

ECSTATIC ECONOMICS · VOLUME III

The Economy That Keeps Becoming

Complexity Economics, Read as a Living System, for Leaders and Households

Ammanuel Santa Anna

Santa Anna, A. (2026). The Economy That Keeps Becoming: Complexity Economics, Read as a Living System, for Leaders and Households. Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989451
Luminous Library Licence 1.0 · business.hautelumiere.life/library/the-economy-that-keeps-becoming

THE LICENCE

The Economy That Keeps Becoming
Complexity Economics, Read as a Living System, for Leaders and Households

Ecstatic Economics · Volume III. The Haute Lumière edition, 2026. © 2026 Ammanuel Santa Anna.

Released under the Luminous Library Licence 1.0 (hautelumiere.life/library-licence): free to read, share, quote, translate and teach, non-commercially, for ever. The Author keeps every other right, including every commercial right.

Published by Luminous Prosperity Inc. for Haute Lumière. Author ORCID 0009-0008-3244-0306.

DOI 10.5281/zenodo.23989451 · this record's page: business.hautelumiere.life/library/the-economy-that-keeps-becoming/

Please cite as: Santa Anna, A. (2026). The Economy That Keeps Becoming: Complexity Economics, Read as a Living System, for Leaders and Households. Luminous Prosperity Inc. https://doi.org/10.5281/zenodo.23989451

A NOTE BEFORE YOU READ

This work is offered for education and leadership development. It is not legal, financial, accounting, investment, medical or psychological advice; take the advice of a qualified professional who knows your circumstances. Companies and events are described from the public record as the author understands it; characterisations are the author's analysis. Examples are illustrative, and no particular outcome is promised to any reader.

ABSTRACT

An economy is not a machine that settles but a system that keeps forming itself — and a leader can work with its loops instead of pulling at levers. Complexity economics, read as a living system, for leaders and households.

CONTENTS

CHAPTER ONE Loops, Not Levers 4

CHAPTER TWO The Economy That Keeps Becoming 17

CHAPTER THREE Where the Model Runs Thin 32

CHAPTER FOUR Nobody Designed the Trap 46

CHAPTER FIVE The Valley and the Highway 62

CHAPTER SIX The Week After the Fire 76

CHAPTER SEVEN A Firm Read as Weather 89

CHAPTER EIGHT The Field Manual 103

CHAPTER NINE Knowing Whether It Worked 123

CHAPTER TEN The Hard Conversations 136

CHAPTER ELEVEN The Household as a Living System 149

CHAPTER TWELVE The Path Across the Lawn 164

CHAPTER THIRTEEN What Grows by Itself 178

WORKBOOK Every Tool, and a Ninety-Day Plan 187

NOTES References 227

CHAPTER ONE

Loops, Not Levers



Loops, Not Levers

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ONE

A desk, and a question

Somewhere this week a leader will sit at a desk with a decision in front of her and a plan that assumes the world will hold still while she carries it out. The plan will have a target, a budget, a timeline and a set of levers: raise this price, cut that cost, hire into this team, close that office. Each lever is expected to move one thing, by roughly the amount the spreadsheet says, and to leave the rest of the organisation where it was.

Most plans are written this way because most of the economics that leaders were taught was written this way. It pictured a market as a machine that settles. Push it, and it returns to balance at a new point. Change an input, and the output moves in proportion. Find the optimum, and stay there.

This book is about a different picture, one that economists have been building carefully for about forty years and that most organisations have not yet been given in a form they can use. In this picture an economy is not a machine but a system that keeps forming itself out of the choices of the people inside it. Those people watch one another, learn, and adjust. Their adjustments change the conditions everyone else is adjusting to. Some changes feed on themselves and grow; others are corrected and fade. Small events early on can decide where a whole market ends up, and arrangements nobody would choose today can hold firm for decades because everything else has grown up around them.

The school of thought that studies economies this way is called complexity economics. Its central figure is the economist W. Brian Arthur, whose work at Stanford and at the Santa Fe Institute in the 1980s and 1990s gave the field its vocabulary: increasing returns, path dependence, lock-in, the economy as an evolving complex system. Chapter 2 tells that story from its beginning.

The question the book asks of it is a practical one. If an economy really does behave like this, what should a leader do differently on Monday? And a household?

The short answer, which the rest of the book earns, is this: look for loops, not levers. A lever moves one thing. A loop is a chain of cause and effect that comes back to where it started, so that a change at one point returns, amplified or corrected, a little while later. Organisations are made of loops. Leaders who see them make different decisions from leaders who see only levers, and the difference compounds.

Why this, and why now

Three things make this a good moment for the idea to leave the seminar room.

The first is that the evidence has matured. When Arthur's paper on competing technologies was published in *The Economic Journal* in 1989, it was a mathematical argument about what could happen under increasing returns. Since then the argument has been tested, disputed and refined. Historians have argued over whether the typewriter keyboard is a real case of inefficient lock-in. Economic geographers have shown

how industries cluster. Central bankers, after 2008, borrowed the methods of ecologists to study how banks fail together. The critics have done their work too, and some of them are right. Chapter 3 takes their objections one at a time. What survives is narrower than the early enthusiasm and much more useful, because it can be acted on without exaggeration.

The second is that the ordinary experience of running something now looks more like complexity than like equilibrium. Markets with strong network effects, where a product becomes more valuable the more people use it, are no longer a curiosity at the edge of the economy. Supply chains are long and coupled, so that a fire in one plant can stop production in another country within days. Reputations spread and collapse through networks rather than through advertising budgets. A leader who manages with levers in this environment keeps being surprised, and the surprises are not random. They have a shape, and the shape can be learned.

The third reason is the one closest to this house. Ammanuel Santa Anna's book *Ecstatic Economics* argues that an economy is a living system, and that its health depends on whether value circulates, whether the people who act feel the consequences of what they do, and whether the arrangements people live within carry vitality or extract it. That is a strong claim. It deserves strong support, and complexity economics is one of the most rigorous bodies of work that supplies it. The two traditions use different words. Arthur writes about positive feedback and nonergodic processes; Santa Anna writes about circulation, severed feedback loops and the felt quality of economic life. They are describing the same kind of system. Where the science gives the mechanism, the living-systems reading asks what the mechanism carries.

Santa Anna puts the circulatory principle in one line, near the end of his chapter on the economic traditions that came before markets:

> *wealth that does not circulate is not wealth. Prosperity that does not connect is not prosperity. The gift must always move.*[1]

A complexity economist would put it differently, and less warmly, but the claim underneath is close: the properties that matter in an economy are properties of flow and connection, not only of stock. A firm with large reserves and severed feedback is less healthy than its balance sheet suggests. A household with modest savings and good loops may be stronger than one with more money and none.



Who this book is for

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ONE

Who this book is for

It is written for three readers.

The leader of an organisation of almost any size, from a founder with thirty people to the executive team of a large firm, who suspects that the plan on the desk is missing something about how the organisation actually behaves. The book gives this reader a set of instruments, each written out in full at the end of its chapter, and a field manual (chapter 8) for running them in sequence over twenty-six weeks.

The person responsible for a household's money and work, who is entitled to the same instruments at a smaller scale. A household is a complex adaptive system too: its savings, skills, habits and commitments reinforce or correct one another over years. Chapter 11 is written for this reader, and most chapters close with a note on how their tool works at a kitchen table.

The curious reader who wants to understand what complexity economics actually claims, where it came from, what its critics say, and how it connects to the wider questions of the Ecstatic Economics series: Elinor Ostrom's work on the commons, the ecological limits drawn by Kate Raworth's doughnut, the evidence on gainsharing and employee ownership, and Karl Polanyi's account of markets embedded in society. Chapter 12 makes those joins.

It is not written for the reader who wants a forecast. Complexity economics is modest about prediction, and this book follows it. Nothing here is investment advice, and nothing here promises what any particular organisation or household will find when it looks at its own loops. What the book does offer is a way of looking that tends to find things worth finding.



How the book is built

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ONE

How the book is built

The book moves in four parts.

The idea (chapters 2 to 4). Chapter 2 tells where complexity economics truly began: not in Santa Fe but in Alfred Marshall's *Principles* of 1890, which already worried about increasing returns, and then in Arthur's long-delayed paper and the institute that gave it a room. It sets out the four properties of lock-in and corrects two myths. Chapter 3 states five fair critiques as their best advocates would, and offers for each a way

through. Chapter 4 reads the whole through the house's living-systems lens: emergence without a villain, the severed feedback loop, and the banking system read as an ecosystem.

The cases (chapters 5 to 7). Two public cases, cited from their sources, and one composite labelled as such. Chapter 5 compares the two great American electronics regions of the late twentieth century, Silicon Valley and Boston's Route 128, through AnnaLee Saxenian's study of why one adapted and the other locked in. Chapter 6 follows the Toyota group in the days after a fire destroyed the only plant making a small brake part in February 1997, and the self-organising recovery that followed. Chapter 7 works the instruments through on Larkfield, a composite software firm, from a Tuesday in February to a Tuesday in May.

The practice (chapters 8 to 11). Chapter 8 is the field manual: the full sequence, week by week, with roles, decisions and the mistakes to catch early. Chapter 9 is about measurement: how to know whether any of it worked, what to count, and how to read the numbers honestly in a system where every measure changes the behaviour it measures. Chapter 10 is about the hard conversations, with boards, partners, unions, families and sceptics, and gives scripts. Chapter 11 brings the whole thing home, to one person's money and work.

The whole (chapters 12 to 13, and the workbook). Chapter 12 joins complexity economics to the rest of the Ecstatic Economics series and says what to do on Monday. Chapter 13 gives the living-systems reading of everything before it, in one image. The workbook at the back gathers every tool in the book and sets out a ninety-day plan.

The instruments

Every chapter ends with a tool, written out in full so that it can be used without any software. Taken together they form one practice:

- The First Map of Loops (chapter 2) names the reinforcing loops, lock-ins and balancing loops in a system.
- The Lock-In Test (chapter 3) prices any arrangement suspected of holding the system in place.
- The Loop Audit (chapter 4) asks what each loop carries, and whether the people acting in it ever see the consequence.
- The Cluster Reading (chapter 5) reads a region, an industry or a network of firms for whether its loops are open or closed.
- The Shock Rehearsal (chapter 6) asks how the system would reorganise if one critical node failed tomorrow.
- The Loop Ledger (chapter 7) gathers the first three into a twelve-week cycle.
- The Twenty-Six Week Plan (chapter 8) is the field manual on one page.
- The Honest Dashboard (chapter 9) chooses a small set of measures that the system cannot easily game.

- The Conversation Cards (chapter 10) prepare the hard conversations.
- The Kitchen-Table Map (chapter 11) runs the whole sequence for a household.
- The Monday Sheet (chapter 12) is a fifteen-minute weekly practice.
- The Desire-Line Walk (chapter 13) is a single habit to carry when everything else has been put away.

No tool needs the others, and each is more useful with them. A reader with one free hour should begin with the First Map of Loops. A reader with a quarter should begin with chapter 8.



Three habits before beginning

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ONE

Three habits before beginning

The book asks three habits of its reader, and it is fair to name them now.

Write predictions before acting. Complexity makes hindsight unusually persuasive. Almost any outcome can be explained after the event by a story about loops, which is exactly the weakness critics of the field point to. The only defence is to write down, before a decision, what you expect to happen and what would prove you wrong. Every probe in this book has a prediction, and most have a stop condition. Chapter 9 explains why this matters more than any measure.

Treat every claim of lock-in as a hypothesis with a price on it. Not everything that persists is a trap. Many arrangements persist because they are adequate and switching would cost more than it returned. Chapter 3 gives the test; the habit is to apply it before calling anything a lock-in.

Put the consequence in front of the person who acts. This is the habit that joins complexity economics to the living-systems reading, and it is the one most often skipped. A feedback loop restored to a manager's dashboard is not restored to the people whose daily choices drive it. Santa Anna's book names the structural shift directly: when people feel the impact of their economic choices, it says,

> *When you feel the impact of your economic choices — in your body, not just your mind — you make different choices.*[2]

The sentence is about individuals and their spending, but it holds for organisations as well. A team that hears, unedited, what its work did for the people it served behaves differently from a team that reads a quarterly summary of the same facts.

What the book will not do

It will not reduce complexity to a formula. Some writing about complex systems offers a model that promises to compute the right decision; this book does not, and it is suspicious of any that does. The instruments here are ways of seeing and of testing, not engines for producing answers.

It will not treat complexity as a reason to do nothing. A theory that leaves its readers less willing to act has failed them. The response to complexity is small, deliberate, well-recorded action, repeated, with the results put in front of the people who can learn from them.

And it will not claim more than the evidence allows. Where a source is a model, it says so. Where a case is a composite, it says so in the same paragraph. Where a figure is illustrative, it says whose figure it is. The notes at the end of each chapter say what each source cannot show.



An image to begin with

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ONE

An image to begin with

Picture a lawn on a university campus, crossed at its edges by paved walks set at right angles, and across the middle by a thin line of worn earth that nobody laid. The first walker cut across because it was shorter. The second followed the faint trace the first had left. By the hundredth, the grass had given up.

That line is increasing returns, path dependence and emergence, written on a lawn. It is also an instruction. It tells whoever is responsible for the lawn where people actually go. The book ends on the same lawn, in chapter 13, and by then the reader will have the instruments to see the lines in their own organisation and household, to ask what each one carries, and to decide which to pave.

The Tool: The One-Hour Start

For the reader who wants to begin before reading further. One sheet, one hour, pencil.

1 · The system

The system I am responsible for, in one line (a firm, a team, a product, a household):

The decision currently on my desk:

2 · The levers I was planning to pull List the levers in the current plan. For each, write what you expect it to move.

LEVER	WHAT IT IS EXPECTED TO MOVE	BY HOW MUCH, AND ON WHAT EVIDENCE

3 · The loops each lever enters For each lever, ask: what will the people affected by it do in response, and how will that come back to me?

Lever 1 enters this loop:

Lever 2 enters this loop:

Lever 3 enters this loop:

4 · One thing that is already growing by itself

A practice, relationship or habit in the system that seems to be reinforcing itself:

Who knows most about it:

5 · One prediction, written now

What I expect to happen in the next ninety days if I carry out the plan unchanged:

What would show me I was wrong:

The date I will read this sheet again:

Keep the sheet. Chapter 2 turns it into the First Map of Loops.

A note on method

The dates and sources for Arthur's work, the Santa Fe Institute, Saxenian's study and the Aisin fire are given in the chapters that discuss them. Nothing in this book is investment advice.

CHAPTER TWO

The Economy That Keeps Becoming



The Economy That Keeps Becoming

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWO

A paper in a drawer

In 1983 an economist at Stanford finished a paper about why one technology wins over another. W. Brian Arthur was born in Belfast, trained in electrical engineering and operations research, and came to economics with an engineer's habit of asking what a system actually does rather than what it ought to do. The paper argued something that sounds harmless today. When a technology becomes more attractive the more people adopt it, small and accidental early events can decide which technology the whole market ends up using, and the winner need not be the better one.

By Arthur's own later account, the leading American journals declined it for years. It appeared in *The Economic Journal* in March 1989, under the title "Competing Technologies, Increasing Returns, and Lock-In by Historical Events".[3] The six-year wait is his story, told in the preface to his 1994 collection, and it is worth holding loosely; journals reject papers for many reasons. What is not in doubt is the paper itself, and what it asked economists to give up.

Most of the economics taught in the twentieth century rested on a quiet assumption: that markets tend toward an equilibrium, a single resting point that the forces of supply and demand will find, whatever the starting conditions. Arthur's paper described markets with more than one possible resting point, and with history choosing among them.

Ammanuel Santa Anna, the author of *Ecstatic Economics*, makes the same refusal from the side of ecology. Writing about what a regenerative enterprise might learn from a forest, he sets the old picture aside in two sentences:

> *A healthy ecosystem is not in equilibrium. It is in generative surplus*[4]

That sentence is the subject of this book. Complexity economics is the school of thought that examines it carefully, with models and evidence, and it is one of the strongest scientific supports the house has for reading an economy as a living thing. This chapter is about where the idea came from, what it changed, and the two myths that have grown around it. The chapters that follow take the fair critique (chapter 3), the living-systems reading (chapter 4), two public cases (chapters 5 and 6), a worked composite (chapter 7), and then the practice.

The claim of the book can be stated in one sentence. An economy is a system that keeps forming itself out of the choices of the people inside it, and a leader or a household that reads it that way will look for loops, not levers.



Older than Santa Fe

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWO

Older than Santa Fe

The first myth is that complexity economics began in New Mexico in the late 1980s. It did not. Its central idea is older than any institute.

Alfred Marshall's *Principles of Economics*, first published in 1890, already discussed what he called the law of increasing return: the tendency, in many industries, for costs per unit to fall as output grows, because experience accumulates, suppliers cluster and skills concentrate in a district. Marshall saw the difficulty it made for his own system. If a firm's advantage grows with its size, there is no obvious reason for any rival to survive, and the tidy picture of many small competitors finding a common price begins to wobble. He kept increasing returns at the edge of his theory, as a fact he respected and could not fully fit.

For most of the following century the profession did roughly the same. Diminishing returns, where each extra unit of input yields a little less, produce stable equilibria and solvable equations. Increasing returns

produce several possible outcomes and an awkward dependence on history. The mathematics preferred the first, and so, for a long time, did the textbooks.

What changed in the 1980s was not the discovery of increasing returns. It was the arrival of tools that could take them seriously. Arthur and his collaborators used probability theory, in the form of what mathematicians call generalised urn processes, to show exactly how a market with self-reinforcing adoption behaves. Computers made it possible to run economies of many interacting agents and watch what happened, rather than solving for a single answer in advance.

The Santa Fe Institute, founded in 1984 as a home for the study of complex systems across disciplines, gave the work a room. In 1987 it convened physicists and economists together, with Kenneth Arrow and the physicist Philip Anderson among the organisers, and the following year Arthur directed its first economics programme. The proceedings were published as *The Economy as an Evolving Complex System*. The title is the whole argument, set in type.

Santa Anna, writing about economic thought in *Ecstatic Economics*, makes the same move from the other direction. He reads the history of economic paradigms not as a series of errors corrected but as a sequence still under way:

> *The neoliberal paradigm is not the end of economic evolution. It is a stage*[5]

The sentence is about ideologies rather than models, but it shares the complexity economist's first instinct. Nothing is final. The resting point is a moment, not a destination.

Four properties of lock-in

The 1989 paper is short and dense, and its most useful contribution for a reader who runs something is a list. Arthur showed that when adoption is self-reinforcing, the market acquires four properties that a market with diminishing returns does not have.[3]

It is not predictable in advance. Knowing the technologies and the preferences of the people who will choose among them is not enough to say which will win. The outcome depends on the order in which early adopters happen to arrive.

It can be inefficient. The technology that wins is not guaranteed to be the one that would have served everyone best. An early lead, gained by accident, can carry an inferior option to dominance.

It becomes inflexible. Once one option has a large enough lead, no reasonable subsidy or improvement to the other can reverse the outcome. The market is locked in.

It is path-dependent. Small events in the history of the market are not averaged away. They are remembered, in the structure of who uses what.

Arthur used the technical word *nonergodic* for the last property. An ergodic process forgets its past; given time, it ends up in the same place whatever happened along the way. A nonergodic one does not. That single word is perhaps the clearest line between the economics he inherited and the economics he helped to build.

The finer point for a leader is this. The four properties are not claims that every market locks in. They are a description of what follows *if* adoption reinforces itself strongly enough. Whether it does is an empirical question, asked market by market. The paper gives you the shape to look for, not the verdict.

Arthur took the argument to a wider audience in *Scientific American* in February 1990, in an essay called "Positive Feedbacks in the Economy".[6] There the engineering vocabulary is plain. Diminishing returns are negative feedback: a disturbance creates forces that push the system back. Increasing returns are positive feedback: a disturbance creates forces that push it further. An economy with both, he argued, is not a balance but a pattern of reinforcements and corrections, some of them local, some of them able to reshape an industry.



The economy that reacts to itself

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWO

The economy that reacts to itself

In 1999 Arthur wrote a short essay in *Science* called "Complexity and the Economy".[7] It is the text most often credited with putting the phrase *complexity economics* into common use, and it moves the argument from technologies to the whole of economic life.

The core of it is a change of picture. In the standard picture, agents face a world and optimise against it. In Arthur's, agents face a world that their own choices are creating, and they adjust their strategies as the world they have created changes. Prices, expectations and institutions are not given in advance. They

emerge from the interaction of many agents, each acting on a partial view, and then feed back into how those same agents act.

The most famous illustration is a puzzle Arthur set out in 1994, about a bar in Santa Fe called El Farol.[8] One hundred people decide each week whether to go to the bar. If fewer than sixty go, it is pleasant and those who went are glad. If more than sixty go, it is crowded and those who went would rather have stayed home. No one can know in advance how many will attend. If everyone uses the same forecasting rule, the rule defeats itself: if it predicts a quiet night, everyone goes and the night is crowded. So the people in Arthur's simulation each hold several rough rules of thumb, use whichever has lately worked best, and drop the ones that fail.

What emerged in his computer runs was an attendance that hovered around sixty without settling there, an ecology of forecasting rules that kept changing, and no single correct way to think. Arthur called the reasoning involved *inductive*: people form hypotheses, test them against experience, and revise, rather than deducing the one right answer from complete information. The economy, in this view, is a population of hypotheses in competition.

Santa Anna's book gives the earlier economics its due before it departs from it. Of Adam Smith's insight into decentralised coordination, he writes:

> *The intelligence in Smith's observation was real: decentralized decision-making by individuals pursuing their own interests can coordinate complex economic activity more efficiently than central planning. Markets do possess a kind of distributed intelligence. This insight remains valuable.*[9]

That is the complexity economist's position almost exactly. Arthur did not argue that markets fail to coordinate. He argued that they coordinate in a way richer than equilibrium theory can describe: continuously, with learning, and sometimes into places nobody would have chosen.

The second myth, and a correction

The second myth runs the other way. It holds that complexity economics proves markets do not work, or that it is a fashionable word for chaos, or that it replaces the rest of economics with metaphor.

Arthur's own most complete statement, "Foundations of complexity economics", published in *Nature Reviews Physics* in 2021, is careful on all three counts.[11] He presents complexity economics as a broader framework in which the economy is not necessarily in equilibrium, and in which equilibrium, where it does occur, becomes one case among several rather than the assumption every model must begin with. Standard methods remain useful where their conditions hold. What changes is that those conditions are now something to be checked rather than presumed.

Nor is complexity the same as chaos. Chaos, in the mathematical sense, describes simple deterministic systems whose behaviour is acutely sensitive to their starting point. Complexity describes systems of many adaptive parts whose interactions produce patterns at a larger scale: patterns that are often quite stable, often recognisable, and not reducible to any single part. A city's neighbourhoods, a market's price history and a language's grammar are all complex in this sense, and none of them is chaotic.

The house's own writing uses the same distinction without the jargon. In the Luminous 100, his upgraded reading of the classical management frameworks, Santa Anna describes how a network of relationships comes to have a purpose of its own, which is emergence stated as a practitioner would meet it: living systems evolve

> *through mutual adaptation between organism and environment, where purpose emerges from the quality of relationship rather than being imposed in advance.*[10]

What the research cannot show, and the house should not claim on its behalf, is that complexity economics has a better record of forecasting than the economics it criticises. It has a different aim. It is more modest about prediction and more ambitious about understanding: about why outcomes cluster where they do, why some arrangements persist long after anyone would have chosen them, and where a small intervention can matter more than a large one. The next chapter takes the limits seriously, because they are real.



Reading your own organisation as still forming

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWO

Reading your own organisation as still forming

The finer points of implementation begin here, with one change of habit.

A leader trained in the older picture asks of any decision: what is the right answer? A leader reading the economy as Arthur does asks first: what kind of system is this decision going into, and what will the system do with it? The second question is slower and more useful. It treats your organisation, your market and your household finances as things still forming, in which your choice becomes one of the conditions everyone else then adapts to.

Three practical consequences follow, and they recur through the rest of this book.

Look for what reinforces itself. In any organisation there are a handful of loops where success breeds success: a product that becomes more useful as more customers use it, a team whose reputation attracts the people who make it better, a supplier relationship that deepens with every order. These are your increasing returns. They are also where early, small choices carry the most weight. A loop of this kind deserves attention long before it shows up in the accounts.

Look for what is locked in. Every organisation also carries arrangements nobody would choose today: a file format, a pricing convention, a meeting, a system that everything else now depends on. Arthur's four properties apply to them. They are the product of history rather than design, and they are inflexible for reasons that have nothing to do with their merit. Naming them is not the same as removing them, but it changes the conversation from "this is how it is done" to "this is where we happen to have landed".

Act as an agent among agents. Your customers, rivals and staff are all forming hypotheses about what you will do, as you are about them. A strategy that works only if nobody responds to it is not a strategy. The El Farol puzzle is a good discipline here: before you commit to a forecast, ask what happens if everyone else holds it too.

The same habits apply at a kitchen table. A household's savings, its skills and its reputation in a community are all capable of reinforcing themselves, and a household's fixed costs, subscriptions and habits are all capable of locking in. A family that reviews its money once a year as a list of balances sees a snapshot. A family that asks which arrangements are growing on their own, and which ones it would not choose again, sees a system.

The instrument below is the first form of that habit. Chapter 7 works it through in full on a composite firm. For now it is enough to take an hour and draw the map.



The paper, and the drawer

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWO

The paper, and the drawer

There is one more reason the story of the 1989 paper matters to this house. Arthur's argument is itself an example of what it describes. An idea that did not fit the dominant equilibrium found no early adopters in the journals that mattered; it survived because a small, separate institution gave it a place to grow, and because the people who read it there began to use it. Today the Santa Fe Institute's economics work is cited across policy, finance and ecology. History chose, and history could have chosen otherwise.

Santa Anna's book is built on a related wager: that an economics organised around vitality rather than extraction is not a utopia waiting for permission but a set of practices that, once adopted by enough people in enough places, begin to reinforce themselves. The rest of the lineage examines that wager from several sides, from Elinor Ostrom's commons to the ecological limits of the doughnut. Complexity economics supplies the grammar of how such practices spread, and the honest warning of how they can stall.

The book from which this one grows is *Ecstatic Economics*.

In Stanford, in 1983, the paper sat in a drawer while the journals decided. Somewhere, in most organisations, there is a drawer like it: a small practice, a habit of one team, an arrangement on one account, that is quietly beginning to reinforce itself. The first task is to find the drawer.

The Tool: The First Map of Loops

One page, one hour, done alone first and then with two colleagues or the other adults in the household. Write in pencil; the map will change.

1 · The system you are mapping

Name the system in one line (a firm, a team, a product, a household):

Its edge: what is inside, and what is the environment it adapts to?

The date of this map:

2 · Reinforcing loops: what grows by itself List up to five. For each, write the loop as a sentence of the form *the more A, the more B, the more A*.

LOOP, WRITTEN AS A SENTENCE	EVIDENCE IT IS RUNNING (A FIGURE, WITH ITS SOURCE)	HOW LONG IT HAS RUN	IS IT HELPING OR HARMING?

3 · Lock-ins: what history chose List up to five arrangements that everything else now depends on.

ARRANGEMENT	WHEN AND HOW IT BEGAN	WHAT DEPENDS ON IT	WOULD YOU CHOOSE IT TODAY? YES / NO / UNSURE	COST OF SWITCHING, ROUGH

4 · Balancing loops: what corrects itself List up to three. *The more A, the less B, the less A.*

…

…

…

5 · The El Farol question

Your main forecast for the next twelve months:

If everyone you compete or cooperate with held the same forecast, what would they do?

- Would that make your forecast wrong? yes / no / unsure

6 · The drawer

One small practice that may be starting to reinforce itself, which nobody has yet named:

Who knows most about it:

The date you will look at it again:

Keep the page. Chapter 3 adds a column to it and chapter 4 a second page; chapter 7 fills it in for a worked case, and the workbook at the back gathers every page in one place.

A note on method

The Santa Fe Institute dates (founded 1984; the 1987 workshop; the 1988 economics programme) and the proceedings, Anderson, P. W., Arrow, K. J., & Pines, D. (eds.) (1988), *The Economy as an Evolving Complex System*, Addison-Wesley, are given from the institute's published histories and the proceedings' title page as commonly catalogued; they were not re-read from an archival record for this book. Marshall's discussion of increasing return is in *Principles of Economics* (Macmillan, 1890), Book IV.

CHAPTER THREE

Where the Model Runs Thin



Where the Model Runs Thin

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER THREE

The keyboard on the desk

There is a keyboard on almost every desk in the English-speaking world whose top row of letters begins Q, W, E, R, T, Y. For a generation of economists it was the most famous object in their field.

In 1985 the economic historian Paul David published a short paper in the *American Economic Review* called "Clio and the Economics of QWERTY".[12] He told the story of the typewriter keyboard: designed in the 1870s, partly to keep the bars of early machines from jamming, adopted by the first successful manufacturers, taught in typing schools, and then fixed in place because every typist trained on it and every employer bought for it. A rival layout, patented by August Dvorak in the 1930s, was said to be faster. It never displaced the old one. David's point was the one W. Brian Arthur was making in theory at the same time: under increasing returns, history can lock a market into an arrangement that is not the best available.

Five years later two economists, Stan Liebowitz and Stephen Margolis, published a reply in *The Journal of Law and Economics* with a pointed title, "The Fable of the Keys".[13] They went back to the evidence for Dvorak's superiority and found it thin. Some of the most cited trials, they argued, had been run under conditions that favoured the Dvorak layout, and later, more careful comparisons found little difference in speed. If the alternative was not actually better, the keyboard was not an example of inefficient lock-in at all. It was an example of a market settling on something adequate and staying there because switching would cost more than it would return.

The dispute has never been fully settled, and this chapter does not try to settle it. It begins here because it is the honest place to begin. Complexity economics makes strong claims, and its strongest illustration was challenged within five years by careful people with good evidence. Ammanuel Santa Anna's book is unusually direct about the value of this kind of disagreement:

> *Those differences are not problems to be resolved but productive tensions that keep the Ecstatic Economics framework honest and adaptive. The goal is polyphony, not unison.*[14]

Chapter 2 set out where the idea came from. This one takes five critiques in turn, states each as fairly as its best advocates would, and offers for each a way through: a practice that keeps what is true in the critique and still lets a leader or a household use the idea.

The first critique: lock-in is rarer than it sounds

The critique. Liebowitz and Margolis went further than one keyboard. Across a series of papers and a 1999 book, *Winners, Losers and Microsoft*, they argued that the cases offered as proof of inefficient lock-in, from keyboards to video formats, mostly dissolve on close inspection. Where a better alternative really exists, they argued, someone usually has an incentive to organise the switch: a firm can subsidise early adopters, a standard can be redesigned, a coalition can move together. True lock-in into an inferior option, in their

view, requires that a better choice be both known and blocked, and that is less common than the stories suggest.

What is fair in it. A great deal. Arthur's 1989 model showed that inefficient lock-in is *possible* under increasing returns; it did not show how often it happens. Much of the popular writing that followed treated possibility as frequency. A leader who sees lock-in everywhere will waste effort trying to escape arrangements that are in fact the sensible equilibrium of a market.

The way through. Treat every suspected lock-in as a hypothesis with a price on it. Write down what the alternative is, what evidence says it is better, and what switching would cost in money, time and retraining, with a source for each figure. If the evidence for the alternative is thin, the arrangement is probably not a trap; it is just a habit with a sensible reason. If the evidence is strong and the switching cost is the only barrier, the next question is who could organise the move. That is a practical question with a practical answer, and it often turns out to be you.



The second critique: models that can fit anything

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER THREE

The second critique: models that can fit anything

The critique. Complexity economics relies heavily on agent-based models: computer simulations in which many agents follow rules, interact and adapt, and the analyst watches what emerges. The difficulty is that such models have many moving parts. Each agent's rules, the network that connects them and the way they learn are all choices the modeller makes. With enough choices, a model can be tuned to reproduce almost any pattern in the data, and a model that can reproduce anything has explained nothing.

Giorgio Fagiolo, Alessio Moneta and Paul Windrum set out the problem with care in 2007, in a survey of how agent-based economists validate their models.[15] They found a field with promising methods and no agreed standard: some models were checked against many statistical features of real data, others against one or two, and the choices made in building a model were often under-reported.

What is fair in it. Most of it. The critics are describing a genuine weakness, and the more careful complexity economists say so themselves. J. Doyne Farmer and Duncan Foley's 2009 essay in *Nature*, "The economy needs agent-based modelling", made the case for these models after the financial crisis while acknowledging that they were less developed than the models they hoped to supplement.[16]

The way through. For a reader who is not building models but reading them, three habits help. Ask how many independent patterns in real data the model reproduces, not merely whether it reproduces one. Ask what the model predicts that its builders did not put in on purpose. And prefer the simple model you can understand to the elaborate one you cannot. In an organisation, the same discipline applies to the informal models every leader carries: a story about why sales rose can usually be made to fit the facts after the event. Write the story down before the quarter starts, with what would prove it wrong.



The third critique: everything depends on everything

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER THREE

The third critique: everything depends on everything

The critique. If outcomes depend on small historical accidents, and the economy is always adapting, then perhaps nothing useful can be said in advance. A theory in which anything might happen is a theory that cannot be tested. The charge is sometimes made with a shrug: complexity economics describes the world beautifully and tells you nothing to do on Monday.

What is fair in it. The point about prediction is real. Arthur himself is modest about forecasting specific outcomes. A theory that says "it depends on history" does owe its users an account of what, in the meantime, they can rely on.

The way through. Complexity economics predicts shapes, not points. It does not say which of two technologies will win; it says that one of them will probably win by a wide margin, that the market will become hard to reverse once one leads, and that early adoption matters more than late. Those are testable claims about the distribution of outcomes, and they are useful. A firm entering a market with strong network effects can know, without knowing who will win, that a narrow early lead is worth more than a better product launched later. A household deciding which of two neighbourhood schools to support with its time can know that small early contributions to a growing institution often carry further than larger contributions to a settled one.

Santa Anna's book names the habit of mind that turns this critique into an excuse. Writing about how new economic arrangements are judged, he observes:

> *This is important because perfectionism is one of scarcity consciousness's most effective weapons.*[17]

The critique, pressed too far, becomes exactly that: a demand that a theory either forecast perfectly or be set aside. If the perfect answer cannot be computed in advance, the task becomes building the capacity to learn quickly from what does happen.

The fourth critique: lock-in as an excuse

The critique. This one comes less from economists than from those who watch how ideas are used. Increasing returns and winner-takes-all dynamics, told carelessly, become a story of inevitability. The dominant platform won because it had to; the incumbent is unmovable; the arrangement that profits a few is simply what the system produced. Complexity becomes a vocabulary for accepting concentration rather than questioning it.

What is fair in it. The misuse is common, and the house should take particular care with it, because the opposite of an extractive economy cannot be built by people who believe extraction is a law of nature. A lock-in described as fate is a lock-in nobody will try to open.

The way through. Read lock-in as a design variable. Gregory Unruh's paper "Understanding carbon lock-in", published in *Energy Policy* in 2000, is a model of how to do it.[18] Unruh described the fossil-fuel economy as a techno-institutional complex: technologies, firms, infrastructure, regulation and habits that had grown up together through increasing returns and now held one another in place. He did not conclude that the arrangement was permanent. He concluded that escape requires acting on several reinforcing loops at

once, because any single loop changed on its own is pulled back by the others. That is a strategy, not a lament.

Robin Cowan's 1990 study of nuclear reactors, in *The Journal of Economic History*, is a sobering companion.[19] He traced how one reactor design, the light-water reactor, came to dominate civilian power, partly through early naval programmes and the learning they generated, while alternatives that some engineers considered promising received far less development. Cowan's evidence is historical and his judgement about the alternatives is necessarily contestable. But the paper shows the mechanism clearly, and the mechanism is not destiny. It is the accumulated weight of early decisions, which different early decisions could have shifted.

Santa Anna's book names the same thing in plain language, and refuses the fatalism in the same breath:

> *structures have their own inertia, their own logic, their own capacity to reproduce themselves regardless of the intentions of the people inside them.*[20]

The structures reproduce themselves. That is why the book argues for changing structures as well as minds, and why chapter 4 reads lock-in as a living pattern that can be redirected.



The fifth critique: "it's complex" as a reason to stop

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER THREE

The fifth critique: "it's complex" as a reason to stop

The critique. The last critique is about behaviour, not theory. In many organisations the word *complex* has become a way of ending a conversation. A problem described as complex is a problem nobody is expected to solve, and a leader who says it often enough has found a polite way to do nothing.

What is fair in it. It is accurate about how the word is often used. A theory of complexity that leaves its readers less willing to act has failed them.

The way through. Complexity calls for a different kind of action, not for less. Dave Snowden and Mary Boone, in their 2007 *Harvard Business Review* article "A Leader's Framework for Decision Making", described the Cynefin framework, which sorts situations by the relationship between cause and effect.[21] In what they called the complex domain, cause and effect can be understood only in retrospect, so the right response is to *probe, sense, respond*: run small, safe experiments, watch what happens, amplify what works and dampen what does not. That is a programme of action, sized for uncertainty.

The Luminous 100, the house's set of upgraded management frameworks, reads Ralph Stacey's related matrix in a way that answers this critique directly. Where the classical tool tends to produce alarm at the discovery that an organisation is in the complex zone, the upgrade asks:

> *What if the complexity we're experiencing is not a threat but evidence that we're alive and growing?*[22]

The critique is answered by using the right tools for the zone you are in. A leader who says "it's complex" and then designs three small experiments with a date to review them has turned the word back into what it should be: a description of the ground, not a reason to stay off it.

What remains standing

Five critiques, all serious, and each leaves complexity economics narrower and stronger than its most enthusiastic advocates made it.

What remains is this. Increasing returns are real in many markets, though not in all, and their presence is a question to answer with evidence. Lock-in into an inferior arrangement is possible and probably less frequent than the stories suggest, which makes the cases where it does occur more worth finding. Models of adaptive agents are promising and hard to validate, so their findings deserve the same scrutiny as any other model's. The theory predicts shapes rather than points, and shapes are enough to act on. And complexity calls for small, deliberate experiments, never for paralysis.

A reader who holds the theory this way will use it better than one who holds it as a revelation. Santa Anna's book draws a related line in its own subject, warning against the temptation to blend every tradition into one smooth story. The same discipline applies here: complexity economics is one voice in the lineage, alongside Elinor Ostrom's work on governing shared resources and Karl Polanyi's account of markets embedded in society, and it is more useful for having its limits named.

The keyboard is still on the desk. Whether it is a trap or a sensible habit is an argument economists are still having, and either answer leaves the same useful question for whoever is typing on it: what would it actually cost to change, and who would have to move first?

The Tool: The Lock-In Test

Use this for any arrangement you suspect is holding your organisation or household in place. One sheet per arrangement. It adds a column to the map from chapter 2.

1 · The arrangement

What it is, in one line:

When it began, and how:

What depends on it now (systems, people, contracts, habits):

2 · The alternative

The specific alternative you believe is better:

The evidence that it is better, with its source and date:

- Grade the evidence: one anecdote / several anecdotes / a measured comparison / several measured comparisons

Who would disagree, and their best argument:

3 · The price of switching

COST	AMOUNT	SOURCE OF THE FIGURE	FALLS ON WHOM
Money			
Time			
Retraining or relearning			
Lost compatibility with others			

4 · The verdict, provisionally

If the evidence for the alternative is thin: a sensible habit. Leave it, and note the date you will look again.

If the evidence is strong and switching is cheap: switch. Owner and date:

- If the evidence is strong and switching is costly: a lock-in. Continue to section 5.

5 · The loops holding it in place List every reinforcing loop that keeps the arrangement where it is (training, contracts, suppliers, other people's expectations).

…

…

…

Which loops could be changed together, at the same time?

Who else would need to move, and what would make it worth their while?

6 · The probe

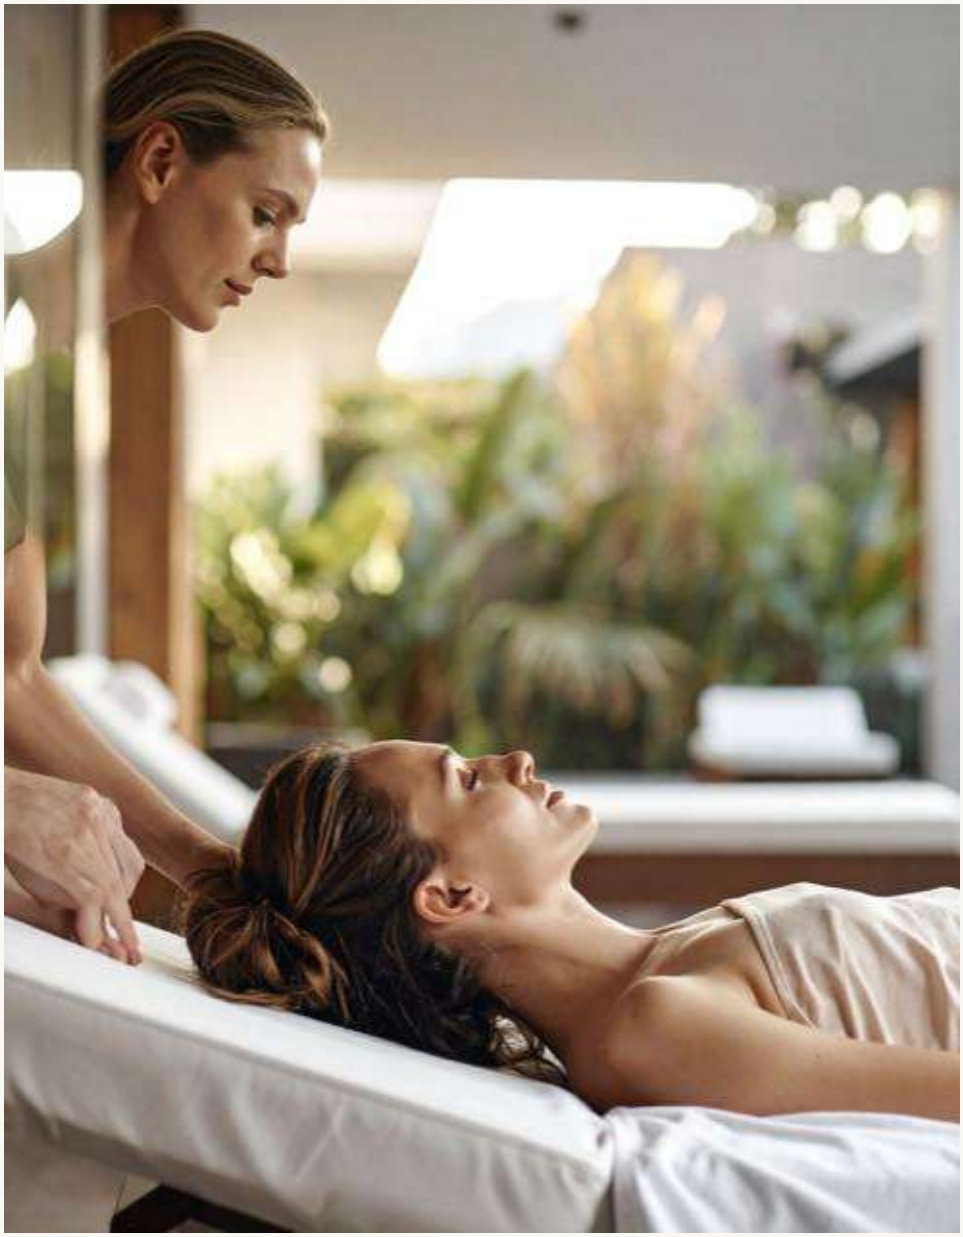
One small, safe experiment to test the alternative:

What result would make you continue? · stop?

Review date:

CHAPTER FOUR

Nobody Designed the Trap



Nobody Designed the Trap

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FOUR

A chequerboard in 1971

In 1971 the economist Thomas Schelling published a paper in *The Journal of Mathematical Sociology* that he had worked out, by his own later account in *Micromotives and Macrobehavior* (1978), with coins on a board.[23] He placed two kinds of coin on a grid and gave each a single, mild preference: a coin was content as long as some modest share of its neighbours were of its own kind, and it would move to a nearby empty square if they were not. No coin wanted a segregated board. Every coin would have been happy in a mixed one.

Then he let them move. A few coins relocated, which changed the neighbourhoods of the coins around them, which prompted a few more to move. Within a handful of rounds the board had sorted itself into large regions of one kind, far more separated than any single coin had wished.

"Dynamic models of segregation" is now one of the founding texts of complexity science, and it is cited far beyond economics, because it shows something uncomfortable with great economy. A pattern at the level of the whole can be produced by the parts without being intended by any of them. Nobody designed the segregated board. Nobody would have chosen it. It emerged.

Ammanuel Santa Anna's *Ecstatic Economics* makes the same observation about the economy he wants to change. Writing about the anxiety that runs through modern economic life, and the way it keeps people overworking and overspending, he refuses the easy villain:

> *This is not conspiracy. It's emergence — the self-organizing property of complex systems that happen to be organized around pathological assumptions. No one designed this trap.*[24]

That passage is the hinge of this chapter. Chapters 2 and 3 set out where complexity economics came from and where it runs thin. This one reads it through the house's own lens, the living-systems reading that Santa Anna's book and the Luminous 100 develop. The claim is simple to state. Complexity economics explains *how* economic patterns emerge from loops. The living-systems reading asks *what those loops carry*: whether they circulate vitality and information back to the people who act, or whether they have been cut.



Emergence without a villain

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FOUR

Emergence without a villain

The value of Schelling's chequerboard, for a reader who runs something, is that it moves blame from people to structure without excusing anyone.

Most organisations explain their worst patterns by finding someone responsible. Meetings run long because one manager likes to talk; the sales team and the delivery team distrust each other because of an old dispute; good people leave because of a particular director. Sometimes the explanation is right. More often, complexity economics suggests, the pattern is produced by many small, reasonable preferences interacting:

each meeting runs a little long because each attendee has one more point, each team protects itself slightly because the other team does, each person who leaves makes the remaining workload heavier for the next one.

Santa Anna's book applies this at the scale of a whole economy. The extractive system, as he describes it, does not need anyone to want it:

> *The extractive economy runs primarily on sympathetic activation. It needs you slightly anxious, slightly driven, slightly afraid that you don't have enough and need to do more, earn more, produce more. This isn't a conspiracy—it's an emergent property of a system built on scarcity assumptions.*[25]

The finer point is in the last clause. An emergent pattern is not random. It is produced by the rules the parts are following, and in a human system those rules include assumptions: that there is not enough, that rest is risky, that a rival's gain is your loss. Schelling's coins followed one rule about neighbours. The people in an economy follow many rules, most of them unwritten, and the patterns that emerge carry the shape of those rules.

This is where the two traditions meet. Complexity economics gives the mechanism: local rules, interaction, feedback, emergence. The living-systems reading adds that the local rules are not fixed facts about human nature. They are learned, often in the body, and they can be changed. Change enough of them, in enough places, and a different pattern emerges from the same mechanism.

The representative agent, retired

A second line of research makes the same point from inside economics.

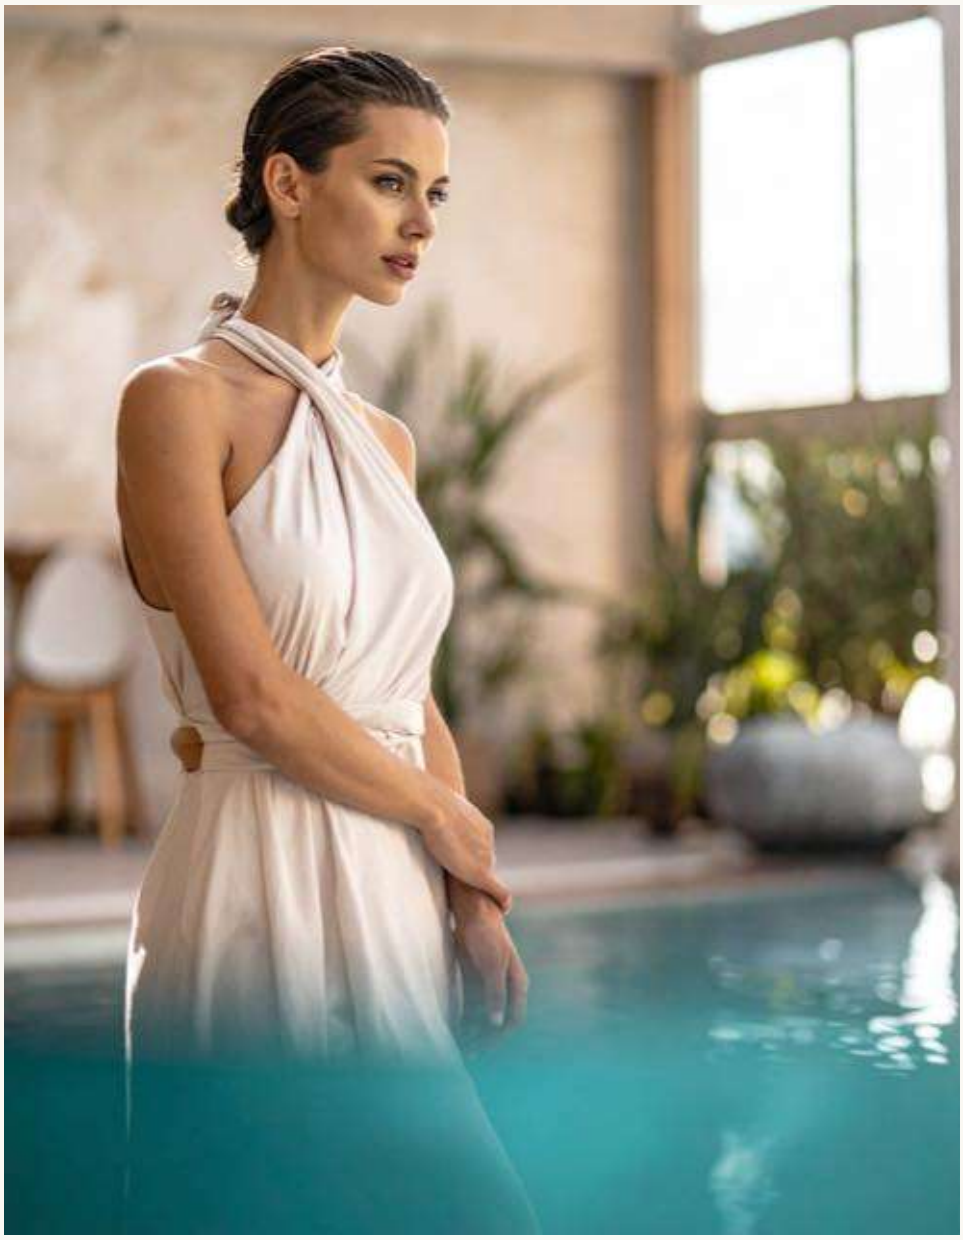
For several decades, much macroeconomic modelling rested on a convenience: the representative agent, a single imagined household whose choices stand in for those of the whole population. Alan Kirman's 1992 paper in the *Journal of Economic Perspectives*, "Whom or What Does the Representative Individual Represent?", is the classic statement of why this is a problem.[26] Kirman showed that the behaviour of a population can differ sharply from the behaviour of any one member of it; that a representative individual's preferences can point the opposite way from the population's; and that the very things economists most want to explain, such as booms, panics and the spread of new practices, arise from interaction among different agents. A model with one agent has no interaction. It cannot produce them.

The living-systems reading takes this further. An organism is not a large cell; a forest is not a big tree. What makes a living system alive is the relationship among its parts, and the information moving between them. Santa Anna's notes reach for the old image of the ant colony to make the point, and his book draws the

conclusion for organisations. When the fearful apparatus of control is lifted and trust is put in its place, he writes,

> *they self-organize toward excellence. Not because they are forced to, but because it is in the nature of living systems to move toward greater coherence, creativity, and vitality when the conditions are right.*[27]

A reader trained in complexity economics will notice that this is a claim about emergence, and should be held to the same standard as any other. Self-organisation is real. Whether it tends toward excellence depends on the rules the parts follow and the information they receive. Santa Anna, to his credit, says so in the next paragraph, and the third movement below takes it up.



The severed loop

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FOUR

The severed loop

If emergence is the mechanism, the house's distinctive diagnosis concerns feedback.

A complex adaptive system learns by feedback. Agents act, the world responds, and the response changes how the agents act next. Arthur's El Farol drinkers improved their forecasts because they could see each week whether the bar had been crowded. Remove that information, and they could not have adapted at all.

Ecstatic Economics argues that much of what goes wrong in modern economic life comes from feedback loops that have been cut, so that the people making decisions no longer receive the consequences of them. The sharpest statement is about investment:

> *When investors cannot feel the impact of their investments, they make decisions that a somatically connected human being would never make. Not because they're evil. Because the system has severed the feedback loop that would normally connect action to consequence.*[28]

This is a complexity economist's sentence written in another vocabulary. A system in which the signal from an outcome never reaches the agent who caused it cannot correct itself. It will keep producing the outcome, however harmful, because nothing in the loop tells it to stop. The book's examples of what repair looks like are deliberately small and concrete:

> *The CSA member who knows the farmer. The worker-owner who shares in the consequences of business decisions. The community currency user who can trace the circulation of value through local relationships. The impact investor who visits the enterprise. In each case, the feedback loop that extractive economics severed has been sutured back together.*[29]

Each example is a loop restored. In each, information about the consequences of a choice travels back to the person who made it, and the choice can then adapt. This is the house's contribution to the complexity reading, stated without any special machinery: an economy is healthiest where its feedback loops are short, honest and felt, and least healthy where they have been lengthened, obscured or cut.



Banks as an ecosystem

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FOUR

Banks as an ecosystem

The most striking evidence that ecologists and economists are describing the same kind of system comes from finance, and it arrived after the crisis of 2008.

In 2008, before the worst of the crisis, Robert May, Simon Levin and George Sugihara published a short commentary in *Nature* titled "Ecology for bankers".[30] May was one of the twentieth century's most important theoretical ecologists. The three authors argued that the banking system had become a network whose stability could be studied with the tools ecologists use for food webs and disease, and that the industry's

own risk models, which treated each bank largely in isolation, could not see the danger in the connections between them.

Three years later Andrew Haldane, then executive director for financial stability at the Bank of England, and May published "Systemic risk in banking ecosystems" in *Nature*.[31] They modelled banks as nodes linked by lending and showed how features that make each bank look safer on its own, such as holding similar assets and diversifying in the same ways, can make the network as a whole more fragile, because the banks become alike and fail together. An ecosystem with one dominant species is efficient until a single disease arrives.

The papers are models, not measurements of the next crisis, and they should be read as such. But they make the house's point with the authority of a central bank. Santa Anna's book draws the same lesson about money itself:

> *Just as a diverse ecosystem is more resilient than a monoculture, a diverse monetary ecosystem — with local currencies, time banks, mutual credit systems, and other complementary instruments operating alongside conventional money — may be more resilient than a monetary monoculture.*[32]

Note the word *may*. The book does not claim more than the evidence allows, and the evidence, from Haldane and May, is a model of how diversity and connection interact, not a proof that any particular complementary currency would have softened a particular crisis.



The loops worth amplifying

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FOUR

The loops worth amplifying

The Luminous 100, the house's set of upgraded management frameworks, reads Peter Senge's systems thinking in the same spirit. Senge's causal loop diagrams distinguish reinforcing loops, where change feeds more change, from balancing loops, where change is corrected. The usual practice is diagnostic: find the vicious cycles. The Luminous upgrade turns the question round:

> *"Where are the virtuous cycles already alive in this system? What reinforcing loops are generating vitality, trust, and developmental momentum — and how do we amplify them until they become the dominant pattern?"*[33]

This is increasing returns read as an opportunity rather than a trap. Arthur showed that under positive feedback, small early advantages can come to dominate a market. The same mechanism can carry a good practice. A team in which trust has started to compound, a supplier relationship that grows more generous with each exchange, a household habit of reviewing its money together that becomes easier each month: each is a reinforcing loop, and each, if noticed early and fed, can become the pattern the rest of the system organises around.

The upgrade of the Cynefin framework makes a related point about where to look. Snowden's complex domain, where cause and effect are clear only afterwards, is usually treated as the difficult zone. The Luminous reading notes that it is also the most generative:

> *Complex domains are where the system is most creatively alive*[34]

The practical consequence for a leader is to stop treating the complex parts of the organisation only as risks to be contained. They are also where new patterns can start, which is to say where small, well-placed choices carry the furthest.

The honest shadow

None of this guarantees a good outcome, and the house's book says so in the passage that follows its praise of self-organisation. Without practices for tension, decision and accountability, it warns, distributed authority can produce chaos or the rule of the loudest voice, and

> *The transition from traditional hierarchy to distributed authority is a developmental process, not a structural flip*[35]

That is the complexity economist's caution in the house's own terms. Emergence does not choose good patterns over bad ones; Schelling's coins produced a pattern nobody wanted. Restoring feedback is necessary and not sufficient. The people receiving the feedback must be able to bear it, interpret it and act on it, and that capacity is built over time.

The way through is the same as in chapter 3: small probes, honest feedback, patience with what is slow. A leader who removes a layer of approval and also teaches the team how to disagree well is changing two loops together, which is how lock-ins are escaped. A leader who removes the layer alone may find the old pattern re-forming within a quarter.

The board, changed

Return to Schelling's chequerboard. The same model that produced segregation from mild preferences can produce integration if the rules change slightly: if coins are content with a lower share of their own kind, or if a few coins actively prefer mixed neighbourhoods, the board stays mixed. Schelling's lesson was never that bad outcomes are inevitable. It was that outcomes depend on rules no one is looking at, and that small changes in those rules can produce large changes in the whole.

Santa Anna's book is a long argument about which rules to change: scarcity for sufficiency, distance for felt consequence, control for trust held with skill. Chapters 5 and 6 test the reading on two public cases, chapter 7 works it through on a composite firm, and chapter 12 joins it to the lineage's work on shared resources and on gainsharing, where the people who create a gain receive part of it and the loop is closed.

The coins are still on the board. Nobody designed the pattern they form, which means nobody has to defend it. That is the opening.



The Tool: The Loop Audit

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FOUR

The Tool: The Loop Audit

A second page for the map begun in chapter 2. It takes the reinforcing loops you listed there and asks what each one carries. Allow ninety minutes with the people who know the system best.

1 · The loops, from your first map Copy each reinforcing loop you listed, as a sentence of the form *the more A, the more B, the more A.*

2 · What each loop carries

LOOP	WHO ACTS IN IT	WHO RECEIVES THE CONSEQUENCE	HOW LONG BEFORE THE CONSEQUENCE IS FELT (DAYS, MONTHS, YEARS)	DOES THE PERSON ACTING EVER SEE IT? YES / NO / PARTLY	DOES THE LOOP CARRY VITALITY, EXTRACTION, OR BOTH?

3 · Severed loops For each loop where the person acting does not see the consequence:

What information would have to travel back, and to whom?

What is the shortest route it could take (a report, a visit, a conversation, a shared figure)?

Who could open that route this month?

4 · Emergent patterns nobody chose List up to three patterns in your organisation or household that nobody seems to want (long meetings, mistrust between two teams, a recurring end-of-month scramble).

PATTERN	THE LOCAL RULES THAT MIGHT BE PRODUCING IT	WHICH ONE RULE, CHANGED SLIGHTLY, MIGHT CHANGE THE PATTERN

5 · Virtuous cycles already alive

One reinforcing loop that is already generating trust or vitality:

What feeds it now:

One small way to feed it more: · owner: · date:

6 · Two loops together

If you change one rule, which other loop is likely to pull it back?

What will you change at the same time, so that the first change holds?

7 · The capacity check

- Are the people who will receive restored feedback able to act on it? yes / not yet

If not yet, what practice would build that capacity (a conversation format, a decision rule, time)?

CHAPTER FIVE

The Valley and the Highway



The Valley and the Highway

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FIVE

Two regions in 1980

At the start of the 1980s there were two places in the United States that a reasonable observer would have named as the centres of the electronics industry. One was the Santa Clara Valley south of San Francisco, which had been called Silicon Valley for about a decade. The other was the ring of suburbs around Boston strung along Route 128, a highway built in the 1950s, where a cluster of minicomputer makers had grown out of the laboratories of the Massachusetts Institute of Technology and the military research of the Second World War.

On most measures an investor would have used at the time, the Boston region looked at least as strong as the Valley. It had older universities, deeper pools of engineering talent, established firms with famous names (Digital Equipment Corporation, Data General, Wang Laboratories, Prime Computer) and a record of turning research into products. Both regions had grown on defence contracts. Both had access to venture capital. Both were full of very able people.

By the early 1990s the two had diverged. Both had been hit hard in the 1980s: the Valley's semiconductor makers by Japanese competition in memory chips in the middle of the decade, Route 128's minicomputer firms by the arrival of the personal computer and the workstation toward its end. The Valley recovered and went on to produce the next generations of the industry. Route 128's leading firms struggled. Wang filed for bankruptcy protection in 1992; Digital, once the second-largest computer company in the world, was sold to Compaq in 1998.

The question of why is the subject of one of the most widely cited books about an industrial region, and it is the first case in this book because it shows the ideas of the earlier chapters at the scale of a whole regional economy, over thirty years, with the evidence gathered by someone who went and asked.

Saxenian's answer

AnnaLee Saxenian, then at the University of California, Berkeley, published *Regional Advantage: Culture and Competition in Silicon Valley and Route 128* in 1994.[36] She had spent years interviewing engineers, managers, founders and investors in both regions, and she set out to explain the divergence without reaching for the usual answers.

The usual answers were about inputs. The Valley had better weather, or Stanford, or more venture capital, or more defence money at the right moment. Saxenian did not dismiss any of these, but she argued that none of them explained the difference, because Route 128 had comparable inputs. What differed was the way the parts of each region were connected to one another.

Route 128, as she described it, was dominated by a small number of large, vertically integrated firms. Each tried to do most things itself: design its own chips, build its own components, write its own software, train its own people. Information was held inside the firm. Engineers tended to stay at one company for a long

time, and loyalty was expected and rewarded. Relations between firms were formal and often guarded. The region was, in her phrase, a collection of independent firms that happened to be near one another.

Silicon Valley, by contrast, was organised as a network. Firms specialised and bought from one another. Engineers moved between companies frequently, carrying knowledge with them, and the boundaries between firms were porous. People who had worked together at one firm met at another, or at a bar, or at a trade association, or in a new venture founded by colleagues who had left. Failure was common and not shameful; a failed start-up released its people and ideas back into the network. Information moved fast, through relationships rather than through formal channels.

When the shocks of the 1980s arrived, the two structures responded differently. A vertically integrated firm on Route 128 that had bet on a particular architecture found it hard to change course, because everything inside the firm had been built around that bet. In the Valley, the same shock released people and ideas into a network in which they recombined quickly, and new firms formed around the new opportunities. The region as a whole adapted even when individual firms did not.

Read with Arthur

Read with the ideas of chapter 2, Saxenian's account is a story about loops.

Both regions had strong increasing returns. In both, success attracted talent, talent produced products, products produced revenue, revenue funded more success. The difference was in what the loops were made of and where they ran.

On Route 128, the reinforcing loops ran mostly inside firms. A firm's success deepened its own capabilities, its own supplier relationships and its own culture. This made each firm very good at what it already did, and it made each firm locked in to what it already did, in exactly the sense Arthur described: its future was increasingly determined by its past, and the structure that made it strong also made it inflexible.

In the Valley, many of the reinforcing loops ran between firms, through the movement of people and information. Success anywhere in the network strengthened the network, which made success elsewhere more likely. The region's capabilities were held less in any single firm than in the pattern of relationships among them. When one firm locked in to a failing design, the region did not, because the region's knowledge was not stored in that firm alone.

There is a useful point here about the level at which lock-in happens. Arthur's model describes a market locking in to one technology. Saxenian's case shows that organisations and regions can lock in to a structure: a way of holding knowledge, of relating to suppliers, of treating people who leave. A region of closed firms can be very productive for a long time and then unable to change; a region of open loops can look disorderly and wasteful and then reorganise itself in a few years.



Read with the living-systems lens

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FIVE

Read with the living-systems lens

Chapter 4 argued that the distinctive question of the living-systems reading is what each loop carries, and whether feedback reaches the people who act. Saxenian's case answers both questions with unusual clarity.

The loops in the Valley carried information, and they carried it to the people who could use it: engineers, who heard within weeks from former colleagues what was working elsewhere. The feedback was short and it was felt. On Route 128 the equivalent information travelled, if at all, upward through a firm's hierarchy

and then down again as a decision, a route that was slow and that filtered out much of what an engineer would have wanted to know.

Santa Anna's book names the principle as one of its design rules for vitality-centred organisations:

> *Organizations that suppress feedback — through hierarchy, blame cultures, or information hoarding — create the economic equivalent of a body that can't feel pain: they lose the ability to adapt and self-correct.*[37]

The passage is written about the inside of an organisation. Saxenian's case shows it holding at the scale of a region. Route 128 did not lack intelligence. It lacked circulation. Its information was held, as capital is held, in reserves that could not easily move to where they were needed.

The Luminous 100, the house's upgraded set of management frameworks, makes the same point about Stafford Beer's Viable System Model, which describes the functions any self-maintaining system must perform. Its upgrade reads every unit of an organisation as itself a viable system, nested inside larger ones:

> *every team is a viable system nested within a viable department nested within a viable organization nested within a viable industry*[38]

Read this way, the Valley was viable at more levels at once. Its firms could fail without the region failing. Route 128's viability was concentrated in a handful of large firms, and when they faltered there was less beneath them to carry the region's capability forward.



A law that kept the loops open

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FIVE

A law that kept the loops open

One part of the story deserves its own section, because it shows a rule nobody designed for the purpose shaping a whole region's dynamics.

In 1999 the legal scholar Ronald Gilson published an article in the *New York University Law Review* titled "The Legal Infrastructure of High Technology Industrial Districts: Silicon Valley, Route 128, and Covenants Not to Compete".[39] Gilson took Saxenian's account and asked what might explain the difference in how often engineers moved between firms. His answer pointed to a provision of California law, section

16600 of the Business and Professions Code, under which contracts restraining a person from engaging in a lawful profession are generally void. In practice, this meant that California courts would not enforce most covenants not to compete: an engineer who left a Valley firm could go to work for a competitor the next week. Massachusetts courts, by contrast, generally enforced reasonable non-compete agreements.

Gilson's argument was that this difference, which predated the electronics industry and had nothing to do with it, may have helped produce the high mobility of Valley engineers, and with it the knowledge spillovers that Saxenian described. Later empirical work has tested the idea, with mixed and debated results, and Gilson himself presented it as a hypothesis about one contributing cause rather than as the whole explanation. It is cited here because the mechanism, whatever its size, is exactly the kind of thing complexity economics teaches a reader to look for: a small, old rule in the background, which changes the cost of moving between two points in a network, and which, compounded over thousands of individual careers, shapes the pattern of a regional economy.

The general lesson for a leader is close to home. Every organisation has rules, written and unwritten, that set the cost of information and people moving across its internal boundaries. Who may talk to whom without permission. Whether a person who moves between teams is treated as a loss or a gain. Whether a project that ends is treated as a failure to be hidden or a release of knowledge to be shared. These rules rarely appear on any strategy document, and they may matter more than most of what does.

What the case cannot show

The case is strong and it should not be overstated. Three limits are worth stating plainly.

First, it is two regions. A comparison of two cases can suggest a mechanism; it cannot prove one. There are many differences between Northern California and Massachusetts, and Saxenian's critics have pointed to several she gave less weight to, including the particular timing of the personal computer and the very large flows of venture capital into the Valley in the 1990s.

Second, Saxenian's account is interpretive. It rests on interviews and on her reading of how two cultures worked. It is careful and widely respected, but it is not a measurement of information flows, and a reader should hold it as a well-argued account rather than as a quantified result.

Third, the outcome was not inevitable. Route 128 did not disappear; the Boston region later became a strong centre for biotechnology and other industries, and some of its firms adapted well. The divergence Saxenian described was real for the computing industry in a particular period. It is not a law about all regions for all time.

Paul Krugman's 1991 paper on increasing returns and economic geography, in the *Journal of Political Economy*, offers a useful frame for holding the case.[40] Krugman showed formally how increasing returns and transport costs can produce a core-periphery pattern, in which industry concentrates in one region rather

than spreading evenly, and how small differences in starting conditions can decide which region becomes the core. His model is about where industries cluster, not about the internal culture of a cluster, so it complements Saxenian rather than confirming her. Taken together they suggest two levels of path dependence: history decides where a cluster forms, and the structure of the cluster's loops decides how well it adapts once formed.



The finer points for a leader

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FIVE

The finer points for a leader

Five practical consequences follow, and each can be tested in an ordinary organisation within a quarter.

Map where your learning is stored. Ask, of each important capability in the organisation, where the knowledge that sustains it actually lives. In one person? In one team? In a set of documents nobody reads? In relationships with three suppliers? A capability stored in one place is a Route 128 capability: strong until that place falters. A capability held across a network of people who talk to one another is closer to the Valley's.

Price the cost of movement. For people moving between teams, and for information moving between them, ask what it costs today in time, permission and reputation. If a capable engineer who moves to another team is quietly marked as disloyal, the cost of movement is high, and the organisation is paying for it in learning it never sees.

Treat endings as releases. When a project ends, a team is disbanded or a product is retired, the people and knowledge involved go somewhere. In a closed system they go into a drawer. In an open one they recombine. A short, honest retrospective, published to everyone who might use it, and a deliberate effort to place the people where their knowledge is wanted, turn an ending into a release.

Specialise and connect, rather than integrate. The case for doing everything in-house is usually made on grounds of control and cost. Saxenian's case adds a cost the usual analysis misses: integration closes loops that might have brought in learning from outside. This does not mean outsourcing everything. It means asking, for each activity, whether the organisation learns more by doing it or by connecting to someone who does it better and talking to them often.

Watch the rules nobody designed for this. Every non-compete clause, approval policy, travel budget, confidentiality default and meeting convention sets the cost of some flow. Most were written for good local reasons. Some, compounded over years, may be shaping the organisation's capacity to adapt more than any strategy. Chapter 3's Lock-In Test applies to them as well as to technologies.



The household reading

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER FIVE

The household reading

The same distinction holds for one person's working life. A career built entirely inside one organisation, with all of its knowledge and relationships held there, is strong while that organisation is strong. A career that also maintains open loops (former colleagues kept in touch with, a skill practised in a community outside work, a habit of sharing what one learns) is less efficient in any single year and much better at reorganising itself when a shock arrives. Neither is wrong. The question is whether the balance is chosen or merely inherited.

The two regions now

Saxenian closed her book with an argument rather than a forecast: that regions, like firms, can learn, and that the institutions and habits which keep loops open can be built deliberately. That argument is this book's as well. The Valley's open network was not designed by anyone. It emerged, partly from a legal accident, partly from the culture of a few early firms, partly from luck. But once its mechanism is understood, a leader does not need to wait for luck. The cost of movement, the treatment of endings, the place where learning is stored and the rules that set the price of a conversation are all within reach.

The highway is still there, carrying commuters past buildings that once held the most famous computer companies in the world. So is the valley. The difference between them, in the years that mattered, was not what either region had. It was how what each had was allowed to move.

The Tool: The Cluster Reading

For any network you belong to or lead: a firm and its suppliers, a set of teams, an industry association, a professional community, a neighbourhood of small businesses. Two hours, with three people who know the network from different positions.

1 · The network

Name it, and draw its edge: who is inside?

The main nodes (firms, teams, people), up to twelve:

2 · Where learning is stored

CAPABILITY THAT MATTERS	WHERE THE KNOWLEDGE LIVES (ONE PERSON, ONE TEAM, DOCUMENTS, RELATIONSHIPS ACROSS SEVERAL NODES)	IF THAT PLACE FAILED TOMORROW, WHO ELSE COULD CARRY IT?

3 · The cost of movement

FLOW	WHAT IT COSTS TODAY (TIME, PERMISSION, REPUTATION)	WHO SET THAT COST, AND WHEN	WOULD YOU SET IT THE SAME WAY TODAY?
A person moving between nodes			
Information moving between nodes without a manager			
A failed project's lessons reaching others			
An outside expert talking to an inside team			

4 · Loops inside, loops between

Three reinforcing loops that run inside a single node:

Three reinforcing loops that run between nodes:

Which set is stronger? What would happen to the network if the strongest single node faltered?

5 · Rules nobody designed for this List up to five rules (contracts, policies, defaults, habits) that change the cost of movement in the network, though they were written for another purpose.

…

...

...

Which one could be changed by the people in this room? · owner: · date:

6 · One release

One recent ending (a project, a product, a team) whose knowledge went into a drawer:

What it would take to release it now:

A note on method

The dates for Wang Laboratories' bankruptcy filing (1992) and Digital Equipment Corporation's sale to Compaq (1998) are matters of public record and are given from standard business histories; they were not re-read from court or company filings for this book.

CHAPTER SIX

The Week After the Fire



The Week After the Fire

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SIX

Saturday, the first of February

Early on Saturday, 1 February 1997, a fire broke out at a plant in Kariya, in Aichi Prefecture, belonging to Aisin Seiki, one of the largest suppliers in the Toyota group. The plant made, among other things, a small brake component called a proportioning valve, known in the industry as a P-valve. It is not a part anyone outside the trade has heard of. It regulates pressure to the rear brakes, it is about the size of a fist, and every Toyota vehicle then in production needed one.

Aisin was the sole source. The line that made P-valves was destroyed along with much of the specialised machinery that made them. And Toyota, the company that had given the world just-in-time production, held only a small stock of the part, because holding small stocks was the point of its system. Within days its assembly plants would have nothing to put into the brakes of the cars they were building.

What followed is one of the most studied episodes in the history of supply chains. It was written up in 1998 by Toshihiro Nishiguchi of Hitotsubashi University and Alexandre Beaudet of the Mitsubishi Research Institute in an article in *Sloan Management Review* titled "The Toyota Group and the Aisin Fire".[41] It is the second case in this book because it shows, more clearly than almost any other public record, a network of firms reorganising itself around a sudden failure without a central plan, and doing it in days.



The shape of the problem

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SIX

The shape of the problem

The fire exposed what looked, in hindsight, like a textbook vulnerability. A single supplier for a critical part. A single plant. Almost no inventory, by design. A conventional risk analysis would have flagged it, and after the event many commentators did: here was the danger of efficiency taken too far, of a system so lean that it had no slack to absorb a shock.

That reading is not wrong, and the episode is often taught as a warning about single sourcing. But it misses what Nishiguchi and Beaudet found most interesting, which was not the fragility of the arrangement but

the speed and manner of its recovery. Toyota's assembly plants did stop. According to their account, they were reopened after only two days of shutdown. Within a few weeks, P-valve supply was back to normal. In the meantime, the valves had been produced by firms that had never made them before, using machinery that had not been designed for them, coordinated by no single authority.

What happened in the network

The story, as the authors tell it, runs roughly as follows.

Aisin, with its own plant lost, distributed the design drawings for the P-valves to other firms in and around the Toyota group, along with whatever specialised equipment and know-how it could recover or describe. Firms that ordinarily made quite different components took the drawings, set up improvised production on whatever machinery they had that could be adapted, and began making valves. Some firms worked out the processes themselves. Others borrowed tools or staff from one another. Toyota sent hundreds of its own staff to Aisin and to the firms helping it. The well-known supplier Denso was among those that took on production; the authors name others, including firms that had no direct business relationship with the P-valve at all.

The striking thing, in Nishiguchi and Beaudet's account, is how little of this was ordered from the top. There was coordination, and Toyota and Aisin played central roles in it, but much of the recovery was organised laterally, between firms that knew one another, trusted one another, and were used to sharing information and solving problems together. Questions of payment and compensation were, by their account, largely left to be settled afterward. Firms began work first.

The authors called what they observed self-organisation, and they were careful to explain why it was possible. The Toyota group had spent decades building the conditions for it: dense relationships among suppliers, regular exchanges of engineers, supplier associations in which firms shared practices, and a shared problem-solving culture that the Toyota Production System had spread across the network. None of that had been designed to cope with a fire at Aisin. All of it turned out to be what made coping possible.



Read with Arthur

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SIX

Read with Arthur

This book has spent several chapters on lock-in, and the Aisin case could be read as a lock-in story: the network had locked in to a single source, and paid for it. That is part of the truth. But the more useful reading, in Arthur's terms, is about what the network's history had built.

Increasing returns operate on relationships as well as on technologies. Every time two Toyota-group suppliers exchanged engineers, solved a shared problem, or met in a supplier association, the cost of their next collaboration fell and their ability to work together rose. These are reinforcing loops of exactly the kind

chapter 2 asked the reader to look for: the more firms cooperated, the more they trusted one another; the more they trusted, the more cheaply they could cooperate. Over decades those loops accumulated into a capacity that appeared on no balance sheet and in no contingency plan. When the fire came, the network spent it.

Path dependence cuts both ways here. The same history that made the group dependent on Aisin for P-valves had also made the group able to replace Aisin, temporarily, within days. A network without that history, faced with the same fire, would not have had the second option.

Read with the living-systems lens

The living-systems reading asks of any system whether its feedback reaches the people who can act, and whether its loops carry vitality or extraction. The Aisin case is one of the cleanest public examples of a system answering yes to both.

The information needed to solve the problem (the designs, the process knowledge, the location of usable machines, the availability of people) reached the people who could act on it within hours, through channels that already existed. Nobody had to escalate a request through several layers and wait for a decision to come back. And the loops that carried the response had been built on a long practice of mutual obligation, so that a firm asked to drop its own work and make an unfamiliar part for a neighbour did so without first negotiating the price.

Santa Anna's book, writing about policy and governance, makes an observation that applies with some force:

> *This is the feedback loop we must understand and interrupt.*[42]

The passage is about the loop by which policy shapes consciousness and consciousness shapes policy, and it argues that the loop can be changed deliberately. The Aisin case shows the same principle in a supply chain. Toyota's way of working had shaped how its suppliers thought about one another; how they thought about one another shaped how they worked together; and the loop, sustained over decades, had produced a network in which a stranger's emergency was treated as a shared one. That loop had not been left to chance. It had been built.

The Luminous 100's upgrade of the fractal organisation, the idea that the same patterns should recur at every scale of a system, asks a question that fits the Toyota group well:

> *What is most alive at the smallest unit that, if amplified fractally, would transform the whole?*
> [43]

In the Toyota group, the smallest unit was the shared habit of solving a problem together on the shop floor. Amplified across firms through supplier associations and engineer exchanges, it became the network's capacity to rebuild a production line in a week.



What the case cannot show

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SIX

What the case cannot show

Three cautions.

First, the main source is one article by two researchers who had close access to the group. It is careful and has been widely cited, but it is a qualitative account. The exact number of firms involved, the precise timings and the costs of the recovery are given in the article to the degree the authors could establish them, and this chapter repeats only what is stated there in general terms.

Second, the episode is not a proof that the Toyota approach to suppliers is better than every alternative. Single sourcing remained a real vulnerability, and other firms in the region were affected: Nishiguchi and Beaudet note that Mitsubishi also had to close some assembly lines in the days after the fire. After later disasters, notably the earthquake and tsunami of March 2011, the Japanese automotive industry, Toyota included, faced much longer disruptions, which suggests that the self-organising capacity shown in 1997 has limits when the shock is large enough or when it strikes many nodes at once.

Third, the cultural conditions that made the recovery possible were built over decades in a particular national and industrial context. They cannot be copied by adopting a policy. What can be learned is the mechanism, and the mechanism is general: a network whose members have practised solving problems together can reorganise around a failure faster than any central plan.

The finer points for a leader

Five practices follow, each within reach of an ordinary organisation.

Count the relationships that would carry you through a failure. For each critical part, supplier, system or person, ask not only what the backup is, but who would actually pick up the telephone, and whether they have ever worked with you on anything difficult. A backup supplier with whom you have never solved a problem is a name on a list. A neighbour with whom you have solved twenty is a capacity.

Practise together before you need to. The Toyota group's engineer exchanges and supplier associations were not emergency measures. They were ordinary practice, repeated for years. A leader who wants a resilient network should look for the ordinary, low-stakes occasions on which people from different parts of it can solve real problems together: a shared improvement project with a key supplier, a rotation of people between two teams, a quarterly session in which one unit brings its hardest problem to another.

Let information travel sideways. In a crisis, the speed of the response is set by how quickly information reaches the people who can act. If every request must go up a hierarchy and come back down, the response will be slow. If people at the working level already know whom to call in other units and are trusted to call them, it will be fast. This is set long before the crisis, by everyday norms about who may talk to whom.

Settle the money afterward, and say so in advance. One of the most striking features of the Aisin recovery was that firms began work before the terms were agreed. That is only possible where there is trust that the terms will be fair. An organisation can make this explicit: a standing agreement with key partners that in a

defined kind of emergency, work starts first and the accounting follows on a stated principle. It costs nothing until it is needed.

Rehearse the failure of one node. Pick the single point in your system whose failure would hurt most, and walk through, with the people who would have to respond, what would happen in the first forty-eight hours. Do not write a plan; ask who would call whom, what they would need to know, and where it is. The rehearsal will reveal which loops are open and which exist only on paper. The tool below is written for that rehearsal.

The household reading

A household has its single points of failure too: one income, one car, one person who knows where the documents are, one neighbour who has a key. The Aisin case suggests that the most useful preparation is not only a reserve of money, though that matters, but a set of relationships in which help has already flowed in both directions. A family that has helped its neighbours through small difficulties, and been helped by them, holds a capacity that no insurance policy provides. It is built in ordinary weeks, by ordinary reciprocity, and it is spent on the bad ones.

Santa Anna's line about the gift, quoted in chapter 1, applies here literally. Help that circulates in ordinary times is the reserve that a household or a network draws on when something burns.



The plant, rebuilt

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SIX

The plant, rebuilt

Aisin rebuilt its P-valve line, and in time Toyota and its suppliers reconsidered how many sources they wanted for parts like it. The lesson most often drawn from the fire is about the danger of putting everything in one basket, and it is a fair lesson.

The lesson this book draws is a second one, and it is less often told. For a week in February 1997, the basket did break, and the network caught what fell out of it. It did so because, for decades before, the firms in it

had been doing the slow, unglamorous work of learning to solve problems together. That work had no line in anyone's budget. It was there when it was needed.

The Tool: The Shock Rehearsal

Ninety minutes, with the people who would actually have to respond. One sheet per critical node.

1 · The node

The single part, supplier, system, person or account whose failure would hurt most:

Why it is critical (what depends on it):

How much buffer exists today (stock, time, money, a second person who knows):

2 · The first forty-eight hours Walk through them aloud. Write down what actually happens, not what the plan says.

HOUR	WHO NOTICES	WHO THEY CALL	WHAT THAT PERSON NEEDS TO KNOW	WHERE THAT INFORMATION IS NOW
0–2				
2–12				
12–24				
24–48				

3 · The relationships that would carry you

WHO COULD HELP	HAVE YOU EVER SOLVED A HARD PROBLEM TOGETHER? YES / NO	WOULD THEY START BEFORE TERMS WERE AGREED? YES / NO / UNSURE	WHAT WOULD MAKE THAT MORE LIKELY?

4 · Loops that exist only on paper

Contacts, agreements or procedures that the rehearsal showed nobody would actually use:

Information that the right person would not be able to find:

5 · Practice, in ordinary weeks

One low-stakes occasion in the next quarter to solve a real problem with one of the relationships above: · owner: · date:

One standing agreement to make in advance (for example, "in an emergency of this kind, work starts first and we settle on this principle afterward"):

6 · The household version

The household's single point of failure:

The person or neighbour who would help in the first forty-eight hours:

One way help could flow toward them this month:

A note on method

The account of the 2011 disruptions is given in general terms from public reporting and is not a finding of the 1998 article. Nothing in this chapter is a recommendation about any particular supplier strategy.

CHAPTER SEVEN

A Firm Read as Weather



A Firm Read as Weather

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SEVEN

A Tuesday in February

The firm in this chapter is a composite. It is called Larkfield here, it makes scheduling and records software for independent veterinary clinics, and it does not exist. It was assembled for teaching from the patterns of several small software companies, and every figure attached to it below was set by the writer to make the arithmetic clear. Where a figure comes from a public source, the source is named. Where it does not, it is Larkfield's, and should be read as an illustration.

On a Tuesday in February, Larkfield's founder sits at a pale oak table in a converted mill with her two most senior colleagues and the page described in chapter 2: the First Map of Loops. Larkfield has sixty people, four hundred and ten clinic customers, and annual revenue of nine point two million dollars. Ninety-four in every hundred clinics renew each year. The founder's question is the one every leader of a company that size eventually asks. Growth has slowed, the board would like a plan, and she has three possible directions and money for one.

The earlier chapters set out the theory: Brian Arthur's increasing returns and lock-in, the fair critiques of both (chapter 3), and the living-systems reading that asks what each loop carries (chapter 4). This chapter works them through on one firm, in the order a leader would use them, with the finer points that only show up in practice.

Ammanuel Santa Anna's book has a sentence worth keeping beside the page before beginning, because two of the three lock-ins on Larkfield's map turn out to be about how the firm charges:

> *Money is not neutral. It is a design system, and the design shapes the behavior.*[44]

The living-systems reading does not offer an answer to the founder's question. What it offers is a better question, and on that Tuesday the better question turns out to be worth more than any of the three directions on the table.



The first map

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SEVEN

The first map

The founder starts with the reinforcing loops, written as sentences.

Referral. The more clinics use Larkfield, the more clinic managers recommend it to colleagues at regional veterinary meetings, and the more clinics use it. Larkfield's own records show that thirty-eight of every hundred new customers in the past two years named a referral from an existing clinic.

Learning. The more clinics Larkfield serves, the more edge cases its support team sees, the better the product handles the odd realities of veterinary scheduling (emergency surgeries, boarding animals, mobile vets), and the more clinics it can serve well.

Hiring. The more the firm is known as a calm place to work, the easier it is to hire experienced engineers from larger companies, the better the product becomes, and the more the firm's reputation grows. The founder is less sure of this loop; she writes it down with a question mark.

Then the lock-ins.

The integration. Nine years ago, to win its first large customer, Larkfield built a connection to one particular practice-management system used by many clinics for billing. Everything since has been built on top of it. Seventy of every hundred support tickets now involve that integration in some way.

The annual contract. Larkfield sells only annual contracts, paid in advance, because that is how its first customers wanted to buy. Smaller clinics have asked for monthly terms for years.

The Thursday release. The product ships updates every Thursday evening, a habit from when the firm had four engineers and one of them preferred to work late on Thursdays. Nobody now remembers why, and the support team spends every Friday dealing with the consequences.

Finally, the El Farol question from the first map. The founder's main forecast is that independent clinics will keep consolidating into larger groups, which would shrink her market. She writes the second half of the question: if every competitor holds the same forecast, they will all chase the consolidating groups, and the independent clinics that remain will be served by nobody in particular. That observation, written in pencil on the Tuesday, changes the rest of the plan.

Testing the lock-ins

Chapter 3 offered a Lock-In Test, and its first rule matters here: treat every suspected lock-in as a hypothesis with a price on it.

The Thursday release fails the test as a lock-in. The alternative, releasing on Tuesday mornings when the full team is present, has no credible case against it, and switching costs one conversation and a calendar change. It is a habit with no reason left. The founder moves it that week. It is the smallest decision on the page and the support team mentions it for months.

The annual contract is more interesting. The evidence that smaller clinics want monthly terms is strong: requests logged by the sales team, and three lost deals in the past year where the clinic named payment terms as the reason. The switching cost is real but bounded: a change to billing, a change to how revenue is forecast, and some risk that existing customers will move to monthly and the firm will receive less cash in advance. The founder decides to test it rather than switch, which is the subject of the probes below.

The integration is a true lock-in in Arthur's sense. It began by accident, when the first large customer happened to use that billing system. Everything depends on it. It is inflexible: replacing it would mean rebuilding the core of the product. And it is plausibly inefficient: the support load suggests a more modern connection would serve clinics better. But the Lock-In Test's fifth section asks which loops hold it in place, and the answer is instructive. The integration is held in place by the learning loop. Nine years of edge cases have been handled inside it. To replace it is to lose much of what the firm has learned.

This is the point at which many firms make an expensive mistake. They see the lock-in, decide it is bad, and fund a large rebuild. The living-systems reading suggests something slower: change the loops that hold the arrangement, rather than the arrangement itself. The founder decides that every new piece of clinic logic will be written so that it does not depend on the old integration, and that the integration will be retired gradually as less depends on it. The lock-in will loosen as the loops that feed it are redirected.



What the accounts cannot see

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SEVEN

What the accounts cannot see

Chapter 4 introduced the Loop Audit, which asks of each reinforcing loop who acts, who receives the consequence, and whether the person acting ever sees it. On Larkfield's page, one loop stands out.

The referral loop is Larkfield's most valuable asset, and nobody inside the firm can see it working. The clinic managers who recommend the product do so at regional meetings, in conversations Larkfield never hears. The sales team records "referral" as the source of a new customer and nothing more. The engineers who build the features that clinic managers praise never learn which features those are. The loop is run-

ning, generating thirty-eight of every hundred new customers, and its feedback reaches nobody who could feed it.

Santa Anna's book describes the same blindness at the scale of a national economy, in a passage about forests:

> *In standard economic accounting, a forest has value only when it's cut down. Standing, it contributes nothing to GDP.*[45]

Larkfield's accounts treat its referral loop the same way. The loop has no line in the budget, no owner and no measure. It appears only after the fact, as revenue attributed to sales. If it weakened, the first sign would be a slower year, months after the cause.

The founder's response follows the book's prescription: suture the loop. She asks the support team to call ten clinics that referred others in the past year and ask one question, what they told their colleague. She asks that the answers be read at the next engineering meeting, unedited. It costs perhaps fifteen hours of staff time. It is the second probe below.

The evidence on what to build

That leaves the founder's original question: which of three directions to fund. The three are a version of the product for large clinic groups, a version for horse and farm veterinarians, and a payments service for clinics.

Here a piece of public research is useful, though it was not written about software firms. In 2007 César Hidalgo, Bailey Klinger, Albert-László Barabási and Ricardo Hausmann published "The Product Space Conditions the Development of Nations" in *Science*.[46] They mapped which products countries export and found that products cluster: countries that make one product tend to make others requiring similar capabilities, and countries tend to move into new products that are close, in this sense, to the ones they already make. Two years later Hidalgo and Hausmann extended the argument in the *Proceedings of the National Academy of Sciences*, treating a country's productive knowledge as a set of capabilities that combine, and showing that the diversity of what a country makes carries information about its future growth.[47]

The research is about nations, it is built on export data, and its predictive claims have been debated; it cannot tell a sixty-person firm what to build. What it supplies is a principle consistent with Arthur's path dependence: capabilities grow by adjacency. Each thing an organisation learns to do makes some next things easier and leaves others as far away as ever.

Read that way, the founder's three options sort themselves. The large-group product is adjacent to what Larkfield knows, but every competitor is chasing it, as the El Farol question showed. The payments service is far from Larkfield's capabilities; it would need new expertise in regulation and risk. The farm-vet

product is adjacent, since scheduling and records are much the same, and it serves exactly the population of independent practices the competitors are leaving. It also feeds the learning loop with a new set of edge cases, and the referral loop through a new set of regional meetings.

The founder does not commit to it. She probes it.



Three probes, twelve weeks

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SEVEN

Three probes, twelve weeks

The Cynefin framework, discussed in chapter 3, recommends that in complex conditions a leader *probe, sense, respond*: run small, safe experiments and amplify what works. The Luminous 100's reading of solutions-focused practice asks the question that sets the tone for doing so, which is to look first

> *where in the living system is the desired future already expressing itself, even in small ways?*[48]

Larkfield's founder runs three probes over twelve weeks, each with a prediction written down in advance, as chapter 3 recommended.

Probe one: monthly terms. Offered to new clinics with fewer than four vets, in two regions only. Prediction: at least six new small clinics in twelve weeks, against a baseline of two in the same regions over the same period last year. Stop condition: if existing customers in those regions ask to switch in numbers that would reduce advance cash by more than five per cent.

Probe two: the referral loop. The ten calls, read aloud to the engineers, then repeated monthly. Prediction: at least two features named by referrers that the engineers did not know were valued. Measure: the list itself.

Probe three: the farm-vet product. Not a product, but six clinics invited to use the existing software with a small set of changes, for free, in exchange for a weekly fifteen-minute call. Prediction: at least four of the six still using it at twelve weeks, and at least one referral from among them.

None of these is large. The combined cost, in Larkfield's arithmetic, is about four per cent of a year's engineering time and some forgone revenue on the free trial. Santa Anna's book makes a case for exactly this scale of response, writing about economic practices more broadly:

> *Perhaps, as E.F. Schumacher suggested, small really is beautiful — and the appropriate response to the scale of our problems is not bigger solutions but more solutions, connected through networks of mutual support.*[49]



What happened, in the composite

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER SEVEN

What happened, in the composite

The results, like everything else about Larkfield, are illustrative. Monthly terms brought nine small clinics in two regions, and one existing customer asked to switch; the founder extends the probe to two more regions. The referral calls named a feature the engineers had considered minor, a single screen showing every animal boarding overnight, as the reason seven of the ten clinics had recommended the product; the next quarter's work includes making it better. Five of the six farm-vet clinics were still using the software at twelve weeks, and two had referred colleagues. The founder takes a proposal for the farm-vet product to

the board, with the probe's figures and the El Farol reasoning, and with a budget a fraction of what any of the three original directions would have cost.

The book's account of open-source software suggests why the referral finding mattered as much as it did:

> *The open-source movement is particularly significant for Ecstatic Economics because it provides empirical evidence for one of our core propositions: vitality-based motivation produces superior outcomes in complex creative work.*[50]

The engineers who heard what clinic managers valued in their work did not need to be told to improve the boarding screen. The restored loop carried the motivation with it.

The finer points, for a leader working this through in a real firm, are three. Write predictions before probes, or every result will look like success afterward. Change two loops at once where a lock-in is involved; Larkfield redirected the learning loop while loosening the integration. And put the feedback in front of the people who can act on it, unedited, rather than summarised by a manager on the way.

A household can run the same sequence at smaller scale. Map the loops (a savings habit that compounds, a commute that drains), test the lock-ins (the subscription nobody would choose again, the phone contract renewed by default), find the severed loop (spending whose consequence is felt only when the statement arrives), and probe one change for twelve weeks with a prediction written down. The tool below is written for a firm and works at a kitchen table.

The tools Larkfield used have instrument counterparts in the Alpha Papers: *Alpha Real Options* for pricing the probes as options, and *Alpha Wardley Map* for reading which capabilities sit next to which.

On a Tuesday in May, the founder sits at the same oak table with the same page. The pencil lines are still there, a little smudged. Beside the referral loop, in ink now, is a list of seven clinic names.

The Tool: The Loop Ledger

The three earlier pages (the First Map, the Lock-In Test, the Loop Audit) gathered into one ledger for a twelve-week cycle. Fill in sections 1 to 4 in week 1, section 5 in week 2, and section 6 in week 12.

1 · The system and the question

System (firm, team, product, household):

The decision on the table, in one line:

The options under consideration:

2 · Loops, with evidence

REINFORCING LOOP (THE MORE A, THE MORE B, THE MORE A)	EVIDENCE, WITH SOURCE	WHO ACTS	WHO FEELS THE CONSEQUENCE	DOES THE ACTOR SEE IT?	CARRIES VITALITY / EXTRACTION / BOTH

3 · Lock-ins, tested

ARRANGEMENT	EVIDENCE FOR THE ALTERNATIVE (GRADED)	SWITCHING COST	VERDICT: HABIT / SWITCH / TRUE LOCK-IN	LOOPS HOLDING IT IN PLACE

4 · The El Farol line

Your forecast:

What everyone else will do if they share it:

Where that leaves room nobody is serving:

5 · Adjacency and probes

OPTION	CAPABILITIES IT NEEDS THAT YOU HAVE	CAPABILITIES IT NEEDS THAT YOU LACK	CROWDED BY OTHERS' FORECASTS?	PROBE (SMALL, SAFE, TWELVE WEEKS)	PREDICTION, WRITTEN NOW	STOP CONDITION

One severed loop to suture during the twelve weeks: · route the feedback will travel: · read unedited to:

Two loops to change together, if a lock-in is involved: and

6 · Week twelve

PROBE	PREDICTED	HAPPENED	AMPLIFY / ADJUST / STOP

What the system now shows that the week-one map did not:

The next decision, and the next twelve-week date:

A note on method

Larkfield is a composite written for teaching. It is not a real company, its figures (sixty staff, four hundred and ten clinics, nine point two million dollars in revenue, ninety-four per cent renewal, thirty-eight per cent of new customers by referral, seventy per cent of support tickets on the integration, the probe results) were set by the writer, and its outcomes are illustrations, not evidence. The twelve-week cycle is this book's recommendation, not the finding of any study. Nothing here is investment advice.

A NOTE FROM THE HOUSE

Larkfield's ledger shows what a firm can learn in twelve weeks by working with its loops rather than pulling at levers. Where a loop has been severed, the gain is usually already there, waiting for the feedback to travel. Luminous Prosperity Inc. reads those loops inside organisations as a gainshare, paid as a share of the gain it helps find.

REQUEST A CONVERSATION

CHAPTER EIGHT

The Field Manual



The Field Manual

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER EIGHT

What this chapter is for

The chapters before this one gave the ideas, the critiques, two public cases and one composite. This chapter gives the sequence. It is written so that a leader can run the whole practice in an organisation over twenty-six weeks, with named roles, a decision at each stage, and the mistakes most likely to occur, in the order they tend to occur.

Twenty-six weeks is this book's recommendation, not a finding from research. It is long enough for two twelve-week probe cycles and short enough to fit inside a planning year. A small organisation can compress it; a large one may run it first in one division. The sequence matters more than the timing.

The manual assumes an organisation of between twenty and a few thousand people. For a household, chapter 11 gives a shorter version of the same sequence.

The roles

Five roles, which in a small organisation may be held by three people and in a large one by a small team for each.

The sponsor. The person with authority over the system being examined, usually the chief executive or the head of the division. The sponsor's job is not to run the practice but to protect it: to give it time, to read its findings unedited, and to make the decisions it brings forward. A sponsor who delegates the reading of the findings has turned the practice into a report, and a report is the thing chapter 12 warns against: feedback routed to one person rather than to the people who act.

The steward. The person who runs the practice week to week. The steward keeps the pages (the maps, tests, audits and ledgers), convenes the sessions, and holds the predictions so that nobody can rewrite them later. The steward should be trusted across the organisation, curious, and senior enough to be heard, but need not be an executive. In Larkfield, the composite of chapter 7, this role might be held by the head of support.

The mapping group. Five to nine people drawn from different parts of the organisation and different levels of it, chosen because between them they see most of how the system actually works. At least two should be people whose daily work puts them in direct contact with customers, suppliers or the physical operation. The group meets for the mapping sessions and for the reviews.

The probe owners. One person for each probe, who designs it with the steward, runs it, records what happens and brings the result to the review. A probe owner should be the person closest to the work the probe touches.

The second reader. One person, or a small rotating group, whose job is to read the findings and try to falsify them. This is the role most often omitted and the one that most protects the practice from its own en-

thusiasm. The second reader asks of every claimed loop: what is the evidence it is running? Of every claimed lock-in: has it been priced? Of every probe result: was the prediction written before or after?



Before week one: three decisions

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER EIGHT

Before week one: three decisions

The sponsor makes three decisions before anything begins.

Which system. The whole organisation, a division, a product line or a single process. The smaller the system, the faster the first cycle; the larger, the more likely the most important loops are inside it. A reasonable rule: choose the smallest system that contains the decision currently on the sponsor's desk.

Which question. The practice works best with a real decision in view. "Should we enter this market?" "Why has retention fallen?" "What should we do about this ageing system?" A practice with no question tends to produce a beautiful map and no change.

What will be shared, and with whom. The mapping will surface things that are uncomfortable: arrangements a past leader chose, loops that carry extraction, places where information is held back. Decide in advance that the maps will be read by everyone in the mapping group, that the findings will be shared with the people who work in the loops they describe, and that nobody will be blamed for what history chose. Chapter 4's argument about emergence without a villain is the governing principle here: if the map becomes a way of finding someone to blame, people will stop telling the truth to it.

Weeks 1 to 2: the first map

What happens. The steward and the mapping group fill in the First Map of Loops from chapter 2. They begin alone, each person writing their own map in pencil for one hour, and then meet for half a day to compare.

The finer points.

- Begin alone. If the group maps together from the start, the most senior or most confident voice will set the shape and the others will fill it in. Separate maps, compared, show where people see the system differently, and those differences are often the most useful finding of the first fortnight.
- Write every loop as a sentence: *the more A, the more B, the more A*. A loop that cannot be written this way is usually a list of things that are related, not a loop.
- For each reinforcing loop, ask for one piece of evidence that it is running: a figure, with its source. Loops without evidence are hypotheses; mark them with a question mark, as the founder in chapter 7 did with her hiring loop.
- Limit the first map to five reinforcing loops, five lock-ins and three balancing loops. The temptation to map everything is strong and the result is a diagram nobody uses.

The decision at the end of week 2. The sponsor reads the combined map and chooses: one reinforcing loop to audit, and up to three suspected lock-ins to test. The choice should favour the loops closest to the question on the desk.

The mistake to catch. A map that describes the organisation as its leaders wish it were. The second reader's question for week 2 is: which of these loops would the people at the working level recognise? If the answer is "few", the map needs another session with them in the room.



Weeks 3 to 4: the lock-in tests

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER EIGHT

Weeks 3 to 4: the lock-in tests

What happens. For each suspected lock-in, the steward and the relevant probe owner complete the Lock-In Test from chapter 3: the arrangement, the alternative, the grade of evidence, the price of switching and the loops holding it in place.

The finer points.

- Grade the evidence honestly. One anecdote, several anecdotes, a measured comparison, several measured comparisons. Most suspected lock-ins rest on anecdotes. That is not a reason to ignore them; it is a reason to probe rather than switch.
- Price the switch in four currencies: money, time, relearning and lost compatibility with others. The last is the one most often missed, because it falls on people outside the room: customers who would have to change their habits, partners whose systems connect to yours.
- Name who would disagree with switching, and write their best argument. If nobody can state it, the test is not finished.

The decision at the end of week 4. For each arrangement, one of three verdicts. *A habit*: move it now, because it costs little, and because a habit moved is evidence to everyone that the organisation can move. *Switch*: the evidence is strong and the switch is cheap; name an owner and a date. *A true lock-in*: the evidence is strong and the switch is costly; it goes forward to the probe stage, and the loops holding it are named so that two can be changed together.

The mistake to catch. Funding a large rebuild of a true lock-in at this stage. Chapter 7's composite firm faced this choice with its old integration and took the slower route. The living-systems reading suggests redirecting the loops that feed a lock-in before replacing the arrangement itself.



Weeks 5 to 6: the loop audit

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER EIGHT

Weeks 5 to 6: the loop audit

What happens. The steward and the mapping group complete the Loop Audit from chapter 4 on the reinforcing loop the sponsor chose: who acts in it, who receives the consequence, how long the consequence takes to arrive, and whether the person acting ever sees it.

The finer points.

- Go and ask. The audit cannot be completed from a desk. For each loop, at least two people who act in it should be asked, in their own words, what they know about its consequences. In Larkfield, the engineers did not know which features clinic managers valued; the only way to learn this was to ask both.
- Distinguish severed loops from slow ones. A loop in which the consequence arrives after a year is not severed, but it is too slow to steer by. A loop in which the consequence never reaches the actor is severed. The repairs differ: a slow loop needs a faster proxy; a severed loop needs a new route.
- Find the virtuous cycles already alive. The Luminous 100's reading of systems thinking asks this directly, and it is the half of the audit most often skipped. A loop that is already generating trust or learning is usually cheaper to feed than a broken loop is to repair.

The decision at the end of week 6. The sponsor chooses one severed loop to suture during the first probe cycle, and the route by which the consequence will travel back. The route must end with the people who act, not with a manager.

The mistake to catch. Repairing the loop by sending the information to the leadership team. Chapter 12 names this as one of the three common mistakes of the first month: a new dashboard on the leadership agenda restores the loop for the people at that table and leaves it severed for everyone else.

Weeks 7 to 8: designing the probes

What happens. The steward and the probe owners design two or three probes, using section 5 of the Loop Ledger from chapter 7. Each probe is small, safe and informative. Each has a prediction written before it begins and a stop condition agreed in advance.

The finer points.

- *Small* means it can be run within existing budgets and authority, or close to them. *Safe* means that if it fails, nothing important is damaged and nobody is harmed. *Informative* means that whichever way it comes out, the organisation learns something it did not know.
- Write the prediction as a number where possible, with a baseline. "At least six new small clinics in twelve weeks, against two in the same period last year" is a prediction. "Increased interest from small clinics" is not.
- Write the stop condition as a trigger, not a judgement. "If advance cash falls by more than five per cent" is a trigger. "If it isn't working" is a judgement, and judgements made mid-probe tend to favour continuing.
- Ask the El Farol question of each probe. If competitors or partners noticed and responded, what would they do, and would it change the result?

- Include at least one probe that tests the suture chosen in week 6. Restoring a feedback loop is itself an intervention whose effect should be predicted and watched.

The decision at the end of week 8. The sponsor approves the probes, their predictions and their stop conditions, and signs the sheet. The steward keeps the signed copy. Nobody, including the sponsor, may change a prediction after this point.

The mistake to catch. A pilot disguised as a probe. Chapter 12 gives the test: if the team cannot say in advance what result would make it stop, it is running a pilot.

Weeks 9 to 20: the first cycle

What happens. Twelve weeks of running the probes. The steward convenes the probe owners for thirty minutes every fortnight; the mapping group meets once, at week 15, to review the map in the light of what the probes are showing.

The finer points.

- Record what happens weekly, in the probe owner's words, before anyone summarises it. Summaries made later tend to fit the outcome.
- Watch for stop conditions and act on them when they are met. A probe stopped on its trigger is a success of the practice, not a failure of the probe.
- Keep the suture running and visible. If the consequence that was supposed to travel back to the people acting stops arriving, find out why within a week. Restored loops tend to decay unless someone tends them.
- Use the Monday Sheet from chapter 12 as the steward's weekly page.

The decision at week 15. The mapping group reviews the map. Have any loops appeared that were not on it? Has any lock-in turned out to be a habit? The sponsor may add one new probe at this point if the review shows a gap, but may not change the existing ones.

The mistake to catch. Probe drift: the gradual expansion of a probe's scope as it begins to succeed, until it is no longer small or safe and its result can no longer be read against its prediction. If a probe is succeeding, let it finish, then amplify it in the next cycle.

Weeks 21 to 22: reading the results

What happens. Each probe owner brings the result to a half-day review with the sponsor, the steward, the mapping group and the second reader. Section 6 of the Loop Ledger is completed: predicted, happened, and a verdict of amplify, adjust or stop.

The finer points.

- Read the predictions aloud before reading the results. This is the single most protective habit in the practice, because it makes rewriting history visible.
- Let the second reader speak before the sponsor. The second reader's task is to find the reasons the result might not mean what it seems to: a confounding event, a selection effect, a measure that moved because people knew it was being watched. Chapter 9 gives the questions.
- Ask what the system showed that the week-one map did not. This is often the most valuable line of the whole review.

The decision at the end of week 22. For each probe: *amplify* (extend it, with a new prediction, in the second cycle), *adjust* (change it and run it again) or *stop* (record what was learned and end it). And one decision about the map: what must be redrawn?

The mistake to catch. Treating a probe that failed clearly as a waste. A clear failure, read honestly, often tells the organisation more than a modest success.



Weeks 23 to 26: the second map, and the decision on the desk

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER EIGHT

Weeks 23 to 26: the second map, and the decision on the desk

What happens. The mapping group redraws the First Map of Loops with what the first cycle taught. The sponsor takes the question that began the practice, the decision on the desk, and decides it, with the map, the tests, the audit and the probe results in front of them.

The finer points.

- Compare the second map with the first. The differences are a record of what the organisation learned about itself, and they should be kept. A year of maps, read in order, is the most honest account an organisation has of how it actually changed.
- Write the decision with the reasoning, including the El Farol reasoning: what others are likely to do, and where that leaves room.
- Set the date for the next cycle. The practice is meant to recur. A second cycle runs faster than the first, because the map exists and the habits have begun.

The decision at the end of week 26. The decision on the desk, and the date of the next cycle.



The rhythm, on one page

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER EIGHT

The rhythm, on one page

WEEKS	STAGE	WHO	OUTPUT	THE DECISION
0	Set-up	Sponsor	System, question, sharing rule	Three set-up decisions
1–2	First map	Steward, mapping group	First Map of Loops	Which loop to audit; which lock-ins to test
3–4		Steward, probe owners	Lock-In Tests	Habit / switch / true lock-in

WEEKS	STAGE	WHO	OUTPUT	THE DECISION
	Lock-in tests			
5–6	Loop audit	Steward, mapping group	Loop Audit	Which loop to suture, by what route
7–8	Probe design	Steward, probe owners	Probes with predictions	Approve and sign
9–20	First cycle	Probe owners, steward	Weekly records; Monday Sheets	Week 15: one probe may be added
21–22	Reading	All, second reader first	Loop Ledger section 6	Amplify / adjust / stop
23–26	Second map	Mapping group, sponsor	Second map; decision	The decision on the desk; next cycle

Time and cost, honestly

The practice costs time, and it is fair to say how much. In a firm of about sixty people, as in the composite of chapter 7, a reasonable estimate is two half-day sessions for the mapping group, about a day of the steward's time each week through the first eight weeks and half a day a week thereafter, and whatever the probes themselves cost, which by design should be small. These are this book's estimates, not measurements from a study, and a larger organisation will need more.

The usual objection is that leaders do not have the time. Two observations answer it, though neither settles it. The first is that most of the time spent is spent on the decision that was already on the desk; the practice changes how that time is used more than it adds to it. The second is that the costliest mistakes in most organisations, the large rebuild of a lock-in that should have been loosened slowly, the market entry that every competitor made at once, the loop that decayed for a year before anyone noticed, cost far more than a quarter's attention.

Scaling up and down

In a small organisation (under twenty people), the sponsor and steward may be the same person, the mapping group may be everyone, and the cycle may be eight weeks rather than twelve. Keep the second reader; in a small organisation it may be someone outside it, a trusted peer or adviser.

In a large organisation, run the practice first in one division with a real decision in view. Resist the temptation to launch it everywhere at once; a practice launched as a programme tends to become a reporting requirement. Let the second division ask for it after seeing the first.

In a network of organisations (a supply chain, an association, a cluster of firms like those in chapters 5 and 6), the sponsor role is usually shared, and the mapping group should include people from several organisations. The Cluster Reading from chapter 5 and the Shock Rehearsal from chapter 6 are the natural starting tools.



A word about the people in the loops

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER EIGHT

A word about the people in the loops

Santa Anna's book warns that the transition from hierarchy to distributed authority is developmental rather than structural, and the warning applies here. Restoring feedback to people who have not been given it before can be uncomfortable for them. They may not trust it at first, may not know what to do with it, or may hear it as blame. The practice should be introduced to them before it begins, with the sharing rule stated plainly, and the first feedback they receive should, where possible, include something that is working. The Luminous 100's reading of appreciative inquiry, and the habit of looking first for the virtuous cycles already alive, are the practical form of this care.

The Tool: The Twenty-Six Week Plan

The whole field manual on one page, to be filled in by the steward at week 0 and kept beside the Monday Sheet.

Set-up (week 0)

System:

Question on the desk:

Sponsor: · Steward: · Second reader:

Mapping group (five to nine names, at least two close to the work):

Sharing rule, in one sentence:

Weeks 1–2 · First map

Dates of individual mapping and of the comparison session:

Loop chosen for audit:

Lock-ins chosen for testing (up to three):

Weeks 3–4 · Lock-in tests

ARRANGEMENT	EVIDENCE GRADE	SWITCHING COST	VERDICT	OWNER, DATE

Weeks 5–6 · Loop audit

Severed loop to suture:

Route the consequence will travel, ending with:

Weeks 7–8 · Probes (signed by sponsor on)

PROBE	OWNER	PREDICTION, WITH BASELINE	STOP TRIGGER	EL FAROL NOTE

Weeks 9–20 · First cycle

Fortnightly review dates:

Week 15 map review date:

Weeks 21–22 · Reading

Review date:

- Predictions read aloud first: yes / no

Second reader's three strongest objections:

Weeks 23–26 · Second map and decision

What the second map shows that the first did not:

The decision on the desk:

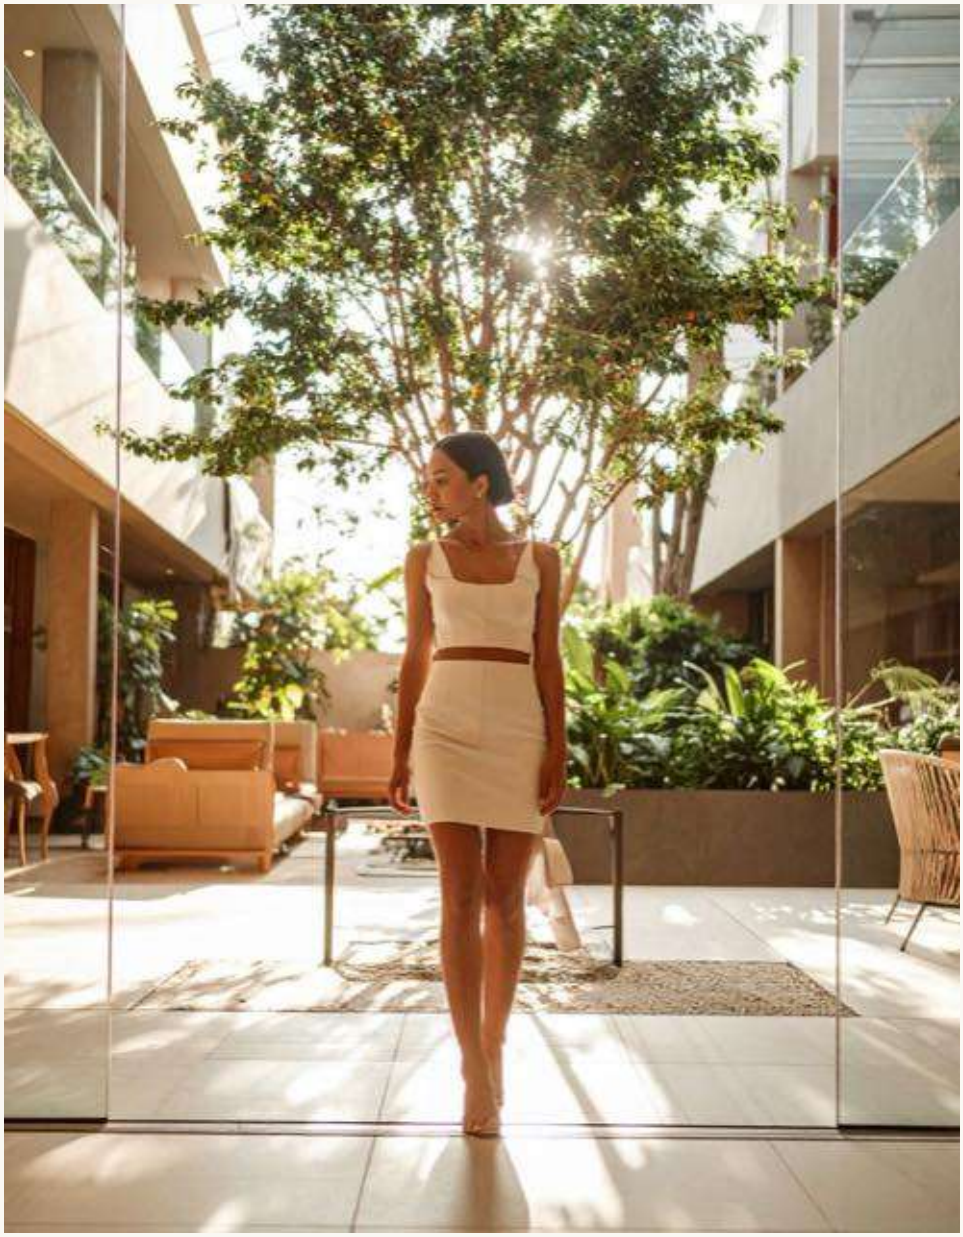
Date of the next cycle:

A note on method

The roles, timings and sequence in this chapter are this book's recommendations, developed from the instruments of chapters 2 to 7 and the composite of chapter 7. They are not the findings of a controlled study, and the time estimates are illustrative. The warning about developmental transition paraphrases Santa Anna, A., *Ecstatic Economics* (Luminous Prosperity, 2026), chapter 4 ("The Vitality-Centered Workplace", on self-management and distributed authority), quoted in part in chapter 4 of this book.

CHAPTER NINE

Knowing Whether It Worked



Knowing Whether It Worked

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER NINE

The measure that moves

There is a problem at the centre of measuring anything in a complex adaptive system, and it is worth stating before any advice is given. The people inside the system can see the measure. They adapt to it. And once they adapt to it, the measure no longer tells you what it told you before.

Two social scientists stated this independently in the 1970s, and their formulations are still the clearest. The psychologist Donald Campbell, in a 1979 paper in *Evaluation and Program Planning* titled "Assessing the Impact of Planned Social Change", put it as a law of social indicators: the more any quantitative indicator is used for social decision-making, the more subject it will be to corruption pressures, and the more apt it will be to distort and corrupt the social processes it was meant to monitor.[51] The economist Charles Goodhart, writing about monetary policy at the Bank of England a few years earlier, made a narrower observation that has since been generalised under his name: that a statistical regularity tends to collapse once pressure is placed on it for control purposes.[52]

Neither of them was writing about complexity economics, but both were describing the El Farol problem from chapter 2 in another setting. A forecast that everyone uses defeats itself. A measure that everyone is judged by changes what everyone does. In a system of agents who learn, measurement is not observation from outside. It is one more loop.

This chapter is about how to measure anyway, honestly, in a way that helps the practice of chapter 8 learn rather than fool itself. It begins with Santa Anna's book, because his critique of the most famous measure in economics is a case study in exactly this problem.



What the largest measure misses

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER NINE

What the largest measure misses

Ecstatic Economics opens its diagnosis of the extractive economy with a criticism of gross domestic product, and it makes the point with two short examples that are hard to forget:

> *A car accident generates GDP.*[53]

And, a few paragraphs later:

> *A forest standing contributes zero. Cut it down, and GDP increases.*[53]

The point is not that GDP is a bad statistic. Its designers were clear about what it does and does not count. The point is about what happens when a measure built for one purpose becomes the target of a whole society's decisions. A measure that counts the flow of market transactions and is blind to the stock of living capital will, once it is the thing everyone is judged by, reward the conversion of living capital into transactions. The forest is cut not because anyone wants it gone but because the measure sees it only when it falls.

Every organisation has its own version of this. A support team judged on tickets closed per hour learns to close tickets quickly rather than to solve the problems behind them, and the number of tickets rises. A sales team judged on new contracts signed learns to sign contracts that do not renew. A household that judges its month by the balance on one account learns to move money out of sight rather than to spend less. In each case the measure is accurate, and in each case it is steering the system toward something nobody wanted.

The book is also honest about the limits of the alternative. Writing about national wellbeing budgets, which measure health, connection and environmental quality alongside output, it says plainly:

> *The framework is necessary but not sufficient.*[54]

A better measure does not, on its own, change the loops that produce the outcomes. It changes what can be seen. The practice of this book needs both.

Five principles for measuring a living system

The principles below are this book's synthesis of Campbell's warning, the living-systems reading and the practice of chapter 8. They are not a standard, and they are not the only way to proceed. They have the merit of being usable by a leader without a statistician, and of failing visibly when they fail.

1 · Predict first

The single most protective habit in measurement is writing down, before acting, what you expect to see and what would prove you wrong. Chapter 8 makes this a rule for every probe. The reason is that complex systems are unusually generous to hindsight. Almost any result can be explained afterward by a story about loops, which is exactly the weakness the critics of chapter 3 identify. A prediction written in advance and kept by someone else is the only reliable defence.

A prediction should name a measure, a baseline, a direction, a size and a time. "At least six new small clinics in twelve weeks, against two in the same regions over the same period last year" names all five. If one of the five is missing, the prediction can usually be satisfied by whatever happens.

2 · Measure the loop, not only its output

A reinforcing loop has several links, and the output is only the last. Larkfield's referral loop, in chapter 7, produced new customers; that was its output. But the loop ran through several earlier links: clinics using the product, clinic managers finding something worth praising, managers meeting colleagues at regional events, colleagues trying the product. A leader who measures only new customers will learn that the loop has weakened months after it began to weaken. A leader who also measures one earlier link, for instance how many existing customers would recommend the product and what they would say, will see the change sooner.

The general rule: for each loop the practice cares about, choose one measure near its end and one near its beginning. The early measure gives warning; the late one confirms.

3 · Pair every measure with its shadow

Campbell's law predicts that any measure used for decisions will be pushed on. The defence is to pair each measure with a second that would move the wrong way if the first were being gamed. Tickets closed per hour, paired with the share of tickets reopened within a week. New contracts signed, paired with first-year renewal. Monthly spending, paired with the balance of money owed. Neither number on its own is trustworthy. Together they are much harder to fool, because improving one by gaming it tends to worsen the other.

A useful test for any proposed measure is to ask: if I were judged by this number and wanted to look good without doing the work, what would I do? The answer names the shadow measure.

4 · Bring the measure to the people in the loop

Chapter 4 argued that the house's distinctive diagnosis concerns severed feedback: consequences that never reach the people who act. Measurement is one of the main ways feedback is restored or kept severed. A measure that is collected from the people in a loop, summarised by an analyst and reported to leadership has not restored any feedback to those people. A measure that is shown to them, in a form they can use, close enough in time to steer by, has.

Santa Anna's book, writing about the trust that wellbeing measures can build when citizens feel their whole lives are being attended to, puts the economic case in two short sentences:

> *Trust, it turns out, is extraordinarily efficient. Suspicion is extraordinarily expensive.*[54]

A measure experienced by the people in the loop as surveillance will generate suspicion, and Campbell's corruption pressures with it. A measure they experience as information for their own use is less likely to be gamed, because they are its first users.

5 · Read the distribution, not the average

Complexity economics predicts shapes, not points (chapter 3). In systems with increasing returns, outcomes tend to be unevenly distributed: a few products, customers, teams or months account for most of the result. An average hides this. A probe that produced a modest average improvement may have produced a large improvement in a few places and none in the rest, and the few places are where the next cycle should look.

The practical habit is to look at every result as a list before looking at it as a number. Which clinics came? Which teams changed? Which weeks were different? The list shows the shape; the number only summarises it.



Questions for the second reader

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER NINE

Questions for the second reader

The second reader in chapter 8's practice exists to falsify the findings. These are the questions that do most of the work, in the order to ask them.

1. Was the prediction written before the probe began? If not, set the result aside as a hypothesis for the next cycle.

2. What else changed during the probe? A new competitor, a change in the season, a staff change, a price change elsewhere. If any of these could explain the result, say how much.
3. Who chose the participants? If the people running the probe chose which customers, teams or regions took part, they may have chosen the ones most likely to succeed.
4. Did people know they were being measured? If so, the measure may have moved because of the attention rather than the intervention. That is not always bad, but it should be named.
5. What is the shadow measure doing? If the main measure improved and the shadow worsened, the improvement is suspect.
6. Is the result an average, and what is the distribution? See principle 5.
7. Would the result surprise anyone who works in the loop? If the people closest to the work are not surprised, the probe may have measured something they already knew. That is useful confirmation, but it is not new learning.

Measuring a lock-in loosening

Some of the most important changes in this book's practice are slow: a lock-in loosened over several cycles, as Larkfield set out to loosen its old integration in chapter 7. They need their own measures, because the arrangement itself will look unchanged for a long time.

The useful measure is usually dependency: how much of the system now depends on the arrangement, and is that share falling? For Larkfield, it would be the share of support tickets involving the old integration, and the share of new clinic logic written so as not to depend on it. For a household, it might be the share of monthly spending that renews automatically. For a regional network like those in chapter 5, it might be the share of a capability held in a single firm. Measure the share at the start, predict where it will be after two cycles, and read it then.



Measuring a restored loop

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER NINE

Measuring a restored loop

A suture, in this book's sense, is a route by which a consequence travels back to the people who act. Three measures tell whether one is working.

- Arrival. Is the consequence reaching the people it was meant to reach, on schedule? Count it. Restored loops decay; an arrival rate that falls is the first sign.
- Use. Are the people who receive it doing anything differently? Ask them, in their own words, and record the answers as a list.
- Effect. Is the loop the suture was meant to strengthen running more strongly? This is the slowest and least certain measure, and it should be predicted in advance like any other.

In Larkfield's case, arrival was the monthly reading of referrers' words to engineers; use was the engineers' decision to improve the boarding screen; effect, still uncertain at twelve weeks, would be whether referrals rose in the following year.

What not to measure

Measurement costs attention, and attention is the scarcest thing in most organisations. A few warnings.

Do not measure what nobody will act on. A figure that no one will use to decide anything is a cost with no return, and it adds to the noise in which useful figures are lost.

Do not build a dashboard before the practice has run once. The first cycle of chapter 8 will show which measures matter. A dashboard built before then encodes the guesses of week one.

Do not measure people when you mean to measure loops. A loop is a pattern among people; measuring individuals within it tends to produce the corruption pressures Campbell described, and to turn a shared pattern into individual blame. Chapter 4's principle of emergence without a villain applies to measurement as much as to diagnosis.

The household reading

A household needs very few measures. The Kitchen-Table Map in chapter 11 suggests four: the share of monthly income that is committed before it arrives (a lock-in measure), the amount set aside in the month (a reinforcing-loop measure), the balance owed (the shadow), and one measure that is not money at all, chosen by the household, of whether its members feel the month went well. The last is the household's version of the wellbeing budget. It cannot be gamed without the household noticing, because the people measured are the only people reading it.



The Tool: The Honest Dashboard

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER NINE

The Tool: The Honest Dashboard

One page, completed by the steward with the second reader. No more than six measures in total.

1 · The loops being measured

Loop 1 (from the Loop Ledger):

Loop 2:

2 · Each loop, measured near its beginning and near its end

LOOP	EARLY MEASURE (A LINK NEAR THE START)	LATE MEASURE (THE OUTPUT)	SOURCE OF EACH FIGURE	HOW OFTEN READ

3 · Shadows

MEASURE	IF I GAMED THIS, WHAT WOULD I DO?	THE SHADOW MEASURE THAT WOULD REVEAL IT

4 · Who sees each figure

MEASURE	SEEN BY THE PEOPLE IN THE LOOP? YES / NO	HOW SOON AFTER IT HAPPENS?	IN WHAT FORM?

5 · Lock-in dependency

Arrangement being loosened:

Share of the system that depends on it today: · source:

Predicted share after two cycles:

6 · Restored loop

Arrival: what is supposed to reach whom, how often? · actual arrival so far:

Use: what have the recipients done differently?

Effect: predicted · observed

7 · The second reader's seven questions Tick each one asked of the latest result: prediction first ☐ · what else changed ☐ · who chose participants ☐ · knew they were measured ☐ · shadow measure ☐ · distribution ☐ · surprising to the people in the loop ☐

A note on method

The five principles and the second reader's questions are this book's synthesis and are not drawn from a single source.

CHAPTER TEN

The Hard Conversations



The Hard Conversations

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TEN

Why the conversations are the work

A map of loops changes nothing until someone acts on it, and in most organisations acting on it requires persuading people who did not draw it. A board that approves budgets. A partner whose systems are tied to yours. A union or staff council that represents the people whose work will change. A family that shares the household's money. A sceptic who has seen fashionable ideas come and go.

Each of these conversations is itself a loop. What you say changes what they expect; what they expect changes what they do; what they do changes what you can say next time. A conversation that goes badly early can lock in a pattern of mistrust that costs more than any lock-in on the map. A conversation that goes well can start a reinforcing loop of trust that makes the next five easier.

This chapter gives the reasoning and then a script for each of six conversations. The scripts are starting points, written in plain words, to be changed to fit the people in the room. They share three features: each begins with what is working, each states what is not known, and each asks for something small.

The principle beneath every script

Amy Edmondson's study of work teams in hospitals and manufacturing, published in *Administrative Science Quarterly* in 1999 as "Psychological Safety and Learning Behavior in Work Teams", found that teams whose members believed it was safe to take interpersonal risks, to admit mistakes, ask questions and raise concerns, engaged more in the behaviours by which teams learn.[55] Her study was of teams, not boards or families, and it measured learning behaviour rather than long-run results. But it states the condition every conversation in this chapter is trying to create. A complex adaptive system learns by feedback, and people will not give honest feedback where it is unsafe to do so.

Santa Anna's book makes the same point about organisations in general, in its list of the conditions for vitality: feedback richness requires, in its words,

> *multiple channels, rapid loops, and the psychological safety to deliver and receive honest information without punishment.*[56]

The scripts below try to make each conversation a place where honest information can travel in both directions.

The Luminous 100's reading of stakeholder mapping gives the posture. The classical tool sorts stakeholders by power and interest and assigns each a strategy for managing them. The upgrade asks instead what each party brings and needs, and imagines a change designed as

> *a web of mutual contribution rather than a campaign to manage resistance and leverage support?*[57]

That is the question behind every script that follows. Not how to get these people to agree, but what each party brings, what each needs, and what the change would look like if it served both.



1 · The board

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TEN

1 · The board

What the board is worried about. Risk, return and accountability. A board hearing about complexity economics for the first time may suspect it is being asked to accept vagueness in place of a plan: "it's complex" as a reason not to commit to targets.

What helps. Lead with a decision, not a philosophy. Show that the practice produces sharper commitments, not looser ones: predictions with baselines, stop conditions, dates. Show that it costs little. Name what it cannot do.

The script.

> *"We have a decision in front of us about [the question on the desk]. Before we commit the full budget, I'd like to run a twelve-week cycle that tests the three assumptions the plan depends on most. Here is what we already know is working: [one reinforcing loop, with its evidence]. We'd like to understand it well enough to feed it deliberately. Here is what we don't know: [the key uncertainty]. Each probe has a written prediction and a stop condition, which I'll share with you now, signed, so that nobody can reinterpret the results afterward. The cost is [the figure], about [the share] of what the full plan would commit. At week twelve I'll bring you the results against the predictions, and a recommendation. What I'm asking for today is approval for the cycle, not for the plan."*

The finer point. Give the board the predictions in writing before the cycle begins. Nothing builds a board's confidence faster than a leader who commits in advance to what would prove her wrong, and nothing protects the practice better.

When the board asks about sharing gains. If the map has shown that the people who produce improvements never see any of the gain (chapter 12 discusses this), the board may hear a proposal to change that. The evidence on gainsharing and employee ownership is examined in two companion books in this series (gainsharing, employee ownership), and the honest answer at this stage is that the evidence is encouraging under certain conditions and that the design matters. A board should be offered a probe, not a policy.



2 · The partner

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TEN

2 · The partner

What the partner is worried about. That a change on your side will cost them: a new integration to build, a changed process, a renegotiation. Partners are inside your loops, and the lost-compatibility cost in chapter 3's Lock-In Test falls on them.

What helps. Bring the map, not the conclusion. Show them the loops that run between you, and ask what they see that you don't. Partners often know things about your system that you cannot see from inside it.

The script.

> *"We've been mapping how our business actually works, and several of the most important loops run through you. I'd like to show you what we've drawn and ask where it's wrong. One thing we've found is [a loop that benefits both parties]. We'd like to strengthen it. Another is [a lock-in that affects both]. We think it may be costing both of us, but we're not sure, and we don't want to change anything that falls on you without understanding your side. Would you be willing to spend an hour on the Lock-In Test with us, on this one arrangement? If the evidence says it's a sensible habit, we'll leave it. If it's a real lock-in, we'd want to change it together, and only at a pace that works for you."*

The finer point. The Aisin case in chapter 6 is worth recalling. The partners who reorganised around the fire could do so because they had solved problems together for years. A hard conversation with a partner is also an occasion to practise exactly that.

3 · The union or staff council

What they are worried about. That "mapping loops" and "running probes" is a new language for intensifying work or cutting jobs. That restored feedback means more surveillance. That the people whose work changes will bear the risk while others take the gain.

What helps. Name these fears before they are raised. State the sharing rule from chapter 8 plainly. Invite representatives into the mapping group as full members, not observers. Commit that no probe will reduce anyone's pay or job security, and that measures of loops will not be used to judge individuals (chapter 9).

The script.

> *"We're going to spend the next six months trying to understand how the organisation actually works, as opposed to how the organisation chart says it works. The people who know that best are the people doing the work, so we'd like two of your members in the group that does the mapping. I want to be clear about three things. First, nothing that comes out of this will be used to judge individuals. We're mapping patterns, not people. Second, no probe in this cycle will reduce anyone's pay, hours or job security; if a probe would, it isn't small or safe enough and we won't run it. Third, the findings will be shared with everyone whose work they describe, unedited. One of the things we expect to find is places where people's work produces improvements that they never see the results of. If we find that, we'd like to talk with you about how those results could come back to the people who produced them."*

The finer point. The third commitment is the one most often broken, usually not from bad faith but because a manager edits the findings to be "constructive". Agree in advance who will check that the findings were shared as written.

4 · The family

What they are worried about. That a conversation about money is a conversation about blame. That changing a habit means losing something they value. That one person will make the decisions.

What helps. Start with what is going well. Draw the map together, each person first on their own sheet. Ask for one twelve-week probe, chosen together, with a prediction everyone can see.

The script.

> *"I've been reading about how money and habits in a household reinforce each other over time, and I'd like to try something with all of us. It isn't about cutting back or about anyone having done anything wrong. Each of us draws, on our own, the things in our household that seem to grow by themselves, good or bad, and the things we keep doing that none of us would choose again if we were starting fresh. Then we compare. At the end we pick one small change to try for twelve weeks. We write down now what we think will happen. In twelve weeks we look again and decide whether to keep it."*

The finer point. Choose the first probe to be something everyone will notice going well: a weekly look at the household's money together instead of a monthly surprise, for example. Chapter 11 develops the household practice in full.



5 · The sceptic

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TEN

5 · The sceptic

What they are worried about. That this is another management fashion. Often they have good reasons: they have watched several arrive and leave, each with its vocabulary, its consultants and its slides.

What helps. Agree with them where they are right. Complexity economics has serious critics (chapter 3), and some of them are right. Ask the sceptic to be the second reader.

Santa Anna's book names the failure mode a leader should avoid in this conversation, which is the opposite of the one the sceptic fears. Writing about how alternative arrangements are judged, it describes the way any flaw becomes a reason to dismiss the whole:

> *The moment an alternative economic structure reveals any flaw, any tension, any compromise, the scarcity mind seizes on it*[58]

The sceptic is useful precisely when they resist that move as well: when they hold the practice to its own standard rather than to a standard of perfection.

The script.

> *"You've seen a lot of these come and go, and I'd like your help making sure this one isn't another. The ideas behind it have real critics, and I've read them. Some of the famous examples have been challenged, the models can be tuned to fit almost anything, and 'it's complex' gets used as an excuse. Those are fair points. What I'm proposing is narrower than the theory. A map, a few tests, and some small probes with predictions written down in advance and stop conditions agreed before we start. I'd like you to be the person who reads the results and tries to prove they don't mean what we think. If you can, we'll stop."*

The finer point. Mean the last sentence. A sceptic who is given the second reader's role and then overruled when they find something will not be persuaded again.

6 · Yourself

The last hard conversation is the one a leader has with herself, usually at the point where a probe has failed, a lock-in has turned out to be harder than it looked, or the map has shown a loop she built that carries extraction.

The Luminous 100's reading of polarity management, which treats apparent opposites in organisational life as the two movements of a single system, offers an image for this conversation. It describes the creative power of a polarity as living in the movement between its poles,

> *the same way breathing's power lives not in inhale or exhale but in the rhythm itself.*[59]

The practice of this book has the same rhythm: committing and testing, acting and reading, confidence in the prediction and readiness to be wrong. A leader who can hold both does not need to win every probe. She needs to keep breathing.

The script, for writing in the steward's book.

> *"What I predicted: _____. What happened: _____. What I learned about the system that I did not know: _____. What I learned about how I read it: _____. What I will do differently in the next cycle: _____."*



Common mistakes in all six

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TEN

Common mistakes in all six

- Leading with the theory. Nobody needs to hear about nonergodic processes to agree to a twelve-week probe. Lead with the decision and the loop.

- Promising outcomes. The practice cannot promise what any organisation or household will find. Promise the process: predictions written, results read honestly, findings shared.
- Asking for too much. Every script ends with a small request. The first agreement, kept, starts a reinforcing loop of trust; an over-large request, refused, starts the opposite one.
- Not closing the loop. Each conversation creates an expectation. Return to the people at the date you named, with the result, whether it went well or not.

The Tool: The Conversation Cards

One card per conversation. Fill it in before the conversation and again after it.

Before

Who:

What they are most likely worried about:

What is working, that I will name first:

What I do not know, that I will say plainly:

What each party brings, and what each needs:

The one small thing I am asking for:

What I will commit to in return (a date to report back, a sharing rule, a stop condition):

During

What they told me that was not on my map:

Their strongest objection, in their words:

After

What was agreed:

The date I will return with the result:

Did this conversation start a loop of trust or of mistrust? What would strengthen or repair it?

A note on method

The scripts are this book's and are offered as starting points. They are not legal, employment or financial advice, and any change affecting pay, contracts or employment should be taken with appropriate advice in the jurisdiction concerned.

CHAPTER ELEVEN

The Household as a Living System



The Household as a Living System

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ELEVEN

The smallest economy

A household is an economy in miniature, and a complex one. Two or three or six people, each with their own habits and hopes, sharing money, time, space and attention. Commitments made years ago (a mortgage, a career, a city, a school) shape what is possible now. Habits reinforce themselves: a weekly routine becomes easier each week it is kept, a debt becomes heavier each month it is carried. Small, accidental early choices, such as which job was taken first or which neighbourhood rented in, compound into a life.

Everything in the earlier chapters applies here, at a scale where one person can see most of it. This chapter runs the whole practice for a household: the map, the lock-ins, the audit, the probes and the measures. It is written for whoever is responsible for a household's money and work, alone or with others. It is not financial advice, and it says nothing about what anyone should buy, sell or invest in. It is about seeing the household as a system, and changing it through small, well-recorded steps.



The power of the default

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ELEVEN

The power of the default

The clearest public evidence that household finances behave like a complex adaptive system, with lock-in and path dependence, comes from a single, much-cited study of retirement saving.

In 2001 Brigitte Madrian and Dennis Shea published a paper in *The Quarterly Journal of Economics* titled "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior".[60] They studied a large American company that changed the default for its retirement savings plan. Before the change, new employees had to choose to join; after it, new employees were enrolled automatically, at a set contribution

rate and in a set investment, unless they chose to leave. Nothing else about the plan changed. The money on offer was the same, the choices available were the same, and anyone could opt in or out at any time.

The result was a very large rise in participation among employees hired under the automatic default. And a striking share of those employees stayed with the default contribution rate and the default investment, even though few would likely have chosen exactly those settings on their own. The default, a small administrative arrangement, had become the arrangement that history chose, and it held.

This is Arthur's lock-in at the scale of one person. The arrangement was not chosen on its merits. It persisted because changing it required effort, attention and a decision, and the cost of those, small as it was, was enough. The finding cuts both ways, which is the useful thing about it. Defaults can trap people in arrangements they would not choose. Defaults can also carry good practices forward without effort.

Richard Thaler and Shlomo Benartzi made the second half of that point into a design. Their 2004 paper in the *Journal of Political Economy*, "Save More Tomorrow", described a programme in which employees committed in advance to increase their saving rate at the time of future pay rises.[61] The design used inertia deliberately: having committed once, most participants stayed in the programme, and their saving rose with each rise in pay. The paper reports results from a small number of firms, and it is a study of one designed intervention, not a general law. But the mechanism is the one this book has followed throughout: a reinforcing loop, started by a small early choice, and then left to compound.

The household's loops

A household has the same three kinds of loop as a firm.

Reinforcing loops grow by themselves. Some carry the household well: a habit of setting money aside, which becomes easier each month it is kept and builds a reserve that reduces the next month's anxiety; a skill practised over years, which raises earnings, which buys time to practise further; a set of neighbours with whom help has flowed both ways, which makes the next request for help easier. Others carry it badly: borrowing to cover a shortfall, which raises next month's costs, which makes the next shortfall likelier; overwork to cover rising commitments, which leaves less time to manage those commitments well.

Lock-ins are arrangements everything else now depends on. A housing cost set years ago. A phone contract renewed by default. A set of subscriptions nobody reviews. A commute fixed by a choice of home. A family expectation about who earns and who cares. Some are sensible habits; some are traps; the Lock-In Test from chapter 3 tells the difference.

Balancing loops correct. A budget that is actually read each week is one: overspending in one week shows up, and the next week adjusts. A household conversation about money that happens regularly is another. Where balancing loops are slow or absent, a household drifts, and the drift is noticed only at a crisis.



The severed loop at home

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ELEVEN

The severed loop at home

Chapter 4 argued that the house's distinctive diagnosis is the severed feedback loop: the consequence of an action never reaching the person who acted. Households have their own versions, and they are very common.

Spending on a card whose statement arrives a month later is a slow loop. Automatic payments are a severed one: the money leaves, and nobody in the household feels it leave. A financial decision made by one partner and borne by another is severed as well. None of these is wrong in itself, and some are conveni-

ences worth having. But each one means that the household is acting without the information it would need to steer.

Santa Anna's book proposes a weekly practice for one person that is, in this book's terms, a suture. Once a week, it suggests, review the economic events of the past seven days and ask of each one:

> *Did I feel more alive or less alive afterward?*[62]

The question is not a measure a bank would recognise, and it is not meant to be. It brings back to the person who acted a consequence that no statement reports: whether the household's economic life is carrying vitality or draining it. Over weeks, the book argues, this inventory builds a map of which activities serve the person and which extract from them. A complexity economist would recognise the method: a fast, felt feedback signal attached to a slow, abstract system.

Path dependence in a family

The longest loops in a household run across generations. Habits about money, work and security are learned early, often without words, and carried into adult life. Santa Anna's book discusses this at length in its chapter on the somatic dimension of economics, and it refuses the fatalism that the idea can invite:

> *This isn't determinism.*[63]

That is the complexity economist's position too. A path-dependent process remembers its history; it is not bound by it. The household that inherited a pattern of anxiety about money, or of avoiding it, can change the pattern, but it will usually need to change more than one loop at once, for the reason Gregory Unruh gave about carbon lock-in in chapter 3: a single loop changed on its own is pulled back by the others. A household that decides to look at its money weekly, and also agrees that nobody will be blamed for what the first look shows, has changed two loops together.

The kitchen-table sequence

Here is the practice of chapter 8, compressed for a household. It takes about thirteen weeks: one to map, twelve to probe.

Week 1, evening one: the map. Each adult draws their own First Map of Loops, alone, in pencil, in thirty minutes. Five things that grow by themselves, good or bad. Five arrangements nobody would choose today. Three things that correct themselves. Then compare. Where the maps differ, ask why, without deciding who is right. Older children can be included in a simpler version: what in our family gets easier the more we do it, and what do we keep doing that none of us likes?

Week 1, evening two: the tests and the audit. Take the three arrangements from the map that the household most suspects. For each, run the Lock-In Test: what is the alternative, what is the evidence it is better, what would switching cost in money, time and effort, and who would it fall on? Most will turn out to be habits that can be changed in an afternoon; change them that week. Then take the most important reinforcing loop and ask the Loop Audit's questions: who acts in it, who feels the consequence, and how soon?

Week 1, evening three: the probe. Choose one probe, small and safe, for twelve weeks. Write the prediction down where everyone can see it. Write the stop condition. Some probes worth considering for a household: a fifteen-minute weekly look at the household's money together, at the same hour; a week's pause before any purchase above an agreed amount; one automatic payment converted to a manual one, so that its leaving is felt; one hour a week spent on a skill that could raise earnings; one way, each month, that help flows toward a neighbour.

Weeks 2 to 13: the cycle. Keep the Monday Sheet from chapter 12, used for the household, once a week. Watch for the stop condition. Do not change the prediction.

Week 13: the reading. Read the prediction aloud first. Then the result. Decide together: keep, adjust or stop. Redraw the map. Set the date for the next cycle.



Measures for a household

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ELEVEN

Measures for a household

Chapter 9 recommends very few measures, each paired with a shadow. For a household, four are usually enough.

- The committed share. The share of monthly income committed before it arrives, to housing, debts, contracts and subscriptions. This is the household's lock-in measure. A falling share means the household is regaining room to adapt.

- The amount set aside. Whatever the household sets aside in the month, for any purpose. This is the reinforcing-loop measure.
- The balance owed. The shadow of the second. If the amount set aside rises while the balance owed rises faster, the loop is not what it looks like.
- The felt month. One question, answered by each adult at the end of the month, on whatever scale the household chooses: did the month go well? This is the household's own wellbeing measure, and the only one of the four that cannot be read off a statement.

None of these figures should be compared with other households. The comparison class for each is the household's own last month.



The El Farol question at home

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER ELEVEN

The El Farol question at home

Chapter 2's El Farol question asks, before acting on a forecast, what would happen if everyone held it. It has a household form, and it is worth asking before any large decision: a move, a career change, a large purchase. If everyone in your situation made the same choice for the same reason, what would happen? Would housing in the area everyone is moving to become dearer? Would the skill everyone is training in become less scarce? Would the time everyone is taking off at once leave the opportunity somewhere else?

The question does not say the choice is wrong. It asks where the room is that the crowd will leave behind, and whether the household might be better placed there.

What a household can do that a firm cannot

A household has one great advantage over a firm in this practice: the people who act and the people who feel the consequences are the same small group, sitting at the same table. The severed loops are shorter and easier to repair. The cost of a probe is lower. The map can be redrawn in an evening.

It also has one great difficulty: the conversation is closer to the bone. Money in a household is tied to love, fear, fairness and history in ways that money in a firm usually is not. Chapter 10's family script, and its principle of beginning with what is working, matter more here than anywhere.

The Tool: The Kitchen-Table Map

Three evenings, then twelve weeks. Each adult fills in sections 1 and 2 alone first.

1 · Loops (alone, thirty minutes)

- Five things in our household that grow by themselves, good or bad, each written as *the more A, the more B, the more A*:

…

…

…

…

…

Three things that correct themselves (*the more A, the less B, the less A*):

2 · Arrangements nobody would choose today (alone)

ARRANGEMENT	WHEN IT BEGAN	WHAT DEPENDS ON IT	WOULD WE CHOOSE IT TODAY?	ROUGH COST OF CHANGING IT

3 · Compare (together)

Where our maps agree:

Where they differ, and why:

4 · The severed loop

One payment, habit or decision whose consequence nobody in the household feels:

How we could bring it closer (a weekly look, a manual payment, a conversation):

5 · The probe

One small, safe change for twelve weeks:

What we predict, written now:

What would make us stop:

Start date: · reading date:

6 · The four measures

MEASURE	TODAY	WEEK 13
Share of income committed before it arrives		
Amount set aside this month		
Balance owed		
The felt month (each adult)		

7 · The El Farol question

The largest decision ahead of us:

If everyone in our situation made the same choice, what would happen?

Where is the room the crowd would leave?

A note on method

The probes and measures in this chapter are this book's suggestions. Nothing here is financial, tax or investment advice; a household making a decision about debt, saving or investment should take advice suited to its circumstances.

CHAPTER TWELVE

The Path Across the Lawn



The Path Across the Lawn

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWELVE

A line in the grass

On most university campuses there is a lawn crossed by a path nobody laid. The paved walks go round the edges at right angles, and across the middle, from the library door to the corner of the street, runs a narrow line of worn earth. Planners call it a desire line.

No committee designed it. The first student who cut across did so because it was shorter. The second followed the faint trace the first had left, because a slightly flattened line is easier to walk than unbroken grass. By the hundredth, the grass had given up, and the line had become the way everyone goes. A small advantage, gained by accident, reinforced by every walker who came after: it is W. Brian Arthur's increasing returns written on a lawn, and it is path dependence in the literal sense.

This book has followed that idea from its origin in Arthur's 1989 paper (chapter 2), through its fair critics (chapter 3), into the living-systems reading of Ammanuel Santa Anna's *Ecstatic Economics* (chapter 4), through two public cases (chapters 5 and 6), and onto the table of a composite firm (chapter 7), and the four chapters before this one turned it into practice. This chapter looks outward. Complexity economics on its own explains how patterns form. It says less about which patterns to want, how to govern them, or what to do with the gains they produce. For those questions it needs company, and the Ecstatic Economics lineage supplies it.

The book itself is direct about the shape that company should take. Writing about the practices that fail when scaled up, Santa Anna proposes not a larger version of each but a different geometry:

> *Perhaps the goal is not one giant cooperative but a network of small ones. Perhaps the goal is not a single global currency but an ecosystem of local currencies.*[64]

That sentence is where complexity economics and the rest of the lineage meet. A network of small, connected, adapting units is what a complex adaptive system is. The question is how to govern one.



Ostrom, and governance at many scales

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWELVE

Ostrom, and governance at many scales

The first join is to the work of Elinor Ostrom, which has its own book in this series.

In 2009 Ostrom became the first woman awarded the Nobel memorial prize in economic sciences. Her prize lecture, published the following year in the *American Economic Review*, carried a title that could serve for this whole series: "Beyond Markets and States: Polycentric Governance of Complex Economic Systems".[65] Ostrom had spent decades studying how communities manage shared resources such as fisheries, irrigation systems and forests, and she had found that the two standard prescriptions, private ownership or

state control, were often outperformed by a third arrangement: many overlapping centres of decision, at different scales, each able to learn and adjust, nested inside one another.

Her design principles for long-enduring commons, which the companion book takes in turn, read in this light as rules for a healthy complex system. Clear boundaries tell the agents who is inside the loop. Monitoring by the users themselves keeps feedback short. Graduated sanctions correct small deviations before they compound. Local rule-making lets each part adapt to its own conditions. And nesting allows the whole to coordinate without any single centre controlling it.

Santa Anna's book reads the same principles through the body:

> *They describe a system that is relational, responsive, graduated, and nested*[66]

The complexity economist and the house are describing one thing from two sides. A polycentric commons is a system designed so that its feedback loops stay intact. Its users see the consequences of their choices because they are the ones who monitor, and they can change the rules because they are the ones who make them. It is the opposite of the severed loop the third piece described.

The book also names what happens when that structure is removed:

> *When shared resources become private property, the relational web of mutual obligation that governed their use is severed.*[67]

Read with complexity economics, enclosure is not only a transfer of ownership. It is the removal of a set of feedback loops, and with them the system's capacity to learn.

The doughnut, and the walls of the room

The second join is to ecological economics, and in particular to Kate Raworth's doughnut, which has its own book.

Raworth's figure, in her 2017 book *Doughnut Economics*, draws two rings. The inner ring is a social foundation, the minimum every person needs; the outer ring is an ecological ceiling, the planetary limits an economy should not breach. The space between is where an economy can thrive. Complexity economics adds a warning that the doughnut's two rings are not fences the system will politely respect. An economy with strong reinforcing loops will push outward wherever a loop rewards it for doing so, and it will push quickly, because positive feedback accelerates. The ceiling has to be built into the loops themselves, as balancing feedback, or the system will cross it before anyone has agreed it should.

Santa Anna's book takes the doughnut and asks one more question of it, a question no ring can measure:

> *The question is: How alive do people feel in the economy we're creating?*[68]

This is the living-systems reading applied to the space inside the doughnut. Meeting the foundation and respecting the ceiling are necessary. What the book adds is that an economy can do both and still be organised around anxiety, and that the quality of life inside the rings is itself an emergent property of the loops people live within. Change the loops, and the felt quality of the economy changes with them.



Closing the loop on gain

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWELVE

Closing the loop on gain

The third join is to the evidence on sharing the gains an organisation creates, which this series examines in two books: the history and research on gainsharing and the evidence on employee ownership and profit sharing.

The connection is direct. Chapter 4 argued that a healthy system is one where the consequence of a choice reaches the person who made it. In most firms, the people whose daily choices create an improvement in cost, quality or service see little of the gain that follows; it travels upward and outward, and the loop that would reward and reinforce the behaviour is cut. Arrangements that return a share of measured gain to the people who produced it are, in complexity terms, feedback loops restored. They make an increasing-returns loop out of improvement itself: the more people improve the work, the more they share in it, the more they improve.

Santa Anna's book describes what changes in the people inside such a loop, writing of Mondragon's worker-owners:

> *The wage employee's neuroception is scanning for threats from the organization — layoffs, pay cuts, arbitrary management decisions. The worker-owner's neuroception is scanning for threats to the organization — and mobilizing creative energy to address them.*[69]

The companion books take the evidence seriously, including its limits. The point for this book is narrower. Complexity economics explains why such loops compound once they start; it is the gainsharing and ownership research that shows whether, and under what conditions, they do.

The fourth join, briefly, is to Karl Polanyi's argument that markets are embedded in society. Polanyi's history and Arthur's models arrive at a shared conclusion from opposite ends: the economy is not a separate machine running beside social life but a pattern that social life produces, and that produces social life in return.



What to do on Monday

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWELVE

What to do on Monday

The joins are only useful if they change what happens on an ordinary working day. Chapter 8 sets out the full sequence week by week; here it is in its shortest form, for a leader and for a household.

For a leader, in the first week. Take the First Map of Loops from chapter 2 and fill it in with two colleagues. Do not try to change anything yet. The act of naming the reinforcing loops and the lock-ins is itself a restored feedback loop: the organisation begins to see its own shape.

In the second week. Choose one suspected lock-in and run the Lock-In Test from chapter 3. Price the switch honestly. Most will turn out to be habits; move those at once, because a habit moved is evidence to everyone that the organisation can move.

In the third week. Run the Loop Audit from chapter 4 on your most valuable reinforcing loop. Ask who acts in it and who feels the consequence. If those are different people, find the shortest route by which the consequence could travel back, and open it. Read what comes back to the people who can act on it, unedited.

Over the following twelve weeks. Run two or three small probes, as the composite firm did in chapter 7, each with a prediction written beforehand and, where one can be set, a stop condition. Amplify what works. Stop what does not. Write down what the system showed you. Keep the pages. A year of them, read in order, is the closest thing an organisation has to a record of how it actually learns, and it is usually more honest than the strategy documents written over the same period.

For a household. The same sequence, at kitchen-table scale. Map the loops in the household's money, time and attention. Test the arrangements nobody would choose again. Find the spending, the commitment or the habit whose consequence arrives too late to steer by, and bring it closer: a weekly look instead of a monthly statement, a conversation instead of a direct debit. Probe one change for twelve weeks.

For both, a governing habit. Before any large decision, ask the El Farol question: if everyone around you made the same forecast, what would they do, and would it still hold? It is the shortest discipline complexity economics offers, and the one most often skipped. A forecast that would defeat itself if widely shared is not a reason to abandon the decision; it is a reason to look for the room the crowd will leave behind, as the composite firm in chapter 7 did when its competitors all turned toward the same customers.

The instruments in the Alpha Papers carry several of these steps in more depth: *Alpha Scenario Planning* for holding several futures at once, *Alpha Viable System Model* for designing an organisation whose feedback loops reach every level, and *Alpha Real Options* for pricing probes as options rather than commitments.

Three mistakes at the start

Leaders who take up this way of reading their organisation tend to make the same three mistakes in the first month. Each is worth naming, because each is easy to avoid once seen.

Mapping everything. The first map of loops invites completeness, and a team that tries to draw every loop in the organisation will produce a diagram nobody can read and nobody will use. Five reinforcing loops, five lock-ins and three balancing loops are enough for a first page. The aim is to see the shape of the system, not to model it. A map that fits on one sheet is a map that gets looked at again.

Treating the probe as a pilot. A pilot is a smaller version of a decision already made; it is designed to succeed, and its results are read generously. A probe is a question put to the system; it is designed to be in-

formative, and a probe that fails clearly is as valuable as one that succeeds. The difference shows in the paperwork. A probe has a prediction written before it begins and a stop condition agreed in advance. A pilot usually has neither. If the team cannot say in advance what result would make it stop, it is running a pilot.

Changing the loop without changing who sees it. A leader who finds a severed feedback loop often repairs it by sending the missing information to themselves: a new report, a new dashboard, a new monthly figure on the leadership agenda. The loop is restored for one person. The people whose daily choices drive it still see nothing. The repair that works is the one where the consequence reaches the person acting, in a form they can use, close enough in time to steer by. This is the point the book makes about investors who never visit the enterprise, and it applies to every manager who reads a summary of what the front line already knows.

A fourth mistake belongs to households more than to firms: waiting for a crisis to look. A family that maps its loops only when money is tight sees them under the worst conditions, when every loop looks like a threat. The same map drawn in an ordinary month shows which loops are already carrying the household well, and those are the ones worth feeding first. The companion book on regenerative economics makes the same argument at the scale of an enterprise: the work begins with what is already alive.

What this book cannot say

A last honest word. Complexity economics is a framework for understanding, and its practical recommendations, including those in this book, are reasoned from its models and from cases rather than proven by controlled trials of whole organisations. The twelve-week cycle, the tools and the sequence above are this book's advice. They are offered as a way to see more clearly, and the house makes no promise about what any particular organisation or household will find when it looks. What the research does support, across Arthur, Ostrom, Haldane and May and the others cited in this book, is the shape of the claim: that systems with intact feedback learn, and systems without it repeat themselves.

The parent book, *Ecstatic Economics*, argues the larger case from which this one draws, and the companion books named above take its other threads further.



The lawn, again

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER TWELVE

The lawn, again

Go back to the lawn. A university that wants to remove the desire line has two choices. It can fence the grass, in which case students will walk round the fence and wear a new line at its end. Or it can pave the line, and accept that the walkers knew something the planners did not.

There is a third thing a planner can do, and it is often told as a planner's custom, though no single campus is documented here. Plant the lawn, wait a season, watch where the lines appear, and only then lay the paths. Complexity economics, read through the living-systems lens, asks the same patience of anyone who

leads: watch where the walking goes, read the loops, and lay the stone where the grass has already been worn.

The Tool: The Monday Sheet

One page that gathers the book into a single weekly practice. Keep it beside the Loop Ledger from chapter 7. Fifteen minutes, once a week, at the same hour.

1 · This week's loop

The reinforcing loop I am watching this week:

What it did since last week (a figure, with its source):

- Did the people acting in it see the consequence? yes / no / partly

2 · This week's lock-in

The arrangement under test:

- Its status: habit, moved / switching, under way / true lock-in, loops being changed together

The two loops being changed together, if a lock-in: and

3 · This week's restored feedback

The consequence that travelled back this week:

To whom: · unedited? yes / no

What they did with it:

4 · The probes

PROBE	WEEK OF TWELVE	PREDICTION (WRITTEN AT START)	SO FAR	STOP CONDITION MET?

5 · The joins

Governance: who monitors this loop, and can they change its rules? (Ostrom)

Limits: what balancing loop keeps this from pushing past a boundary we care about? (the doughnut)

Gain: when this loop produces a gain, who receives it, and do the people who produced it see any of it?

Embedding: whose relationships does this loop depend on, outside the firm or household?

6 · The El Farol line

The forecast I am acting on:

What everyone else would do if they shared it:

7 · The desire line

One place this week where people went their own way, around the official path:

What it tells me about where the path should be:

A note on method

Raworth, K. (2017). *Doughnut Economics: Seven Ways to Think Like a 21st-Century Economist*. Random House Business. Described here from the book's well-known framework; no page is quoted. The desire-line practice of planting first and paving later is told as a planner's custom; no single documented campus is cited, and it should be read as an illustration rather than a case. Nothing in this book is investment advice.

CHAPTER THIRTEEN

What Grows by Itself



What Grows by Itself

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER THIRTEEN

The whole, read once

This book has followed one idea from a paper in a drawer at Stanford to a kitchen table. It is worth reading the whole of it once, from above, in the terms of the living-systems lens that Ammanuel Santa Anna's *Ecstatic Economics* develops, before putting it down.

Complexity economics, in Brian Arthur's hands and those of the people who worked beside him, made a scientific case that an economy is not a machine settling into balance but a system that keeps forming itself. Its agents learn and adapt to a world their own choices are creating. Some of their interactions reinforce themselves and grow; others correct themselves and fade. History matters, because early accidents are remembered in the structure of who does what. Arrangements nobody would choose can hold for decades, and arrangements that once seemed fixed can dissolve within a few years when the loops beneath them change.

The critics, fairly read, narrowed the claim. Inefficient lock-in is possible and probably less common than the stories suggest. Models of adaptive agents are promising and hard to validate. The theory predicts shapes rather than points. None of this weakens what remains; it makes it usable.

The living-systems reading added a question that the science, by its nature, leaves open: what each loop carries. A loop can circulate value, information and trust back to the people who act, or it can be severed, so that the people who act never feel what they cause. Santa Anna's book places its diagnosis there. An extractive economy, in its account, is not one designed by villains. It is one in which the loops have been lengthened, obscured or cut, so that the system keeps producing outcomes nobody would choose, because nothing in the loop tells it to stop. And it reads the alternative not as a plan to be imposed but as a set of practices that, once adopted by enough people in enough places, begin to reinforce themselves.

The two cases showed what this looks like at scale. In the Santa Clara Valley, loops that ran between firms rather than inside them carried information to the engineers who could use it, and a region reorganised itself after a shock that left its neighbour struggling. In the Toyota group, decades of ordinary practice in solving problems together became, in the week after a fire, the capacity to rebuild a production line without a central plan.

And the practice, from the field manual to the kitchen table, came down to a small number of habits. Map the loops before pulling the levers. Price every suspected lock-in before calling it one. Put the consequence in front of the person who acts. Probe small, with a prediction written first. Read the results honestly, the second reader first. Change two loops together when one is held by the other. Look for what is already growing, and feed it.

The house's reading

Santa Anna's book states its central principle near the start of its account of the economic traditions that preceded the market, and it is the right sentence to read the whole of this book against:

> *An economy that serves life must be grounded in the qualities of life itself — in flow rather than accumulation, in relationship rather than isolation, in sufficiency rather than craving, in reciprocity rather than extraction*[70]

Read with complexity economics, each pair in that sentence names a kind of loop. Flow rather than accumulation: value that circulates through a network rather than pooling in one node, as the Valley's knowledge circulated where Route 128's was held. Relationship rather than isolation: interaction among agents, which the representative-agent models of chapter 4 could not represent and which produces everything interesting an economy does. Sufficiency rather than craving: a balancing loop, which tells a system when enough is enough, and which an economy built only of reinforcing loops lacks. Reciprocity rather than extraction: a loop that returns something to the people who produced it, as the Toyota suppliers' help to one another returned in kind, and as arrangements that share gains with those who create them close a loop that most firms leave open.

The sentence is not a scientific finding, and this book has not treated it as one. It is a statement of what to look for. Complexity economics supplies the method of looking.



What remains uncertain

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER THIRTEEN

What remains uncertain

A book that ends with certainty would contradict everything in it. Three things remain open, and a reader should carry them.

Whether the practices this book recommends will produce good outcomes in any particular organisation or household is not known in advance, and cannot be. The tools are reasoned from models, cases and the experience of practitioners, not tested in controlled trials of whole organisations. They are offered as ways of seeing more clearly and learning faster.

How far restored feedback improves the decisions of the people who receive it depends on whether those people can bear it, interpret it and act on it. Santa Anna's book is explicit that this capacity is developed over time, and the practice of chapter 8 builds it slowly.

And which patterns an economy should grow toward is a question complexity economics, on its own, does not answer. For that, the reader needs the company described in chapter 12: Elinor Ostrom's work on governing what is held in common, the ecological limits of the doughnut, the evidence on gainsharing and ownership (gainsharing, employee ownership), Karl Polanyi's account of markets embedded in society, the economics of gift and reciprocity, the case for patient capital, regenerative and circular economics, and the argument that scarcity is often a design choice. The parent book of all of them is *Ecstatic Economics*. Together they make a series, and the series is a network: each book is stronger for the others, and the loops between them are where the most interesting questions live. The house's pages at hautelumiere.life hold them together.



For a leader who would like company

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · CHAPTER THIRTEEN

For a leader who would like company

The practice of this book can be run alone, and every tool in it is written out in full so that it can be. Some leaders prefer to run it with someone outside the organisation beside them: a second reader who has no stake in the map coming out one way or another, and who has seen the same loops in other firms. That is the work of the house's coaching and consulting practice, and for organisations where the map shows that the people who create a gain never see any of it, the conversation can extend to how that loop might be closed. The first conversation begins at hautelumiere.life.

One image

Go back, one last time, to the lawn.

It is early in the morning, before anyone has crossed it. The paved walks run round the edges at right angles, exactly where the planners put them. Across the middle runs the thin line of worn earth, from the library door to the corner of the street. Nobody laid it. Nobody voted for it. It was made by the first walker who cut across because it was shorter, and by everyone after who found the faint trace easier to follow than the open grass.

Now picture the person responsible for the lawn, standing at its edge with a coffee, looking at the line. She could fence it. She could let it widen until the grass is gone. Or she could do what this book has asked of every reader: see that the line is information, read what it tells her about where people actually go, ask what the path carries and whom it serves, and then, with care, and only where the grass has already been worn, lay the stone.

In every organisation and every household there are lines like this one. Most are invisible to the people who make them. They are the loops that grow by themselves. The work is to see them, to ask what they carry, and to tend the ones worth tending until they become the way everyone goes.

The Tool: The Desire-Line Walk

A single habit to carry when the other tools have been put away. Once a month, thirty minutes, alone.

1 · Walk the system Spend twenty minutes watching where people in your organisation or household actually go, rather than where the official path says they should: which conversations happen outside meetings, which workarounds everyone uses, which shortcuts have worn smooth.

2 · Write down three lines

A desire line (a path people take that nobody designed):

What it tells me about where the official path is wrong:

Whom it serves, and whom it leaves out:

3 · Choose one

Pave it (make the informal path official):

Leave it (watch another month):

Fence it (the line carries harm, and must be redirected, with a better path offered):

4 · The question to end on

What is growing by itself here that is worth tending?

A note on method

The desire-line image is used throughout this book as an illustration of path dependence; no particular campus is cited.

WORKBOOK

Every Tool, and a Ninety-Day Plan



Every Tool, and a Ninety-Day Plan

HAUTE LUMIÈRE · THE ECONOMY THAT KEEPS BECOMING · WORKBOOK

How to use this workbook

Every chapter of this book ends with a tool written out in full. They are gathered here in the order a practitioner would usually meet them, so that the workbook can be copied, printed or kept beside a desk without the chapters. Each tool names the chapter that explains it.

The tools are designed to work at a pale table with a pencil. None needs software, and none needs all the others. A reader with one hour should begin with the One-Hour Start and the First Map of Loops. A reader with a quarter should follow the ninety-day plan below. A leader with half a year should follow the Twenty-Six Week Plan from chapter 8.

The ninety-day plan

Ninety days is the shortest version of the practice that includes a full probe cycle. It is written for one leader with a small mapping group, or for a household. Days are working days in a firm and ordinary weeks in a household.

Days 1 to 5 · See the system

- Day 1: the One-Hour Start (chapter 1), alone.
- Days 2 to 3: the First Map of Loops (chapter 2), each member of the group alone, then compared in one session. Limit: five reinforcing loops, five lock-ins, three balancing loops.
- Day 4: choose one loop to audit and up to three suspected lock-ins to test.
- Day 5: state the sharing rule (chapter 8) to everyone whose work the map describes.

Days 6 to 15 · Test and audit

- Days 6 to 10: the Lock-In Test (chapter 3) on each suspected lock-in. Move every habit at once.
- Days 11 to 15: the Loop Audit (chapter 4) on the chosen loop, asking at least two people who act in it. Choose one severed loop to suture and the route the consequence will travel, ending with the people who act.
- Optional, for a network or a supply chain: the Cluster Reading (chapter 5) and the Shock Rehearsal (chapter 6).

Days 16 to 20 · Design the probes

- Two or three probes, using section 5 of the Loop Ledger (chapter 7). Each small, safe, informative.
- A prediction with a baseline, and a stop trigger, for each. The El Farol question asked of each.
- The Honest Dashboard (chapter 9): no more than six measures, each loop measured near its beginning and its end, each measure paired with its shadow.
- Prepare the conversations the probes need, with the Conversation Cards (chapter 10). Hold them.

- The sponsor signs the predictions. The steward keeps them.

Days 21 to 80 · Run the cycle

- The Monday Sheet (chapter 12), every week, fifteen minutes, at the same hour.
- Fortnightly, thirty minutes with the probe owners.
- At day 50, one review of the map. One probe may be added; none may be changed.
- Act on every stop trigger when it is met.

Days 81 to 85 · Read the results

- Predictions read aloud first. The second reader speaks before the sponsor and asks the seven questions (chapter 9).
- For each probe: amplify, adjust or stop.
- What the system showed that the first map did not.

Days 86 to 90 · Redraw and decide

- The second First Map of Loops. Keep the first beside it.
- The decision that was on the desk on day 1, made with everything in front of you.
- The date of the next cycle.
- The Desire-Line Walk (chapter 13), and once a month thereafter.

For a household, the same plan runs on the Kitchen-Table Map (chapter 11): three evenings in the first week, then twelve weeks of the Monday Sheet used for the household, then one evening to read and redraw.

The One-Hour Start

From chapter 1.

For the reader who wants to begin before reading further. One sheet, one hour, pencil.

1 · The system

The system I am responsible for, in one line (a firm, a team, a product, a household):

The decision currently on my desk:

2 · The levers I was planning to pull List the levers in the current plan. For each, write what you expect it to move.

LEVER	WHAT IT IS EXPECTED TO MOVE	BY HOW MUCH, AND ON WHAT EVIDENCE

3 · The loops each lever enters For each lever, ask: what will the people affected by it do in response, and how will that come back to me?

Lever 1 enters this loop:

Lever 2 enters this loop:

Lever 3 enters this loop:

4 · One thing that is already growing by itself

A practice, relationship or habit in the system that seems to be reinforcing itself:

Who knows most about it:

5 · One prediction, written now

What I expect to happen in the next ninety days if I carry out the plan unchanged:

What would show me I was wrong:

The date I will read this sheet again:

Keep the sheet. Chapter 2 turns it into the First Map of Loops.

The First Map of Loops

From chapter 2.

One page, one hour, done alone first and then with two colleagues or the other adults in the household. Write in pencil; the map will change.

1 · The system you are mapping

Name the system in one line (a firm, a team, a product, a household):

Its edge: what is inside, and what is the environment it adapts to?

The date of this map:

2 · Reinforcing loops: what grows by itself List up to five. For each, write the loop as a sentence of the form *the more A, the more B, the more A*.

LOOP, WRITTEN AS A SENTENCE	EVIDENCE IT IS RUNNING (A FIGURE, WITH ITS SOURCE)	HOW LONG IT HAS RUN	IS IT HELPING OR HARMING?

3 · Lock-ins: what history chose List up to five arrangements that everything else now depends on.

ARRANGEMENT	WHEN AND HOW IT BEGAN	WHAT DEPENDS ON IT	WOULD YOU CHOOSE IT TODAY? YES / NO / UNSURE	COST OF SWITCHING, ROUGH

4 · Balancing loops: what corrects itself List up to three. *The more A, the less B, the less A.*

…

…

…

5 · The El Farol question

Your main forecast for the next twelve months:

If everyone you compete or cooperate with held the same forecast, what would they do?

- Would that make your forecast wrong? yes / no / unsure

6 · The drawer

One small practice that may be starting to reinforce itself, which nobody has yet named:

Who knows most about it:

The date you will look at it again:

Keep the page. Chapter 3 adds a column to it and chapter 4 a second page; chapter 7 fills it in for a worked case, and this workbook gathers every page in one place.

The Lock-In Test

From chapter 3.

Use this for any arrangement you suspect is holding your organisation or household in place. One sheet per arrangement. It adds a column to the map from chapter 2.

1 · The arrangement

What it is, in one line:

When it began, and how:

What depends on it now (systems, people, contracts, habits):

2 · The alternative

The specific alternative you believe is better:

The evidence that it is better, with its source and date:

- Grade the evidence: one anecdote / several anecdotes / a measured comparison / several measured comparisons

Who would disagree, and their best argument:

3 · The price of switching

COST	AMOUNT	SOURCE OF THE FIGURE	FALLS ON WHOM
Money			
Time			
Retraining or relearning			
Lost compatibility with others			

4 · The verdict, provisionally

If the evidence for the alternative is thin: a sensible habit. Leave it, and note the date you will look again.

If the evidence is strong and switching is cheap: switch. Owner and date:

- If the evidence is strong and switching is costly: a lock-in. Continue to section 5.

5 · The loops holding it in place List every reinforcing loop that keeps the arrangement where it is (training, contracts, suppliers, other people's expectations).

…

…

...

Which loops could be changed together, at the same time?

Who else would need to move, and what would make it worth their while?

6 · The probe

One small, safe experiment to test the alternative:

What result would make you continue? · stop?

Review date:

The Loop Audit

From chapter 4.

A second page for the map begun in chapter 2. It takes the reinforcing loops you listed there and asks what each one carries. Allow ninety minutes with the people who know the system best.

1 · The loops, from your first map Copy each reinforcing loop you listed, as a sentence of the form *the more A, the more B, the more A*.

2 · What each loop carries

LOOP	WHO ACTS IN IT	WHO RECEIVES THE CONSEQUENCE	HOW LONG BEFORE THE CONSEQUENCE IS FELT (DAYS, MONTHS, YEARS)	DOES THE PERSON ACTING EVER SEE IT? YES / NO / PARTLY	DOES THE LOOP CARRY VITALITY, EXTRACTION, OR BOTH?

3 · Severed loops For each loop where the person acting does not see the consequence:

What information would have to travel back, and to whom?

What is the shortest route it could take (a report, a visit, a conversation, a shared figure)?

Who could open that route this month?

4 · Emergent patterns nobody chose List up to three patterns in your organisation or household that nobody seems to want (long meetings, mistrust between two teams, a recurring end-of-month scramble).

PATTERN	THE LOCAL RULES THAT MIGHT BE PRODUCING IT	WHICH ONE RULE, CHANGED SLIGHTLY, MIGHT CHANGE THE PATTERN

5 · Virtuous cycles already alive

One reinforcing loop that is already generating trust or vitality:

What feeds it now:

One small way to feed it more: · owner: · date:

6 · Two loops together

If you change one rule, which other loop is likely to pull it back?

What will you change at the same time, so that the first change holds?

7 · The capacity check

- Are the people who will receive restored feedback able to act on it? yes / not yet

If not yet, what practice would build that capacity (a conversation format, a decision rule, time)?

The Cluster Reading

From chapter 5.

For any network you belong to or lead: a firm and its suppliers, a set of teams, an industry association, a professional community, a neighbourhood of small businesses. Two hours, with three people who know the network from different positions.

1 · The network

Name it, and draw its edge: who is inside?

The main nodes (firms, teams, people), up to twelve:

2 · Where learning is stored

CAPABILITY THAT MATTERS	WHERE THE KNOWLEDGE LIVES (ONE PERSON, ONE TEAM, DOCUMENTS, RELATIONSHIPS ACROSS SEVERAL NODES)	IF THAT PLACE FAILED TOMORROW, WHO ELSE COULD CARRY IT?

3 · The cost of movement

FLOW	WHAT IT COSTS TODAY (TIME, PERMISSION, REPUTATION)	WHO SET THAT COST, AND WHEN	WOULD YOU SET IT THE SAME WAY TODAY?
A person moving between nodes			
Information moving between nodes without a manager			
A failed project's lessons reaching others			
An outside expert talking to an inside team			

4 · Loops inside, loops between

Three reinforcing loops that run inside a single node:

Three reinforcing loops that run between nodes:

Which set is stronger? What would happen to the network if the strongest single node faltered?

5 · Rules nobody designed for this List up to five rules (contracts, policies, defaults, habits) that change the cost of movement in the network, though they were written for another purpose.

…

…

…

Which one could be changed by the people in this room? · owner: · date:

6 · One release

One recent ending (a project, a product, a team) whose knowledge went into a drawer:

What it would take to release it now:

The Shock Rehearsal

From chapter 6.

Ninety minutes, with the people who would actually have to respond. One sheet per critical node.

1 · The node

The single part, supplier, system, person or account whose failure would hurt most:

Why it is critical (what depends on it):

How much buffer exists today (stock, time, money, a second person who knows):

2 · The first forty-eight hours Walk through them aloud. Write down what actually happens, not what the plan says.

HOUR	WHO NOTICES	WHO THEY CALL	WHAT THAT PERSON NEEDS TO KNOW	WHERE THAT INFORMATION IS NOW
0–2				
2–12				
12–24				
24–48				

3 · The relationships that would carry you

WHO COULD HELP	HAVE YOU EVER SOLVED A HARD PROBLEM TOGETHER? YES / NO	WOULD THEY START BEFORE TERMS WERE AGREED? YES / NO / UNSURE	WHAT WOULD MAKE THAT MORE LIKELY?

4 · Loops that exist only on paper

Contacts, agreements or procedures that the rehearsal showed nobody would actually use:

Information that the right person would not be able to find:

5 · Practice, in ordinary weeks

One low-stakes occasion in the next quarter to solve a real problem with one of the relationships above: · owner: · date:

One standing agreement to make in advance (for example, "in an emergency of this kind, work starts first and we settle on this principle afterward"):

6 · The household version

The household's single point of failure:

The person or neighbour who would help in the first forty-eight hours:

One way help could flow toward them this month:

The Loop Ledger

From chapter 7.

The three earlier pages (the First Map, the Lock-In Test, the Loop Audit) gathered into one ledger for a twelve-week cycle. Fill in sections 1 to 4 in week 1, section 5 in week 2, and section 6 in week 12.

1 · The system and the question

System (firm, team, product, household):

The decision on the table, in one line:

The options under consideration:

2 · Loops, with evidence

REINFORCING LOOP (THE MORE A, THE MORE B, THE MORE A)	EVIDENCE, WITH SOURCE	WHO ACTS	WHO FEELS THE CONSEQUENCE	DOES THE ACTOR SEE IT?	CARRIES VITALITY / EXTRACTION / BOTH

3 · Lock-ins, tested

ARRANGEMENT	EVIDENCE FOR THE ALTERNATIVE (GRADED)	SWITCHING COST	VERDICT: HABIT / SWITCH / TRUE LOCK-IN	LOOPS HOLDING IT IN PLACE

4 · The El Farol line

Your forecast:

What everyone else will do if they share it:

Where that leaves room nobody is serving:

5 · Adjacency and probes

OPTION	CAPABILITIES IT NEEDS THAT YOU HAVE	CAPABILITIES IT NEEDS THAT YOU LACK	CROWDED BY OTHERS' FORECASTS?	PROBE (SMALL, SAFE, TWELVE WEEKS)	PREDICTION, WRITTEN NOW	STOP CONDITION

One severed loop to suture during the twelve weeks: · route the feedback will travel: · read unedited to:

Two loops to change together, if a lock-in is involved: and

6 · Week twelve

PROBE	PREDICTED	HAPPENED	AMPLIFY / ADJUST / STOP

What the system now shows that the week-one map did not:

The next decision, and the next twelve-week date:

The Twenty-Six Week Plan

From chapter 8.

The whole field manual on one page, to be filled in by the steward at week 0 and kept beside the Monday Sheet.

Set-up (week 0)

System:

Question on the desk:

Sponsor: · Steward: · Second reader:

Mapping group (five to nine names, at least two close to the work):

Sharing rule, in one sentence:

Weeks 1–2 · First map

Dates of individual mapping and of the comparison session:

Loop chosen for audit:

Lock-ins chosen for testing (up to three):

Weeks 3–4 · Lock-in tests

ARRANGEMENT	EVIDENCE GRADE	SWITCHING COST	VERDICT	OWNER, DATE

Weeks 5–6 · Loop audit

Severed loop to suture:

Route the consequence will travel, ending with:

Weeks 7–8 · Probes (signed by sponsor on)

PROBE	OWNER	PREDICTION, WITH BASELINE	STOP TRIGGER	EL FAROL NOTE

Weeks 9–20 · First cycle

Fortnightly review dates:

Week 15 map review date:

Weeks 21–22 · Reading

Review date:

- Predictions read aloud first: yes / no

Second reader's three strongest objections:

Weeks 23–26 · Second map and decision

What the second map shows that the first did not:

The decision on the desk:

Date of the next cycle:

The Honest Dashboard

From chapter 9.

One page, completed by the steward with the second reader. No more than six measures in total.

1 · The loops being measured

Loop 1 (from the Loop Ledger):

Loop 2:

2 · Each loop, measured near its beginning and near its end

LOOP	EARLY MEASURE (A LINK NEAR THE START)	LATE MEASURE (THE OUTPUT)	SOURCE OF EACH FIGURE	HOW OFTEN READ

3 · Shadows

MEASURE	IF I GAMED THIS, WHAT WOULD I DO?	THE SHADOW MEASURE THAT WOULD REVEAL IT

4 · Who sees each figure

MEASURE	SEEN BY THE PEOPLE IN THE LOOP? YES / NO	HOW SOON AFTER IT HAPPENS?	IN WHAT FORM?

5 · Lock-in dependency

Arrangement being loosened:

Share of the system that depends on it today: · source:

Predicted share after two cycles:

6 · Restored loop

Arrival: what is supposed to reach whom, how often? · actual arrival so far:

Use: what have the recipients done differently?

Effect: predicted · observed

7 · The second reader's seven questions Tick each one asked of the latest result: prediction first ☐ · what else changed ☐ · who chose participants ☐ · knew they were measured ☐ · shadow measure ☐ · distribution ☐ · surprising to the people in the loop ☐

The Conversation Cards

From chapter 10.

One card per conversation. Fill it in before the conversation and again after it.

Before

Who:

What they are most likely worried about:

What is working, that I will name first:

What I do not know, that I will say plainly:

What each party brings, and what each needs:

The one small thing I am asking for:

What I will commit to in return (a date to report back, a sharing rule, a stop condition):

During

What they told me that was not on my map:

Their strongest objection, in their words:

After

What was agreed:

The date I will return with the result:

Did this conversation start a loop of trust or of mistrust? What would strengthen or repair it?

The Kitchen-Table Map

From chapter 11.

Three evenings, then twelve weeks. Each adult fills in sections 1 and 2 alone first.

1 · Loops (alone, thirty minutes)

- Five things in our household that grow by themselves, good or bad, each written as *the more A, the more B, the more A*:

…

…

...

...

...

Three things that correct themselves (*the more A, the less B, the less A*):

2 · Arrangements nobody would choose today (alone)

ARRANGEMENT	WHEN IT BEGAN	WHAT DEPENDS ON IT	WOULD WE CHOOSE IT TODAY?	ROUGH COST OF CHANGING IT

3 · Compare (together)

Where our maps agree:

Where they differ, and why:

4 · The severed loop

One payment, habit or decision whose consequence nobody in the household feels:

How we could bring it closer (a weekly look, a manual payment, a conversation):

5 · The probe

One small, safe change for twelve weeks:

What we predict, written now:

What would make us stop:

Start date: · reading date:

6 · The four measures

MEASURE	TODAY	WEEK 13
Share of income committed before it arrives		
Amount set aside this month		
Balance owed		
The felt month (each adult)		

7 · The El Farol question

The largest decision ahead of us:

If everyone in our situation made the same choice, what would happen?

Where is the room the crowd would leave?

The Monday Sheet

From chapter 12.

One page that gathers the book into a single weekly practice. Keep it beside the Loop Ledger from chapter 7. Fifteen minutes, once a week, at the same hour.

1 · This week's loop

The reinforcing loop I am watching this week:

What it did since last week (a figure, with its source):

- Did the people acting in it see the consequence? yes / no / partly

2 · This week's lock-in

The arrangement under test:

- Its status: habit, moved / switching, under way / true lock-in, loops being changed together

The two loops being changed together, if a lock-in: and

3 · This week's restored feedback

The consequence that travelled back this week:

To whom: · unedited? yes / no

What they did with it:

4 · The probes

PROBE	WEEK OF TWELVE	PREDICTION (WRITTEN AT START)	SO FAR	STOP CONDITION MET?

PROBE	WEEK OF TWELVE	PREDICTION (WRITTEN AT START)	SO FAR	STOP CONDITION MET?

5 · The joins

Governance: who monitors this loop, and can they change its rules? (Ostrom)

Limits: what balancing loop keeps this from pushing past a boundary we care about? (the doughnut)

Gain: when this loop produces a gain, who receives it, and do the people who produced it see any of it?

Embedding: whose relationships does this loop depend on, outside the firm or household?

6 · The El Farol line

The forecast I am acting on:

What everyone else would do if they shared it:

7 · The desire line

One place this week where people went their own way, around the official path:

What it tells me about where the path should be:

The Desire-Line Walk

From chapter 13.

A single habit to carry when the other tools have been put away. Once a month, thirty minutes, alone.

1 · Walk the system Spend twenty minutes watching where people in your organisation or household actually go, rather than where the official path says they should: which conversations happen outside meetings, which workarounds everyone uses, which shortcuts have worn smooth.

2 · Write down three lines

A desire line (a path people take that nobody designed):

What it tells me about where the official path is wrong:

Whom it serves, and whom it leaves out:

3 · Choose one

Pave it (make the informal path official):

Leave it (watch another month):

Fence it (the line carries harm, and must be redirected, with a better path offered):

4 · The question to end on

What is growing by itself here that is worth tending?

A last page

Keep every completed page, dated, in one folder. A year of them, read in order, is the most honest record an organisation or a household has of how it actually learned. That record is itself a reinforcing loop: the more pages are kept, the easier the next map is to draw, and the better the next map, the more worth keeping the pages become.

Loops, in your own firm

An economy is not a machine that settles but a system that keeps forming itself. A firm is the same, and so is the gain inside it.

Luminous Prosperity Inc. takes this work into organisations as a gainshare. The flow is opened from the inside, with the people who already hold what the firm knows. The gain is counted from records the firm already keeps, against a measure both sides sign before anyone knows the answer, and the house is paid as a share of what is found: between thirty and ten per cent of it.

The workbook gives twelve weeks of probes. Which loops in your firm carry vitality and which carry extraction is read from your own accounts and your own people. That begins with a conversation.

REQUEST A CONVERSATION

References

1. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 2, "The Gift Economy: Wealth as Flow, Not Stock". The italics are the author's.
2. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 3, "The Somatic Dimension of Wealth". The italics are the author's.
3. Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116–131. https://doi.org/10.2307/2234208. The four properties (nonpredictability, potential inefficiency, inflexibility, nonergodicity) are set out in the paper. It is a theoretical model of technology adoption; it does not measure how often real markets lock in. The account of the paper's delay is Arthur's own, in the preface to *Increasing Returns and Path Dependence in the Economy* (University of Michigan Press, 1994), and is not independently verified here.
4. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 4, "The Regenerative Design Imperative".
5. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on integral and developmental economics.
6. Arthur, W. B. (1990). Positive feedbacks in the economy. *Scientific American*, 262(2), 92–99. https://doi.org/10.1038/scientificamerican0290-92. An essay for general readers summarising the increasing-returns argument; it is not itself new evidence.
7. Arthur, W. B. (1999). Complexity and the economy. *Science*, 284(5411), 107–109. https://doi.org/10.1126/science.284.5411.107. A perspective essay; it states a programme rather than testing one.
8. Arthur, W. B. (1994). Inductive reasoning and bounded rationality. *American Economic Review*, 84(2), 406–411 (Papers and Proceedings). https://www.jstor.org/stable/2117868. The El Farol figures (one hundred people, a threshold of sixty) are from this paper. It reports computer simulations, not observations of a real bar. No DOI is registered for this issue of the *Proceedings*; the stable JSTOR address is given instead.
9. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 1, on Adam Smith.
10. Santa Anna, A. *The Luminous 100*, "Effectuation (Saras Sarasvathy)", Luminous Upgrade. Luminous Prosperity, 2026.
11. Arthur, W. B. (2021). Foundations of complexity economics. *Nature Reviews Physics*, 3, 136–145. https://doi.org/10.1038/s42254-020-00273-3. A review of the framework by its principal author; a reader should weigh it as a statement of position as well as of method.
12. David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2), 332–337 (Papers and Proceedings). https://www.jstor.org/stable/1805621. A historical narrative; its account of Dvorak's superiority rests on earlier trials that were later disputed (see note 2). No DOI is registered for this issue; the stable JSTOR address is given.
13. Liebowitz, S. J., & Margolis, S. E. (1990). The fable of the keys. *The Journal of Law and Economics*, 33(1), 1–25. https://doi.org/10.1086/467198. A re-examination of the evidence on keyboard layouts; it challenges one case and the general frequency of inefficient lock-in, but does not show that inefficient lock-in never occurs. The later book is Liebowitz, S. J., & Margolis, S. E. (1999). *Winners, Losers and Microsoft*. Independent Institute.
14. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on integral economics, "Premature synthesis".
15. Fagiolo, G., Moneta, A., & Windrum, P. (2007). A critical guide to empirical validation of agent-based models in economics: Methodologies, procedures, and open problems. *Computational Economics*, 30(3), 195–226. https://doi.org/10.1007/s10614-007-9104-4. A methodological survey; it describes the state of practice up to its date and does not settle which validation standard is right.
16. Farmer, J. D., & Foley, D. (2009). The economy needs agent-based modelling. *Nature*, 460, 685–686. https://doi.org/10.1038/460685a. A commentary arguing a case; it is not an empirical test of agent-based models' accuracy.
17. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 4, "Common Pitfalls".

18. Unruh, G. C. (2000). Understanding carbon lock-in. *Energy Policy*, 28(12), 817–830. https://doi.org/10.1016/S0301-4215(00)00070-7. A conceptual paper; it proposes the techno-institutional complex as a framework and does not measure the strength of each loop.
19. Cowan, R. (1990). Nuclear power reactors: A study in technological lock-in. *The Journal of Economic History*, 50(3), 541–567. https://doi.org/10.1017/S0022050700037153. A historical study; whether any alternative reactor design would in fact have performed better is a judgement the evidence cannot fully settle.
20. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on policy and governance.
21. Snowden, D. J., & Boone, M. E. (2007). A leader's framework for decision making. *Harvard Business Review*, 85(11), 68–76. https://hbr.org/2007/11/a-leaders-framework-for-decision-making. A practitioner framework; it is not an empirical study of decision outcomes. No DOI is registered for this article.
22. Santa Anna, A. *The Luminous 100*, "Complexity Adaptive Systems (Stacey Matrix)", Luminous Upgrade. Luminous Prosperity, 2026. The quotation crosses a line break in the source.
23. Schelling, T. C. (1971). Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2), 143–186. https://doi.org/10.1080/0022250X.1971.9989794. A theoretical model; it shows that mild preferences *can* produce strong segregation, not that any particular real pattern of segregation arose this way.
24. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on the somatic dimension of economic life.
25. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 1.
26. Kirman, A. P. (1992). Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2), 117–136. https://doi.org/10.1257/jep.6.2.117. A theoretical critique of a modelling convention; it does not itself supply a replacement model.
27. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on vitality-centred workplaces.
28. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on money and finance.
29. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on money and finance, "Restored feedback loops".
30. May, R. M., Levin, S. A., & Sugihara, G. (2008). Ecology for bankers. *Nature*, 451, 893–894. https://doi.org/10.1038/451893a. A commentary proposing an approach; it presents no new data on bank failures.
31. Haldane, A. G., & May, R. M. (2011). Systemic risk in banking ecosystems. *Nature*, 469, 351–355. https://doi.org/10.1038/nature09659. A network model of banking stability; its findings are properties of the model and are not a forecast of any particular crisis.
32. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on vitality-centred workplaces and complementary currencies.
33. Santa Anna, A. *The Luminous 100*, "Systems Thinking / Causal Loop Diagrams (Senge)", Luminous Upgrade. Luminous Prosperity, 2026.
34. Santa Anna, A. *The Luminous 100*, "Cynefin Framework (Dave Snowden)", Luminous Upgrade. Luminous Prosperity, 2026.
35. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on vitality-centred workplaces.
36. Saxenian, A. (1994). *Regional Advantage: Culture and Competition in Silicon Valley and Route 128*. Harvard University Press. An interpretive comparative study based on extensive interviews; it argues for the importance of network structure but does not measure information flows directly. No DOI is registered for the book.
37. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 4, "Design Principles for Vitality-Centered Workplaces", on feedback richness.
38. Santa Anna, A. *The Luminous 100*, "Viable System Model (Stafford Beer)", Luminous Upgrade. Luminous Prosperity, 2026.
39. Gilson, R. J. (1999). The legal infrastructure of high technology industrial districts: Silicon Valley, Route 128, and covenants not to compete. *New York University Law Review*, 74(3), 575–629. A legal and economic argument proposing one contributing mechanism; it is not an empirical test, and later studies of non-compete enforcement have reached mixed conclusions not reviewed here. No DOI is registered.

40. Krugman, P. (1991). Increasing returns and economic geography. *Journal of Political Economy*, 99(3), 483–499. https://doi.org/10.1086/261763. A formal model of industrial concentration; it explains where clusters form under stated assumptions, not how a particular cluster's culture works.
41. Nishiguchi, T., & Beaudet, A. (1998). The Toyota group and the Aisin fire. *Sloan Management Review*, 40(1), 49–59. A qualitative account based on interviews and company sources; it reports that assembly plants reopened after two days of shutdown and names participating firms, but its figures for the scale and cost of the recovery are the authors' and are not independently verified here. No DOI is registered for this article.
42. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 5, "Policy and Governance in Ecstatic Economics".
43. Santa Anna, A. *The Luminous 100*, "Snowflake Model / Fractal Organization", Luminous Upgrade. Luminous Prosperity, 2026.
44. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 4, "Community Currencies and Alternative Exchange".
45. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on ecology and the measure of wealth.
46. Hidalgo, C. A., Klinger, B., Barabási, A.-L., & Hausmann, R. (2007). The product space conditions the development of nations. *Science*, 317(5837), 482–487. https://doi.org/10.1126/science.1144581. A study of national export data; it says nothing directly about firms, and its use here is as an analogy for capability adjacency.
47. Hidalgo, C. A., & Hausmann, R. (2009). The building blocks of economic complexity. *Proceedings of the National Academy of Sciences*, 106(26), 10570–10575. https://doi.org/10.1073/pnas.0900943106. Country-level analysis; its measures of complexity and their predictive power have been debated in later literature not reviewed here.
48. Santa Anna, A. *The Luminous 100*, "Solutions-Focused Coaching", Luminous Upgrade. Luminous Prosperity, 2026.
49. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on money and finance, "The Scalability Trap".
50. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on practical applications.
51. Campbell, D. T. (1979). Assessing the impact of planned social change. *Evaluation and Program Planning*, 2(1), 67–90. https://doi.org/10.1016/0149-7189(79)90048-X. A methodological essay drawing on programme evaluation; its "law" is a generalisation from experience, not a measured regularity. The wording above paraphrases it.
52. Goodhart, C. A. E. (1975). Problems of monetary management: the U.K. experience. *Papers in Monetary Economics*, Reserve Bank of Australia, vol. 1. The observation was made about monetary aggregates; its later generalisation to all targets is a convention of the literature rather than Goodhart's own claim. No DOI is registered; the bibliographic details are given from the standard citation and were not re-read from the original for this book.
53. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 1, "What the GDP Measures — and What It Misses Entirely".
54. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 5, "Wellbeing Budgets: When Governments Measure What Matters".
55. Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly*, 44(2), 350–383. https://doi.org/10.2307/2666999. A field study of fifty-one teams in one manufacturing company; it links team psychological safety to learning behaviour, and does not test the conversational scripts offered here.
56. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 4, "Design Principles for Vitality-Centered Workplaces", on feedback richness.
57. Santa Anna, A. *The Luminous 100*, "Stakeholder Mapping", Luminous Upgrade. Luminous Prosperity, 2026.
58. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 4, "Common Pitfalls".
59. Santa Anna, A. *The Luminous 100*, "Polarity Management (Barry Johnson)", Luminous Upgrade. Luminous Prosperity, 2026.
60. Madrian, B. C., & Shea, D. F. (2001). The power of suggestion: Inertia in 401(k) participation and savings behavior. *The Quarterly Journal of Economics*, 116(4), 1149–1187. https://doi.org/10.1162/003355301753265543. A study of one large US employer's change of default; it shows the strength of defaults in that setting and does not say what contribution rate or investment is right for anyone.

61. Thaler, R. H., & Benartzi, S. (2004). Save More Tomorrow: Using behavioral economics to increase employee saving. *Journal of Political Economy*, 112(S1), S164–S187. https://doi.org/10.1086/380085. Reports the results of a designed programme at a small number of firms; it is evidence for the mechanism, not a recommendation of any saving rate.

62. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 3, "Practice 5: The Vitality Inventory (weekly)".

63. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 3, "The Intergenerational Transmission of Economic Somatics".

64. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on money and finance, "The Scalability Trap".

65. Ostrom, E. (2010). Beyond markets and states: Polycentric governance of complex economic systems. *American Economic Review*, 100(3), 641–672. https://doi.org/10.1257/aer.100.3.641. Ostrom's Nobel lecture; it summarises a large body of field and laboratory research, but its design principles describe what long-enduring commons have in common rather than guaranteeing that any commons applying them will endure.

66. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on practical applications, "The Commons Renaissance".

67. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on practical applications, "The Commons Renaissance".

68. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on ecology and the measure of wealth.

69. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter on practical applications, on Mondragon.

70. Santa Anna, A. *Ecstatic Economics*. Luminous Prosperity, 2026. Chapter 2, "The Common Thread: Vitality as the Foundation of Genuine Prosperity". The bold is the author's.

The Economy That Keeps Becoming

Ammanuel Santa Anna · Luminous Prosperity Inc. · Haute Lumière

Set in Cormorant Garamond and Manrope.

Read the Business Library at business.hautelumiere.life/library.

Luminous Library Licence 1.0 · hautelumiere.life/library-licence

REQUEST A CONVERSATION

business.hautelumiere.life/book/?kind=gainshare

The Business Library

business.hautelumiere.life/library/

The packages

hautelumiere.life/pricing/