development path.

Exercise 35.2. Compute the item counts for preventive maintenance at 90 and for deferral to condition 60 followed by restoration, over 20 months.

Exercise 35.3. Prove that for any strictly increasing reward of the item count, the deferral path in the chapter ranks above the preventive one.

Exercise 35.4. Select a reward system, such as a bonus on tickets closed or on repair visits completed, and describe a way it can invert the intended ranking.

Part VIII

Ecology, Care, Technology, and Development

Part VIII applies the framework to ecological reproduction, care, productive infrastructure, development, and coordinated transitions, and closes the economic text.

Housing thread. An energy retrofit is a transition with a viable path, and damp and mould are ecological conditions of the dwelling.

The last part of the economic text takes the same discipline to settings where the maintained thing is a living system, a person, a technology stack, a community's capacities, or a coordinated transition. The housing case contributes two additions: damp and mould as ecological conditions of the dwelling, and an energy retrofit that has to be financed while the building continues to be let.

Chapter 36

Ecological Reproduction and Economic Time

A community harvests a stock of fish worth 10 per unit at the dock. The stock is at 80 and the question is how much to take now. Selling down to 20 pays 600 at once. The decision is hard because the stock regenerates on its own timescale, which is not the timescale of the community's finances, and because below a threshold it does not regenerate at all.

36.1 Regeneration, thresholds, and horizons

Let the stock s regenerate by $g(s) = \lfloor s(100 - s)/400 \rfloor$ for $s > 20$ and $g(s) = 0$ for $s \le 20$, a logistic growth with a depensation threshold at 20. The values are $g(30) = 5$, $g(40) = g(50) = g(60) = 6$, $g(70) = 5$, $g(80) = 4$. The largest constant harvest that keeps the stock fixed is 6 per period, at stocks of 40 to 60. A constant harvest of 10 from a stock of 50 drives it to $46, 42, 38, 33, 28, 23, 17$ in seven periods and to a negative value in the ninth: the community has fallen below the threshold and cannot recover on its own, so the restoration time is not long but unbounded. The calculation is reproduced in the companion (check BOOK, Chapter 36).

Two plans for the stock of 80 differ in their financial and ecological consequences. The liquidation plan harvests down to 20 for 600 and has no further harvest. The managed plan harvests down to 50 for 300 and then takes 6 per period, worth 60. With a discount rate ρ per period the managed plan is worth $300 + 60/\rho$ and is better than liquidation exactly when $60/\rho > 300$, that is $\rho < 1/5$. A community whose financial horizon implies a discount rate above twenty percent per period will choose to liquidate, and it will be rational given its horizon. The economic problem is that the horizon is imposed by the financing, not by the ecology.

36.2 Damp and the dwelling

The dwelling has an ecology of its own. Damp and mould are conditions that arise when moisture is not removed and that accelerate deterioration. In the housing model, add one tenth of extra deterioration per month in every tier. A unit at 90 that is left alone is lettable for 28 months without damp and 18 with it, and a unit at 60 for 8 months without damp and 6 with it. The effect is the same as raising the hold cost of every unit by 100 per month, so that the holding costs become 200, 300, 300, and 500 for the four units, a total of 1300 against a ceiling of 400 and unit 4 beyond the ceiling alone. A condition that reproduces itself, as damp does, shifts the thresholds of Chapter 8 and is a regeneration problem in miniature.

What has been established. A depensation threshold turns an unsustainable harvest into an irreversible one, the managed plan beats liquidation only for discount rates below one fifth per period, and a self-reproducing condition such as damp shifts the viability thresholds of a dwelling. The decision now enabled is to check whether a financing horizon is compatible with the regeneration time before accepting it. Assumptions still required are a regeneration function and a threshold, which are empirical. The question passed on to Chapter 37 is how the maintenance of persons is similarly constrained.

Exercise 36.1. Compute the stock path under a constant harvest of 7 from a stock of 50 and the period at which it falls below 20.

Exercise 36.2. Prove that the largest sustainable constant harvest is $\max_{s>20} g(s)$ and find the set of stocks that attain it.

Exercise 36.3. Select a managed resource, such as a fishery or aquifer, and say whether its financing horizon is compatible with its regeneration time.

Chapter 37

Care and the Maintenance of Persons

A family member cares for a relative who needs 150 hours of care a month. The caregiver can work 120 hours a month without accumulating exhaustion, and each hour above that adds to an exhaustion level that, at 60, leaves the caregiver unable to continue. The decision is how to arrange the care. The difficulty is that two different things are at stake: the recipient's wellbeing, and the services the caregiver's activity supplies, and a framework that tracks only the second misses the caregiver's constraints.

37.1 The caregiver's viability

Let exhaustion x evolve by $x' = \max\{0,\, x + H - 120\}$ where H is the hours worked in the month, and let the acceptable set be $x < 60$. At $H = 150$ the exhaustion rises by 30 per month, from 0 to 30 and then to 60, so the caregiver is outside the acceptable set after two months. The viability kernel for the caregiver alone allows $H \leq 120$ forever and no more, and the recipient's need of 150 cannot be met from the caregiver's own hours in a viable way. The gap of 30 hours per month must be filled by someone: at a paid rate of 25 per hour this is 750 per month. The unpaid hours of the caregiver, valued at the same rate, are $120 \times 25 = 3000$ per month of services that appear in no conventional account. The calculation is reproduced in the companion (check BOOK, Chapter 37).

Proposition 37.1 (Sustainable care hours)**.** *Let exhaustion evolve by* $x' = \max\{0, x + H - h^*\}$ *with sustainable hours* h^* *and acceptable set* $x < \bar{x}$*. The caregiver can supply* H *hours every month forever if and only if* $H \leq h^*$*, and for* $H > h^*$ *and starting exhaustion* x_0 *the number of months before breakdown is* $\lceil (\bar{x} - x_0)/(H - h^*) \rceil$*.*

The proof is that exhaustion rises by $H - h^*$ each month while positive, so it crosses $\bar{x}$ after the stated number of months, and it does not rise when $H \leq h^*$. In the example $\lceil 60/30 \rceil = 2$. The hypotheses are the linear accumulation and the absence of recovery

from rest except through the floor at zero; if rest hours reduced exhaustion below the floor of recent work, the sustainable hours would depend on schedule and not only on the monthly total.

37.2 Dependency and the danger of valuing persons by contribution

Care is skilled, materially supported work that sustains capacities and relationships. Recognizing it in accounts, by valuing the unpaid hours at a replacement wage, makes it visible and understates what it is. The recipient's wellbeing is not the value of the services supplied, and a recipient who cannot contribute is not thereby of lesser worth. A framework that values persons solely through productive contribution has no place for dependency and cannot say that the arrangement above, which burns out the caregiver in two months, is a failure. The viability formalism in this chapter records the caregiver's constraint and the gap to be covered; it does not say how the gap should be covered or by whom, which is a question of distribution and of the institutions of the previous part.

What has been established. A caregiver has a viability kernel that limits sustainable hours, the example reaches breakdown in two months at 150 hours, and the gap of 30 hours per month is a cost of 750 that no account currently records. The decision now enabled is to plan care against the caregiver's constraint as well as the recipient's need. Assumptions still required are the exhaustion dynamics, which are empirical and individual. The question passed on to Chapter 38 is how ownership of the infrastructure used in production, including care, shapes who controls it.

Exercise 37.2. Compute the breakdown time with starting exhaustion 20 and $H = 140$.

Exercise 37.3. Modify the dynamics so that a day of rest reduces exhaustion by 5 and find the sustainable work schedule at 150 hours of need.

Exercise 37.4. Describe a care arrangement and list the services it supplies that no account records and the constraints on the caregiver.

Chapter 38

Local AI and the Ownership of Productive Infrastructure

A small property office uses a language model for correspondence and record keeping. It can rent the service from a cloud supplier or run a model locally on its own equipment. The decision is which, and the difficulty is that the comparison involves a purchase price, a running cost, a maintenance burden, a dependence on a supplier whose terms can change, and a distribution of control over the data and the tool. The case is used here as a sustained example of the framework, with stylized numbers that are assumptions of the example and not market data.

38.1 Total costs and dependence

Suppose a local installation costs 2400 to buy, amortized over 36 months, with electricity of 20 per month and maintenance work of four hours per month valued at 25 per hour, which is 100. The monthly cost is $2400/36 + 20 + 100 = 186.67$. A cloud service costs 150 per month in the first year, and the supplier raises the price by 30 per month each year, an assumption reflecting the switching-cost mechanism of Chapter 28. Over the first three years the cloud totals $12 \times (150 + 180 + 210) = 6480$ and the local installation totals $2400 + 36 \times 120 = 6720$, so the cloud is cheaper by 240. Cumulative costs cross in month 38 at 6960 each, and beyond that the local installation is cheaper. The comparison is sensitive to the assumed growth of the supplier's price and to the hardware lifetime, both of which are stated as assumptions. The calculation is reproduced in the companion (check BOOK, Chapter 38).

The cost comparison omits the quantities that Chapter 7 showed matter. The cloud service is an exposed complement: if the supplier withdraws or changes terms, the office's capability falls to zero until a replacement is found, at a switching cost that grows with incumbency. The local installation is a controllable complement but it carries a maintenance burden that falls on whoever has the skill, and that skill erodes if it is not practised

(Chapter 23). Persistent agents that act on the office's behalf add a further dependence on the interface between the agent and the records, and open software reduces the exposure to a single supplier at the cost of requiring a maintainer.

38.2 Distribution of control

Ownership of productive infrastructure is a bundle of permissions in the sense of Chapter 27: the right to use, modify, move, and exclude. A rented service gives use without modification or portability; a local installation gives modification and portability at the cost of maintenance. Neither is better in general. The framework asks, for each option, which permissions are held, which capacities support them, and who bears the cost of the failure of a complement. The distribution of these across offices, households, and communities is an issue of the distribution of control over productive infrastructure and is not settled by the cost comparison.

What has been established. Under the stated assumptions the cloud is cheaper for 37 months and the local installation afterwards, supplier dependence and maintenance burden are costs that the total omits, and the distribution of control is a separate question. The decision now enabled is to compare options on cost, dependence, maintenance burden, portability, and control together. Assumptions still required are the price path, the hardware life, and the maintenance time. The question passed on to Chapter 39 is how capabilities of this kind are established in communities that do not hold them.

Exercise 38.1. Recompute the break-even month if the supplier raises its price by 15 per month each year.

Exercise 38.2. Add an exit cost of 600 to the cloud option and a hardware failure cost of 400 at month 24 to the local option and recompute the three-year totals.

Exercise 38.3. Select a technology stack and describe which parts are exposed complements and which are controllable.

Chapter 39

Development as the Establishment of Capabilities

A village with a budget of 800 is deciding between a road to the market town and a school. Households benefit from different things, and the road reaches only those within a short distance of its route. The decision is how to spend, and the difficulty is that the outputs the two options produce are different kinds of capability, and that the village's authority over the conditions of its life may be reduced by the way the project is financed.

39.1 Establishing domains and who can reach them

P02 says that expenditure on access, infrastructure, learning, or recognition can change the choice problem itself. Take five households, with distances to the new road of 1, 1, 2, 4, and 6 kilometres. The road costs 600 and establishes market access for households within 2 kilometres, which are the first three. A feeder track costing 200 extends access to a household at 4 kilometres; a second feeder costing 300 extends to 6. With a budget of 800 the village can build the road and one feeder, 800 for the road and the first feeder, and the fourth household is reached; the fifth is not. The domain of choice has been enlarged for four households of five, and the village then chooses within the domains established. The distribution of access to the newly established opportunity is a variable of the model, and in this example it is unequal in a way that follows geography. The calculation is reproduced in the companion (check BOOK, Chapter 39).

The budgeted inclusion of P02 applies: outcomes under establish-then-select contain those under selection only, strictly whenever an affordable establishment makes an otherwise inaccessible outcome reachable. Here the road makes market selling reachable for three households with 600, for four with 800, and for all five only with 1100.

39.2 Output and authority

Development raises output and can lower communities' authority over their own conditions, for instance when financing carries conditions or when the project is owned elsewhere. Consider two plans: an externally owned road that raises monthly output by 300 and gives the community an authority index of 0.5, and a community-managed track that raises output by 200 and keeps the index at 1.0. In the profile of Chapter 16 the two are incomparable: the first is larger in output and smaller in authority. A choice between them is a choice of weights, and the choice belongs to the community. The framework requires that the authority coordinate be recorded and that the weights be stated by those who bear the effects. It also requires the corrective freedom of Chapter 15: a project that cannot be reversed or revised reduces the community's future authority even if the present index is high.

What has been established. Establishment of access changes who can choose, the budget of 800 reaches four of five households, and the output-authority tradeoff is an incomparability that the community must resolve. The decision now enabled is to report the distribution of access and the authority coordinate for each project. Assumptions still required are the cost and reach of establishment actions and an agreed measure of authority. The question passed on to Chapter 40 is how a destination of this kind is reached by a path that can be financed and revised.

Exercise 39.1. Find the budget at which all five households are reached and the budget at which the first four are reached with a school of cost 250.

Exercise 39.2. Prove the budgeted inclusion of P02 for this village, assuming a zero-cost null action.

Exercise 39.3. Select a development project and describe who decided the weights between output and community authority.

Chapter 40

Crisis, Transition, and Institutional Repair

The owner wants to retrofit the building to cut running costs by 100 per unit per month. A retrofit costs 2000 per unit and takes the unit out of letting for a month. The owner has 2500 in cash, a net monthly income of 1700 from the building after maintenance and the mortgage, and no other resources. The destination, a fully retrofitted building, is clearly better. The decision is how to get there, and the difficulty is that a viable destination does not make every path to it viable.

40.1 A viable destination and a viable path

Doing all four units in the first month costs 8000 and takes all four out of letting for the month. With no rent and no running costs the owner's net for the month is $-400 - 700 = -1100$, since maintenance and the mortgage continue, so cash after the month is $2500 - 1100 - 8000 = -6600$. The path leaves the acceptable set, nonnegative cash, though the destination is good. Doing one unit at a time, at months 0, 3, 6, and 9, gives in month 0 revenue of 3000 from the three units let, running costs of 900, maintenance of 400, and the mortgage of 700, a net of 1000, and after the outlay of 2000 the cash is $2500 + 1000 - 2000 = 1500$. From month 1 the retrofitted unit saves 100 and all four are let, the monthly net is 1800, and cash rises to 3300 and then 5100 before the next outlay. The minimum cash along the sequential path is 1500, at the end of month 0. The sequential path is viable and the simultaneous one is not. The companion recomputes the whole path (check BOOK, Chapter 40).

The tolerance of the path to cost overruns is the largest outlay at month 0 for which cash stays nonnegative: $2500 + 1000 - x \geq 0$ gives $x \leq 3500$, an overrun of 1500, or 75%, on the first unit. The tolerance is the quantity that matters for revising a plan that is going wrong: a plan with a tolerance of 75% can absorb a bad estimate on the first step, and one with a tolerance of 5% cannot.

40.2 Waiting for a subsidy

Suppose a subsidy of 800 per unit may be announced next month with probability one half. A unit's retrofit has present value $\mu = 2400$ in savings and costs $I = 2000$, so acting now gives $V_{\text{act}} = (2400 - 2000) + 100 = 500$ with first-month savings of 100. Waiting gives $V_{\text{wait}} = \beta\,\mathbb{E}[(2400 - I')^{+}]$ with $I' = 1200$ if the subsidy is announced and 2000 otherwise, so $\mathbb{E}[(2400 - I')^{+}] = \frac{1}{2} \cdot 1200 + \frac{1}{2} \cdot 400 = 800$, and with $\beta = 0.9$ and no expiry $V_{\text{wait}} = 720 > 500$. If the programme might close, survival with probability q independent of the announcement gives waiting better exactly when $q > 500/720 = 25/36$. This is the two-period structure of P26, applied to a transition with an explicit information date; its value is that it makes the owner's reasons for sequencing explicit and the decision to wait revisable at the date of the announcement.

40.3 Mechanisms for revising a plan

A transition plan should state its milestones, the tolerance of the path to each, the temporary protections it requires, and the rule by which a missed milestone triggers revision: a restoration duty with a rank, a reserve of funds, and a reopening procedure of the kind in Chapter 30. The housing example gives all of these in miniature. At scale, transitions involving energy, public finance, and industrial change have the same logic and larger consequences, and the failure the framework identifies most often is a plan with a viable destination and a path that leaves the acceptable set.

What has been established. A viable destination does not make a path viable, the simultaneous retrofit leaves the acceptable set and the sequential one does not, the tolerance of the path to overruns can be calculated, and waiting for information is worth it when the survival probability exceeds 25/36. The decision now enabled is to require a viable path and a revision rule in every transition plan. Assumptions still required are the cash model, the cost estimates, and the information date. This closes the economic text; the next part states the results precisely.

Exercise 40.1. Compute the minimum cash along the sequential path when the retrofit slips so that each unit is offline for two months.

Exercise 40.2. Show that the tolerance to overrun is nonincreasing in the opening outlay and give the maximal first outlay.

Exercise 40.3. Select a public transition programme and state its destination, its path, its tolerance, and its revision rule.

Part IX

Mathematical Foundations and Economic Results

Part IX consolidates the formal machinery and states the principal results with hypotheses, organized by theorem family. It is the formal centre of the book.

HOUSING THREAD. The housing model is stated formally, its viability kernel is computed, and the finite-state companion reproduces the results and their counterexamples.

Part IX fixes the statements. Each result is given with its hypotheses, separated into finite, countable, and general-state cases where they differ, and each worked number from the earlier parts that depends on it is listed. The formal setting is the one of the notation chapter: a state space S, actions $A(s)$, disturbances $W(s, a)$, a transition T, and an acceptable set $G \subseteq S$. The housing unit of Chapter 6 is the running instance, with the companion reproducing every example, and the stable labels P01 to P29 of the proof programme are used throughout.

Chapter 41

Transition Systems and History Compression

41.1 Setting and the factorization lemma

A *transition system* is a tuple (S, A, W, T, G) with $T : S \times A \times W \to S$ defined where $a \in A(s)$ and $w \in W(s, a)$. A *history* is a finite alternating sequence of states, actions, and disturbances consistent with T. Given an obligation language Λ of acceptable full histories, the continuation set of a history is $L(h) = \{u : hu \in \Lambda\}$, and a *summary* is a function q on histories.

Theorem 41.1 (Factorization lemma, P11 and P01)**.** *Let $F : S \to Y$ be a representation and $P : S \to \{0, 1\}$ a property. Then $P = \bar{P} \circ F$ on $F(S)$ for some $\bar{P}$ if and only if $F(s) = F(s')$ implies $P(s) = P(s')$, and the coarsest representation through which P factors is the quotient of S by $\ker P = \{(s, s') : P(s) = P(s')\}$. For histories with the continuation property $L(h)$, a summary q is exact if and only if it refines the equivalence $h \sim h' \Leftrightarrow L(h) = L(h')$; if $L(he) = e^{-1}L(h)$ for every event e then $\sim$ is a right congruence and defines an update map on classes.*

Proof. If $P = \bar{P} \circ F$ and $F(s) = F(s')$ then $P(s) = \bar{P}(F(s)) = \bar{P}(F(s')) = P(s')$. Conversely if $F(s) = F(s')$ implies $P(s) = P(s')$, define $\bar{P}(y) = P(s)$ for any s with $F(s) = y$, which is well defined. Any representation through which P factors identifies only states that P does not separate, hence refines nothing finer than $\ker P$, and $\ker P$ itself factors it. The history statement is the case $P(h) = L(h)$ valued in sets of futures, and the congruence is the computation $L(he) = e^{-1}L(h) = e^{-1}L(h') = L(h'e)$. □

The hypotheses are minimal. The dynamic version needs the residual property, which fails if the rules change without being recorded as events. When the number of classes of $\sim$ is finite the minimal summary is a deterministic automaton, which is the Myhill–Nerode construction [17]; when it is infinite no finite exact summary exists. For the housing unit of Chapter 6, started at condition 60 with the spends $\{0, 200\}$ and the obligation that

the unit stay lettable, the automaton has the eleven conditions reachable from 60 and one dead state, a total of twelve states, and the twelve states are pairwise inequivalent, so nothing can be forgotten: every distinct condition is a distinct state. The syntactic monoid of the same language has 108 elements. The count is reproduced in the companion (check BOOK, Chapter 41).

41.2 What a representation forgets

A representation's loss is exactly its kernel. The tier grade of Chapter 21 identifies conditions with equal holding cost, and it forgets the number of months of lettability remaining; the two quantities are different properties and the grade factors through one and not the other. The accounting map of Chapter 20 identifies states with equal reported profit and forgets the backlog. A history-dependent constraint, such as the cumulative-excess condition for bounded backlog (Chapter 8), is a property of the whole history that no function of the present state captures, and its summary is the running supremum $\sup_{s \le t} \int_s^t (\rho_M - \rho_R)$, which is again a state once adjoined. The general principle is that a history-dependent constraint can always be made Markov by enlarging the state to the minimal exact summary, and that the enlargement is finite exactly when the summary is.

Theorem 41.2 (P02, establishment precedes selection, formal statement)**.** *Let D be an opportunity domain, $E(e, D) \supseteq D$ the domain established by an acquisition action e with cost $\kappa(e)$, and $C(D')$ the outcomes obtainable by selection inside D' with selection costs σ. If there is a zero-cost identity establishment e_0 with $E(e_0, D) = D$, then for every budget b the outcomes obtainable by establish-then-select policies within b contain those obtainable by selection only, and the inclusion is strict if and only if some e with $\kappa(e) \le b$ and some x with $\kappa(e) + \sigma(x) \le b$ satisfy $x \in C(E(e, D)) \setminus C(D)$.*

The proof is the one in Chapter 4. For the numbers there the threshold budget is 550.

Accounting as a functor. By P19, an accounting interpretation F from the path category to a monoid is a functor if $F(g \circ f) = F(g) + F(f)$ and $F(\mathrm{id}) = 0$. For the ledger of Chapter 20 this holds by the posting rules, and the companion checks it on random concatenations (check BOOK, Chapter 49). Functoriality is not sufficiency, by Theorem 41.1.

Exercise 41.3. Compute the Nerode classes of the unit automaton started at 60 with the single spend 200 and explain the result.

Exercise 41.4. Prove that a right congruence of finite index on the histories defines a deterministic automaton recognizing the continuation language.

Exercise 41.5. Give a history-dependent obligation whose minimal exact summary is infinite.

Chapter 42

Orders, Dominance, and Incomparability

42.1 Dominance orders

Let X be a set of continuation profiles, each a vector in $\mathbb{R}^m$ with coordinates oriented so that larger is better. Define $y \preceq x$ when $y_i \leq x_i$ for every i. This is the product order, a partial order, and it is the dominance relation of Chapter 16. Three different comparisons are often run together. *Inclusion* compares sets of options by containment, $A \subseteq B$. *Cardinality* compares them by size. *Scalar utility* compares by a real number. The sets $A = \{51, \ldots, 100\}$ and $B = \{40, \ldots, 90\}$ of conditions, with $|A| = 50$ and $|B| = 51$, are not comparable under inclusion and are ranked by cardinality in favour of B, a ranking that depends on how conditions are individuated and that no inclusion argument supports.

Theorem 42.1 (P13, scalarization and multi-utility)**.** *(a) Let $\preceq$ be a partial order on X with incomparable x, y. There is no $v : X \to \mathbb{R}$ such that $x' \preceq y' \Leftrightarrow v(x') \leq v(y')$ for all $x', y' \in X$. (b) If the biconditional is weakened to monotonicity, $x' \preceq y' \Rightarrow v(x') \leq v(y')$, scalar functions exist (for example $v = \sum_i x_i$). (c) For a finite partial order there is a family of d real functions $v_1, \ldots, v_d$, with d the order dimension of Dushnik and Miller [9], such that $x' \preceq y'$ if and only if $v_i(x') \leq v_i(y')$ for all i, and no family of fewer functions has this property.*

Proof. (a) For reals $v(x) \leq v(y)$ or $v(y) \leq v(x)$, so the biconditional would make x, y comparable. (b) The sum is strictly monotone for the product order. (c) A family of linear extensions whose intersection is the order, called a realizer, yields the functions by position, and conversely a family of functions with the property gives linear extensions by breaking ties arbitrarily, whose intersection is the order. The minimal size of a realizer is the dimension by definition. □

The preservation requirement that generates the impossibility in (a) is order reflection, the converse direction. A scalar can order-preserve without reflecting; what it cannot do

is declare incomparable profiles incomparable. In the finite case the dimension can be computed. The three plans of Chapter 16 form an antichain of three elements, whose dimension is two, with the realizer "A, B, C" and "C, B, A". The standard example S_3 of six elements $a_1, a_2, a_3, b_1, b_2, b_3$ with $a_i < b_j$ if and only if $i \neq j$ has dimension three, which the companion confirms by exhaustive search over linear extensions (check BOOK, Chapter 42). The infinite case, where multi-utility representations can require infinitely many functions or fail without additional topological hypotheses, is not settled here and is stated as a hypothesis to be fixed for each application.

Economic reading. Prices and utilities are not ruled out, and (b) says that useful monotone indicators exist. What is ruled out is exactness: no real-valued function can represent all the distinctions a partial order records.

Exercise 42.2. Find the order dimension of the poset of three plans in Chapter 16 with the added profile of the zero-spend plan.

Exercise 42.3. Prove that a finite order has dimension at most its width (Dilworth's theorem may be used) and compute both for S_3.

Exercise 42.4. Give an infinite poset with no finite multi-utility representation.

Chapter 43

Reachability, Robustness, and Recovery

43.1 The kernel and its algorithms

For a finite system, the iteration $K_0 = G$, $K_{n+1} = \Phi(K_n)$ with $\Phi(X) = \{s \in G : \exists a\ \forall w,\ T(s,a,w) \in X\}$ terminates at K in at most $|S|$ steps, with each step costing $O(|S|\,|A|\,|W|)$ operations. Correctness: the sequence is decreasing since Φ is monotone and $K_1 \subseteq G = K_0$, a decreasing sequence of subsets of a finite set is eventually constant, and a constant value $X = \Phi(X)$ is a fixed point containing every post-fixed point by induction, so it equals the greatest fixed point of P03. The witnessing action at $s \in K$ gives a state-feedback policy. In infinite systems the iteration need not terminate; the infinite chain of the companion is a system in which every K_n is nonempty and the greatest fixed point is empty, so termination claims carry the finiteness hypothesis.

The comparative statics of P23 and its corollaries P04 and P14 are as in Chapter 13. Corrective reachability P08 is the dual least fixed point Reach, and the safe-degradation and restoration results P06 and P07 are as stated. The depth theorem of Chapter 15 shows that a review of depth n certifies K_n, with $K_n \supsetneq K$ for n below the stabilization step; for the unit under the ceiling 100 the stabilization step is 34.

43.2 Worst-case, almost-sure, and risk-budget guarantees

The robust kernel treats every allowed disturbance. It is monotone decreasing in the disturbance set, which has a practical consequence for model uncertainty: a modeller who assumes a smaller disturbance set than the true one overstates the kernel. In the unit under the ceiling 400, ignoring the one-tenth shock gives a kernel of every lettable condition and including it gives conditions from 51 upward. The almost-sure kernel in the countable case, P25, coincides with the robust kernel over the support, and in the example the unit at 50 is not almost-surely viable. Its failure probability within T months, for a shock of

probability 1/10 per month, is $P(\mathrm{Bin}(T, 1/10) \geq 11)$, which is 1.09×10^{-10} at $T = 12$, about 7.2×10^{-6} at 24, about 0.034 at 60, about 0.66 at 120, and 0.9993 at 240 (check BOOK, Chapter 43). A risk-constrained notion therefore needs a stated horizon and a budget that is spent as the horizon runs, which is the augmented-state formulation of P25; a fixed per-period risk budget does not suffice because failure probability accumulates. The general-state case, with measurability and selection hypotheses, remains a research specification.

43.3 The waiting theorem

Theorem 43.1 (Finite-horizon waiting, general form of P26)**.** *Let a finite Markov chain with states s, stopping payoffs $g(s)$, transition probabilities $p(s,t)$, and discount factor $\beta \in (0,1]$ be given, with absorbing states of payoff zero representing expiry or a declined opportunity. The value of the optimal stopping problem with horizon N is the function V_N defined by $V_0 = g$ and $V_{n+1}(s) = \max\{g(s),\ \beta \sum_t p(s,t) V_n(t)\}$, and acting at the first date at which $g(s) \geq \beta \sum_t p(s,t) V_n(t)$ is optimal. For $N = 1$ with an undecided state u that moves to a high or low state, the formula reduces to the two-period comparison of Chapter 9.*

The proof is backward induction. With $g(u) = 600$, high and low states paying 2300 and 0, a survival probability of 4/5, and $\beta = 0.9$, the value is $\max\{600, 0.9 \cdot \frac{4}{5} \cdot \frac{1}{2} \cdot 2300\} = 828$, and with survival probability one it is 1035, which the companion reproduces (check BOOK, Chapter 43). The continuous-time and infinite-horizon versions need the usual regularity conditions and are not proved here.

Exercise 43.2. Compute K_n for $n = 1, 2, 3$ for the unit under the ceiling 100 and the smallest n at which K_n excludes condition 70.

Exercise 43.3. Prove that the robust kernel is nonincreasing in the disturbance set.

Exercise 43.4. Extend Theorem 43.1 to two information stages and compute the value for the investment of Chapter 9.

Chapter 44

Composition and Systemic Failure

44.1 The obstruction and the positive results

The static obstruction is P09: for divisible simultaneous demands d_j on capacity c, $\sum_j d_j \leq c$ is necessary and sufficient for joint feasibility, and $d_j \leq c$ for each j does not imply it. The dynamic obstruction is the pair of units at $(50, 50)$ under the shared ceiling, each in its own kernel and the pair outside the joint kernel. The positive results are the product theorem of Chapter 31, the scarcity-price decentralization of Chapter 17 under separability, and the circular assume-guarantee composition of Chapter 32 under a lag.

Proposition 44.1 (Share partitions give a sufficient condition)**.** *Let two systems share a resource of capacity c with ceilings c_1 and c_2 satisfying $c_1 + c_2 \leq c$, and let $K_i(c_i)$ be the kernel of system i when its action set is restricted to draws at most c_i. Then $K_1(c_1) \times K_2(c_2) \subseteq K_{12}(c)$.*

Proof. A pair of policies, each respecting its own ceiling, respects the shared one, and each keeps its component in G_i against all disturbances; the pair therefore keeps the product in $G_1 \times G_2$. □

In the housing pair, the union over all splits of the ceiling of 400 into multiples of 100 of the products $K_1(c_1) \times K_2(c_2)$ equals the joint kernel exactly, all 2500 states, which the companion verifies (check BOOK, Chapter 44). The equality is a fact about this model, in which a unit's viability depends on a constant holding spend, and is not a general theorem: when dynamic reallocation across time can help, the union of static splits can be strictly smaller.

44.2 Accounting consistency and operational progress

A transactional outbox records the obligation to send an event in the same transaction as the order, which prevents the omission of an obligation for a committed order. It does not by itself guarantee eventual delivery or prevent duplicate effects at the recipient.

Enumerating the delivery schedules of one order with a lossy channel and at most two attempts shows the possibilities: without retry the order can be committed with the obligation recorded and the effect never delivered; with retry and no idempotence at the recipient the effect can occur twice; with retry and idempotence the effect occurs at most once, but all attempts can still be lost. Eventual delivery requires a fairness assumption on the channel. The separation is the same as that between P20's preservation and progress: the invariant that every committed order has a recorded obligation is preserved, and progress needs environmental assumptions.

Exercise 44.2. For three units at $(60, 60, 60)$ and a ceiling of 400, list the splits of the ceiling that give a viable triple.

Exercise 44.3. Find a system in which the union of static splits is strictly smaller than the joint kernel.

Exercise 44.4. State the fairness assumption under which the outbox with retry delivers eventually and prove delivery under it.

Chapter 45

Measurement and Empirical Identification

45.1 From quantities to indicators

Each theoretical quantity in the book needs an observable counterpart. The maintenance backlog needs an inspection standard and a restoration price. The hold-cost ratio needs condition records. The switching cost S of Chapter 28 can be read from quotes through $S = Nu(q^* - q_0)$ at the point where a customer is indifferent. Recovery time is the rank of Chapter 14, observed from completed repairs. Exit capacity is the comparison of liquid resources with the exit cost. Obligation fulfillment is the record of discharge events of Chapter 20. Service interruptions are observed directly, and the retained competence of Chapter 23 is the hardest to observe and is inferred from response times in drills.

45.2 Which causal claims need which designs

The housing model supplies the point in a form that can be computed. Suppose the owner follows the rule of spending exactly the holding cost on each unit. Across units from 40 to 100, spend is negatively related to condition, since worse units cost more to hold, and the condition change from one month to the next is zero for every unit, so spend and change have no variation to correlate. A regression of the change in condition on spend would find nothing, though the causal effect of an extra 100 of spend on next month's condition is exactly one tenth. The observational association is produced by the rule, not by the technology. Three designs identify the effect. A randomized comparison, adding 100 to a random subset, gives a difference of exactly one tenth. A natural experiment, such as an exogenous change in the crew ceiling, identifies the effect on units whose spend it constrains. Longitudinal data identify it only with a structural assumption about the decision rule. The companion reproduces the calculation (check BOOK, Chapter 45).

The general lesson is that a claim about what a maintenance policy does cannot be

read from the correlation between spending and condition in data generated by managers who respond to condition. Claims about switching costs and exit capacity have the same structure and need either a source of variation not chosen by the unit studied or an explicit structural model.

Quantity	Observable	Identification
Backlog	inspection against a standard	census, no causal claim
Switching cost S	quotes at indifference	randomized offers or exogenous supplier exit
Recovery time (rank)	completed repair durations	longitudinal, with policy rule
Exit capacity	resources against exit cost	survey plus exogenous cost shocks
Competence	drill response times	randomized practice

What has been established in Part IX. The representation, order, viability, composition, and waiting results have been stated with hypotheses and the finite cases proved; the infinite order dimension, the general-state almost-sure kernel, and the continuous-time waiting theorem remain specifications. The next three parts reinterpret the same objects.

Exercise 45.1. Design a randomized comparison to identify the effect of maintenance spending on condition when owners choose spending.

Exercise 45.2. Show that under the hold-cost rule the covariance of spend and condition change is zero for any population of units.

Exercise 45.3. Select a claim in the book that rests on an assumed rate (erosion, deterioration, growth of switching cost) and describe the study that would estimate it.

Part X

Categorical Interpretations

Part X interprets economic processes as morphisms and states what the interpretation preserves. Its stated payoff is that parallel composition fails to be independent under shared sinks, and the structure needed to repair that is identified.

Housing thread. Repair processes become morphisms between building interfaces, and accounting becomes a functor that forgets the backlog.

Part X gives the structures a categorical form where the form does something. The test applied to each construction is whether it states a theorem that the unstructured version does not, and where it does not, the chapter says so. Every example is the housing unit or the four-unit building, and the corresponding companion entries are named.

Chapter 46

Economic Processes as Morphisms

46.1 The syntactic category

Let Σ be a finite alphabet of events, for the housing unit the spends $\{0, 200\}$ of a month, and let $\Lambda \subseteq \Sigma^*$ be the language of admissible histories from a given initial state, for the unit started at 60 the words that keep the unit lettable throughout. The *syntactic congruence* $\equiv_\Lambda$ identifies u, v when $xuy \in \Lambda \Leftrightarrow xvy \in \Lambda$ for all contexts x, y. The quotient $\Sigma^*/\equiv_\Lambda$ is the syntactic monoid, and when it is regarded as a one-object category its morphisms are the economically distinguishable processes.

Proposition 46.1 (Syntactic category, P18)**.** *The quotient $M(\Lambda) = \Sigma^*/\equiv_\Lambda$ is a monoid under concatenation, the projection $\Sigma^* \to M(\Lambda)$ is a monoid homomorphism, and Λ is a union of congruence classes. It is the smallest quotient with these properties, and it is finite if and only if Λ is regular.*

Proof. If $u \equiv u'$ and $v \equiv v'$ then $xuvy \in \Lambda \Leftrightarrow xu'vy \in \Lambda$ by the first equivalence in context (x, vy), and $xu'vy \in \Lambda \Leftrightarrow xu'v'y \in \Lambda$ by the second in context (xu', y), so the operation is well defined and associative because concatenation is. Λ is saturated, since taking $x = y = \varepsilon$ gives $u \in \Lambda \Leftrightarrow u' \in \Lambda$. Any congruence saturating Λ refines $\equiv_\Lambda$, so the latter is the coarsest. Finiteness of the quotient characterizes regular languages by the Myhill–Nerode theorem in its two-sided form [17]. $\square$

For the unit started at 60 with spends $\{0, 200\}$ the monoid has 108 elements, and membership in Λ depends only on the class, which the companion confirms by comparing every word up to length seven (check BOOK, Chapter 46). The set Λ of admissible histories is *not* closed under composition: the word of seven zero spends is admissible and its square of fourteen is not, so admissibility is a property of the whole word and not a property inherited from the factors. An economic process is accordingly a morphism of the syntactic category together with the admissibility of the contexts in which it is used, not a free-standing map.

46.2 What the construction adds

The construction replaces the informal phrase “this process can be followed by that one” by a decidable relation on classes, which in the finite case is the transition monoid of the minimal automaton. It does not add economic content beyond what the automaton already records, and a claim that a categorical structure *explains* the viability of the unit would exceed what has been shown. What is gained is a place to state composition theorems: Chapters 47 and 50 state them over the same objects.

Exercise 46.2. Compute the syntactic monoid of the language “the unit never falls below 50” for the unit started at 60 with the single spend 200.

Exercise 46.3. Show that Λ above is not closed under composition by exhibiting the factors and verifying the failure by hand.

Exercise 46.4. Prove that the syntactic monoid of a regular language is isomorphic to the transition monoid of its minimal automaton.

Chapter 47

Parallel Production and Monoidal Structure

47.1 Independence fails under a shared sink

Two processes performed side by side are represented by a pair of words (u, v) run in lockstep. Independence would say that the pair is admissible whenever each word is, so that the monoidal product of two admissible processes is admissible. This is the interchange condition: performing u and v together should be equivalent to performing them in either order with the same effect.

Proposition 47.1 (Interchange fails under a shared resource)**.** *Let two units at condition* 50 *share a crew ceiling of* 400*. There are words w_1, w_2 such that each is admissible for its unit when the unit has the whole ceiling, and the pair run in parallel violates the ceiling in every month; the product of the two admissible processes is therefore not defined, and the parallel product is a partial operation defined only on draw-compatible pairs.*

Proof. A unit at 50 deteriorates by four tenths a month, so holding it costs 400 per month, the whole ceiling; the word of six monthly spends of 400 is admissible for the unit when it alone draws on the crew. Run in parallel, the two words draw 800 in each month against a ceiling of 400. The companion exhibits the words and checks both claims (check BOOK, Chapter 47). □

The repair is to index the products by the *draw profile*: a process carries the resource it draws month by month, and the parallel product (u, v) is defined exactly when the pointwise sum of the profiles stays within the capacity. This is a monoidal structure on a category in which composition and product are partial, and it recovers the unrestricted product when the resource is not scarce. In the example, the pair of words with spend 200 every month for six months is defined, since it draws exactly 400 (check BOOK, Chapter 47); it is the draw profile, not admissibility of the factors, that decides. The

connection with Chapter 31 is that independence, the hypothesis of the product theorem, is the statement that the profiles never conflict.

Exercise 47.2. Describe the monoidal structure on draw profiles for three units sharing a ceiling of 400, and say when the unit object exists.

Exercise 47.3. Show that the interchange law holds for partially defined products whenever both sides are defined.

Exercise 47.4. Give an economic example in which the shared sink is an ecological stock and state the profile.

Chapter 48

Resources, Obligations, and Enriched Categories

48.1 Costs that compose

Restoring a unit from condition 38 to 90 can be done by different plans, each with a time and a money cost. Let a morphism from condition a to condition b carry the pair $(\text{months}, \text{cost})$ and compose by addition of both. The order on the pair is dominance in both coordinates: smaller time and smaller cost are better. This is an enrichment over the quantale of costs ordered by product dominance, and the enriched hom is a set of Pareto-optimal pairs rather than a single number.

For the restoration of Chapter 14 the options from 38 to 90 form a set of four Pareto-optimal pairs, one for each ceiling $500, 600, 800, 1200$ of the plan. Composition is checked: the Pareto set of options from 38 to 60 composed with that from 60 to 90 gives exactly the options from 38 to 90 (check BOOK, Chapter 48), so intermediate stops cost nothing in expressive power. A single scalar distance suppresses the distinction: with a weight b per month of delay added to money, the weights $0, 10, 20, 30$ select four different plans, those with ceilings $1200, 800, 600, 500$ respectively. A fixed scalar therefore fixes a preference between delay and money, which the Pareto structure leaves open.

Proposition 48.1 (Tropical scalarization)**.** *For every weight $b \geq 0$ the map sending an option to $\text{cost} + b \cdot \text{months}$ is a monotone functor from the Pareto enrichment to the nonnegative reals with addition, and its minimizer lies in the Pareto set. The converse, that every Pareto option is a minimizer for some weight, holds for the convex lower envelope only; in general some Pareto options are never selected by a linear weight.*

The proof is immediate from the definitions together with the supporting-line characterization of the convex envelope. The relation to Chapter 17's scarcity price is that the weight is the shadow price of time.

Exercise 48.2. Identify which of the four options in the example lie on the convex enve-

lope.

Exercise 48.3. Prove that the composite of two Pareto sets, reduced to its Pareto subset, is associative.

Exercise 48.4. Add a risk coordinate to the enrichment and show where the order dimension of Chapter 42 enters.

Chapter 49

Accounting as a Structure-Preserving Map

49.1 The ledger functor

A ledger assigns to each monthly allocation of spend across the four units of the building $(90, 70, 55, 50)$ the profit it books, and to a history of allocations the sum of the monthly profits. The assignment is a functor from the free monoid of allocations, regarded as a one-object category, to the integers under addition, because the booked profit of a concatenation is the sum of the booked profits, provided the second history is run from the state at which the first one ended. The companion checks this on 200 random concatenations (check BOOK, Chapter 49).

Proposition 49.1 (Accounting insufficiency)**.** *A functor F from histories to a monoid is a sufficient statistic for a property P of end states only if $F(h) = F(h')$ implies $P(h) = P(h')$. The ledger above fails this for the property "restoration cost to the initial condition".*

Proof. The first claim is Theorem 41.1. For the second, the allocation $(0, 0, 0, 400)$ held for twelve months and the allocation $(0, 0, 200, 200)$ held for twelve months both book 24600 over the year. The first ends in the conditions $(76, 42, 13, 50)$ and the second in $(76, 42, 55, 26)$, with restoration costs 8400 and 6600 to bring the units back to $(90, 70, 55, 50)$. The companion recomputes the two histories and both costs (check BOOK, Chapter 49). □

The proposition is the precise form of the claim in Chapter 20 that profit does not record the backlog: the information lost is a function of the end state that the functor does not see. The functor is faithful on no nontrivial quotient of histories that includes the backlog, so what is needed is a second functor, to the condition vector, and the pair of functors is the accounting representation sufficient for the task.

Exercise 49.2. Find another pair of allocations with equal twelve-month profit and different restoration costs.

Exercise 49.3. Show that adding the end-state condition to the ledger gives a sufficient representation for restoration cost.

Exercise 49.4. Prove that a functor to a group cannot distinguish a history from its inverse composed with it.

Chapter 50

Open Systems, Institutions, and Compositional Guarantees

50.1 Interfaces and the assume-guarantee form

An open system is a transition system with an interface of inputs it assumes and outputs it guarantees. For the building, each unit assumes a draw c_i from the crew and guarantees a condition floor; the crew assumes that the sum of draws is within the ceiling and guarantees the draws it supplies. The circular composition rule of Chapter 32 says that if each component's guarantee holds provided the others' guarantees held up to the previous step, then all guarantees hold at every step. This is proved by induction on time and remains an induction, not a categorical identity.

Proposition 50.1 (Assume-guarantee induction, P27)**.** *Let components $C_1, \dots, C_n$ with guarantees G_i and assumptions A_i over a shared trace satisfy: G_i holds at step 0 without assumptions, and G_i holds at step $t+1$ whenever A_i holds up to step t; and the conjunction of the G_j up to step t entails A_i up to step t. Then every G_i holds at every step.*

Proof. By induction on t: the claim at step 0 is the first hypothesis, and if every G_j holds up to step t then every A_i holds up to step t by the third hypothesis, so every G_i holds at step $t+1$ by the second. □

Applied to four units at 90 supplied by the crew, the simulation runs for 24 months with every guarantee holding and the units still at 90, as the companion records. Applied to the building $(90, 70, 55, 50)$ the third hypothesis fails at the first month: the units' hold costs of $100, 200, 200, 400$ sum to 900, above the ceiling, and the induction does not start (check BOOK, Chapter 50). The categorical language organizes the proof obligations as the commutation of the interface diagram; the empirical premises, which are the costs and the ceiling, remain independent inputs and nothing in the diagram supplies them.

What has been established in Part X. The syntactic category, the partial monoidal structure on draw profiles, the Pareto enrichment, the ledger functor, and the assume-guarantee induction are stated, with the first, third, and fifth proved and the finite cases computed. No claim is made that the categorical form proves viability; it organizes where the hypotheses sit.

Exercise 50.2. Compute the first month at which the guarantee fails for the building under a ceiling of 400 when the units start at $(90, 90, 90, 50)$.

Exercise 50.3. State an assume-guarantee decomposition of the household budget of Chapter 3 and identify the circular dependency.

Exercise 50.4. Show that dropping the lag in the third hypothesis makes Proposition 50.1 false by constructing a two-component counterexample.

Part XI

Type Theoretic Interpretations

Part XI interprets states, permissions, and commitments as types. Its stated payoff is the precise gap between formal inhabitation and real-world executability, and the separate roles of preservation and progress.

HOUSING THREAD. Repair commitments are typed with deadlines and explicit failure branches, and degraded operating modes carry evidence.

Part XI reads the same objects as types. The aim is modest and specific: to show, with small systems whose rules are fully stated, where a type discipline prevents a class of error, where it cannot, and where well-typedness and executability come apart. The examples use the repair of the housing unit of Chapter 6.

Chapter 51

Economic States and Admissible Actions as Types

51.1 Actions indexed by state

Let State be the type of contexts $\Gamma = (m, c, \pi)$ of money m, crew capacity c, and a set π of permissions. The action family Act : State $\to$ Type assigns to each context the actions whose preconditions it satisfies. For the unit, Act(Γ) contains the repair with spend M at condition k exactly when repair $\in \pi$, $M \le m$, and $M \le c$. A repair with spend M is then a term of Act(Γ), and its typing judgment produces the successor context $(m - M,\, c - M,\, \pi)$ and the successor condition $\max(0, \min(100, k - d(k) + M/100))$. The companion implements the judgment: at $k = 51$, spend 300 from the context $(400, 400, \{\text{repair}\})$ is accepted, returns condition 52, and leaves the context $(100, 100, \{\text{repair}\})$; the same repair without the permission is rejected by the check, which raises the error before any effect (check BOOK, Chapter 51).

Proposition 51.1 (Preservation for the repair calculus, P20)**.** *If* $\Gamma \vdash \text{repair}(M) : \text{ok}$ *and the repair executes in a context that satisfies* Γ*, the resulting context satisfies the successor context given by the judgment, and no resource is negative.*

Proof. The judgment requires $M \le m$ and $M \le c$, so the successor amounts $m - M$ and $c - M$ are nonnegative; the permission set is unchanged. Execution applies the same subtraction, so the actual and the judged successor coincide. □

51.2 Inhabitation is not executability

Preservation says that typed steps stay typed. It does not say that a typed step can be taken. The type of the repair above assumes that the crew capacity recorded in the context is available when the work is to start. If the environment withdraws the crew, so that the available capacity is 0 while the context still records 400, the term still inhabits

Act(Γ) and the execution fails; the companion exhibits the stuck state (check BOOK, Chapter 51). The type system is correct about the model and silent about whether the model's context matches the world. Closing the gap is an empirical task of keeping the context synchronized with observations, and the next chapters show which parts of the gap can be shrunk by protocol design.

Exercise 51.2. Add a deadline to the context and state the typing rule for a repair that must complete by month t.

Exercise 51.3. Prove progress for the calculus under the assumption that the available resources are never below the recorded ones, and show by example why the assumption is needed.

Exercise 51.4. Give a dependent type for the action "evict" that makes the notice rule of Chapter 20 a typing requirement.

Chapter 52

Linear and Affine Resources

52.1 Which resources may be copied

A *linear* resource must be used exactly once, an *affine* resource at most once, and an unrestricted one any number of times. Money and crew-hours are consumed by use and are modelled as linear quantities that are split and not duplicated: the context $(400, 400)$ is split into the part $(300, 300)$ passed to a repair and the remainder $(100, 100)$ left behind, and the part consumed cannot be spent again. A permission, such as the right to repair, is affine: it can be exercised repeatedly in the model if it is not consumed, and it can be dropped, but a dropped permission cannot afterwards be used. Information is unrestricted, since copying a record does not deplete it, though its value to a competitor may fall, an effect the type system does not capture.

The companion implements a context with these three disciplines. Splitting 400 into 300 and 100 leaves 100 to the original owner; an attempt to pay 200 from the original after the split fails because the funds are held by the sub-context; and using a permission after dropping it raises an error (check BOOK, Chapter 52).

Proposition 52.1 (No duplication, no use after drop)**.** *In the split-and-pay calculus, the sum of the amounts held by all contexts derived from a root context never exceeds the root amount, and an affine permission is used at most once in the sense that no use occurs after a drop in any well-typed derivation.*

Proof. Splitting transfers amounts between contexts without creating any, and paying removes them, so the total is nonincreasing by induction on derivations. The affine rule marks a dropped permission as unavailable, and the typing of a use requires availability. □

No single discipline fits all economic objects. Obligations are neither linear nor affine in the strict sense: a debt cannot be silently dropped, which is the opposite of affine, and it can be discharged by one of several payments, which is neither. Treating obligations as linear resources that must be consumed by a discharge event captures the first property

and needs an explicit default branch for the second, which is the subject of the next chapter.

Exercise 52.2. Show that a debt as a linear resource does not model transfer to a third party without an additional rule, and supply the rule.

Exercise 52.3. Classify as linear, affine, or unrestricted: a lease, a licence to a patent, a recipe, a warranty, and a maintenance crew-hour.

Exercise 52.4. Prove that the total held across derived contexts is conserved when no payment occurs.

Chapter 53

Obligations, Deadlines, and Protocols

53.1 Escrow as a protocol with failure branches

A protocol is a finite labelled graph whose states are stages of an exchange and whose edges are the messages or events that advance it. The escrow protocol of the companion has the states *init*, *funded*, *delivered*, *disputed*, *settled*, and *refunded*. A deposit moves *init* to *funded*. From *funded* either delivery moves it to *delivered* or a timeout moves it to *refunded*. From *delivered*, acceptance settles and a dispute moves to *disputed*, from which the arbiter's resolution either pays, reaching *settled*, or refunds. The terminal states are *settled* and *refunded*.

Proposition 53.1 (Protocol progress, P20)**.** *A protocol has progress when every nonterminal state has an outgoing edge and every state can reach a terminal state. The escrow protocol has progress, and it loses progress when the outgoing edges of* delivered *are removed.*

Proof. For the escrow graph, reachability of a terminal state is read off: *funded* reaches *refunded* by timeout, *disputed* reaches *settled* or *refunded*, *delivered* reaches *settled*, and *init* reaches *funded*. If both outgoing edges of *delivered* are removed, the state is nonterminal and has no edge, so progress fails there and for every state whose only route passes through it. The companion computes both cases (check BOOK, Chapter 53). □

Progress is a property of the protocol graph. Whether the parties take the edges is a property of their incentives and behaviour: the buyer who never accepts or disputes keeps the protocol in *delivered* forever unless the protocol adds a deadline edge, which is why every state with a pending obligation needs a timeout branch. The deadline fixes a rule and is not a theorem about people. The relation to P28 is that the interval within which an exit remains available is itself a deadline edge, and the closed-form interval of Chapter 29 is the length of the edge.

Exercise 53.2. Add a timeout edge to *delivered* and state what is traded away.

Exercise 53.3. Show that a protocol with progress can still allow an infinite run, and give a condition that rules it out.

Exercise 53.4. Write the repair protocol of a tenant's notice of defect with the statutory periods of Chapter 20 as deadline edges.

Chapter 54

Proof-Carrying Commitments and Verified Degradation

54.1 A certificate that can be checked

A proof-carrying commitment ships with evidence that a third party can check without redoing the search. For restoration, the evidence is a map from conditions to a pair of a spend and a rank. The checker accepts when, for every certified condition, the spend is within the ceiling, the successor is certified and lettable, and the rank decreases by one at each step until rank zero, which must lie at or above the target. A certificate for restoring a unit to condition 60 under the ceiling 400, starting from 51, has fifty entries, with rank 5 at 51, and the checker accepts it. Changing the entry at 51 to rank 3 is rejected, and no certificate exists for the unit at 50 (check BOOK, Chapter 54).

Proposition 54.1 (Soundness of the rank checker, P07)**.** *If the checker accepts a certificate then from every certified condition the policy that follows the certified spends reaches a condition at or above the target in at most the certified rank steps, remaining lettable throughout.*

Proof. By induction on the rank. At rank zero the condition is at or above the target. At rank $r > 0$ the certified spend is feasible, and the successor is certified with rank $r - 1$ and lettable, so by the induction hypothesis the target is reached in at most $r - 1$ further steps. □

The checker is sound with respect to the model, which is a hypothesis: the certificate does not show that the repair price, the crew ceiling, or the deterioration schedule are those of the world. The proposition is the finite fragment of the safe-restoration result; completeness, that every restorable condition has a certificate, holds by construction in the finite model, since the rank function of Chapter 14 supplies one.

54.2 Replay integrity and historical authority

A record system can be reconstructed by replaying its log of operations under a canonical encoding, checking at each operation that its parent is the preceding one and its tick is valid, ignoring duplicates, and recording refusals. Success of the replay shows *integrity*: the log is internally consistent and reproduces the state. It does not show *authority*: that the operations were issued by parties entitled to issue them, or that the log is the log which the community accepts. A replayed state can be internally perfect and historically unauthorized, for example a complete forged log. The two properties are verified by different evidence, the first by recomputation and the second by signatures or institutional acceptance, and a claim of "verified history" must state which is meant. The distinction is used again in Chapter 55 and in the audit of Chapter 25. This is a candidate import from the author's repository review and the claim here is the conceptual one only; the implementation was not tested.

Exercise 54.2. Construct a certificate for restoring to 60 at the ceiling 600 from condition 45 and compare the ranks with those at the ceiling 400.

Exercise 54.3. Give a forged certificate that passes the structural checks but violates the ceiling, and say which check catches it.

Exercise 54.4. List the evidence needed to move a replayed record from integrity to authority.

Chapter 55

Refinement, Interpretation, and the Limits of Formal Assurance

55.1 A verified scheduler that fails its purpose

A scheduler allocates the crew ceiling of 400 among the four units. The refinement type of its output is $\{a \in \mathbb{Z}^4_{\geq 0} : \sum a_i \leq 400\}$. The greedy scheduler that funds units in increasing order of holding cost until the ceiling is spent has this type, and the companion tests it on 500 random buildings (check BOOK, Chapter 55), in addition to the one-line argument that it never allocates beyond the remaining capacity. On the building $(90, 70, 55, 50)$ it returns $(100, 200, 0, 0)$: it meets its specification and leaves the two worst units unfunded, with holding costs 200 and 400. The implementation is verified relative to the specification and the specification is the wrong one for the purpose, which was to keep the building viable. The holding costs total 900 against the ceiling 400, so no allocation holds all four units at their current conditions in that month, and the right specification would also say what to do when this is so: triage by the loss function of Chapter 17, or an explicit declaration that the target is unattainable.

The general point is the one the proof programme reserves under P15 and P17: a formal assurance has a premise that the specification expresses the purpose, and that premise is outside the formalism. Consistency of a model is not fairness of its allocation, a verified computation does not make its inputs true, and a check that a stock is positive does not show that the ecological stock it represents is adequate. A formal assurance statement should name its specification, its trusted inputs, and the purpose it does not address.

What has been established in Part XI. Preservation is proved for the repair calculus and soundness for the rank checker, both relative to the stated models; progress is a property of the protocol graph and is not a behavioural prediction; the gap between inhabitation and executability is exhibited, not closed.

Exercise 55.1. Write a specification of the scheduler under which the greedy rule is not

correct, and give a correct rule for it.

Exercise 55.2. Identify the trusted inputs in the rank certificate of Chapter 54.

Exercise 55.3. Describe a case in which a formally verified allocation is unfair by a stated criterion.

Part XII

Sheaf Theoretic Interpretations

Part XII interprets local knowledge and plans over overlapping contexts. Its stated payoff is the distinction between a presheaf of plans and a sheaf of feasible allocations, together with one genuine cohomological obstruction.

Housing thread. Landlord, city, and contractor each hold local records, and gluing succeeds or fails depending on whether the shared crew is represented.

Part XII reads the book's local-to-global questions through presheaves and sheaves, with one cohomological obstruction. As in Part X the discipline is that a structure appears only when it states something the unstructured version does not. The examples are the records of the four-unit building held by a landlord, a city, and a contractor, and a small exchange graph.

Chapter 56

Local Economic Knowledge and Overlapping Contexts

56.1 A base of contexts

Let $X = \{1, 2, 3, 4\}$ be the units of the building. A context is a subset $U \subseteq X$ of units a party can see. The landlord sees $\{1, 2, 3\}$, the city inspector $\{2, 3, 4\}$, and the contractor $\{3, 4\}$. The base is chosen for the economic reason that each party's records cover the units it is responsible for, and the overlaps are the units on which two parties' records can be compared. A different base, such as jurisdictions or fiscal years, gives a different structure and is a modelling choice, not a derived fact.

A *presheaf* F assigns to each context U a set $F(U)$ of descriptions over U and to each inclusion $V \subseteq U$ a restriction map $F(U) \to F(V)$, respecting identities and composition. A presheaf is a *sheaf* for a cover $\{U_i\}$ of U if locality holds (two descriptions over U that agree on every U_i are equal) and gluing holds (every family of descriptions on the U_i that agree on the overlaps comes from one description over U).

Proposition 56.1 (Unconstrained assignments form a sheaf, P21)**.** *For a set V of values, let $F(U)$ be the set of functions $U \to V$ and let restriction be restriction of functions. Then F satisfies locality and gluing for every cover of every U.*

Proof. If two functions on U agree on every member of a cover they agree at every point, since each point lies in some member. If a family of functions f_i on U_i agree on overlaps, define $f(x) = f_i(x)$ for any i with $x \in U_i$; this is well defined because two such i give the same value on $U_i \cap U_j$, and it restricts to each f_i. □

For the landlord's and inspector's cover of the units $\{1, 2, 3, 4\}$ by $\{1, 2, 3\}$ and $\{2, 3, 4\}$ with spends in $\{0, 200\}$, the companion enumerates all matching families and confirms that each has exactly one glued section (check BOOK, Chapter 56). The result is nearly trivial, and its use is as a baseline: the failures in the next chapters arise when the descriptions carry a constraint that is not local.

Exercise 56.2. Give a cover of the building for which the sheaf condition for assignments is trivial and say why that cover is uninformative.

Exercise 56.3. Show that locality can fail for a presheaf of reported profits that forgets the unit labels.

Exercise 56.4. Choose a base of jurisdictions for a river with three riparian owners and state the restriction maps.

Chapter 57

Presheaves of Plans, Accounts, and Commitments

57.1 Restriction and what it forgets

The records of the building restrict naturally. A record of spends by unit restricts to any subset of units by discarding the other entries. A record of the total spend does not: the total over $\{1, 2\}$ cannot be recovered from the total over $\{1, 2, 3\}$. Two histories with spends $(100, 300)$ and $(200, 200)$ on units 1 and 2 have the same total 400 and different restrictions to unit 1 (check BOOK, Chapter 57). An account that aggregates is a function of the record and not a presheaf on the same base unless restriction is defined for it, and the departmental consistency question, whether the accounts agree on overlaps, has no content for aggregates that cannot be restricted.

Obligations raise the same question. If the landlord's record of repair duties and the inspector's record of the same defect are compared on the overlap, the duty period of Chapter 20 is the shared datum and a discrepancy in it is a failure of agreement on an overlap. Whether the discrepancy is the inspector's error, the landlord's, or a difference of definition is not decided by the presheaf structure; it records that some party's description is not a restriction of a single global description.

A third failure is the omission of a shared dependency. If the landlord's plan and the contractor's plan each fund their own units and neither mentions the crew they both draw on, each local plan is internally consistent and the missing object is outside the base. The remedy in Chapter 58 is to add the crew to the model, not to strengthen the gluing.

Exercise 57.1. Define a restriction map for monthly profit that makes profit a presheaf and say what is lost.

Exercise 57.2. Give a presheaf of obligations in which two parties' descriptions agree on the overlap but cannot be glued.

Exercise 57.3. Show that the family of total spends over contexts does not form a

presheaf on the subsets of the building under any restriction maps that are compatible with the sums.

Chapter 58

Gluing and Global Feasibility

58.1 Matching plans that cannot be carried out

Let $F(U)$ be the set of spend assignments $U \to \{0, 200\}$ whose total is at most the crew ceiling 400. Restriction preserves feasibility, since a smaller domain has a smaller total, so F is a presheaf. It is not a sheaf. On the disjoint cover $\{1, 2\}, \{3, 4\}$ there are 16 pairs of local assignments (all four assignments on each half are locally feasible), and each matches trivially because there is no overlap; of these, 11 have total at most 400 over the four units and 5 do not, so five matching families have no glued section in $F(X)$. On the overlapping cover $\{1, 2, 3\}, \{3, 4\}$ there are 14 matching families, 11 of which glue and 3 of which do not (check BOOK, Chapter 58). The five and three are the proposed assignments that every local party can accept and that no one can carry out.

Proposition 58.1 (Gluing needs the scarce resource in the model)**.** *Let $F(U)$ be the assignments over U whose total draw is at most c. Then F is a presheaf. It satisfies the gluing condition for a cover $\{U_i\}$ of U if and only if every matching family has total draw at most c over U; this holds if some U_i contains U, and fails in general when the cover is proper and c is binding.*

Proof. Restriction of an assignment with nonnegative draws does not increase the total, so F is a presheaf. A matching family has a unique glued assignment on U by Proposition 56.1, and the family glues in F exactly when this assignment is feasible. If some U_i equals U the glued assignment is the member f_i, which is feasible. The example above is a proper cover on which a matching family has a total of 600. □

The honest repair is to add to the model an object representing the crew, for example a context that contains the total draw, so that the cover includes a member on which the constraint is checked. Two phenomena are worth separating from gluing. A local repair condition, such as every pair of neighbouring plans being adjusted until it is feasible, need not produce a global section, since local fixes can pass the violation around. And when exact matching is replaced by penalized reconstruction, minimizing a data-fit term plus

a penalty per mismatch on an overlap, a candidate with one mismatch can have lower energy than an exact one: with a data term of 10 for the exact candidate, and 3 plus a penalty of 2 for one mismatch for the imperfect candidate, the imperfect candidate has energy 5 and is selected. The numbers are stipulated to show the logic, not estimated from data, and the comparison in the companion confirms the arithmetic. The conclusion is that a reconstruction method that tolerates mismatches does not return a section of the sheaf, and its output is a recommendation.

Exercise 58.2. Compute the unglued families for the cover $\{1,2\},\{2,3\},\{3,4\}$ and the ceiling 400.

Exercise 58.3. Add a crew context to the base and restate gluing so that infeasible families are excluded by the cover.

Exercise 58.4. Construct an instance of local repairs that cycles without reaching a global section.

Chapter 59

Obstructions, Inconsistency, and Institutional Boundaries

59.1 An exchange graph

Take three currencies or goods A, B, C with multiplicative exchange rates, so that one unit of A buys r_{AB} of B, and so on around the cycle $A \to B \to C \to A$. If the rates are $1/40$, $5/2$, and 20, the product around the cycle is $\frac{1}{40} \cdot \frac{5}{2} \cdot 20 = \frac{5}{4}$, which is not 1. A trader who goes round the cycle ends with $5/4$ of the initial amount. The rates are consistent with a single price system p, in the sense $r_{XY} = p_Y / p_X$, if and only if every cycle product is 1.

Proposition 59.1 (Potential versus cycle product, P22)**.** *On a connected graph with positive exchange rates r_{xy} satisfying $r_{yx} = 1/r_{xy}$, there is a price system $p > 0$ with $r_{xy} = p_y / p_x$ for every edge if and only if the product of the rates around every cycle is 1. The map from cycles to cycle products is a homomorphism from the cycle space to the positive reals, and the rates are consistent if and only if its class in $H^1(\text{graph}; \mathbb{R}_{>0})$ vanishes. With a proportional fee f on each exchange, no arbitrage around a cycle of length n holds if the product times $(1-f)^n$ is at most 1.*

Proof. Necessity is the telescoping product $\prod p_{x_{i+1}} / p_{x_i} = 1$. For sufficiency fix a spanning tree and a root, and define p_x as the product of rates along the tree path from the root to x; each non-tree edge closes a cycle with the tree, and the cycle product being 1 gives $r_{xy} = p_y / p_x$ for it. The class statement is the definition of the first cohomology of the graph with these coefficients: a one-cochain is a choice of rates, a coboundary is a choice of the form p_y / p_x, and the cocycle condition is automatic for a graph. The fee statement is the multiplication of the cycle product by $(1-f)^n$. $\square$

In the example a fee of 8% gives $\frac{5}{4} \cdot 0.92^3 \approx 0.97 < 1$, so the cycle is no longer profitable, and a fee of 5% gives $\frac{5}{4} \cdot 0.95^3 \approx 1.07 > 1$, which still allows it (check BOOK, Chapter 59). Friction therefore moves the obstruction from an equality to an inequality.

59.2 A twisted cover

A second obstruction is of a different kind. Cover a circle by three arcs, each with a local section of a sign-valued variable, say the orientation in which a claim is read, and let the transition between neighbouring arcs be $+1$ or -1 for same or opposite orientation. If the transitions are $(+1, +1, -1)$ no choice of signs a_1, a_2, a_3 on the arcs gives $g_{ij} = a_i a_j$ on all three overlaps, because the product of the right sides around the circle is $a_1^2 a_2^2 a_3^2 = 1$ while the product of the transitions is -1. The class in $H^1(S^1; \mathbb{Z}/2)$ is nontrivial, there is no global orientation, and the nonexistence is a mathematical fact that does not depend on any of the parties' errors (check BOOK, Chapter 59). Cohomology is used here because the coefficient structure and the interpretation of the obstruction were both stated first. A claim that organizational disagreement is "cohomological" without such a structure is a metaphor and is not used in this book.

Exercise 59.2. Find the minimal proportional fee that eliminates the arbitrage in the example for a cycle of length 3 and verify the claim at the boundary.

Exercise 59.3. Show that the cycle product is multiplicative under concatenation of cycles through a common vertex.

Exercise 59.4. Give a twisted cover with four arcs and signs having a trivial class, and exhibit the trivialization.

Chapter 60

Descent, Federation, and the Organization of Collective Futures

60.1 Federated records

Independently governed bodies coordinate by comparing records on overlaps, not by absorbing one into another. The landlord holds $\{1 : 90, 2 : 70, 3 : 55\}$, the city holds $\{2 : 70, 3 : 55, 4 : 50\}$, and the contractor holds $\{3 : 55, 4 : 50\}$. The records agree on every overlap and glue to the building record $(90, 70, 55, 50)$, the same result in whatever order the contexts are merged. Suppose the city instead records 53 for unit 3. The merge now stops with a conflict at unit 3 and the federation reports it without choosing a winner (check BOOK, Chapter 60). The merge cannot say which record is right; descent supplies the compatibility condition and the construction of the global record when it holds, and an institutional procedure, such as inspection, supplies the resolution when it does not.

60.2 The three interpretations together

The categorical interpretation of Part X organizes processes and their composition, the type-theoretic interpretation of Part XI says which actions are admissible in a context and where admissibility departs from executability, and the sheaf-theoretic interpretation of this part says when local descriptions can be assembled. They answer three different questions about the same building. Whether the repair of unit 4 can follow the restoration of unit 3 is a question about composition and uses the syntactic category. Whether the repair is permitted and funded is a question for the typing judgment. Whether the landlord's, the city's, and the contractor's records of the repair agree is a question of gluing. None of the three answers whether the building is viable, which is the question of the kernel of Part IX under the actual resource constraints. The interpretations organize obligations of proof and records of evidence, and the economic content comes from the models and the data.

What has been established in Part XII. The sheaf property is proved for unconstrained assignments and the failure of gluing is exhibited for feasible allocations; the graph theorem is proved for connected exchange graphs, and the twisted cover is a worked example of a nontrivial class. The reconstruction results are stipulated examples; no empirical claim about institutions is made.

Exercise 60.1. Define a federation protocol for three bodies with a conflict-resolution rule and state which of the properties of the merge above it preserves.

Exercise 60.2. Show that the merge of agreeing records is independent of the order of contexts and give the general statement.

Exercise 60.3. Take a decision in the book and say which of the three questions, composition, admissibility, or compatibility, each of its premises belongs to.

Appendix A

Prerequisites

The appendices supply background so that Parts I to VIII need no outside reading beyond calculus and elementary probability.

A.1 Optimization

Convex programs, Lagrangian duality, Slater's condition, complementary slackness, and shadow prices (used for P24). Dynamic programming and the Bellman equation for the benchmark cases.

A.2 Dynamical Systems

Discrete-time systems, invariant sets, attractors, and thresholds. Continuous-time viability and Hamilton–Jacobi reachability [4, 16].

A.3 Probability

Transition kernels, almost-sure events, measurability and selection, and the facts about Borel and analytic sets that P25 requires [5].

A.4 Order Theory

Partial orders, lattices, the Knaster–Tarski theorem [21], order dimension, and multi-utility representations [9].

A.5 Category Theory

Categories, functors, monoidal categories, enrichment [13], resource theories [7], and compositional models [10, 14].

A.6 Dependent and Linear Types

Dependent types, linear and affine disciplines [11], session types, and refinement types.

A.7 Sheaves

Presheaves, sites, the sheaf condition, torsors, and Čech cohomology in degree one [15].

Appendix B

The Computational Companion

The directory `companion/` holds a finite-state implementation in Python using only the standard library. Its output is recorded in `companion/EXPECTED_OUTPUT.txt`. The checks below are *reported checks*: they have been run by the author of the skeleton, and the package and recorded output should be inspected before any of them is cited as verified. Run `python3 -m vcc.run all` from `companion/`, or `python3 -m unittest discover tests`.

Three kinds of check are distinguished. An *exhaustive* check enumerates every case of a small model and so establishes the claim for that model only. A *random* check samples systems from a recorded seed. It can find counterexamples and give implementation evidence, but it cannot prove a universal theorem. A *counterexample* check exhibits a specific system in which a hypothesis has been removed.

Result	Claim checked	Kind	Seed	Hypo
P03	iteration = Knaster–Tarski union	random, 300 systems	1	–
P03	no termination on an infinite chain	counterexample	–	finite
P05	local acceptability without viability	exhaustive, 3 states	–	–
P09	each demand fits, the sum does not	random, 500 draws	2	–
P10	$K_{12} = K_1 \times K_2$	random, 200 pairs	3	–
P10	$K_{12} \subsetneq K_1 \times K_2$	counterexample	–	indep
P23	monotone in G, actions, disturbances	random, 300 systems	4	–
P12	housing condition and deferral (H3)	exhaustive	–	–
P09	shared crew, two units (H4)	exhaustive	–	–
Chapters 1–60	62 worked calculations recomputed (BOOK)	exact arithmetic and exhaustive search	–	–
P22	potential vs. negative cycle	exhaustive, 3 nodes	–	frictio
P28	lost-exit interval in closed form	random, 400 parameter sets	5	–
P29	availability, timing, sufficiency	random, 600 cases	6	–

Random checks confirm closed forms that are proved in the text; they are not part of

any proof. The housing-case numbers in Theorems H1 and H2 are reproduced exactly by the last two rows.

The simulator

The file `simulator/housing_v2_simulator.html` is a single self-contained page that embeds the specification `spec/housing_v2.json` verbatim, lists the assumptions in force, runs constant, holding-cost, greedy, and kernel-feedback policies under the month order of the specification, computes one- and two-unit viability kernels, and re-runs eight of the companion's chapter calculations. The command `node simulator/test_simulator.js` confirms that the embedded specification equals the file, that those calculations pass, and that 300 random traces and fifteen kernels generated by the Python companion are reproduced exactly. The terminal styling is optional and does not affect any computation. The simulator implements the housing model of the book and none of the later parts' abstract constructions.

Appendix C

Framework Concordance

This appendix records which earlier text supplies each concept, how it has been modified for economic use, and which connections remain proposals. Provenance is preserved: where material is attributed to the author or stored in an imported collection, it is cited to its original source and not to the collection. All entries are provisional until the author confirms them.

Concept	Source text	Economic modification	Status
Domain establishment	The Domain and Its Establishment	Chapter 4, P02: establishment as an economic activity	proposal
Distinction and continuation	Distinction and Continuation	Chapter 26: classifications as allocators of futures	proposal
Verification without correspondence	Verification Without Correspondence	Chapter 25, P17: what a check establishes	proposal
Aspect relegation	Aspect Relegation Theory	Chapters 23, 24, P16: attention and return margin	proposal
Residuals and replaced things	Economy of Replaced Things	Chapter 10: disposition duties	proposal
Admissible degradation	source to be assigned	Chapter 12, P06	proposal
Continuous intervention	source to be assigned	Chapter 28: cumulative dependence	proposal
Semantic frameworks	source to be assigned	Chapter 21: task-relative adequacy	proposal

Candidate cases from documentation-level review

Two further reviews read project README files only, without inspecting implementations or running tests. A civic-maintenance verification design supplies a case for Chapters 20, 25, and 30. A commerce system with stock reservations and a transactional outbox supplies a resource-commitment case for P09, P20, and P27 (Chapters 31 and 44). An agent-automation tool bears on Chapter 24. Thematic overlap with the book's concerns establishes no influence or borrowing, and promotional claims such as zero risk or general Sybil resistance are excluded.

Candidate imports from a source review

A targeted review of material in the author's imported repositories identified four imports and several counterexamples. The review was not exhaustive, no files were copied, and implementation claims were not tested.

1. Replay integrity versus historical authority (Chapters 22, 25, 54).
2. The identity $D_{\mathrm{KL}}(P(\cdot \mid B) \,\|\, P) = -\log P(B)$ for restricted distributions (Chapter 16), and the distinction between expenditure and damage to regenerative capacity (Chapters 7, 23).
3. Allocation of finite maintenance capacity through explicit response functions (Chapters 17, 23).
4. Conditional platform revenue $R(q) = (\alpha q + \beta)M(q)$ (Chapter 33), restricted to the narrow revenue proposition.

Counterexamples and exercises: a ledger verifier that does not recompute stored roots or aggregate repeated postings (Chapter 25); a bounded-debt claim corrected by the cumulative-excess condition (Chapter 8); sheaf repair claims that do not yield global sections (Chapter 58). Claims about consciousness, universal entropy laws, cosmological foundations, or unverified performance are not imported.

Bibliography

[1] Martín Abadi and Leslie Lamport. Conjoining specifications. *ACM Transactions on Programming Languages and Systems*, 17(3):507–535, 1995.

[2] Kenneth J. Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.

[3] Kenneth J. Arrow and Anthony C. Fisher. Environmental preservation, uncertainty, and irreversibility. *Quarterly Journal of Economics*, 88(2):312–319, 1974.

[4] Jean-Pierre Aubin. *Viability Theory*. Birkhäuser, 1991.

[5] Dimitri P. Bertsekas and Steven E. Shreve. *Stochastic Optimal Control: The Discrete-Time Case*. Academic Press, 1978.

[6] Ronald H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.

[7] Bob Coecke, Tobias Fritz, and Robert W. Spekkens. A mathematical theory of resources. *Information and Computation*, 250:59–86, 2016.

[8] Avinash K. Dixit and Robert S. Pindyck. *Investment under Uncertainty*. Princeton University Press, 1994.

[9] Ben Dushnik and E. W. Miller. Partially ordered sets. *American Journal of Mathematics*, 63(3):600–610, 1941.

[10] Brendan Fong and David I. Spivak. *Seven Sketches in Compositionality*. Cambridge University Press, 2019.

[11] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–101, 1987.

[12] J. R. Hicks. *Value and Capital*. Clarendon Press, Oxford, 1939.

[13] F. William Lawvere. Metric spaces, generalized logic, and closed categories. *Rendiconti del Seminario Matematico e Fisico di Milano*, 43:135–166, 1973.

[14] Saunders Mac Lane. *Categories for the Working Mathematician*. Springer, 2 edition, 1998.

[15] Saunders Mac Lane and Ieke Moerdijk. *Sheaves in Geometry and Logic*. Springer, 1992.

[16] Ian M. Mitchell, Alexandre M. Bayen, and Claire J. Tomlin. A time-dependent Hamilton–Jacobi formulation of reachable sets for continuous dynamic games. *IEEE Transactions on Automatic Control*, 50(7):947–957, 2005.

[17] Anil Nerode. Linear automaton transformations. *Proceedings of the American Mathematical Society*, 9:541–544, 1958.

[18] Elinor Ostrom. *Governing the Commons*. Cambridge University Press, 1990.

[19] Amartya Sen. *Commodities and Capabilities*. North-Holland, 1985.

[20] Steven Shavell. The judgment proof problem. *International Review of Law and Economics*, 6(1):45–58, 1986.

[21] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.