

Business Models as Cash-Flow Structure
Formalization, a Catalogue of Types, and Survey Design

v0.28.9 7 September 2026

This text is under construction. The structure of the theory, the propositions, and the empirical conclusions may all change. Claims that have been withdrawn, and the reasoning that led to each withdrawal, are recorded in Appendix E. Version numbers correspond to the stages listed there.

This text formalizes a business model as a map Φ from a schedule of deliveries to a schedule of settlements, and derives from it the credit position κ , the self-financeable growth rate g^* , and a three-way decomposition of surplus. Part I builds the theory. Part II generates seven families and twenty-eight types from the degrees of freedom of Φ , then deduces the subspaces that remain once constraints on capital and labour are imposed, the limit in which the operator is removed, and the subspace obtained by replacing the supplier of the information structure. Part III sets out the five stages that connect the theory to observation, together with the treatment of selection bias, the pre-registration procedure, and a taxonomy of verification obstacles. Part IV reports the empirical work: families 2 and 5, the Financial Statements Statistics of Corporations by Industry, platform fees, and the records of a public distributed ledger; its opening chapter (Chapter 14) tabulates which theoretical quantities could be measured and which could not. Twenty-two implications derived from the theory were fixed in writing before the observations were made; six were supported, eleven rejected, and five could not be tested, and the record here includes the parts that were wrong. Because the empirical work is not complete, the deductions about Φ under constraints (Chapters 8, 9 and 10) are not presented as a separate part on applications but placed in Part II, which treats the space of Φ . Proofs appear in the body and in Appendix A, notation in Appendix B, and correspondences with existing theory in Appendix C.

Contents

Part I	Theory	9
Chapter 1	Introduction: Choosing the Time Structure of Cash as the Variable	10
1.1	The problem	10
1.2	Structure of the text, and how to read it	10
1.3	What this text does not claim	11
Chapter 2	The Business Model: Φ and the Generation of Cash Flow	13
2.1	Primitive sets	13
2.2	Schedules and spaces	13
2.3	The map from delivery to settlement	14
2.4	Generating the cash flow	14
2.5	Degrees of freedom of Φ	15
2.6	Three risk operations	17
2.7	Divergence between right and exercise	21
Chapter 3	The Credit Position: κ and the Lag τ	23
3.1	The gap between delivery and settlement	23
3.2	Decomposing the cash balance	28
Chapter 4	Information Structure and the Objective	31
4.1	Verifiable shared information	31
4.2	The objective: a cash constraint imposed path by path	32
4.3	Heterogeneity of discount factors	33
Chapter 5	The Growth Constraint: CCC and the Self-Financeable Growth Rate	35
5.1	The cash conversion cycle	35
5.2	The fundamental inequality	35
5.3	Growth and free cash flow	36
5.4	Level and change are separate constraints	37
5.5	CCC is a sum of lags	39
5.6	The path constraint and the peak	40

5.7	Comparative statics	41
Chapter 6	Surplus and Layers: Source, Durability, Mutability	43
6.1	The three-way decomposition of surplus	43
6.2	Surplus and cash	44
6.3	Layers: five levels ordered by mutability	46
Part II	The Space of Maps	48
Chapter 7	A Catalogue of Φ : Seven Families, Twenty-Eight Types	49
7.1	How the catalogue is built	49
7.2	Family 1: one-off exchange	49
7.3	Family 2: continuing access	50
7.4	Family 3: renting an asset over time	50
7.5	Family 4: recovery through a complement	50
7.6	Family 5: outcome- and state-contingent	51
7.7	Family 6: intermediation	51
7.8	Family 7: separating beneficiary from payer	52
7.9	Assigning an observation to a type	53
7.10	Limits of the catalogue	54
7.11	A real firm is a bundle of Φ	55
Chapter 8	A Subspace Under Constraints: Alone, Without Assets, Without Credit	56
8.1	The premises imposed in this chapter	56
8.2	Translating the premises	57
8.3	What the cloud means structurally	58
8.4	Sieving the types	59
8.5	The credit bootstrap problem	59
8.6	Relaxing a premise: those who hold assets	60
8.7	The formation of credit	64
8.8	The trade-off between contract work and publishing	67
8.9	Employment contracts reread	69
8.10	Limits of this chapter	70
Chapter 9	Unmanned Operation: The Limit With the Operator Removed	71
9.1	The question	71
9.2	Formalizing “unmanned”	71
9.3	Decomposing the bootstrap	74
9.4	Limits of this chapter	75
Chapter 10	The Supplier of the Information Structure: Settlement by Protocol	77

10.1	The premises imposed in this chapter	77
10.2	Translating the premises	78
10.3	Identifiers diverge from parties	80
10.4	Sieving the types	81
10.5	The collateral constraint	84
10.6	Observability	86
10.7	Limits of this chapter	88
 Part III Method		 89
Chapter 11	Design of the Empirical Work: Connecting to Observation	90
11.1	What is to be tested	90
11.2	Five stages to pass	90
11.3	What to check at each stage	91
11.4	Reconciliation versus verification	92
11.5	Where each stage is handled in this text	94
11.6	Priorities for the survey	94
11.7	Variables to record	95
11.8	What must be pre-registered	96
Chapter 12	Selection Bias: Decomposing the Distortions That Enter Observation	97
12.1	The problem	97
12.2	Conditioning on survival: a circularity in the framework	97
12.3	Selection into disclosure: correlated with the axis of interest	98
12.4	Age: a problem of domain, not of bias	98
12.5	Age, period and cohort are not identified	99
12.6	Retrospective narrative	99
12.7	Mistaking the unit: what a row denotes versus what its values denote	99
12.8	Selection of variables: the gaps are not random	100
12.9	The observation frame in practice	100
12.10	How to proceed	100
Chapter 13	A Taxonomy of Verification Obstacles: Six States	103
13.1	Why a taxonomy is needed	103
13.2	Six categories	104
13.3	The taxonomy also serves as a test before starting	105
 Part IV Empirics		 106
Chapter 14	Measuring the Theoretical Quantities: What Could and Could Not Be Measured	107

14.1	Purpose of this chapter	107
14.2	The primitive sets made concrete	108
14.3	The observed range of the map Φ	108
14.4	Cumulants and the credit position	109
14.5	Measuring the growth constraint	110
14.6	Right and exercise, and cognitive surplus	117
14.7	The information structure $\mathcal{G}$ made concrete	120
14.8	The three-way decomposition of surplus	121
14.9	The terms of the launch cost	122
14.10	Quantities relating to selection and identification	122
14.11	Quantities that could not be measured	123
Chapter 15	Empirics of families 2 and 5: prepayment and outcome-contingent forms	124
15.1	The propositions tested in this chapter	124
15.2	Family 2: prepaid payment instruments	125
15.3	Family 5: fee-charging employment placement	129
15.4	Comparison of the two families and its methodological implications	146
15.5	The structure of the data source	146
15.6	What was reached and what was abandoned	147
15.7	A quantity that is not tabulated: the credit position of sole proprietors	149
Chapter 16	Evidence From the Corporate Statistics: κ and CCC Across Firms	150
16.1	Position of this chapter	150
16.2	Hot and cold are not one-dimensional	150
16.3	Data source	151
16.4	Size and the direction of credit	152
16.5	Is CCC a leading indicator of the cycle?	155
16.6	An observation not anticipated: the long-run rise in CCC	156
16.7	Levels by industry (for reference)	158
16.8	Decomposing the rise in CCC	158
16.9	Separating the contemporaneous negative correlation	160
16.10	The effect of the scope of inventories on the conclusions	162
16.11	An observation the simplified capacity constraint does not fit	166
16.12	Limitations of this chapter	167
16.13	On both implications being rejected	168
Chapter 17	The Price of Externalization: Platform Fees	169
17.1	The implication tested	169
17.2	Result: a ladder of rates	170
17.3	A natural experiment within one platform	171
17.4	A counterexample: price discrimination by size	171

17.5	Verdict	172
17.6	Practical implications	174
17.7	Limits of the measurement	174
17.8	How the institutional layer acts	174
17.9	Assessing feasibility of further investigation	175
17.10	Limits of this chapter	177
Chapter 18	Evidence From Protocol Records: $M(t)$ and the Rate of Modification	178
18.1	The implications tested	178
18.2	Results	179
18.3	Quantities measurable from the same records	180
Part V	Conclusion	183
Chapter 19	Where This Stands: What Was Established and What Remains Open	184
19.1	What was established and what was not	184
19.2	Fixing the implications in advance	187
19.3	Open problems	187
19.4	Known weaknesses of the framework	194
Part VI	Appendices	198
Chapter A	Supplement: proofs concerning method	199
A.1	Correlation induced by conditioning on a collider	199
A.2	Non-identification of age, period and cohort	201
A.3	Preference for variance under a deadline	202
A.4	Substitutability of capital and credit	203
Chapter B	Notation	205
Chapter C	Correspondences with existing theory	208
Chapter D	List of pre-registered implications	212
D.1	The list	212
D.2	Where registration changed the verdict	214
D.3	Defects that arose in operating the registry	214
Chapter E	Revision history	215
E.1	History of versions	215
E.2	Claims withdrawn	218
E.3	Corrections deriving from measurement	219

E.4	Corrections deriving from reasoning	220
E.5	A record of pre-registration at work	220
E.6	Making assumptions explicit	221
E.7	Introducing the lag	221
E.8	Re-evaluating the relation to existing theory	222
E.9	Filling in definitions	222
E.10	Examining reversibility, and a correction	223
E.11	Examination from the worker's viewpoint	224
E.12	Extending the taxonomy of obstacles	224
E.13	Extending the taxonomy of biases	225
E.14	Reorganizing the list of open problems	226
E.15	The means of reproduction were not included	226
E.16	Removing a description resting on a single observation	227
E.17	Setting theorem-type body text in roman	227
E.18	Testing implication Y, and a defect in the registration	227
E.19	A route for testing the deductions of Part II	228
E.20	Registering implications Y and Z	229
E.21	Checking the scope of inventories	229
E.22	Filling in operationalization and reorganizing the list of open problems	230
E.23	Placing the empirical design first	231
E.24	Implication Z tested	231
E.25	Giving deductive grounds for choosing families 2 and 5	233
E.26	Removing references to places that do not exist	233
E.27	Resolving a contradiction in the description of ϕ_{cog}	233
E.28	Extending the scope of inventories to 62 industries	234
E.29	Stating the handling of per-firm data	234
E.30	Multiple definitions of a symbol	234
E.31	Correcting a count	235
E.32	Considered but not adopted in the body	235
E.33	Changes of structure	236
Chapter F	Data sources	237
References		239

Part I

Theory

Chapter 1

Introduction: Choosing the Time Structure of Cash as the Variable

1.1 The problem

The standard variables for describing a firm are revenue, margin, and market share. This text uses none of them. It describes a business model through the **time structure of cash**.

There is a reason for the choice. What keeps a firm in existence is not a condition on profit but a constraint on its cash balance, and that constraint is imposed path by path rather than in expectation (Section 4.2). Any framework that treats survival and growth must take the time structure of cash as a first-class object.

The standard variables are not discarded. **The income statement, the balance sheet, and the cash-flow statement appear as projections of the object constructed here, each collapsing one of its indices** (Table 2.1). This text does not posit quantities separate from accounting; it keeps the indices that accounting collapses.

1.2 Structure of the text, and how to read it

The text has six parts.

Part	Title	Content
I	Theory	Definitions, propositions, proofs. The basis for everything that follows
II	The space of maps	The whole space of Φ (28 types), the subspaces left under constraints, their limit, and the case where the information structure is supplied differently
III	Method	The five stages that connect theory to observation, the treatment of selection bias, and a taxonomy of verification obstacles
IV	Empirics	How each theoretical quantity was measured, and tests in four domains
V	Conclusion	What was established and what remains open
A	Appendices	Proofs of methodological claims, notation, correspondences with existing theory, data sources, pre-registration, revision history

Table1.1: The parts of this text.

Parts I and II are deductive; Part III supplies the procedure that connects theory to evidence.

Part II proceeds in three stages. It first enumerates the space generated by the degrees of freedom of Φ (Chapter 7), then derives the subspace that survives constraints on capital and labour (Chapter 8), and finally treats the limit in which the operator is removed (Chapter 9). All three are deduction, not evidence.

The absence of a separate part on applications is deliberate. An application is warranted once the theory has passed through evidence, and the empirical work here is not complete. Deduction under constraints is not application; it is **the operation of cutting the space of maps with a constraint**, and it therefore belongs in Part II.

Chapter 14 opens Part IV and tabulates, for every symbol introduced in the theory, the value that could be measured and the fact that could not. Each later chapter also opens with a table of the theoretical quantities it handles, so that one can trace where each symbol was measured.

Depending on the reader's interest, the following routes are suggested.

- **Theory alone:** Parts I and II. Proofs of the propositions are in the body; only the proofs of methodological claims are collected in Appendix A
- **Practical design:** Chapter 5 and Part II
- **The magnitudes the theoretical quantities actually take:** Chapter 14
- **Survey method:** Parts III and IV

Numerical examples accompany most sections. The notation is listed in Appendix B.

1.3 What this text does not claim

To forestall misreading, the scope is limited in advance.

First, this is **not a predictive theory**. Most of the implications fixed in advance and tested in Part IV were rejected. The framework works as a descriptive instrument, but this text has no established result about what the quantities it describes predict.

Twenty-two implications were nevertheless fixed in advance. There are two reasons.

The first is that a descriptive framework **can supply an after-the-fact account of any fact whatever**. That something can be explained is no evidence that the framework is right. Unless the implications are fixed in writing before the observations and the ones that miss are left standing as misses, there is no way to distinguish the framework working from the author having reasoned backwards. The places where an after-the-fact story was nearly adopted are listed in Section 19.2. **The fixed implications are not claims of the framework; they are a constraint imposed on the author's own interpretation.**

The second is that **an implication failing and the framework being rejected are different events**. Whether the descriptions CCC, κ and Φ work is one question; what those quantities predict is another. A failed implication refutes a conjecture about the latter, not the former (Section 16.13).

Second, this text is **not normative**. It does not argue which Φ is desirable, only which Φ is feasible under the constraints.

Third, **most of the individual components have counterparts in existing theory**. Appendix C is divided into three layers separating the existing theory that explains each term, the part that connects them, and the organization contributed here. The contribution of this text lies in the connection, not in the novelty of the elements.

Chapter 2

The Business Model: Φ and the Generation of Cash Flow

2.1 Primitive sets

Definition 2.1 (Primitive sets). Four sets are taken as primitive.

$$N = \{0, 1, \dots, n\} \quad \text{the set of parties; 0 is the firm in question} \quad (2.1)$$

$$T = [0, H] \subset \mathbb{R}_{\geq 0} \quad \text{time} \quad (2.2)$$

$$(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in T}, \mathbb{P}) \quad \text{state space and filtration} \quad (2.3)$$

$$\mathcal{D} \quad \text{the space of deliveries (goods and services)} \quad (2.4)$$

The parties are grouped into customers C , suppliers S , capital providers K , labour L , and third parties R who assume risk. These are subsets of N and cover $N \setminus \{0\}$. An individual party is written as an element of N , as in $i \in C$. When a group symbol is used directly as an index it denotes the sum over the parties in that group; for instance $\kappa_C = \sum_{i \in C} \kappa_i$. R is where degree of freedom (3) moves the paying party, which is the transfer operation of Section 2.6.

Regulation and commercial custom are not counterparties to cash and are therefore not in N . They constrain the range from which Φ may be chosen and are treated as the institutional layer in Section 6.3.

2.2 Schedules and spaces

The relationship with party i is described as a pair of schedules.

$$\delta_i : T \times \Omega \rightarrow \mathcal{D} \quad \text{delivery (value handed from the firm to } i) \quad (2.5)$$

$$\pi_i : T \times \Omega \rightarrow \mathbb{R} \quad \text{settlement (payment from } i \text{ to the firm)} \quad (2.6)$$

A valuation map $v : \mathcal{D} \rightarrow \mathbb{R}$ measures deliveries in monetary units.

The collections of such tuples are taken as spaces.

$$\Delta = \{ (\delta_i)_{i \in N} \mid \delta_i : T \times \Omega \rightarrow \mathcal{D} \}, \quad \Pi = \{ (\pi_i)_{i \in N} \mid \pi_i : T \times \Omega \rightarrow \mathbb{R} \} \quad (2.7)$$

Δ is called the *space of delivery schedules* and Π the *space of settlement schedules*.

2.3 The map from delivery to settlement

Definition 2.2 (Business model). Using Δ and Π from (2.7), a business model is defined as the map

$$\Phi : \Delta \longrightarrow \Pi \quad (2.8)$$

Remark 2.3 (Φ and ϕ are different objects). Capital Φ denotes the map (2.8) itself, that is, the **structure**. Lower-case ϕ denotes the **size of the surplus** that arises in one period under that structure; it is decomposed into three sources in (6.1) of Chapter 6. Φ fixes the shape of the correspondence between delivery and settlement; ϕ measures the share the firm retains under that shape. No general relation between the two is given here. How the choice of Φ moves each term of ϕ is discussed term by term in Chapter 6.

2.4 Generating the cash flow

Once Φ in (2.8) is fixed, the settlement with each party is fixed. This is called the cash flow.

Definition 2.4 (Cash flow). Given a delivery $\delta \in \Delta$, the map Φ returns as its image the tuple of settlements $\pi = \Phi(\delta) \in \Pi$. An element of Π is a tuple of functions indexed by party, so its i -th component π_i is a function on $T \times \Omega$. Write

$$X(i, t, \omega) \equiv \pi_i(t, \omega), \quad \pi = \Phi(\delta) \quad (2.9)$$

Then $X(i, t, \omega)$ is the net movement of cash between the firm and party i at time t in state ω . Inflows to the firm are positive. Payments to suppliers or to labour appear as $\pi_i < 0$.

Φ takes exactly one argument, the tuple of deliveries δ ; it does not take t or ω . The dependence on time and state is carried by the π_i that Φ returns.

X is the image of Φ . Replacing Φ produces a different X from the same delivery δ . This is the central operation of the text.

Equation (2.9) carries three indices: who, when, and in which state. The usual financial statements are projections that collapse this object along one index.

Statement	Operation	Content
Income statement	aggregation over t	includes accrual adjustments
Balance sheet	a cross-section at fixed t	unsettled balances
Cash-flow statement	differences along t	the actual movement of cash

Table2.1: The financial statements are projections that collapse the three indices of (2.9).

Remark 2.5 (Every quantity here is a level). Every quantity introduced after (2.9) — including κ_i , CCC and ϕ — is a level. Each is defined as a value at a time t , and no notation for change is provided.

This is not a presentational choice but a **limitation**. As Part IV confirms repeatedly, the levels of these quantities differ by two orders of magnitude across industries and product forms, and **cross-sectional comparison cannot be interpreted**. What can be read is change over time within one object.

The text has no theoretical apparatus for change and deals with it case by case in the empirical work (Section 19.4.3).

Example 2.6 (The three indices made concrete). Consider a service priced at 1,000 yen a month and supplied to 100 customers. If customer $i \in C$ pays at the end of each month, then

$$X(i, t, \omega) = \begin{cases} 1,000 & t \in \{1, 2, \dots\} \text{ (month end)} \\ 0 & \text{otherwise} \end{cases}$$

which does not depend on ω . If cancellation is possible, ω contains the cancellation time and X does depend on ω .

With a single supplier, $S = \{s\}$, and a server cost of 30,000 yen a month, $X(s, t, \omega) = -30,000$. Written with the group index, $\sum_{i \in S} X(i, t, \omega) = -30,000$.

2.5 Degrees of freedom of Φ

The features that specify Φ in (2.8) are organized into six degrees of freedom. They divide into those that correspond to the three indices of (2.9) and those that do not.

■ Those that move an index.

- (1) **Timing**: the order of $\text{supp}(\pi)$ and $\text{supp}(\delta)$. Moves t . Fixes the sign of κ .
- (3) **Paying party**: the beneficiary, a third party, or the opposite side of the market. Moves the index i .
- (4) **State dependence**: whether π is measurable with respect to ω . Moves ω , and fixes where risk resides. This is exactly the $\mathcal{G}$ -measurability of Chapter 4; the range the designer of Φ may choose from is bounded by $\mathcal{G}$ (Proposition 4.1, Remark 4.2).

■ That which fixes the functional form of the map.

- (2) **Form of dependence**: how π depends on δ . The linear forms collapse into a single family.

$$\pi = F + p \delta \tag{2.10}$$

F is the part independent of the quantity delivered and p is the unit price. $p = 0$ gives a flat fee, $F = 0$ gives usage pricing, and both non-zero gives a two-part tariff. A non-linear form is outcome-contingent pricing $\pi = h(y)$, where y is the outcome and h turns the outcome into a settlement; whether y is observable belongs to degree of freedom (4) (Chapter 4).

Remark 2.7 (The typeface of F , and the use of c). The F in (2.10) is the fixed part of a tariff and is a different quantity from the σ -algebra $\mathcal{F}$ on the state space. The two are distinguished by typeface.

Throughout this text c denotes marginal cost and is always written with its argument, $c(\delta)$ (Definition 2.14). A flat fee is written F .

■ Properties on the side of the domain Δ .

- (5) **Separation of right from exercise**: whether the contractual right $\bar{\delta}$ and realized use δ can diverge.
- (6) **Repetition**: one-off, auto-renewing, or with a committed term. Fixes the transaction frequency ν .

Remark 2.8 (Three groups). (1), (3) and (4) move the indices of (2.9): they **change where X is placed for the same δ** . (2) fixes the form of the map itself (Remark 2.7), and (5) and (6) are properties of the domain Δ .

The three risk operations of Section 2.6 are the operations that move (1), (3) and (4).

Of these, (1) and (2) respectively fix the sign of κ and the room for cognitive surplus. These two axes give a coarse partition of the space of Φ (Figure 2.1). **The two are taken from different groups**: (1) moves an index and (2) fixes a functional form.

The two axes of the partition and the degrees of freedom used to assign a family are different things. The assignment order of Section 7.9 uses (3), (4) and (6); (1) and (2) do not appear. The former are axes for viewing the space of Φ coarsely, the latter a rule for assigning an observed Φ uniquely; the purposes differ.




Figure2.1: A two-axis partition of Φ . The vertical axis fixes the sign of κ , the horizontal axis the room for cognitive surplus.

2.6 Three risk operations

Because there are three indices, there are only three operations on X . The correspondence is an organization contributed here.

Index moved	Name	Invariant	Examples
t	timing	total	advances, instalments, leases
i	transfer	total and variance	factoring, buying insurance
ω	pooling	total	underwriting, diversified investment

Table2.2: Three risk operations: the index moved, the invariant, and typical examples.

As Remark 2.8 states, these three are the operations that move degrees of freedom (1), (3) and (4): timing moves t , transfer moves i , and pooling moves ω .

Of the three, only pooling actually reduces variance. This is a consequence of probability theory and holds independently of the structure of Φ . **It is quoted here, not derived.**

Proposition 2.9 (Variance reduction by pooling). For identically distributed risks $X_1, \dots, X_n$ with variance σ^2 and pairwise correlation ρ , **provided ρ does not depend on n** ,

$$\text{Var} \left[\frac{1}{n} \sum_{j=1}^n X_j \right] = \frac{\sigma^2}{n} (1 + (n-1)\rho) \xrightarrow{n \rightarrow \infty} \rho \sigma^2. \quad (2.11)$$

Proof. By bilinearity of covariance,

$$\text{Var} \left[\sum_j X_j \right] = \sum_j \text{Var}[X_j] + \sum_{j \neq k} \text{Cov}[X_j, X_k] = n\sigma^2 + n(n-1)\rho\sigma^2.$$

Dividing both sides by n^2 gives the first equality. As $n \rightarrow \infty$, $\frac{\sigma^2}{n} \rightarrow 0$ and $\frac{(n-1)\rho\sigma^2}{n} \rightarrow \rho\sigma^2$, which gives the limit. $\square$

Remark 2.10 (Independence is a finite resource). The limit in (2.11) is $\rho\sigma^2$, which is non-zero whenever $\rho > 0$. Pooling consumes independence, and independence is a finite resource. Underwriting limits on earthquake insurance, the sensitivity of platforms to the business cycle, and systemic risk among financial institutions all come from the same term of this equation.

Remark 2.11 (Endogeneity of ρ). Proposition 2.9 treats ρ as an exogenous constant. The assumption is hard to sustain.

The mechanism that makes ρ endogenous is not the exhaustion of poolable risks but the fact that **the act of pooling itself creates correlation**. The example of [30] is clear: if two parties each take half of two risks and diversify perfectly, each party's individual probability of failure falls, but **their portfolios become perfectly correlated**. The nature of the risk is unchanged, yet correlation between the parties appears.

Hence ρ is a function of the number of poolers and the overlap of their holdings, and the limit $\rho\sigma^2$ in (2.11) can rise as pooling proceeds. The proof is formally correct, but **its range of application is confined to cases with few poolers and little overlap**. This dependence is not treated here (Section 19.4.1).

Remark 2.12 (Social optimality of pooling). The mechanism of Remark 2.11 implies that the individual and the collective optimum can diverge. [19] derives conditions under which the diversification optimal for an intermediary is not socially optimal and restricting risk sharing is preferable.

This text is written from the standpoint of an individual operator and does not treat the effect of the choice of Φ on the system. Family 5-4 (underwriting) and family 6-2 (escrow) in Part II make pooling their business and are the types where this point bears directly.

Example 2.13 (How fast correlation bites). With $\sigma = 100$ and $\rho = 0$, at $n = 100$ the standard deviation falls to $100/\sqrt{100} = 10$. With $\rho = 0.3$, however,

$$\text{Var} = \frac{100^2}{100}(1 + 99 \times 0.3) = 305.5, \quad \text{standard deviation} \approx 17.5.$$

Even as $n \rightarrow \infty$ the standard deviation does not fall below $\sqrt{0.3} \times 100 \approx 54.8$. At a correlation of 0.3 — a modest value — the benefit of pooling is more than halved.

2.6.1 Cost fixed by technology

Producing a delivery δ costs something. Because this text does not treat the composition of inputs, **production technology appears only as a correspondence C_{prod} from delivery to cost**.

C_{prod} **is given independently of Φ** . Different Φ can be placed on the same technology, and the same Φ on different technologies. What this text varies is Φ , not the technology.

Definition 2.14 (Splitting the cost). Total cost is split in two.

$$C(\delta, \Phi) = \underbrace{C_{\text{prod}}(\delta)}_{\text{fixed by technology}} + \underbrace{C_{\text{admin}}(\Phi, \delta)}_{\text{created by the contract}} \quad (2.12)$$

For the first term, write

$$C_{\text{prod}}(0) \quad \text{fixed cost} \quad (2.13)$$

$$c(\delta) = \frac{dC_{\text{prod}}}{d\delta} \quad \text{marginal cost} \quad (2.14)$$

Remark 2.15 (No shape is assumed for $c(\delta)$). Neither monotonicity nor convexity is assumed for $c(\delta)$. Cost structures differ greatly across industries, and imposing one shape would narrow the range of application.

The price is that **the optimal quantity delivered δ cannot be fixed at an interior point**. The capacity constraint (2.17) supplies an upper bound, but the point short of it at which cost

balances price does not follow from the assumptions made here. Where the level of δ matters later — the allocation in Chapter 8, for instance — it is treated as a corner of the constraint set.

The split is meaningful because the second term is real. C_{admin} arises from all six degrees of freedom of Φ (Section 2.5).

Kind	Route	Content
Setup	conclusion	the cost of agreeing Φ with the counterparty; small for standard contracts, large for negotiated ones
Operation	measurement	under $\pi = p\delta$, δ must be metered and billed
	rights management	administering $\bar{\delta}$: memberships, tracking balances
	ongoing handling	auto-renewal, cancellation
Securing performance	collection and bad debt	under $\tau_C > 0$, invoicing, chasing, and losses on non-collection
	verification	under $\pi = h(y)$, confirming y
	dispute	disputes arise even when the contract is written
Outsourcing	intermediation	settlement, delivery, customer acquisition performed by others

Table2.3: Routes through which C_{admin} arises.

Remark 2.16 (Advances cost something too). $\tau_C < 0$ (an advance) requires no collection cost, but it does create custody of the funds held, refund handling, and regulatory compliance. The deposit obligation under the Payment Services Act, treated in Section 8.5, is an example. $C_{\text{admin}} > 0$ **whatever the sign of τ .**

Remark 2.17 (Conditions under which C_{admin} is ignored). From here on C_{admin} is not written explicitly. Two conditions make it negligible.

- (1) **Operations are systematized.** Once the machinery is built, the marginal cost per transaction of measurement, rights management and ongoing handling approaches zero.
- (2) **No intermediary is used.** Intermediation fees are proportional to transaction value and do not vanish with systematization. Chapter 17 measures this cost and maps it onto the kinds of C_{admin} in Remark 17.1.

The range in which the conditions hold is limited. They hold for the solo business of Chapter 8 and for family 2 in Part IV, but **not for the small firms treated in Chapter 16**. Where invoicing and chasing are done by hand, C_{admin} cannot be ignored.

The price of the omission should be stated. Because C_{admin} is dropped, this text **cannot treat the route by which the choice of Φ feeds back into ϕ through cost**. Equation (2.12) is not used again; only C_{prod} appears. A relation such as “lengthening τ_i raises collection cost” cannot be written within this framework.

What is called a business model here is therefore the freedom to choose different contractual forms on an identical cost structure. That freedom is not complete: the choice of Φ feeds back into cost through C_{admin} . Where C_{prod} appears is collected in Remark 2.19.

Remark 2.18 (Switching cost and the value of waiting). Treating Φ as a choice variable amounts to **setting the switching cost to zero**. The cost of changing from one Φ to another belongs to the setup term of C_{admin} and is not generally zero.

In the theory of irreversible investment ([8]), a decision is irreversible not because it cannot physically be undone but because **undoing it is expensive and locks in a particular path**. Writing C_{sw} for the switching cost and $\Delta\phi$ for the gain in surplus from changing Φ :

$C_{\text{sw}} = \infty$	no change is possible; the first choice is final
$C_{\text{sw}} \ll \Delta\phi$	the assumption made here: Φ is freely chosen
$C_{\text{sw}} \sim \Delta\phi$	there is value in delaying the choice

Table2.4: The size of the switching cost C_{sw} and the freedom to choose Φ .

In the third case **waiting has value**: it can be optimal not to fix Φ until the uncertainty resolves. This implication is not treated here. For the solo business of Chapter 8, where the switching cost is relatively large, the point may not be negligible.

Remark 2.19 (Where C appears). The quantities of Definition 2.14 are used repeatedly. ϕ_{prod} is defined as an advantage in $c(\delta)$ (Chapter 6); a near-zero marginal cost is the condition that relaxes the capacity constraint (Chapter 8); and the fact that both C and Φ can be specified in advance is the condition for replicability (Chapter 9). The physical layer refers to these properties of C (Section 6.3).

Example 2.20 (Identical C_{prod} , different Φ). The same business software can be supplied in three ways.

Φ	Settlement	κ_C (mid-period, one customer)
Perpetual licence	240,000 yen on installation	−120,000 yen
Annual	60,000 yen at the start of each year	−30,000 yen
Monthly in arrears	5,000 yen at each month end	+2,500 yen

Table2.5: Three choices of Φ on an identical C_{prod} , and the resulting κ_C .

The technology is identical and so is the marginal cost. Only Φ differs, yet the sign of κ reverses and the working capital required changes.

2.7 Divergence between right and exercise

Degrees of freedom (2) and (5) are linked. When π is proportional to δ the customer pays for what is used, so right and exercise coincide. Only when π is independent of δ can the two diverge.

When degree of freedom (5) of (2.8) is present, define the following.

Definition 2.21 (Right and exercise). Let $\bar{\delta}$ be the ceiling the contract grants the customer and δ the use actually realized. Write the divergence as $\bar{\delta} - \delta \geq 0$ and define the gross margin arising from it as

$$\phi_{\text{cog}} = p \cdot \mathbb{E}[\bar{\delta} - \delta] \quad (2.15)$$

Remark 2.22 (Which contractual forms admit divergence). Equation (2.15) can be positive only when the contract states $\bar{\delta}$ explicitly.

Under usage pricing $\pi(t) = p\delta(t)$, the customer's total payment is $p \int \delta$. No use means no payment, so the contract need not state a ceiling $\bar{\delta}$ and the difference $\bar{\delta} - \delta$ is undefined.

Under a flat fee, where π does not depend on δ , the consideration is fixed irrespective of use, so the contract states the extent of the right $\bar{\delta}$. Here $\delta < \bar{\delta}$ can be realized. Under a two-part tariff $\pi = F + p\delta$, divergence arises only over the range corresponding to the fixed part F .

This is a consequence of Definition 2.21, not an independent claim. It presumes that $\bar{\delta}$ is given exogenously. When the customer chooses $\bar{\delta}$ from a menu, the two are determined jointly and the meaning of the difference changes (Section 19.4.1).

Remark 2.23 (The general form of the capacity constraint, and the simplification used here). Capacity on the supply side covers the occupancy of a facility, the number of machines, the disposable time of staff, and so on. Since δ_i is a function on $T \times \Omega$ by the definitions of Chapter 2, the general form of the capacity constraint is evaluated at each instant.

$$\sum_i \delta_i(t, \omega) \leq \text{Cap}(t, \omega) \quad \forall t \in T, \forall \omega \in \Omega \quad (2.16)$$

Capacity itself depends on time because opening hours, equipment run-times and staffing all vary over time.

From here on Cap is treated as a constant, that is, (2.16) is used in the form integrated over the whole period,

$$\sum_i \delta_i \leq \text{Cap} \quad (2.17)$$

where δ_j denotes a total.

The simplification is justified when demand is distributed over time in roughly the same way across members and capacity does not vary by time of day. When neither holds, only the peak instant binds and (2.17) is stronger than necessary. Chapter 16 treats an observation with which this simplification does not agree.

Proposition 2.24 (Membership under a capacity constraint). Suppose the capacity constraint (2.17) is present and **the number of members n and average use $\mathbb{E}[\delta]$ are determined independently** (the number of members is written n because m is reserved for the margin). Then a flat membership $\pi_j = F$ is feasible only if

$$\sum_j \mathbb{E}[\delta_j] \ll \sum_j \bar{\delta}_j \quad (2.18)$$

that is, low utilization is a condition of feasibility.

Proof. Let n be the number of members and $\bar{\delta}$ each member's contractual ceiling. If every member exercised the right, demand would be $n\bar{\delta}$. Feasibility requires $n\mathbb{E}[\delta] \leq \text{Cap}$. On the other hand, since members join because they value $\bar{\delta}$, normally $n\bar{\delta} > \text{Cap}$ (otherwise the capacity constraint does not bind). Combining the two gives $\mathbb{E}[\delta] \leq \text{Cap}/n < \bar{\delta}$. $\square$

Example 2.25 (Two kinds of membership). A facility with room for 100 has 1,000 members and charges 8,000 yen a month. If $\bar{\delta}$ is “use it every day”, $n\bar{\delta}$ far exceeds capacity. At an actual utilization of 10%, $n\mathbb{E}[\delta] = 100$ and the constraint is just satisfied. **If utilization rises to 30%, the business fails physically.**

By contrast, where the fee grants a right to purchase discounts there is no capacity constraint, and higher utilization is favourable because it enlarges the purchasing volume. Outwardly both belong to family 2-1, but the shape of the feasible region differs.

Chapter 3

The Credit Position: κ and the Lag τ

3.1 The gap between delivery and settlement

Write the cumulants of the delivery δ_i of Section 2.2 and of the settlement π_i fixed by Φ in (2.8) (with $\pi = \Phi(\delta)$) as

$$D_i(t) = \int_0^t v(\delta_i(s)) ds, \quad P_i(t) = \int_0^t \pi_i(s) ds \quad (3.1)$$

D_i is value given and P_i is consideration received.

Definition 3.1 (Credit position). The credit position towards party i is defined as

$$\kappa_i(t) = D_i(t) - P_i(t) \quad (3.2)$$

By (3.1),

$$\kappa_i(t) = \int_0^t v(\delta_i(s)) ds - \int_0^t \pi_i(s) ds$$

so κ_i is a function of the Φ of (2.8): replacing Φ for the same delivery δ changes κ_i . A positive κ_i means the firm is extending credit to i (delivered, not yet collected); a negative κ_i means i is extending credit to the firm (received, not yet delivered).

The sum of credit positions plus inventory is used repeatedly below.

Definition 3.2 (Working capital). With inventory $\text{Inv}(t) \geq 0$, define

$$W(t) = \sum_{i \in N} \kappa_i(t) + \text{Inv}(t) \quad (3.3)$$

Since κ_i is a function of Φ , so is W .

The W of (3.3) is the amount of money tied up in the business; dividing it by the rate of sales gives CCC in Chapter 5.

The merit of (3.2) is that it unifies.

Account	Party	Sign	Meaning
Receivables	C	$\kappa_C > 0$	the firm extends credit to the customer
Advances received, contract liabilities	C	$\kappa_C < 0$	the customer extends credit to the firm
Payables	S	$\kappa_S < 0$	the supplier extends credit to the firm
Advances paid	S	$\kappa_S > 0$	the firm extends credit to the supplier

Table3.1: Four accounts unified by (3.2).

These are not four separate accounts but one quantity with different signs and indices.

3.1.1 Existing theory for each term

Each term of κ_i has an established theory behind it. This text takes them as its foundation.

Party and sign	Phenomenon	Explanatory theory
$\kappa_C > 0$ (inter-firm)	selling on account	four motives for trade credit: financing advantage, price discrimination, quality assurance, liquidation advantage [14, 15]
$\kappa_C < 0$ (inter-firm)	advance payment	reverse trade credit: supplier finance, transaction assurance, bargaining power [12, 7]
$\kappa_C < 0$ (to consumers)	prepaid contracts	mental accounting, evidence on breakage [13]
$\kappa_S < 0$	payables	the mirror image of trade credit
$\kappa_L < 0$	wages in arrears	the bonding argument: discipline through deferred pay [1]
$\kappa_K < 0$	borrowing	standard corporate finance
κ_K^i	equity (residual claim)	party-dependent; Definition 3.21

Table3.2: Existing theory explaining each term of κ_i .

Remark 3.3 (The four literatures do not cite one another). The theories in Table 3.2 developed in different fields to answer different questions.

Theory	Field	Central question
Trade credit	corporate finance	why a supplier can lend more cheaply than a bank
Bonding argument	labour economics	why involuntary unemployment exists
Signalling	economics of information	how ability is disclosed
Work on prepaid contracts	consumer behaviour	why consumers pay first

Table3.3: The fields and central questions of the theories explaining each term of κ_i .

Because the questions differ, these literatures do not cite one another. Yet in (3.8) **the customer's prepayment, the supplier's trade credit, the worker's deferred wage, and the investor's equity all appear as terms of the same sum.**

Each theory explains one term of the sum; none has a framework that treats the sum. The contribution here is that integration, not the explanation of any individual term.

The relation between the cumulants is shown in Figure 3.1.




Figure3.1: The credit position measured against cumulative delivery $D(t)$. Where the settlement curve $P(t)$ lies above $D(t)$, $\kappa < 0$; where it lies below, $\kappa > 0$.

3.1.2 Both sides of a transaction: κ is antisymmetric

Equation (3.2) is written on the firm's own books. Writing the same transaction on the counterparty's books exchanges value given for consideration received.

Proposition 3.4 (Antisymmetry). If two parties i and j share the valuation map v , and $\kappa_j^{(i)}$ denotes the credit position towards j on the books of i , then

$$\kappa_j^{(i)} = -\kappa_i^{(j)} \quad (3.4)$$

Proof. Value given by i to j is value received by j from i , and consideration received by i from j is consideration given by j to i . Hence $D_j^{(i)} = P_i^{(j)}$ and $P_j^{(i)} = D_i^{(j)}$; substituting into (3.2) gives the result. $\square$

Remark 3.5 (A shared v is required). Equation (3.4) presumes both parties use the same v . As Remark 3.23 notes, v may be subjective; but the settlement amount π is $\mathcal{G}$ -measurable and agreed, so any disagreement is confined to the D side. Differences in acceptance criteria and disputed receivables are the cases in point.

Antisymmetry has a consequence for aggregation.

Corollary 3.6 (Inter-firm credit vanishes on aggregation). Let $\mathcal{N}$ be the set of parties in the corporate sector. The sum of credit positions closed within $\mathcal{N}$ is zero, so

$$\sum_{i \in \mathcal{N}} W^{(i)} = \kappa_{\text{vs. non-corporate}} + \sum_{i \in \mathcal{N}} \text{Inv}^{(i)} \quad (3.5)$$

where $\kappa_{\text{vs. non-corporate}}$ is the net position of the corporate sector against households, government and the rest of the world.

Proof. For each pair (i, j) inside $\mathcal{N}$, (3.4) gives $\kappa_j^{(i)} + \kappa_i^{(j)} = 0$. Summing the W of Definition 3.2 over all parties in $\mathcal{N}$ cancels the internal pairs and leaves the pairs with the outside together with inventory. $\square$

Corollary 3.7 (Not every stratum can be a net provider of credit). Partition the corporate sector into disjoint strata. It is impossible for every stratum to be a net provider of inter-firm credit. If one stratum is a net provider, another is a net recipient.

Proof. By Corollary 3.6 the sum of inter-firm positions across strata is zero. If the net position of every stratum were strictly positive, the sum would be strictly positive, contradicting zero. $\square$

Corollary 3.7 places a structural constraint on comparisons across size strata. When the direction of credit is measured by size in Part IV, part of the observed sign pattern follows automatically from this constraint.

Remark 3.8 (What the aggregate CCC measures). Dividing (3.5) by total sales gives the aggregate CCC. Since inter-firm positions have dropped out of the numerator, **the aggregate CCC is not an indicator of inter-firm credit**. What remains is the position against the non-corporate sector, plus inventory. Chapter 16 reports the CCC of all industries, and interpreting its level requires this distinction.

3.1.3 Additivity over parallel Φ

One party may run several Φ in parallel.

Proposition 3.9 (Additivity of κ). If $\Phi^1, \dots, \Phi^K$ are run in parallel on disjoint deliveries $\delta^1, \dots, \delta^K$, then

$$\kappa = \sum_{k=1}^K \kappa^k, \quad W = \sum_{k=1}^K W^k \quad (3.6)$$

Proof. D_i and P_i in (3.1) are integrals and therefore additive in the integrand. Equation (3.2) is their difference and preserves additivity. $\square$

What is additive is κ and W , not CCC. Since $\text{CCC} = W/r$ is a ratio, and ratios are not additive (Corollary 5.19).

Remark 3.10 (The decomposition is not unique). Equation (3.6) asserts that a decomposition exists, not that it is unique. Several families $\{\Phi^k\}$ give the same κ . The types of Chapter 7 are a generating set for Φ , not a basis.

3.1.4 Expressing the position as a lag

κ_i is a difference of cumulants, and what it contains is **the time gap between delivery and settlement**. Make this explicit.

Definition 3.11 (Settlement lag). For party i , write τ_i for the average lag from delivery to settlement. Its sign follows that of κ_i : settlement later than delivery is positive.

Sign	State	Examples
$\tau_i > 0$	settlement later than delivery (the firm extends credit)	receivables, selling on account
$\tau_i < 0$	settlement earlier than delivery (i extends credit)	advances received, payables

Table3.4: The sign of the settlement lag τ_i and the corresponding state.

Express the κ_i of (3.2) through the lag τ_i .

Proposition 3.12 (Decomposition of the credit position). In a steady state with a constant rate of delivery, as a time average

$$\bar{\kappa}_i = r \tau_i \quad (3.7)$$

Proof. If settlement lags delivery uniformly by τ_i , the unsettled cumulant at time t is the delivery of the most recent τ_i of time:

$$\kappa_i(t) = \int_{t-\tau_i}^t r(s) ds.$$

If r is constant the right-hand side equals $r\tau_i$. If the lag has a distribution, take τ_i to be its mean. $\square$

Remark 3.13 (It holds only on average). Actual settlement is discrete. Under monthly billing paid at the end of the following month, π is a pulse once a month and $\kappa_i(t)$ moves in a sawtooth. Equation (3.7) holds only for its time average.

In statistics based on period-end balances, if fiscal year-ends cluster in particular months the phases of the sawtooth line up and the figures can deviate systematically from the mean. No correction for this bias is made here.

Example 3.14 (κ of an annually billed subscription). A service priced at 1,000 yen a month is paid annually, 12,000 yen received on 1 January. Delivery is spread evenly over twelve months, so after t months

$$\kappa_C(t) = 1,000 t - 12,000.$$

This is $-12,000$ at $t = 0$, $-6,000$ at $t = 6$, and 0 at $t = 12$. With 100 customers, the mid-period average is $\kappa_C = -600,000$ yen. That amount is interest-free funding.

3.2 Decomposing the cash balance

Use the κ_i of (3.2) to decompose the cash balance.

Proposition 3.15 (Credit decomposition of cash). With equity $E(t)$, inventory $\text{Inv}(t)$ and fixed assets $A(t)$,

$$M(t) = E(t) + \sum_{i \in N} \max\{-\kappa_i(t), 0\} - \sum_{i \in N} \max\{\kappa_i(t), 0\} - \text{Inv}(t) - A(t) \quad (3.8)$$

Proof. By the balance-sheet identity, total assets equal total liabilities plus net assets. Split assets into cash M , positive credit positions (receivables, advances paid), inventory Inv and fixed assets A ; map liabilities to negative credit positions (payables, advances received) and net assets to E . Then

$$M + \sum_i \max\{\kappa_i, 0\} + \text{Inv} + A = \sum_i \max\{-\kappa_i, 0\} + E.$$

Solving for M gives (3.8). $\square$

Equation (3.8) is only a rearrangement of an identity, but it fixes how to read it. The second term is the total credit drawn from others and the third the total credit extended to others.

Corollary 3.16 (What cash is). A firm's cash is the total credit drawn from all counterparties, less the credit extended and the assets fixed in place.

3.2.1 The net form

The second and third terms of (3.8) are written separately in order to show the gross amounts. Looking only at the net, for any real x we have $\max\{-x, 0\} - \max\{x, 0\} = -x$, so the two terms cancel.

Proposition 3.17 (Cash is the reverse side of working capital). With the working capital W of Definition 3.2, (3.8) can be written

$$M(t) = E(t) - W(t) - A(t) \quad (3.9)$$

Proof. Summing $\max\{-\kappa_i, 0\} - \max\{\kappa_i, 0\} = -\kappa_i$ over i gives $\sum_i \max\{-\kappa_i, 0\} - \sum_i \max\{\kappa_i, 0\} = -\sum_i \kappa_i$. Substitute into (3.8) and collect Inv into W . $\square$

Checking against Example 3.19: $\sum_i \kappa_i = -600 - 50 + 200 = -450$ and $\text{Inv} = 0$, so $W = -450$, and $100 + 450 - 30 = 520$ agrees.

Equation (3.9) says the same thing as (3.8), but **it connects, through W , to the CCC of Chapter 5**. How the constraint $M \geq 0$ of Section 4.2 bears on the choice of Φ is settled by way of this form.

Remark 3.18 (The identity is a standard one). Equation (3.9) is the identity used in financial analysis, cash equals equity minus net operating assets. Its significance here is that W is a function of Φ through (3.2), so the cash constraint becomes a constraint on Φ .

Example 3.19 (Funding from three sources). A business started with 1,000,000 yen of equity receives 6,000,000 yen of annual prepayments from customers ($\kappa_C = -6,000,000$), owes 500,000 yen to a cloud provider ($\kappa_S = -500,000$), and has 2,000,000 yen of receivables from a large customer ($\kappa_{C'} = +2,000,000$). There is no inventory and fixed assets are 300,000 yen. In units of 10,000 yen,

$$M = 100 + (600 + 50) - 200 - 0 - 30 = 520.$$

Against 1,000,000 yen of equity, cash on hand is 5,200,000 yen. The difference of 4,200,000 yen is credit drawn from customers and suppliers, and it has a due date.

3.2.2 Equity and capital providers

The E of Proposition 3.15 appears as a residual; its composition is made explicit here. Equation (3.8) uses this E .

Definition 3.20 (Equity). With paid-in capital $P_K(t)$ and cumulative dividends $\text{Div}(t)$, define

$$E(t) = P_K(t) + \int_0^t \phi(s) ds - \text{Div}(t) \quad (3.10)$$

The second term is retained earnings.

Definition 3.20 links the ϕ of Chapter 6 to the E of this chapter. Substituting the three-way decomposition of ϕ ,

$$\text{retained earnings} = \int (\phi_{\text{prod}} + \phi_{\text{barg}} + \phi_{\text{cog}}) ds - \text{Div}$$

Chapter 6 treats the erosion of ϕ_{cog} by learning and regulation, but **what is eroded is the future flow; what has already accumulated remains in E .**

Definition 3.21 (Credit position with capital providers). For a capital provider $i \in K$ with probability measure $\mathbb{P}_i$, let $V_i(t) = \mathbb{E}_i[\text{present value of future cash flows}]$ be the value of the residual claim, and define

$$\kappa_K^i(t) = V_i(t) - P_K(t) \quad (3.11)$$

Proposition 3.22 (κ_K does not enter (3.8)). Because κ_K^i depends on the capital provider i , it cannot be included in the identity (3.8).

Proof. Equation (3.8) is a rearrangement of the balance-sheet identity and consists only of quantities on whose value all parties agree. Receivables and payables agree in amount between the parties, but the V_i of (3.11) depends on $\mathbb{P}_i$ and does not. Including a party-dependent quantity would make the identity differ by party, and it would cease to be an identity. $\square$

Remark 3.23 (Three degrees of subjectivity). Several unobservable quantities appear in this text, and they differ in character. Using the $\mathcal{G}$ of Chapter 4 they fall into three degrees.

Degree	Condition	Examples
Shared	$\mathcal{G}$ -measurable	the settlement amount π , receivables
Private information	$\mathcal{F}_t$ -measurable but not $\mathcal{G}$ -measurable	the valuation map v , effort a
Measure-dependent	$\mathcal{F}_t$ -measurable but dependent on $\mathbb{P}_j$	the value of the residual claim V_j

Table3.5: Three degrees of subjectivity among unobservable quantities.

The subjectivity of v is informational asymmetry and can be resolved by disclosure. The subjectivity of V_j is **disagreement about the probability measure**; the same information yields different forecasts, so disclosure does not resolve it.

Two consequences follow. First, the value of κ_K^i does not agree between the parties and so differs in character from the other κ_i . Second, **it is precisely that disagreement that makes equity investment possible**: an investor invests because they value V more highly than the manager does. This has the same structure as the heterogeneity of discount factors in Proposition 4.10.

The maturity of capital is discussed in Remark 3.24.

Remark 3.24 (τ_K and insolvency). For borrowing, τ_K is fixed as the repayment date; for equity it is not fixed. Capital is the longest-dated credit and has no maturity.

Making dividends obligatory would fix κ_K , but (4.3) would then break immediately upon non-payment. Leaving them non-obligatory as a residual claim, and forcing liquidation at the insolvency time of Section 4.2, avoids this. Bankruptcy is the mechanism that fixes κ_K .

Remark 3.25 (Funds held on behalf of others). Among positions with $\kappa_i < 0$, funds that legally belong to someone else — escrow, for instance — occupy the same place in (3.8) but differ in character. They drain rapidly when volume falls, and must not be treated as equivalent to other negative κ .

Chapter 4

Information Structure and the Objective: When a Contract Can Be Written, and the Path Constraint

4.1 Verifiable shared information

Write $\mathcal{G}_i \subseteq \mathcal{F}$ for the information party i can observe, and $\mathcal{G} = \bigcap_i \mathcal{G}_i$ for the shared information verifiable between the parties.

Proposition 4.1 (Implementability). A necessary condition for a settlement schedule π to be implementable as a contract is that π be $\mathcal{G}$ -measurable.

Proof. Suppose π is not $\mathcal{G}$ -measurable. Then there exist two states ω, ω' indistinguishable under $\mathcal{G}$ with $\pi(\omega) \neq \pi(\omega')$. The payment obligation depends on which state obtains, yet the parties cannot jointly verify which one does. They therefore cannot reach agreement on the amount, and no third party can enforce performance. $\square$

This condition constrains the design of Φ severely. When the effort level a is not in $\mathcal{G}$, no contract conditioned on a can be written and only the observable outcome y can be used.

What is observable	Contract form	Examples
Effort a	input-linked	hourly wages, cost plus
Only the outcome y	outcome-contingent	success fees, royalties, deductibles
Both, partially	mixed	fixed plus commission

Table4.1: What is observable and the corresponding contract form.

This is a standard result of contract theory ([17]) and is not a contribution of this text.

Remark 4.2 (What $\mathcal{G}$ bounds is degree of freedom (4)). Proposition 4.1 bears directly on degree of freedom (4), state dependence, in Section 2.5. π may be made to depend on ω only within $\mathcal{G}$; distinctions among ω that lie outside $\mathcal{G}$ cannot be written into a contract.

So $\mathcal{G}$ bounds the codomain of Φ , not its domain. Any delivery δ may be chosen; what is constrained is how π may be made to move with it.

The constraint appears in the assignment order of Part II. The family decided by degree of freedom (4) is family 5 (outcome- and state-contingent), and that family cannot be implemented unless the y in $\pi = h(y)$ is $\mathcal{G}$ -measurable. The impossibility result for family 7 in Chapter 10 (Proposition 10.6) has the same shape.

Remark 4.3 (Endogeneity of $\mathcal{G}$). This chapter treats $\mathcal{G}$ as given. In practice **an operator can invest in $\mathcal{G}$** : follow-up surveys make the outcome y observable and widen the range of verifiable information.

The staffing-service disclosures treated in Section 15.3.5 include a field for the number of placements whose separation could not be ascertained, which shows that the breadth of $\mathcal{G}$ differs by operator and is itself a choice.

Proposition 4.1 should therefore be read as a short-run condition with $\mathcal{G}$ held fixed. In the long run $\mathcal{G}$ becomes part of Φ .

Remark 4.4 (A distinction of direction). Risk transfer comes in two kinds that run in opposite directions. Success fees and royalties leave risk with the agent *because effort is unobservable*. Cost plus returns risk to the principal *because effort is observable but the outcome is uncertain*. Grouping both under “risk sharing” erases the difference.

Example 4.5 (Optimal strength of outcome contingency). Consider the linear contract $\pi = \alpha + by$. Writing γ for the agent’s absolute risk aversion and σ^2 for the noise variance of the outcome (r is reserved for the rate of sales in Chapter 5, hence γ here), the optimal outcome-contingency coefficient takes the form

$$b^* = \frac{1}{1 + \gamma \sigma^2 k} \quad (4.1)$$

where k is a constant relating to the marginal cost of effort ([18]). The larger σ^2 , the smaller b^* : **the noisier the outcome, the weaker the outcome contingency should be.**

As a numerical illustration, with $\gamma k = 1$: $b^* = 0.67$ at $\sigma^2 = 0.5$, $b^* = 0.33$ at $\sigma^2 = 2$, and $b^* = 0.10$ at $\sigma^2 = 9$.

4.2 The objective: a cash constraint imposed path by path

The $M(t, \omega)$ in the constraint is the cash balance of Proposition 3.15.

Definition 4.6 (The optimization problem). With period profit ϕ_t and discount factor β ,

$$\max_{\Phi} \quad \mathbb{E} \left[\sum_{t \in T} \beta^t \phi_t \right] \quad (4.2)$$

$$\text{s.t.} \quad M(t, \omega) \geq 0 \quad \forall t \in T, \forall \omega \in \Omega \quad (4.3)$$

The character of (4.3) is the centre of this text.

Proposition 4.7 (The path constraint does not commute with expectation). $\mathbb{E}[M(t)] \geq 0$ does not imply (4.3). Moreover the insolvency time $T_{\text{ins}} = \inf\{t : M(t) < 0\}$ is absorbing, and no ϕ_t for $t > T_{\text{ins}}$ is realized. Because τ_i denotes the settlement lag, the insolvency time is written T_{ins} .

Proof. A counterexample suffices for the first claim. If $M(t)$ takes +100 with probability 1/2 and −100 with probability 1/2, then $\mathbb{E}[M(t)] = 0 \geq 0$, yet the constraint is violated on a path of probability 1/2.

The second claim holds because insolvency occurs where $M(T_{\text{ins}}) < 0$ and no further delivery δ can be performed. The objective should therefore be written $\mathbb{E}[\sum_{t < T_{\text{ins}}} \beta^t \phi_t]$, and T_{ins} depends on Φ . $\square$

Corollary 4.8 (Insolvency while profitable). Even if $\phi_t > 0$ for every t , (4.3) can be violated. Insolvency while profitable is thus a failure of the constraint (4.3), not of the objective (4.2).

Proof. By (3.9) in Chapter 5, $M = E - W - A$. A positive ϕ raises E through (3.10), but if the increase in W exceeds it, M falls. Proposition 5.9 gives the depletion time in that case. $\square$

The condition for insolvency while profitable is rewritten as $g > g^*$ in Chapter 5 (Remark 5.8).

Example 4.9 (Insolvency while profitable, numerically). A business with margin $m = 10\%$ grows from 100 million yen of annual revenue to 200 million. Its profit is 20 million yen.

If credit to customers runs at 20% of revenue and inventory at 5%, the sum of the credit positions of (3.2) and inventory is 25 million yen at 100 million of revenue and 50 million at 200 million. The increase of 25 million yen leaves in cash.

Against a profit of 20 million yen, 5 million is short. The books show a profit and the cash is short, and (4.3) can be violated.

4.3 Heterogeneity of discount factors

Proposition 4.10 (Gain from rearranging the settlement schedule). If $\beta_C < \beta_0$ (the customer discounts the future more heavily than the firm), there exists a rearrangement of the timing of π alone, leaving the delivery δ unchanged, that improves both parties' subjective valuations simultaneously.

Proof. Let the current contract call for a payment p at time 0, and consider deferring it to a payment p' at time 1. The customer's valuation changes from $-p$ to $-\beta_C p'$, so the customer

improves if $\beta_C p' < p$, that is $p' < p/\beta_C$. The firm's valuation changes from p to $\beta_0 p'$, so the firm improves if $\beta_0 p' > p$, that is $p' > p/\beta_0$. Since $\beta_C < \beta_0$ we have $p/\beta_0 < p/\beta_C$, so the interval is non-empty. Any p' in it improves both. $\square$

Example 4.11 (The gain from instalments). Suppose the customer's annualized discount rate is 20% ($\beta_C = 0.833$) and the firm's cost of funds is 5% ($\beta_0 = 0.952$). For an item with a cash price of 100,000 yen, a price p' payable in one year improves both parties for $105,000 < p' < 120,000$. Taking $p' = 112,000$: the firm receives $112,000 \times 0.952 = 106,624$ yen in present value, a gain of 6,624 yen; the customer values it subjectively at $112,000 \times 0.833 = 93,296$ yen, a gain of 6,704 yen.

Remark 4.12 (Two readings). $\beta_C < \beta_0$ can have two different causes. For a customer facing a liquidity constraint, present funds are objectively more valuable. Alternatively the cause may be cognitive, as with hyperbolic discounting ([21]). Meeting the first and exploiting the second are almost indistinguishable from outside, yet their durability (Chapter 6) is opposite. This framework cannot settle the distinction.

Chapter 5

The Growth Constraint: CCC and the Self-Financeable Growth Rate

5.1 The cash conversion cycle

Let $r(t)$ be the rate of sales and $W(t)$ the working capital of Definition 3.2. Since W is a sum of the κ_i of (3.2), it depends on Φ .

Definition 5.1 (Cash conversion cycle). In a steady state, define

$$\text{CCC} = \frac{W}{r} \quad (5.1)$$

Its dimension is time.

5.2 The fundamental inequality

Substituting (5.1) into (3.9) of Proposition 3.17, the cash balance can be written

$$M(t) = E(t) - \text{CCC} \cdot r(t) - A(t) \quad (5.2)$$

The constraint $M \geq 0$ of Section 4.2 therefore takes the following form.

Proposition 5.2 (The fundamental inequality). The cash constraint (4.3) is equivalent to

$$\text{CCC} \cdot r(t, \omega) \leq E(t) - A(t) \quad (5.3)$$

Proof. Substitute $W = \text{CCC} \cdot r$ into (3.9) and rearrange $M \geq 0$. □

Equation (5.3) reads: **working capital cannot exceed equity net of fixed assets**. The left-hand side is fixed by Φ , the right-hand side by the capital structure. The rest of this chapter, and each chapter of Part II, treats a special case of this inequality.

Corollary 5.3 (Upper bound on scale). When $\text{CCC} > 0$, the feasible rate of sales is bounded above by

$$r \leq r^{\max} = \frac{E - A}{\text{CCC}} \quad (5.4)$$

When $CCC \leq 0$ and $E \geq A$, (5.3) holds for every r .

The dimension is yen per year, matching the rate of sales. Where g^* is a bound on the growth rate, $r^{\max}$ is **a bound on the level**. The two are separate constraints and neither implies the other (Remark 5.8).

Corollary 5.4 (Without capital, only Φ with $CCC \leq 0$ can be chosen). When $E = 0$, (5.3) requires $CCC \cdot r \leq -A \leq 0$, that is $CCC \leq 0$.

Corollary 5.4 is the starting point of Chapter 8. What remains available to a party without capital is only those Φ in which settlement precedes delivery.

5.3 Growth and free cash flow

The fundamental inequality is a constraint on the level. Change is treated next.

Proposition 5.5 (Growth and free cash flow). Writing m for the margin and I for capital expenditure, and **taking CCC to be constant and independent of m** ,

$$FCF = m r - CCC \cdot \dot{r} - I \quad (5.5)$$

Proof. Free cash flow is gross margin less the increase in working capital and capital expenditure:

$$FCF = m r - \dot{W} - I.$$

By (5.1), $W = CCC \cdot r$, and if CCC is constant then $\dot{W} = CCC \cdot \dot{r}$. Substituting gives (5.5). $\square$

Remark 5.6 (FCF is the time derivative of the cash balance). Differentiate (5.2). By (3.10), $\dot{E} = \phi - \text{Div}$, and fixed assets satisfy $\dot{A} = I - d r$ (investment less depreciation). With no dividends and no external funding,

$$\dot{M} = \phi + d r - CCC \dot{r} - I = FCF$$

So FCF is not an independently introduced quantity but **the time derivative of (3.9)**. The g^* of the next section is therefore the boundary of $\dot{M} \geq 0$.

Solve (5.5) for g . The CCC of (5.1) appears in the denominator.

Corollary 5.7 (Self-financeable growth rate). Writing d for depreciation as a fraction of sales and $g = \dot{r}/r$ for the growth rate, the condition $FCF \geq 0$ is

$$g \leq g^* = \frac{m + d - I/r}{CCC} \quad (5.6)$$

When $CCC < 0$ there is no g^* : $FCF > 0$ at any growth rate.

Proof. If m is the operating margin, depreciation has already been deducted, yet depreciation involves no cash outlay. The cash-basis gross margin is therefore $(m + d)r$. Correcting (5.5) gives $FCF = (m + d)r - CCC \dot{r} - I$. Substituting $\dot{r} = g r$ and dividing by r , $FCF/r = (m + d) -$

$CCCg - I/r$. Solving $FCF \geq 0$ for g gives (5.6). When $CCC < 0$ the left-hand side is positive for every $g > 0$. $\square$

5.4 Level and change are separate constraints

Remark 5.8 (The constraint on the level and the constraint on change are independent). Corollary 5.3 follows from $M \geq 0$ and Corollary 5.7 from $\dot{M} \geq 0$. Neither implies the other. Even with $M(0) > 0$, if $g > g^*$ then M falls and (4.3) is violated in finite time.

Proposition 5.9 (Time to cash depletion). Let $I = 0$, let g be constant, and let $r(t) = r_0 e^{gt}$. If $g > g^* = (m + d)/CCC$, then starting from $M(0) = M_0 > 0$ the cash runs out at

$$T_{\text{ins}} = \frac{1}{g} \ln \left(1 + \frac{g M_0}{(CCCg - m - d) r_0} \right) \quad (5.7)$$

and $T_{\text{ins}} \rightarrow \infty$ as $g \downarrow g^*$.

Proof. By Remark 5.6, $\dot{M} = a r_0 e^{gt}$ with $a = (m + d) - CCCg$. Integrating, $M(t) = M_0 + (a r_0 / g)(e^{gt} - 1)$. Since $g > g^*$ implies $a < 0$, M falls, and solving $M(T_{\text{ins}}) = 0$ gives (5.7). As $g \rightarrow g^*$, $a \rightarrow 0$ and the argument of the logarithm diverges. $\square$

Equation (5.7) puts a value on the insolvency time that Proposition 4.7 introduced only as a symbol, and fixes when the insolvency of Corollary 4.8 actually arrives.

Example 5.10 (g^* is a threshold, not a cliff). With $m = 5\%$, $d = 2\%$, $CCC = 0.25$ years, $M_0 = 100$ and $r_0 = 1,000$, we have $g^* = 28\%$.

Growth rate	T_{ins}
$g^* + 0.001$	16.8 years
$g^* + 0.01$	8.7 years
$g^* + 0.05$	3.9 years
$g^* + 0.2$	1.4 years
$g^* + 0.5$	0.6 years

Table 5.1: Excess over g^* and the time to cash depletion.

A slight excess leaves more than a decade; a large one exhausts the cash within a year. **What matters is not whether g^* has been exceeded but by how much.**

Remark 5.11 (CCC is also an amplifier of demand variation). Equation (5.2) is linear in r with coefficient $-CCC$. If E and A are independent of r in the short run,

$$\text{Var}[M] = CCC^2 \text{Var}[r]$$

So CCC **not only fixes the bound on the level; it amplifies variation in demand into variation in cash.** The larger a sector's CCC, the more readily (4.3) is violated for the same variation in demand.

Because no process is specified for r , no $\mathbb{P}(T_{\text{ins}} \leq t)$ is given here. That the amplification coefficient is CCC can be stated without specifying the process.

Remark 5.12 (No assumption that $I \approx dr$). With a constant asset base, $I = dr$ and (5.6) reduces to m/CCC . That assumption is not made here. The measurements of Section 14.5 put I/r between -1.8% and 6.7% across industries and size classes, which does not agree with d .

Remark 5.13 (Divergence when CCC is small). Equation (5.6) diverges as $\text{CCC} \rightarrow 0$. Even for positive CCC, if it is only a few days then a small movement in the numerator moves g^* a great deal.

At $\text{CCC} = 5$ days, for instance, a one-point change in the numerator moves g^* by about 73 points. **In sectors where CCC is small, g^* does not work as an indicator.** Section 14.5 shows this by reporting standard deviations alongside.

g^* is the rate of growth that can be financed internally and is in substance identical to the self-financeable growth rate of [4]. The relation between FCF and the growth rate is shown in Figure 5.1.




Figure 5.1: Growth rate and free cash flow. The vertical intercept is the numerator of (5.6); the sign of CCC fixes the sign of the slope.

Example 5.14 (The magnitudes g^* actually takes). Consider a business with operating margin $m = 5\%$, depreciation $d = 2.5\%$ of sales, and capital expenditure $I/r = 3.5\%$ — close to the all-industry figures measured in Section 14.5. The numerator is

$$m + d - I/r = 5.0 + 2.5 - 3.5 = 4.0\%$$

Computing g^* for several values of CCC:

CCC (days)	CCC (years)	g^*	For reference: m/CCC
120	0.329	12.2%	15.2%
90	0.247	16.2%	20.3%
60	0.164	24.3%	30.4%
30	0.082	48.7%	60.9%
−30	−0.082	does not exist	does not exist

Table5.2: g^* and m/CCC for several values of CCC (illustrative).

A business with $\text{CCC} = 120$ days necessarily needs outside funding to grow faster than 12% a year. A business with $\text{CCC} = -30$ days accumulates cash the faster it grows. **The sign fixes the qualitative constitution; the magnitude fixes the quantitative room.**

The right-hand column gives the value when d and I are ignored. Here $d < I/r$, so m/CCC overstates g^* . **For businesses investing above depreciation, an assessment using m alone is optimistic.**

Remark 5.15 (Domain of definition). Equation (5.1) presumes a steady state and is undefined immediately after founding, where $r \approx 0$. This is taken up again in Section 12.4.

5.5 CCC is a sum of lags

Substituting Proposition 3.12 into (5.1) eliminates r .

Proposition 5.16 (Structure of CCC). In a steady state,

$$\text{CCC} = \sum_i \tau_i + \tau_{\text{inv}}, \quad \tau_{\text{inv}} \equiv \frac{\text{Inv}}{r} \quad (5.8)$$

That is, CCC does not depend on the flow; it is a sum of lags.

Proof. From $W = \sum_i \kappa_i + \text{Inv}$ and (3.7), $W = r \sum_i \tau_i + \text{Inv}$. Divide both sides by r . Since (3.7) holds as a time average (Remark 3.13), this proposition is likewise a statement about averages. $\square$

Three consequences follow from (5.8). All are deductions and use no observation.

Corollary 5.17 (Source of differences in level). τ_i is the quantity specified by degree of freedom (1) of Section 2.5, the order of $\text{supp}(\pi)$ and $\text{supp}(\delta)$. The level of CCC is therefore fixed by Φ .

Corollary 5.18 (Cross-sectional comparison is meaningless). Comparing the level of CCC across objects with different Φ is comparing different contractual forms, and carries no information about demand or efficiency.

Corollary 5.19 (Aggregation is an r -weighted mean). In the parallel arrangement of Proposition 3.9, writing r_k for the rate of sales of each part,

$$\text{CCC} = \sum_k w_k \text{CCC}_k, \quad w_k = \frac{r_k}{\sum_j r_j} \quad (5.9)$$

so CCC is not additive.

Proof. From $W = \sum_k W_k$ (Proposition 3.9) and $W_k = \text{CCC}_k r_k$, $\text{CCC} = \sum_k \text{CCC}_k r_k / \sum_j r_j$. $\square$

Equation (5.9) means that computing CCC for a firm running Φ from several families in parallel yields a weighted mean of terms with opposite signs. This is the ground for taking Φ , not the firm, as the unit of observation.

Corollary 5.20 (Decomposition of change). Differencing (5.8) splits it as

$$\Delta \text{CCC} = \underbrace{\sum_i \Delta \tau_i}_{\text{change in contracts}} + \underbrace{\Delta \tau_{\text{inv}}}_{\text{change in operations}} \quad (5.10)$$

In (5.10), $\Delta \tau_i \neq 0$ means a change in contractual terms, whereas $\Delta \tau_{\text{inv}}$ arises independently of contracts.

5.6 The path constraint and the peak

Equation (4.3) is imposed for every t and every ω . Yet, as Remark 3.13 notes, $\bar{\kappa}_i = r\tau_i$ is a time average. **What binds is the peak, not the average.**

Proposition 5.21 (Working capital at the peak). With a constant rate of sales, a billing period T_b and a payment lag τ_i , the time average of $\kappa_i(t)$ is $r(\tau_i + T_b/2)$ and its maximum is $r(\tau_i + T_b)$. Equation (5.3) therefore becomes

$$r \leq \frac{E - A}{\text{CCC} + T_b/2} \quad (5.11)$$

Proof. At each billing date kT_b the deliveries of $[(k-1)T_b, kT_b)$ are invoiced and settled at $kT_b + \tau_i$. The unsettled cumulant is the sum of what has accrued since the last invoice and what has been invoiced but not paid; it is a sawtooth with minimum $r\tau_i$ just after a billing date and maximum $r(\tau_i + T_b)$ just before one. Its average over a cycle is $r(\tau_i + T_b/2)$. The peak exceeds the CCC measured as an average by $T_b/2$, so evaluating the left-hand side of (5.3) at the peak gives (5.11). $\square$

With monthly billing, $T_b/2$ is about fifteen days. **Estimating the capital required from a CCC computed as an average understates it by half a month.** The measurements in Part IV rest on period-end balances or within-period averages and do not include this difference.

5.7 Comparative statics

Every quantity in this chapter is a function of parameters. Setting the derivatives side by side shows the designer of Φ what to move and by how much.

Quantity	Derivative	Content
g^*	$\partial \ln g^* / \partial \tau_i = -1/\text{CCC}$	independent of i
g^*	$\partial g^* / \partial m = 1/\text{CCC}$	improving the margin bites harder the smaller CCC is
Eq. (5.4)	$\partial \ln r^{\max} / \partial \tau_i = -1/\text{CCC}$	the same coefficient as for g^*
Eq. (5.4)	$\partial \ln r^{\max} / \partial E = 1/(E - A)$	raising equity bites harder the larger A is
$\text{Var}[M]$	$\partial \text{Var}[M] / \partial \text{CCC} = 2\text{CCC} \text{Var}[r]$	Remark 5.11

Table 5.3: Comparative statics of the quantities in this chapter.

Proposition 5.22 (Lags are perfect substitutes at the margin). By (5.8), $\partial \ln g^* / \partial \tau_i = -1/\text{CCC}$ for every i , independently of i .

Proof. From $g^* = (m + d - I/r)/\text{CCC}$ and $\text{CCC} = \sum_j \tau_j + \tau_{\text{inv}}$ we have $\partial \text{CCC} / \partial \tau_i = 1$, hence $\partial \ln g^* / \partial \tau_i = -\partial \ln \text{CCC} / \partial \tau_i = -1/\text{CCC}$. $\square$

Collecting one day earlier and paying one day later are equivalent for g^* . By the third row of Table 5.3, the same coefficient appears for $r^{\max}$. Whether one looks at the level or at change, the design guidance for Φ is the same.

Corollary 5.18 supplies a deductive ground for the limitation stated in Remark 2.5. That the quantities here are defined as levels is a presentational choice, but **the impossibility of cross-sectional comparison is a consequence of the definitions.**

Remark 5.23 (No comparable quantity can be constructed). Corollary 5.18 states why comparison fails; it does not supply a way to compare.

The natural candidate is the deviation from the τ_i^{contract} that the contract specifies,

$$\tau_i^{\text{observed}} - \tau_i^{\text{contract}}$$

This quantity does not depend on Φ , so it can be compared across industries, and it measures whether performance follows the contract.

But τ_i^{contract} is the payment term of an individual contract and does not exist in published statistics. **The theoretical problem is solved and the measurement problem takes its place** — category (iv), access constraint, in the taxonomy of Section 13.2.

Remark 5.24 (Independence of m and CCC). Corollary 5.7 gives g^* as a ratio but does not guarantee that numerator and denominator move independently.

The measurements of Section 14.5 show CCC rising 49% and m rising 91% between fiscal 2000 and fiscal 2024 — **both increasing**. If holding more working capital and earning a higher margin are positively related, a change in g^* cannot be read as a loosening of the constraint.

This text does not formalize the relation between the two. g^* is reported as the value of a ratio and given no causal interpretation.

Chapter 6

Surplus and Layers: Source, Durability, Mutability

6.1 The three-way decomposition of surplus

The surplus ϕ announced in Remark 2.3 is defined here. Period profit is decomposed by source into three terms; the decomposition is an organization contributed by this text. The third term is the cognitive surplus of (2.15) and can be positive when degree of freedom (5) of (2.8) is present. The first two are defined in Definition 6.2.

$$\phi = \phi_{\text{prod}} + \phi_{\text{barg}} + \phi_{\text{cog}} \quad (6.1)$$

Symbol	Source	Character and decay
ϕ_{prod}	advantage in marginal cost $c(\delta)$	technical advantage; does not reduce anyone else's share; eroded by competition but reproducible through investment
ϕ_{barg}	the map Φ	transfer of shares; zero-sum with $\sum_i \phi_i = 0$; constrained by the institutional layer
ϕ_{cog}	divergence between $\bar{\delta}$ and δ	eroded by learning, regulation and better interfaces, and in principle not recovered

Table6.1: The three-way decomposition of surplus ϕ , with the source and durability of each component.

Remark 6.1 (Where the return to bearing risk belongs). A party that makes a business of the transfer and pooling of Section 2.6 — an insurer or a payment processor, say — earns a return for bearing risk. Where does it belong in (6.1)?

The part by which variance is actually reduced through pooling is ϕ_{prod} . By Proposition 2.9, raising n removes the σ^2/n term; this is a technical advantage and reduces nobody else's share.

The part corresponding to the limit $\rho\sigma^2$ of (2.11), by contrast, cannot be removed. The return for bearing it is a transfer and satisfies $\sum_i \phi_i = 0$; it belongs to ϕ_{barg} .

In practice the two are not separated and appear as a single rate. Where convenient this text writes ϕ_{risk} (Section 17.5.1), which is not a new component but the undivided sum of ϕ_{prod} and ϕ_{barg} .

Definition 6.2 (Production surplus and bargaining surplus). Using the marginal cost $c(\delta)$ of Definition 2.14 and the price p^* that would obtain without market power, define

$$\phi_{\text{prod}} = (p^* - c(\delta)) \delta \quad (6.2)$$

$$\phi_{\text{barg}} = (p - p^*) \delta \quad (6.3)$$

where p is the actual price. The character of p^* is discussed in Remark 6.3.

Remark 6.3 (p^* is a counterfactual). Definition 6.2 requires p^* , which is not observed. It is a counterfactual price posited for a state without market power, and has the same structure as the standard procedure for measuring market power in industrial organization.

Separating ϕ_{prod} from ϕ_{barg} therefore requires an additional assumption identifying p^* .

One assumption suffices: **that a control is observable which is identical except for market power**. Given such a control, its price can be read as p^* and the difference attributed to ϕ_{barg} . Chapter 17 satisfies this in the form of two rates on the same platform that differ only in function. Where the assumption fails, this text goes no further than the sum $\phi_{\text{prod}} + \phi_{\text{barg}}$.

Model the decay of cognitive surplus as a function of the transaction frequency ν :

$$\dot{\phi}_{\text{cog}} = -\lambda(\nu) \phi_{\text{cog}}, \quad \lambda'(\nu) > 0 \quad (6.4)$$

Proposition 6.4 (Implication of the decay). The solution of (6.4) is $\phi_{\text{cog}}(t) = \phi_{\text{cog}}(0)e^{-\lambda(\nu)t}$, with half-life $\ln 2/\lambda(\nu)$. The higher the transaction frequency, the faster cognitive surplus disappears.

Proof. Direct integration of a linear first-order ordinary differential equation. Solving $\phi(t) = \phi(0)e^{-\lambda t} = \phi(0)/2$ gives $t = \ln 2/\lambda$. Since $\lambda'(\nu) > 0$, the half-life is decreasing in ν . $\square$

Example 6.5 (Frequency and half-life). Suppose λ is proportional to the number of transactions a year, $\lambda = 0.02\nu$.

Good	Transactions a year ν	λ	Half-life
Daily necessities	200	4.0	0.17 years
Monthly service	12	0.24	2.9 years
Annual contract (insurance, memberships)	1	0.02	34.7 years

Table6.2: Transaction frequency ν and the half-life of cognitive surplus (illustrative).

6.2 Surplus and cash

ϕ affects cash only through (3.10), and the three components do not affect it in the same way.

Proposition 6.6 (Recognizing cognitive surplus does not move cash). Suppose $p\bar{\delta}$ is received under a prepaid contract and delivery ends at $\delta < \bar{\delta}$. When the residual obligation lapses at expiry and ϕ_{cog} is recognized, the M of (3.9) does not change.

Proof. With the valuation $v(\delta) = p\delta$, the credit position before recognition is $\kappa_C = p\delta - p\bar{\delta} = -p(\bar{\delta} - \delta) = -\phi_{\text{cog}}$. On recognition κ_C becomes 0, so W rises by ϕ_{cog} . At the same time E rises by the same amount by (3.10). In (3.9) the two cancel. $\square$

Cognitive surplus generates no cash at the moment of recognition. The cash arrived when the prepayment was received, and its effect there appeared not as ϕ_{cog} but as $\kappa_C < 0$. The ϕ_{cog} of (2.15) and the undelivered balance of Chapter 3 are two cross-sections of the same transaction.

Component	When the cash moves
$\phi_{\text{prod}}, \phi_{\text{barg}}$	at settlement, lagged by τ_C
ϕ_{cog}	at contracting, not at recognition

Table6.3: The three components of (6.1) and when the cash moves.

Remark 6.7 (This agrees with the accounting treatment). Proposition 6.6 says the same thing as the transfer from contract liabilities (breakage) under revenue-recognition standards. Its significance here is the translation into the language of κ : the transfer moves both E and W in (3.9) by the same amount, so they cancel.

Corollary 6.8 (The equity that can be built from cognitive surplus is bounded). Under the decay (6.4), the total equity that can be accumulated from ϕ_{cog} is bounded by

$$\int_0^\infty \phi_{\text{cog}}(0) e^{-\lambda(\nu)s} ds = \frac{\phi_{\text{cog}}(0)}{\lambda(\nu)} \quad (6.5)$$

Since $\lambda'(\nu) > 0$, the higher the transaction frequency the lower the bound.

Substituting (6.5) into Corollary 5.3, for a business with no source of surplus other than cognitive surplus,

$$r \leq \frac{E(0) + \phi_{\text{cog}}(0)/\lambda(\nu) - A}{\text{CCC}} \quad (6.6)$$

Equation (6.6) shows that the transaction frequency fixes the upper bound on the scale a business can reach. Equation (6.4) does not stop at a discussion of half-lives; it connects to the constraint of Chapter 5. The coefficient values have no empirical basis and only illustrate the structure of the ratio. What follows is an ordering: goods bought daily retain almost no cognitive surplus, while contracts on an annual cycle retain it for a long time.

6.3 Layers: five levels ordered by mutability

The foregoing is organized into five layers ordered by mutability. The higher the layer, the more readily it is rewritten by outside forces.

Layer	Object	Mutability
Institutional	regulation, commercial custom, interface norms	changes exogenously and rewrites the four layers below at once
Cognitive	$\bar{\delta} - \delta$	eroded monotonically by learning and regulation
Credit	the allocation of κ_i	depends on bargaining power; constrained by the institutional layer
Timing	the relation between $\text{supp}(\pi)$ and $\text{supp}(\delta)$	largely determined by the physical layer
Physical	C_{prod} , marginal cost, asset intensity	technology and physical law; given in the short run

Table6.4: Five layers ordered by mutability.

This order coincides with the order of durability of surplus. A ϕ that rests on the cognitive and institutional layers is eroded continuously from two directions, legal change and customer learning. An advantage originating in the physical layer is harder to erode.

Remark 6.9 (The institutional layer does not act uniformly). “The institutional layer imposes constraints” covers at least three distinct forms.

Form	Effect	Example in this text
Entry requirement	removes the feasibility of a Φ	deposit obligations, licensing (Section 8.5)
Price regulation	hollows out an option	capped fees (Part IV)
Manipulation of competitive conditions	erodes ϕ_{barg} indirectly	litigation pressure, fee deregulation (Remark 17.4)

Table6.5: Three forms in which the institutional layer imposes constraints.

The three erode differently. The first two make a particular Φ unavailable; the third leaves Φ in place and lowers the price. **The dominant form in recent years is the third; forms that regulate price directly are in decline** (Section 17.8).

Saying that something is “eroded by the institutional layer” therefore says nothing about how fast or how surely. Chapter 17 treats a case in which fees arising from market power can be separated from fees arising from function, but what is observed there is a single reduction of the third form, not continuous decay.

Part II

The Space of Maps

Chapter 7

A Catalogue of Φ : Seven Families, Twenty-Eight Types

7.1 How the catalogue is built

Within the space generated by the six degrees of freedom of Section 2.5, the points and regions that recur are organized into seven families and listed. For each type the sign of the credit position κ , the character of settlement, and representative examples are given. The catalogue claims no exhaustiveness; it is a working classification used to select objects of study.

7.2 Family 1: one-off exchange

Forms in which one δ corresponds to one π .

No.	Type	κ / character of settlement	Examples
1-1	immediate exchange	≈ 0 / simultaneous with δ	retail, restaurants, vending
1-2	one-off advance	< 0 / precedes δ	made-to-order, weddings, tickets, crowdfunding
1-3	one-off arrears	> 0 / follows δ	B2B trade credit, invoiced sales
1-4	instalments	> 0 (long) / one δ split into many π	hire purchase, BNPL, mortgages

Table7.1: Family 1: one-off exchange.

Only 1-4 is qualitatively different in structure. The delivery δ occurs at a point in time while π is spread over a long horizon; it commercializes the difference in discount factors of Section 4.2 on its own. Its surplus is weighted towards ϕ_{barg} and ϕ_{cog} .

7.3 Family 2: continuing access

Forms in which a right $\bar{\delta}$ is sold and δ is exercised against it. The principal home of cognitive surplus.

No.	Type	κ / character of settlement	Examples
2-1	flat access	annual < 0 / monthly in arrears > 0 , $\pi = F$	SaaS, membership retail, facility use
2-2	usage	≈ 0 / $\pi = p\delta$	electricity, cloud, telecoms
2-3	two-part tariff	mixed / $\pi = F + p\delta$	mobile, dues plus goods, booking discounts
2-4	prepaid	< 0 / π first, timing of δ at the customer's discretion	gift cards, ticket books, transit cards
2-5	options and warranties	< 0 / π first, δ only on exercise	extended warranties, cancellable bookings, credit lines

Table7.2: Family 2: continuing access.

2-1 divides in two according to whether a capacity constraint is present (the proposition of Section 2.7). 2-4 and 2-5 are formally sales of options, leaving the operator short.

7.4 Family 3: renting an asset over time

No.	Type	κ / character of settlement	Examples
3-1	lease and rental	the asset comes first / proportional to time	equipment leasing, property rental
3-2	shared asset	the asset comes first / tied to utilization	lodging, car sharing, coworking
3-3	transfer plus long collection	> 0 (large) / π continues after title passes	vendor finance, equipment finance

Table7.3: Family 3: renting an asset over time.

What characterizes this family is that the dominant term is not κ but the A of (3.8). In 3-2 a capacity constraint is explicitly present, so the utilization rate decides feasibility.

7.5 Family 4: recovery through a complement

Forms that split δ and deliberately skew π .

No.	Type	κ / character of settlement	Examples
4-1	lock-in consumables	early > 0 / δ_1 below cost, recovered on δ_2	printers, coffee machines, games consoles
4-2	freemium	early > 0 / $\pi = 0$ for some, recovered from converts	storage, business chat, mobile games
4-3	hardware subsidy plus finance	> 0 / device discounted, recovered through recurring charges	mobile carriers, auto sales finance

Table7.4: Family 4: recovery through a complement.

By definition this family deliberately creates $\kappa > 0$ early on. If the payback period on customer acquisition lengthens, the whole family becomes dependent on capital markets.

7.6 Family 5: outcome- and state-contingent

Forms in which π is measurable with respect to ω . The arguments of Chapter 4 bear directly here.

No.	Type	κ / character of settlement	Examples
5-1	success fee	> 0 (large) / $\pi = h(\text{outcome})$	recruitment, M&A advisory
5-2	revenue share	> 0 / $\pi = \text{customer revenue} \times \text{rate}$	franchising, app distribution, agencies
5-3	cost plus	intermediate / $\pi = \text{cost} + \text{margin}$	some construction, consulting, public procurement
5-4	underwriting	< 0 / π taken first, δ depends on ω	life and non-life insurance, guarantees
5-5	IP licensing	mixed / royalties, zero marginal cost of δ	semiconductor IP, character licensing, patents

Table7.5: Family 5: outcome- and state-contingent.

5-1 and 5-2 transfer risk because effort is unobservable; 5-3 transfers it in the opposite direction because the outcome is uncertain (the remark in Chapter 4). Only 5-4 has a mirrored time structure within the family.

7.7 Family 6: intermediation

Forms that stand between two other parties' transaction.

No.	Type	κ / character of settlement	Examples
6-1	transaction fee	≈ 0 / volume $\times$ rate	securities, property brokerage, marketplaces
6-2	escrow	< 0 (other people's funds) / held then remitted	payment processing, lodging platforms, travel
6-3	spread	> 0 (inventory) / the gap between bid and ask	wholesale, currency exchange, market making
6-4	charging for opportunity	< 0 / charged for listing or contact, not the transaction	job advertising, search advertising

Table7.6: Family 6: intermediation.

The κ of 6-2 is not the operator's own money. It occupies the same place in (3.8) as other negative κ but differs in that it drains rapidly when volume falls. The difference between 6-1 and 6-3 is whether inventory risk is held, which is again the transfer question of Chapter 4.

7.8 Family 7: separating beneficiary from payer

The operation of moving the index i itself.

No.	Type	κ / character of settlement	Examples
7-1	advertising	mixed / $\pi = 0$ for the beneficiary, a third party pays	broadcast, search, social
7-2	cross-subsidy	mixed / one side free, recovered from the other	single-sided charging in two-sided markets
7-3	public payment	> 0 / an insurer or government pays	healthcare, long-term care, education
7-4	donation and support	< 0 / no explicit consideration	support-based charging, tipping, non-profits

Table7.7: Family 7: separating beneficiary from payer.

7-3 is special: because the payer sets the price, the design freedom of Φ is absorbed into the institutional layer. A revision of the official price changes κ and m exogenously and at the same time.

7.9 Assigning an observation to a type

To use the catalogue empirically, a rule is needed for assigning an observed Φ to a type. A rule decided after the fact allows the assignment to be chosen to fit the conclusion. The rule is fixed below.

The information needed for assignment is the degrees of freedom of Section 2.5, which coincide with the variables listed in Chapter 11.5 as those to be recorded for each Φ .

The conditions defining the families are not mutually exclusive, so the order of the tests is fixed.

Order	Test	If it applies
1	Do beneficiary and payer differ (degree of freedom (3))?	family 7
2	Does it stand between two other parties' transaction?	family 6
3	Is π measurable with respect to ω (degree of freedom (4))?	family 5
4	Is delivery split and π deliberately skewed?	family 4
5	Does an asset come first, recovered in proportion to time or utilization?	family 3
6	Is there repetition (degree of freedom (6))?	family 2; otherwise family 1

Table 7.8: Order of the family tests. Applied from the top, the first family that applies is taken.

Once the family is fixed, the number within it is specified by degree of freedom (1) (timing), (2) (form of dependence) and (5) (separation of right from exercise).

Remark 7.1 (Some assignments depend on the order). There exist Φ whose assignment changes if the order in Table 7.8 is changed.

An app-distribution fee applies a rate to the volume distributed, which fits 6-1 (transaction fee), and moves with the seller's revenue, which fits 5-2 (revenue share). Under the order above, family 6 is tested first.

There is no deductive ground for the order. It follows a policy of testing the composition of parties before the form of settlement, and is a convention. Being a convention is precisely why it must be fixed in advance.

Remark 7.2 (Handling composite contracts). As Section 7.10.2 notes, actual contracts are often composites of several types. In that case Φ is first decomposed into components and Table 7.8 applied to each.

The unit of decomposition is the range over which one correspondence between delivery and settlement closes. Selling a device and supplying consumables are separate correspondences, so

they split into two components, assigned to 1-1 and 2-2 respectively and recorded as the composite 4-1.

7.10 Limits of the catalogue

The twenty-eight types above are not exhaustive. Of three omissions considered, two can be described with existing types.

7.10.1 The Φ of labour: no family is added

An employment contract is also a map from delivery to settlement, but the direction is reversed: the worker supplies δ and the firm pays π . Because the sign convention for X in Chapter 2 makes inflows to the firm positive, all twenty-eight types presume the firm is the side supplying δ .

But **seen from the worker's side, the worker is an operator supplying δ** . The structure is the same as the solo business of Chapter 8, and the catalogue applies unchanged.

Contract form	Type	κ_L
Monthly salary in arrears	repeated 1-3, one-off arrears	> 0 (the worker extends credit)
Daily or weekly pay	1-1, immediate exchange	≈ 0
Advances and signing payments	1-2, one-off advance	< 0
Bonuses	5-1, success fee	> 0
Equity compensation	5-2, revenue share	> 0 (extreme)
Retirement payments	1-4, instalments	> 0 (very long)

Table7.9: Contract forms seen from the worker's side and the corresponding types.

No new family is needed. A worker holds a composite contract combining families 1 and 5.

Remark 7.3 (Not a special case but an object under different constraints). Table 7.9 shows that the catalogue applies once the viewpoint is reversed, but a worker cannot be regarded as a special case of a solo business. The constraints differ in both Φ and ϕ .

Constraints on Φ . Five of the six degrees of freedom of Section 2.5 are constrained by law. The Labour Standards Act requires wages to be paid in currency, directly, in full, at least once a month, and on a fixed date; timing is fixed to arrears and the paying party to the employer. State dependence is also permitted only to a limited degree, so tying wages to performance is restricted. **It is not that there is no choice, but that the choice set is small.**

Constraints on ϕ . Defining ϕ_{prod} for a worker would require a marginal cost, and none is determined. For ϕ_{cog} no counterpart is even clear. The three-way decomposition of Chapter 6 does not apply to workers as it stands.

Furthermore the insolvency time of Section 4.2 is absorbing, whereas a worker does not cease to exist upon being unable to pay. This text does not treat workers as an independent object.

Two points are worth recording. First, because monthly pay in arrears is the norm, $\kappa_L > 0$ dominates. Payroll lags are usually 15–45 days, so workers extend credit to firms on the same order as trade payables. That κ_L belongs in the third term of (3.8) should be made explicit.

Second, an employment contract is also a trade in which the stability of κ is bought at the price of assigning credit to the organization (Section 8.7). In the language of the catalogue:

The worker avoids 5-1 (success fee, high variance) and chooses 1-3 (one-off arrears, low variance), giving up the accumulation of public credit in exchange.

This has the same form as Proposition 4.10: whoever chooses stability discounts a future claim and gives it away.

7.10.2 Composite contracts: combinations of existing types

Actual B2B contracts are often composites of several types. This is **not a limit of the catalogue but the correct way to use it**. By Proposition 3.9, κ is additive over parallel Φ , so a description that lists several types is meaningful at the level of κ . But, as Remark 3.10 states, **the decomposition is not unique**. Several sets of types give the same κ , so the types are a **generating set** for Φ , not a basis. The assignment procedure is fixed in Remark 7.2.

7.10.3 Claims that circulate: outside the framework

Forms in which π itself circulates as a claim on future δ — token issuance, for instance — still cannot be handled.

The map $\Phi : \Delta \rightarrow \Pi$ of (2.8) presumes that delivery and settlement correspond between the same counterparty. When a claim is resold to a third party, **the party that receives delivery changes after the fact**. This is a problem of the domain and cannot be expressed in the setting of Chapter 2, which fixes N . An extension of the framework is required.

7.11 A real firm is a bundle of Φ

A large firm normally runs Φ from several families in parallel. Because the sign of κ differs by family, computing CCC at the level of the firm yields a mixed average of terms with opposite signs. By Corollary 5.19 that average is weighted by the rate of sales, not a simple average of the families' CCC. To restate the conclusion of Section 12.10.3: **the unit of observation should be Φ , not the firm**.

Chapter 8

A Subspace Under Constraints: Alone, Without Assets, Without Credit

8.1 The premises imposed in this chapter

Chapter 7 enumerated the whole space of Φ . This chapter imposes constraints on it and derives a subspace: which Φ are feasible under given conditions, derived from the equations of Part I.

The object is an IT service of a scale one person can run. The premises are as follows.

- (1) There is one operator, and elements requiring human coordination are kept to a minimum
- (2) No outside funding is raised (including crowdfunding)
- (3) Assets are few and the business starts with no commercial credit

This chapter is deduction, not evidence. The validity of the conclusions depends on the validity of the premises.

Theoretical quantities in this chapter

E, A, Inv	Eq. (3.8)	set to 0 by the premises
κ_i	Eq. (3.2)	its sign is the design variable
Cap	Eq. (8.2)	one person's time. not measured
$M(t)$	Eq. (3.8)	the object of the path constraint. not measured
T_{ins}	Section 4.2	bounded above as a deadline by $E(0)/c_{\text{liv}}$
$T_{\text{dev}}, T_{\text{acq}}, T_{\text{credit}}$	Eq. (8.11)	priced in Chapter 17

Table8.1: Theoretical quantities in this chapter: the symbols relevant to a constrained solo business, and their measurement status.

8.2 Translating the premises

8.2.1 The financial constraint

From premises (2) and (3), set $E(0) \approx 0$ and $A \approx 0$. For software, $\text{Inv} \approx 0$. These are the conditions of Corollary 5.4, which requires $\text{CCC} \leq 0$, that is $\sum_i \kappa_i \leq 0$. Written in gross form,

$$\sum_i \max\{-\kappa_i, 0\} \geq \sum_i \max\{\kappa_i, 0\} \quad (8.1)$$

Equation (8.1) states that **the credit drawn must always exceed the credit extended**. This is not an assumption but the special case of the fundamental inequality (5.3) with $E = A = 0$. It is also the substitution of Proposition A.15 seen at the corner $E = 0$. The assumption $A \approx 0$ agrees with data on capital expenditure by sole proprietors (Remark 14.11). Because there is no cushion of equity, the path constraint of Section 4.2 must hold at every instant.

8.2.2 The capacity constraint

Premise (1) puts a ceiling on total delivery.

$$\sum_i \delta_i(t) \leq \text{Cap}, \quad \text{Cap} = \text{one person's time} \quad (8.2)$$

Cap is the same quantity that governed the feasibility of memberships in Proposition 2.24. There it was the capacity of facilities on the supply side; here it is the operator's disposable time. Its value differs by operator and is not disclosed, so it is not measured in this text.

8.2.3 Which one binds

Equations (8.1) and (8.2) were introduced separately but are imposed together. Converting both into upper bounds on the scale r makes them comparable.

Proposition 8.1 (The binding constraint switches). By Corollary 5.3, when $\text{CCC} > 0$ the cash constraint gives $r \leq (E - A)/\text{CCC}$. Converting (8.2) into value, the capacity constraint gives $r \leq v(\text{Cap})$. Hence, with

$$\text{CCC}^\dagger = \frac{E - A}{v(\text{Cap})} \quad (8.3)$$

as the boundary, the cash constraint binds where $\text{CCC} > \text{CCC}^\dagger$ in (8.3) and the capacity constraint binds where $\text{CCC} < \text{CCC}^\dagger$.

Proof. The smaller of the two bounds binds. $(E - A)/\text{CCC} \leq v(\text{Cap})$ holds exactly when $\text{CCC} \geq (E - A)/v(\text{Cap})$. $\square$

Effort to reduce CCC increases scale only in the range $\text{CCC} > \text{CCC}^\dagger$. Below that, capacity binds and changing contractual terms does not increase r .

For a solo business with $E \approx 0$, $CCC^\dagger \approx 0$. In Part I, which treated firms, the financial constraint $M(t, \omega) \geq 0$ bound; here the binding constraint moves to capacity as soon as CCC is made negative. As customers increase, time runs out before cash does. **That this chapter runs into a capacity constraint after choosing an advance-payment form is due to this move.**

The design principles therefore come in two.

Constraint	Requirement	Meaning
Eq. (8.1)	$\kappa_C < 0$	choose a form paid in advance
Eq. (8.2)	$\partial\delta/\partial(\text{customers}) \approx 0$	choose a delivery that is replicable

Table8.2: The binding constraint and the corresponding design principle.

The region satisfying both is the feasible set (Figure 8.1).




Figure8.1: The feasible region of Φ under the premises of one operator, no assets and no credit.

8.3 What the cloud means structurally

A by-product is the following correspondence.

Before the cloud, procuring servers required capital and A was large. After the cloud $A \approx 0$, and because usage billing is in arrears, $\kappa_S < 0$ arises instead: credit extended by the infrastructure provider.

$$\text{Effect of the cloud: } A \longrightarrow \kappa_S < 0 \quad (8.4)$$

The requirement moved from the physical layer to the credit layer, and its magnitude fell by orders of magnitude — from the purchase price of a server to one month's usage fee. That one person can now start an IT business is explained by this.

8.4 Sieving the types

The 28 types of Part II are passed through the sieve of (8.1) and (8.2).

Family	Survives	Drops out	Why it drops out
1 one-off exchange	1-1, 1-2	1-3, 1-4	cannot extend credit
2 continuing access	2-1, 2-2, 2-3	2-4, 2-5	regulation / cannot absorb variance
3 renting assets	—	all	requires A
4 complement recovery	4-2 (conditionally)	4-1, 4-3	deliberately creates early $\kappa > 0$
5 outcome-contingent	5-5	5-1 to 5-4	requires credit or pooling
6 intermediation	6-1, 6-4	6-2, 6-3	regulation / inventory
7 payer separation	7-4	7-1, 7-2, 7-3	scale or credit

Table8.3: Sieving the 28 types: those that survive under the constraints of one operator, no assets and no credit.

About ten of the 28 types remain.

Remark 8.2 (Removal by the institutional layer). Two of the reasons for dropping out are institutional. Type 2-4 (prepaid) triggers a deposit obligation under the Payment Services Act once unused balances exceed a threshold, and 6-2 (escrow) requires registration as a funds-transfer business. **The institutional layer takes particular Φ away from small operators.** Part IV records cases where the institutional layer regulates price and thereby extinguishes an option (capped fees); this section is the route by which the same thing happens through entry requirements.

8.5 The credit bootstrap problem

Equation (8.1) requires $\kappa_C < 0$. But being paid in advance requires credit, and premise (3) says there is none. The contradiction is structural.

There are only three solutions.

- Make the amount small.** If the absolute amount paid in advance is negligible to the counterparty, no credit is needed. The consequence is that **many small transactions are structurally required**. This is also consistent with the path constraint of Section 4.2: dependence on one large customer reaches T_{ins} on that customer's single delay.
- Borrow credit.** If a payment platform provides escrow and a refund guarantee, the customer trusts the platform rather than the operator. Performing 6-2 oneself is prohibited

by regulation, but riding on someone else's is possible. A $\kappa > 0$ arises to the extent of the fee and the payout lag.

- (c) **Complete the delivery first.** Sell a finished product. The question of what supports $M(t)$ during development then remains.

A consequence follows from (c). Writing T_{dev} for the development period, T_{acq} for the acquisition period, and c_{liv} for the rate of living and infrastructure costs,

$$E_{\text{required}} = c_{\text{liv}} (T_{\text{dev}} + T_{\text{acq}}) \quad (8.5)$$

Under premise (3) this cannot be met, so a route is needed for injecting E from outside. Salaried income is that route. That is, **starting while holding another job is not a matter of attitude but a requirement of the constraint.** Designing κ is a question only after this valley has been crossed.

8.5.1 Cash flow of the surviving types

Annually billed SaaS is the typical case. Both $\kappa_C < 0$ (annual prepayment) and $\kappa_S < 0$ (cloud billed in arrears) are negative, so $W < 0$ and hence $\text{CCC} < 0$, and the g^* of (5.6) does not exist. Growth generates cash.

Even when financially unconstrained, however, (8.2) remains. In a design where the support load is proportional to the number of customers, the capacity constraint breaks first. In practice this, and not funding, is what stops a solo business.

The design guidance in summary:

- (1) Be paid in advance ($\kappa_C < 0$)
- (2) Keep amounts small and numbers large (no credit needed, and it lowers variance under the path constraint)
- (3) Choose a form whose delivery is not proportional to the number of customers (capacity)
- (4) Avoid Φ that the institutional layer prohibits
- (5) Secure from outside the E needed to cross the valley

Item (3) is the hardest, and failure at (3) despite satisfying (1) and (2) is likely common.

8.6 Relaxing a premise: those who hold assets

Relax the asset part of premise (3) and set $E(0) > 0$. The case in mind is people who hold accumulated salary assets after retirement.

8.6.1 The social background

The average age at founding rose from 39.7 in 1979 to 49.7 in 2012. The share of founders aged 60 and over rose from 19.3% in 1992 to 35.0% in 2012 for men, and from 7.2% to 20.3% for

women. Teikoku Databank's 2024 survey of newly incorporated companies records an all-time high average age of 48.4 for company representatives. The premise is realistic.

8.6.2 The constraint does not loosen; it changes type

The effect of $E(0) > 0$ is not simply a loosening of the financial constraint. What is decisive is **whether E is a flow or a stock**.

$$\text{Salaried: } E(t) = E(0) + (w - c_{\text{liv}})t \quad (\text{increasing if } w > c_{\text{liv}}) \quad (8.6)$$

$$\text{Retired: } E(t) = E(0) - c_{\text{liv}}t \quad (\text{monotonically decreasing}) \quad (8.7)$$

For the salaried, the $M \geq 0$ of (4.3) is a **floor**: on approaching it one can pull back, and time is an ally. For the retired it is a **deadline**. With no revenue the insolvency time is bounded above deterministically by

$$T_{\text{ins}} \leq \frac{E(0)}{c_{\text{liv}}} \quad (8.8)$$

and arrives regardless of the choice of Φ .

Proposition A.11 supplies this.

Proposition 8.3 (Preference over variance under a deadline). Under the deadline constraint (8.8), a strategy that lowers variance is not necessarily optimal. If the expected time to completion of the safe route exceeds $E(0)/c_{\text{liv}}$, that route's probability of success is nearly zero.

The advice “keep amounts small and numbers large to reduce variance” derived in Section 8.5(a) was a conclusion under a floor constraint. It can reverse when the type of constraint changes.

8.6.3 The objective bifurcates

There is a more important bifurcation.

R1 the business must support living costs; the deadline constraint (8.8) applies

R2 living is separately provided for by pension and savings

Table8.4: The bifurcation in the objective: R1 (dependent on living costs) versus R2 (self-sufficient through pension and savings).

Under R2 the business need not cover c_{liv} .

The bifurcation is also observed across age groups. Analyses of the labour economy by the Ministry of Health, Labour and Welfare show that economic reasons for working among those aged 55–64 have declined over the long run while purpose and social participation have risen.

8.6.4 The change in the objective can be expressed through parameters

The observed change in motives can be written either as a change in the *form* of the objective or as a change in its *parameters*. The former is the stronger claim.

Socioemotional selectivity theory ([3]) takes the latter. It holds that as time horizon contracts, goals shift from acquiring knowledge to emotional meaning, and that what drives the shift is not age itself but **perceived remaining time**. Experimentally inducing a limited time horizon in younger people reportedly produces preference patterns close to those of older people.

Translated into the formulation of Section 4.2, what changes is H and β .

$H \rightarrow \text{small}$	fewer terms in the sum; distant future ϕ_t do not enter the objective
$\beta \rightarrow \text{small}$	the future is weighted less; the present state is valued relatively more

Table8.5: Effect of a contracting time horizon on the objective: what changes in H and β .

In the limit of both, the objective approaches an evaluation of the present state itself rather than of accumulated future ϕ . That is what the theory means by prioritizing present-oriented emotional goals over future-oriented ones.

This formulation requires no rewriting of the objective. The framework of Section 4.2 is used unchanged.

Remark 8.4 (Why a formulation that changes the form of the objective is not adopted). R2 could be written as $\max(\text{some other objective})$ subject to $\phi \geq 0$, moving ϕ from the objective into a constraint. That is not adopted here, for two reasons.

First, the grounds are insufficient. What is observed is the distribution of stated motives, which says nothing about the *form* of the objective. Deriving a structural change from a distribution of motives is a leap.

Second, and more seriously, $\max(\text{some other objective})$ subject to $\phi \geq 0$ has no content. Unless “some other objective” is specified, the formulation can explain any behaviour and therefore predicts nothing. It has the same defect as the one noted for the measurability hypothesis in Section 15.3.5.

Written as changes in H and β , the form of the objective is preserved, and perceived remaining time even has a measurement scale (Future Time Perspective). **Where the same observation can be explained under a weaker assumption, that is the one to take.**

Remark 8.5 (The state of the theory). Socioemotional selectivity theory is itself contested. There are reports that a limited future time horizon does not necessarily contribute to emotional well-being and can predict worse well-being, and the structure of the scale itself is being reconsidered. This text does not adopt the theory; it only notes that **a route exists that explains the observation without changing the form of the objective.**

It has also been noted that founders in their fifties frequently cite securing income, while those in their sixties more often seek connection with society.




Figure8.2: Three asset trajectories. For the salaried the constraint acts as a floor, for R1 as a deadline. R2 halts at a lower bound because living costs lie outside the business.

8.6.5 Under R2 the feasible region is reversed

The upper left of Figure 8.1 was assessed as “capped”, but under R2 being capped is not a defect.

Equation (8.2) binds only where expansion of scale is required. If H is short and β small, accumulated future ϕ carries less weight in the objective and stopping at a handful of customers causes no difficulty (Section 8.6.4). A delivery that takes personal effort in fact contributes directly to present emotional goals.

Solo business without assets or credit	the feasible region is the upper right (replicable $\times$ paid in advance)
A retiree under R2	the feasible region is the upper left (labour-dependent $\times$ paid in advance)

Table8.6: Comparison of feasible regions: a solo business without assets or credit versus a retiree under R2.

Within one and the same framework, the optimal quadrant reverses. The guidance “choose a replicable delivery” was correct only when ϕ is being maximized.

8.6.6 Family 3 revives, and a reconciliation with data

In the sieve, family 3 (renting assets over time) dropped out entirely for requiring A . Introducing $E(0) > 0$ should revive it.

By field of business at founding by age in the 2014 White Paper on Small and Medium Enterprises, among those aged 60 and over the leading fields are services 39.1%, real estate and goods rental 12.2%, and scientific research and professional and technical services 11.9%, while **real estate and goods rental is only 0.2% among those under 30.**

A sixtyfold difference. The deduction that the presence of $E(0)$ decides the availability of family 3 agrees with the distribution by age. This is an after-the-fact reconciliation and shows no causation, but because the sieve was stated in advance it has value as a check.

8.6.7 The structural role of pensions

The following correspondence holds.

Device	Effect	Where the requirement moves
The cloud	converts fixed assets A into negative κ_S	physical layer $\rightarrow$ credit layer
Pensions	move living costs c_{liv} outside the business	objective $\rightarrow$ exogenous

Table8.7: Devices that remove a constraint: the cloud and pensions.

Both remove the binding constraint on a solo business. **Pensions generate the very option that R2 represents.** This is an instance of the institutional layer widening the feasible set of Φ , the opposite of the removal in the previous section.

8.6.8 Restating the path constraint

The constraint of Section 4.2 has been written $M(t) \geq 0$, which is inappropriate for this group.

$$M(t, \omega) \geq M_{\text{floor}}, \quad M_{\text{floor}} \gg 0 \quad (8.9)$$

$E(0)$ cannot be re-earned through labour. Even for an identical loss, the asymmetry of utility differs from that of the working-age population. In practice the design is to carve out in advance a business budget B that may be spent and leave $E(0) - B$ untouched. The R2 trajectory of Figure 8.2 halts at a lower bound where this discipline operates.

Note that support services and information aimed at this group face an incentive to have E spent beyond B . Fixing M_{floor} in advance is not a modelling device but a practical defence.

8.7 The formation of credit

8.7.1 Why credit attributable to an individual is zero

Independently of E , consider whether a retiree can bring credit with them. The conclusion is $S_{\text{ind}} \approx 0$, and this is not a rule of thumb but a consequence of the measurability condition of Chapter 4.

Credit is the ground on which a counterparty accepts $\kappa_C < 0$, that is, a verifiable record of performance, and it must be in $\mathcal{G}$. But **an individual's contribution inside an organization is by definition not measurable from outside the organization**. This is no accident: it is precisely because it is not measurable that a principal–agent problem exists inside the organization.

$$\begin{aligned} & \text{the organization is observed as the performing party} \\ \Rightarrow & \text{the record of performance attaches to the organization} \\ \Rightarrow & \text{the individual's contribution is outside } \mathcal{G} \end{aligned} \tag{8.10}$$

An organization is a device that absorbs measurability from outside, and losing that on departure is the natural consequence of having enjoyed it.

Remark 8.6 (Not a sale but original attribution). Saying that credit is “continuously sold” during employment is inaccurate. A sale implies transferring something once held.

Just as the author of a work made for hire is the legal person from the outset, credit **attaches originally to the organization at the moment it arises**. It does not arise with the individual and then transfer.

The difference reaches practice. If it were a sale, one could choose to stop selling; with original attribution, **the terms of attribution themselves must be changed**. Publishing under one's own name while employed is an operation that changes where credit attaches.

8.7.2 Separating three quantities

What becomes discontinuous on leaving employment is settled once the quantities are separated.

Quantity	Content	On leaving
Usable credit	the degree to which counterparties actually accept advance payment	discontinuous downwards
Held credit S_{ind}	what is in $\mathcal{G}$ under one's own name	continuous, remains near zero
Convertible stock	personal networks, technical knowledge	continuous, but decays thereafter

Table8.8: Three quantities that behave differently around departure.

While employed, an individual *borrow*s the organization's credit rather than holding it. What falls on departure is the return of the borrowed part; the held part was zero to begin with and does not change. **It looks discontinuous, but what is lost was never one's own.**

Remark 8.7 (Leaving employment is an act of measurement). While employed, the borrowed and the held parts are bundled and cannot be separated. Leaving is an act of measurement that makes visible a decomposition previously outside $\mathcal{G}$. Nobody can tell how much credit was one's own until the organization is removed. That it cannot be estimated in advance is not a lack of information but the fact that it has, in principle, never been measured.

8.7.3 Two discontinuities, and a lag between them

Not everything falls on departure. During employment, non-compete and confidentiality obligations restrict publication, and leaving removes them: **the freedom to convert is discontinuous upwards**. But non-compete obligations run for a period and in most cases do not lift immediately.




Figure 8.3: The three quantities around departure. The two discontinuities point in opposite directions and are separated in time.

Remark 8.8 (What is at a maximum is not credit). The statement “one’s credit assets are at a maximum just after leaving” is wrong. Credit is zero, and remains zero after leaving.

What is at a maximum is the **convertible stock**: contacts and technical knowledge. That is not credit but material that can be converted into credit through the operation of publishing. Properly stated, what is at a maximum is not credit but the stock of convertibility, and it decays thereafter as $e^{-\lambda t}$.

The conclusion — move early — is unchanged, but the reason differs. It is not that a held asset is depleting but that **the material for turning something into an asset is depleting**.

The shaded region of Figure 8.3 is the structural problem. The non-compete period coincides with the period in which the stock decays fastest, and conversion is impossible during it. The higher the decay rate λ in a field, the greater the loss.

The only remedy for this lag is to produce results under one’s own name while still employed. Structurally, remedies after leaving arrive too late.

8.7.4 Capital and credit are substitutes

Restating (8.1): when $E \approx 0$, one has no choice but to earn $\sum \max\{-\kappa_i, 0\}$, and **credit is the only means of funding**. If $E > 0$, the inequality holds even without that term.

This is the substitution of Proposition A.15 seen at this chapter’s corner $E \approx 0$. Hence **a retiree’s advantage is capital, not credit**. In the dimension of credit they start from zero, like a young founder without assets.

Remark 8.9 (Confusing experience with credit). Popular accounts of founding a business later in life repeatedly cite experience, skills and contacts. But **experience and credit are different things**. Experience lowers C_{prod} ; what moves the sign of κ is credit. Experience that cannot be verified from outside does not obtain advance payment.

Confusing the two obscures the correct strategy of using capital in place of credit, and leads instead to the wrong question: “why is there no work despite the experience?”

8.7.5 Starting with the sign reversed

The guidance of Section 8.5 required $\kappa_C < 0$, which cannot be executed initially under the premise of zero credit. The following trajectory is taken instead.

Early	accept $\kappa \geq 0$ (allow payment in arrears and bridge with capital)
Middle	publish and accumulate each engagement in a form that enters $\mathcal{G}$
Late	shift to $\kappa < 0$ (advance payment becomes obtainable)

Table8.9: The trajectory from zero credit to advance payment, in three stages.

The role of capital is to buy the period in which credit is built. Rewriting (8.5),

$$E_{\text{required}} = c_{\text{liv}} (T_{\text{dev}} + T_{\text{acq}} + T_{\text{credit}}) \quad (8.11)$$

adds the term T_{credit} . Without assets this term cannot be covered, which is why credit had to be borrowed from a platform as in Section 8.5(b). With capital it can be built without borrowing. That is the concrete content of the substitution.

8.8 The trade-off between contract work and publishing

8.8.1 The conflict is not peculiar to individuals; the inability to parallelize is

Contract work is the efficient way to obtain assets; publishing results is what is needed to obtain credit efficiently. Organizations face the same conflict: development contractors and consulting firms alike find it hard to accumulate a public record because of confidentiality.

The difference lies in the form of the capacity constraint.

$$\text{Organization: } \sum \delta \leq \text{Cap}, \quad \text{Cap proportional to headcount} \quad (8.12)$$

$$\text{Individual: } \delta_{\text{contract}} + \delta_{\text{public}} \leq \text{Cap}, \quad \text{Cap fixed} \quad (8.13)$$

An organization can run both in parallel and cross-subsidize. It performs family 7-2 (cross-subsidy) not between external customers but between internal activities. An individual has no such means. **What is peculiar to the individual is not the conflict but the inability to parallelize.**

8.8.2 Two kinds of credit

“Contract work generates no credit” is inaccurate. There are two kinds of credit.

	Referral credit	Public credit
Where it resides	$\mathcal{G}_{\text{local}}$ (the customer’s private information)	$\mathcal{G}$ (verifiable public information)
Cost of accumulating	simultaneous with the work, no extra cost	consumes Cap separately
Growth	linear, slow	can be non-linear
Reach	only within the network	reaches unknown counterparties
Effect on $\kappa < 0$	strong	weak but broad

Table8.10: Comparison of the two kinds of credit: referral and public.

Contract work produces referral credit as a by-product at no extra cost, and its effect on $\kappa < 0$ is in fact more direct than that of public credit.

8.8.3 Whether the conflict arises depends on the quadrant aimed at

Aiming at the upper left	a handful of customers suffices and referral credit is enough; no conflict
Aiming at the upper right	one must reach many unknown counterparties, for which referral credit is in principle insufficient; the conflict bites

Table8.11: Whether the conflict arises depends on the quadrant aimed at.

By the conclusion of Section 8.6, the optimum for a retiree under R2 was the upper left. **For that group, no conflict between contract work and publishing arises.** The conflict becomes a structural problem for those who must maximize ϕ , that is, those aiming at the upper right.

8.8.4 Partial conversion and the optimal allocation

Even when aiming at the upper right, the substitution is not one for one. Writing Cap_{cl} for the part of the Cap of (8.2) devoted to contract work and $\text{Cap} - \text{Cap}_{\text{cl}}$ for the part devoted to publishing,

$$\text{accumulation of public credit} = \gamma (\text{Cap} - \text{Cap}_{\text{cl}}) + \theta \text{Cap}_{\text{cl}}, \quad 0 < \theta < \gamma \quad (8.14)$$

The second term is the leakage. Of the problems solved on a client engagement, **the part that generalizes can be published.** The engagement itself cannot be written about, but the general

technical problem encountered in it can. And because the thinking is already done, the marginal cost of writing is low.

Equation (8.14) is linear in Cap_{cl} and, since $0 < \theta < \gamma$, strictly decreasing. On its own this makes the optimal allocation $\text{Cap}_{\text{cl}} = 0$, that is, no contract work at all. The allocation is settled at an interior value only because contract work brings in cash.

Proposition 8.10 (Contract work stays at the minimum the constraint requires). Let w be the hourly return on contract work and c_{liv} the rate of living costs. The cash required by the fundamental inequality (5.3) is $w \text{Cap}_{\text{cl}} \geq c_{\text{liv}}$, and the allocation maximizing (8.14) is

$$\text{Cap}_{\text{cl}}^* = \frac{c_{\text{liv}}}{w} \quad (8.15)$$

Proof. The objective is strictly decreasing in Cap_{cl} and the feasible set is $\text{Cap}_{\text{cl}} \geq c_{\text{liv}}/w$, so the smallest feasible Cap_{cl} is optimal. $\square$

The optimum is not an interior point of the objective but the edge of the constraint. If $E(0) > 0$ the constraint loosens to $w \text{Cap}_{\text{cl}} \geq c_{\text{liv}} - E(0)/T$ and Cap_{cl}^* falls. The statement in Section 8.6 that capital buys the period in which credit is built can be written as a fall in the right-hand side of (8.15).

Separately from the allocation, there is room to raise θ . **The criterion for choosing engagements moves from the fee to θ .** At equal rates, choose the engagement containing a generalizable problem. The wall of confidentiality is cleared at the level of abstraction, not at the level of the engagement.

8.9 Employment contracts reread

From this chapter's argument, an employment contract can be written as follows.

$$\begin{aligned} &\text{the individual supplies } \delta \rightarrow \text{credit attaches to the organization} \\ &\rightarrow \text{the individual receives stability of } \kappa_L \end{aligned} \quad (8.16)$$

An employment contract is a trade in which stability of κ is bought at the price of placing the attribution of credit with the organization. In exchange for a certain monthly receipt, the credit asset accumulated does not become one's own.

This has the same form as the argument about discount factors in Section 4.2, and the same structure as the choice between families 5-1 and 1-3 discussed in Chapter 7.

The terms of attribution are a negotiable variable. Securing what may be produced under one's own name while employed — public talks, contributions to public repositories, writing — belongs to the same dimension as negotiating pay. In many cases it is granted at no cost.

8.10 Limits of this chapter

- (1) **This is deduction, not evidence.** Apart from the reconciliation of family 3's revival with the distribution by age, the conclusions of this chapter are deductions from premises. If the premises are wrong, so are the conclusions.
- (2) $S_{\text{ind}} \approx 0$ **depends on the field.** In fields that are small enough that participants know one another as individuals — academic disciplines, particular technical communities, small professions — an individual's performance is observed without passing through an organization. There the separation of the borrowed from the held part is already advanced during employment, and the discontinuity on departure is small. The size of S_{ind} depends not only on individual effort but on **the measurability structure of the field.**
- (3) **The objective under R2 is not identified.** It was shown that the change can be written through H and β , but not verified that this is sufficient. Strictly, therefore, no claim of optimality under R2 has been made.
- (4) **The decay rate λ of the convertible stock is not measured.** How serious the lag of Figure 8.3 is in practice depends on λ , and there is no basis here for its value.

Chapter 9

Unmanned Operation: The Limit With the Operator Removed

9.1 The question

Chapter 8 treated a solo business. This chapter asks about its limit: whether a business with no operator at all can stand. Since it is obvious that launching requires people, the question of unmanned operation is confined to steady running, after which the launch cost itself is examined to see how far it can be reduced.

Theoretical quantities in this chapter

$\Omega, \hat{\Omega}$	Chapter 4	the state space, and the subset envisaged at design time
λ_{novel}	this chapter	the arrival rate of $\Omega \setminus \hat{\Omega}$
$\mathcal{G}$	Chapter 4	automation demands a stronger condition than this
$\kappa_C > 0$	Eq. (3.2)	credit issued by the operator as a free period

Table9.1: Theoretical quantities in this chapter.

9.2 Formalizing “unmanned”

9.2.1 Automation is a stronger condition than measurability

Chapter 4 made $\mathcal{G}$ -measurability the condition for a contract to be implementable. Automation is stronger.

Definition 9.1 (Automation). Writing the response down *in advance* as a measurable function $a : \Omega \rightarrow \mathcal{A}$ on the state space. $\mathcal{A}$ is the set of available responses (a different quantity from fixed assets A , distinguished by typeface).

Proposition 4.1 required π to be $\mathcal{G}$ -measurable. Automation adds the requirement that the response be written down *in advance*. A contract need only be verifiable after the fact; automation

does not have that luxury. An algorithm is nothing other than a measurable function specified in advance. Thus

$$\text{implementable} : \pi \text{ is } \mathcal{G}\text{-measurable} \quad \subsetneq \quad \text{automatable} : a \text{ specified in advance on } \hat{\Omega} \quad (9.1)$$

The difficulty is that Ω is not closed. Writing $\hat{\Omega} \subset \Omega$ for the state space envisaged at design time, a is undefined when $\omega \in \Omega \setminus \hat{\Omega}$ arrives. Legal change, a discovered vulnerability, a payment provider's change of terms, and unforeseen use are such cases.

Proposition 9.2 (Unmanned operation is a function of duration). Writing λ_{novel} for the arrival rate of unforeseen states, the period over which operation can be unmanned is about $1/\lambda_{\text{novel}}$. Fully unmanned operation holds for a finite period but not indefinitely.

Remark 9.3 (An upper bound on λ_{novel} can be measured). In the “inside” partition of Chapter 10, *responses* to unforeseen states leave a record. Modification runs at 0.048 and cessation at 0.030 a year, and arrivals requiring no response cannot be counted, so these are a lower bound on λ_{novel} (Section 18.3, Remark 18.6). The $1/\lambda_{\text{novel}}$ of this proposition is therefore shorter than 20 years.

The distribution is bimodal. Those that are touched are touched with a median of 42 days; those that are not go untouched for a median of 4.6 years. **A mean $1/\lambda$ conceals this bimodality.**

9.2.2 What is required is not labour but availability

Responding to $\Omega \setminus \hat{\Omega}$ is not steady labour. Normally nothing is done, and action occurs only on arrival. The structure is the same as family 2-5 (options and warranties) of Part II.

$$\text{the operator of an unmanned business is short an option on } \Omega \quad (9.2)$$

Delivery δ in normal times can be made zero, but **availability cannot**. Labour can be reduced without limit; standing by remains. What is called “unmanned” is unmanned in labour as a flow; in availability as an option it is manned.

This is a different kind of constraint from the capacity constraint of Chapter 8. Equation (8.2) bounds the *total* of delivery; availability bounds **the length of an interval over which one cannot respond**.

Definition 9.4 (Response deadline and the availability constraint). Write T_{resp} for the time allowed between the arrival of $\omega \in \Omega \setminus \hat{\Omega}$ and the response, and call it the *response deadline*. Writing u for the length of an interval over which the operator cannot respond, the availability constraint is

$$\sup u \leq T_{\text{resp}} \quad (9.3)$$

Equations (8.2) and (9.3) both have the dimension of time, but they bind different things: the former acts on an integral, the latter on a supremum. **Hence no reduction in the amount of labour loosens (9.3).** This offers an account of where the practical burden of a solo business lies.

Imposing both together makes the content of that burden more concrete.

Proposition 9.5 (Availability restricts the granularity of capacity). Under (9.3), the time available for delivery is fragmented into intervals of length at most T_{resp} . A delivery requiring an uninterrupted stretch longer than T_{resp} therefore cannot be performed even when the Cap of (8.2) has room to spare.

Proof. $\sup u \leq T_{\text{resp}}$ requires that no interval without response exceed T_{resp} . An interval devoted to delivery is an interval without response, so its length is also at most T_{resp} . A delivery requiring an uninterrupted stretch longer than T_{resp} does not fit into such an interval. $\square$

The availability constraint does not reduce capacity; it restricts its granularity. The T_{dev} of Chapter 8 requires long uninterrupted stretches, whereas operation is a repetition of short responses. Proposition 9.5 derives the difficulty of combining development with operation in a solo business as the intersection of two constraints.

Remark 9.6 (They loosen in different directions). Equation (8.2) loosens by adding people. Equation (9.3) does not loosen without a relief operator, and so long as one is alone the only route is to lengthen T_{resp} itself. A longer response deadline is obtained by contractual agreement or by shrinking $\Omega \setminus \hat{\Omega}$.

Remark 9.7 (The level of T_{resp} differs by component). T_{resp} depends on the content of $\Omega \setminus \hat{\Omega}$. Chapter 17 decomposes λ_{novel} into four components, and the grace allowed for a published vulnerability differs by orders of magnitude from that for a change in the law. Evaluating (9.3) in practice requires T_{resp} by component, which is not measured here.

Remark 9.8 (Correspondence with Knight’s distinction). Variation over $\hat{\Omega}$ has a measure defined on it and is therefore risk; both automation and insurance (family 5-4) are possible. $\Omega \setminus \hat{\Omega}$ is uncertainty, and neither is.

Remark 9.9 (A degree of unmannedness cannot be defined). “How unmanned is it?” cannot be expressed as a ratio.

Measuring on the delivery side does not work. δ carries monetary units through the valuation map v , but v measures value, not effort. The ratio of the value of automated to manual delivery does not express a degree of unmannedness.

Measuring on the response side does not work either. Automation means that a is specified in advance on $\hat{\Omega}$, so a degree would be the ratio of the sizes of $\hat{\Omega}$ and Ω . But by Remark 9.8 no measure is defined on $\Omega \setminus \hat{\Omega}$.

The obstacle is therefore not the absence of a unit but the absence of a measure, and sharpening the definition does not remove it. Unmanned operation can only be stated as a function of duration, not of degree (Proposition 9.2).

Remark 9.10 (Unmanned operation and profit). If Φ and C_{prod} can be fully specified in advance, they are fully replicable. Competition therefore presses $\phi_{\text{prod}} \rightarrow 0$. The closer to fully unmanned, the thinner the source of profit.

One may read profit as the return to the residual that cannot be automated. Network effects and brand create exceptions, so this cannot be asserted flatly, but the mechanism is plausibly at work.

9.2.3 Where a free period fits

A free trial is a state in which δ is supplied with $\pi = 0$, that is, the operator extends credit to the customer, $\kappa > 0$. It works because **the marginal cost is zero, so the cost of extending that credit is near zero**. A free trial of a physical good loses inventory; a free period for software costs only infrastructure.

$$\text{free trial} = \text{credit that only a zero-marginal-cost business can issue cheaply} \quad (9.4)$$

Section 8.5 listed three solutions to the credit bootstrap problem, and this is a fourth: **extend the credit oneself in place of having credit**. T_{credit} is being purchased at the cost of $\kappa > 0$.

9.3 Decomposing the bootstrap

9.3.1 Each term and where it can be externalized

For each term of (8.11), who can absorb it. The period required for registration with the institutional layer is added as T_{inst} .

Term	Content	Externalized to	Price
T_{dev}	building the delivery system	off-the-shelf platforms, payment SDKs, hosting	usage fees
T_{inst}	registration with the institutional layer	merchant screening by a payment processor, merchant of record	fees
T_{credit}	accumulating credit	a platform's refund guarantee; or self-funded through a free period	fees, $\kappa > 0$
T_{acq}	acquiring customers	distribution through stores and marketplaces	sales commission
Enumerating $\hat{\Omega}$	designing the envisaged states	not possible	—
Choosing Φ	deciding the contract form	not possible	—

Table9.2: The terms of the bootstrap cost, where each can be externalized, and at what price.

Proposition 9.11 (The irreducible residue). The only two things that cannot be externalized are the choice of Φ and the enumeration of $\hat{\Omega}$.

Externalization has a consistent price: **permanently giving up** ϕ_{barg} . A one-off launch cost is converted into a perpetual cost in the form of fees.

9.3.2 The dominant term by family

Among the types that survived in Chapter 8, the dominant component of the launch cost differs.

Type	T_{dev}	T_{credit}	T_{acq}	Dominant
1-1 one-off digital goods	high	low	low	T_{dev}
2-1 flat access	high	medium	medium	variance
6-4 charging for opportunity	medium	low	high	T_{acq} (two-sided)
7-4 donation and support	low	high	medium	T_{credit}
5-5 IP licensing	medium	high	medium	T_{credit}

Table9.3: The dominant component of launch cost for each surviving type.

Noting that the dominant components of 1-1 and 7-4 are complementary suggests a design that splits the launch cost as a transition across families.

Stage	Content
Publishing	accumulate T_{credit} . No revenue. T_{dev} advances as a by-product
7-4 support	the first π stands on accumulated credit alone; demands almost no T_{dev}
1-1 one-off sale	the launch is light because T_{credit} is done. $\kappa \approx 0$
2-1 flat access	$\kappa < 0$, and growth generates cash in the sense of Chapter 5

Table9.4: Splitting the launch cost through a transition across families.

The leakage term θ of Chapter 8 ((8.14)) here works in the opposite direction, as a contribution from publishing to T_{dev} . **The launch cost need not be paid all at once; it can be split by moving across families.**

9.4 Limits of this chapter

- (1) **There is no quantity expressing a degree of unmannedness.** As Remark 9.9 states, the obstacle is the absence of a measure, not of a unit, and sharpening the definition does not remove it. This chapter states only what concerns duration.
- (2) **No level is given for T_{resp} .** Equation (9.3) can be written, but its right-hand side differs by component of $\Omega \setminus \hat{\Omega}$ and no value is identified (Remark 9.7).

- (3) **The inverse relation between unmannedness and profit is only a mechanism noted.** No verification, including the treatment of counterexamples such as network effects and brand, has been carried out.
- (4) **No price is given for externalizing the launch cost.** Table 9.2 lists where each term can be externalized but gives no numbers. Chapter 17 measures them.

Chapter 10

The Supplier of the Information Structure: Settlement by Protocol

10.1 The premises imposed in this chapter

Chapter 8 imposed constraints on resources and Chapter 9 took their limit. Both take the information structure $\mathcal{G}$ of Chapter 4 as given. This chapter varies $\mathcal{G}$ itself and asks **how the set of feasible Φ changes when the supplier of $\mathcal{G}$ is replaced.**

Two things should be said first.

First, **this chapter is deduction, not evidence.** The validity of the conclusions depends on the validity of the premises.

Second, **nothing here is new.** Every relation stated has a counterpart in the existing literature, and a correspondence table is given in Appendix C. The purpose of the chapter is to show how the theory built in Part I works under specific conditions.

What is replaced is the $\mathcal{G}$ of Chapter 4; the question is who guarantees the measurability condition of Proposition 4.1.

Definition 10.1 (Protocol). A set of rules published in advance, whose execution every participant can verify independently, is called a *protocol* and written $\mathcal{P}$. The tuple of variables $\mathcal{P}$ holds is called its *state*. State transitions are taken to be atomic: no part of a transition is applied on its own.

The premises are as follows.

- (1) Settlement π is expressed as a state transition on $\mathcal{P}$
- (2) The execution of $\mathcal{P}$ can be verified independently by every participant (Definition 10.1)
- (3) Whether delivery δ is contained in the state of $\mathcal{P}$ depends on the object

The implementation in wide use that satisfies premises (1) and (2) is settlement on a public distributed ledger, that is, a blockchain. Not everything called by that name satisfies the premises, however. Where verifiers are restricted, or where the state is held by a particular operator, (2) fails and the deductions of this chapter do not apply.

Premises (1) and (2) are a relaxation. The $\mathcal{G} = \bigcap_i \mathcal{G}_i$ of Chapter 4 is normally supplied by a third party outside the transacting parties, on whom verification and enforcement of π depended. Premises (1) and (2) make that unnecessary.

Only (3) cannot be fixed as a premise, and that is what forces this chapter's case split.

Theoretical quantities in this chapter

$\mathcal{G}$	Chapter 4	the object whose supplier is replaced
The index i of κ_i	Eq. (3.2)	becomes an identifier on the protocol. additivity is lost
$\bar{\delta} - \delta$	Definition 2.21	observable as protocol state in the “inside” partition
E	Eq. (3.8)	appears as collateral
Degrees of freedom (2), (3)	Section 2.5	the sieve cuts along these two

Table10.1: Theoretical quantities in this chapter: the symbols involved when the supplier of the information structure is replaced.

10.2 Translating the premises

10.2.1 Three functions the third party performed

The role of the supplier of $\mathcal{G}$ in a contract divides into three.

Function	Content	Usual bearer
(a) Recording	holding the state of κ_i	banks, payment processors
(b) Verification	judging whether the condition triggering π has occurred	courts, auditors
(c) Enforcement	transferring value when the condition holds	courts, banks

Table10.2: Three functions a third party performs in a contract.

By premises (1) and (2), $\mathcal{P}$ can always bear (a). It can bear (b) so long as the condition can be judged from the state of $\mathcal{P}$, and (c) so long as what is transferred is the state of $\mathcal{P}$.

10.2.2 A case split by where delivery resides

Premise (3) produces two cases.

Partition	Where δ resides	What enters $\mathcal{G}$	Functions $\mathcal{P}$ can bear
Inside	the state of $\mathcal{P}$	both δ and π	(a), (b), (c)
Outside	outside $\mathcal{P}$	only π	(a), and (b), (c) concerning π

Table10.3: The partition by where delivery resides, and the functions the protocol can bear.

The difference between the partitions bears directly on degree of freedom (1) of Section 2.5, timing, and fixes the sign of κ_i in (3.2).

Proposition 10.2 (Simultaneity). In the “inside” partition, for any Φ there is an implementation with $\kappa_i = 0$. In the “outside” partition, $\kappa_i = 0$ cannot be achieved for a Φ that involves δ .

Proof. In the “inside” partition both δ and π are state transitions of $\mathcal{P}$. By the atomicity of Definition 10.1 the two can be composed as a single transition, and D_i and P_i then move together, so the κ_i of (3.2) is identically 0.

In the “outside” partition δ is not a state of $\mathcal{P}$, so $\mathcal{P}$ has no time for δ . $\kappa_i = 0$ requires $\tau_i = 0$, but τ_i is not determined on $\mathcal{G}$. $\square$

That is, **in the “inside” partition degree of freedom (1) becomes a full design variable, while in the “outside” partition $\kappa \neq 0$ is forced.** Immediate exchange in the real world achieves simultaneity because the parties observe each other at the same moment; $\mathcal{P}$ has no such observation. This is why 1-1 (immediate exchange) disappears in the “outside” partition in Section 10.4.

Remark 10.3 (The partition is a property of Φ , not of the object). For one and the same object the partition changes with what is taken to be delivery. A Φ exchanging a protocol balance for another balance is inside; a Φ exchanging the same balance for currency outside the protocol is outside.

That a record of ownership sits on the protocol does not make the partition inside. Where a non-fungible token is tied to a physical artwork or a membership, what moves is the record and the delivery is a handover in the real world. Because δ lies outside $\mathcal{P}$, the partition is outside. That the two cannot be identified shows up as the fact that a change of record-holder does not cause a handover. For the same token, a Φ whose delivery is only the exercise of a right on the protocol is inside.

The question “is that object inside or outside?” therefore does not arise. What can be asked is “is that Φ inside or outside?” This has the same form as the conclusion of Section 12.10.3: here too the unit of observation must be Φ rather than the object.

Remark 10.4 (Intermediation in the “outside” partition). To bring a condition concerning δ into $\mathcal{P}$ in the “outside” partition, some party must write an external fact into the state. What $\mathcal{P}$ can verify is only that a prescribed signature is attached; it does not verify that the value written

is factual. That party therefore fails premise (2), and **the third party supposedly removed returns.**

The writing party may be a single entity or a quorum of several, so removal is continuous rather than discrete. This chapter does not treat that continuum and handles only the two ends (Section 10.7). A classification of implementations is in [35].

10.2.3 The binding constraint switches

In Chapter 8, which of the financial and capacity constraints binds first switches with the premises. The same happens here.

Because premises (1) and (2) widen $\mathcal{G}$, the measurability condition of Proposition 4.1 is looser in the “inside” partition than anywhere else in this text. $\mathcal{G}$ **therefore ceases to bind and something else does.** Section 10.5 identifies what.

10.3 Identifiers diverge from parties

This section is a by-product, placed where Chapter 8 recorded the correspondence $A \rightarrow \kappa_S < 0$ for the cloud.

The κ_i of (3.2) is defined per party i and summed as $\sum_i \kappa_i$ in (3.8). That summation presumes **that the index denotes a party uniquely.**

What can be identified on a protocol is an identifier. $\mathcal{P}$ cannot rule out one party holding several identifiers, because generating an identifier costs nothing. This is a known problem [32].

Remark 10.5 (Loss of additivity). Writing j for an identifier, i for a party, and $j \mapsto \iota(j)$ for the correspondence, what is observed is κ_j while what is needed is

$$\kappa_i = \sum_{j: \iota(j)=i} \kappa_j \quad (10.1)$$

Since ι is not observable, (10.1) cannot be evaluated.

This is a limit of the definitions of Chapter 2, not a consequence of this chapter’s premises. It shows that a setting which fixes N and ties the index to a party does not fit an environment where identifiers are freely created. It originates in the same place as the circulation of claims in Section 7.10.3, namely the setting of the domain. The theory is not amended.

Remark 10.5 bears directly on family 7.

Proposition 10.6 (A Φ requiring identification of the beneficiary cannot be implemented). When ι is not observable, a π conditioned on the beneficiary and the payer being different is not $\mathcal{G}$ -measurable.

Proof. The defining property of family 7 is that the party receiving δ differs from the party paying π . $\mathcal{P}$ identifies an identifier j and $\iota(j)$ is not observable, so $\iota(j_1) \neq \iota(j_2)$ cannot be judged on $\mathcal{G}$. By Proposition 4.1, a π depending on this condition cannot be written. $\square$

Proposition 10.6 does not depend on the partition: even when δ is a state of $\mathcal{P}$, the *party* receiving it cannot be identified. This is why 7-1 and 7-2 drop out in both partitions. 7-4 survives because it makes no explicit exchange and therefore needs no identification of the beneficiary.

That mechanisms presuming parties to be distinct fail when identifiers can be created freely is treated in mechanism design as false-name-proofness [38]. Proposition 10.6 is its manifestation at the level of contract form.

10.4 Sieving the types

The 28 types of Part II are passed through the sieve of the two partitions of Table 10.3.

The criterion differs from that of Chapter 8. There the cut was feasibility; here it is whether the protocol contributes. A type that receives no contribution in the “outside” partition does not disappear; enforcement simply returns to a third party. Below, a type “receives a contribution” when the condition governing π enters $\mathcal{G}$ and what is enforced is a state of $\mathcal{P}$. The judgment is made per Φ , not per type (Remark 10.3).

No.	Type	Inside	Outside	Grounds
Family 1: one-off exchange				
1-1	immediate exchange	yes	—	simultaneity is available only when δ and π are the same state transition
1-2	one-off advance	yes	yes	π comes first and does not refer to δ
1-3	one-off arrears	yes	—	outside, completion of δ cannot be judged
1-4	instalments	yes	partly	outside, collection of π works but default does not reach the underlying asset
Family 2: continuing access				
2-1	flat access	yes	yes	$\pi = F$ does not refer to δ
2-2	usage	yes	—	outside, metering δ needs the party of Remark 10.4
2-3	two-part tariff	yes	partly	outside, only the fixed part F
2-4	prepaid	yes	yes	inside, the unused balance is itself the protocol state
2-5	options and warranties	yes	—	outside, the exercise condition lies outside the protocol
Family 3: renting an asset over time				
3-1	lease and rental	yes	—	inside only where the asset itself is a protocol state
3-2	shared asset	conditionally	—	the range in which utilization can be written as a state transition is narrow
3-3	transfer plus long collection	yes	—	inside with collateral; outside it does not reach the underlying asset
Family 4: recovery through a complement				
4-1	lock-in consumables	conditionally	—	the family presumes the device lies outside the protocol, which fails inside
4-2	freemium	yes	yes	$\kappa > 0$, but it is self-issued credit and asks no enforcement of the protocol
4-3	hardware subsidy plus finance	yes	—	outside, the device lies outside the protocol

Table 10.4: Sieving the 28 types (families 1–4). “yes” marks those that receive a contribution.

No.	Type	Inside	Outside	Grounds
Family 5: outcome- and state-contingent				
5-1	success fee	conditionally	—	only where the y of $\pi = h(y)$ is a protocol quantity
5-2	revenue share	yes	—	inside, the counterparty's flow is observable on the protocol
5-3	cost plus	—	—	actual costs lie outside the protocol either way
5-4	underwriting	yes	—	inside, ω is confined to protocol events
5-5	IP licensing	yes	—	inside, where use is observable on the protocol
Family 6: intermediation				
6-1	transaction fee	yes	—	inside, completed in combination with the simultaneity of 1-1
6-2	escrow	yes	partly	outside, only one side is held; release needs the party of Remark 10.4
6-3	spread	yes	—	inside, inventory becomes a protocol balance
6-4	charging for opportunity	conditionally	—	only where contact can be defined as a protocol event
Family 7: separating beneficiary from payer				
7-1	advertising	—	—	that beneficiary and payer differ cannot be verified (Section 10.3)
7-2	cross-subsidy	—	—	as above
7-3	public payment	—	—	the payer sets the price and design freedom sits in the institutional layer
7-4	donation and support	yes	yes	no explicit exchange, so no identification of the beneficiary is required

Table 10.5: Sieving the 28 types (families 5–7). “yes” marks those that receive a contribution.

Twenty types receive a contribution in the “inside” partition and five in the “outside” partition. Where the sieve of Chapter 8 cut 28 types down to about ten, the “inside” sieve barely cuts at all.

This chapter's premises are a relaxation, not a constraint.

That 1-1 disappears in the “outside” partition follows from Proposition 10.2: if δ lies outside the protocol, the form degrades into either an advance or arrears.

Receiving a contribution and being executable are, however, different things. The advance forms (1-2, 2-4) pass in the table for the “inside” partition, but executing them requires collateral towards the counterparty and they are not open to a party without assets (Corollary 10.12, Corollary 10.11). Section 10.5 treats this.

Remark 10.7 (The axis along which the sieve cuts). The five types surviving in the “outside” partition are 1-2, 2-1, 2-4, 4-2 and 7-4, and all they share is that π is not measurable with respect to δ . Of the six degrees of freedom of Section 2.5, only (2) carries information for this sieve.

This is not a proposition. Since the “outside” partition was defined as the case where δ does not enter $\mathcal{G}$, the impossibility of writing a δ -measurable π follows from the definition. The

information lies not in the claim but in the fact that applying the sieve to 28 types individually turned out to be explained by a single degree of freedom.

Remark 10.8 (The direction in which the institutional layer acts). Types 2-4 (prepaid) and 6-2 (escrow) were the ones removed in Section 8.5 by deposit obligations and licensing. In the “inside” partition they are implemented as protocol.

One and the same type both drops out and returns, depending on the direction in which the institutional layer acts. This is an instance showing that the first form (entry requirement) of the three in Remark 6.9 does not act in only one direction.

10.5 The collateral constraint

Section 10.4 showed that $\mathcal{G}$ ceases to bind. This section identifies what does.

Consider a Φ in the “inside” partition with $\kappa_i > 0$, that is, one in which the operator extends credit. To enforce on non-payment of π , what is enforced must be a state of $\mathcal{P}$. But $\mathcal{G}$ consists only of states of $\mathcal{P}$, so **neither the counterparty’s ability to pay nor their past record of performance is in $\mathcal{G}$.**

Credit can therefore be provided for only through states locked in advance, that is, collateral. Writing B_i for the collateral one has locked for counterparty i and B'_i for the collateral i has locked for oneself, the κ_i of (3.2) satisfies

$$\max\{\kappa_i, 0\} \leq B'_i, \quad \max\{-\kappa_i, 0\} \leq B_i \quad (10.2)$$

Collateral is provided from one’s own assets without duplication, so $\sum_i B_i \leq E$, and summing (10.2) over i gives the following.

Proposition 10.9 (Collateral constraint). In the “inside” partition, the total credit that can be received does not exceed equity.

$$\sum_i \max\{-\kappa_i, 0\} \leq E \quad (10.3)$$

Proof. Summing the second inequality of (10.2) over all i gives $\sum_i \max\{-\kappa_i, 0\} \leq \sum_i B_i$. The same asset cannot be locked for several counterparties at once, so $\sum_i B_i \leq E$; combining the two gives the result. $\square$

Compare (10.3) with (8.1) of Chapter 8. There, under $E \approx 0$, the requirement was $\sum \max\{-\kappa_i, 0\} \geq \sum \max\{\kappa_i, 0\}$: the credit drawn had to exceed the credit extended. Here the credit that can be drawn is itself bounded by E .

Corollary 10.10 (The substitution degenerates to one direction). In the “inside” partition, the substitution of Proposition A.15 becomes one-directional: capital substitutes for credit, but credit does not substitute for capital.

Proof. By Proposition 10.9 the terms with $\kappa < 0$ are bounded above by E . In (3.8), terms with $\kappa < 0$ cannot raise funding beyond E . The reverse substitution — E substituting for $\kappa < 0$ — holds as in Proposition A.15. $\square$

Corollary 10.10 is the third place in this text where Proposition A.15 appears.

Place	E	Credit	What binds there
Chapter 8, solo business	0	the only means of funding	the credit bootstrap problem
Section 8.6, those with assets	positive	built from zero	capital buys the period in which credit is built
This chapter, “inside”	the only means of funding	0	there is no period in which to build

Table10.6: The three places Proposition A.15 appears. The two ends line up.

Section 8.7 described the route by which credit forms through accumulated public performance. That route does not operate in the “inside” partition, because $\mathcal{P}$ identifies only identifiers (Section 10.3) and cannot attribute a past record of performance to a party.

Corollary 10.11 (Exclusion of parties without assets). When $A = \text{Inv} = 0$ and $E = 0$, the only feasible Φ are those with $\kappa_i = 0$ for every i .

Proof. Substituting $E = 0$ into (10.3) gives $\kappa_i \geq 0$ for every i . On the other hand, by Proposition A.15, when $A = \text{Inv} = E = 0$ the requirement is $\sum_i \max\{-\kappa_i, 0\} \geq \sum_i \max\{\kappa_i, 0\}$, whose left-hand side is 0, giving $\kappa_i \leq 0$. $\square$

Corollary 10.12 (Growth through advance payment requires capital). In the “inside” partition, achieving $\text{CCC} < 0$ requires $E > 0$.

Proof. By (5.1), $\text{CCC} = W/r$ with $W = \sum_i \kappa_i + \text{Inv}$. Substituting $E = 0$ into Proposition 10.9 gives $\kappa_i \geq 0$ for every i , which together with $\text{Inv} \geq 0$ gives $W \geq 0$ and hence $\text{CCC} \geq 0$. $\square$

Chapter 8 showed a route by which a solo business without assets creates $\text{CCC} < 0$ through annual prepayment — a state in which growth generates cash. **That route closes in the “inside” partition**, because receiving an advance requires collateral towards the counterparty and collateral is provided out of capital.

That is, **what remains to a party without assets is confined to types with κ identically 0**: 1-1 (immediate exchange), 2-2 (usage) and 6-1 (transaction fee), which Part II recorded as $\kappa \approx 0$, and their composites. In Chapter 8, $E = 0$ required $\kappa_C < 0$; here it forces $\kappa = 0$. Solutions (a), (b) and (c) of Section 8.5 all presumed a shortage of credit; what is short here is capital.

Remark 10.13 (Correspondence with existing theory). Nothing in this section is new. The structure by which net worth constrains funding through collateral capacity is in [33], and the formulation in which collateral constraints govern a firm’s financing and risk management jointly is in [36, 37]. Proposition A.15 itself has counterparts there too (Appendix C).

That the contractible space widens under premises (1) and (2) is treated in [31]. The point corresponding to Corollary 10.11 — that a mechanism demanding collateral excludes parties without capital — also has a literature [34].

The role of this section is to place these on one and the same inequality, that of (3.8).

10.6 Observability

Premise (2) has a secondary consequence. Because the state of $\mathcal{P}$ is verifiable by every participant, the following quantities are individually observable in the “inside” partition.

Quantity	Measurement status in this text	Status in the “inside” partition
$M(t)$	requires per-firm data and is meaningless in aggregate	observable per identifier
κ_j	as above	as above
$\bar{\delta} - \delta$	only estimable from aggregates	directly observable as the unused balance
B	this text has no corresponding quantity	observable as locked state

Table10.7: Quantities that become observable in the “inside” partition.

What is observable, however, is per identifier, and by Remark 10.5 the per-party κ_i cannot be recovered. **The obstacle to measurement changes from the absence of aggregation to the impossibility of it.** This differs from category (v) of Section 13.2 (Remark 13.2) and is a new form relative to the taxonomy of Part III.

10.6.1 A population frame becomes available

It is not only quantities that become observable.

The conditioning on survival of Chapter 12 arises because the objects observed are limited to paths that did not violate (4.3). In the “inside” partition **the record of the protocol is itself the population**. Every Φ that operated remains, including those that stopped, so no conditioning occurs.

Chapter 17 abandoned direct measurement of the unmanned operating period for want of a population frame. That reason disappears here.

10.6.2 What can and cannot be tested

Being observable does not mean being testable. Three cases must be separated first.

Case	Content
Quantities the protocol enforces	observing them is not a test; it is reading the rules
Quantities that result from design	these are what can be tested
Per-party quantities	cannot be recovered (Remark 10.5)

Table10.8: Three cases among the quantities observable in the “inside” partition.

The collateral constraint of Proposition 10.9 belongs to the first case. Since the protocol requires collateral, observing that collateral is posted does not test the proposition. **What can be tested is confined to quantities that appear as the result of an operator’s choice.**

Remark 10.14 (The population differs). A distribution measured in the “inside” partition cannot be generalized to firms at large. In the sense of Chapter 12 it is a different population. **What can be tested here is whether a deduction is correct, not whether that deduction applies to real solo businesses.** The propositions of Chapter 8 can be tested; nothing is said about that chapter’s objects.

Chapter 18 observes these quantities in practice.

10.7 Limits of this chapter

- (1) **This is deduction.** The conclusions depend on premises (1), (2) and (3). The tests of Section 18.2 measure the implications of the deduction and do not verify the premises themselves.
- (2) **A rule was needed to narrow what is judged.** Table 7.8 takes Φ as given and cannot be applied directly to ledger identifiers. Section 18.2 places the question “is this a Φ ?” before the assignment order (Remark 18.5). That step is not part of the rule of Chapter 7.
- (3) **Only the two ends are treated.** As Remark 10.4 notes, the degree to which the third party is removed in the “outside” partition is continuous. This chapter treats only full removal and no removal, not the middle. Most existing arrangements are in the middle.
- (4) **A limit of the definitions remains.** The unobservability of ι in Section 10.3 originates in fixing N in Chapter 2. This chapter records it and does not extend the domain. It is the same place as the circulation of claims in Section 7.10.3.
- (5) **ϕ is not treated.** This chapter argues only about the feasibility of Φ and does not apply the three-way decomposition of surplus of Chapter 6. An argument of the same form as the remark in Chapter 9 — that the replicability of a protocol puts pressure on ϕ_{prod} — could hold, but judging it requires measured fees and is not attempted here.

Part III

Method

Chapter 11

Design of the Empirical Work: Connecting to Observation

11.1 What is to be tested

Part I defined quantities and propositions, and Part II deduced the subspaces that remain under constraints. Neither used observation. From this part onwards they are connected to observation. What is connected divides in two.

Object	Content	Question
Quantities of Part I	κ , CCC, g^* , ϕ , $M(t)$, $\mathcal{G}$, and others	what values do they take
Deductions of Part II	the sieve of types, substitution, unmanned operation, collateral constraints	do they hold under the premises

Table11.1: The two objects the empirical work connects to.

The first is measurement, the second verification. Because they need different preparation, they are treated in stages below.

11.2 Five stages to pass

Connecting one quantity, or one proposition, to observation passes through five stages. The organization is this text's own.

Stage	Question	What is needed if it fails
1 Operationalization	does the quantity have a unit and an object of observation	sharpening the definition
2 Population	what set is to be enumerated	reselecting the object
3 Measurement	can values be obtained at the granularity required	opening a route of acquisition
4 Reconciliation	does it agree with existing aggregates	—
5 Verification	are the implications fixed in advance supported	—

Table11.2: Five stages from theory to observation.

The first three are prerequisites in series: failing any one blocks the stages that follow. **The last two are ways of using the values obtained, not a continuation of the sequence.** They are separated by whether pre-registration has taken place.

The benefit of separating the stages is that the work required differs by where the process stalls. A gap at stage 1 is not filled by collecting any amount of data, and a gap at stage 3 is not filled by refining any number of definitions. **Without first judging where the process is stuck, effort is misallocated.**

11.3 What to check at each stage

11.3.1 Stage 1: operationalization

Three checks.

- (1) Does the quantity have a **unit**?
- (2) Is the **object of observation** fixed — the firm, a Φ , or an identifier?
- (3) For a proposition, is the **sign uniquely fixed**?

Some quantities of Part I are operationalized by their definitions and some are not. For the former, Chapter 11.5 lists the variables to be recorded for each Φ .

For the latter — quantities that are defined but whose unit or object of observation is not fixed — that has to be settled before proceeding to measurement. **A defect at this stage cannot be detected by adding data.** Values do come out, so a result appears without it being clear what was measured.

Item (3) is specific to propositions. If the theory does not fix the sign, either result is consistent with it and no test is possible.

11.3.2 Stage 2: population

Three checks.

- (1) Have the reference date, the scope and the lower bound been fixed?
- (2) What falls outside the sample, and is it correlated with the axis of interest?
- (3) Do objects that have stopped or exited remain in the sample?

Items (2) and (3) are the central theme of this part and are treated in Chapter 12. In this framework (3) weighs particularly heavily: the insolvency time of Section 4.2 depends on Φ , so an attempt to estimate Φ makes the composition of the sample itself a function of the object of estimation.

11.3.3 Stage 3: measurement

Three checks.

- (1) Is the quantity measured anywhere?
- (2) If it is, can it be obtained at the granularity required?
- (3) If it cannot, which category of obstacle applies?

The judgment at (3) governs the design of the survey, because the category determines whether spending effort produces progress or nothing at all. Since the categories were constructed through the survey work itself, they are placed in Chapter 13.

Remark 11.1 (Why this chapter does not supply the taxonomy itself). The judgment at stage 3 needs the taxonomy of obstacles, but that taxonomy was constructed through the empirical work. Placing it here would put findings not yet obtained ahead of their basis.

What this chapter supplies is the stages to check, not a classification of why a stage was not passed. The former is a design procedure, the latter a result of the survey. The two can be stated independently: the stages say where the process stalled, the taxonomy says why.

11.3.4 Stages 4 and 5: how the values obtained are used

Passing stage 3 yields values. Their use divides in two, treated in the next section.

11.4 Reconciliation versus verification

Comparing a deduction with existing aggregates after the fact is called *reconciliation*. Judging an implication fixed in advance against data is called *verification*.

	Reconciliation	Verification
When the implication is fixed	after	before
Declaring the population	not required	required
If they disagree	can be explained by another factor	the implication is rejected
Falsifiable	no	yes

Table11.3: Reconciliation versus verification.

Reconciliation cannot be falsified. If the figures disagree one can say another factor was at work, and nothing in the sample rules that explanation out. Reconciliation therefore neither supports nor refutes the theory.

It still has a use. It confirms that a deduction is not wildly at odds with reality, and it helps in choosing what to verify next. **It has value as a check and none as evidence.**

The danger lies elsewhere.

Remark 11.2 (Reconciliation can skip stages). Because reconciliation uses existing aggregates, **it works without declaring one's own population:** stage 4 can be reached without passing stages 2 and 3.

This is the route by which reconciliation and verification are confused. An implication supported after passing stages 1 to 3, and a direction that happens to agree with existing aggregates after skipping stages 2 and 3, differ by an order of magnitude in the strength of the evidence obtained. Outwardly both can be written up as “consistent with the theory”.

The comparison in Chapter 8 between the revival of family 3 and the fields of business started by age is of this form. That section states explicitly that it is an after-the-fact reconciliation and shows no causation.

Hence the following discipline.

When a reconciliation is performed, say that it is one. Do not use it in place of an implication fixed in advance.

This has the same form as the rule in Chapter 12 that well-known firms are used for illustration and not for inference. Both require that **material not be used beyond the strength of the evidence it carries.**

11.5 Where each stage is handled in this text

Stage	Where handled	Content
1 Operationalization	Part I, Chapter 11.5	the definitions, and the variables to record for each Φ
2 Population	Chapter 12	decomposition of selection bias and how to handle it
3 Measurement	Chapters 14 and 13	what could and could not be measured; the taxonomy of obstacles
4 Reconciliation	this chapter	the discipline of Section 11.4
5 Verification	Chapter 11.5, Part IV	the pre-registration procedure and the judgments on 22 implications

Table11.4: The five stages and where this text handles them.

Chapter 14 opens Part IV in order to tabulate the results of stage 3. Making explicit which quantities could not be measured is that chapter's main purpose, and that is a record of where between stages 1 and 3 the process stalled.

Remark 11.3 (The stage reached differs by quantity). The stage reached is not a single figure for the theory as a whole; it differs by quantity and by proposition. Within one chapter, quantities that reached stage 5 sit alongside quantities stalled at stage 1.

There is therefore no single answer to the question of how far the empirical work has come. The stage reached can only be stated quantity by quantity.

11.6 Priorities for the survey

The objects are narrowed to families 2 and 5. The selection can be derived deductively.

■ ϕ_{prod} does not depend on Φ . Of the three terms of (6.1), $\phi_{\text{prod}} = (p^* - c(\delta))\delta$ is fixed by marginal cost and the competitive price and contains nothing about the shape of the map from delivery to settlement. Replacing Φ alone, holding C_{prod} and marginal cost fixed, changes the sign of κ and the working capital required, not the production surplus (Table 2.5).

The shape of Φ bears on surplus through the other two terms. To test the framework, it suffices to choose Φ that isolate each of those two terms.

■ The two terms arise from different degrees of freedom. $\phi_{\text{cog}} = p \cdot \mathbb{E}[\bar{\delta} - \delta]$ can be positive only when the contract states a ceiling $\bar{\delta}$, that is, when π does not refer to δ (Remark 2.22). That is degree of freedom (5). Relating the divergence to the transaction frequency ν additionally requires the repetition of degree of freedom (6): in a one-off, $\bar{\delta} - \delta$ is a single realization with no

distribution, so no relation to ν can be asked. **The family in which both (5) and (6) are present is family 2**, which Part II designates the principal home of cognitive surplus.

$\phi_{\text{barg}} = (p - p^*)\delta$ is a zero-sum transfer whose allocation is fixed by whether the y in $\pi = h(y)$ is verifiable. That is degree of freedom (4). In the order of Table 7.8, **the only family decided by degree of freedom (4) is family 5**.

■ The two families sit diagonally in the quadrant diagram. Degrees of freedom (1) and (2) dominate, and those two axes give the four-way split of Figure 2.1. Family 2 is advance and flat (π independent of δ , divergence possible), family 5 is arrears and linked (π tied to δ , no divergence): **both axes take opposite values**. Choosing only one family would allow only one of the two axes to be varied.

■ What can be tested. Family 2 permits direct tests of cognitive surplus ((2.15)) and of the capacity proposition. In particular, that one and the same flat membership divides in two according to whether a capacity constraint is present is the clearest instance of outwardly similar Φ having different durability.

Family 5 permits a direct test of the measurability condition (Chapter 4). The claim that success fees and cost plus transfer risk in opposite directions is testable as a correspondence between contract form and observability.

■ Why they are studied together. The two families rest on different institutional foundations, so data become available in different ways. Family 2 has an exhaustive frame through registration under the Payment Services Act; family 5 has per-firm filings published under the disclosure obligation of the Employment Security Act. **The former supplies a comprehensive frame, the latter individual variables**.

Looking at one family alone, when verification fails to progress **one cannot tell whether to attribute it to a defect in the frame or a defect in the variables**. Setting side by side two families with different foundations makes it possible to separate which side is stuck. That is a further reason for taking both.

11.7 Variables to record

For each Φ , at minimum the following are recorded.

Variable	Content
$\text{sgn } \kappa$	whether settlement precedes or follows delivery
CCC	only where a steady state holds; otherwise record that it does not
Form of dependence	(F, p) of (2.10), or $h(y)$
Separation of $\bar{\delta}$ and δ	whether they separate, and an estimate of the divergence
Capacity constraint	whether present, and whether binding
ν	transaction frequency ((6.4))
Measurability	whether effort or outcome is verifiable
Paying party	the beneficiary or a third party
Estimated split of ϕ	a qualitative allocation across $\phi_{\text{prod}} / \phi_{\text{barg}} / \phi_{\text{cog}}$

Table11.5: Minimum variables to record for each Φ .

11.8 What must be pre-registered

Following the conclusions of Part III, the following are fixed and recorded before the survey begins.

- (1) The definition of the population (reference date, geographic scope, industry scope, size floor)
- (2) The threshold for treating a state as steady (the criterion for stability of r)
- (3) How failure cases are collected, and the range in which they cannot be
- (4) Which of the APC effects is assumed zero, and on what grounds
- (5) The criterion distinguishing primary sources from retrospective accounts

Deciding these after the fact lets the choices of Sections 12.2 and 12.3 contaminate the conclusions.

Chapter 12

Selection Bias: Decomposing the Distortions That Enter Observation

12.1 The problem

Applying the framework of Part I to real firms introduces systematic distortion at the very start of any comparison. The distortion is not a single “bias” but a mixture of at least eight problems of different kinds. They are decomposed below, and those that can be handled are separated from those that cannot be, in principle.

12.2 Conditioning on survival: a circularity in the framework

An observable firm is the realization of a path that did not violate the constraint (4.3). Writing the insolvency time as $\tau = \inf\{t : M(t) < 0\}$, what is observed is

$$\mathbb{P}(X \mid \tau > t) \quad (12.1)$$

and not $\mathbb{P}(X)$. And what fixes τ is precisely Φ , the object of estimation. The structure is shown in Figure 12.1.




Figure12.1: Survival S is a common consequence of Φ and ϵ — a collider. Conditioning on S induces a spurious correlation between them.

S is a common consequence of Φ and the exogenous shock ϵ , that is, a collider. Conditioning on a collider induces correlation between two causes that are independent (Berkson’s paradox).

12.2.1 The direction of the distortion reverses with what is conditioned on

The distinction matters in practice.

■ **Conditioning on survival.** Among surviving firms, those that persisted despite an unfavourable Φ are those whose ε was favourable. In a surviving sample the disadvantage of Φ therefore appears offset by an advantage in ε , and **the effect of structure is understated**. The disadvantage of $\text{CCC} > 0$ looks smaller than it is when only survivors are examined.

■ **Conditioning on success (fame).** Here the situation differs. Collecting only successful firms removes from the sample every firm that failed with the same Φ , so there is no control group. The problem is not understatement but that **estimation is impossible at all**: the sample contains no information that separates Φ from ε .

Treating both as one and the same “survivorship bias” hides this difference.

12.3 Selection into disclosure: correlated with the axis of interest

Listing and disclosure are not random selections. By (5.6), a firm with $\text{CCC} < 0$ has no g^* and can finance growth internally, so in principle it has little need to go to outside capital markets. Conversely a firm with a large positive CCC needs outside capital and is pushed towards listing.

That is, **the probability of selection is itself a function of the object of estimation**. The selection probability $\mathbb{P}(S_{\text{disc}} = 1 \mid \text{CCC})$ tends to be increasing in CCC. The distribution of CCC observed in a database of listed companies is therefore displaced from the population distribution systematically and in a predictable direction.

12.4 Age: a problem of domain, not of bias

That a firm on its first day cannot be compared with one several years old is a problem of a different kind. Equation (5.1) was derived assuming a steady state and is undefined at $r \approx 0$, where it divides by zero.

This is therefore not a bias but **a question of the framework’s range of application**. The remedy differs accordingly: a bias is corrected, whereas a question of domain can only be declared. Concretely, one of the following is taken.

- (1) Restrict the objects to a region where r is stable, and state the threshold.
- (2) Model the transient explicitly with a separate system of variables.
- (3) Lower the unit of analysis from the firm to Φ and judge maturity per Φ (Section 12.10.3).

Introducing a correction factor and forcing the steady-state formula onto the data is the remedy most to be avoided.

12.5 Age, period and cohort are not identified

Three clocks act on an observed firm at once. Writing c for the founding cohort, p for the year of observation and a for age (this section follows the notation of the Age–Period–Cohort literature; a, p, c have these meanings in this section only and are unrelated to their use elsewhere),

$$a = p - c \quad (12.2)$$

holds identically. By (12.2) the three variables are linearly dependent, and **identifying the three effects simultaneously is impossible without adding an exogenous restriction**. This has been discussed at length in demography and epidemiology as the Age–Period–Cohort problem, and no general solution exists.

Hence the question “is a given industry’s favourable CCC in a given period due to Φ , to youth, or to the financial environment of the time?” cannot be answered from data alone. All one can do is state an identifying assumption — setting one of the effects to zero, say — and show how much the conclusion depends on it.

12.6 Retrospective narrative

The description of a successful firm’s Φ is itself contaminated. The structure is retold so as to explain the success, so bias enters the *measurement* of the explanatory variable Φ . No correction on the dependent-variable side removes it.

As Part I stated, most Φ are not products of design but arise as properties of a line of business, their financial function recognized after the fact. Distinguishing primary sources — contemporaneous disclosures, contracts, reporting — from retrospective accounts is the minimum remedy.

12.7 Mistaking the unit: what a row denotes versus what its values denote

The preceding sections concerned distortions in the *selection* of the sample. Failures of a different kind occur at the stage of acquisition.

What a row denotes and what the values in that row denote need not coincide. When a published list has rows at a lower unit while the figures are reported at a higher unit, the value is replicated across every row belonging to the same higher unit. Treating the list as observations lets an object with many lower units count as many times over.

Missing values can be detected in the output; this mistake cannot. The values are complete, so it passes any post-collection check. Section 15.3.5 records a case in which it reversed the conclusion.

There is only one remedy: **confirm the unit at which the values are reported and consolidate at that unit**. Consolidation is not preprocessing; it is a step that decides the conclusion.

12.8 Selection of variables: the gaps are not random

The foregoing concerned distortions in the sample and in identification. The same kind of problem arises for *which variables are observable*.

There can be a correlation whereby the areas of greatest theoretical interest are the hardest to measure. A field in which institutions have intervened heavily is evidence that the authorities judged there to be a problem there, and theoretical interest is drawn to the same place. Yet in a field whose transaction forms are unstable enough to require intervention, that very instability can make the required quantity impossible to compute from published aggregates (Remark 15.3).

Treating gaps in the data as random systematically removes from the sample the area one most wants to know about. The other sections of this chapter treat selection of the sample; this is selection of variables, and no correction on the sample side addresses it. All one can do is describe where the gaps are concentrated.

12.9 The observation frame in practice

Some orders of magnitude for Japan.

Population	Size	Source and notes
Firms (SME Agency basis)	about 3.6 million	recompiled from the Economic Census; excludes primary industries
Firms (National Tax Agency basis)	about 6.4 million	includes about 3.79 million sole proprietors
Listed companies	3,975	as of 27 August 2026
Firms mentioned routinely	a few hundred	of order 10^{-4} of the whole

Table12.1: The size of the corporate population in Japan. The order of magnitude changes with each stage of definition.

Listed companies are about 0.1% of the whole, well-known firms less than 0.01%. When one says “a variety of firms”, which population is meant decides everything.

But **exhaustiveness does not guarantee size**. Even with a complete frame, if the population itself is small there is no power. Section 15.3.5 treats a population that can be obtained exhaustively yet numbers only 60 to 70 entities. **Exhaustiveness and size are independent requirements, and both must be checked before a survey’s prospects can be assessed.**

12.10 How to proceed

What can be handled is listed separately from what cannot.

12.10.1 Declare the frame in advance and track forward

Rather than “collect well-known firms”, define “all operators in a given industry existing at a reference date t_0 ” and track forward from there. Survival then becomes an observed dependent variable rather than a hidden filter, and the collider conditioning of Section 12.2 is avoided.

Available frames in Japan include the Economic Census and the databases of Teikoku Databank and Tokyo Shoko Research. In Europe, Orbis covers unlisted firms broadly. Disclosure obligations in the United States are looser, so its unlisted data are relatively weak; in this respect Japan and Europe are better placed.

12.10.2 Collect the failure side explicitly

Insolvencies are observable through the official gazette and the tabulations of credit research firms. If a denominator can be constructed, survivorship bias becomes an estimable selection model.

12.10.3 Lower the unit of analysis from the firm to Φ

This is the most effective remedy. By the definitions of Chapter 2, Φ is an object at the level of a contract and a firm is a bundle of them. A large firm normally runs several Φ in parallel whose κ differ even in sign. Computing CCC at the level of the firm is looking at an average that mixes them.

Lowering the unit of analysis to Φ

- partially resolves the age problem (a new Φ inside a mature firm can be handled)
- increases the number of observations
- reduces the influence of the institutional artefact called “the firm”

12.10.4 Group by the shape of the constraint, not by industry classification

Industry classifications fix a particular era’s view of industry and differ by country. Grouping instead by the structural features the framework of Part I names —

- presence of a capacity constraint (Section 2.7)
- transaction frequency ν ((6.4))
- observability of effort (Chapter 4)
- capital intensity (the weight of A in (3.8))

makes comparison meaningful across industries and eras. This is the return on having constructed the theory.

12.10.5 Where a distortion cannot be removed, state its sign

The non-identification of Section 12.5 and the retrospective narrative of Section 12.6 cannot be removed. Their direction can still be stated. Saying “this estimate understates the effect of structure” or “this description of Φ likely contains after-the-fact rationalization” makes the robustness of a conclusion assessable.

12.10.6 A rule for using well-known firms

Hence the following rule.

Well-known firms are used for illustration, not for inference.

Citing an individual firm to explain a structure is permitted. Concluding that a structure is superior from the existence of a successful example is not.

Chapter 13

A Taxonomy of Verification Obstacles: Six States

13.1 Why a taxonomy is needed

“Cannot be verified” covers states of different kinds. Because the remedies differ, they must be treated separately. This taxonomy is an organization contributed by this text.

When work stalls at stage 3 of the five stages of Chapter 11, availability, one cannot decide whether to spend effort or abandon the attempt without being able to say what exactly is stuck.

13.2 Six categories

Category	State	Correct remedy
(i) No measurement exists	nobody anywhere measures the quantity	conduct an original survey; secondary data will never fill the gap
(ii) Ill-posed hypothesis	the theory does not fix the sign	refine the theory; more data does not help
(iii) Not identified	the data exist but confounding cannot be broken	better design; find exogenous variation or a control group
(iv) Access constraint	it exists but cannot be obtained	open a route, or abandon on legal or technical grounds
(v) A choice in aggregation	it is surveyed but does not appear in the published tables	request the tabulation, or apply to use anonymized microdata
(vi) The adjustment does not operate	data are ample, but the explanatory variable does not vary over the range the theory assumes	no route resolves it

Table13.1: Six categories of verification obstacle: the state and the correct remedy.

The role of the taxonomy is to separate what effort can move from what it cannot.

(iii), (iv) and (v) can move with effort or ingenuity: (iii) with design, (iv) with a route, (v) with a change in what is published.

(i), (ii) and (vi) do not move. (i) cannot be filled without bringing the measurement into existence; (ii) does not fix a sign however much data is added; and under (vi), so long as the explanatory variable does not vary, no increase in sample size makes a test possible.

Remark 13.1 ((vi) is not a negative result). (vi) is a defect neither of the theory nor of the measurement. It is **the fact that the adjustment mechanism the theory assumes does not operate in reality.**

That a theoretical prediction fails to appear is itself an observation. In Part IV the statement “fee rates are not adjusted to the risk of the outcome” is such a case (Remark 15.6).

In terms of the layers of Section 6.3, prices being set by custom is a phenomenon of the institutional layer. The picture is one in which a theory premised on optimization fails to operate in the face of custom.

Remark 13.2 (Impossibility of aggregation differs from a choice in aggregation). Under (v) the quantity is merely not tabulated; it is in the sample. By contrast, there are cases where the unit of observation does not correspond to a party, so a quantity per party cannot be recovered. The identifiers of Chapter 10 are such a case (Remark 10.5).

That case is not resolved by a change in what is published. It is closer to (i) than to (v), but differs in that the quantity is not unmeasured; **the unit of measurement differs from the unit of the question.**

13.3 The taxonomy also serves as a test before starting

The taxonomy organizes, after the fact, why a verification did not complete; it also **serves as a test before a survey is begun.**

At stage 3 of Chapter 11 it is the criterion for judging whether the quantity can be obtained at the granularity required. The feasibility assessment of Chapter 17 is a worked instance of this use.

Part IV

Empirics

Chapter 14

Measuring the Theoretical Quantities: What Could and Could Not Be Measured

14.1 Purpose of this chapter

This chapter tabulates which of the symbols introduced in Part I could actually be measured, and what values they take where they could. Each chapter of Part IV treats a separate domain, so the correspondence with the theoretical quantities is easily scattered. This chapter supplies the index.

Its main purpose is to make explicit what could not be measured. Fewer than half the quantities the framework defines could be connected to real data.

The stage reached differs by quantity (Remark 11.3). The list here is a record of where between stages 1 and 3 of Section 11.2 each one stalled.

14.2 The primitive sets made concrete

Symbol	Domain	The set realized, and its cardinality
N	family 2	{issuers, purchasers, merchants}. 2,034 issuers (829 third-party, 1,205 own-use)
N	family 5	{agencies, job seekers, employers}. 30,561 establishments, 13,946 with placements
N	industry	the population of the corporate statistics: 62 industries $\times$ 16 capital classes
T	family 2	{FY2020–FY2024}, $ T = 5$
T	family 5	FY2024 alone; the panel covers FY2020–FY2025, $ T = 6$
T	industry	FY2000–FY2024, $ T = 25$
$\mathcal{D}$	family 2	four media of prepaid instruments (paper, magnetic, IC, server)
Ω	all domains	not observed. Only the arrival rate is proxied in Chapter 9
$\mathcal{F}_t$	all domains	not observed

Table14.1: How the primitive sets N , T , $\mathcal{D}$, Ω and $\mathcal{F}_t$ were realized in each domain.

Remark 14.1 (Consequence of Ω being unmeasurable). The state space itself cannot be observed. The dependence of π or δ on ω therefore cannot be verified directly, and every empirical result here concerns quantities marginalized over ω . The one exception is Chapter 9, where the arrival rate λ_{novel} of unforeseen states is proxied by counts of vulnerabilities.

14.3 The observed range of the map Φ

Part II enumerated 28 types; those that could be connected to real data are limited to the following (Table 14.2).

Type	Instance	Quantities observed
2-4 prepaid	prepaid payment instruments	D , P , κ , an upper bound on $\bar{\delta} - \delta$
5-1 success fee	fee-charging employment placement	π per placement, y , the extent of $\mathcal{G}$
6-1 transaction fee	payment and distribution platforms	the price of ϕ_{barg}
6-4 charging for opportunity	marketplace customer acquisition	the price of T_{acq}
1-3 one-off arrears	trade credit in the corporate statistics	$\kappa_C > 0$

Table14.2: The five of the 28 types that could be connected to real data, and the quantities observed.

That is **five of the 28 types**. Family 3 (renting assets), family 4 (complement recovery) and family 7 (payer separation) were not measured at all.

14.4 Cumulants and the credit position

14.4.1 D , P and κ in family 2

Substituting FY2024 values into the definitions of Chapter 2, in millions of yen:

$$P(\text{FY2024}) = 28,150,350 \quad (\text{issued: settlement from customers}) \quad (14.1)$$

$$D(\text{FY2024}) = 28,241,983 \quad (\text{redeemed: delivery by the firm}) \quad (14.2)$$

$$-\kappa_C(\text{FY2024}) = 2,895,327 \quad (\text{unused balance: customers credit the firm}) \quad (14.3)$$

Against (3.2), $\kappa_C = -2.90$ trillion yen. Being negative, customers extend credit to the issuer. Applying (5.1),

$$\text{CCC} = \frac{W}{r} = \frac{2,895,327}{28,150,350} \times 365 = 37.5 \text{ days.}$$

14.4.2 CCC by medium

The same calculation by medium shows a difference of two orders of magnitude.

Medium	$-\kappa_C$ (million yen)	r (million yen)	CCC (days)
Magnetic	562,980	107,925	1,904
Paper	1,127,149	486,732	845
Server	670,552	14,331,607	17
IC card	534,645	13,224,087	15

Table14.3: $-\kappa_C$, r and CCC by medium.

14.4.3 κ in the corporate statistics

Applying (3.2) to the corporate sector as a whole, FY2024, all industries excluding finance and insurance.

$$\kappa_C = \text{DSO} = 71 \text{ days (capital 1bn yen and over), } 35 \text{ days (under 10m yen)} \quad (14.4)$$

$$\kappa_S = -\text{DPO} = -44 \text{ days (1bn and over), } -21 \text{ days (under 10m)} \quad (14.5)$$

$$\sum_i \kappa_i = 27 \text{ days (1bn and over), } 14 \text{ days (under 10m)} \quad (14.6)$$

The sign of κ depends on both size and industry, and the industry effect dominates.

In retail, $\sum_i \kappa_i = -16$ days for firms of 1bn yen and over: the sign reverses.

14.5 Measuring the growth constraint

The $g^* = (m + d - I/r)/\text{CCC}$ of Corollary 5.7 needs four quantities. The measurement proceeds in two stages.

Object	Quantities used	Period
Long-run trend	m and CCC only; the ratio m/CCC	FY2000–FY2024
Level	$m, d, I, \text{CCC}; g^*$ itself	FY2016–FY2024

Table14.4: The two stages of measuring the growth constraint: object and period.

d and I are not used on the long-run side because depreciation and fixed-asset balances were not obtained back to the 1990s. For m , the operating margin is used: the m of Chapter 5 is what funds increases in working capital, so a figure net of selling and administrative expenses is closer to the definition.

14.5.1 The corporate sector over time (m/CCC)

Fiscal year	CCC (days)	m	m/CCC
2000	24.9	2.62%	38.5%
2005	24.4	3.16%	47.3%
2009	28.0	2.01%	26.2%
2015	30.6	3.95%	47.1%
2020	35.4	3.06%	31.5%
2024	37.2	5.01%	49.2%

Table14.5: CCC, m and m/CCC for the corporate sector over time (all industries).

All industries excluding finance and insurance, all sizes. d and I not included

The trend is shown in Figure 14.1.




Figure 14.1: m/CCC over time (all industries, on an operating-margin basis). It differs from the g^* of Corollary 5.7 by d and I .

Proposition 14.2 (m/CCC has not declined). Between FY2000 and FY2024, CCC worsened by 49%, from 24.9 to 37.2 days. But m improved by 91%, from 2.62% to 5.01%, so

$$\frac{m}{CCC} : 38.5\% \rightarrow 49.2\%$$

and the ratio in fact rose. The mean over 25 years is 41.1% with a standard deviation of 7.3 points. Against a mean of 38.0% over the first twelve years, the last thirteen average 44.0%. No declining trend is present.

Remark 14.3 (This proposition is not about g^* itself). The g^* of Corollary 5.7 is $(m + d - I/r)/CCC$. Proposition 14.2 concerns only m/CCC , which excludes d and I .

The reason for using m/CCC in the long time series is given in Table 14.4. The g^* including d and I is computed for FY2016–FY2024 in Section 14.5.

The proposition is therefore **a claim about the trend of a ratio**, not about the level of the self-financeable growth rate. If I has trended, the path of g^* can differ from that of m/CCC .

Remark 14.4 (Numerator and denominator must be measured separately). Looking at the deterioration of CCC alone suggests a worsening growth constraint. But Corollary 5.7 is a ratio, and **no conclusion follows from a change in the denominator alone**. Arguing from CCC without obtaining m reaches the opposite of the truth.

What drives the movement of m/CCC is mainly m . The troughs are FY2008–FY2009 (around 26%) and FY2020 (31.5%), both due to falling margins rather than to CCC.

14.5.2 g^* including depreciation and investment

Corollary 5.7 needs, besides m , the depreciation rate d and investment I . As Remark 5.12 states, $I \approx dr$ is not assumed. I is estimated as the change in fixed-asset balances plus depreciation. Land is excluded because it is not depreciated.

$$I \approx (\text{tangible} + \text{intangible} - \text{land})_t - (\text{tangible} + \text{intangible} - \text{land})_{t-1} + \text{depreciation}_t \quad (14.7)$$

Size (capital)	CCC	m	d	I/r	g^* (2024)	9-yr mean	SD
1bn yen and over	47	7.8%	3.1%	4.2%	52.0%	49.3%	6.0
100m – 1bn	35	4.8%	2.1%	4.3%	27.1%	39.6%	12.2
50m – 100m	30	3.9%	1.8%	1.9%	45.8%	33.0%	15.2
20m – 50m	25	3.5%	2.1%	1.2%	64.3%	35.0%	17.0
10m – 20m	37	1.8%	2.3%	5.5%	–13.8%	15.5%	30.8
Under 10m	26	1.5%	2.9%	2.0%	34.7%	3.1%	41.4

Table14.6: CCC, m , d , I/r and g^* by capital class (FY2024, nine-year mean, SD).

All industries excluding finance and insurance; CCC in days; SD is the nine-year standard deviation of g^* . Ten years of data (FY2015–FY2024) were obtained, but because the I of (14.7) needs the previous year's fixed-asset balance, g^* can be computed only from FY2016, and the mean and SD are taken over the nine years FY2016–FY2024

Remark 14.5 (Single-year values differ in reliability by size). The nine-year standard deviation of I/r is 0.3 for firms of 1bn yen and over and 2.9 for those under 10m — an order of magnitude apart. Because (14.7) uses changes in balances, retirements and disposals and **movement between capital classes** contaminate it. Capital changes with equity issuance irrespective of performance, so firms move between classes.

The direction of movement is not constant. The I/r of the under-10m class recorded -1.6% in 2016 and $+7.5\%$ in 2022, so this is not a one-directional bias. Averaging therefore largely cancels the effect, but single-year values are unreliable for the small classes.

This is why the table reports both single-year and nine-year figures with standard deviations.

Following Remark 14.5, judgments use the nine-year means.

Proposition 14.6 (The constraint on small firms is not CCC). On nine-year means, g^* for firms under 10m yen of capital is 3.1%, more than an order of magnitude below the 49.3% of firms of 1bn and over. The difference does not come from CCC: the CCC of the small classes (26–37 days) is in fact shorter than that of large firms (47 days).

The cause is the numerator. $m + d - I/r$ is 6.7% for firms of 1bn and over and 2.4% for those under 10m; besides a low m of 1.5%, an investment burden exceeding d contributes.

In the single year, the under-10m class shows a high 34.7%, but with a standard deviation of 41.4 no meaning should be read into that value.

14.5.3 g^* by industry

Industry	CCC	m	d	I/r	9-yr mean g^*	SD
Manufacturing	42	5.4%	2.8%	4.2%	34.8%	5.6
All industries	37	5.0%	2.5%	3.5%	37.7%	7.5
Wholesale	33	2.2%	0.7%	0.4%	20.9%	6.9
Construction	40	4.5%	1.5%	1.4%	47.1%	11.3
Retail	23	2.7%	1.5%	1.6%	29.0%	12.2
Employment placement and dispatch	33	3.9%	0.6%	−0.3%	42.7%	13.9
Information and communications	58	8.2%	4.8%	5.2%	61.2%	16.5
Other services	34	5.4%	3.0%	5.5%	47.9%	32.3
Accommodation	22	6.6%	3.5%	6.3%	−204.2%	872
Food services	5	1.2%	2.4%	−0.4%	266.9%	3334

Table 14.7: CCC, m , d , I/r and the nine-year mean g^* by industry, with standard deviations.

The bottom two industries in Table 14.7 exhibit the divergence of Remark 5.13. The CCC of food services is 5 days, so a one-point movement in the numerator moves g^* by about 73 points. The standard deviation of 3334 is the consequence of that amplification: **g^* is not functioning as an indicator.**

Remark 14.7 (Two meanings of a short CCC). In the types of Part II, food services is family 1-1 (immediate exchange) and $\tau_i \approx 0$ is a structural property (Corollary 5.17). A short CCC indicates sound liquidity, but **it simultaneously makes the ratio g^* meaningless.** Corollary 5.7 is useful only for objects whose CCC is some tens of days or more.

14.5.4 The m of small firms is decided by directors' remuneration

Proposition 14.6 showed that m is the constraint but not why m is low. Decompose the cost structure.

Size (capital)	Gross margin	SG&A ratio	Personnel ratio	m
1bn yen and over	23.6%	15.8%	7.7%	7.8%
100m – 1bn	21.5%	16.8%	10.3%	4.8%
50m – 100m	25.9%	22.0%	12.3%	3.9%
20m – 50m	26.0%	22.4%	15.0%	3.5%
10m – 20m	33.8%	32.0%	17.7%	1.8%
Under 10m	46.3%	44.8%	22.7%	1.5%

Table14.8: Gross margin, SG&A ratio, personnel ratio and m by capital class (FY2024).

All industries excluding finance and insurance, FY2024

Small firms have high gross margins: 46.3% against 23.6% for large firms. Their SG&A ratio, however, is 44.8% against 15.8%. As Remark 2.17 notes, part of that difference may be C_{admin} : at a scale where invoicing and chasing are done by hand, the operating cost of contracts is booked in SG&A. This text does not separate it. The low m is therefore not due to a lack of pricing power; the drop occurs at the SG&A stage.

Moving personnel costs back into cost of sales narrows the gap in gross margin from 22.7 points to 7.7 (15.9% against 23.6%). Two-thirds of the gap was a difference of accounting classification. Even after the adjustment, however, the small classes remain higher.

Decomposing personnel costs into directors and employees identifies the cause.

Size (capital)	Directors' remuneration	Employee wages	Directors' share	m
1bn yen and over	0.16%	7.6%	2.0%	7.8%
100m – 1bn	0.40%	9.9%	3.9%	4.8%
50m – 100m	0.99%	11.3%	8.1%	3.9%
20m – 50m	2.35%	12.7%	15.6%	3.5%
10m – 20m	3.98%	13.7%	22.5%	1.8%
Under 10m	8.04%	14.6%	35.5%	1.5%

Table14.9: Directors' remuneration, employee wages, directors' share and m by capital class.

The directors' share is directors' remuneration as a fraction of personnel costs

Proposition 14.8 (Directors' remuneration explains the difference in m by size). Directors' remuneration as a fraction of sales is 0.16% for firms with capital of 1bn yen and over and 8.04% for those under 10m — a **fiftyfold difference**. The gap in m between the largest and smallest classes is 6.3 points, while the gap in directors' remuneration is 7.88 points, large enough to account for it.

Adding directors' remuneration back into operating profit removes the monotone ordering.

Size	m	After adding remuneration back
1bn yen and over	7.8%	8.0%
20m – 50m	3.5%	5.9%
Under 10m	1.5%	9.5%

Table14.10: m before and after adding directors' remuneration back into operating profit, by size.

The smallest class becomes the highest. The same holds within industries: m is higher for the small classes in 1 of 9 industries as reported, but in 7 after remuneration is added back.

Remark 14.9 (A discrepancy with the definition of m). Chapter 5 defined m as what funds increases in working capital. If directors’ remuneration is discretionary, what a small operator can actually use is not operating profit but **operating profit plus directors’ remuneration**. On that basis the numerator for the under-10m class is not the 2.4% of $m + d - I/r$ but 10.4% once remuneration of 8.04% is added. Dividing by $CCC = 26$ days gives 146%, far above the 51.7% of large firms (Section 14.5; nine-year mean 49.3%).

That is, Proposition 14.6’s claim that small firms are constrained by m is correct only so long as reported operating profit is taken to be m ; **with an m faithful to the definition, the conclusion reverses**.

14.5.5 Comparison with unincorporated businesses

Proposition 14.8 rests on the corporate statistics. Whether the same structure appears in a different statistic, with different size classes, and in a sample including **entities without corporate personality**, is checked here.

The Basic Survey on Small and Medium Enterprises was used: a sample of about 110,000 firms in which four corporate size classes by headcount and unincorporated businesses appear in the same table. Unincorporated businesses have no line item for directors’ remuneration, so the proprietor’s return to labour remains in operating profit.

Registered implication: the operating margin of unincorporated businesses is close to the corporate small classes’ figure “after adding remuneration back”. The two are substantively the same economic agent and differ only in accounting classification.

Class	2019	2020	2021	2022	2023	2024	Mean	SD
Corporate, 5 or fewer	2.7	1.1	2.6	3.0	2.6	3.0	2.5	0.7
Corporate, 6–20	2.4	1.5	2.5	2.8	3.0	3.2	2.6	0.6
Corporate, 21–50	2.8	2.1	3.2	2.8	3.4	3.7	3.0	0.6
Corporate, 51 and over	3.4	2.8	3.2	3.6	3.6	4.0	3.4	0.4
Unincorporated	17.0	18.5	18.2	15.8	15.4	16.3	16.9	1.3

Table14.11: Operating margin by corporate size class and for unincorporated businesses (FY2019–FY2024).

Operating margin (%), all industries, by fiscal year

The implication is supported. The operating margin of unincorporated businesses is about five times the corporate figure in all six years, with a standard deviation of 1.3. This is not a single-year peculiarity. The corporate figure after adding remuneration back in Section 14.5.4 was 9.5%. The 16.9% of unincorporated businesses exceeds that, but is far closer to it than to the 1.5–3.0% before adding back.

The personnel-cost ratio explains the difference.

Class	Operating margin	Personnel ratio	Total
Corporate, 5 or fewer	2.5%	14.7%	17.2%
Corporate, 6–20	2.6%	16.9%	19.5%
Unincorporated	16.9%	9.8%	26.7%

Table 14.12: Operating margin and personnel ratio, corporate and unincorporated (six-year means).

Six-year means. The personnel ratio is labour cost in cost of sales plus personnel cost in SG&A

Unincorporated businesses have the lowest personnel ratio, because the proprietor's labour is not booked as a cost. What appears in personnel costs as directors' remuneration in a company remains in operating profit for an unincorporated business.

Remark 14.10 (The sign reverses during the pandemic). In FY2020 the operating margin fell in every corporate class (2.7% $\rightarrow$ 1.1% for firms of five or fewer). **Only unincorporated businesses rose** (17.0% $\rightarrow$ 18.5%).

This has the same form as the observation in Remark 14.12. Because the operating profit of an unincorporated business contains the proprietor's own take, the ratio does not fall when sales fall unless the proprietor reduces that take. If employee wages and outsourcing are cut first, it rises.

That directors' remuneration in companies moves independently of profit, and that the operating margin of unincorporated businesses rises in a downturn, are **two appearances of one and the same phenomenon**.

Remark 14.11 (Capital expenditure by unincorporated businesses). The same survey gives I/r for unincorporated businesses: a six-year mean of 1.9% with a standard deviation of 1.0. That is lower and less stable than the 3.1% (SD 0.1) of corporate firms with 51 or more employees. The share of firms making any capital expenditure is 8.1% (FY2024): more than nine in ten invest nothing.

Chapter 8 assumed $A \approx 0$ for a solo business. The assumption agrees with the data on unincorporated businesses.

Remark 14.12 (What cannot be decided). Whether directors' remuneration is a discretionary appropriation of profit or a return to labour cannot be decided here.

An observation supporting discretion: the directors' remuneration ratio of the under-10m class stayed in the narrow band 7.6–9.1% throughout FY2015–FY2024, while m over the same period ranged from -2.0% to $+2.1\%$. **Profit moves and remuneration does not**.

But the labour-return account cannot be rejected either. Directors of small firms work in substance as employees, and it is equally possible that the remuneration is the return to that. Distinguishing the two needs remuneration per director and hours worked, but **directors are not employees under the Labour Standards Act and no statistics on their hours**

exist. Both the Monthly Labour Survey and the Basic Survey on Wage Structure are confined to employees. This belongs to category (i), no measurement exists, of Section 13.2.

14.5.6 By industry (FY2024)

Industry	CCC (days)	m	g^*
Pure holding companies	12	57.7%	1710%
Water transport	7	4.8%	261%
Mining, quarrying and gravel extraction	29	19.0%	236%
Motor vehicles and parts	9	5.1%	201%
Leasing	210	6.6%	11.4%
Petroleum and coal products	20	0.6%	11.4%
Textiles	78	2.0%	9.2%
Fisheries	40	−1.5%	negative
Agriculture and forestry	22	−4.0%	negative

Table14.13: CCC, m and g^* by industry (FY2024; the extremes).

Where the numerator $m + d - I/r$ is negative, $g^* < 0$, and by Proposition 5.5 **free cash flow is negative even at zero growth.**

A long CCC and a small g^* do not coincide (Table 14.13). Leasing has the longest CCC at 210 days but a high $m = 6.6\%$, so on m alone it sits at the same level as petroleum and coal products (CCC = 20 days, $m = 0.6\%$). **CCC alone cannot decide the growth constraint.** The g^* including d and I is given in Section 14.5.

14.6 Right and exercise, and cognitive surplus

The $\phi_{\text{cog}} = p \cdot \mathbb{E}[\bar{\delta} - \delta]$ of (2.15) is a measurable quantity, and measurements exist. Four are given below. **What this text obtained itself, however, is only an upper bound; the value is unknown.**

14.6.1 ϕ_{cog} is measurable

■ This text's measurement: family 2 recoveries on expiry.

$$p \cdot \mathbb{E}[\bar{\delta} - \delta] \leq 20,634 \text{ million yen} \quad (14.8)$$

$$\frac{20,634}{2,895,327} = 0.71\% \quad (\text{of unused balances}) \quad (14.9)$$

$$\frac{20,634}{28,150,350} = 0.073\% \quad (\text{of issuance}) \quad (14.10)$$

The inequality holds because this figure also contains recoveries through refunds. **What was obtained is an upper bound, not ϕ_{cog} itself.** The three examples below all give the value itself.

■ **Gym memberships.** [6] matched the contract choices and daily attendance records of 7,752 members across three US gyms, measuring the $\bar{\delta} - \delta$ of (2.15) directly. Members predicted 9.50 visits a month and made 4.17. Members on plans above \$70 a month attended on average only 4.3 times, paying more than \$17 a visit where a ten-visit pass would cost \$10. The loss over a membership averaged \$600 against total payments of about \$1,400, so ϕ_{cog} is about 43% of the amount paid.

■ **Subscription renewals.** [9] covers a wider range. Using comprehensive data from a payment-card network, it exploits as an exogenous shock the fact that card renewal forces an active renewal decision. Across ten subscriptions, **inattention raises sellers' revenue by 89% on average** (estimates ranging from 14% to over 200%). The attention parameter λ is estimated at 0.18 on average, ranging from 0.04 to 0.50 across services.

■ **Breakage on prepaid balances.** [13] treats family 2-4 (prepaid) directly. Analysing transaction data on over four million people for a scheme adding a bonus to prepaid restaurant balances, it records that **about 40% of prepaid value goes unused**. Operators earn a median of \$5.50 of breakage profit per dollar of bonus, and the profit from breakage exceeds the revenue from increased spending.

Proposition 14.13 (Three conditions for measuring ϕ_{cog}). Estimating ϕ_{cog} requires all three of the following simultaneously.

- (1) **The contractual right $\bar{\delta}$:** the tariff and any usage ceiling
- (2) **Realized use δ :** individual records for the same party; aggregates will not do
- (3) **An exogenous event forcing attention:** card renewal, a change of terms, a price revision

Grounds. Without (1) and (2) the difference $\bar{\delta} - \delta$ cannot be defined. Without (3), observed continuation cannot be distinguished between “continued because it was worth it” and “continued through inattention”, and λ is not identified. [6] satisfies (1) and (2) through contract data and attendance records; [9] satisfies (3) through card renewal. $\square$

This text's family 2 measurement stopped at an upper bound **because (2) and (3) were missing**. Aggregate statistics contain no individual records and there is no exogenous shock. In the taxonomy of Section 13.2 this is not (i), no measurement exists, but (iv), access constraint: the method of measurement is established and what is needed is reaching the data.

14.6.2 A detour through the menu

Condition (2) of Proposition 14.13, individual realized use, is not available in published statistics. But a **substitution of definitions can sometimes provide a detour**.

When a flat rate F and a per-use price p coexist on the same menu, the break-even number of uses can be computed.

$$\delta^* = F/p \quad (14.11)$$

Even where $\bar{\delta}$ is “unlimited” and unmeasurable, (14.11) follows **from the menu alone**. Using it as a proxy for $\bar{\delta}$, all that is needed is the industry-average number of uses δ .

Since (14.11) is equivalent to $p\delta^* = F$, substituting into (2.15) gives

$$\phi_{\text{cog}} \approx p \mathbb{E}[\delta^* - \delta] = F - p \mathbb{E}[\delta] \quad (14.12)$$

Under this proxy, cognitive surplus is the difference between the flat fee and the per-use equivalent. The registered implication $\delta^* - \delta > 0$ becomes the same claim as (14.12) being positive, that is, the flat fee exceeding the per-use equivalent of average use.

Registered implication: for flat-rate members, $\delta^* - \delta > 0$; the average member does not reach break-even.

■ **Result.** From the Survey on Selected Service Industries, δ is 89 visits a year per individual member, that is 7.42 a month. δ^* was computed from the published prices of one large full-service operator (2026, tax included).

Facility category	Unlimited plan	Per visit	δ^*	$\delta^* - \delta$
I	12,680	2,530	5.01	−2.40
II	13,980	2,860	4.89	−2.53
III	16,580	3,190	5.20	−2.22
IV	19,080	3,630	5.26	−2.16

Table14.14: Unlimited and per-visit prices by facility category, and the break-even number of visits δ^* .

The implication is rejected. The sign is negative in all four categories (Table 14.14): for the average member the flat contract is rational.

The effective unit price — total dues divided by total visits — points the same way, falling 29% over 25 years from 1,147 yen in 2000 to 815 yen in 2023. Against a per-visit market rate of 2,530–3,630 yen, **members are using the facility more cheaply than paying per visit.**

Remark 14.14 (When the detour applies). Equation (14.11) is uniquely determined only when the menu is a binary choice between flat and per-use. Actual tariffs are stepped: the operator above offers four visits a month, eight a month, and unlimited. The threshold then changes with the comparison chosen. Choosing the unlimited plan is rational at 5.0 visits a month against per-visit pricing, but requires 9.1–9.5 against the eight-visit plan. **Average use of 7.42 is rational on the former and irrational on the latter.**

The detour therefore applies only to binary menus; with a stepped menu the proxy for $\bar{\delta}$ is not uniquely determined.

Remark 14.15 ($\phi_{\text{cog}} = 0$ does not follow from this section). The rejection is a result about the mean. The findings of [6] concern the distribution, and a structure in which the lower tail loses is not excluded by a mean above break-even. Mean and distribution are different claims.

Remark 14.16 (What remains unmeasured within this framework). The prior work concentrates on subscriptions and memberships, that is, family 2-1 of Part II. No comparable estimates were found for family 2-4 (prepaid, gift cards) or family 2-5 (options and warranties). Nor was any study found estimating the *share* of ϕ_{cog} within total ϕ across industries. What [9] measures is the increment to revenue, not a composition of profit.

The reason for purchase by medium serves as a proxy for the transaction frequency ν .

Medium	For own use	As a gift	CCC (days)
IC card	95.5%	0.0%	15
Server (online)	78.1%	14.0%	17
Paper	41.5%	30.1%	845

Table14.15: Reason for purchase by medium (own use versus gift) and CCC.

The higher the gift share, the lower ν and the longer CCC. The $\lambda(\nu)$ of (6.4) itself was not measured.

14.7 The information structure $\mathcal{G}$ made concrete

The $\mathcal{G}$ of Chapter 4 — verifiable shared information — is observable in family 5 as a concrete set of fields. The items published on the Employment Service Information Site are what $\mathcal{G}$ consists of.

Symbol	Instance	Status
y	separations within six months among those placed on open-ended contracts	numeric, six years
Breadth of $\mathcal{G}$	the count of “separation not ascertained” (the smaller, the wider $\mathcal{G}$)	numeric; zero in all years for both observations
b (outcome-contingency)	refund schemes	presence is binary; rates and periods are not in a uniform format
Proxy on the price side	fee rates achieved by occupation	continuous; 307 firms (Section 15.3.5)
σ^2 (noise in the outcome)	baseline separation rates by occupation	not measured

Table 14.16: The information structure $\mathcal{G}$ made concrete, and the observation status of each symbol (family 5).

Measured values of y at the two establishments that could be confirmed:

$$\frac{y}{\text{open-ended placements}} = \frac{124}{4,348} = 2.85\%, \quad \frac{3}{189} = 1.59\%.$$

Testing the $b^* = 1/(1 + \gamma\sigma^2 k)$ of Example 4.5 requires both b and σ^2 as continuous quantities, but refund rates and periods are not in a uniform format and the uncontrollable part cannot be extracted from σ^2 .

A test on the price side was possible, however. Fee rates achieved by occupation were obtained as a continuous quantity for 307 firms and correlated with separation rates (Section 15.3.5).

14.8 The three-way decomposition of surplus

The three components of (6.1) can be separated only in the case of Chapter 17.

Component	Measured	Grounds
ϕ_{prod} (price of function)	20.0 points	within-platform comparison on Gumroad, differing only in whether customer acquisition is included
ϕ_{barg} (rent from market power)	15.0 points	the App Store’s two size tiers; function identical
ϕ_{risk} (price of transfer and pooling)	—	≈ 0 for these two firms; positive where cases are screened individually
ϕ_{cog}	—	not measured in this domain

Table14.17: Measured values of the three-way decomposition of surplus, and their grounds (platform fees).

This is **the only case in which two of the three components could be separated numerically for a single transaction**. No procedure for separating the components exists in the other domains.

The two firms are intermediaries operating at uniform rates published in a price list, so the risk component ϕ_{risk} was negligible and separation was possible. Where an intermediary sets rates by individual screening, a third term is added and the separation fails (Remark 17.2).

14.9 The terms of the launch cost

For each term of (8.11), the price of externalization was measured, as the difference in effective rate on a \$50 transaction.

Term	Means of externalization	Price
T_{inst}	payment processing	3.6%
T_{inst} (tax)	merchant of record	+1.4–2.4 points
T_{dev} (delivery)	a sales platform with hosting	+8–9 points
T_{acq}	marketplace customer acquisition	+20 points
T_{credit}	self-funded through a free period	the cost of $\kappa > 0$

Table14.18: The terms of the launch cost, the means of externalization, and the price.

T_{acq} **exceeds the sum of all the other items**. T_{dev} itself (the development period) and the enumeration of $\hat{\Omega}$ cannot be externalized and therefore have no price.

14.10 Quantities relating to selection and identification

Symbol	Measured	Source
$\mathbb{P}(S = 1)$ (listed)	$3,975/3,600,000 = 0.11\%$	Section 12.9
$\mathbb{P}(S = 1)$ (mentioned)	of order 10^{-4}	as above
τ (insolvency time)	not observed	outside the sample through conditioning on survival
(a, p, c)	not identified	Proposition A.7
Upper bound on λ_{novel}	48,185 a year (published CVEs)	Chapter 9
Effective λ_{novel}	245 a year (confirmed exploitation) = 0.51%	as above

Table14.19: Measured quantities relating to selection and identification, and their sources.

14.11 Quantities that could not be measured

The following are defined by the framework but could not be connected to real data (Table 14.20).

Symbol	Name	Why it cannot be connected
$M(t)$	cash balance	needs per-firm data; meaningless in aggregate
$\Omega, \mathcal{F}_t$	state space, filtration	not observable in principle
v	valuation map	subjective; the definition cannot be fixed
β_i	discount factor by party	requires experimental methods
Cap	capacity ceiling	differs by operator and is not disclosed
σ^2	noise in the outcome	variance by occupation is not tabulated (Section 15.3.5)
$\lambda(\nu)$	decay rate of cognitive surplus	the numerator ϕ_{cog} is itself unmeasured
ϕ_{cog}	cognitive surplus	only an upper bound here; the method is established (Section 14.6.1)
$\hat{\Omega}$	envisaged state space	exists only inside the designer
τ	insolvency time	outside the sample through conditioning on survival (Proposition 4.7)

Table 14.20: Quantities defined by the framework that could not be connected to real data.

ϕ_{cog} is measurable in principle (Proposition 14.13); that this text could not reach it is a matter of data granularity. The rest belong to (i), no measurement exists, or are unobservable in principle.

Remark 14.17 (The framework is larger than what can be measured). About half the quantities defined could not be connected to real data. That is to say the framework describes the world more finely than the world is recorded, and it shows that precision of description and testability are independent. In the taxonomy of Section 13.2, most of what cannot be connected belongs to (i), no measurement exists, and effort does not resolve it.

Chapter 15

Empirics of families 2 and 5: prepayment and outcome-contingent forms

15.1 The propositions tested in this chapter

Among the types enumerated in Part II, families 2 (continuing access) and 5 (outcome- or state-contingent) are taken as subjects, and Remark 2.22 (conditions for divergence) together with the measurability condition of Chapter 4 is tested. **The reason for choosing these two families is given in Section 11.5:** each contains, singly, one of the two terms of the three-way decomposition of surplus that depend on the shape of Φ , and they sit on the diagonal of Figure 2.1. Following the policy of Part III, the implications were fixed before the investigation, and whether they were supported is recorded as it stands, including the parts that failed.

The two implications fixed in advance are as follows.

Registered implication (family 2): the lower the transaction frequency ν of a product, the larger the divergence between entitlement and exercise. Concretely, gift vouchers show a larger divergence than payment instruments used daily.

Registered implication (family 5): the harder the ex post verification of the outcome in an occupation, the stronger the refund scheme (a higher refund rate, or a longer covered period).

The conclusions first. The family-2 implication was supported as to level, but what could be measured was κ , not ϕ_{cog} . The family-5 implication could not be tested: the necessary variable does not exist in the aggregate statistics.

Theoretical quantities treated in this chapter

D, P, κ_C	equation (3.2)	directly observed as issuance, redemption and unused balance
$CCC = W/r$	equation (5.1)	computed by medium and by industry; 15–1,904 days
$\bar{\delta} - \delta$	Section 2.7	the portion reaching expiry; an upper bound only
ν	equation (6.4)	proxied by the reason for purchase (share bought as gifts)
$\mathcal{G}$	Chapter 4	observed in family 5 as the set of published items
y	Chapter 4	number separating within six months
b	Example 4.5	presence of a refund scheme; rates and periods are not standardized
Realized fee rate	Section 15.3.5	a proxy on the price side; continuous

Table15.1: Theoretical quantities treated in this chapter (families 2 and 5) and the measurements corresponding to them.

15.2 Family 2: prepaid payment instruments

15.2.1 Data source and population

The 27th Survey of the Actual State of Issuing Businesses (FY2024 edition) of the Japan Payment Service Association was used. As set out in Section 12.9, it is a complete frame covering listed and unlisted issuers alike.

Item	Content
Population	2,034 registered or notified issuers at the end of FY2024 (829 third-party type, 1,205 own-use type)
Responses	812 issuers, a response rate of 39.9%; on an issuance basis, however, about 90% of all issuance is covered
Scale	issuance ¥28.2 trillion, unused balance ¥2.90 trillion
Caveat	the association states explicitly that “because respondents differ from survey to survey, the figures are not continuous”

Table15.2: Data source and population for family 2 (prepaid payment instruments), FY2024.

15.2.2 Measuring the residence period

The quantities of this section correspond to the notation of the theory as follows: issuance is cumulative settlement P , redemption is cumulative delivery D , and the unused balance is $-\kappa_C$ in equation (3.2). The FY2024 values are

$$P = 28,150,350, \quad D = 28,241,983, \quad \kappa_C = -2,895,327 \quad (\text{millions of yen}) \quad (15.1)$$

so $\kappa_C < 0$: the customer extends credit to the issuer. Taking working capital as $W = -\kappa_C$ and sales velocity as $r = P$, equation (5.1) gives

$$\text{CCC} = \frac{W}{r} = \frac{2,895,327}{28,150,350} \times 365 = 37.5 \text{ days},$$

which converted to days is called the residence period. The results by industry are given in Table 15.3.

Industry	Issuance (¥mn)	Residence days
Hotels and inns	652	2,111
Travel	42,603	1,948
Department stores	72,890	1,223
Real estate	1,230	650
Telecommunications	237,299	463
Cooperatives, chambers of commerce, etc.	187,364	140
Sports and leisure	12,361	115
Food services	142,327	107
Credit and instalment sales	4,301,983	59
Transport	3,123,308	26
Supermarkets	2,198,404	17
Total	28,150,350	38

Table15.3: Issuance and residence days by industry (family 2, FY2024).

The ordering is consistent with the registered implication. Transport (transit IC cards, daily use) and supermarkets (everyday shopping) are around 20 days, whereas accommodation vouchers, travel vouchers and department-store gift certificates run from 1,200 to 2,100 days.

As corroboration, the tabulation of reasons for purchase serves as a proxy for the transaction frequency ν of equation (6.4). For paper instruments, 41.5% are for the purchaser's own use and 30.1% are gifts; for IC instruments, 95.5% are for own use and 0.0% gifts. In a gift the agent holding $\bar{\delta}$ and the agent performing δ are separated, so ν is structurally low.

15.2.3 Separating the alternative

There is a strong alternative. If the contraction of the paper-voucher market reduces issuance (a flow) while the unused balance (a stock) remains as an accumulation from the past, the ratio rises mechanically. This is precisely the period effect of Section 12.5.

A cross-section by industry cannot separate the two. But a panel by medium **contains growing and contracting media at the same time**, which makes identification possible.

Medium	FY20	FY21	FY22	FY23	FY24	Issuance, FY20→FY24
Magnetic	1,307	1,630	1,338	1,916	1,904	−32.2%
Paper	775	758	668	764	845	−6.0%
Server	16	15	17	16	17	+45.3%
IC	15	14	16	15	15	+23.5%

Table15.4: Residence days by medium and the change in issuance (FY2020–FY2024, family 2).

(Figures are residence days; the rightmost column alone is the change in issuance.)

The verdict splits in two.

- (1) **The difference in levels cannot be explained by market contraction.** Paper instruments lost only 6.0% of issuance over five years, yet already showed a residence of 775 days in FY2020 — about fifty times the 15 days of IC instruments. Since the difference exists at a starting point where contraction had not proceeded, the level derives from structure. Moreover the residence of the two growing media (IC and server) is almost immobile over five years: no mechanical rise in the ratio occurs.
- (2) **Part of the variation can be explained by market contraction.** Magnetic instruments lost 32.2% of issuance and their residence rose from 1,307 to 1,904 days. Paper instruments have also returned from 668 days in FY2022 to 845. Here the alternative bites.

The path of the levels is shown in Figure 15.1.

The separation is thus **level from ν , part of the variation from a period effect**. Against the problem described as generally unidentified in Section 12.5, the simultaneous presence of a growing segment permitted partial identification. This is not a general solution but a piece of luck specific to this subject matter.

As an auxiliary observation, the share with no expiry date set is 71.2% for paper, 52.0% for magnetic, 37.9% for IC and 19.0% for server (physical stores), which apart from magnetic orders the same way as residence days. The direction of causation cannot be pinned down, however: is it that no expiry is set and so the balance sits, or that products whose balance sits — gifts, for instance — are the ones given no expiry? Probably the same ν determines both, but these data cannot settle it.




Figure 15.1: Residence days by medium (logarithmic scale). A two-order-of-magnitude difference in level exists from the starting point.

15.2.4 What could not be tested

Proposition 15.1 (The limits of what was measured). What this section measures is the size of the credit position κ , not the cognitive surplus ϕ_{cog} of equation (2.15).

A long residence does not imply that the entitlement is ultimately never exercised. The only figure corresponding to ϕ_{cog} is ¥20,634 million recovered through expiry and the like — 0.71% of the unused balance and 0.073% of issuance. The association further notes that this amount is thought to include recoveries through refunds, so it is not a pure non-exercise figure.

The conclusion for family 2 is therefore confined to the following.

In a prepayment Φ , the lower the transaction frequency ν of a product, the larger $|\kappa|$, and the difference reaches two orders of magnitude. Whether this residence converts into cognitive surplus, however, cannot be judged from public data.

Chapter 5 (κ and residence) is supported, while Section 2.7 (cognitive surplus) remains untested. The former is only a necessary condition for the latter.

15.2.5 The absence of per-firm data

The association states explicitly that officials of the Local Finance Bureaux and other administrative bodies are under a statutory duty of confidentiality and cannot pass matters reported by issuers to the association as a third party. The Financial Services Agency publishes a list of registered and notified issuers, but it does not include unused balances. In addition, issuers themselves have asked that the state of non-redemption (expiry and the like) also be disclosed — that is, the breakage rate is not currently tabulated.

As a result, no dispersion within an industry or within a medium can be observed at all. Only means are visible.

15.2.6 The nature of the remaining obstacle

ϕ_{cog} cannot be measured in this domain because the breakage rate is not tabulated. Issuers themselves have requested of the association’s statistics that “the state of non-redemption (expiry and the like) also be disclosed”, so measurement is not being carried out within the frame of this statistic.

This looks like (i) absence of measurement. But the obstacle is more narrowly located. Of the three conditions of Proposition 14.13, what is missing in this domain is (2) realized use at the individual-record level and (3) an exogenous event that forces attention; (1) the contractual entitlement is observable as the expiry date by medium. That is, **the method of measurement is established, and what is missing is the granularity of the data.**

Under the classification of Section 13.2, then, ϕ_{cog} for family 2 belongs to (iv) access constraint. Unlike an ordinary access constraint, however, data of the required granularity exist only inside the issuer, and there is no route by which publication can be awaited.

That per-firm unused balances are confidential is secondary. Even if per-firm data were obtained, ϕ_{cog} could not be estimated so long as the reported items contain no record of individual usage.

Principal obstacle for family 2	(iv) access constraint; individual records and an exogenous shock are lacking
---------------------------------	---

Table15.5: Verdict on the principal obstacle for family 2.

15.3 Family 5: fee-charging employment placement

15.3.1 Data source and population

The Ministry of Health, Labour and Welfare’s “Tabulation of Employment Placement Business Reports, FY2024 (preliminary)” was used. It is a complete enumeration of all operators under the Employment Security Act and, like family 2, escapes listing bias.

Item	Content
Population	30,561 fee-charging placement offices filing a business report; 13,946 of them recorded placements
Scale	888,993 regular placements (fee-charging), fee income about ¥983.5 billion
Unit price	about ¥1.03 million in fees per regular placement

Table15.6: Data source and population for family 5 (fee-charging employment placement), FY2024.

15.3.2 The choice of contractual regime

The Employment Security Act provides two methods of charging fees: the capped scheme and the notified scheme. The capped scheme is a formula-based contract fixed at a set rate of the wage; the notified scheme is a contract the operator sets freely and files.

Type of fee	Amount (¥thousand)	Share
Notified-scheme fees	980,821,781	99.73%
Capped-scheme fees	1,804,568	0.18%
Other (application fees, etc.)	827,808	0.08%
Total	983,454,157	100%

Table15.7: Breakdown by method of charging fees (family 5, FY2024).

The capped scheme has all but disappeared.

Remark 15.2 (This fact is not evidence about measurability). The capped scheme is fixed at about 11% of the wage (six months' worth), which converts to only about 5% of annual income. The prevailing notified-scheme level of around 30% of annual income is an order of magnitude different. The market did not reject formula-based contracts; the price the formula sets is simply too low. This is a phenomenon in which the institutional layer prescribed a price and the option consequently fell out of use, and it has nothing to do with the measurability of Chapter 4. The disappearance of the capped scheme is a consequence of the price level and provides no material for the measurability hypothesis.

15.3.3 Fees per placement by occupation

Fees mix in temporary and day work. Nationally, fees relating to regular placements account for ¥913.6 billion, about 93% of the whole, but the breakdown by occupation does not separate the two. Defining the degree of contamination as the number of temporary and day placements (person-days) per regular placement, the calculation is restricted to occupations where this is below one.

Occupation	Placements	Contamination	¥10k/pl.
Corporate and organizational officers	1,392	0.42	437
Corporate and organizational managers	16,007	0.16	288
Professionals in management, finance and insurance	11,346	0.09	268
Legal occupations	2,217	0.05	225
Researchers	3,167	0.27	188
General affairs, HR and planning	25,716	0.38	183
Development engineers	19,339	0.40	165
IT and communications engineers (software development)	39,780	0.20	150
Sales occupations	89,563	0.28	142
Architecture, civil engineering and surveying	26,230	0.12	139
Designers	3,969	0.38	116

Table15.8: Fees per placement by occupation (occupations with contamination below one, family 5).

The ordering corresponds to seniority, and if the fee rate is roughly constant this merely reflects annual income. **Without a denominator of annual income, the table has no power to identify the measurability hypothesis.** Whether officers' fees are high because outcomes are hard to observe or simply because incomes are high cannot be distinguished from these data.

15.3.4 Why the measurability hypothesis cannot be tested

Data on refund schemes do not exist in this tabulation. They appear only firm by firm on the Jinzai Service General Site. The measurability hypothesis is therefore untested.

There is a more serious problem.

Remark 15.3 (The gap in the data coincides with the interest). Medicine, care and childcare are the fields into which regulation, refund schemes included, has been concentrated. The very fact that the institutional layer intervened is evidence that the authorities judged the measurability problem to be serious there. Yet these occupations are heavily contaminated by temporary and day work (spot shifts), and fees per placement cannot be computed from the aggregates. The degree of contamination is 35 for physicians, 12 for nurses and assistant nurses, 5 for institutional care and 4 for childcare workers.

That is, **the domain of greatest theoretical interest is precisely the gap in the data.** This may not be a coincidence: fields in which spot work develops are fields in which employment relations shorten, which is exactly where ex post verification of outcomes is hard.

15.3.5 Fees and turnover

Data source

On checking the office detail pages of the Jinzai Service General Site, the problem of unstructured data feared in Part IV did not arise, except for the content of refund schemes.

Variable	Published form	Verdict
Placements (open-ended / fixed-term / day)	numeric, six years FY2020–FY2025	structured
Separations within six months (open-ended)	numeric, same six years	structured
Separation not ascertained (open-ended)	numeric, same six years	structured
Realized fee rate by occupation	percentage, from FY2025	structured
Presence of a refund scheme	yes / no	structured
Refund rate and period	a PDF attached by the operator in a free format	unstructured

Table15.9: Variables published on the Jinzai Service General Site and their degree of structure.

Moreover **the detailed search allows occupation, realized fee rate and turnover rate to be specified as conditions, and displays all three together in the results.** This removes the need to open offices one at a time and made the retrieval of the 307 firms of Section 15.3.5 possible. The occupations that can be specified as conditions, however, are limited to nine in medicine, care and childcare.

The source is favourable in three respects.

First, **regular and temporary or day work are separated.** The contamination that made the aggregates for physicians and nurses unusable in Part IV does not arise at the level of the individual firm.

Second, **fees are published as rates (percentages).** The problem stated in Part IV — that without a denominator of annual income there is no identifying power — does not arise, since the figures are published already divided, and no external income data need be procured. In one office, for instance, the FY2025 realized rates were published by occupation as 40.1% for software development, 37.4% for designers and 37.7% for sales.

Third, **there is a column for “separation not ascertained”.** This is a figure showing whether the operator can observe what becomes of those placed, and is very nearly a direct observation — not a proxy — of the breadth of $\mathcal{G}$ (verifiable shared information) of Chapter 4. It may answer the problem recorded in Part IV as “the greatest weakness of the design is the failure to construct a proxy on the B1 side”. It is self-reported, however, and the possibility that operators who do not follow up enter zero cannot be excluded.

The individual pages of certified operators also carry a field for “MHLW Jinzai Service General Site URL / licence number”, confirming that a join key exists from the list of certifications to the office detail pages.

The extent to which the regressor is prescribed by regulation

The regressor to be used in the test is **the strength b of the refund scheme**, concretely the refund rate and the covered period. Of these, the *covered period* is subject to a regulatory floor in some fields.

Under the certification scheme for medicine, care and childcare, certification from FY2024 requires that a refund on early separation cover the case in which the jobseeker separates within at least six months of placement.

For certified operators in this field alone, the floor on the covered period is a product of regulation rather than of market choice. Keeping certified operators in the sample while using the covered period as a regressor measures compliance with regulation, not the choice of a contract. The reach of this point is limited, however (Remark 15.4).

Remark 15.4 (The limits of application). The reach of this point is limited.

First, it applies only to *certified* operators in medicine, care and childcare. Certified operators number roughly 60–70 actual firms, a very small part of the 30,561 offices.

Second, only the *floor* on the covered period is prescribed; the level of the refund rate, and the setting of a period beyond the floor, are unregulated.

Third, certification is voluntary and may be declined. In August 2026, however, use of a certified operator was included among the requirements for the special treatment relaxing staffing standards, so economic value has been attached to certification. Pressure to comply is strengthening.

The claim that “regulation prescribes the regressor” therefore applies **to part of the sample, and only to the floor of one component of the variable**. It cannot be extended to the field as a whole, or to refund schemes in general.

This belongs to (iii) non-identification, but it is at the same time an opportunity for identification. Because the regulation bears on one field only, a three-group comparison can be constructed.

Group	Field	Certified	What determines the refund
A	medicine, care, childcare	yes	regulation (six months as a floor)
B	medicine, care, childcare	no	market choice
C	other fields	—	market choice

Table15.10: Three groups (A/B/C) by what determines the refund scheme.

B against C is the hypothesis test proper. A against B answers a different question: **how far did the regulation move the market?** If most of group B also offers six months, the regulation is non-binding and merely ratifies what the market had already chosen. If group B centres on

one to three months, the regulation is binding and has imposed a level away from the market equilibrium.

A against B is distinct from the measurability hypothesis, but it is a direct measurement of how far the institutional layer can rewrite Φ , and thus a test of the claim of Section 6.3.

The two sources record at different granularities

There is a more practical but serious problem. The refund scheme and the fee rate must be taken from different sources, and the two record at different granularities.

Source	Content	Granularity
List of certified operators	the roster of certified operators, published by MHLW	firm $\times$ field
Jinzai Service General Site	placements, separations, fee rates by occupation	office

Table15.11: The refund scheme and the fee rate: their sources and granularities.

In the list of certified operators, one firm is registered as separate records for the three fields of medicine, care and childcare, and the certification number differs by field. The licence number on the Jinzai Service General Site, by contrast, denotes a particular office.

For a firm with several offices it is therefore unclear whether the licence number tied to the certification covers the firm's performance as a whole. One large operator is certified for nine occupations in care and six in medicine, but there is no guarantee that office-level fee rates by occupation cover that granularity (the office checked in Section 15.3.5 had only five occupations).

This error correlates with size: the larger the operator, the greater the mismatch, and certified operators are disproportionately large. It is a systematic bias.

The three-way split of occupational classifications

Three classifications must be joined, and no correspondence table exists.

System	Division	Example
Employment Placement Business Report	about 110 occupations	nurse, assistant nurse
Jinzai Service General Site	numeric codes	009, 010, 017, 033, 048
Certification scheme	fine divisions by field	nursing occupations, rehabilitation professionals

Table15.12: The three systems of occupational classification and how each divides.

The certification scheme notes that "nursing occupations" includes nurses, assistant nurses, public health nurses and midwives, but the others are unclear. The system of numeric codes does not match the ordering of occupations in the Employment Placement Business Report. Building a correspondence table is within the range of what labour can solve.

Constraints on an analysis restricted to certified operators

The list of certified operators runs to about 100 entries. Since certification is by firm $\times$ field and one firm may appear in several fields, the number of actual firms is estimated at 60 to 70.

An analysis using certification as a regressor cannot be carried out at this scale. Only part of the 307-firm sample of Section 15.3.5 is certified, and no power for a between-group test is obtained. As stated in Section 12.9, the completeness of a frame does not guarantee the size of a sample.

This is, however, **a constraint only on analyses using certification as a variable.** The relation between the realized fee rate and the turnover rate can be measured for all operators irrespective of certification, and was in fact measured across 307 firms.

The sign is not uniquely determined

What family 5 sets out to test is the claim that the measurability of Chapter 4 determines the contractual form. Put naively, it becomes “the harder the ex post verification of the outcome in an occupation, the stronger the refund scheme”, but this form ignores a force running the other way.

Two forces follow from the optimal strength $b^* \propto 1/(1 + r\sigma^2 k)$ of Example 4.5.

- **B1 (numerator, the incentive side):** the more the operator’s effort bears on retention in an occupation, the more valuable it is to load π onto y . The refund becomes *stronger*.
- **B2 (denominator, the risk side):** the more separation is determined by factors outside the operator’s control, the more loading it onto y merely transfers risk. The refund becomes *weaker*.

The naive formulation sees only B1. In many occupations the two covary in the same direction, so **the sign is not uniquely determined.**

Remark 15.5 (This is an unresolved problem of the field as a whole). The indeterminacy of the sign is not a defect of setup specific to this work.

The risk–incentive trade-off (RIT) has been the subject of a quarter-century of dispute and is not empirically supported. [24, 25] examine 26 empirical studies and find that only four report the negative relation the theory predicts; rather, a *positive* correlation between uncertainty and outcome contingency is observed. Prendergast’s own explanation is that in uncertain situations the firm cannot ascertain how workers spend their time, so it delegates authority and uses output-based pay to constrain that discretion. The force called B1 in this section thus already exists as a formalized rival theory.

The dispute continues. Supporting results are beginning to be reported in the laboratory, and it has been argued that RIT holds robustly once expected-utility theory is abandoned for rank-dependent utility.

This hypothesis therefore belongs to (ii) ill-posedness of the hypothesis under the classification of Section 13.2. To put it in testable form, B1 and B2 must be separated and operationalized.

A test from the price side

The formula for b^* cannot itself be tested, since the required quantities are not available.

Required quantity	Symbol	State
Strength of outcome contingency	b	the refund rate and period; the format differs by operator
Noise in the outcome	σ^2	only uncontrollable variation must be extracted
Informativeness of effort	numerator	no proxy could be constructed
Risk aversion	r	requires experimental methods

Table15.13: Quantities required to test b^* , and their state.

σ^2 is not something aggregation suffices for. What is needed is **the variance of the part the operator's effort cannot control**. The variance of turnover by occupation mixes in the operator's screening and its skill at supporting retention: controllable and uncontrollable variation are not separated.

Instead, **the response on the price side** is examined. If the measurability hypothesis is correct, then in occupations where the outcome is more uncertain the operator should be forced to set a lower fee.

Registered implication: the realized fee rate by occupation is negatively correlated with the turnover rate of that occupation.

Alternative: the fee rate is determined by the income level of the occupation and by commercial custom, and is unrelated to turnover.

Result of the test

Using the detailed search of the Jinzai Service General Site, occupation, realized fee rate and turnover rate were retrieved together. The sample is Tokyo, licence numbers 13-Yu, four occupations.

Occupation	n	Fee rate	Median	SD	Turnover	r	Share at custom
Nurses	82	23.0%	21.9	6.9	12.8%	-0.171	40%
Physicians	101	22.9%	20.0	6.2	2.3%	+0.015	64%
Institutional care	67	22.2%	22.1	5.3	6.4%	-0.103	45%
Childcare workers	57	25.6%	25.0	5.3	5.8%	+0.025	54%
Total	307	23.3%	22.2	6.1	6.7%	-0.095	—

Table15.14: Realized fee rates and turnover rates by occupation, and the result of the test (307 firms).

The last column is the share of operators whose fee rate is exactly 20/25/30%.

The coverage of the sample is recorded in Remark 15.7, and a secondary observation in Remark 15.8.

The implication is not supported. The regressor does not move. This is case (vi) of Section 13.2 (Remark 13.1). In all four occupations the correlation is insignificant ($|t| \leq 1.55$) and the sign is not consistent. Pooling the 307 firms gives $r = -0.095$, $t = -1.67$.

Controlling for size does not change this. The partial correlation controlling for the log of open-ended placements is -0.084 ($t = -1.19$), and on residuals after removing occupation means, -0.055 ($t = -0.78$).

The asymmetry across occupations

The contrast across occupations matters more.

	Range	Ratio
Fee rate (occupation means)	22.2 – 25.6%	1.16×
Turnover rate (occupation means)	2.3 – 12.8%	5.49×

Table15.15: Dispersion across occupations: fee rate against turnover rate.

Although turnover differs by a factor of 5.5, the fee rate differs by only 1.16. The correlation at the level of occupations is $r = -0.102$, essentially zero.

Fee rates also concentrate on a few values. Operators at exactly 20%, 25% or 30% make up 40 to 64% of the whole, and 90% fall within 15–35%. Concentration is highest (64%) for physicians, where retention is most stable.

Remark 15.6 (The adjustment the theory posits is not operating). The $b^* = 1/(1 + r\sigma^2k)$ of Example 4.5 predicts that the strength of outcome contingency is adjusted to the noise in the outcome. The observation of this section shows that **this adjustment does not appear in prices**.

This differs from every one of (i)–(iv) in Section 13.2; it is case (vi).

(i) absence of measurement	does not apply; figures for 307 firms are published
(ii) ill-posedness of the hypothesis	applies separately, but the observation here is independent of it
(iii) non-identification	does not apply; confounders are controlled
(iv) access constraint	does not apply; the data were reached

Table15.16: The observation against the taxonomy of verification obstacles (none of (i)–(iv) applies).

The regressor does not vary within the range the theory posits. If price is set by custom, then however precisely σ^2 is measured, b does not respond. In being soluble **neither by refining the theory nor by adding data**, this obstacle differs in kind from all the preceding ones.

It is at the same time a descriptive finding. In placement for medicine, care and childcare, the fee rate is not adjusted to the risk in the outcome and concentrates on the customary values of 20–30%.

Remark 15.7 (Limits of the sample in this section). The sample is confined to **Tokyo, licence numbers 13-Yu, four occupations**. It does not represent the country.

First, the head offices of nationwide operators concentrate in Tokyo, so large operators are over-represented; several have more than 1,000 open-ended placements.

Second, Tokyo offices of operators licensed by other prefectural labour bureaux are excluded.

Third, **the detailed search allows occupations to be specified only for medicine, care and childcare**. The same test cannot be constructed for other fields such as information and communications or sales.

Fourth, operators that display fees as amounts rather than rates are excluded; across the four occupations this is about a quarter of the 307 firms.

That the sample is Tokyo, however, is unlikely to weaken the conclusion. It is the region with the highest density of operators in the country, so **concentration on customary values is observed even in the most competitive market**.

Remark 15.8 (A secondary observation: size and turnover). The correlation between size (the log of open-ended placements) and the turnover rate differs greatly by occupation.

Occupation	r (size, turnover)
Nurses	−0.009
Physicians	+0.191
Institutional care	+0.431
Childcare workers	+0.331

Table15.17: Correlation of size (log of open-ended placements) with turnover, by occupation.

Care and childcare show a strong positive correlation: the larger the operator, the higher the turnover. Both fields fall under the certification scheme (Section 15.3.5), which suggests a structure in which placing in volume and achieving retention are hard to combine. This work has not tested that.

A candidate proxy for B1: size and turnover

Table 15.13 records that no proxy for the informativeness of effort — the B1 side — could be constructed. The observation of Remark 15.8 is a candidate.

■ **Mechanism.** If growth in size thins the input per placement, then **turnover rises only in occupations where the operator’s effort bears on retention**. Where effort does not bear, turnover is unchanged even as input thins. The occupation-level “correlation of size with turnover” therefore gives an ordering of the magnitude of B1.

The observed ordering is as follows (Remark 15.8).

Occupation	r (size, turnover)	Implication for B1
Institutional care	+0.431	effort bears
Childcare workers	+0.331	the same
Physicians	+0.191	intermediate
Nurses	−0.009	effort does not bear

Table 15.18: The ordering of occupations if the correlation of size with turnover is read as a proxy for B1.

Remark 15.9 (The mechanism is not identified). The reading above depends on **the assumption that growth in size thins the input per placement**. If capacity expands in proportion to size, input does not thin and no correlation arises. This work has not tested that assumption.

The same sign can also arise from looser screening: a larger operator pushes through more jobseekers, and the check on fit per placement becomes coarser. **Thinning of input and looser screening cannot be distinguished by this observation**. Both work in the same direction as B1, so the variable can be used as a proxy, but no claim about the mechanism can be made.

There is a literature in labour economics on the relation between size and the quality of placements, which this work has not surveyed. The point here is not novelty but whether an observation already in hand can serve as a proxy for B1.

■ **Registration of the implications.** Following the procedure of Section 11.5, the following are fixed before testing.

Registered implication (1): the occupation-level “correlation of size with turnover” preserves the ordering of Table 15.18 in samples from other prefectures — that is, positive for institutional care and childcare workers, near zero for nurses.

Alternative: the correlation is an artefact of circumstances specific to Tokyo, in particular the concentration there of the head offices of nationwide operators, and does not appear in other prefectures.

Registered implication (2): the stronger the correlation of size with turnover in an occupation, the stronger its refund scheme, because in occupations with large B1 it is more valuable to load π onto y .

Alternative: the strength of the refund corresponds to the level of turnover; that is, B2 (the risk side) dominates, and occupations with higher turnover have weaker refunds.

If (2) holds, a route is obtained for separating B1 from B2 in the measurability hypothesis, whose sign Section 15.3.5 found indeterminate. (1) is a precondition, confirming that the correlation is not an artefact of the sample.

■ Matters fixed in advance.

Item	Content
Population	offices publishing a business report on the Jinzai Service General Site; a prefecture is specified and identified by the prefix of the licence number
Occupations	the medicine, care and childcare fields specifiable in the detailed search, including at least the four occupations of Table 15.14
Definition of size	the log of open-ended placements, as in Section 15.3.5
Floor	offices with zero open-ended placements are excluded, the turnover rate being undefined
Exclusion	offices displaying fees as amounts are excluded for (2) only; they are used in (1)
Identification	records are consolidated by licence number, taking the firm rather than the office row as the unit (Remark 15.11)

Table15.19: Matters fixed in advance for the two implications.

(1) is tested in Section 15.3.5 and (2) in Section 15.3.5.

Remark 15.10 (The test must not use the same sample). The correlations of Table 15.18 were observed *after the fact* in the 307-firm sample. Testing (1) on the same sample would be an error of the same form as the retrospective narrativization of Chapter 12. What (1) demands is reproduction out of sample; until that is met, the correlation remains a comparison (Section 11.4).

The ordering of occupations by the correlation of size with turnover

The detailed search of the Jinzai Service General Site was run with the prefecture set to the whole country for the four occupations. Since the sample of Section 15.3.5 was confined to Tokyo, **the part excluding Tokyo is out of sample.**

Occupation	Search results	Rows extracted	After consolidation
Institutional care	1,075	1,075	464
Childcare workers	552	552	200
Physicians	479	479	269
Nurses and assistant nurses	1,322	1,322	469

Table15.20: The sample used for (1). The unit is office rows; after consolidation, firms.

The rows extracted agree with the search results. The licence-number letters are three: Yu (fee-charging) 3,221, Mu (free placement office) 200, and Toku 7.

The total of 1,402 after consolidation is a count of firm–occupation pairs, not of firms. Distinct firms number 1,130, with 226 appearing in more than one occupation. The duplication does not affect the correlations, which are taken by occupation, but **the figure must not be cited as a number of firms.**

Remark 15.11 (Rows are offices; values are firms). Each row of the search results denotes an office, but **placements, turnover rates and fees are values at the level of the licence number, that is the firm, replicated across all of that firm’s offices.**

Among physicians, a single company accounts for 134 of the 479 rows. Taking correlations without consolidating gives that company 134 times the weight: the median of open-ended placements moves from 14 to 6,430 and the sign of the correlation turns from positive to negative. **Consolidation by licence number is not preprocessing but a procedure that determines the conclusion.**

■ **Comparison with the reference values.** The reference values of Table 15.18 were computed on a composition excluding amount-displayed fees. On the Tokyo portion under the same composition, all four occupations agree to within 0.02 in correlation and 2 in sample size. **Under the registered rule (including amount-displayed fees), however, institutional care gives +0.079, which does not agree with the reference value of +0.431 (Remark 15.12).**

■ **Results.** Both the sample excluding Tokyo and the nationwide sample including it are computed. The former is the pre-registered out-of-sample test; the latter is the best estimate of the ordering.

	Excluding Tokyo			Nationwide		
Occupation	<i>n</i>	<i>r</i>	<i>t</i>	<i>n</i>	<i>r</i>	<i>t</i>
<i>As registered (amount-displayed fees included)</i>						
Institutional care	176	+0.246	+3.34	244	+0.187	+2.97
Childcare workers	59	+0.163	+1.25	109	+0.265	+2.85
Physicians	54	+0.159	+1.16	114	+0.139	+1.48
Nurses	150	+0.067	+0.81	242	+0.029	+0.44
Order	care > child. > phys. > nurse			child. > care > phys. > nurse		
<i>Excluding amount-displayed fees (the composition of the reference values)</i>						
Institutional care	122	+0.221	+2.48	169	+0.228	+3.03
Childcare workers	42	+0.645	+5.34	82	+0.484	+4.94
Physicians	43	+0.126	+0.81	93	+0.146	+1.41
Nurses	116	+0.070	+0.75	186	+0.045	+0.61
Order	child. > care > phys. > nurse			child. > care > phys. > nurse		

Table15.21: The correlation of size with turnover, under four combinations of sample and exclusion rule.

■ **Verdict.** The facts are stated separately.

- (1) **The ordering agrees with the reference values in only one of the four combinations.** There are four combinations of sample (excluding Tokyo / nationwide) and exclusion rule (including / excluding amount-displayed fees), and the reference ordering is reproduced only under “excluding Tokyo” with “including amount-displayed fees”. In the other three the top two swap.
- (2) **The ordering is not reproduced nationally.** The best estimate of the ordering is the nationwide sample including Tokyo, where childcare workers exceed institutional care under both rules.
- (3) **The qualitative division depends on neither sample nor rule.** That institutional care and childcare workers are positive and nurses near zero holds in all four combinations. Nationally both institutional care and childcare are significant ($|t| > 2.8$) and nurses have $|t| < 0.7$.

The verdict is partial support. What is supported is (3): **the qualitative division of positive, positive and near-zero alone.** The ordering itself is not reproduced.

As recorded in Remark 15.12, the registered rule did not match the composition of the reference values, so it cannot be said that the one combination agreeing in (1) was chosen in advance.

Remark 15.12 (The registered exclusion rule did not match the composition of the reference values). Table 15.19 laid down that offices with amount-displayed fees would be used in (1). But the reference values of Table 15.18 were computed on a sample excluding them (Remark 15.7).

The registered rule was not the composition that produced the reference values being compared against.

This is why both are reported and the verdict confined to the part independent of the rule. **The defect would not have arisen had the composition of the reference values been checked at the time of registration.**

Remark 15.13 (Limits of the sample). Two points are recorded.

First, the displayed turnover rate cannot be computed from the placements and separations in the same row. For physicians the two agree in only 115 of 479 rows, the periods differing. The turnover rate is used as the site's own aggregate. This matches the composition of Section 15.3.5 and does not impede comparison, but **it is recorded that size and turnover are not quantities over the same period.**

Second, the distribution of size differs between Tokyo and the rest. The median of open-ended placements is 23 in Tokyo and 5 elsewhere. The national estimate combines the two, and the distribution of size is bimodal.

The strength of the refund

■ How strength is measured was fixed first. At registration (Table 15.19) the regressor was written only as “the refund scheme is strong (a higher refund rate, or a longer covered period)”, with no measurement laid down. Deciding after reading the documents would be fitting after the fact, so the following was fixed before reading them.

Let t be the number of days to separation and $r(t)$ the refund rate at that point, and take as strength

$$\text{strength} = \frac{1}{180} \int_0^{180} r(t) dt. \quad (15.2)$$

The window is 180 days because **the site's turnover rate is defined over six months from placement**, aligning the observation windows of regressor and regressand. Where a rate is written as a range the lower bound is taken, and where intervals overlap the higher rate is taken. The details are recorded in `data/raw/jinzai/henreikin/README.md`.

■ Reading the documents. The 907 PDFs referenced from the refund field of the office detail pages were read. The 132 carrying no text layer were put through OCR, and places OCR could not settle were checked by eye. No document was illegible.

Category	Count	Content
Tiers legible	839	strength computed
No scheme, stated explicitly	11	strength 0
Scheme declared, rate 0%	4	strength 0
Not disclosed, “as set out in the contract”	46	missing
Link does not mention refunds	7	missing
Total	907	854 with a computable strength

Table15.22: Result of reading the 907 refund documents.

Remark 15.14 (Not disclosing is not not refunding). The 46 that say “details are set out in the contract” and the 7 whose link points to a fee schedule mean **that the strength is unknown, not that it is 0**. Setting them to 0 would understate strength. By contrast the 11 that state “no scheme” and the 4 that declare a scheme while filing a rate of 0% are *observations* of strength 0. The two are counted separately.

Everything below was tested both excluding the missing values and setting them to 0; the verdict did not change.

■ **Results.** The 854 with a computable strength are tabulated by occupation. The treatment of the 53 missing values is as in Remark 15.14.

Occupation	n	Median	Mean	B1
Institutional care	359	26.2	27.8	+0.246
Childcare workers	155	25.0	23.9	+0.163
Physicians	199	33.3	33.4	+0.159
Nurses	367	25.8	28.0	+0.067

Table15.23: The strength of refunds by occupation (median and mean), and the B1 measured in (1).

Implication (2) is rejected. The ordering by B1 is institutional care, childcare workers, physicians, nurses; the ordering by strength is physicians, institutional care, nurses, childcare workers. The rank correlation is $\rho = 0.000$.

The alternative is not supported either. The correlation of strength with turnover is $r = -0.037$ (1,080 observations) and with size (the log of open-ended placements) $r = -0.079$ (571 observations), both near zero. Neither B1 (thinning of input) nor B2 (the risk side).

The difference across occupations appears only for physicians, and is not explained here (Remark 15.15). The difference is also visible only once strength is measured; the presence or absence of a scheme cannot detect it (Remark 15.16).

Remark 15.15 (On physicians alone standing out). In Table 15.23 the strength for physicians is about 30% higher than for the other three occupations. Since physicians rank third of four on B1, this cannot be explained by B1.

This work does not explain it. It is an observation that was not registered in advance, and supplying an explanation would be of the same form as the retrospective narrative listed in Section 19.2. Only the fact of the observation is recorded.

Remark 15.16 (Presence is no proxy for strength). Looking only at presence, among the 904 single-occupation firms, 84.3% of physicians, 84.0% of nurses, 86.5% of institutional care and 87.0% of childcare report “yes”, with $\chi^2(3) = 1.13$: no difference by occupation. **Having a scheme is the norm, and as a binary the variable has almost no variance.**

Only on measuring strength does the difference across occupations (Table 15.23) appear. That difference is likewise unrelated to B1, but **proxying by presence cannot even detect whether a difference exists**. That four offices file a rate of 0% while displaying the scheme as “yes” is direct evidence that the two are independent.

15.3.6 Assignment of obstacles

The principal obstacle for family 2 is (iv) access constraint: data of the required granularity cannot be reached. For family 5 it is (ii) ill-posedness of the hypothesis as regards the formula for b^* , and (vi) non-operation of the adjustment as regards the test on the price side.

For family 5 the data source proved favourable once the route was checked (Section 15.3.5). Indeed the test from the price side was carried out on 307 firms (Section 15.3.5). Verification nevertheless fails to advance because, for the formula for b^* , the proposition does not determine the sign uniquely, and on the price side the regressor does not vary within the range the theory posits.

This distinction has a practical implication. Where the obstacle is (iii), (iv) or (v), investing labour produces progress; where it is (i), (ii) or (vi), the same labour produces nothing. **Unless which one it is is determined first, labour is misallocated.**

15.4 Comparison of the two families and its methodological implications

	Family 2	Family 5
Implication	supported as to level	untested
What the aggregate layer shows	a two-order-of-magnitude difference in κ	almost nothing
Alternative	partly separable	not begun
Per-firm data	unavailable (confidentiality)	available (Jinzai Service General Site)
Remaining task	measuring ϕ_{cog}	sampling design and retrieval of per-firm data

Table15.24: Families 2 and 5 compared: state of verification, availability of per-firm data, remaining tasks.

Family 5 is in a good position if per-firm data can be reached, but at present, with only the aggregate layer examined, it yields a weaker conclusion than family 2. **The completeness of a frame and the richness of its variables are independent, and the prospects of an investigation cannot be assessed without checking both.**

One further structure common to both families was observed.

Proposition 15.17 (A trade-off between the aggregate and per-firm layers). In the two domains treated here, a complete frame founded on regulation supplies the structural variables (κ , contractual form) only as aggregates, while the surplus variables (ϕ_{cog} , ϕ_{barg}) can be observed only in a selected subsample.

To reach ϕ_{cog} in family 2 one must descend to the annual securities reports of listed firms that disclose breakage income separately; but that means returning to the 0.1% sample of Section 12.9, and the disclosure-selection bias of Section 12.3 revives. The constraint that **structure is visible in the complete frame while surplus is visible only in a selected sample** holds at least in these two domains.

15.5 The structure of the data source

15.5.1 The detailed search

The site's detailed search allows occupation, realized fee rate and turnover rate to be given as search conditions, and displays them in the result listing when they are given. The samples of Sections 15.3.5 and 15.3.5 come from this route.

The occupations specifiable as search conditions are limited to nine in medicine, care and childcare. The same retrieval is impossible in other fields such as information and communications or sales. The restriction of this work to four occupations follows from this constraint.

The detail page of each office carries more than the listing: placements and separations by fiscal year, all occupations the office has filed, and a reference to the refund-scheme filing. The refund data of Section 15.3.5 came from there.

15.5.2 Outcomes of the three routes

The outcomes of the three routes listed in the investigation plan (Part III) are recorded.

■ **Measuring the publication rate.** In the search results for the four occupations, about three quarters of operators display fees as rates and about a quarter as amounts. Publication itself is carried out by nearly all operators —good progress for the first year of the requirement.

■ **Describing turnover by field. Achieved.** As in Section 15.3.5, turnover rates for the four occupations were obtained at the level of the operator. The occupation means are 2.3% for physicians, 5.8% for childcare workers, 6.4% for institutional care and 12.8% for nurses — **a spread of a factor of 5.5.**

■ **Measuring how binding the regulation is. Achieved.** The strength of refunds was measured for 907 documents and implication (2) tested in Section 15.3.5. As to fee rates, the standard deviation among certified operators is 2.2, smaller than the 7.0 of the uncertified (Section 15.3.5). Since the certification conditions say nothing about fee rates, this is an indirect effect.

15.5.3 The mismatch of recording granularity

The mismatch between firm and office granularity noted in Section 15.3.5 appears in stark form in the national sample of Section 15.3.5 (Remark 15.11). Rows of the search results denote offices, but placements, turnover rates and fees are values at the level of the firm, replicated across all of that firm's offices.

15.6 What was reached and what was abandoned

15.6.1 Subjects abandoned

- (1) Measuring ϕ_{cog} for family 2 **with the data of this work.** Lacking (2) and (3) of Proposition 14.13, it cannot be estimated without reaching data carrying individual usage records and an exogenous shock. Quantities of the same kind have been measured by others (Section 14.6.1); what is abandoned is *measurement from published statistics*, not measurement as such.

- (2) Testing the formula for b^* in family 5. Obtaining b as a continuous quantity, and extracting from σ^2 only the uncontrollable part, are both impossible ((ii) ill-posedness of the hypothesis).
- (3) Continuing the test from the price side. Section 15.3.5 established the absence of correlation. Since the regressor concentrates on customary values and does not vary, enlarging the sample will not change the result ((vi) non-operation of the adjustment).

15.6.2 Subjects possible under conditions

- (5) Extension to other prefectures. The sample of Section 15.3.5 is confined to Tokyo (Remark 15.7). Since the detailed search takes a prefecture, reproduction elsewhere is possible by the same procedure.
- (6) Testing the relation between size and turnover. Two implications were registered in Section 15.3.5; the test requires samples from other prefectures.
- (7) Collecting refund rates and periods. Since the published format differs by operator, this is a matter of a person reading them one at a time. Restricted to the 60–70 certified operators it is of workable scale.

15.6.3 What families 2 and 5 reached

	Family 2	Family 5
Verified	a two-order-of-magnitude difference in κ ; the alternative partly separated	fee rate and turnover uncorrelated; concentration on customary values
Not verified	ϕ_{cog} (with the data of this work)	the formula for b^*
Principal obstacle	(iv) access constraint	(ii) ill-posedness and (vi) non-operation of the adjustment
Secondary obstacles	confidentiality of per-firm data	recording granularity, small population, failure of identification
Condition for resuming	reaching individual-record data	none; (vi) is not solved by labour

Table15.25: Summary of what families 2 and 5 reached.

For family 2, the part of the registered implication concerning levels was supported.

For family 5, the formula for b^* was not reached, but **the test on the price side yielded a descriptive result**. In placement for medicine, care and childcare, although occupation means of turnover differ by a factor of 5.5, fee rates differ by only 1.16 and concentrate on the customary values of 20–30%.

Beyond identifying the structural defect of the hypothesis in Section 15.3.5, showing in Section 15.3.5 that **the adjustment the theory posits does not operate in reality** is a result that could not have been reached by adding data.

15.7 A quantity that is not tabulated: the credit position of sole proprietors

In attempting to measure κ for sole proprietors, two statistics were consulted.

Statistic	State	Verdict
Survey of the Economy of Individual Enterprises	surveyed assets and liabilities, but discontinued in FY2019	<i>loss</i> of measurement
Basic Survey of Small and Medium Enterprises	a questionnaire for sole proprietors exists, but the tables of assets and liabilities cover corporations only	surveyed but <i>not tabulated</i>

Table15.26: The state of the two statistics bearing on κ for sole proprietors.

The first is (i) absence of measurement, but differs from the usual case in that **it was measured in the past**. Data through FY2018 exist, so the past can be analysed.

The second is more peculiar. Since a questionnaire exists, responses are very likely collected; but they do not appear in the published tables. This is neither (i) nor (iv) but a fifth obstacle: **the choice of what to tabulate**.

Remark 15.18 (The choice of tabulation determines what is observed). The selection bias treated in Part III concerned the selection of the sample and the measurement of variables. What this section shows is **selection at the stage of tabulation**.

Which cross-sections are published is decided by the compiler of the statistic. The judgement to tabulate assets and liabilities for corporations and not for individual enterprises is presumably grounded in the needs of SME policy, but the result is that **the credit position of sole proprietors is not observed**.

To “the completeness of a frame does not guarantee the size of a sample” of Section 12.9 must be added: **being in the sample does not guarantee being tabulated**.

Chapter 16

Evidence From the Corporate Statistics: κ and CCC Across Firms

16.1 Position of this chapter

The κ and CCC defined in Part I are applied to the industry and size tabulations of the Financial Statements Statistics of Corporations by Industry. Two cross-sections are taken.

The cross-section by size examines how the direction of credit corresponds to size class. Section 6.3 writes the credit layer as an image of bargaining power; if so, smaller firms should be the providers of credit.

The cross-section over time examines whether change in CCC indicates the state of an industry — **whether demand is tight**. But “the industry is hot” is not one-dimensional, and CCC touches only one of its axes. That is settled in the next section.

To state the conclusion first, **both implications failed**: the direction of credit has the opposite sign, and the leading property vanished entirely.

16.2 Hot and cold are not one-dimensional

Before measuring, what is to be judged must be settled. The $g^* = (m + d - I/r)/\text{CCC}$ of (5.6) consists of two separate axes, a numerator expressing profitability and a denominator expressing liquidity. Likewise, **an industry being hot and entering it being profitable are different things**.

Even where demand is tight, if an intermediary controls access to that demand, the surplus does not remain with the entrant. Chapter 17 finds the price of T_{acq} reaching 20 points. Two axes are therefore needed.

Axis	Indicator	Measurability
Tightness of demand	change in κ , growth rate	measurable in the corporate statistics
Ease of taking surplus	the price of T_{acq} , concentration of ϕ_{barg}	no cross-industry indicator exists

Table16.1: Two axes for judging hot and cold, and their measurability.

Only the former can be handled here. Of the four quadrants, up and down can be distinguished but not left and right.

Remark 16.1 (Why the horizontal axis cannot be measured, structurally). This is not a shortage of sources but an asymmetry of classification.

Industry classifications are cut from the entrant's point of view: a firm in information and communications is tabulated there. But **intermediaries span industries**. A payment platform belongs to information and communications while intermediating for retail, food service, travel and education alike.

Summing $\phi_{\text{barg}}^{\text{intermediary}}$ by industry therefore only books the intermediary's take in one industry, and **which industry it was taken from cannot be recovered**. In the notation of Section 2.6, an intermediary is one of the party indices i ; it is neither a type of Φ nor an industry.

The entrant's margin already reflects the leakage to intermediaries, but whether it is due to that leakage or to something else cannot be separated from industry aggregates. Separating it needs a cross-section seen from the intermediary side, and why that is unavailable is set out in Section 17.5.1.

Theoretical quantities in this chapter

κ_C	Eq. (3.2)	DSO; 35–71 days by size
κ_S	Eq. (3.2)	–DPO; DPO is 21–44 days by size
W, r	Chapter 5	working capital, sales
CCC	Eq. (5.1)	37 days for all industries; 5–210 days by industry
g^*	Eq. (5.6)	illustrated with assumed values, m not obtained here
m	Chapter 5	not obtained

Table16.2: Theoretical quantities in this chapter.

16.3 Data source

The time-series data of the Financial Statements Statistics of Corporations by Industry (e-Stat table 0003060791) were used. It is a designated statistical survey under the Statistics Act, and

the annual survey covers all profit-making corporations. Capital classes extend down to below two million yen, so the bias towards listed companies raised in Section 12.9 does not arise.

Six items were obtained: notes receivable, accounts receivable, finished goods and merchandise, notes payable, accounts payable and sales. Section 16.10 adds inventory (period end) for a reconciliation. Raw figures rather than ready-made turnover indicators were obtained so that the definitions of Chapter 5 could be applied directly, without depending on whether the statistic's own denominator is sales or cost of sales.

$$\text{DSO} = \frac{\text{notes receivable} + \text{accounts receivable}}{\text{sales}} \times 365 \quad (16.1)$$

$$\text{DPO} = \frac{\text{notes payable} + \text{accounts payable}}{\text{sales}} \times 365 \quad (16.2)$$

$$\text{CCC} = \text{DSO} + \text{DIO} - \text{DPO} \quad (16.3)$$

These express the credit positions of (3.2) in days. That is,

$$\kappa_C \leftrightarrow \text{DSO}, \quad \kappa_S \leftrightarrow -\text{DPO}, \quad \sum_i \kappa_i \leftrightarrow \text{DSO} - \text{DPO} \quad (16.4)$$

A firm with $\sum_i \kappa_i > 0$ is a provider of credit. The working capital of (5.1) is $W = (\text{DSO} + \text{DIO} - \text{DPO}) \times r/365$.

16.4 Size and the direction of credit

16.4.1 The implication tested

Registered implication: the smaller the size class, the larger $\text{DSO} - \text{DPO}$; that is, smaller firms are the providers of credit.

The ground was the claim of Section 6.3 that the credit layer is an image of bargaining power: small firms without bargaining power should be the ones kept waiting for payment.

16.4.2 Result

Size (capital)	DSO	DIO	DPO	DSO–DPO
1bn yen and over	71	20	44	27
100m – 1bn	61	20	47	15
50m – 100m	47	20	37	10
20m – 50m	49	15	39	10
10m – 20m	53	18	35	18
Under 10m	35	12	21	14

Table16.3: DSO, DIO, DPO and DSO–DPO by capital class.

All industries excluding finance and insurance, FY2024, in days




Figure 16.1: DSO and DPO by capital class (all industries, FY2024). Both are short for the small classes.

The implication is rejected; the sign is the opposite. Firms of 1bn yen and over sit at 27 days and those under 10m at 14, so large firms are the providers of credit. The ordering is consistent over the ten years 2015–2024, and the large-firm figure has in fact widened from 19 to 27 days.

16.4.3 The error in the reasoning

The error can be located. **It was overlooked that DSO and DPO both move.**

Small firms are kept waiting on receivables (DSO works against them). But at the same time they are not trusted by suppliers either and cannot stretch payables. As Figure 16.1 shows, the DPO of the under-10m class is 21 days, less than half the 44 of firms of 1bn and over.

A lack of bargaining power acts on both sides: collection is slow but payment is fast. On balance, **weak bargaining power does not appear in the net credit position.**

Proposition 16.2 (Extending credit requires capital). A net provider of credit, that is a party with $\sum_i \kappa_i > 0$, satisfies $E > A$.

Proof. By the fundamental inequality of Proposition 5.2, $W \leq E - A$. Since $W = \sum_i \kappa_i + \text{Inv}$ with $\text{Inv} \geq 0$, $\sum_i \kappa_i > 0$ implies $W > 0$ and hence $E - A > 0$. $\square$

Under the $E \approx 0$ of Chapter 8, the requirement was the reverse, $\kappa < 0$ (Corollary 5.4). Tolerating collection as long as a DSO of 71 days is possible because there is capital behind it.

16.4.4 The sign splits by industry

Industry	1bn and over	Under 10m	Difference
Information and communications	64	15	+49
Construction	67	19	+48
All industries	27	14	+13
Manufacturing	26	24	+2
Employment placement and dispatch	24	25	−1
Retail	−16	6	−22

Table16.4: DSO–DPO by industry, comparing capital of 1bn yen and over with under 10m. Only retail reverses sign.

DSO–DPO, FY2024, in days

Only retail reverses sign. The −16 days of large retailers is exactly the negative credit position described in Part I: cash collected immediately and purchases paid in arrears. Here large firms are the recipients of credit and small retailers the providers.

The +49 days of information and communications and construction shows large firms bearing long advances in the position of contractor — the territory of families 1-3 (one-off arrears) and 5-3 (cost plus) of Part II.

The claim that “larger firms are the providers of credit” is therefore itself only an average across industries. What decides is not size but **the Φ dominant in that industry**. The lesson recorded in Part IV — that levels of κ differ structurally across industries so cross-sectional comparison is meaningless — holds on the size axis too.

16.4.5 Implication for Chapter 8

Chapter 8 deduced that a solo business without assets requires $\kappa_C < 0$. The data show that requirement is not met: the under-10m class has DSO of 35 days and DPO of 21, a net 14 days as a provider.

The form differs from what was envisaged at registration, however. The small classes are not extremely $\kappa > 0$; rather **both DSO and DPO are short**. The transactions themselves are small and short. Rather than lying outside the feasible region, they appear to sit at different coordinates.

16.5 Is CCC a leading indicator of the cycle?

16.5.1 The implication tested

Registered implication: the rate of change of CCC *leads or is synchronous with* the rate of change of sales. In particular, a worsening receivables turnover *leads* a slowdown in sales.

Correlations between the annual change $\Delta CCC(t)$ and the sales growth rate $g(t+k)$ were computed by industry and aggregated over 45 industries, excluding aggregate and short series.

16.5.2 Result

Lag	Mean correlation	Median	Share negative	Sign test
Leading $k = -1$	0.036	0.058	47%	$p = 0.72$
Synchronous $k = 0$	-0.159	-0.228	69%	$p = 0.008$
Lagging $k = +1$	0.072	0.118	31%	—
Lagging $k = +2$	0.085	0.035	44%	—

Table 16.5: Correlation between $\Delta CCC(t)$ and sales growth $g(t+k)$ by lag, over 45 industries.

Synchrony is supported. A worsening CCC and a slowdown in sales occur in the same fiscal year: 31 of 45 industries show a negative correlation, with $p = 0.008$ on a sign test.

Leading is rejected outright. 47% negative is no information at all ($p = 0.72$). Change in CCC carries nothing about the following year's sales.

16.5.3 The synchronous negative correlation may be a denominator effect

The mechanism is likely the reverse of what was assumed. $CCC = W/r$ has sales in the denominator. If sales fall, then unless inventory and receivables fall at the same rate, CCC worsens *by definition*.

When the panel by medium for family 2 was treated in Part IV, this was stated explicitly as the alternative: issuance shrinks and the ratio rises mechanically. Section 16.9 separates it here and shows that **the denominator effect contributes about 17%**.

Claim	Verdict
Change in CCC is synchronous with change in sales	supported; the denominator effect contributes 17% (Section 16.9)
Change in CCC leads change in sales	rejected ($p = 0.72$)
CCC is a leading indicator of the cycle	rejected

Table16.6: Verdicts on three claims about the leading property.

16.5.4 That the sign splits by industry is a finding

Thirteen of the 45 industries have $|r| > 0.4$ in the synchronous correlation.

Strongly negative	r	Strongly positive	r
Pulp, paper and paper products	−0.78	Agriculture and forestry	+0.80
Fabricated metal products	−0.67	Other transport equipment	+0.69
Information and communications	−0.58	Accommodation and food services	+0.54
Production machinery	−0.53	Agriculture, forestry and fisheries	+0.52

Table16.7: Industries with $|r| > 0.4$ in the synchronous correlation (four strongly negative, four strongly positive).

The strongly negative industries are all make-to-order with large inventories and receivables, matching the conditions under which the denominator effect is strong.

What matters is that positive industries exist. **A pure denominator effect would make every industry negative.** The positive industries have a structure in which CCC grows as sales grow — they build working capital in order to grow — which is the region where the g^* of (5.6) binds.

Change in CCC therefore contains components other than the denominator effect. They cannot be separated in these data.

16.6 An observation not anticipated: the long-run rise in CCC

Separately from the tests, a fact not envisaged in advance was observed.

As Figure 16.2 shows, the CCC of all industries rose from 24.9 days in FY2000 to 37.2 in FY2024, almost monotonically over 25 years. The net credit position also doubled, from 9.8 to




Figure16.2: CCC and the net credit position, all industries excluding finance and insurance.

18.5 days. **The Japanese corporate sector as a whole is moving towards holding more working capital.**

Remark 16.3 (What the aggregate measures). By Corollary 3.6, credit positions closed within the corporate sector cancel on aggregation. The CCC here is therefore not an indicator of inter-firm credit but the sum of the position against the non-corporate sector and inventory, divided by sales. Indeed, $DSO - DPO$ for all industries in Table 16.3 is 13 days, and most of the gap to $CCC = 37.2$ days is inventory. **The dominant component of the aggregate CCC is inventory.** The same reason explains why negative CCC in individual firms does not appear in aggregate.

16.6.1 The growth constraint has not worsened, however

It is tempting to conclude from this that the growth constraint has worsened, since the denominator of (5.6) has risen by half. **That inference is wrong.**

Over the same period the operating margin m improved by 91%, from 2.62% to 5.01%, out-running the worsening denominator. The ratio m/CCC consequently rose from 38.5% to 49.2% (Proposition 14.2).

	FY2000	FY2024	Change
CCC	24.9 days	37.2 days	+49%
m	2.62%	5.01%	+91%
m/CCC	38.5%	49.2%	+28%

Table16.8: Change in CCC, m and m/CCC from FY2000 to FY2024.

What is used here is m/CCC , not the $g^* = (m + d - I/r)/CCC$ of Corollary 5.7 (Remark 14.3), because d and I were not obtained far enough back.

Remark 16.4 (Discipline in handling ratios). CCC is the denominator of a ratio and says nothing on its own about the growth constraint. Drawing a conclusion from a rise in CCC alone is looking at one half of a ratio. **Concluding from one half of a ratio is the same kind of error as the retrospective narrative that Section 12.6 warns against.** The difference is that the latter is an interpretation added after the fact, whereas this is filling a shortfall in measurement with inference.

The cause of the rise in CCC itself remains untested. Which of industrial composition, regulation of payment terms, the decline of trade bills, or the level of interest rates contributes is not separated here. The importance of the question is limited, however, since g^* has not worsened.

16.7 Levels by industry (for reference)

As the lesson of Part IV has it, cross-sectional comparison of levels cannot be interpreted; the figures are recorded as a reference on the distribution of Φ .

Long CCC	Days	Short CCC	Days
Leasing	210	Food services	5
Real estate	89	Water transport	7
Chemicals	80	Other transport	9
Textiles	78	Entertainment	9
General-purpose machinery	77	Motor vehicles and parts	9
Production machinery	64	Electricity	10

Table16.9: Industries with the longest and shortest CCC (top six each).

FY2024, excluding aggregate series

The ordering agrees with the types of Part II. Leasing (family 3-1) and real estate acquire an asset first and recover it over time; food services (family 1-1, immediate exchange) sits at $\kappa \approx 0$. The structure of the differences by industry described earlier appears directly.

16.8 Decomposing the rise in CCC

The rise in CCC is decomposed into its causes. The following implication was fixed in advance.

Registered implication: the rise in CCC owes more to the composition effect than to within-industry change, because industries with long CCC have gained weight.

Alternative: if the within-industry effect dominates, working capital has increased in each industry irrespective of composition.

16.8.1 Shift-share decomposition

By Corollary 5.19, the overall CCC is a sales-weighted mean of industry figures, $CCC = \sum_k w_k CCC_k$. Its difference splits in two.

$$\Delta CCC = \underbrace{\sum_k \bar{w}_k \Delta CCC_k}_{\text{within-industry}} + \underbrace{\sum_k \Delta w_k CCC_{\bar{k}}}_{\text{composition}} \quad (16.5)$$

Taking $\bar{\cdot}$ to be the midpoint of the two periods, (16.5) **holds exactly with no residual**. Substituting $\bar{w} = (w_1 + w_0)/2$ and $CCC_{\bar{k}} = (CCC_1 + CCC_0)/2$, the right-hand side equals $\sum_k (w_{k1} CCC_{k1} - w_{k0} CCC_{k0})$. This is not an approximation introduced for measurement but a consequence of CCC being a ratio (Corollary 5.19).

Period	ΔCCC	Within-industry	Composition
2000→2024 (33 common industries)	+11.4 days	+9.5 days (82.9%)	+1.9 days (17.1%)
2010→2024 (45 common industries)	+10.1 days	+9.9 days (98.1%)	+0.2 days (1.9%)

Table16.10: Shift-share decomposition of ΔCCC into within-industry and composition effects.

The implication is rejected. The within-industry effect accounts for 83% to 98%. Since 2010, where the classification is stable, the composition effect is nearly zero. An explanation by the shift towards services does not hold.

The reason the composition effect is small is also clear. Information and communications rose in weight from 0.2% to 5.8%, pushing CCC up by +3.05 days, while electrical machinery shrank from 6.3% to 2.1% (−1.62 days) and advertising from 5.6% to 0.8% (−1.59 days): **the pushes up and down nearly cancel**.

16.8.2 Components of the within-industry effect

The within-industry effect splits into DSO, DIO and DPO.

Period	DSO	DIO	DPO	Main cause
2000→2024	−0.93	+3.49	+6.90	shortening DPO
2010→2024	+5.26	+4.02	+0.62	lengthening DSO

Table16.11: The within-industry effect split into DSO, DIO and DPO.

In days of contribution to CCC

Against Corollary 5.20, DSO and DPO correspond to τ_C and τ_S and DIO to τ_{inv} . Hence:

Period	Change in contracts $\sum \Delta\tau_i$	Change in operations $\Delta\tau_{\text{inv}}$
2000→2024	+5.97 days	+3.49 days
2010→2024	+5.88 days	+4.02 days

Table16.12: The split into change in contracts and change in operations, by Corollary 5.20.

In both periods, change in contracts accounts for about 60%. The rise in CCC originates in changes to contractual terms rather than in demand or inventory management.

The main cause differs by period, and the path of trade bills explains it.

	2000	2010	2024
Notes receivable	12.1 days	6.1 days	4.7 days
Notes payable	15.3 days	7.2 days	5.1 days

Table16.13: Notes receivable and payable, FY2000, FY2010 and FY2024.

Bills fell to below a third during the 2000s. Notes payable lost 8.1 days, from 15.3 to 7.2, and the 6.9-day shortening of DPO is almost entirely explained by this. **Bills are a device for creating $\kappa_S < 0$, and their disappearance pushed CCC up.**

In the 2010s the decline of bills had run its course and DSO lengthened by 5.26 days. The breakdown is notes receivable -1.37 days and accounts receivable $+6.62$ days. Although bills keep falling, accounts receivable grow by more than enough to offset them. By industry, DSO rose in 28 of 45 industries, with the largest contributions from wholesale (59→70 days), retail (22→28 days), information and communications (62→78 days) and construction (65→71 days).

Remark 16.5 (Not separable from a shift between accounts). The growth of accounts receivable admits two readings: a genuine lengthening of the collection period, and a shift between accounts accompanying the move to electronically recorded monetary claims. Under the latter, the rise in CCC would be only apparent.

Separating them requires the balance of electronically recorded claims, which is not among the survey items of the Financial Statements Statistics. Under the classification of Section 13.2 this belongs to (iii) non-identification.

16.9 Separating the contemporaneous negative correlation: the contribution of the denominator effect

For the contemporaneous negative correlation observed in Section 16.5, we separate information about demand from the denominator effect inherent in the definition $\text{CCC} = W/r$.

In log differences, $\Delta \ln \text{CCC} = \Delta \ln W - \Delta \ln r$. Under a pure denominator effect W would not move. Estimating the elasticity

$$\beta = \frac{\Delta \ln W}{\Delta \ln r} \quad (16.6)$$

therefore makes $1 - \beta$ the contribution of the denominator effect.

The theoretical value of β is 1. By (5.1), $W = \text{CCC} \cdot r$, so if CCC is constant then $\Delta \ln W = \Delta \ln r$. Hence

$$\Delta \ln \text{CCC} = (\beta - 1) \Delta \ln r \quad (16.7)$$

and the departure of β from 1 is exactly the sales elasticity of CCC. Equation (5.10) of Corollary 5.20 decomposes this movement into change in contracts and change in operations.

Registered implication: $\Delta \ln W$ and $\Delta \ln r$ are positively correlated; that is, working capital falls when sales fall.

16.9.1 Result

Pooled estimate (44 industries, 946 observations)	$\beta = 0.830$
Coefficient of determination	$R^2 = 0.076$
Median of industry-level β	0.625
Standard deviation of industry-level β	1.161
Contribution of the denominator effect $1 - \beta$	0.170

Table 16.14: Elasticity of $\Delta \ln W$ with respect to $\Delta \ln r$ (equation (16.6)).

The implication is supported. Working capital moves almost proportionally with sales, and only about 17% is attributable to the denominator effect. The contemporaneous negative correlation of Section 16.5 carries, for the most part, information from the demand side.

Substituting $\beta = 0.830$ into (16.7) gives a sales elasticity of CCC of -0.170 : **CCC rises in phases where sales fall**. The amplification factor of Remark 5.11 is then larger, so the degree to which a downturn in demand is transmitted to cash strengthens in recessions. Calling $1 - \beta$ the “denominator effect” is a computational label, not a claim that CCC has not moved.

16.9.2 Lag structure

By Proposition 3.12, $W(t)$ corresponds to the integral of r over the past τ periods. In industries with long CCC it should also depend on the previous year’s sales. The following was registered in advance.

Registered implication (1): the coefficient β_1 on $\Delta \ln r(t - 1)$ is larger the longer that industry’s CCC.

Registered implication (2): the total elasticity $\beta_0 + \beta_1$ including the lag exceeds the contemporaneous-only estimate of 0.830, so the contribution of the denominator effect is smaller than 17%.

The distributed-lag model $\Delta \ln W(t) = \beta_0 \Delta \ln r(t) + \beta_1 \Delta \ln r(t-1) + \varepsilon$ was estimated on 886 observations across 43 industries.

β_0 (contemporaneous)	0.836
β_1 (one-period lag)	-0.015
Total elasticity	0.821
R^2	0.081

Table16.15: Estimates of the distributed-lag model.

Both implications are rejected. β_1 is nearly zero, and its correlation with CCC is -0.101 — the sign is also reversed. In industries with $\text{CCC} < 30$ days, $\beta_1 = +0.217$; in those with $\text{CCC} \geq 60$ days, -0.267 . The total elasticity falls slightly from 0.830 to 0.821, leaving the contribution of the denominator effect essentially unchanged at 17% to 18%.

Remark 16.6 (Proposition 3.12 is not rejected). Failure to detect a lag structure does not mean equation (3.7) is wrong. The level of CCC is at most 210 days and mostly 30–60 days, which fits within a year at the granularity of annual data. For β_1 to be significant, CCC would have to exceed one year, and no such industry exists.

Quarterly data might allow detection, but the quarterly survey rests on the provisional closings of sampled corporations and is therefore affected by closing adjustments.

Remark 16.7 (The confidence is limited). $R^2 = 0.076$ is extremely low: 92% of the variance of $\Delta \ln W$ is unexplained by $\Delta \ln r$. The dispersion across industries is also large — $\beta < 0.5$ in 14 industries, $\beta > 1.0$ in 13, and negative in some (retail -0.43 , leasing -0.01).

Annual $\Delta \ln W$ mixes in year-end inventory adjustment, timing shifts of large transactions, and changes in accounting policy. **One may say that the denominator effect is not dominant, but this is insufficient to assert positively that the movement is information about demand.**

16.10 The effect of the scope of inventories on the conclusions

The DIO of this chapter uses only “finished goods or merchandise” and excludes work in process and raw materials. Since the Financial Statements Statistics carries “inventories (end of period)” as a separate item, the difference between the two measures the size of the understatement. This section measures that size.

This section is not a test of an implication. It substitutes a different measurement definition on the same data and checks how much the conclusions move.

16.10.1 Two cross-sections

“Inventories (end of period)” was obtained along two cuts. **The two identify different variation**, and must be read separately.

Cut	Coverage, and the variation identifying β	Obs.
Size cut	10 industries $\times$ 11 size strata $\times$ FY2015–2024. β is identified by variation between size strata	1,100
Industry cut	62 industries $\times$ all sizes $\times$ FY2000–2024. β is identified by variation between industries and years	1,364

Table16.16: The two cuts used to obtain inventories.

The estimate of this chapter (Section 16.9) uses 946 observations across 45 industries, **the same design as the industry cut**. The figures that correct this chapter are therefore those of the industry cut; the size cut answers a different question. Unless stated otherwise, the figures below are from the industry cut.

16.10.2 The level is understated roughly twofold

The ratio of DIO has a median of 1.99, and the median CCC rises from 28.0 to 43.9 days. For all industries (excluding finance and insurance), FY2000 goes from 24.9 to 38.2 days and FY2024 from 37.2 to 53.5 days.

The size cut gives a median ratio of 1.89 and a difference with median 4.1 days and mean 8.4 days — comparable in magnitude.

Industry	DIO goods	DIO inv.	CCC goods	CCC inv.
Manufacturing	19.2	50.4	41.6	72.8
Construction	9.7	33.8	39.8	63.9
All industries (excl. finance and insurance)	18.7	35.0	37.2	53.5
Information and communications	7.5	14.6	58.1	65.2
Other services	8.0	12.2	34.4	38.6
Wholesale	24.2	27.6	32.9	36.3
Food services	3.4	6.1	4.9	7.7
Accommodation	6.5	8.9	21.8	24.2
Retail	24.4	25.4	23.3	24.2
Employment placement and worker dispatch	0.2	0.5	32.8	33.0

Table16.17: Difference in DIO and CCC by the scope of inventories (all sizes, FY2024, in days).

The difference is of a different order across industries. CCC rises by more than 30% in manufacturing and construction, while retail and employment placement barely move. **Divergence arises only in industries that hold work in process** — an obvious structure, but large enough to change the ordering of the industry levels in Section 16.7.

16.10.3 The effect on differences is small, but not negligible

Comparing the two definitions for ΔCCC , the industry cut gives a correlation of 0.882 over 1,302 observations, with signs agreeing 86.8% of the time. The size cut (990 observations) gives a correlation of 0.976 and sign agreement of 91.0%: **agreement is lower across industries**. Because the weight of work in process differs by industry, including cross-industry variation produces a gap.

16.10.4 Correcting the definition worsens the elasticity slightly

Equation (16.6) is estimated on a common sample. Since the observations with $W > 0$ differ by definition, **the comparison is restricted to observations with $W > 0$ under both definitions**.

Cut	Definition of DIO	β	R^2	Denom. effect $1 - \beta$
Industry cut	finished goods or merchandise only	0.840	0.073	0.160
$n = 955$	inventories (total)	0.828	0.166	0.172
Size cut	finished goods or merchandise only	0.900	0.092	0.100
$n = 733$	inventories (total)	0.960	0.213	0.040
Estimate in the text	finished goods or merchandise only	0.830	—	0.170

Table 16.18: Difference in the elasticity β by the scope of inventories, on a common sample.

In both cuts, taking inventories at their total more than doubles the coefficient of determination. **Working capital is more nearly proportional to sales when work in process and raw materials are included.**

But the direction in which the denominator effect moves is opposite across the two cuts. In the size cut it falls from 0.100 to 0.040, while in the industry cut it rises slightly from 0.160 to 0.172.

Remark 16.8 (Change the cut and the direction changes). Correction shrinks the denominator effect in the size cut and enlarges it in the industry cut. **The two are not error: they identify different variation** (Table 16.16). Between size strata, larger firms hold more work in process,

so taking the total makes W more nearly proportional to sales. Between industries, the weight of work in process itself is dispersed as a characteristic of the industry, so taking the total does not improve proportionality.

The estimate of this chapter is identified by variation between industries and years, so **it is the industry cut that applies here**. A direction measured in one cut cannot be carried into the other. **The range a sample covers is not the same as the range a conclusion reaches**.

Remark 16.9 (Without a common sample the sign comes out reversed). Since the observations with $W > 0$ differ by definition, estimating without aligning the sample gives β of 0.894 and 0.787 in the size cut, yielding the opposite conclusion that the denominator effect *increases*. Restricted to the common sample the values are 0.900 and 0.960. **When measuring the effect of a definition, note that changing the definition changes the composition of the sample**.

16.10.5 Verdict

The conclusion of Section 16.9 is maintained, but **it is merely maintained, not strengthened**.

Claim	Verdict
The absolute level cannot be interpreted	Supported. The median CCC moves from 28.0 to 43.9 days
The effect on analyses of differences and correlations is small	Supported. But the correlation of ΔCCC is 0.882, below the 0.976 of the size cut
The denominator effect is not dominant	Supported. But correction raises it slightly from 0.160 to 0.172

Table16.19: Verdicts on three claims concerning the scope of inventories, by the industry cut.

All three are supported, but none strongly. That $1 - \beta$ is 0.172 means 17% of the increase in working capital is not proportional to sales: not dominant, but not negligible either.

16.11 An observation the simplified capacity constraint does not fit

Remark 2.23 introduced the simplification of treating Cap as a constant. This section records a case in which it does not fit reality. It is a record of a limitation, not a revision of the theory.

16.11.1 The observation

The price schedule of one major full-service fitness club chain was checked as of 2026. The company offers a weekday-daytime-only plan for those aged 60 and over alongside its ordinary plan. The two have identical structures — once a week (up to four visits a month), twice a week (up to eight), and unlimited — and differ only in price.

Facility category	Once/wk	Twice/wk	Unlimited
I	4.0%	10.3%	22.1%
II	4.8%	9.9%	21.5%
III	4.1%	9.6%	21.7%
IV	4.6%	10.0%	21.5%

Table16.20: Discount rate of the 60-and-over plan, by facility category and usage tier.

Discount of the 60-and-over plan relative to the ordinary plan

In all four categories the discount increases monotonically with usage. It stays around 4% for once a week but reaches about 22% on the unlimited plan. The price ratio from once a week to unlimited is 1.67–1.74 on the ordinary plan and 1.36–1.43 on the 60-and-over plan.

16.11.2 Inconsistency with the simplification

Under equation (2.17), a large discount on the unlimited plan is hard to explain, since it presses on the capacity constraint. Under the general form (2.16), however, it is explicable: restricting the plan to weekday daytime places $\text{supp } \delta_i$ within the idle part of $\text{Cap}(t)$, so total volume can be increased without worsening the constraint at the peak instant.

That is, **the operator manipulates the distribution over time, not the total.** The simplification of this work cannot express this device.

16.11.3 Why the theory is not revised from this observation

Remark 16.10 (No generalization from a case). The observation of this section rests on one company's price list. Part I is an abstraction intended to cover all Φ , while Part IV is a concrete study confined to particular industries and firms. **The concrete carries no claim of exhaustiveness.**

Rewriting equation (2.17) from an $n = 1$ observation would propagate a single case across all 28 types. This is of the same form as the retrospective narrativization warned against in Section 12.6: revising the structure of a population from one sample.

A revision of the theory is warranted when a deductive error is found, when the same limitation appears in several independent domains, or when it affects the deductions of Part II. None of these applies here. The observation is therefore recorded, and no revision is made.

16.11.4 The simplified capacity constraint

In connection with Section 8.6, the following implication was fixed in advance.

Registered implication: older cohorts choose pay-per-use over flat-rate contracts. This follows from the preference predicted by socioemotional selectivity theory — that under a limited time horizon one chooses the certain. A flat-rate contract is a bundle of options; pay-per-use is certain.

Alternative: older cohorts have more disposable time and use the facility more often, so the flat rate is more advantageous. Under this account the preference for flat rates *rises*.

The signs are opposite, so the form is identifiable.

However, **no data on consumer choice could be obtained, and the implication is untested**. What is needed is the distribution of contract-form choice by age, which operators hold but do not publish.

Only the supply side's price design was obtained. The discount structure above rests on the premise that older cohorts use the facility more, which is consistent with the alternative and points against the registered implication. But **this is the operator's belief, not the consumer's choice**. It requires the premise that the firm reads demand correctly, and so remains indirect evidence.

16.12 Limitations of this chapter

- (1) **The separation of the denominator effect is of low confidence.** The elasticity estimate puts the contribution at about 17%, but $R^2 = 0.076$, and most of the variation in $\Delta \ln W$ is unexplained.
- (2) **The horizontal axis of the four quadrants is not measured.** No cross-industry indicator exists for the ease of capturing surplus, so the classification remains one-dimensional, on the demand side alone.
- (3) **The interpretation of the growth in accounts receivable is not settled.** Section 16.8 excludes industrial composition and explains half by the disappearance of bills, but whether the growth of accounts receivable in the 2010s is a genuine lengthening or a shift between accounts cannot be separated (Remark 16.5).

- (4) **The contract choice of older cohorts is untested.** As in Section 16.11, no consumer-side choice data could be obtained. The supply side's price design is consistent with the alternative but remains indirect evidence.
- (5) **Inventories are “finished goods or merchandise” only.** Work in process and raw materials are excluded, so DIO is understated and the level of CCC comes out shorter than it is. Section 16.10 measures the difference for ten industries and confirms that the effect on analyses of differences and correlations is small, but **correcting the 45-industry estimate would require obtaining that item for all industries.**

16.13 On both implications being rejected

The direction of credit had the opposite sign; the leading property vanished. Both registered implications failed.

Part IV recorded the effect of pre-registration, and it operated more clearly in this chapter. Had the implications not been fixed beforehand, the contemporaneous negative correlation would very likely have been read as “CCC works as a business-cycle indicator” and the credit position of large firms as “large firms support their counterparties”. Both are retrospective narratives.

The failure of the implications is not a rejection of the framework. The descriptions CCC, κ and Φ function effectively; what failed were **conjectures about what these quantities predict**. Effectiveness as a tool of description and effectiveness as a tool of prediction are different things.

Chapter 17

The Price of Externalization: Platform Fees

17.1 The implication tested

The claim of the previous section — that the price of externalization is a transfer of ϕ_{barg} — can be tested directly. Following the discipline of Part III, the implication was fixed before the survey.

Registered implication: a platform’s fee rate increases monotonically in the number of terms the platform absorbs.

The alternative is stated explicitly: fee rates are set by market power rather than by the functions absorbed. That is, the very decomposition of Chapter 6 into ϕ_{prod} and ϕ_{barg} is at issue. Two comparisons were specified in advance for identification.

- (1) Same function, different scale. If the rates are close, the function account holds; if they differ greatly, the market-power account does.
- (2) Differences in rate within one operator. Whether there is a difference the function account cannot explain.

This survey falls under none of the obstacles of Section 13.2: the price lists are public, and neither per-firm data nor confidentiality is involved.

Theoretical quantities in this chapter		
ϕ_{prod}	Eq. (6.1)	the price of function; 20.0 points
ϕ_{barg}	Eq. (6.1)	rent from market power; 15.0 points
$T_{\text{inst}}, T_{\text{dev}}, T_{\text{acq}}$	Eq. (8.11)	measured as prices of externalization
λ_{novel}	Chapter 9	48,185 CVEs a year, 245 a year effective

Table17.1: Theoretical quantities in this chapter.

17.2 Result: a ladder of rates

Effective rates were computed for a \$50 transaction. Where payment processing is charged separately, $2.9\% + \$0.30$ is added.

Platform	Effective rate	Terms absorbed
Stripe (domestic cards)	3.6%	payment
Polar Starter	5.0%	payment + tax (merchant of record)
Paddle / Lemon Squeezy	6.0%	payment + tax (merchant of record)
Gumroad direct	14.5%	payment + tax + delivery and hosting
App Store small business	15.0%	payment + tax + delivery + review + distribution
App Store standard	30.0%	as above
Gumroad Discover	34.5%	the above + customer acquisition

Table17.2: Effective rates on a \$50 transaction, by platform.

The same ordering is shown in Figure 17.1.




Figure17.1: Effective rate on a \$50 transaction. It rises with the number of terms absorbed.

Remark 17.1 (The ladder is a price list for C_{admin}). What appears here is the price of the outsourced part of the C_{admin} of Definition 2.14.

Increment	Corresponding kind of C_{admin}	Price
Stripe	payment processing (operation)	3.6%
→ merchant of record	tax compliance (securing performance, setup)	+1.4–2.4 points
→ sales platform	delivery and rights management (operation)	+8–9 points
→ marketplace	customer acquisition (not C_{admin} but T_{acq})	+20 points

Table17.3: The increments of Table 17.2 and the corresponding kinds of C_{admin} , with prices.

Only the last term differs in character. Customer acquisition is not the operation of a contract but the winning of customers, and belongs not to C_{admin} but to the T_{acq} of (8.11). **The externalization price of C_{admin} is about 11 points; T_{acq} alone is 20.**

The ladder is broadly monotone. Stripe absorbs payment only, at a flat 3.6% for domestic cards, debit, wallets and convenience-store payment. The merchant-of-record providers Paddle and Lemon Squeezy charge 5% + 50¢ and collect and remit sales tax in more than 200 jurisdictions. The increment from Stripe to a merchant of record can be read as the price of T_{inst} . For an individual selling into Europe, the alternative is registering for VAT in each country or paying an accountant monthly, so a price of 1.4–2.4 points is consistent with a return to function.

17.3 A natural experiment within one platform

This is the most important observation. Gumroad’s 30% Discover fee applies only when the customer arrived through Gumroad’s marketplace; purchases through one’s own link or an external route are charged the standard 10%. Discover can also be switched off.

The same platform, the same market power, the same other functions. The only difference is whether T_{acq} was supplied.

$$\text{price of } T_{\text{acq}} = 20.0 \text{ points} \quad (17.1)$$

The market-power account cannot explain this difference: one and the same firm charges the same customer different rates according only to whether it performed the acquisition. Equation (17.1) is the strongest evidence for the implication.

Moreover the standalone price of T_{acq} , 20 points, exceeds the sum of all the other items (about 11 points, from 3.6% to 14.5%). **Customer acquisition is the most expensive item to externalize.** The judgment in family 6-4 that T_{acq} is the dominant component is supported in direction.

17.4 A counterexample: price discrimination by size

At the same time, a case the function account cannot explain was obtained.

The App Store Small Business Program cuts the fee from 30% to 15% for developers earning under one million dollars a year. Paid apps, in-app purchases and subscriptions are all covered.

The functions supplied are exactly the same and only the rate differs, by a factor of two. A 15-point difference cannot be a return to function; it is price discrimination by size.

The circumstances of its introduction are also suggestive. The programme was announced in 2020, when antitrust actions by the US Department of Justice and the European Commission coincided with the dispute with Epic Games. **Pressure from the institutional layer moved the rate.**

17.5 Verdict

Claim	Verdict
The rate rises with the number of terms absorbed	supported (the ladder is monotone)
Differences in rate are a return to function	partly supported ((17.1))
Rates are fully explained by function	rejected (the App Store's 15-point gap)

Table17.4: Verdicts on three claims about the price of externalization.

In the decomposition of (6.1), the level of platform fees can therefore be written

$$\phi_{\text{fee}} = \underbrace{\phi_{\text{prod}}}_{20.0 \text{ pt}} + \underbrace{\phi_{\text{barg}}}_{15.0 \text{ pt}} \quad (17.2)$$

the first from the within-platform comparison on Gumroad and the second from the App Store's two size tiers. **This is the only case in this text where two of the three components were separated numerically for a single transaction.**

Remark 17.2 ((17.2) is incomplete). The two firms above are intermediaries operating at uniform rates from a published price list. Where an intermediary sets rates by individual screening, a third term is added to (17.2).

Merchant fees for payment processing vary with the form of transaction, volume, the goods sold, and the incidence of fraud and chargebacks; for domestic card payments 3–6% is the usual guide. That variation is not rent from market power but the price of the *transfer* and *pooling* of Section 2.6. From the merchant's side the index i is moved and risk handed over; the intermediary bundles many merchants along ω .

Properly, then,

$$\phi_{\text{fee}} = \phi_{\text{prod}} + \underbrace{\phi_{\text{risk}}}_{\text{price of transfer and pooling}} + \phi_{\text{barg}} \quad (17.3)$$

where ϕ_{risk} is not a new component but, as Remark 6.1 states, the undivided sum of ϕ_{prod} and ϕ_{barg} . For the two firms in this section $\phi_{\text{risk}} \approx 0$, which is why it reduced to two terms.

17.5.1 The risk component: untested

For the second term of (17.3), an implication follows from Proposition 2.9.

As Remark 2.10 states, the limit of (2.11) bottoms out at $\rho\sigma^2$. An intermediary can pool away risks that are independent within an industry but cannot remove the part correlated through a common factor.

Proposition 2.9 assumes ρ does not depend on n (Remark 2.11). If an intermediary takes on industries with higher correlation as it adds merchants, ρ rises and the effect appears more strongly than stated here. **Dropping the assumption does not change the sign.**

Registered implication: differences in payment-processing rates across industries correspond not to the mean chargeback rate μ of the industry but to **the within-industry correlation** ρ . An industry with a merely high mean can pass the cost through, whereas an industry structured so that the whole of it deteriorates at once requires an additional premium.

This splits the explanation “the rate is high because the risk is high” in two. Whether the high rates for information products or travel are due to an average incident rate or to the whole industry deteriorating together has different implications for the design of Φ : the former can be passed through in price, the latter cannot.

Testability is constrained as set out in Remark 17.3.

Remark 17.3 (Why this implication cannot be tested). Estimating ρ requires time series of chargeback rates by industry, which are internal data of intermediaries and the card networks and are not published. Even μ is not published.

Furthermore, **industries with high ρ are screened out at merchant onboarding**. The rates observed are a distribution after industries with high ρ have been excluded, so the conditioning on survival of Section 12.2 operates here too. The obstacle is double, and in the taxonomy of Section 13.2 it is a combination of (iii), not identified, and (iv), access constraint.

The proposition that correlation fixes the limit of pooling is itself not new. Solvency regulation for insurers and capital regulation for banks compute capital requirements with explicit correlations between risk categories, implementing the same structure as (2.11). This text’s contribution is only the perspective of applying it to the setting of rates for a Φ .

In the three-way decomposition of Chapter 6, a fee contains both ϕ_{prod} (the functions actually performed) and ϕ_{barg} (transfer), and Gumroad’s design gave the means of separating them.

This is consistent with the durability claim of Section 6.3: **the part deriving from market power is eroded by the institutional layer**. The App Store’s cut from 30% to 15% was a product of regulatory pressure, whereas the part deriving from function (Gumroad’s 20 points for acquisition) has not been eroded.

17.6 Practical implications

Prices now enter the launch strategy of Chapter 8.

Term externalized	Additional cost
Payment only	3.6%
+ tax	+1.4–2.4 points
+ delivery and hosting	+8–9 points
+ customer acquisition	+20 points

Table17.5: Externalization cost for each term of (8.11).

Each term of (8.11) now carries a price. Only acquisition is an order of magnitude more expensive: T_{acq} alone at 20 points exceeds the sum of all other items (about 11 points). The guidance follows that T_{acq} **is often better done in-house**. Placing the “publish and build credit” stage first in the transition across families in the previous chapter was an investment of time in doing T_{credit} and T_{acq} oneself. Against a price of 20 points, that investment of time is readily justified.

17.7 Limits of the measurement

The rates cover only major platforms in Japan and the United States, and no population was declared. By the discipline of Section 12.9 this is a comparison of cases, not an estimate. The within-platform comparison on Gumroad is an identification that does not depend on the population, however, and that part alone is robust.

Effective rates also depend on transaction size. Because of the fixed components (50¢, 30¢), the effective rate of merchant-of-record models rises at low transaction values. The table is computed at \$50.

17.8 How the institutional layer acts

Section 6.3 stated that the ϕ_{barg} deriving from market power is eroded by the institutional layer. This section checks how it acts, in four cases.

Domain	Action of the institutional layer	Current state
Property brokerage	statutory cap on commission (1970 notification, 3% + 60,000 yen)	in force; the cap was raised in 2024
Fee-charging placement	capped fee (about 11% of wages)	in force but used in 0.18% of cases
Securities commissions	fixed commissions under exchange rules	abolished entirely in 1999
App distribution	pressure through antitrust litigation	the rate cut from 30% to 15%

Table 17.6: Four forms in which the institutional layer acts on fee rates, and their current state.

Capping a rate is rare, and the trend is towards abolition. Securities commissions were liberalized above 1bn yen in 1994, above 50m in 1998, and entirely in October 1999. The United States did the same in 1975 and the United Kingdom in 1986.

In property brokerage the cap can be seen to bind. The July 2024 amendment raised the cap for properties up to 8m yen from 180,000 to 300,000 yen, and the stated background was that 247 municipalities — 14% of the total — had no licensed brokerage at all, and that the business was not viable at the existing commission. **Supply disappeared because the cap was too low**, which is direct evidence that the cap binds.

Remark 17.4 (From imposing caps to manipulating competitive conditions). Of the four cases, only the first two regulate the rate directly. In securities commissions, the *abolition* of fixed commissions cut rates to between one seventh and one tenth of their pre-liberalization level. In app distribution, litigation pressure prompted a voluntary reduction.

All four cases here concern rates, so of the three forms of Remark 6.9 they fall under price regulation and manipulation of competitive conditions; entry requirements do not appear. **The latter dominates in recent years.** Forms that regulate rates directly are in decline and are being replaced by manipulation of competitive conditions.

17.9 Assessing feasibility of further investigation

The result of applying the taxonomy of Section 13.2 to the other questions of this chapter is recorded here.

Candidate	Principal obstacle	Verdict
The price of externalization	none	done (Chapter 17)
Decomposing λ_{novel}	volume of linking work	runner-up
Transitions between families	no frame; conditioning on success	weak
Direct measurement of unmanned duration	no population frame	abandoned
The rate of unmannedness itself	(i) + (ii)	abandoned

Table17.7: Applying the taxonomy of verification obstacles (Section 13.2) to the other questions of this chapter.

17.9.1 Runner-up: decomposing λ_{novel}

λ_{novel} can be decomposed, and its components map onto the layers of Section 6.3.

Component	Origin	Data source	Layer
λ_{sec}	discovery of vulnerabilities	vulnerability databases	physical
λ_{dep}	breaking changes in dependencies	version histories in package registries	physical
λ_{inst}	legal change, changes of terms	the official gazette, providers' change notices	institutional
λ_{dem}	change in demand	not measured	cognitive

Table17.8: Four components of λ_{novel} , their origins, data sources and layers.

Public data exist for three of the four, and an implication follows.

Registered implication: the fewer the dependencies in a configuration, the smaller λ_{dep} and the longer unmanned operation can run.

This bears directly on design decisions: if unmanned operation is the aim, few dependencies should be preferred to rich functionality. Connecting a measurement of λ to cessation events must be built oneself, however, and that is where the work concentrates.

17.9.2 Questions abandoned

Direct measurement of the unmanned operating period lacks a population frame for this chapter's objects. No list of solo-developed services exists, and collecting those that stopped after the fact reproduces the conditioning on survival of Section 12.2 exactly.

The abandonment is specific to the objects, however. In the “inside” partition of Chapter 10, the protocol record is itself the population, and duration and cessation rates can

be measured on the full population (Section 18.3). What cannot be measured is solo-developed services, not the unmanned operating period as such.

Transitions between families are similar. Cases that transitioned can be collected, but cases that did not and cases that failed cannot be observed, so nothing can be said about the effectiveness of the route.

The rate of unmannedness itself is an ill-posed proposition. Defining what fraction of δ is automated requires a unit for δ , and none exists. The defect has the same form as that of the measurability hypothesis in Section 15.3.5.

17.10 Limits of this chapter

- (1) λ_{novel} was not measured. Estimating the unmanned operating period depends on that value, and there is no basis for it.
- (2) Chapter 17 is a comparison of cases with no declared population. Apart from the within-platform identification, it must not be treated as an estimate.
- (3) The implication registered in this chapter states the same relation as one registered in Chapter 10. This was not noticed at the time of registration. The test of $\Phi 1$ **rejected** it (Section 18.2). It was not tested on this chapter's objects.
- (4) λ_{novel} was not decomposed into its four components. What Section 18.3 measured is the rate of response, not a breakdown by component.

Chapter 18

Evidence From Protocol Records: $M(t)$ and the Rate of Modification

In the “inside” partition of Chapter 10, $M(t)$, κ_j , $\bar{\delta} - \delta$ and B are published per identifier (Section 10.6). These are quantities that other chapters cannot measure for want of per-firm data. This chapter observes them.

The population is every Φ operating at the reference date, and the sample was not chosen after the fact.

18.1 The implications tested

Following the procedure of Chapter 11.5, four implications were fixed before acquisition. They are numbered with Φ to distinguish them.

Implication $\Phi 1$: the fewer other procedures a procedure calls, the longer until it is modified. This is the λ_{dep} side of Proposition 9.2.

Alternative: the frequency of modification is fixed by usage, not by the number of dependencies. Few dependencies but heavy use still leads to modification.

Implication $\Phi 2$: for a Φ whose transactions cease, $M(t)$ declines monotonically before the cessation. This follows from $\tau = \inf\{t : M(t) < 0\}$ of Section 4.2 being absorbing.

Alternative: cessation occurs independently of $M(t)$, which does not change until immediately before.

Implication $\Phi 3$: the distribution of CCC among Φ operating in the “inside” partition is skewed non-negative, because by Corollary 10.12 realizing $\text{CCC} < 0$ requires $E > 0$.

Alternative: the distribution of CCC resembles that of the corporate sector in Chapter 16, with no skew specific to the “inside” partition.

Implication $\Phi 4$: assigning operating Φ by the rule of Section 7.9, family 7 barely appears except for 7-4. This follows from Proposition 10.6.

Alternative: family 7 types appear with about the same weight as other families.

18.1.1 What was fixed in advance

Item	Content
Population	every Φ operating in the protocol record at the reference date; not chosen after the fact
Floor	those with fewer than a set number of transactions in the reference period are excluded; the threshold is fixed before acquisition
Definition of cessation	no transaction over a set period. Because the protocol itself does not stop in the “inside” partition, cessation rather than stoppage is used
Steady state	CCC presumes a steady state, so only periods with stable r are used (Section 12.4)
Unit of observation	the identifier. No aggregation to parties (Remark 10.5)
Assignment of types	the order of Section 7.9, unchanged

Table18.1: What was fixed in advance for the four implications.

Remark 18.1 (λ_{dep} and usage can be correlated). The alternative to implication $\Phi 1$ is hard to reject, because procedures with many dependencies tend to be feature-rich and also heavily used. The test must control for usage, and whether an effect of dependencies survives in the residual cannot be known in advance. **The control procedure is fixed before acquisition.**

18.2 Results

One implication was supported.

Implication	Sample	Result	Verdict
$\Phi 1$	1.93m modifiable procedures	the number of dependencies does not explain time to modification	rejected
$\Phi 2$	427 with both fees and balances	declines over the 30 days before cessation ($p = 3.2 \times 10^{-9}$)	supported
$\Phi 3$	52 operating Φ	the alternative is not identified	untestable
$\Phi 4$	as above	family 7 appears once, but as 7-2; 7-4 does not appear	rejected

Table18.2: Results for the four implications.

Remark 18.2 (Concentration of clones dominates the estimate). The 1.93 million identifiers point to only 76,609 distinct implementations, and 1.55 million of them — 80% of the population — point to a single one. These are not independent Φ but copies of one Φ .

Over the whole sample the coefficient on the number of dependencies is positive and significant ($p = 0.004$), but removing that one cluster reverses the sign and destroys significance ($p = 0.45$). **Cluster-robust standard errors do not protect against this.** They correct the correlation within a cluster but not the bias in a point estimate when one cluster accounts for 95% of the exposure.

The loss of additivity in Remark 10.5 concerned the divergence between parties and identifiers; the concentration of identifiers pointing to one implementation is a different form of divergence. **When assembling a sample from a ledger, the concentration of clones must be counted first.**

Remark 18.3 (The alternative is not identified). $CCC < 0$ means payment in advance, which by Corollary 10.12 requires $E > 0$. But the CCC of the comparison group, the corporate sector, is non-negative in all 58 industries of Chapter 16 (minimum 4.9 days). A negative CCC appears in individual firms and cancels out on aggregation by industry.

The implication and the alternative therefore say the same thing, and no observation distinguishes them. The ill-posedness has the same form as that of the measurability hypothesis in Section 15.3.5. In addition, τ in the “inside” partition is measured in seconds against days in the corporate sector — five orders of magnitude apart.

Remark 18.4 (Being implementable and appearing are different). Family 7 is 1 of 52, so the quantitative part — that it is rare — was right. But what appeared was 7-2 (cross-subsidy), and 7-4 appeared zero times.

The implication derived “only 7-4 remains” from Proposition 10.6, but the proposition speaks of whether implementation is possible, not of how often something appears. **Deriving appearance from implementability was the error.**

The single 7-2 selects its beneficiaries by signature, outside the protocol. That is consistent with Proposition 10.6, but since the wording of the implication missed, it is not treated as supporting evidence.

Remark 18.5 (A step is needed before the assignment order). Table 7.8 takes Φ as given. Applied to ledger identifiers, 131 of a sample of 200 were not Φ at all: 74 were token ledgers (a medium of settlement with no pair of δ and π), 27 were infrastructure with no consideration, and **30 were identifiers imitating the name of a settlement medium.**

The last satisfy Definition 10.1 but have no delivery. The unobservability of ι in Section 10.3 appears here as misrepresentation of a name.

18.3 Quantities measurable from the same records

Three quantities that other chapters could not measure can be measured from the records used for these tests.

18.3.1 The rate of response to unforeseen states

Proposition 9.2 put the unmanned operating period at $1/\lambda_{\text{novel}}$. Two rates can be measured in the “inside” partition.

Quantity	Content	λ (per year)	$1/\lambda$
Modification rate λ_{fix}	rate at which the implementation was replaced; 1.93m records, excluding the clone concentration	0.048	20.6 years
Cessation rate λ_{stop}	rate at which operation ceased; all 1,718	0.030	33.8 years

Table18.3: Response rates measured in the “inside” partition.

Remark 18.6 (What is measured is not λ_{novel}). Both entries in Table 18.3 count **responses to arrivals**, not **arrivals of unforeseen states**. A modification is a case where the operator responded; a cessation is a case where they did not, or could not.

Arrivals that required no response appear in neither. Hence

$$\lambda_{\text{novel}} \geq \lambda_{\text{fix}} + \lambda_{\text{stop}}$$

and what was measured is a lower bound. The two are not additive because their populations differ, but an upper bound follows: **the $1/\lambda_{\text{novel}}$ of Proposition 9.2 is shorter than 20 years.**

The distribution of time to modification is extremely skewed. The median of the 20,732 that were modified is 42 days; the median of the 1.9 million that ended the observation without modification is 4.6 years. **What gets touched is touched at once; what does not goes untouched for years.** A mean $1/\lambda$ conceals this bimodality.

18.3.2 Duration of the procedures

Chapter 17 abandoned direct measurement of the unmanned operating period for want of a population frame. As Section 10.6.1 notes, that reason disappears in the “inside” partition.

Of all 1,718 that operated, 143 (8.3%) ceased and 1,575 continue. Among those that ceased, the median duration was 1.62 years, the quartiles 0.64–2.82 years, and the maximum 6.0 years. Those continuing have a median of 2.75 years. Because the ceased were counted from the full population rather than collected after the fact, no conditioning on survival (Section 12.2) arises.

Cessation rates are ordered by kind: services $1/\lambda = 11.1$ years, high-leverage operation 11.5, issuance support 15.3, indices 18.1, bridging 20.1, insurance 22.3, lending 22.6, and yield 22.7. **The direction — kinds with lighter dependence on other procedures last longer — can be read off, but since implication $\Phi 1$ was rejected, this ordering is not explained by the number of dependencies.**

18.3.3 Locked state B

Table 10.1 recorded that this text has no quantity corresponding to B . In the “inside” partition it is directly observable as assets placed on the protocol, and daily series for 1,718 records are available. This is what served as the proxy for $M(t)$ in implication $\Phi 2$.

18.3.4 Rates

Once B is measurable, a rate can be defined as annual fees divided by B . But **the denominator differs by family**. For family 3 (renting assets) B is the right denominator, whereas for family 6 (intermediation) the denominator is transaction value and B is merely inventory. Mixing them makes the ratio diverge.

Restricting to family 3 types with B of at least one million dollars, 144 records give the following.

Kind	n	Median	Quartiles
Arbitrage	7	4.97%	2.00–9.31%
Real assets	20	3.85%	2.35–5.84%
Collateralized debt	15	3.53%	0.99–11.86%
Liquid staking	19	3.12%	2.58–4.21%
Lending	30	2.97%	1.77–4.68%
Risk management	21	2.83%	2.02–4.47%
Yield	15	2.72%	0.76–8.01%
Yield aggregation	14	1.69%	0.90–4.19%

Table18.4: Rates for family 3 types (annual fees $\div$ locked state).

The overall median is 3.07% with quartiles of 1.77–5.84%. The ordering has the same form as the ladder of rates in Chapter 17, but **the two cannot be compared**: the rates there are on transaction value, the rates here on balances. The dimensions differ.

Part V

Conclusion

Chapter 19

Where This Stands: What Was Established and What Remains Open

19.1 What was established and what was not

19.1.1 Derived as theory

Table19.1: Claims derived as theory, and their status.

Claim	Content	Status
Description by Φ	can be written as a map from delivery to settlement	worked
Unification through κ	receivables, advances and payables become one quantity	works
Integration through κ_i	four literatures enter the same sum	the central contribution (Remark 3.3)
$CCC \cdot r \leq E - A$	the cash constraint is a constraint on working capital	the fundamental inequality (Proposition 5.2)
$r^{\max} = (E - A)/CCC$	a bound on scale, separate from g^*	derived (Corollary 5.3)
$g^* = (m + d - I/r)/CCC$	self-financeable growth rate; the boundary of $\dot{M} \geq 0$	derived (Corollary 5.7)
Time to cash depletion	how long before $g > g^*$ bites	closed form (Proposition 5.9, Example 5.10)
$CCC = \sum \tau_i + \tau_{\text{inv}}$	CCC is a sum of lags	derived (Proposition 5.16)

(continued)

Claim	Content	Status
Perfect substitution of τ_i	collecting sooner and paying later are equivalent	derived (Proposition 5.22)
Antisymmetry of κ	inter-firm credit vanishes on aggregation	derived (Proposition 3.4)
Peak versus average	the capital required is fixed by $CCC + T_b/2$	derived (Proposition 5.21)
Cash neutrality of cognitive surplus	no cash moves at recognition	derived (Proposition 6.6)
Bound on accumulated cognitive surplus	transaction frequency fixes the bound on scale	derived (Corollary 6.8)
Switch of the binding constraint	from cash to capacity at $CCC^\dagger$	derived (Proposition 8.1)
Availability and granularity	deliveries that cannot be performed even with spare capacity	derived (Proposition 9.5)
Membership under capacity	low utilization is the condition of feasibility	derived (Proposition 2.24)
Allocation in a solo business	contract work stays at the minimum the constraint requires	derived (Proposition 8.10)
Linking E to ϕ	what retained earnings accumulate	the accumulation of ϕ , by Definition 3.20
Party dependence of κ_K	is equity the same as other κ	it depends on the measure and does not enter the identity (Proposition 3.22)

19.1.2 Judged empirically

Table 19.2: Claims judged empirically, including those rejected and untested.

Claim	Content	Status
Sign of CCC	two orders of magnitude across industries and media	confirmed (Part IV)
Path of g^*	change over 25 years	has not declined (Proposition 14.2)

(continued)

Claim	Content	Status
Numerator of g^*	m alone or investment too	including I moves the level substantially
m of small firms	the margin gap by size	explained by directors' remuneration (Proposition 14.8)
Comparison with unincorporated	accounting classification or economic substance	a separate statistic supports Proposition 14.8 (Section 14.5.5)
Causes of rising CCC	composition or within-industry change	within-industry 83%; the disappearance of trade bills is the main cause
Contribution of the denominator effect	mechanical variation of $CCC = W/r$	about 17%; the sales elasticity of CCC is -0.17
Price of externalization	fees correspond to the functions absorbed	partly supported (Chapter 17)
How the institutional layer acts	caps or manipulation of competitive conditions	the latter dominates recently (Remark 17.4)
Predictive power of CCC	as a leading indicator of the cycle	rejected (Chapter 16)
Size and the direction of credit	smaller firms as providers of credit	rejected, sign reversed (Proposition 16.2)
Measuring ϕ_{cog}	cognitive surplus measured	not reached here (Section 14.6.1)
Measurability and contract form	observability fixes Φ	untested (Section 15.3.5)
Fee rate and separation rate	is outcome risk reflected in price	rejected ; uncorrelated across 307 firms (Section 15.3.5)
Absorbency of the path constraint	does $M(t)$ fall before cessation	supported (Section 18.2)

(continued)

Claim	Content	Status
Dependencies and unmanned duration	do fewer dependencies mean later modification	rejected ; the concentration of clones drives the result
Implementability of family 7	is 7-4 what remains	rejected ; only the rarity was right
Unmanned duration	the level of $1/\lambda_{\text{novel}}$	shorter than 20 years (Remark 9.3)
Duration of procedures	measured on the full population	8.3% ceased, median 1.62 years (Section 18.3)

As a descriptive instrument the framework worked. As a predictive one it did not. These are different properties, and only the former is claimed here.

19.2 Fixing the implications in advance

The implications derived from the theory were fixed in writing before the observations, so that support and rejection could not be chosen afterwards. Twenty-two were registered: six supported, eleven rejected, five untestable. The wording and the verdicts are listed in Appendix D.

The number of rejections is not a failure of the method. Without the fixing, an after-the-fact account could have been given for any fact observed. Appendix D records four places where a post-hoc interpretation was nearly adopted.

19.3 Open problems

The open problems are organized by the six categories of Section 13.2 (Section 13.1). Because the remedy differs by category, the allocation of effort depends on the distinction.

19.3.1 (i) No measurement exists

The quantity is not measured anywhere. Secondary data will not fill the gap; an original survey is required.

- (1) **Is directors' remuneration a cost or an appropriation of profit?** The cause of the low m in the small classes is directors' remuneration (Proposition 14.8). As a fraction of sales it is fifty times that of large firms, and adding it back to operating profit removes the ordering by size.

The remaining question is whether that remuneration is a return to labour or a discretionary appropriation. That the remuneration ratio does not move when profit does, and that the operating margin of unincorporated businesses rises in a downturn (Remark 14.10), suggest

discretion; but the labour-return account cannot be rejected. Deciding requires directors' hours worked, and **directors are not employees under the Labour Standards Act, so no official statistics on their hours exist** (Remark 14.12).

An alternative route would identify from changes in the corporate tax rate or the thresholds for reduced rates as exogenous shocks, but cutting a sample at a threshold requires microdata.

19.3.2 (ii) Ill-posed hypothesis

The theory does not fix the sign. More data does not help; the theory must be refined.

- (2) **Testing the measurability hypothesis.** The formula for b^* itself cannot be tested because the required quantities are not available (Section 15.3.5).

A test on the price side was carried out. Across 307 firms, fee rates achieved and separation rates were uncorrelated (Section 15.3.5). Separation rates differ 5.5-fold across occupational averages while fee rates differ only 1.16-fold, and prices cluster on the customary 20–30%. The obstacle is therefore not only (ii). It belongs to (vi) of Section 13.1 — **the explanatory variable does not vary over the range the theory assumes** — and is resolved neither by refining the theory nor by adding data.

- (3) **The objective under R2: are H and β enough?** A route was shown that explains the change through H (time horizon) and β (discount factor) without changing the form of the objective (Section 8.6.4). Whether that suffices is untested.

Testing requires observing contract choices among older people, and no distribution of contract-form choice by age is available. The supply side's price design is consistent with the alternative, but it is not consumers' choice.

19.3.3 (iii) Not identified

The data exist but confounding cannot be separated. Better design or exogenous variation is required.

- (4) **Interpreting the lengthening of receivables.** The long-run rise in CCC was decomposed in Section 16.8. Changes in industrial composition are excluded (the composition effect is 17%, and 2% since 2010), and the rise in the 2000s is explained by shorter DPO as trade bills disappeared.

What remains is the lengthening of receivables in the 2010s (+6.62 days). Whether this is a genuine lengthening of collection or a reclassification accompanying the shift to electronically recorded monetary claims cannot be separated. Separating them requires balances of such claims, which the corporate statistics do not survey (Remark 16.5).

- (5) **What drives the variation in $\Delta \ln W$.** The contemporaneous negative correlation was separated in Section 16.9, giving a denominator-effect contribution of about 17%. That the denominator effect is not dominant could be shown.

An estimate allowing for lag structure was attempted but nothing was detectable in annual data (Remark 16.6). The R^2 of the elasticity estimate is only 0.081, so most of the annual variation in working capital is driven by something other than sales. Year-end inventory adjustment, timing shifts on large transactions, and changes in accounting policy plausibly contaminate it, but none is separated.

19.3.4 (iv) Access constraint

It exists but cannot be obtained. A route must be opened, or the attempt abandoned.

- (6) **Measuring ϕ_{cog} .** Of the three conditions of Proposition 14.13, only (1), the contractual right, was reached here. (2), individual realized use, and (3), an exogenous event forcing attention, are missing.

A detour was constructed through $\delta^* = F/p$, but it applies only to binary menus and the proxy is not uniquely determined for a stepped tariff (Remark 14.14).

For family 2-4 (prepaid), [13] shows that about 40% of prepaid value goes unused, so **the quantity this text could not measure for family 2 has already been measured**. The remaining gaps are an estimate for family 2-5 (options) and a cross-industry estimate of the share of ϕ_{cog} within total ϕ .

- (7) **An indicator of how readily surplus is taken.** No indicator corresponding to the horizontal axis of the two in Chapter 16 could be constructed. The reason is structural.

An intermediary is an entity that spans industries, not an attribute of one (Remark 16.1). It cannot in principle be captured by industry aggregates. Attempting a cross-section on the intermediary side fails because rates decided by individual screening are not disclosed. **That a rate is negotiated presupposes that nobody outside the transaction knows it.** Collecting only published rates biases the sample towards thin intermediaries operating from a price list.

An implication was registered for the risk component (Section 17.5.1), but the data needed to estimate ρ lie inside intermediaries, and industries with high ρ are screened out, so it cannot be tested.

- (8) **Curvature of the age effect.** By Proposition A.7 the levels of the three coefficients are not identified, but the second difference — the curvature — is (Remark A.8).

Pursuing this requires a sample carrying firm age. The corporate statistics used here have three axes — industry, size and year — and do not include the year of founding. As Remark A.10 states, an age effect is latent in the CCC time series, and there is no means of separating it.

19.3.5 Matters of the framework's descriptive range

Items that do not fit the taxonomy of obstacles and concern the setting of the object itself.

- (9) **The simplification of the capacity constraint.** As Remark 2.23 states, Cap was treated as a constant. Section 16.11 records an observation with which this does not agree, but with $n = 1$ the theory was not amended (Remark 16.10). Whether the same limit appears in other families is unconfirmed.
- (10) **Social optimality of pooling.** This text is written from the standpoint of an individual operator and does not treat the effect of the choice of Φ on the system (Remark 2.12). Families 5-4 (underwriting) and 6-2 (escrow) make pooling their business, and pooling that is individually optimal may be excessive for the system.
- (11) **The credit position of sole proprietors.** No published statistic measures the κ of sole proprietors. The Survey of Unincorporated Enterprises dropped its assets and liabilities items in FY2019, and the Basic Survey on Small and Medium Enterprises has a questionnaire for sole proprietors but does not tabulate assets and liabilities (Section 15.7).
As a result, the effect of the Freelance Act (in force November 2024) setting an upper bound on τ cannot be tested. Against the three forms of the institutional layer organized in Section 17.8, this is **a fourth form that regulates time directly**, and it is of high theoretical interest as a case in which law imposes a ceiling on the τ_i of Section 3.1.4.
- (12) **Forms in which claims circulate.** The $\Phi : \Delta \rightarrow \Pi$ of Chapter 2 presumes that delivery and settlement correspond between the same counterparty. When settlement is resold to a third party, as with token issuance, the party receiving delivery changes after the fact and cannot be expressed in a setting with N fixed (Section 7.10.3). An extension of the domain is required.
- (13) **Observation through protocol records. The tests were carried out.** They were carried out as reported in Chapter 18. This was the first test of the deductions of Part II, and only $\Phi 2$ was supported.
What remains is per-party quantities. What is observable is per identifier, and because one party may hold several identifiers the per-party κ_i cannot be recovered (Remark 10.5). The obstacle is not the absence of aggregation but **its impossibility**, and it differs from both forms added in Section 13.1. The tests revealed a second form of this divergence: the concentration of identifiers pointing to one implementation (Remark 18.2).
What could be verified, moreover, is the correctness of a deduction, not whether that deduction applies to real solo businesses (Remark 10.14).
This is the same place as item (12): a constraint originating in the setting of the index.
- (14) **A quantity expressing a degree of unmannedness.** “How unmanned is it?” cannot be expressed as a ratio. Measured on the delivery side, the valuation map v measures only value; measured on the response side, no measure is defined on $\Omega \setminus \hat{\Omega}$ (Remark 9.9). The obstacle is the absence of a measure, not of a unit, and sharpening the definition does not remove it.
- (15) **Separating ϕ_{prod} from ϕ_{barg} .** Definition 6.2 requires the price p^* that would obtain without market power, which is a counterfactual and is not observed (Remark 6.3). The numerical

separation achieved in Chapter 17 rests on the exceptional availability of a within-platform comparison differing only in function. For a general object, an additional assumption identifying p^* is required.

- (16) **The degree to which the third party is removed.** Chapter 10 treats only the cases where the protocol bears recording and enforcement of settlement entirely and not at all, not the middle (Remark 10.4). The composition of the party writing external facts into the state varies, and most existing arrangements are in the middle.

19.3.6 Items with no obstacle

The following fall under none of the obstacles of Section 13.2. Either the route of acquisition is established or the work can be done with data already obtained. They have simply not been started.

- (17) **Decomposing λ_{novel} into components.** Section 18.3 measured the response rates: 0.048 a year for modification and 0.030 for cessation, so the $1/\lambda_{\text{novel}}$ of Proposition 9.2 is shorter than 20 years.

What remains is the decomposition into components. What was measured is the total of responses, without separating whether they originate in vulnerabilities, dependencies, institutions or demand. Separating them requires reading the content of each modification, which is possible only where verified sources exist, so the sample is biased.

- (18) **Correcting the scope of inventory.** Chapter 16 uses only “finished goods and merchandise” and omits work in progress and raw materials, so DIO is understated.

This was carried out for 62 industries (Section 16.10). The median CCC moves from 28.0 to 43.9 days, but the correlation of ΔCCC is 0.882 and the denominator effect moves only from 0.160 to 0.172. The chapter’s conclusions stand.

What remains is the cross-section with size classes. “Inventory (period end)” is available for 62 industries only for all sizes together, and by size class for only 10 industries. **A cross-section split by both industry and size exceeds e-Stat’s 100,000-cell limit.** It is possible by splitting the retrieval, but what would be obtained is information linking the two cross-sections of Remark 16.8, which would not move the chapter’s conclusions.

19.3.7 Priorities

Assessed on feasibility, breadth of consequence and cost, separating what can be started from what effort will not advance.

	Item	Category	Feasibility	Consequence
1	(17) components of λ_{novel} (only the breakdown remains)	disclosure	low	medium
2	(3) H and β under R2	(ii)	medium	medium
3	(5) what drives $\Delta \ln W$	(iii)	medium	medium
4	(11) κ of sole proprietors	choice in aggregation	medium	large
5	(13) per-party κ_i	impossibility of aggregation	low	large
6	(18) scope of inventory (crossed with size)	choice in aggregation	low	small

Table19.3: Priorities among the items that can be started.

	Item	Category	Why effort will not advance it
7	(6) ϕ_{cog}	(iv)	individual records and an exogenous shock are out of reach
8	(1) directors' remuneration	(i)	no statistics on hours worked exist
9	(2) measurability hypothesis	(ii)+(vi)	the explanatory variable does not vary
10	(7) how readily surplus is taken	(iv)	an intermediary is not an attribute of an industry
11	(4) interpreting receivables growth	(iii)	balances of electronically recorded claims are not surveyed
12	(8) curvature of the age effect	(iv)	no sample carries firm age
13	(14) degree of unmannedness	—	no measure is defined
14	(15) separating the components of ϕ	—	p^* is a counterfactual
15	(10) social optimality of pooling	—	outside the object as set
16	(9) capacity constraint	—	as above
17	(12) circulation of claims	—	requires extending the domain
18	(16) degree of removal of the third party	—	as above

Table19.4: Items that effort will not advance.

Items in categories (i) and (ii) do not advance with effort.

Among them, (11), the κ of sole proprietors, differs in character. The obstacle is neither absence of measurement nor non-identification but **a choice in aggregation**: responses that may already have been collected are not published. As Section 15.7 notes, the Basic Survey on Small and Medium Enterprises has a questionnaire for sole proprietors. A request to the statistical authority for the tabulation, or an application to use anonymized microdata, is a possible route.

Item (13), protocol records, is one where quantities this text abandoned as requiring per-firm data are in fact published in the “inside” partition of Chapter 10. But because Remark 10.5 blocks aggregation to parties, what can be measured is confined to the distribution per identifier.

Two items have no obstacle, and both can be started with published data or an established acquisition procedure. Item (18) has been done for 62 industries and what remains is the cross-section with size classes (Section 16.10). The kind of work required differs from that of the items in Table 19.4.

19.4 Known weaknesses of the framework

The weaknesses revealed through verification fall into three types. None is an error in an individual proposition; all **originate in the construction of the framework**.

19.4.1 (A) Excessive assumptions of independence

Several propositions treat as independent quantities that are in fact determined jointly.

Proposition	Pair treated as independent	The actual relation
2.9 pooling	n and ρ	scale exhausts weakly correlated objects (Remark 2.11)
2.22 divergence	$\bar{\delta}$ and δ	jointly determined through menu choice
2.24 membership	n and $\mathbb{E}[\delta]$	the fee moves both
4.1 implementability	$\mathcal{G}$ and π	$\mathcal{G}$ is an object of investment (Remark 4.3)
5.5 FCF	m and CCC	both rose together in the measurements (Remark 5.24)
A.15 substitution	E and κ	with capital one can extend credit (a corollary of Proposition 14.8)

Table 19.5: Propositions assuming independence, and the quantities that are in fact jointly determined.

Each proposition states its independence assumption explicitly, but **no consequence is derived for the case where the assumption is dropped**. The proofs are formally correct; the weakness is that their range of application is narrow.

Four of the implications fixed in advance failed because of this weakness.

Implication	Pair treated as independent	Form of the error
A	κ and ϕ_{cog}	identified dwell time with non-exercise
B	b and σ^2	looked only at the numerator of the relation
F	DSO and DPO	overlooked that bargaining power acts on both sides
K	$\bar{\delta}$ and δ	did not anticipate the stepped structure of the menu

Table 19.6: Four implications that failed because of excessive assumptions of independence.

B and F have the same form: **a variable acting on both sides was handled on one side only.**

19.4.2 (B) Excessive assumptions of exogeneity

Institutions, Φ , v , Cap and $\mathcal{G}$ are all treated as exogenous.

- (1) **Subjectivity of the valuation map v .** The v that measures delivery in money was taken as given, but it is itself subjective, and the definition of ϕ_{cog} depends on it.
- (2) **Endogeneity of Φ , and switching cost.** Φ was treated as a firm's choice variable, but it is in fact an equilibrium of bargaining between parties and cannot be chosen unilaterally. In addition the switching cost is set to zero (Remark 2.18). Where that cost is of the same order as the gain in surplus, waiting has value and a question of timing arises — when to fix Φ . This text does not treat it.
- (3) **Exogeneity of the institutional layer.** Institutions were treated only as an exogenous constraint. In fact they change endogenously in response to firm behaviour. The correspondence between antitrust litigation and rate reductions in Chapter 17 is an example.
- (4) **Simplification of Cap.** It was treated as a constant, but it is properly a function of time and is also a design variable through restrictions on hours of use (Remark 2.23).
- (5) **Endogeneity of $\mathcal{G}$.** The information structure was taken as given, but an operator can extend it through follow-up surveys (Remark 4.3).

19.4.3 (C) Level and change are not separated

This is the weakness most easily missed. (A) and (B) can be recognized as “the assumption is too strong”, whereas (C) has no corresponding concept in the theory and is hard to notice as an absence.

As Remark 2.5 states, every quantity here is defined as a level. Yet in the empirical work, four times over, cross-sectional comparison of levels could not be interpreted and only change could be discussed.

Situation	Response
Family 2, panel by medium	levels differ by two orders of magnitude; time effects separated using the coexistence of growing segments
Chapter 16, CCC by industry	5–210 days; stated explicitly that cross-sectional comparison is uninterpretable
Shift-share decomposition	change in levels split into within-industry and composition effects
Elasticity estimation	β estimated as a relation between changes

Table19.7: Situations where level and change were not separated, and the response in each.

The theory has only levels while the empirical work can speak only of change.

Part of this mismatch was resolved by introducing the lag τ_i in Section 3.1.4. Proposition 5.16 gives $CCC = \sum_i \tau_i + \tau_{inv}$, from which it is deduced that the level of CCC is fixed by Φ (Corollary 5.17) and that cross-sectional comparison is meaningless (Corollary 5.18). A tool for splitting ΔCCC into change in contracts and change in operations was also obtained (Corollary 5.20).

That level and change are separate constraints was likewise settled at the boundary of the fundamental inequality of Proposition 5.2: $M \geq 0$ bounds the level and $\dot{M} \geq 0$ bounds change (Remark 5.8). The two are independent, yet their sensitivity to τ_i carries the same coefficient $-1/CCC$ (Proposition 5.22, Table 5.3).

What remains is constructing a comparable quantity. The candidate is the deviation from the contractual τ_i^{contract} , but contractual terms do not exist in published statistics (Remark 5.23). The theoretical problem is solved and the measurement problem has taken its place.

19.4.4 The three types, and mid-level theory

For (A) and (C), the existing theory that ought to be connected can be identified.

Weakness	Where it appears	Theory that could be connected
(A) $\bar{\delta}$ and δ	Remark 2.22, family 2, gyms	mechanism design, self-selection and quantity discounts
(A) b and σ^2	Proposition 4.1, family 5	the screening effect of contracts
(A) DSO and DPO	Proposition A.15, Chapter 16	the theory of trade credit
(C) level and change	all empirical chapters	—

Table19.8: Weaknesses of the framework and the existing theory that could be connected.

These were referred to but not connected systematically. Borrowing the informativeness principle in Chapter 4 while failing to check its contested status when framing an implication in Section 15.3.5 happened because the same theory was handled at different depths when citing it and when framing an implication.

For (C), no theory to connect to has been identified.

Part VI

Appendices

AppendixA

Supplement: proofs concerning method

Those claims about method made in the body that call for proof are gathered here. Each section also states why the formalization takes the form it does, which implications were not treated in the body, and the limits within which the proof applies.

The implications developed here are: the relation between the strength of conditioning and the distortion (Remark A.3); that collecting only failures does not remove the distortion (Remark A.4); that curvature is identified in the APC problem (Remark A.8); that the APC problem is latent in the CCC series of this work (Remark A.10); that the objective becomes discontinuous under a deadline constraint (Remark A.13); and that substitutability is confined to a static constraint (Remarks A.17 and A.18).

Limits of application are recorded in Remarks A.5 and A.14.

A.1 Correlation induced by conditioning on a collider

A.1.1 The choice of formalization

Section 12.2 treats the problem that observed firms are conditioned on survival. To formalize this one must decide what survival is a function of.

This work adopts a **linear additive threshold** $S = \mathbf{1}\{a\Phi + b\varepsilon > \theta\}$, where Φ is the advantage deriving from the contractual form and ε an exogenous shock. There are three reasons for this form.

First, **additivity** is essential. The structure in which a disadvantage in Φ can be compensated by an advantage in ε is the source of the correlation induced by conditioning. A multiplicative form $S = \mathbf{1}\{\Phi\varepsilon > \theta\}$ yields results of the same kind, but loses the reading in terms of compensation.

Second, it must be a **threshold**. Survival is a binary event rather than a continuous quantity, so what is conditioned on is a set. Were S continuous, this would be a problem of regression, not of conditioning.

Third, the sign restriction $a, b > 0$ states that **both contribute positively to survival**. If one were negative the result would reverse. The statement in Section 12.2 that “the direction of the distortion reverses with what is conditioned on” depends on these signs.

Proposition A.1 (Collider bias). Let Φ and ε be independent random variables, and for $a, b > 0$ and a threshold θ define the survival indicator $S = \mathbf{1}\{a\Phi + b\varepsilon > \theta\}$. Then conditioning on $S = 1$ makes Φ and ε negatively correlated:

$$\text{Cov}[\Phi, \varepsilon \mid S = 1] < 0.$$

Proof. Fix a realization x of Φ . The condition $S = 1$ is equivalent to $\varepsilon > (\theta - ax)/b$. The right-hand side is decreasing in x , so the conditional distribution $\varepsilon \mid (\Phi = x, S = 1)$ is stochastically decreasing in x . Hence

$$\mathbb{E}[\varepsilon \mid \Phi = x, S = 1]$$

is decreasing in x . If the joint distribution of Φ and ε is non-degenerate the decrease is strict, and in the definition of the covariance

$$\text{Cov}[\Phi, \varepsilon \mid S = 1] = \mathbb{E}[(\Phi - \mathbb{E}\Phi)(\mathbb{E}[\varepsilon \mid \Phi, S = 1] - \mathbb{E}\varepsilon) \mid S = 1]$$

the two factors inside the bracket move in opposite directions, so the covariance is negative. $\square$

Corollary A.2 (Understatement of the structural effect). In a surviving sample, a disadvantage in Φ is partly offset by an advantage in ε . Estimators of the effect of Φ on outcomes are therefore biased towards zero.

A.1.2 Implications not treated in the body

Remark A.3 (Strength of conditioning and size of the distortion). The smaller $\mathbb{P}(S = 1)$, the larger the absolute conditional correlation. Raising the threshold θ requires ever more extreme combinations of Φ and ε for survival, which strengthens the compensating relation between them.

This has a practical implication: **the more severely the sample is narrowed, the larger the distortion**. The listed firms treated in Section 12.9 (0.11% of the whole) are subject to a double conditioning — listing as well as survival — so the distortion exceeds that of simple survival conditioning.

Remark A.4 (Conditioning in the other direction). The same mechanism operates when only *exited* firms are observed, that is $S = 0$. Since $\mathbb{P}(\varepsilon < (\theta - ax)/b)$ is increasing in x , Φ and ε are negatively correlated under $S = 0$ as well.

Hence **collecting only failures does not remove the distortion**. Only a sample containing both successes and failures — that is, an unconditioned sample — is unbiased. This explains the importance of the fact that in Part IV the media panel for family 2 contained growing and shrinking segments at the same time.

Remark A.5 (Limits of the proof). Proposition A.1 assumes **unconditional independence** of Φ and ε . If the two are correlated to begin with, the negative contribution of conditioning mixes with the original correlation and the sign is indeterminate.

In reality contractual form and exogenous shocks may not be independent. A firm that chooses a Φ robust to the business cycle has, in effect, chosen an environment with small variance of ε . In that case the proposition does not apply.

Example A.6 (Numerical confirmation). Let $\Phi, \varepsilon \sim N(0, 1)$ be independent and $S = \mathbf{1}\{\Phi + \varepsilon > 0\}$. Then $\Phi + \varepsilon \sim N(0, 2)$ and $\mathbb{P}(S = 1) = 1/2$. It is a known result that under $S = 1$, $\text{Corr}[\Phi, \varepsilon \mid S = 1] \approx -0.57$. A correlation that is 0 unconditionally swings sharply negative merely by conditioning on survival.

A.2 Non-identification of age, period and cohort

A.2.1 Why the argument is made in a linear model

The claim of Section 12.5 is stated here in the language of linear algebra. The linear setting is adopted **not because non-identification arises from linearity, but because it appears most clearly in the linear case.**

In a non-linear model (α, π, γ) can sometimes be identified: if, for instance, the age effect is quadratic and the period effect linear, the two can be distinguished. But such identification **depends entirely on the assumed functional form.** When changing the assumption changes the conclusion, one can hardly say the parameters are identified.

The linear case is presented to make clear that the *core* of the non-identification lies in the identity $a = p - c$ and cannot be evaded by contrivances of functional form.

Proposition A.7 (Non-identification in APC). This proposition follows the notation of the age–period–cohort literature. Within this section a, p, c, π, y denote age, period, cohort, the period effect and the observation, and bear no relation to their use elsewhere in this work (allocation received, unit price, marginal cost, settlement, outcome).

Consider the model explaining an observation y as a linear sum of age a , period p and cohort c ,

$$y = \mu + \alpha a + \pi p + \gamma c + \epsilon.$$

When the identity

$$a = p - c \tag{12.2}$$

holds, (α, π, γ) is not identified.

Proof. For any $\eta \in \mathbb{R}$ set $(\alpha', \pi', \gamma') = (\alpha + \eta, \pi - \eta, \gamma + \eta)$. By (12.2),

$$\alpha' a + \pi' p + \gamma' c = \alpha a + \pi p + \gamma c + \eta(a - p + c) = \alpha a + \pi p + \gamma c.$$

That is, the fitted values agree for every η . The columns (a, p, c) of the design matrix are linearly dependent, the rank falling short by one, so the least-squares solution is not unique. $\square$

A.2.2 Implications not treated in the body

Remark A.8 (What is identified and what is not). What is unidentified is the *level* of the three coefficients; not everything is unknown.

Any quantity left invariant by the transformation $(\alpha + \eta, \pi - \eta, \gamma + \eta)$ of Proposition A.7 is identified. For instance $\alpha + \gamma$, and the **second differences** of the coefficients — that is, curvature — do not depend on η .

Hence “does CCC rise monotonically with age?” is not identified, but “is there curvature in the age effect?” is. Abandoning the separation of the three variables in Section 12.5 **does not mean abandoning curvature as well**. This work does not pursue that direction.

Remark A.9 (A constraint for identification). Identification requires adding one exogenous constraint. Typically one imposes $\alpha = 0$, $\pi = 0$ or $\gamma = 0$, or assumes two adjacent coefficients equal. Every such constraint is untestable, and the conclusion depends on the assumption. This work imposes no constraint and adopts the policy of stating the non-identification explicitly (Section 12.5).

Remark A.10 (How it appears in this work). The APC problem is latently at work in several places here.

Chapter 16 treats the time series of CCC; if the age distribution of firms changes over time, the observed rise in CCC is a mixture of period and age effects. The shift-share decomposition of Section 16.8 separates industrial composition but **does not separate the age composition**.

If younger firms have shorter CCC, a rising share of older firms could account for the rise in CCC. This work has not tested that.

A.3 Preference for variance under a deadline

The proposition of Section 8.6 is proved here.

Proposition A.11 (Under a deadline, minimizing variance is not optimal). Let $\Theta > 0$ be the time required to arrive, $T_{\max}$ the deadline, and define success as $\{\Theta \leq T_{\max}\}$. When $\mathbb{E}[\Theta] > T_{\max}$, an increase in variance holding $\mathbb{E}[\Theta]$ fixed can increase the probability of success. In particular the probability of success converges to 0 in the limit $\text{Var}[\Theta] \rightarrow 0$.

Proof. If $\text{Var}[\Theta] = 0$ then $\Theta = \mathbb{E}[\Theta] > T_{\max}$ with probability one, so $\mathbb{P}(\Theta \leq T_{\max}) = 0$. On the other hand, giving the variable variance while preserving the mean — say a two-point distribution with $\Theta = T_1 < T_{\max}$ with probability q and $\Theta = T_2$ with probability $1 - q$, chosen so that $qT_1 + (1 - q)T_2 = \mathbb{E}[\Theta]$ — gives a success probability of $q > 0$. The increase in variance thus improves the probability of success. $\square$

Remark A.12 (Contrast with the floor constraint). This proposition does not contradict the conclusion of Section 8.5, “lower the variance with many small amounts”. The latter holds under the floor constraint $M \geq 0$, the former under the deadline constraint $T_{\text{ins}} \leq E(0)/c_{\text{liv}}$. **Change the type of constraint and the optimal attitude to variance changes.**

Remark A.13 (Implications of taking the probability of success as the objective). Proposition A.11 takes the objective to be $\mathbb{P}(\Theta \leq T_{\text{max}})$, which differs from $\mathbb{E}[\sum \beta^t \phi_t]$ in Section 4.2.

The two diverge because **the loss from exceeding the deadline does not depend on by how much it is exceeded**. Missing T_{max} by a day and by a year give the same result, so the objective is binary. This discontinuity reverses the attitude to variance.

The insolvency time τ of Section 4.2 is likewise an absorbing state with the same property. That is, **the framework of this work contains regions in which the objective is not smooth**. The pathwise constraint of Proposition 4.7 is a manifestation of this.

Remark A.14 (Limits of the proof). The proposition shows that an increase in variance *can* improve the probability of success, not that it always does. If $\mathbb{E}[\Theta] < T_{\text{max}}$, an increase in variance lowers it instead.

The boundary is the comparison of $\mathbb{E}[\Theta]$ with T_{max} . In practice this becomes the judgement **play safe if the prospect suffices, gamble if it does not**, but $\mathbb{E}[\Theta]$ is not usually knowable in advance. This work does not treat that estimation problem.

A.4 Substitutability of capital and credit

Proposition A.15 (The substitution relation). When $A = \text{Inv} = 0$, the constraint $M \geq 0$ is equivalent to

$$E + \sum_i \max\{-\kappa_i, 0\} \geq \sum_i \max\{\kappa_i, 0\}. \quad (\text{A.1})$$

That is, E and negative κ are perfect substitutes in the constraint.

Proof. Substitute $A = \text{Inv} = 0$ into equation (3.8) of Proposition 3.15 and rearrange $M \geq 0$; the claim follows immediately. The two terms on the left enter with the same coefficient 1, so a decrease in one is exactly compensated by an equal increase in the other. $\square$

Corollary A.16 (Implications in both directions). If $E \approx 0$ then $\kappa < 0$ is required (Section 8.5). Conversely, if $E > 0$ is large, $\kappa > 0$ can be held (Part IV, the result on implication F).

A.4.1 The meaning and the limits of “perfect substitution”

Remark A.17 (The substitution holds only in the constraint). Proposition A.15 states that E and $-\kappa$ are equivalent **only within the single inequality** (3.8). It does not mean the two are economically equivalent.

They differ in three respects.

First, **there is a deadline**. $\kappa < 0$ carries an obligation to deliver, and if the τ_i of Section 3.1.4 is finite, delivery equivalent to repayment occurs on the due date. E has no deadline.

Second, **availability differs**. E is raised from the capital market, $\kappa < 0$ from a counterparty. As Section 3.1.1 sets out, the two are explained by different theories and become available under different conditions.

Third, **the claim on ϕ differs**. The provider of E holds a residual claim (Definition 3.21); the provider of $\kappa < 0$ does not.

Remark A.18 (Where the limits of substitution appear). The one-person business of Chapter 8 presupposes $E \approx 0$ and therefore treats only one endpoint of the substitution.

When $E > 0$ was introduced in Section 8.6, **the type of constraint changed from a floor to a deadline**. This is a manifestation of the asymmetry between E and $\kappa < 0$: E is consumed as living expenses and declines, whereas $\kappa < 0$ is not consumed but extinguished by delivery.

Proposition A.15 is thus a claim about a *static* constraint; behaviour through time differs.

AppendixB

Notation

Basic quantities and maps

TableB.1: Notation: basic quantities and maps.

Symbol	Name	Definition or first appearance
$X(i, t, \omega)$	cash flow	equation (2.9)
δ_i, π_i	delivery, settlement	Chapter 2
$\bar{\delta}$	contractual entitlement	Section 2.7
D_i, P_i	cumulative delivery, cumulative settlement	Chapter 2
κ_i	credit position	equation (3.2)
τ_i	average lag from delivery to settlement	Section 3.1.4
τ_{inv}	residence time of inventory, Inv/r	equation (5.8)
W	working capital, $\sum_i \kappa_i + \text{Inv}$	Definition 3.2
$M(t)$	cash balance; $M = E - W - A$	equation (3.8), Proposition 3.17
A, Inv	fixed assets, inventories	Proposition 3.15
Φ	business model	equation (2.8)
$C(\delta), c(\delta)$	total cost, marginal cost; c always carries its argument	Definition 2.14
F, p	fixed part of the fee, unit price	equation (2.10)
h, y	outcome-contingent map, outcome	ibid.
E	equity; paid-in capital plus retained earnings	equation (3.10)

(continued on the next page)

Symbol	Name	Definition or first appearance
κ_K^i	credit position with capital providers; agent-dependent	Definition 3.21
V_i	value of the residual claim under agent i 's measure	ibid.
d, I	depreciation as a rate on sales, capital expenditure	Corollary 5.7
$\mathcal{D}$	the space of deliveries	Part I, primitives
$\mathcal{F}, \mathcal{F}_t$	the σ -field of events and the filtration	ibid.

Derived quantities and notation of the method

TableB.2: Notation: derived quantities and notation of the method.

Symbol	Name	Definition or first appearance
$\mathcal{G}$	verifiable shared information	Chapter 4
H	horizon (planning period)	ibid., primitives
T_{ins}	time of insolvency; distinguish from τ_i	Proposition 4.7, equation (5.7)
β_i	discount factor of agent i	Section 4.2
r, m	sales velocity, gross or operating margin	Chapter 5
CCC	cash conversion cycle	equation (5.1)
$r^{\max}$	upper bound on scale, $(E - A)/\text{CCC}$	equation (5.4)
g^*	self-financeable growth rate	equation (5.6)
$\phi_{\text{prod}}, \phi_{\text{barg}}, \phi_{\text{cog}}$	three-way decomposition of surplus	equation (6.1)
ϕ_{risk}	payment for transfer and pooling; an unseparated sum of ϕ_{prod} and ϕ_{barg}	Remark 6.1
$\nu, \lambda(\nu)$	transaction frequency, decay rate	equation (6.4)
ε	exogenous shock	Section 12.2
S	indicator of survival or disclosure	Section 12.2

(continued on the next page)

Symbol	Name	Definition or first appearance
a, p, c	age, period, cohort; within that section only	equation (12.2)
Cap	capacity limit	equation (8.2)
Cap _{cl}	the part of capacity allocated to commissioned work	equation (8.14)
c_{liv}	rate of living expenses	equation (8.11)
$\mathcal{A}$	the set of responses; distinguish from fixed assets A	Definition 9.4
T_{resp}, u	response deadline, length of an interval of unavailability	equation (9.3)
$T_{\text{dev}}, T_{\text{acq}}, T_{\text{credit}}, T_{\text{inst}}$	the terms of the start-up cost	equation (8.11)
λ_{novel}	arrival rate of unanticipated states	Chapter 9
$\hat{\Omega}$	the state space anticipated at design time	Chapter 9
DSO, DIO, DPO	days sales outstanding, days inventory outstanding, days payable outstanding	Chapter 16

AppendixC

Correspondences with existing theory

Those components of this work that have counterparts in existing theory are set out here. The contribution specific to this work is confined to the lower tier.

The work divides into three layers. The first is the existing theory explaining each term of κ_i , introduced as a foundation. The second connects those theories within a single framework, and is where the contribution of this work lies. The third is the organization specific to this work.

Layer one: existing theory explaining each term

Term	Theory	Field
$\kappa_C > 0$, inter-firm	the four types of trade credit [14, 15]	corporate finance
$\kappa_C < 0$, inter-firm	reverse trade credit [12, 7]	corporate finance
$\kappa_C < 0$, to consumers	prepaid contracts and breakage [6, 9, 13]	consumer behaviour
$\kappa_L < 0$	the bonding argument [1]	labour economics
accumulation of a public record	labour-market signalling [16]	economics of information

TableC.1: Correspondences, layer one: existing theory explaining each term of κ_i .

These literatures do not cite one another (Remark 3.3).

Layer two: the contribution as connection

Connection	Content
Unification through κ_i	the objects of four fields enter the same sum in (3.8)
Explaining the level through τ_i	Proposition 5.16, Corollary 5.17
A typology of Φ	Part II; existing theory has no classification of contractual forms
The three-way decomposition of ϕ	equation (6.1); durability by source

TableC.2: Correspondences, layer two: the contribution of this work as connection.

Layer three: individual correspondences with existing theory

TableC.3: Correspondences, layer three (a): individual correspondences.

Content of this work	Corresponding existing theory
The measurability condition for contracts (Chapter 4)	contract theory, the informativeness principle [17]
The empirical difficulty of the risk–incentive trade-off (Section 15.3.5)	[24, 25]
Measurement of ϕ_{cog} (Section 14.6.1)	[6, 9]
Switching costs of Φ and the value of waiting (Remark 2.18)	irreversible investment and real options [8]
The mechanism by which pooling creates correlation (Remark 2.11)	diversification and systemic risk [30, 19]
Change of goals with the time horizon (Section 8.6.4)	socioemotional selectivity theory [3]
The optimal strength of outcome contingency (Example 4.5)	linear contracts and multitasking [18]
The g^* of equation (5.6)	the self-financeable growth rate [4]; distinct from the Higgins-type sustainable growth rate, which fixes the payout ratio and leverage
Equation (3.9) ($M = E - W - A$)	a standard identity of financial analysis: cash = equity – net operating assets

(continued on the next page)

Content of this work	Corresponding existing theory
Proposition 5.2 and Corollary 5.3 (the bound on scale)	a formalization of overtrading; the structure in which working capital is bounded by net worth is isomorphic to [33]
Proposition 6.6 (cash neutrality of cognitive surplus)	transfer from contract liabilities in revenue recognition; standard accounting treatment
Proposition 5.21 (the gap between peak and average)	the peak funding requirement of practice; a standard topic of working-capital management
Corollary 5.19 (non-additivity of CCC)	a general property of aggregating ratios; the shift-share decomposition (16.5) is its difference form
Conditioning on a collider (Proposition A.1)	Berkson's paradox [2], the framework of causal inference [23], the causal-graph representation of selection bias [11]
Non-identification in equation (12.2)	the age-period-cohort problem [22]
Equation (2.11)	the pooling of variance; a standard result of actuarial science
Heterogeneity of discount factors (Proposition 4.10)	time preference, hyperbolic discounting [21]
The distinction between risk and uncertainty in automation (Chapter 9)	[20]
Organizations absorbing measurability (Section 8.7)	the boundary of the firm [5], property-rights theory [10]
The separation of payer and beneficiary (family 7)	two-sided markets [26]
Proposition A.15 (substitution of capital and credit)	collateral constraints and net worth [33, 36, 37]
The bootstrapping problem for credit (Section 8.5)	ibid.; corresponds to the result that agents with lower net worth face tighter constraints
Enlargement of the contract space by a protocol (premises (1) and (2) of Chapter 10)	distributed consensus and contractibility [31]
The divergence of identifier and agent (Section 10.3)	the Sybil attack [32]
Corollary 10.11 (exclusion of agents holding no assets)	collateral requirements and the limits of financial inclusion [34]

(continued on the next page)

Content of this work	Corresponding existing theory
Remark 10.4 (intermediation “outside” the partition)	importing external states and their verifiability [35]
Proposition 10.6 (identifying the beneficiary)	free generation of identifiers and false-name-proofness [38]
Theory that should be connected but is not	Corresponding existing theory
Joint determination of $\bar{\delta}$ and δ (Section 19.4.1)	mechanism design, self-selection and quantity discounts
The screening effect of contracts (ibid.)	the part of contract theory treating adverse selection
Joint determination of DSO and DPO (ibid.)	the theory of trade credit
The relation of size to the quality of placements (Section 15.3.5)	intermediation and match quality in labour markets

TableC.4: Correspondences, layer three (b): theory that should be connected but is not.

Organization specific to this work	Correspondence
The correspondence of three indices with three operations (Section 2.6)	this work
The definition through Φ (Chapter 2)	this work
The three-way decomposition of surplus (equation (6.1))	this work
The layer structure (Section 6.3)	this work
The catalogue of types of Φ (Part II)	this work
The six-way taxonomy of verification obstacles (Section 13.2)	this work

TableC.5: Correspondences, layer three (c): organization specific to this work.

AppendixD

List of pre-registered implications

This is the list of the advance commitments described in Section 19.2. The main text states only the results; the wording of each registration and its verdict are listed here.

There are 22 registrations: 6 supported (3 of them partially), 11 rejected, and 5 untestable.

D.1 The list

TableD.1: Implications fixed in advance, and their verdicts.

Symbol	Subject	Registered implication and result	Verdict
A	Family 2	Divergence and transaction frequency. The level is supported, but what was measurable was κ	Partial
B	Family 5	Refunds and measurability. The sign is not uniquely determined	Untestable
C	Fees	Payment for externalization. A counterexample not explained by function was also found	Partial
D	Unattended operation	Number of dependencies and length of unattended period. The same implication as $\Phi 1$, rejected with it	Rejected
E	Corporate statistics	The leading property of CCC ($p = 0.72$)	Rejected
F	Corporate statistics	Size and the direction of credit. The sign is reversed	Rejected
J	Fees	Rate differences correspond to the correlation ρ . ρ is not observable	Untestable
K	ϕ_{cog}	$\delta^* > \delta$. The sign is negative in all four categories	Rejected

(continued on the next page)

Symbol	Subject	Registered implication and result	Verdict
L	Capacity constraint	Contract choice of older cohorts. No consumer-side data	Untestable
M	Corporate statistics	The composition effect in the rise of CCC. Within-industry is 83%	Rejected
N	Corporate statistics	Positive correlation of $\Delta \ln W$ and $\Delta \ln r$ ($\beta = 0.83$)	Supported
O	Corporate statistics	Lag structure. Not detectable in annual data	Rejected
P	Corporate statistics	Total elasticity including the lag. $0.830 \rightarrow 0.821$	Rejected
T	Labour side	Movement of κ under the new statute. No statistics exist	Untestable
U	Sole proprietorships	The operating margin is close to the value with officers' remuneration added back	Supported
X	Family 5	Negative correlation of fee rate and turnover rate. No correlation across 307 firms	Rejected
Y	Family 5	Ordering by occupation of the correlation between size and turnover. Not reproduced nationally; only the qualitative division holds	Partial
Z	Family 5	Occupations with larger B1 have stronger refunds. Rank correlation $\rho = 0.000$	Rejected
$\Phi 1$	Procedure	Number of dependencies and time to repair. Concentration of replicates was producing the result	Rejected
$\Phi 2$	Procedure	$M(t)$ declines ahead of discontinuation ($p = 3.2 \times 10^{-9}$)	Supported
$\Phi 3$	Procedure	Non-negativity of CCC. The corporate sector is non-negative in every industry too, so no contrast can be drawn	Untestable
$\Phi 4$	Procedure	The type in family 7. Rarity was right; the surviving type was wrong	Rejected

G, H and I targeted statutory regulation of fee rates, but were withdrawn at the design stage once the population turned out to contain only three cases. Q, R and S were registered in the examination from the worker's viewpoint, but testing has not been begun.

D.2 Where registration changed the verdict

Recorded here are the places where, absent an advance commitment, a retrospective interpretation would likely have been adopted.

- In family 5, the fact that 99.7% of capped-rate fees have disappeared could have been read as circumstantial support for the measurability hypothesis. In fact the price level of the capped scheme is simply an order of magnitude lower, and the two are unrelated.
- In Chapter 16, the contemporaneous negative correlation between CCC and sales could have been read as “CCC works as a business-cycle indicator”, and the credit position of large firms as “large firms support their counterparties”. Both are retrospective narratives.
- In $\Phi 2$, looking only at the fact that the balance of discontinued Φ declines beforehand would read as support. Applying the same window to continuing Φ , the two do not separate at 180 days. Without having decided to place a control, the overall downward trend would have been misread as support.
- In $\Phi 4$, the fact that the sole family-7 instance to appear places the selection of beneficiaries outside the procedure could have been offered as circumstantial support for Proposition 10.6. But the registered wording was “what survives is 7-4”, and what appeared was 7-2. A mechanism agreeing with a proposition does not rescue an implication that failed.

D.3 Defects that arose in operating the registry

D and $\Phi 1$ are the same implication registered twice. They were registered separately at stages 13 and 16, and the duplication went unnoticed at the time. The cause was keeping the register of implications by chapter and never reconciling it as a whole.

AppendixE

Revision history

E.1 History of versions

This is a theory under construction, and its structure and content have been remade several times. The principal stages are recorded here. The names given to the versions are matters of convenience.

The version number corresponds to the stage: stage n is $v0.n.0$. A change of structure, the addition or withdrawal of a proposition, or new empirical work raises the stage; corrections of slips and adjustments of format are handled in the third digit ($v0.8.1$, for instance).

TableE.1: History of versions: the version number, change of structure, and principal content of each stage.

Stage	Version	Structure	Principal content
Stage 1	$v0.1.0$	chronological	the definition of Φ , 27 types, the four-way classification of method
Stage 2	$v0.2.0$	reordered logically	five parts: theory, catalogue, method, empirics, conclusion. Proofs and numerical examples added
Stage 3	$v0.3.0$	change of parts	deductions under constraints moved from “applications” to “the space of maps and constraints”
Stage 4	$v0.4.0$	empirics added	CCC and g^* measured from the Financial Statements Statistics
Stage 5	$v0.5.0$	weaknesses classified	independence, exogeneity, and level versus change organized as three types
Stage 6	$v0.6.0$	τ introduced	differences in level deduced from $CCC = \sum \tau_i + \tau_{inv}$

(continued on the next page)

Stage	Version	Structure	Principal content
Stage 7	v0.7.0	existing theory integrated	existing theory matched to each term of κ_i ; Appendix C made three-layered
Stage 8	v0.8.0	definitions filled in	the split of costs, E , κ_K , and the introduction of d and I into g^*
Stage 9	v0.9.0	applied chapter added	deductions for replacing the supplier of the information structure (Chapter 10); the catalogue corrected to 28 entries
Stage 10	v0.10.0	propositions filled in	simultaneity, identification of the beneficiary, and growth through prepayment added to Chapter 10; indivisibility imposed in Definition 10.1
Stage 11	v0.11.0	method filled in	revision requests recorded in the empirical chapters reflected into Chapter 12; failure of identification and the choice of variables added as sections
Stage 12	v0.12.0	empirical design placed first	the five stages connecting theory to observation placed as Chapter 11 at the head of Part III; the distinction of comparison from testing stated as a discipline
Stage 13	v0.13.0	operationalization filled in	availability written as an equation and rules for assignment to types laid down; open problems reorganized into 21 items
Stage 14	v0.14.0	measurement corrected	the effect of the scope of inventories on CCC and on the elasticity measured for ten industries (Section 16.10)
Stage 15	v0.15.0	further registrations	the correlation of size with turnover read as a proxy for B1, and implications Y and Z registered (Section 15.3.5)
Stage 16	v0.16.0	a route to testing Part II	implications $\Phi 1$ – $\Phi 4$ registered for the procedural records (Section 18.1)
Stage 17	v0.17.0	implication Y tested	a national sample retrieved and tested on the part excluding Tokyo; partial support (Section 15.3.5)
Stage 18	v0.18.0	descriptions removed, format	a description resting on a single observation and an $n = 2$ table removed; theorem-type body text unified in roman

(continued on the next page)

Stage	Version	Structure	Principal content
Stage 19	v0.19.0	means of reproduction included	data and processing placed in the repository, and the description in Appendix F brought into line with the fact
Stage 20	v0.20.0	open problems reorganized	completed items dropped and errors of classification corrected, leaving 19 items
Stage 21	v0.21.0	measurement extended, handling stated	the scope of inventories extended to 62 industries, weakening the verdict of stage 14 (Section 16.10); the handling of per-firm data stated in Appendix F
Stage 22	v0.22.0	grounds for selection, consistency of references	the reason for choosing families 2 and 5 derived from the three-way decomposition and the degrees of freedom (Section 11.5); a contradiction in the description of ϕ_{cog} resolved (Section 14.6.1)
Stage 23	v0.23.0	implication Z tested	907 refund documents read, strength measured, and implication Z rejected (Section 15.3.5); registrations reached 18 and open problems 18
Stage 24	v0.24.0	implications $\Phi 1$ – $\Phi 4$ tested	the deductions of Part II tested for the first time using the records of a public distributed ledger (Section 18.2); only $\Phi 2$ supported; registrations reached 22
Stage 25	v0.25.0	quantities measured incidentally	λ and the duration of procedures measured, giving an upper bound in Proposition 9.2 (Section 18.3); implication D found identical to $\Phi 1$ and its verdict changed to rejected
Stage 26	v0.26.0	organized as a book	22 chapters merged into 20 and the symbols of the implications cleared from the body; the wording of the registrations and their verdicts gathered into Appendix D, and the empirics of the procedure moved to the empirical part as Chapter 18; verification obstacles presented as six classes from the outset

(continued on the next page)

Stage	Version	Structure	Principal content
Stage 27	v0.27.2	Part I systematized	the chain of derivations starting from Φ carried through. Chapters 2 and 4 merged, reducing seven chapters to six; Δ and Π defined from the primitives and $X = \Phi(\delta)$ adopted. The degrees of freedom of Φ organized into three groups and matched to the three ways of operating on risk
Stage 28	v0.28.0	deductions gathered under the fundamental inequality	the cash constraint arranged as $CCC \cdot r \leq E - A$ (Proposition 5.2), from which the bound on scale, the time of depletion, the peak, the comparative statics and the switching of the binding constraint are derived. Antisymmetry and additivity of κ added, and what the aggregate CCC measures identified. The funding constraint and the deadline of Chapter 8, previously introduced separately, recast as special cases of the fundamental inequality

E.2 Claims withdrawn

Claims have been withdrawn several times in the course of this work. What was withdrawn and why is recorded here. Since the body retains only the corrected description, which claims were once thought to hold appears in this appendix alone.

So that a reader partway through is not misled, the body is kept consistent with the corrected content. What follows is a record of how this work erred, and is not needed to understand the body.

E.3 Corrections deriving from measurement

Original description	After correction	Occasion
The rise in CCC lowered g^* by a third over a quarter-century	m rose 91%, exceeding the 49% deterioration of the denominator, and m/CCC rose from 38.5% to 49.2%	obtaining m (Proposition 14.2)
The constraint on small firms is credit (κ)	what actually binds is m ; the CCC of small firms is shorter than that of large firms	measurement by size stratum (Proposition 14.6)
ϕ_{cog} has been measured nowhere in the world	instances of measurement exist, and the method is established	checking prior work (Section 14.6.1)
The low m of small firms is due to insufficient spreading of fixed costs	officers' remuneration is the main cause; adding it back erases the ordering by size	obtaining the breakdown of personnel costs (Proposition 14.8)
In R2, ϕ is demoted from the objective to a constraint	withdrawn; the change is explicable by H and β , and the demotion has no content	checking prior work (Remark 8.4)
Time dependence of the capacity constraint should be taken into the theory	withdrawn; it would rewrite the generating set of 27 types from an $n = 1$ observation	recognizing the asymmetry of the abstract and the concrete (Remark 16.10)
The ill-posedness of implication B is a failure of the author's design	it is the risk–incentive trade-off dispute, unresolved for a quarter-century	checking prior work (Section 15.3.5)

TableE.2: Corrections deriving from measurement: the original description, the corrected content, and the occasion.

The first of these is the weightiest. Having observed only the deterioration of CCC, the work came close to concluding that g^* had fallen; but Corollary 5.7 is a ratio, and no conclusion follows from a change in the denominator alone. The error would have been avoided by obtaining m , and in fact the work had halted with the note that “interpretation is reserved because the change in m has not been measured”. **Reserving judgement turned out to be right.**

E.4 Corrections deriving from reasoning

Original description	After correction
Credit was “being sold continuously” during employment	original attribution, as with works made for hire: it does not arise in the individual and then transfer
The credit asset is greatest immediately after leaving employment	what is greatest is not credit but the resource for conversion (contacts, technical knowledge); credit remains zero
λ_{novel} is the sum of its components	it is not additive: with a lock file, λ_{dep} is conditional on λ_{sec}
Family 5 has favourable data conditions	at the aggregate layer it is weaker than family 2; the claim holds only on descending to the per-firm layer
Difficulty of access to real data is the principal obstacle to verification	the principal obstacle for family 5 is (ii) ill-posedness of the hypothesis; family 2 is (iv) access constraint, but in the peculiar form that the required granularity is never published

TableE.3: Corrections deriving from reasoning: the original description and the corrected content.

E.5 A record of pre-registration at work

The following are places where, had the implication not been fixed in writing beforehand, the work would likely have drifted into retrospective interpretation.

- (1) **The disappearance of capped-scheme fees.** That 99.7% of the capped scheme has disappeared in family 5 could have been read as circumstantial support for the measurability hypothesis. In fact the price level of the capped scheme is simply an order of magnitude lower, and it has nothing to do with measurability.
- (2) **The contemporaneous negative correlation of CCC with sales.** A negative correlation was obtained in 31 of 45 industries ($p = 0.008$). This could have been read as “CCC works as a business-cycle indicator”, but it is not distinguished from the denominator effect.
- (3) **The credit position of large firms.** That DSO – DPO is large for large firms could have been interpreted as “large firms support their counterparties”. In fact it is the result of implication F being rejected, sign and all.

The retrospective narrativization described in Section 12.6 of others' accounts arises equally in the author's own interpretation.

E.6 Making assumptions explicit

Several propositions here were originally stated without making their independence assumptions explicit. In the course of verification it emerged that quantities in fact determined jointly were being treated as independent, and the assumptions were written into each proposition.

Proposition	Assumption added	Occasion
2.9	ρ does not depend on n	examining implication J
2.22	$\bar{\delta}$ is given exogenously	testing implication K (tiered menus)
2.24	m and $\mathbb{E}[\delta]$ are independent	ibid.
5.5	m and CCC are independent	obtaining m (Proposition 14.2)

TableE.4: Independence assumptions added to each proposition, and the occasion for adding them.

The endogeneity of $\mathcal{G}$ (Remark 4.3) was added at the same time. None of the proofs was changed: **only the range of application was narrowed.**

A passage using the undefined symbol $\mathcal{X}$ as the domain of the production function was rewritten without naming a domain. Since this work does not treat the internal structure of f , no space of inputs need be introduced. $\mathcal{D}$, $\mathcal{F}$, H and ϕ_{risk} were added to the notation list.

E.7 Introducing the lag

The τ_i of Section 3.1.4 did not exist in the original version. It was introduced in the course of working on weakness (C) organized in Section 19.4.3, the failure to separate level from change.

The decision to introduce it was made after first confirming that it would be used in several places: the explanation of differences in level (Corollary 5.17), the impossibility of cross-sectional comparison (Corollary 5.18), and the decomposition of change (Corollary 5.20).

Distortion of CCC during growth and a reinterpretation of the β estimate were originally also expected as uses, but the rejection of implications O and P showed that these cannot be detected in annual data (Section 16.9.2). That use was removed from the grounds for introduction.

At the same time, Example 4.9 was corrected where it used CCC and W before their definitions, and rewritten in terms of κ alone.

E.8 Re-evaluating the relation to existing theory

In the draft of this work, some of the terms of κ_i were expected to be territory specific to it, without checking whether existing theory covered them. On checking the literature, theory existed for every one.

Original expectation	Existing theory found	Field
Treating prepayment as trade credit is unexplored	reverse trade credit [12, 7]	corporate finance
There is no theory of κ for deferred labour	the bonding argument [1]	labour economics
Accumulation of a public record is specific to this work	labour-market signalling [16]	economics of information
ϕ_{cog} in family 2-4 is unmeasured	empirical work on breakage of prepaid balances [13]	consumer behaviour

TableE.5: For each term of κ_i : the original expectation, the existing theory found, and the field.

This check moved where the contribution lies. What was originally sought was *territory without existing theory*; in fact the contribution is **introducing existing theory as the explanation of each term and treating them as a sum** (Remark 3.3). This is why Appendix C was recast in three layers.

The same failure has recurred. When implication B was set out in Section 15.3.5, the existence of the risk–incentive trade-off dispute had not been checked.

E.9 Filling in definitions

The following are not corrections of errors but additions where the definitions were thin.

Subject	What was added	Occasion
The technology side	costs split into $C_{\text{prod}} + C_{\text{admin}}$ (Definition 2.14)	the source of ϕ_{prod} was undefined
ϕ_{prod}	defined by marginal cost and a counterfactual price (Definition 6.2)	ibid.
Equity E	defined as paid-in capital plus accumulated ϕ (equation (3.10))	it was implicit as a residual
κ_K	defined as a residual claim (Definition 3.21)	equity was being treated identically to debt
g^*	recast to include depreciation d and investment I (Corollary 5.7)	$I \approx dr$ was implicitly assumed

TableE.6: What was added where definitions were thin, and the occasion.

The revision of g^* changes the measured values. In the stratum below ¥10 million, against a value of 20.9% from m alone, the nine-year average including I is 3.1%. The conclusion of Proposition 14.6 is maintained, but the burden of investment is added to the cause.

As to the split of costs, “ C does not depend on Φ ” was originally proved as a proposition. But that Φ does not appear once $C(\delta)$ has been defined is a consequence of the definition, and the proof was circular. It was withdrawn, rewritten as the definition of the split, and the eight routes by which Φ generates costs enumerated.

Defining κ_K revealed that subjectivity divides into three stages (Remark 3.23). The subjectivity of the valuation map v is asymmetry of information; the subjectivity of the value V_j of the residual claim is disagreement of probability measures. The former can be resolved by disclosure; the latter cannot.

E.10 Examining reversibility, and a correction

The framework was reviewed from the standpoint of irreversibility and checked against existing theory. One correction and two additions resulted.

Correction. In Remark 2.11, the mechanism making ρ endogenous was described as “objects with low correlation are exhausted”. According to [30] the correct mechanism is different: it is not that the objects of pooling run out, but that **the act of pooling itself creates correlation among those who pool**. If two agents each hold half of two risks, each becomes individually safer while the two become perfectly correlated.

Additions. The switching cost of Φ (Remark 2.18) and the social optimality of pooling (Remark 2.12) were recorded. The former corresponds to the theory of irreversible investment, the latter to research on systemic risk.

The framework of classifying reversibility by the size of a cost is itself already held by the theory of irreversible investment and is not a contribution of this work. The formulation that a decision is irreversible not because it cannot physically be undone but because undoing it is costly belongs to the line of [8].

E.11 Examination from the worker's viewpoint

Whether to treat workers as an independent subject was examined, and the conclusion was not to.

The organization initially attempted was “a worker is a special case of the one-person business”: a one-person business whose Φ is given and whose customers concentrate in a single firm. It emerged, however, that the constraints differ in both Φ and ϕ , and this organization was withdrawn (Remark 7.3). The Labour Standards Act constrains five of the six degrees of freedom of Section 2.5, and ϕ_{prod} and ϕ_{cog} have no counterparts.

Four implications were registered in the course of this examination.

Implication	Content	Result
Q	workers with fewer assets have shorter payment cycles	untested
R	workers hold a higher M_{floor} than firms	untested
S	the paucity of types of employment contract is due to statutory constraint	untested
T	the new freelance statute shifts the allocation of κ	untestable
U	the operating margin of individual enterprises is close to the corporate value with officers' remuneration added back	supported

TableE.7: Implications registered in the examination from the worker's viewpoint, and the results.

Implication T is untestable because κ for sole proprietors is not published. In this process an obstacle was found that does not fit the four classes of Section 13.2: the form in which **a survey is taken but not tabulated** (Remark 15.18).

E.12 Extending the taxonomy of obstacles

The four classes of Section 13.2 were constructed from the initial experience of investigation. Later investigation produced two obstacles that do not fit them.

Class	Occasion of discovery	Record
(v) choice of tabulation	seeking κ for sole proprietors: the questionnaire exists but the tables do not	Section 15.7
(vi) non-operation of the adjustment	in testing implication X, the regressor does not move although data for 307 firms exist	Remark 15.6

TableE.8: The obstacles that did not fit the four classes: class, occasion of discovery, and where recorded.

(vi) in particular carries large implications for the method of this work. Part III attributed the failure to verify to four causes: measurement, theory, identification and access. The case in which **reality does not behave as the theory posits** had not been anticipated.

At the stage of setting out implication B for the measurability hypothesis, it was written that b could be obtained only as a binary. Later investigation revealed that the realized fee rate is published as a continuous quantity, making the test possible. The obstacle was not the granularity of the data.

E.13 Extending the taxonomy of biases

The selection biases of Part III originally consisted only of those concerning the selection of the sample. Two distortions that the selection of the sample cannot explain emerged in the course of the investigation.

Class	Occasion of discovery	Record
Mistaking the unit	rows of a listing are offices while the values are replicated at firm level	Remark 15.11
Selection of variables	in family 5, the fields of greatest theoretical interest are those whose aggregates cannot yield the required quantity	Remark 15.3

TableE.9: Two distortions the selection of the sample cannot explain, and the occasions of discovery.

Both were added as sections of Chapter 12 (Sections 12.7 and 12.8), and the opening of that chapter was changed from “six kinds” to “eight kinds”.

The former concerns **the unit of observation rather than the selection of the sample**, the latter **variables rather than samples**. Both lay outside the range Chapter 12 originally treated.

The grounds for the former were replaced at stage 17. They originally rested on the behaviour that, in retrieval specifying a licence number directly, a number that failed to resolve returned the immediately preceding result. This admits interpretation as a measure limiting automated retrieval, and erecting a general obstacle from a single observation was a leap of the same form as in Section 16.11. The grounds were changed to the fact that in Section 15.3.5 the presence or absence of consolidation reversed the sign of the correlation. The description of that behaviour and the accompanying $n = 2$ table were removed.

At the same time, the point that a complete frame does not yield power if the population is small was added to Section 12.9. Section 15.3.5 had referred to that section with “as stated”, but the statement referred to did not exist there.

E.14 Reorganizing the list of open problems

The list, extended from 12 items to 21 at stage 13, was reorganized into 19 at stage 20.

- (1) **A completed item remained.** Geographical extension of the sample was completed when the national sample was retrieved for the test of Section 15.3.5. It nevertheless remained first among the items that could be started.
- (2) **One task was split in two.** Testing implication Z and collecting the refund rates and periods that are its regressor were listed as separate items. A test and the data missing for it are two faces of the same task; they were merged.
- (3) **A classification was wrong.** The curvature of the age effect was placed under “items with no obstacle”, but **pursuing it requires a sample carrying firm age.** The Financial Statements Statistics used here has three axes — industry, size and year — and does not include the year of founding. It was moved to (iv) access constraint.

Items with no obstacle thereby fell from six to three. **A list drifts from the actual state not when items are added but when finished items are left undeleted.**

E.15 The means of reproduction were not included

Appendix F stated that the processing used in the calculations was “attached”. **Not one file was in the repository.** A claim of the body was not satisfied by what was published.

Data and processing were included at stage 19, and three defects were fixed at the same time.

- (1) **The `parse.py` that Appendix F named did not run.** Its inputs were still absolute paths from an old environment, and running it failed with files not found. It was rewritten to read `raw/hojin/`.
- (2) **Some intermediate data had no means of regeneration.** No script existed to build the series for officers’ remuneration and depreciation, and they could not be restored from a fresh clone. `parse.py` was extended to all six series, and the regenerated values were confirmed to agree throughout with the previous ones.

- (3) A collision of output paths was resolved: the results of the test of implication E were being written under the same name as intermediate data.

After inclusion, it was confirmed from a state with `derived/` deleted that `parse.py` and all 14 analyses run.

E.16 Removing a description resting on a single observation

Chapter 13 recorded the behaviour that, in retrieval specifying a licence number directly, a number that failed to resolve returned the immediately preceding result, and generalized this into the obstacle “retrieval succeeds but the object is mistaken”.

It admits interpretation as a measure limiting automated retrieval, and other explanations were not excluded. Erecting a general obstacle from a single observation is a leap of the same form as the one Remark 16.10 warns against. The description and the accompanying $n = 2$ table were removed at stage 18.

The section of Chapter 12 itself was retained. Its grounds were changed to the fact that in Section 15.3.5 the presence or absence of consolidation reversed the sign of the correlation, and its subject changed from “mistaking” to **the mismatch between what a row denotes and what its values denote**. That appears at a magnitude sufficient to reverse the conclusion in a sample of 1,402 firm rows across four occupations (1,130 distinct firms), and is not a single observation.

E.17 Setting theorem-type body text in roman

Propositions and corollaries were being set in the `plain` style, which italicizes the body text. Japanese has no italic form, so it stayed upright, and **only the Latin letters and digits became italic, mixed in among it**. In places even the figures in tables placed inside propositions were italicized.

They were unified in the same roman as definitions, remarks and examples. The headings remain bold and numbered. Emphasis by `\emph` in the body that contained Latin letters and became italic was likewise changed to bold.

The heading of proofs was also changed. The default in `amsthm` is the English **Proof**, which does not suit a Japanese document; `\proofname` was set accordingly.

Italics thereby remain only in the punctuation attached to the proof heading and in journal names in the bibliography.

E.18 Testing implication Y, and a defect in the registration

At stage 17 a national sample was retrieved from the Jinzai Service General Site and implication Y tested on the part excluding Tokyo (Section 15.3.5). The verdict is partial support.

The registration was defective. Table 15.19 laid down that offices with amount-displayed fees would be used for implication Y, but the reference values of Table 15.18 against which the

comparison was made had been computed on a sample excluding them (Remark 15.12). **The composition of the reference values had not been checked at the time of registration.** Both rules are reported, and the matter handled by showing that the verdict does not depend on the rule.

Three points emerged during retrieval.

First, **rows of the search results are offices while the values are at firm level** (Remark 15.11). Taking correlations without consolidating gives, among physicians, a single company 134 times the weight, and the sign of the correlation turns from positive to negative. Consolidation was not preprocessing but a procedure determining the conclusion.

Second, **the displayed turnover rate cannot be computed from the placements and separations in the same row.** For physicians the two agree in only 115 of 479 rows. The composition matches Section 15.3.5 and does not impede comparison, but it was recorded that size and turnover are not quantities over the same period (Remark 15.13).

Third, **the initial tabulation dropped part of the sample.** Extraction was restricted to the licence-number letters Yu and Fu, so 200 Mu (free placement office) and 7 Toku records were missed. This was wrongly described as loss during copying, and up to 11.2% was reported as lost. **The number of data rows agrees with the search results; nothing was lost.**

The error changed the verdict. With the dropped sample, the ordering excluding Tokyo did not agree with the reference values, and it had been recorded that the ordering is not preserved. After correction the ordering agrees perfectly under the registered rule. The verdict was set to partial support not because the ordering is not preserved but because **the agreement of the ordering depends on the rule of tabulation, and that rule had not been aligned with the composition of the reference values at the time of registration.**

The reference values of Table 15.18 agree, on the Tokyo portion of the corrected sample under the same composition, to within 0.02 in correlation; nurses agree to three decimal places.

E.19 A route for testing the deductions of Part II

The deductions of Part II **had almost no pre-registered implications** up to stage 15. Of the 16 implications then registered, only D and T clearly targeted that part; D was supported as to direction only, and T is untestable.

At stage 16, implications $\Phi 1$ – $\Phi 4$ were registered for the “inside” partition of Chapter 10 (Section 18.1). In that partition the record of the procedure is itself the population, and halted procedures remain, so the survival conditioning of Chapter 12 does not arise. The reason for abandonment given in Chapter 17 — “no population frame” — disappears here.

Three points were distinguished before registering (Table 10.8).

- (1) **Quantities the procedure enforces are not objects of verification.** The collateral constraint of Proposition 10.9 is such a quantity: observing that collateral is posted is

merely reading the rules. Only quantities appearing as the result of an operator's choice are objects of verification.

- (2) **The population differs** (Remark 10.14). What can be verified is the correctness of the deduction, not whether the deduction applies to real one-person businesses.
- (3) **Quantities at the level of agents cannot be recovered** (Remark 10.5). What can be measured is only the distribution at the level of identifiers.

For $\Phi 1$ it was noted that, because the number of dependencies may correlate with usage, the alternative is hard to reject (Remark 18.1).

E.20 Registering implications Y and Z

The observation of Remark 15.8 — that the correlation of size with turnover differs greatly by occupation — **remained recorded as a secondary observation with no use** up to stage 14.

At stage 15 a route was set out for reading it as a proxy for the B1 side, the informativeness of effort that Table 15.13 records as lacking a proxy, and implications Y and Z were registered (Section 15.3.5).

Two points were fixed at registration.

First, **the test must not use the same sample** (Remark 15.10). The correlation was observed after the fact in the 307-firm sample, and testing on the same sample would be an error of the same form as retrospective narrativization. What is required is reproduction out of sample.

Second, **the mechanism is not identified** (Remark 15.9). Thinning of input and looser screening give the same sign, so the variable can be used as a proxy but no claim about the mechanism can be made.

There is a literature in labour economics on the relation between size and the quality of placements which this work has not surveyed; it was added to “theory that should be connected but is not” in Appendix C.

E.21 Checking the scope of inventories

Chapter 16 computed DIO from “finished goods or merchandise” alone, excluding work in process and raw materials. The limitation was recorded as “the effect on analyses of differences and correlations is small, but the absolute level cannot be interpreted”, **with no grounds given**.

At stage 14 the difference was measured against “inventories (end of period)” in the Financial Statements Statistics (Section 16.10). The result supports that statement: the level moves by more than 30% in some industries, but the correlation of ΔCCC is 0.976.

Incidentally, the elasticity β rose from 0.900 to 0.960 and the coefficient of determination improved from 0.092 to 0.213. **Working capital is more nearly proportional to sales when work in process and raw materials are included**. The contribution of the denominator effect moves smaller, and the conclusion of Section 16.9 is strengthened.

During estimation it emerged that, because the observations with $W > 0$ differ by definition, the sign comes out reversed unless the sample is aligned (Remark 16.9). Without alignment the values are 0.894 and 0.787, giving the opposite conclusion that the denominator effect increases.

E.22 Filling in operationalization and reorganizing the list of open problems

Classifying each item against the five stages of Chapter 11 revealed that **items halted at stage 1 (operationalization) were not registered in the list of open problems**. Although a defect at stage 1 requires no data retrieval and is filled by definition, the list was skewed towards problems of measurement and access.

The following were done at stage 13.

- (1) **Availability was written as an equation** (Definition 9.4). Previously it was stated only as “consuming the constraint” and had no form as an inequality. Whereas the capacity constraint bounds the total volume of delivery, availability bounds the length of an interval during which no response can be made. Both have the dimension of time.
- (2) **The reason the degree of automation cannot be defined was identified** (Remark 9.9). It had been recorded that “ δ has no unit”, but δ carries monetary units through the valuation map. Correctly, what v measures is value, not effort; and measuring on the response side leaves no measure defined on $\Omega \setminus \hat{\Omega}$. **The obstacle is the absence of a measure, not of a unit, and is not resolved by refining the definition.**
- (3) **Rules for assignment to types were laid down** (Section 7.9). Using the catalogue empirically requires a rule for assigning an observed Φ to a type, and none existed. Since the conditions identifying families are not exclusive, an order was fixed. The choice of order has no deductive grounds and is a convention, so it must be fixed in advance (Remark 7.1).
- (4) **A limitations section was placed in Chapter 9** (Section 9.4). Of the deductive chapters it alone had none, and its limitations were written in Chapter 17. The theoretical limitations were moved to that chapter and the empirical ones left in Chapter 17.
- (5) **Open problems were reorganized from 12 items to 21**. Nine items in the limitations sections and records of abandonment across the chapters were not reflected in the list. Six of them **fell under no obstacle at all and had simply not been started**, so a subsection was separated out for “items with no obstacle”. The table of priorities was likewise split into what can be started and what will not advance.

This changed the close of Chapter 19. It had read “no item remains that can be advanced by retrieving public data alone”; six items have no obstacle.

E.23 Placing the empirical design first

Chapter 13 recorded that the four-way classification should be placed ahead of the design of an investigation, because Part III treated selection bias in samples while having no procedure for judging “is verification possible at all?”.

This was carried out at stage 12 as Chapter 11. The four-way classification itself is not placed first, however. The classification was extended to six through the empirical work, and placing it first would put findings not yet obtained ahead of their derivation. **What was brought forward is only the stages to be checked; the classification of causes for failing to pass them was left in Chapter 13** (Remark 11.1).

The distinction between comparison and testing was also stated as a discipline (Section 11.4). This work contains several consistency checks not resting on pre-registration, such as the family-3 comparison of Chapter 8, but no passage stated their standing; each chapter merely noted individually that “this is a retrospective comparison and does not establish causation”. The discipline is needed because a comparison requires no declaration of a population and so can be carried out while skipping stages 2 and 3 (Remark 11.2).

E.24 Implication Z tested

Up to stage 22, implication Z was left untested as “the sample is already retrieved; what is missing is only the strength of the refund scheme”. It was tested at stage 23 and **rejected**.

E.24.1 Implications $\Phi 1$ – $\Phi 4$: three different failures

Registered at stage 16 and tested at stage 24. Of the four, only $\Phi 2$ is supported, and the forms of failure all differ.

$\Phi 1$: an error in the design of the estimate. At registration it was thought that keeping the identifier as the unit of observation and using standard errors robust to clustering by implementing address would handle the concentration of replicates. This is wrong. Robust standard errors correct for within-cluster correlation but not for bias in the point estimate when one cluster accounts for 95% of the exposure. Only on reporting an estimate excluding the dominant cluster did it emerge that the sign of the coefficient reverses and becomes insignificant (Remark 18.2).

$\Phi 3$: an error in the choice of control. “The distribution of CCC is skewed non-negative” was designed as a contrast with the corporate sector, but the corporate sector’s CCC is also non-negative throughout at the level of industry aggregates. Negative CCC appears in individual firms and cancels out in aggregation. The implication and its alternative say the same thing, and observation cannot tell them apart. This is ill-posedness of the same form as implication B in Section 15.3.5 (Remark 18.3).

Φ4: mistaking the implication. Proposition 10.6 states that a Φ requiring identification of the beneficiary cannot be implemented. From this “only 7-4 survives among family 7” was derived, but the proposition speaks of implementability, not frequency of appearance. The quantitative part — that family 7 is rare — was right; the part about the surviving type was wrong (Remark 18.4).

What the three have in common is that each was noticeable at the time of registration. All are problems of design that could have been examined before measurement; the data were not lacking.

E.24.2 Dropping the symbol for the production function

Up to stage 26 the technology side was called a production function f . But **no functional form was ever given**, and the notation list recorded that “the domain is not stated”. There were five mentions in the body, none of which used f itself; it appeared only through C_{prod} .

At stage 27, f was dropped and the technology side unified into C_{prod} (the correspondence from delivery to cost). The definition of production surplus (Definition 6.2) already referred to $c(\delta) = dC_{\text{prod}}/d\delta$ and holds without f .

A symbol with a name and no substance had been left in place.

E.24.3 Defining the relation between classes of agents and individuals

The primitives merely listed customers C , suppliers S and so on as “the breakdown of agents”, without giving the relation between an individual customer and the κ_C denoting customers as a whole. κ_C occurs frequently in the body, and in Example 2.6 the writing wavered between classes and individuals.

It was added to the definitions that a class is a subset of N , that an individual agent is written as $i \in C$ and as an element of N , and that a class as a subscript denotes the sum $\kappa_C = \sum_{i \in C} \kappa_i$. The index j is used for identifiers in Chapter 10 and is therefore not used for individual agents.

E.24.4 Revising the verdict on implication D

Implication D was registered at stage 13 as “the fewer the dependencies of a configuration, the smaller λ_{dep} and the longer the period of unattended operation”. Implication $\Phi 1$, registered at stage 16, reads “the fewer the other procedures a procedure calls, the longer the period before it is repaired” — **the same relation**. The duplication went unnoticed at registration.

Up to stage 24 the table of verdicts recorded D as “partial support”. But no section of the body supported that verdict, and Chapter 17 recorded that “implication D has not been tested”. **The table and the body contradicted each other.**

The test of $\Phi 1$ at stage 24 is also a test of D. The verdict was changed to **rejected**.

Registering the same relation twice is itself a failure in operating the pre-registration. The cause was keeping the register by chapter and never reconciling it as a whole.

How strength was to be measured had not been laid down at registration. Deciding after reading the documents would be fitting after the fact, so it was fixed as equation (15.2) before they were read.

Extraction by regular expression was abandoned. Checking 8 documents produced errors in 3, and the errors were biased towards overstating strength. Not being random error, they do not vanish on averaging. All 907 were read by a person, and the 132 carrying no text layer were put through OCR. None was illegible.

The reading brought several facts about publication practice to light. Some documents paste the rate table as an image alone, so the body is legible while the content is not. Some offices have a refund-field link pointing to a fee schedule. Some display the scheme as “yes” while filing a rate of 0%. **Filing a document and having its content legible are different things.**

E.25 Giving deductive grounds for choosing families 2 and 5

Section 11.5 said only that “families 2 and 5 contain the theoretical issues most clearly” and **did not show why those two**. What could be derived from the framework was left underived.

It was rewritten in four steps. Since ϕ_{prod} in equation (6.1) does not depend on Φ , the shape of Φ bears only on the remaining two terms. ϕ_{cog} requires degrees of freedom (5) and (6) and ϕ_{barg} requires (4), and the order of family assignment maps these onto families 2 and 5 respectively. The two families sit on the diagonal of Figure 2.1, spanned by the two dominant axes.

The reason for *placing them side by side* was also set out first, as a difference of institutional grounding. It does not anticipate the conclusion of Part IV (completeness and richness are independent) and is confined to the design-level reason that the two need separating.

E.26 Removing references to places that do not exist

Part IV contained “at the end of Part II it was expected that family 5 had the better conditions”, but **no such statement exists in Part II or in Section 11.5**. A sentence from an earlier version had been deleted, leaving only the sentence that answered it. The conclusion was kept and only the reference to the premise dropped.

Also, the conditions for divergence were changed from a proposition to a remark at stage 2, but **two places still referred to it as a proposition**. Both were corrected.

E.27 Resolving a contradiction in the description of ϕ_{cog}

Immediately before Section 14.6.1 it was stated that “the only corresponding measurement is the recovery figure for family 2”, while the section went on to say that “prior work contains direct instances of measurement”. **Of the same quantity it was written both that no measurement exists and that measurements exist.**

Correctly, it is measurable and measurements exist; what this work obtained on its own is only an upper bound. The family-2 recovery figure was folded in as one of the instances, and the passage recast to list four.

E.28 Extending the scope of inventories to 62 industries

Stage 14 measured only ten industries and eleven size strata, and **recorded that “the elasticity improves on correcting the definition”**. Extending to 62 industries and all sizes reverses the direction: the denominator effect rises slightly from 0.160 to 0.172.

The two are not error. Variation between size strata and variation between industries identify different things. The estimate of that chapter is identified by variation between industries and years, so it is the 62-industry figure that applies (Remark 16.8).

The range covered by the sample had been mistaken for the range reached by the conclusion. There was never any guarantee that a direction obtained on ten industries could be applied to a 45-industry estimate. All three verdicts remain “supported”, but each is weaker than at stage 14.

E.29 Stating the handling of per-firm data

The Jinzai Service General Site publishes the filings of individual offices rather than aggregates. These were included in the repository at stage 19, but **using them as input to analysis and redistributing the collection are different things**.

The policy was written into Appendix F and the per-firm data removed from the repository. What remains is only the list of licence and filing-receipt numbers, containing none of the filed values. Re-retrieving with that as a seed reproduces the realized fee rates, turnover rates and placement counts. What had already been published was removed from the history.

E.30 Multiple definitions of a symbol

At stage 28 the whole text was swept and six places where the same letter carried different meanings were fixed.

Letter	Meaning retained	Side changed
c	marginal cost $c(\delta)$	flat fee $\rightarrow F$, living expenses $\rightarrow c_{\text{liv}}$
m	gross or operating margin	number of members $\rightarrow n$
r	sales velocity	risk aversion $\rightarrow \gamma$
τ	settlement lag τ_i	time of insolvency $\rightarrow T_{\text{ins}}$
A	fixed assets	the set of responses $\rightarrow \mathcal{A}$, agents $\rightarrow i, j$
a	response; age (APC)	allocation to commissioned work $\rightarrow \text{Cap}_{\text{cl}}$

TableE.10: Multiple definitions of symbols resolved at stage 28.

c had four meanings. Only the cohort c of age–period–cohort was retained, following the notation of that literature, with a note to that effect in the body.

The three forms of degree of freedom (2) were also gathered into equation (2.10), $\pi = F + p\delta$: $p = 0$ is flat, $F = 0$ is usage-based, and both non-zero is a two-part tariff. What had stood as three forms became one family.

E.31 Correcting a count

The catalogue of Part II was described consistently from stage 1 as “27 types”, but the counts by family are 4, 5, 3, 3, 5, 4, 4, totalling 28. Eleven places in the body were corrected at stage 9. The composition of the families and the content of each type are unchanged.

E.31.1 Calling the types a “basis”

Up to stage 27 it was recorded of the catalogue of Chapter 7 that “the fact that a Φ cannot be reduced to a single type means the types function as a basis”. But no addition was defined on the space of Φ , so the word had no meaning.

At stage 28 the additivity of κ was laid down as Proposition 3.9 and, at the same time, the non-uniqueness of the decomposition shown (Remark 3.10). Without uniqueness the types are a generating set, not a basis. The word in the body was changed.

E.32 Considered but not adopted in the body

In writing Chapter 10 at stage 9, a structure presenting it as a novel contribution was initially considered. On checking the literature it emerged that the restriction of funding by a collateral constraint, the exclusion of agents holding no assets, and the divergence of identifier from agent all have existing counterparts, so it was recast to claim no novelty and a table of correspondences added to Appendix C (Remark 10.13).

The chapter was retained because, in the position of showing how the theory of Part I functions under concrete conditions, an example varying the $\mathcal{G}$ of Chapter 4 was missing. Its role is the same as that of Chapter 8 and Chapter 9.

E.32.1 Describing an allocation that is not an interior solution as one

Up to stage 27 it was recorded of equation (8.14) in Chapter 8 that “an interior solution exists”. The equation is linear in the allocation Cap_{cl} and, under $0 < \theta < \gamma$, strictly decreasing, so taken alone it is maximized at the corner $\text{Cap}_{\text{cl}} = 0$. The error arose from putting the cash that commissioned work brings in neither into the objective nor into the constraints.

At stage 28 the cash the fundamental inequality demands was made explicit as a constraint, and the optimal allocation re-derived as the edge of the constraint, $\text{Cap}_{\text{cl}}^* = c_{\text{liv}}/w$ (Proposition 8.10). The direction of the conclusion —that the criterion for selecting engagements moves from the fee to θ — is unchanged.

E.33 Changes of structure

Apart from corrections of content, the structure was changed twice.

- (1) From the chronological order of the examination to the logical order of theory, catalogue, method, empirics and conclusion.
- (2) The deductions under constraints (Chapters 8 and 9) moved from a separate part on “applications” into Part II. Applications are justified only after a theory has passed through empirical work, and the empirical work here is not complete.

The second change was accompanied by fixing 64 inconsistent references (parts confused with chapters) and two logical inversions (a later chapter referred to in the past tense).

AppendixF

Data sources

Subject	Statistic or source	Where used
Family 2	Japan Payment Service Association, 27th Survey of Issuing Businesses [27]	Part IV
Family 2	Payment Services Act; FSA register of issuers of prepaid payment instruments	Part IV
Family 5	MHLW, Tabulation of Employment Placement Business Reports, FY2024 [28]	Part IV
Family 5	Jinzai Service General Site: office details and refund-scheme filings	Section 15.3.5
Fees	published fee schedules of each platform	Chapter 17
Industry classification	MOF, Financial Statements Statistics of Corporations by Industry, time series [29]	Chapter 16
Age at founding	White Paper on Small and Medium Enterprises 2014; Teikoku Databank survey of newly established corporations	Section 8.6
Comparison with sole proprietorships	SME Agency, Basic Survey of Small and Medium Enterprises, FY2020–2025 final reports	Section 14.5.5
Regulation of payment dates	Act on Ensuring Proper Transactions Involving Specified Entrusted Business Operators (in force November 2024)	Section 15.7
Vulnerabilities	published CVE counts, CISA KEV catalogue	Chapter 9

TableF.1: Data sources used in the empirical work.

■On reproduction. The data obtained and the processing are placed under `data/`: `raw/` holds the inputs as obtained and `scripts/` the analysis. `derived/` is intermediate data, regenerated by `scripts/parse.py`.

The calculations of Chapter 16 use `ccc.py`, `e1.py`, `elast.py` and others; the tests of Section 15.3.5 use `parse_jinzai.py` and `predY.py`. The correspondence is listed in `data/README.md`.

■On handling. All statistics used are public. The Jinzai Service General Site alone, however, publishes **the filings of individual offices** rather than aggregates, and this work uses them as input to statistical analysis.

What this work presents is the result of the analysis, not the collection of per-firm data that was gathered. That collection is not reproduced, published or redistributed. The repository contains the processing of the analysis and the aggregates obtained from it. Per-firm data retrieved from that site — the search-result listings, the office detail pages and the refund-scheme filings alike — are not included.

What remains is only **the list of licence and filing-receipt numbers** (`data/raw/jinzai/permits.tsv`), which contains none of the filed values. Using it as a seed, running `data/scripts/fetch_jinzai_detail.py` reproduces the realized fee rates, turnover rates and placement counts. Retrieval was carried out at six-second intervals per record and spread across separate days.

References

- [1] G. A. Akerlof and J. L. Yellen, Do deferred wages dominate involuntary unemployment as a worker discipline device?, NBER Working Paper No. 2025, 1986. nber.org/papers/w2025
- [2] J. Berkson, Limitations of the application of fourfold table analysis to hospital data, *Biometrics Bulletin*, 2(3):47–53, 1946. doi:10.2307/3002000
- [3] L. L. Carstensen, D. M. Isaacowitz, and S. T. Charles, Taking time seriously: A theory of socioemotional selectivity, *American Psychologist*, 54(3):165–181, 1999. doi:10.1037/0003-066x.54.3.165
- [4] N. C. Churchill and J. W. Mullins, How fast can your company afford to grow?, *Harvard Business Review*, 79(5):135–143, 2001. hbr.org/2001/05/how-fast-can-your-company-afford-to-grow
- [5] R. H. Coase, The nature of the firm, *Economica*, 4(16):386–405, 1937. doi:10.1111/j.1468-0335.1937.tb00002.x
- [6] S. DellaVigna and U. Malmendier, Paying not to go to the gym, *American Economic Review*, 96(3):694–719, 2006. doi:10.1257/aer.96.3.694
- [7] A. Daripa and J. Nilsen, Ensuring sales: A theory of inter-firm credit, *American Economic Journal: Microeconomics*, 3(1):245–279, 2011. doi:10.1257/mic.3.1.245
- [8] A. K. Dixit and R. S. Pindyck, *Investment under Uncertainty*, Princeton University Press, 1994. doi:10.2307/j.ctt7sncv
- [9] L. Einav, B. Klopock, and N. Mahoney, Selling subscriptions, *American Economic Review*, 115(5):1650–1671, 2025. doi:10.1257/aer.20231612
- [10] S. J. Grossman and O. D. Hart, The costs and benefits of ownership: A theory of vertical and lateral integration, *Journal of Political Economy*, 94(4):691–719, 1986. doi:10.1086/261404
- [11] M. A. Hernán, S. Hernández-Díaz, and J. M. Robins, A structural approach to selection bias, *Epidemiology*, 15(5):615–625, 2004. doi:10.1097/01.ede.0000135174.63482.43
- [12] S. Mateut, Reverse trade credit or default risk? Explaining the use of prepayments by firms, *Journal of Corporate Finance*, 29:303–326, 2014. doi:10.1016/j.jcorpfin.2014.09.009
- [13] Pay now, buy never: The economics of consumer prepayment schemes, NBER Working Paper No. 34918, 2026. nber.org/papers/w34918
- [14] M. A. Petersen and R. G. Rajan, Trade credit: Theories and evidence, *Review of Financial Studies*, 10(3):661–691, 1997. doi:10.1093/rfs/10.3.661
- [15] R. A. Schwartz, An economic model of trade credit, *Journal of Financial and Quantitative*

- Analysis*, 9(4):643–657, 1974. doi:10.2307/2329765
- [16] M. Spence, Job market signaling, *Quarterly Journal of Economics*, 87(3):355–374, 1973. doi:10.2307/1882010
- [17] B. Holmström, Moral hazard and observability, *Bell Journal of Economics*, 10(1):74–91, 1979. doi:10.2307/3003320
- [18] B. Holmström and P. Milgrom, Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design, *Journal of Law, Economics, and Organization*, 7:24–52, 1991. doi:10.1093/jleo/7.special_issue.24
- [19] R. Ibragimov, D. Jaffee, and J. Walden, Diversification disasters, *Journal of Financial Economics*, 99(2):333–348, 2011. doi:10.1016/j.jfineco.2010.08.015
- [20] F. H. Knight, *Risk, Uncertainty and Profit*, Houghton Mifflin, Boston, 1921. archive.org/details/riskuncertainty00knig
- [21] D. Laibson, Golden eggs and hyperbolic discounting, *Quarterly Journal of Economics*, 112(2):443–477, 1997. doi:10.1162/003355397555253
- [22] W. M. Mason and S. E. Fienberg (eds.), *Cohort Analysis in Social Research: Beyond the Identification Problem*, Springer, New York, 1985. doi:10.1007/978-1-4613-8536-3
- [23] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed., Cambridge University Press, 2009. doi:10.1017/CB09780511803161
- [24] C. Prendergast, The provision of incentives in firms, *Journal of Economic Literature*, 37(1):7–63, 1999. doi:10.1257/jel.37.1.7
- [25] C. Prendergast, The tenuous trade-off between risk and incentives, *Journal of Political Economy*, 110(5):1071–1102, 2002. doi:10.1086/341874
- [26] J.-C. Rochet and J. Tirole, Platform competition in two-sided markets, *Journal of the European Economic Association*, 1(4):990–1029, 2003. doi:10.1162/154247603322493212
- [27] Japan Payment Service Association, The 27th Survey of the Actual State of Issuing Businesses (FY2024 edition), 2025.
- [28] Ministry of Health, Labour and Welfare, Tabulation of Employment Placement Business Reports, FY2024, 2025.
- [29] Policy Research Institute, Ministry of Finance, Financial Statements Statistics of Corporations by Industry, time series, e-Stat table ID 0003060791.
- [30] W. Wagner, Diversification at financial institutions and systemic crises, *Journal of Financial Intermediation*, 19(3):373–386, 2010. doi:10.1016/j.jfi.2009.07.002
- [31] L. W. Cong and Z. He, Blockchain disruption and smart contracts, *The Review of Financial Studies*, 32(5):1754–1797, 2019. doi:10.1093/rfs/hhz007
- [32] J. R. Douceur, The Sybil attack, in *Peer-to-Peer Systems*, Lecture Notes in Computer Science 2429, pages 251–260, 2002. doi:10.1007/3-540-45748-8_24
- [33] N. Kiyotaki and J. Moore, Credit cycles, *Journal of Political Economy*, 105(2):211–248, 1997. doi:10.1086/262072
- [34] OECD, The limits of DeFi for financial inclusion, OECD Business and Finance Policy Papers,

2024. doi:10.1787/f00a0c7f-en

- [35] A. Pasdar, Y. C. Lee and Z. Dong, Connect API with blockchain: a survey on blockchain oracle implementation, *ACM Computing Surveys*, 55(10):1–39, 2023. doi:10.1145/3567582
- [36] A. A. Rampini and S. Viswanathan, Collateral, risk management, and the distribution of debt capacity, *The Journal of Finance*, 65(6):2293–2322, 2010. doi:10.1111/j.1540-6261.2010.01616.x
- [37] A. A. Rampini and S. Viswanathan, Collateral and capital structure, *Journal of Financial Economics*, 109(2):466–492, 2013. doi:10.1016/j.jfineco.2013.03.002
- [38] M. Yokoo, Y. Sakurai and S. Matsubara, The effect of false-name bids in combinatorial auctions: new fraud in Internet auctions, *Games and Economic Behavior*, 46(1):174–188, 2004. doi:10.1016/S0899-8256(03)00045-9