Strangers at Scale

Strangers at Scale

Abdul Kavungal (AK)

Award Winning Creations

Strangers at Scale

© 2026 Abdul Kavungal (AK). Published by Award Winning Creations.

All rights reserved. No part of this book may be reproduced in any form without written permission from the publisher, except for brief quotations in reviews and scholarly works.

First edition, 2026.

This book discusses ideas from *The Human Swarm* by Mark W. Moffett, *The Origin of Wealth* by Eric D. Beinhocker and *Thinking in Systems* by Donella H. Meadows. It is an independent work of commentary and synthesis and is not affiliated with or endorsed by those authors or their publishers. Case material is drawn from published sources listed in the notes.

Display type set in Nightgilt Display from the TyAk type foundry. Text set in Crimson Text.

Contents

Introduction: The Most Ordinary Miracle 7

Part I: The Anonymous Animal 13

- Chapter 1: The Forest and the Carriage 15
- Chapter 2: Billions Without a Quarrel 25
- Chapter 3: Three Fingers in a Basement 35
- Chapter 4: Born Reading Markers 43
- Chapter 5: How Many People Can You Know? 53
- Chapter 6: Real People 63

Part II: The System of Trust 73

- Chapter 7: Trust as a Stock 75
- Chapter 8: Traps in the Marker System 83
- Chapter 9: Wealth Is What a Society Agrees to Count 91
- Chapter 10: Money, the Marker Everyone Carries 99
- Chapter 11: Kinship's Long Shadow 107
- Chapter 12: Differentiate, Select, Amplify 115
- Chapter 13: Status, Fairness and the Instinct Beneath the Spreadsheet 123

Part III: Markers at Work 131

- Chapter 14: The Badge and the Brand 133
- Chapter 15: Credentials and Licences 141
- Chapter 16: The Great Impostor 149
- Chapter 17: Uniforms and Dress Codes 157
- Chapter 18: The Dialect of the Office 165
- Chapter 19: Markers at a Distance 173
- Chapter 20: When Two Tribes Sign a Deal 181
- Chapter 21: Why Culture Resists Redesign 189

Part IV: Platforms and Digital Strangers 197

- Chapter 22: Reputation Engines 199
- Chapter 23: A Stranger's Spare Room 207
- Chapter 24: Commons Without Kin 213
- Chapter 25: Counterfeit Belonging 221

Part V: Designing Markers on Purpose 229
Chapter 26: The Marker Audit 231
Chapter 27: Designing Markers on Purpose 239
Chapter 28: Repairing Broken Markers 247
Conclusion: Real People 255
A Note on Sources and Method 261
Notes and Sources 263
Glossary 273

Introduction: The Most Ordinary Miracle

At a quarter past eight on a weekday morning, a train pulls into a city station and several hundred people press into a carriage built for fewer. Nobody in the carriage knows more than one or two of the others. Many of them will stand closer to a stranger than they would stand to a cousin at a wedding. A woman reads over a man's shoulder. A teenager's backpack presses into the ribs of a retired accountant. A man in a hi-vis vest eats a pastry while a lawyer tries not to get crumbs on her coat. For twenty minutes, in a sealed metal box, hundreds of unrelated adults tolerate one another so completely that the whole event is boring. Then the doors open and they scatter into a city of several million people, almost none of whom they will ever meet.

We do not think of this as remarkable. We should. Put a few hundred unfamiliar chimpanzees in that carriage and you would not have a commute; you would have a massacre. Chimpanzees are our closest living relatives, intelligent, social and capable of deep loyalty to the animals they know. They are also, as decades of field research have shown, intensely hostile to the animals they do not know. A chimpanzee that meets an unfamiliar chimpanzee from a neighbouring community has, broadly speaking, two choices. It can flee, or, if its side has the numbers, it can attack. There is no third option in which the two of them nod, look at their phones and get off at different stops.

The biologist Mark W. Moffett built a book, *The Human Swarm*, around this contrast. His argument is that very few animals have ever solved the problem of living among strangers. Most social mammals, from meerkats to wolves to elephants, live in what he calls individual-recognition

societies. Every member must know every other member, at least well enough to tell friend from intruder. Such societies work beautifully, and they have a ceiling. They cannot grow beyond the number of faces, voices or smells their members can keep straight. When they get too large, they split.

Humans and certain social insects broke through that ceiling. An ant from a colony of millions does not know the other ants. It knows that they smell right. A citizen of a nation of tens of millions does not know the other citizens. She knows that they carry the same passport, speak with an accent she recognises, eat food she considers normal, salute the same flag and count on their fingers the way she does. Moffett calls these signals markers. Markers are the reason a society can be anonymous without being dangerous. They let you treat a stranger as one of us without having to know who that stranger is.

This book takes Moffett's idea out of the forest and the anthill and into the places where most readers spend their working lives: companies, markets and digital platforms. Each of those is an anonymous society. A multinational with a hundred thousand employees is far too large for anyone to know everyone. A marketplace where millions of buyers send money to sellers they will never see depends on a stranger behaving well. A platform that lets you sleep in a stranger's spare room, get into a stranger's car or merge a stranger's code into software running on a billion devices is betting its entire business on trust among people who have never met. None of these institutions could exist on the scale they do without markers. The brand on the box, the degree on the wall, the lanyard at the security gate, the badge beside a username, the shared vocabulary in a Slack channel, the particular way a team's emoji reactions land: these are the descendants of the ochre, beads and greetings that let hunter-gatherer bands recognise members of their wider society.

Three books in one conversation

The book is a mashup in the honest sense. It puts three bodies of thought into a room and makes them talk to one another.

From Moffett comes the core puzzle and the core concept. Anonymous societies are rare in nature. They depend on markers. Markers are learned early, held deeply and defended fiercely. Even infants read them. Hunter-gatherer bands, often imagined as tiny isolated groups, usually saw themselves as part of a much larger people with a shared language, shared customs and a shared name. Many such peoples called themselves by a word meaning something close to "real people," an indication of how much of human identity is built on the line between those who carry the right markers and those who do not.

From Eric Beinhocker's *The Origin of Wealth* comes a view of the economy as an evolving complex system rather than a machine that settles into equilibrium. Beinhocker argues that economic order emerges through the same broad algorithm that drives biological evolution: variation is generated, some variants are selected, and the successful ones are amplified. Products, business plans, institutions and what he calls social technologies, meaning the ways we organise people, all evolve this way. Beinhocker also insists that wealth is not a fixed physical quantity. What counts as wealth is defined by culture. For Maasai herders, wealth has long been measured in cattle. For a bank customer in Frankfurt, it is a balance on a statement. Each definition only works because a whole society agrees to honour it, and that agreement is itself a kind of marker.

From Donella Meadows's *Thinking in Systems* comes the language for seeing how markers behave over time. Meadows taught generations of readers to see stocks, which are accumulations such as water in a bathtub or money in an account, and flows, which fill and drain them. She showed how feedback loops either reinforce change or resist it, how delays make systems oscillate and overshoot, and how the most powerful places to intervene in a system are rarely the obvious ones. In her famous hierarchy of leverage points, changing a number such as a subsidy or a tax rate sits near the bottom. Changing the rules, the information flows, the goals and, highest of all, the paradigm out of which a system arises sit near the top.

Put these three together and a picture forms. Trust among strangers behaves like a stock. It accumulates slowly through repeated good experience and drains quickly through betrayal. Markers are the

information flows that let people infer that stock without having to measure it directly. When markers are reliable, trust can scale far beyond the circle of people who know each other. When markers are forged, inflated, confused or abruptly changed, trust drains, and the system behaves in ways that look irrational from inside: mergers that make sense on paper fall apart, culture programmes produce cynicism, platforms fill with fraud, and a currency that was solid on Monday is shunned by Friday.

Why leaders should care

Most leaders already work with markers every day without calling them that. A founder choosing a company name, a logo and a dress code is designing markers. A head of HR writing a competency framework is designing markers. A product team building a review system, a verification badge or a seller rating is designing markers. A board approving an acquisition is deciding whether two sets of markers can be blended, replaced or left alone. A regulator setting licensing requirements is deciding which markers will be legally required before one stranger may treat another. Very few of these people have a working theory of what markers do, why they matter so much to the people who carry them, or how they fail.

The absence shows. Mergers are announced as marriages of equals and run as conquests, and the people on the losing side read every changed letterhead as a message. Rebrands that test well in focus groups are greeted with fury by loyal customers who experience the change as a betrayal of membership. Platforms launch rating systems that inflate until every score is excellent and therefore meaningless. Credential systems grow until they keep competent outsiders out and fail to keep incompetent insiders from causing harm. Security teams discover that an attacker did not need to break any code, because they only needed to look like a colleague.

The argument of this book is that these failures are predictable once you see markers as part of the structure of a system. Markers are not decoration. They are the means by which a large population of strangers decides,

thousands of times a day, whom to cooperate with and whom to treat with suspicion. They are shaped by evolution, culture and design, and they can be read, mapped, tested and repaired.

What this book claims, and what it does not

It is worth being clear at the start about the limits of the argument.

The book does not claim that markers are the only thing holding large groups together. Laws, contracts, courts, police, payment systems and plain self-interest all play their part. Markers work alongside these institutions and often make them possible, but they do not replace them.

It does not claim that markers are always benign. The same capacity that lets a nation of strangers cooperate also lets a nation of strangers persecute a minority whose markers differ. The history of human societies includes yellow stars, segregated water fountains and genocides justified by the claim that some people were not real people. A book about designing markers that ignored this would be dangerous. The chapters that follow treat exclusion as a core property of markers, not an unfortunate side effect, and they treat the leader's job as designing markers that widen the circle of trust rather than hardening its edge.

It does not claim that the science is settled. Some of the ideas discussed here are robust and widely replicated. Others, including the famous claim that humans can maintain about 150 stable relationships, are contested, and the book says so. Where research is preliminary, the text treats it as preliminary.

And it does not claim that companies are tribes in any literal sense. The analogy between a hunter-gatherer people and a software company is useful because both face the same structural problem of cooperating beyond personal acquaintance. That shared problem produces similar solutions. It does not make employees into warriors or customers into kin.

How the book is organised

Part I, The Anonymous Animal, sets out the puzzle. It begins with individual-recognition societies and the ceiling they cannot break, moves to the ants that broke it with chemistry, then to the human markers that do the same job in culture. It ends with a careful look at Dunbar's number and what the debate over it does and does not tell us.

Part II, The System of Trust, brings in Meadows and Beinhocker. It treats trust as a stock with inflows, outflows and delays, looks at wealth and money as markers a society agrees to honour, follows the evolution of markers through differentiation, selection and amplification, and examines the instincts of status and fairness that make markers feel so personal.

Part III, Markers at Work, applies the framework to organisations: brands, credentials and licences, uniforms and dress codes, the dialects of working life, the collision of markers in mergers and the reasons culture so stubbornly resists redesign.

Part IV, Platforms and Digital Strangers, turns to the online institutions that have pushed anonymous cooperation further than any society before them: reputation systems at eBay, Airbnb and Stack Overflow, commons such as Wikipedia and open-source software, and the counterfeit markers that now threaten them.

Part V, Designing Markers on Purpose, gathers the lessons into practical tools for leaders who want to design new markers, audit existing ones and repair those that have broken.

The conclusion returns to the idea of real people. Every society draws a line between those it treats as members and those it does not. The great achievement of human civilisation has been to move that line outward, from the band to the people, from the people to the nation, and, imperfectly, toward strangers on the far side of the planet. Organisations and platforms are now among the main places where that line is drawn. The people who design them are, whether they like it or not, in the business of deciding who counts as real.

PART I

The Anonymous Animal

Before markers can be designed, they have to be understood as a biological and cultural achievement. Part I asks why so few animals live among strangers, how ants and humans manage it, and how far our capacity to know one another personally really reaches.

CHAPTER 1

The Forest and the Carriage

In 1974 a community of chimpanzees in Gombe Stream National Park, in what is now Tanzania, began to kill its neighbours. The victims were not strangers in the usual sense. A few years earlier they had belonged to the same community, the Kasekela group studied by Jane Goodall and her colleagues. A set of males and females had drifted south and begun to spend their time apart, until the researchers recognised them as a separate community, which they called Kahama. Over the next four years, Kasekela males made a series of raids into the southern territory. Moving quietly in groups, they found Kahama males alone or in small numbers and attacked them with a violence that shocked the field team. By the end of the episode, the Kahama community had effectively ceased to exist, and Kasekela had absorbed much of its range.

The episode, often called the Four-Year War, upended a gentle picture of our closest relatives. For a while some critics suggested that the violence was an artefact of human interference, perhaps caused by the bananas researchers had provided to bring the animals into view. Later evidence made that explanation hard to sustain. At Ngogo, in Kibale National Park in Uganda, John Mitani, David Watts and Sylvia Amsler documented a large chimpanzee community killing neighbours over roughly a decade and then expanding into the territory it had cleared, taking over an area equal to around a fifth of its former range. In 2014 a large team led by Michael Wilson compiled data from eighteen chimpanzee communities and concluded that lethal attacks were better explained as an adaptive strategy, a way of eliminating rivals when the risk to attackers was low, than as a response to human disturbance.

What matters for this book is not that chimpanzees are violent. Many animals are. What matters is the shape of their tolerance. Within a community, chimpanzees groom, share food, form alliances, reconcile after fights and care for orphans. They know one another as individuals, with histories, reputations and moods. Across the boundary of the community, that knowledge stops. An animal from the next valley is not a potential partner whose qualities can be discovered. It is an outsider, and the default responses to outsiders are avoidance or attack. The community ends where personal recognition ends.

Now return to the commuter carriage described in the introduction. Hundreds of adult humans, most of them complete strangers to one another, stand in an enclosed space without fear. They do not need to know each other's names, alliances or histories. They do not need to have groomed one another or fought and reconciled. Each passenger takes for granted that the others belong to the same broad society and will follow its unspoken rules. The difference between the chimpanzee boundary and the train carriage is the difference this book is about.

Societies built on faces

Mark Moffett, who spent years as a field biologist studying ants before turning to human societies, makes a distinction that is easy to state and surprisingly powerful. Most social animals live in individual-recognition societies. In such a society, membership depends on every member knowing every other member. You belong because the others know you, and you know them. A stranger is anyone you do not recognise, and recognising requires memory.

Individual recognition takes many forms. Some animals recognise faces. Others recognise voices, smells or a combination of cues. Elephants are a striking example. In the Amboseli population studied for decades in Kenya, researchers including Karen McComb showed that adult females respond differently to the recorded contact calls of familiar and unfamiliar elephants, and that they appear to keep track of the calls of a large number of other individuals, on the order of a hundred adult females spread across

many family groups. When a call comes from an animal they know, they relax or call back. When it comes from a stranger, they bunch together defensively.

Wolves offer a different version. A typical pack is a family, often a breeding pair and their offspring from several years. Each animal knows the others intimately, and the pack defends its territory against other packs. An unfamiliar wolf approaching a pack is usually chased away or attacked, unless it is accepted through a slow process of approach and submission, perhaps to replace a lost breeder. Meerkats in the Kalahari live in groups that can run to a few dozen animals, cooperating to watch for predators, babysit pups and teach the young to handle scorpions. A meerkat that wanders into another group's territory meets aggression. Membership is a matter of personal history.

Birds do the same. Several species of jay live in family groups that defend territories and recognise their members. Some primates track not only who is in their group but who is related to whom, and how those relationships rank. Baboon researchers have shown, through clever playback experiments, that baboons respond with surprise to recorded sequences of calls that imply a reversal of the established dominance order. They are keeping score, and the score is personal.

These societies are not primitive. They are elegant solutions to real problems. A group of animals that know one another can share labour, defend young, pass on knowledge and resolve disputes. Individual recognition lets each member build a model of every other member: who is reliable, who is dangerous, who is kin, who owes whom a favour. Because members know each other, cheating is risky, since the cheat will be remembered. The system enforces cooperation through memory.

The ceiling

The trouble with memory is that it runs out. A society that requires every member to recognise every other member cannot grow indefinitely. As group size rises, the number of relationships rises faster. Ten animals have forty-five pairwise relationships. A hundred have nearly five thousand. A

thousand have nearly half a million. No brain on Earth can keep a detailed model of half a million relationships, and none needs to, because no individual-recognition society gets that large.

In practice, these societies reach a size at which they become unstable and split. Tensions grow, subgroups spend more time apart, and eventually a fission occurs. Sometimes the split is peaceful. Sometimes, as at Gombe, the new boundary becomes a front line. Moffett's point is that this ceiling is not a defect of particular species. It is a structural feature of societies based on personal recognition. If belonging requires knowing, then the size of a society is capped by the capacity of its members to know.

There is a second, subtler limit. Individual recognition makes strangers hard to absorb. An animal that arrives from outside starts with no history and therefore no trust. It may be tolerated only after a long period on the margins, or it may never be accepted at all. That makes such societies resistant to migration, trade and the kinds of exchange that allow ideas and goods to move between groups. Chimpanzee communities do exchange members, mostly through young females transferring between groups, but they do not trade or form alliances across the boundary. There is no market among chimpanzee communities, because there is no basis for trusting a chimpanzee you have never met.

Readers who work in organisations will already have noticed the parallel. A start-up of ten people is an individual-recognition society. Everyone knows everyone. Trust is personal, reputations are known, and cheating is quickly noticed. At thirty people the founders still know everyone, but the employees may not. At a hundred, new hires meet colleagues they have never heard of, and teams begin to form their own identities. Somewhere between a hundred and a few hundred, many companies experience a crisis that founders describe in strikingly similar terms. The culture feels as if it is slipping. People who used to solve problems over coffee now need meetings. New arrivals do not understand how things work. Old hands complain that the place has changed. What has happened is that the company has hit the ceiling of personal recognition, and it has not yet built the markers that would let it function as an anonymous society.

The few who broke through

If individual recognition were the only option, large societies would not exist. Yet they do, in two very different branches of the tree of life. Some social insects, especially certain ants, live in colonies of millions or even billions. Humans live in societies of millions and, in the case of the largest nations, more than a billion. Neither ants nor humans can recognise more than a tiny fraction of their fellow members. Both manage to cooperate with those they do not recognise.

Moffett calls these anonymous societies. The word is easy to misread. It does not mean that members are unknown in every sense, or that they have no personal relationships. Humans have friends and families, and some ants may treat nestmates they interact with frequently a little differently. It means that membership does not depend on personal knowledge. You can belong to an anonymous society without being known to most of its members, and you can recognise another member without ever having met them. What makes that possible is a shared set of signals that indicate membership. Ants use chemistry. Humans use culture.

The next two chapters explore those signals in detail. For now it is enough to see the logic. Once membership can be read from a marker rather than recalled from memory, the ceiling on group size disappears. A society can grow as large as its markers can spread and remain legible. A stranger who carries the right marker can be treated as a member from the first moment of contact. A stranger without it, or with the marker of a rival society, can be identified as an outsider just as quickly.

That is the trade at the heart of anonymous societies. They gain scale and lose some of the fine-grained trust that comes from personal knowledge. A wolf knows exactly how much it can rely on each member of its pack. A commuter knows almost nothing about the person standing on her foot. What the commuter has is a set of expectations attached to a broad category of person, a fellow member of her society, and those expectations are good enough to let her ride the train without fear.

Why the trade is not free

It would be comforting to think of anonymous societies simply as an upgrade, the same cooperative instincts applied at greater scale. The truth is more complicated, and the complications run through the rest of this book.

First, markers can be faked. A member of an individual-recognition society is hard to impersonate, because the others know what they look, sound and smell like. A member of an anonymous society can be impersonated by anyone who learns the markers. Spies, con artists, counterfeiters and social parasites all exploit this weakness, and so, as Chapter 2 shows, do certain beetles and caterpillars that infiltrate ant colonies by mimicking their chemistry.

Second, markers can exclude people who ought to be included. If membership is read from a marker, then anyone whose markers differ will be treated with suspicion, however trustworthy they may be as individuals. Immigrants who speak with an accent, employees who dress differently, platform users whose names or photos signal a minority background: all of them may find that the marker speaks before they do.

Third, markers create a new kind of boundary. Individual-recognition societies have boundaries too, but they are fuzzy and personal. Anonymous societies have boundaries defined by category, and categories can be mobilised. The same marker that lets two strangers trust each other on a train can be used to rally a nation against a neighbour, or a workforce against the company that has just acquired it.

Fourth, markers change. A face is relatively stable over a lifetime. A marker such as a flag, a slang word, a logo or a way of dressing can shift in years or months. When markers change, people who carried the old ones may suddenly look like outsiders. Much of the pain of organisational change, as later chapters argue, comes from this effect.

The commuter's hidden curriculum

Consider what a commuter actually reads in those twenty minutes on the train. She notices the language spoken around her and its accents. She notices clothing that signals work, school or leisure, and perhaps religion or subculture. She notices behaviour: who gives up a seat to whom, who keeps their voice down, who stands too close or stares too long. She notices objects such as a briefcase, a tool bag, a university scarf or a football shirt. Most of this processing happens below conscious awareness. If someone behaves in a way that falls outside the expected range, she will feel it before she can explain it. A passenger who shouts, who wears something wildly out of place or who violates the unspoken rules of personal space draws attention, and often a quiet shift of other passengers away.

She also notices the absence of markers. A carriage of a city's morning train contains people of many ethnicities, ages and professions, yet each of them carries enough of the shared markers of the city to be legible. They pay the fare in the expected way. They queue, or at least jostle, according to local norms. They speak the local language when they need to. In a large modern city, the markers that matter for basic trust are surprisingly thin. Much of the work is done by small behaviours: tapping a card, looking away, saying sorry.

This is worth pausing on, because it shows that markers need not be identity badges in the dramatic sense. A flag and a uniform are markers, but so is the knowledge of when to stand on the right of an escalator. In a mature anonymous society, many markers are procedural. They indicate that you know how things are done here, which is the strongest available evidence that you belong.

The same is true in organisations. Employees read one another through dress, vocabulary, meeting habits and the tools they use, but also through procedural knowledge such as how to request a budget, whom to copy on an email and which kinds of disagreement are acceptable in public. A new hire who does not know these things is quickly recognised as a newcomer,

however senior their title. A long-serving employee who suddenly changes them, for instance by adopting a new set of buzzwords after a reorganisation, may be seen as having switched sides.

From the train to the firm

The comparison between chimpanzee communities and organisations can be pushed too far. Companies do not raid one another's offices and leave bodies in the corridors. Yet the comparison is useful precisely because it highlights a structural problem that every large organisation must solve. At some size, personal recognition can no longer hold the group together. Something else must take over, and that something is a system of markers.

The organisations that manage this transition well tend to do three things, often without having a name for them. They build markers deliberately rather than leaving them to accident. They make those markers carry real information about values and behaviour rather than only surface appearance. And they create ways for newcomers to acquire the markers quickly, so that the growing society can absorb strangers without losing its coherence.

The organisations that manage it badly tend to cling to individual recognition long after it has stopped working. Founders try to interview every hire, approve every decision and know everyone by name. When this becomes impossible, they substitute rules for markers, and a culture of trust gives way to a culture of compliance. Alternatively, they let markers grow at random, so that different parts of the company develop incompatible dialects, rituals and loyalties. Either way, the company begins to feel like several smaller societies sharing an office, with all the suspicion that implies.

What the forest teaches

The chimpanzees of Gombe and Ngogo were not monsters. They were acting according to the logic of a society in which trust depends on personal knowledge, and in which a rival community represents a threat to food,

mates and safety. Their violence was the dark side of their loyalty. Within the community, they could be generous and gentle. Beyond it, they could not imagine cooperation.

Human beings carry versions of both tendencies. We are capable of intense loyalty to people we know personally and of suspicion, sometimes hostility, toward people we perceive as outsiders. What distinguishes us is that we can redraw the boundary. By adopting shared markers, we can extend the circle of us to include millions of people we will never meet. By changing or abandoning markers, we can push people out of that circle. The boundary is not fixed by biology. It is drawn and redrawn by culture, and increasingly by design.

That is the opportunity and the responsibility this book addresses. Every leader of a large organisation, every designer of a platform and every architect of a market is in the business of drawing boundaries. The tools they use are markers. They can use them to build trust among strangers, or they can use them carelessly and find that trust collapses at the first contact between groups that had not realised they were different.

A first exercise in noticing

Before moving on, it is worth trying an exercise that will recur in different forms throughout the book. Pick an organisation you know well, perhaps your own. Imagine that a stranger has just walked through its door, physically or digitally. Ask what that stranger would need to know, display or do in order to be treated as a member within the first hour.

Write the answers down. Some will be formal: an access badge, an email address, a login, a job title. Others will be cultural: the right way to dress, the vocabulary to use, the rituals of meetings and messages. Still others will be procedural: knowing whom to ask, how to request help and which tools to use for which purpose.

Then ask a harder question. Which of those markers actually carry information about whether the stranger can be trusted, and which merely signal familiarity? An access badge tells you the person has been hired. It does not tell you whether they are competent or honest. A shared

vocabulary tells you they have absorbed the culture. It does not tell you whether they share its values or are simply good at imitation. The gap between signalling membership and signalling trustworthiness is where many of the problems in this book begin.

A chimpanzee has no such gap to worry about. It knows each member of its community and can judge each one on the basis of personal history. We do not have that luxury. We have chosen, as a species, to live among strangers. The rest of this book is about how we manage it, and how we can manage it better.

CHAPTER 2

Billions Without a Quarrel

If you walk along the coast of California, from the outskirts of San Diego north past Los Angeles and on toward San Francisco, you will pass through what is, by some measures, one of the largest societies on Earth. It is not a human society. It belongs to the Argentine ant, *Linepithema humile*, a small brown species that arrived in the United States as a stowaway in the late nineteenth century and spread across the warm, irrigated landscapes of the state during the twentieth.

In their native range in South America, Argentine ants live in colonies that behave much as ant colonies usually do. Neighbouring colonies are rivals. Workers from one nest attack workers from another. In California something unusual happened. Researchers led by Neil Tsutsui and colleagues found that ants taken from nests hundreds of kilometres apart showed little or no aggression toward one another. Along a stretch of coast roughly nine hundred kilometres long, ants from countless nests behave as members of one enormous colony, sometimes called the Very Large Colony. Its population runs into the billions. Workers and queens move between nests. There is no recognisable boundary inside it, and there is no war.

Later work found that the peace extends even further. In 2009 a team including Eiriku Sunamura reported that Argentine ants from the main supercolonies in Europe, Japan and California did not fight when brought together. Ants that had never shared a continent treated one another as nestmates. The researchers described the result as a global megacolony, the largest known cooperative unit of animals outside our own species.

The Argentine ant is a dramatic case, but it illustrates a principle that runs through Mark Moffett's work. Ants do not know each other. A worker in a colony of millions cannot possibly recognise the others as individuals. What it can recognise is a smell. That smell is a marker, and it is the reason a society of strangers can live without quarrelling.

The chemistry of belonging

An ant's body is covered in a thin waxy layer of compounds called cuticular hydrocarbons. They help prevent water loss, but they also carry information. The precise blend of hydrocarbons differs among species and, within many species, among colonies. Each colony develops a characteristic profile, shaped by the genes of its members, by the food they eat and by the materials of the nest. Ants groom one another and share food mouth to mouth, which spreads the blend among nestmates until the whole colony carries a shared scent.

When two ants meet, they touch antennae. In that brief contact, each samples the other's chemical profile and compares it with an internal template of what home smells like. If the profile matches closely enough, the encounter is calm. If it does not, the ant may become alarmed, bite or spray, or recruit nestmates to attack. Recognition is not based on knowing who the other ant is. It is based on knowing what kind of ant it is.

This is the essential feature of an anonymous society. The marker is attached to the category, not the individual. An ant does not need to have met a nestmate before to treat it as a nestmate. It needs only to check the marker. That is why a colony can grow to millions without losing coherence. Each new worker that emerges from its pupa acquires the colony smell quickly and is accepted without any history.

Markers in ants are not fixed in the genome alone. They are partly learned and partly environmental. Experiments with Argentine ants by Dangsheng Liang and Jules Silverman showed that diet can change an ant's hydrocarbon profile. When ants from a single colony were raised on different prey, they acquired different chemical signatures, and former nestmates sometimes treated one another with aggression when reunited.

Belonging, in other words, can be rewritten by circumstance. Two groups that start with the same marker can drift apart if they live differently for long enough.

Why the Argentine ants made peace

The Californian supercolony raises an obvious question. Why did ants that are rivals at home become peaceful abroad? One leading explanation, developed by Tsutsui and others, points to a genetic bottleneck. The ants that founded the introduced populations were probably few in number, so they carried only a fraction of the genetic variation found in South America. Because the genes that influence recognition became less diverse, colonies descended from the founders smelled much more alike. Ants that would have recognised one another as foreign in Argentina now recognised one another as kin.

Uniform markers, in this account, produced uniform trust. Without the small chemical differences that normally separate colonies, the boundaries dissolved, and the ants stopped spending energy fighting their neighbours. They put that energy instead into foraging, reproducing and expanding. The result was an extraordinary ecological success. Argentine ants have displaced many native ant species in California and in other parts of the world, partly because they do not waste resources on internal war.

The story is not one of permanent harmony, however. In southern California, researchers including Melissa Thomas and David Holway found that the Very Large Colony shares its range with several smaller colonies whose markers differ. Where these colonies meet, they fight with remarkable intensity. Field studies near San Diego documented border battles in which enormous numbers of ants died, with estimates running to millions of deaths over several months along a few kilometres of contested ground. The peace inside the supercolony and the slaughter at its edges are two sides of the same system. Markers that unite also divide.

Traffic without traffic lights

Moffett's own research before he turned to humans concerned an ant very different from the Argentine species. As a doctoral student he studied the marauder ant of South and Southeast Asia, a species whose workers vary dramatically in size, from tiny minors to massive majors many hundreds of times heavier. Marauder ant colonies run to hundreds of thousands of workers. They build and maintain trunk trails, cleared paths that carry traffic between the nest and foraging areas, and they mount sweeping raids in which waves of workers overwhelm insects, seeds and other food.

What fascinated Moffett was the organisation of these operations. Workers of different sizes take on different tasks. Large workers can clear obstacles and defend the trail, while smaller ones carry food. Traffic flows in both directions along the trails without collapsing into gridlock. No ant directs the flow. There is no foreman, no map and no traffic light. Each worker follows simple rules triggered by chemical cues, and the pattern of the whole emerges from countless local interactions.

This is a theme that Donella Meadows would have recognised immediately. A complex system can show coherent behaviour without central control, because its structure, the arrangement of stocks, flows and feedbacks, produces that behaviour. The marauder ants' roads are a physical expression of a feedback loop. Ants that find food lay trail pheromone on the way back. More ants follow the stronger trail and reinforce it. If food runs out, ants stop reinforcing the trail, the chemical evaporates, and traffic shifts elsewhere. The system adjusts itself without anyone being in charge.

For human organisations, the lesson is not that managers are unnecessary. It is that much of the coordination in a large society happens through shared signals rather than direct instruction. When the signals are clear, people can act on them without asking permission. When the signals are confused, coordination requires more and more supervision, until the organisation is spending as much effort managing itself as it spends on its real work.

Farmers before farmers

The leafcutter ants of the American tropics take the logic of anonymous cooperation even further. Species in the genera *Atta* and *Acromyrmex* cut fragments of leaves and carry them back to underground nests. They do not eat the leaves. They chew them into a mulch on which they cultivate a fungus, and the fungus becomes their food. The farming relationship is ancient, with origins estimated at tens of millions of years, long before any human planted a seed.

A mature leafcutter colony can contain millions of workers organised into castes of different sizes. Large soldiers defend the nest. Medium workers cut and carry leaves. Small workers tend the fungus gardens, clean the crop and remove contaminants. Some of the smallest ride on leaf fragments and appear to defend the carriers from parasitic flies. The nest itself can extend many metres underground, with chambers for fungus, waste and ventilation.

The scale of this cooperation depends entirely on markers. Workers do not know one another, yet they divide labour among millions of individuals with great precision. They share food, care for the brood and defend the colony against intruders whose scent differs. The fungus garden itself is part of the colony's identity. Its chemistry contributes to the smell of the workers, and alien fungi are recognised and removed. Membership, in effect, extends to the crop.

It is tempting to draw a direct analogy with human agriculture, and Moffett does note the parallels. The more interesting point for this book is structural. Large-scale division of labour requires that workers trust one another to do their part, without supervising one another directly. In ants that trust is underwritten by markers that are hard to fake and constantly refreshed by contact. In human organisations it is underwritten by a mix of markers, rules and incentives, and the balance among them determines how much supervision the organisation needs.

Counterfeit nestmates

Ant markers are robust, but they are not perfect, and the imperfections reveal much about how markers work. Many creatures have learned to exploit them. Certain beetles, mites, flies and spiders live inside ant nests by mimicking the colony's chemical profile or acquiring it through contact. Once inside, they may feed on the colony's food, its brood or both, while the workers treat them as members.

One of the best-studied examples is the large blue butterfly and its relatives in the genus *Phengaris*, formerly known as *Maculinea*. Their caterpillars drop from the plants on which they hatch and wait to be found by worker ants of a particular species. The caterpillars produce surface chemicals that resemble those of the ants' own larvae, so the workers carry them back to the nest as if rescuing lost brood. Some species then feed on ant larvae, while others are fed by the workers directly. Research by Francesca Barbero and colleagues, published in *Science* in 2009, found that the caterpillars of some species also produce sounds that mimic those of the ant queen, apparently raising their status inside the colony so that workers prioritise them even over their own young when the nest is disturbed.

A caterpillar that pretends to be a queen is about as clear an illustration as one could want of the central weakness of anonymous societies. Because membership is read from a marker, anyone who can produce the marker can obtain the benefits of membership. The better the forgery, the deeper the access. In the case of these butterflies, a forged status signal does more than gain admission. It diverts the colony's resources toward the forger.

Human societies are full of equivalents. Impostors who pose as doctors, pilots or police officers, fraudsters who assume the identity of a trusted supplier, attackers who join an open-source project and win the confidence of maintainers: each exploits the fact that an anonymous society cannot check every member individually. Chapters 16 and 25 explore human counterfeits in detail. For now, note the structural lesson. The value of a marker is directly linked to the cost of forging it. Markers that are cheap to copy invite parasites.

Colonies that cannot absorb

Ant markers have another limitation that matters for organisations. They are good at telling members from non-members, but they make it hard to merge two groups. In most ant species, workers from two colonies will fight if brought together, even if they belong to the same species and live in identical conditions. Their markers differ, and the difference triggers hostility. There are exceptions. Some species tolerate multiple queens and can fuse colonies under certain conditions, and researchers can sometimes blend colonies by mixing their scents. But the default is that two societies with different markers do not combine peacefully.

This is not because the ants are stubborn. It is because their recognition system has no mechanism for accepting an outsider as an insider except by changing the outsider's markers. An ant introduced into a foreign nest may sometimes survive if it is held apart until it acquires the local scent, but that takes time and often fails. The colony's capacity to absorb strangers depends on its capacity to rewrite their markers.

Humans have a far richer toolkit for absorbing strangers, including language, ritual, law and explicit negotiation. Yet the underlying dynamic remains. When two human organisations merge, they bring different markers into contact. Employees from each side read the others' vocabulary, dress, rituals and procedures as foreign. Unless there is a deliberate process for building shared markers, or for allowing distinct markers to coexist within a shared identity, the encounter produces friction. In extreme cases it produces something close to the border wars of San Diego, with talent leaving in large numbers and whole divisions effectively destroyed.

The price of uniformity

The Argentine ant's success carries a warning. Its supercolonies are peaceful partly because their markers became uniform after a genetic bottleneck. That uniformity allowed vast scale and ecological dominance. It may also carry long-term costs. Low genetic diversity can leave a

population more vulnerable to disease, and some researchers have reported local collapses of Argentine ant populations in places where they once dominated. The causes are still being studied, but the possibility illustrates a general principle. A society held together by a single uniform marker may be efficient and peaceful, but it may also be fragile in ways that are not visible until conditions change.

Organisations face a similar trade-off. A strong, uniform culture can make cooperation easy. Everyone speaks the same language, shares the same assumptions and reads the same signals. Decisions are fast, conflict is low and newcomers either fit in quickly or leave. The same uniformity can make the organisation blind to threats that fall outside its shared assumptions and hostile to people whose markers differ. Diversity of markers creates friction, but it also creates the variety that allows a system to adapt. Eric Beinhocker's account of economic evolution, introduced in Part II, makes the same point from a different angle. Without differentiation, there is nothing for selection to work on.

Lessons from the colony

It would be easy to take the wrong lessons from ants. Human organisations are not superorganisms, and people are not workers driven by chemical cues. People have personal relationships, private goals and the capacity to question the markers they are given. The analogy is valuable not because people resemble ants but because both ants and humans face the same structural problem of cooperating at scale without personal knowledge, and the solutions share certain features.

Three of those features are worth carrying forward.

The first is that markers attach to categories, not individuals. They allow a member of an anonymous society to treat a stranger as an insider without knowing anything personal about them. That is what makes scale possible, and it is also what makes forgery possible.

The second is that markers are maintained by constant contact. Ants refresh their colony scent by grooming and sharing food. If contact stops, markers drift. Human organisations maintain their markers through

meetings, rituals, shared tools and stories. Remote work, rapid growth and reorganisation all reduce contact, and they can cause markers to drift in ways that leaders notice only when groups that used to cooperate begin to regard each other as foreign.

The third is that markers both unite and divide. The same signal that lets billions of Argentine ants live in peace triggers mass slaughter at the boundary with a colony whose signal differs. Human markers work the same way. A shared identity creates trust inside and suspicion outside. The art of designing markers, which Part V takes up in detail, is largely the art of drawing the boundary in the right place and making it porous enough for the right strangers to cross.

Moffett observes that ants and humans are the champions of anonymous society. Ants got there first, by many millions of years, and they got there through chemistry. Humans got there through something far more flexible and far more dangerous: culture. The next chapter turns to the markers that humans use and to the surprising fact that we begin reading them almost as soon as we are born.

CHAPTER 3

Three Fingers in a Basement

Midway through Quentin Tarantino's 2009 film *Inglourious Basterds*, a British officer named Archie Hicox sits in a cramped basement tavern in occupied France, disguised as a German officer. He is there with two German-speaking allies to meet a double agent. The plan depends on his being taken for a German, and for a while it works, despite his slightly unusual accent, which he explains away by claiming to come from a remote mountain region. A Gestapo major named Hellstrom joins the table, suspicious but charming. The tension builds through conversation, a card game and a round of drinks. Then Hicox orders three more glasses of whisky. He holds up his hand with the index, middle and ring fingers extended, the thumb and little finger folded.

The major's face changes. As the scene makes clear, a German counting to three would use the thumb, index and middle finger. Hicox has used the British gesture. Everything else about his disguise, the uniform, the language, the rank, the manner, had passed. A tiny habit of the hand, learned in childhood and never consciously examined, gave him away.

Mark Moffett uses this scene to illustrate a point about human markers that is easy to state and hard to overstate. Our societies are held together not only by obvious symbols such as flags and uniforms but by an enormous repertoire of small behaviours that we learn without noticing and display without thinking. These behaviours tell other members of our society that we belong. They also tell outsiders, or people trained to spot outsiders, that we do not.

The repertoire of belonging

Human markers can be sorted roughly into a few kinds, though in practice they overlap.

Some are explicit symbols. Flags, anthems, coats of arms, religious emblems, team colours and corporate logos all exist mainly to signal membership. They are deliberately designed, widely recognised and often treated with reverence. People will die for flags, and they will also burn them, which shows how much meaning those pieces of cloth have been asked to carry.

Some are bodily. Clothing, hairstyles, tattoos, scarification, jewellery and posture all carry information about where a person belongs. Many of these can be changed, which makes them useful for signalling chosen identities as well as inherited ones.

Some are linguistic. Language itself is a marker, and so, at a finer grain, are dialect, accent, vocabulary and idiom. Linguistic markers are among the most powerful because they are hard to fake. Most adults who learn a new language retain traces of their first in their pronunciation, and listeners are remarkably sensitive to those traces.

Some are culinary. What a group eats, how it prepares food and what it refuses to eat are among the oldest and most persistent markers. Food rules such as kosher and halal are explicitly religious. Others are simply habits: whether to eat with hands or utensils, which animals are food and which are pets, what time dinner happens.

Some are procedural. These are the rules of everyday interaction: how close to stand, how long to make eye contact, how to greet, how to queue, how to count on your fingers. They are the least visible markers and in some ways the most revealing, because people rarely think to disguise them.

Moffett's argument is that the sheer number and variety of these markers is what makes human anonymous societies so robust. No single marker needs to carry the whole burden. A stranger who speaks your language with your accent, dresses roughly as you do, eats what you eat and

counts as you count is overwhelmingly likely to be a member of your society, even if you have never met. A stranger who differs on several of these dimensions is likely to be an outsider, even if they wear the right uniform.

Shibboleths

The idea that a small behaviour can reveal membership is ancient. The Hebrew Bible, in the twelfth chapter of the Book of Judges, tells of a conflict between the Gileadites and the Ephraimites. After defeating the Ephraimites, the Gileadites seized the fords of the Jordan River. When fleeing men tried to cross, the guards asked them to say the word *shibboleth*. The Ephraimites, according to the story, could not pronounce the "sh" sound and said *sibboleth* instead. Those who failed the test were killed. The text gives the number of dead as forty-two thousand.

Whatever the historical basis of the story, the word has passed into English as a name for any custom, phrase or habit that distinguishes members of a group from outsiders. Shibboleths have been used in many conflicts since. In 1937, according to many historical accounts, soldiers under the Dominican dictator Rafael Trujillo massacred thousands of Haitians and Dominicans of Haitian descent near the border, and some accounts say that people were asked to pronounce the Spanish word for parsley, *perejil*, with those who rolled the "r" in a way associated with Haitian Creole speakers singled out for death. Historians debate how widely the test was actually used, but the episode is remembered in the Dominican Republic and Haiti as the Parsley Massacre, and the name itself shows how readily people understand the logic of a linguistic marker as a sorting device.

These are extreme examples, and they belong in this book precisely because they are extreme. Markers are not neutral. They evolved, in part, as tools for distinguishing allies from enemies. In conditions of conflict, a marker that was once a harmless badge of identity can become a death sentence. Any leader who designs markers needs to remember that the sorting function never entirely disappears.

Markers in other animals

Humans are not the only animals with learned markers, and a brief look at others helps clarify what is distinctive about us. Several bird species, including white-crowned sparrows studied by Peter Marler and colleagues, have regional song dialects. Young males learn the songs they hear during a sensitive period, so birds from different areas sing recognisably different versions. Some whales show similar patterns. Sperm whales in parts of the Pacific belong to vocal clans, identified by researchers including Hal Whitehead through their distinctive repertoires of rhythmic click patterns called codas. Clans with different codas can overlap geographically but tend not to associate with one another.

These examples suggest that learned group signals have evolved more than once. What sets humans apart is not the existence of markers but their number, their flexibility and their role in creating societies of millions. A sparrow's dialect distinguishes neighbours from distant birds. A human's accent, clothing, diet and gestures together place them in a society that may span a continent. And unlike birds, humans can invent new markers deliberately, debate them, change them by law and adopt them as adults.

Learned early, held deeply

Many human markers are acquired in childhood and become so deeply ingrained that people experience them as natural. An accent is the classic case. Children who grow up speaking a language acquire its sounds with ease, but adults who learn the same language later in life usually retain a foreign accent no matter how fluent they become. The ability to perceive and produce fine phonetic distinctions narrows in the first year of life. By around their first birthday, infants raised in an English-speaking environment have become less sensitive to sound contrasts that do not matter in English but do matter in other languages, a pattern documented by Janet Werker and colleagues in a series of studies beginning in the 1980s.

The same is true of many procedural markers. Children learn how close to stand to other people, how to greet adults, how to use eating utensils and how to count on their fingers well before they can articulate any rules. By the time they are adults, these behaviours feel like part of who they are. They are rarely examined, and they are difficult to change even when a person wants to.

This has two consequences. First, deeply learned markers are reliable. Because they are hard to fake, they provide strong evidence of membership. A spy can learn a language, but an accent or an unconscious gesture may still betray them, as Hicox discovered. Second, deeply learned markers are sticky. People who carry them cannot easily shed them, even when they move to a new society. Immigrants may find that their accent marks them as outsiders for the rest of their lives, whatever their citizenship, loyalty or contribution.

The same stickiness affects organisations. Employees who joined a company under one regime carry its markers: its vocabulary, its assumptions about how decisions are made, its habits of dress and communication. When the company changes, through a merger, a new chief executive or a strategic pivot, those markers do not change overnight. They are not simply opinions that can be updated by a memo. They are habits, sometimes decades old, that feel like part of the person. Expecting people to discard them quickly is like expecting an adult to lose an accent by deciding to.

Food, the oldest flag

Food deserves special attention because it shows how markers combine practical and symbolic roles. Every human society has rules about what may be eaten, how it should be prepared and with whom it should be shared. Some of these rules have practical origins, though attempts to explain them all as health measures have not held up well. Many simply mark the boundaries of a community. Eating together is a sign of trust, and refusing to eat what others eat is a sign of difference.

Food markers can be surprisingly political. In 2003, after France opposed the invasion of Iraq, the cafeterias in the United States House of Representatives renamed French fries as "freedom fries" on their menus. The change was trivial in practical terms and was quietly reversed a few years later, but it attracted wide attention because people understood exactly what it meant. A food associated with one nation had been rebranded to remove the association, as a gesture of national loyalty.

In organisations, food plays a smaller but real role. Shared meals, team lunches and the rituals of the office kitchen are among the ways colleagues signal belonging. Companies that provide free food are not only offering a benefit. They are creating a space in which strangers become familiar through repeated contact. When those rituals change, for example after a move to remote work, people often report a loss that is hard to name. They have lost one of the markers that told them they were part of the same group.

Flags and the sacred

At the other end of the spectrum from food habits lie explicit symbols such as flags. National flags are a relatively modern invention in their current form, and many date only from the nineteenth or twentieth century. Yet within a few generations they can acquire an emotional power that rivals ancient religious symbols. People salute them, fold them according to precise rules, drape them over coffins and argue bitterly about whether burning them should be legal.

The power of a flag does not come from the cloth. It comes from the fact that the flag stands for the whole anonymous society. When you salute a flag, you are expressing loyalty to millions of people you will never meet. When you burn one, you are rejecting that loyalty, or protesting against what you believe the society has become. The emotional charge reflects the stakes. The flag is the marker that defines the boundary of the group, and attacks on it feel like attacks on the group itself.

Corporate logos aspire to a version of the same power, and occasionally achieve it. Customers who tattoo a brand's logo on their bodies, employees who wear company clothing outside work and investors who treat a company's symbol as a sign of quality are all treating the logo as more than a trademark. Chapter 14 explores what happens when companies change their logos and discover, sometimes to their cost, how much meaning those symbols carried.

Many markers, many groups

One complication must be acknowledged early. Humans belong to many groups at once. A person may be a citizen of one country, a member of a religious community, a supporter of a football club, an alumnus of a university, an employee of a company and a member of a profession. Each group has its own markers, and many markers overlap. A person can display different markers in different contexts, switching vocabulary, dress and behaviour as they move between home, work and social life.

Moffett distinguishes between societies, which in his usage are the large groups that define a person's primary membership, such as a nation or, in hunter-gatherer times, a people, and the many smaller groups that exist within them. Not everyone uses the terms the same way, but the distinction is useful. Some markers define the outer boundary of a society, separating members from foreigners. Others define groups within it, separating supporters of one team from another or employees of one company from another. The inner groups compete and cooperate within a shared framework set by the outer markers.

Organisations sit in the middle of this structure. They are smaller than nations but larger than personal networks. They draw on markers from the wider society, such as language and law, while developing their own. A company's culture is, in part, a set of markers that distinguish its members from employees of other companies. Within the company, departments, teams and offices develop their own markers too. A healthy organisation manages these layers so that local identities strengthen rather than undermine the shared one.

The gesture that gives *you* away

Return to the basement tavern. What made Hicox's mistake fatal was not that the gesture was important in itself. Counting to three with one set of fingers rather than another has no practical significance. It mattered because it was a marker that Hicox had not known he carried, and because his opponent was trained to look for exactly such markers.

The scene captures something leaders often miss. When people from different groups meet, they read each other through many small signals, most of them unconscious. A newcomer to an organisation may master the official language of values and strategy, yet still be marked as an outsider by the way they write emails, schedule meetings or respond to criticism. An acquired company's employees may adopt the parent company's logo and job titles, yet still be identified as different by their habits, their jargon and their assumptions about how decisions are made.

None of this is a reason for despair. It is a reason for attention. Markers that are visible can be discussed, taught and, if necessary, changed. Markers that are invisible operate without anyone's consent. The first step in managing markers is to notice them, and the most important markers to notice are often the smallest. The next chapter looks at how early that noticing begins, and at what experiments with infants and arbitrary groups reveal about the minds that carry markers.

CHAPTER 4

Born Reading Markers

A three-month-old baby cannot sit up unaided, cannot reach reliably for a toy and has no words. Yet researchers who study infant perception have found that babies of this age already show preferences that look, in a faint and early form, like the beginnings of group identity.

In a study published in 2005, David Kelly and colleagues showed pairs of faces to infants and measured how long the babies looked at each. Newborns showed no particular preference between faces of their own ethnic group and faces of other groups. By three months, infants raised among people of one ethnicity tended to look longer at faces from that group. A few months later, another pattern appeared. Infants became better at telling apart faces from the group they saw most and worse at telling apart faces from groups they rarely saw. This narrowing of perception, which resembles the way infants tune their hearing to the sounds of their native language, has been documented by several research teams and is often called the other-race effect.

The researchers were careful not to describe these preferences as prejudice. They reflect exposure. A baby who sees mostly one kind of face becomes expert in that kind of face. Infants raised in diverse environments show weaker effects, and the effects can be modified by experience. What the research shows is not that humans are born biased but that we are born ready to learn which cues mark the people around us, and that this learning begins long before we can talk.

Mark Moffett draws on this body of work to make a point about markers. If infants begin sorting the social world by familiar and unfamiliar cues within months of birth, then reading markers is not a sophisticated adult skill. It is a basic part of how human minds develop. Markers are absorbed with the milk.

The voice of home

Faces are only one cue. Language is another, and the evidence here is particularly striking. In a series of experiments reported in 2007, Katherine Kinzler, Emmanuel Dupoux and Elizabeth Spelke studied how infants respond to speakers of their native language and of a foreign language. Five- and six-month-old infants looked longer at people who had previously spoken to them in their native language. Ten-month-olds were more likely to accept a toy offered by a native speaker than by a speaker of a foreign language, even though the toy was offered silently.

Later work by Kinzler and colleagues extended the findings to accent. Children of preschool age preferred to befriend children who spoke with their own accent over children who spoke the same language with a foreign accent. In some studies, accent influenced children's choices of friends more strongly than ethnicity did. A child who looked different but sounded familiar was preferred to a child who looked familiar but sounded different.

This fits neatly with what Chapter 3 said about accent as a marker that is hard to fake. Children appear to treat accent as reliable evidence of group membership because, in the environments in which humans evolved, it was. People who spoke like you had grown up near you, among people like the ones you knew. People who spoke differently came from somewhere else. The preference for familiar speech is a heuristic, a rough guide to who is likely to share your norms.

Liking those who like what we like

Infants also read markers that have nothing to do with appearance or language. In a study published in 2012, Neha Mahajan and Karen Wynn showed infants of about eleven months puppets that expressed preferences for one food or another. Infants who had themselves chosen one food tended to prefer the puppet that shared their choice. A related study led by Kiley Hamlin and colleagues, published in 2013, found something more unsettling. When infants watched a puppet with a different food preference being helped or hindered by other puppets, many preferred the puppet that hindered the dissimilar one.

These studies are small, and infant research is notoriously hard to replicate. Results from a few dozen babies should not be treated as laws of nature. But the pattern they suggest is consistent with a much larger literature on older children and adults. Humans tend to like people who share their tastes, and they can extend that preference into a willingness to see dissimilar people treated badly. Shared preferences, in this light, are markers too. A trivial choice, such as which snack you prefer, becomes evidence of whether you are one of us.

Klee or Kandinsky

The ease with which arbitrary differences become markers was demonstrated most famously in a series of experiments by the social psychologist Henri Tajfel and his colleagues at the University of Bristol in the early 1970s. Tajfel was interested in the minimum conditions needed to produce favouritism toward one's own group. He showed schoolboys pictures by the painters Paul Klee and Wassily Kandinsky and asked which they preferred. The boys were then told, on the basis of their answers, that they belonged to the Klee group or the Kandinsky group. In some versions, the assignment was in fact random.

Next, each boy was asked to divide small sums of money between other boys, identified only by code number and group membership. He never knew who they were and would not receive any of the money himself. The boys consistently gave more to members of their own group. More strikingly, many chose allocations that maximised the difference between groups, even when that meant their own group received less in absolute terms than it could have. Being ahead of the other group mattered more than getting the most.

Tajfel called this the minimal group paradigm. The groups had no history, no shared goals, no contact and no meaningful difference. The only marker was a label based on a passing preference for a painter. Yet the label was enough to produce favouritism. Tajfel and his colleague John Turner later developed social identity theory, which holds that people derive part of their self-esteem from the groups they belong to and are therefore motivated to see those groups as better than others.

The minimal group findings have been replicated many times, with variations that test their limits. The effect is typically modest, and it can be reduced or reversed by context, for instance when people are told that fairness matters or when they expect to be held accountable for their choices. But the core finding has proved durable. Humans can form group loyalties around almost any distinction, provided it is presented as a marker of membership.

Robbers Cave, reconsidered

A more dramatic study is often cited alongside Tajfel's work, and it deserves a more careful treatment than it usually receives. In 1954 the psychologist Muzafer Sherif and his colleagues ran an experiment at a summer camp in Robbers Cave State Park in Oklahoma. Two groups of boys were brought to the camp separately and given time to form identities. They chose names, the Rattlers and the Eagles. The researchers then brought the groups into contact through competitions, and hostility followed. There

were raids, insults and flag burning. Later, the researchers introduced tasks that required the groups to cooperate toward shared goals, such as fixing a water supply, and the hostility reportedly diminished.

The study became a textbook example of how competition creates conflict and shared goals reduce it. More recent historical work has complicated the story. The psychologist Gina Perry, in her book *The Lost Boys*, examined archival records and argued that the researchers intervened more actively than the published account suggested, shaping the conflict they reported. An earlier attempt by Sherif's team, at a camp in Middle Grove, New York, in 1953, had not gone as planned. In Perry's account, the boys in that study largely refused to turn against each other, and the experiment was abandoned.

None of this means that intergroup conflict is a myth. It means that Robbers Cave should be read as a vivid illustration rather than a clean demonstration, and that the boys in these studies were not passive subjects responding automatically to group labels. They were active participants who sometimes resisted the roles the experimenters seemed to want them to play. The lesson for leaders is important. Group markers create a readiness for favouritism and suspicion, but the way that readiness plays out depends heavily on context, leadership and the explicit or implicit messages people receive about how they are expected to behave.

Why arbitrary markers matter at work

Organisations create minimal groups constantly, usually without intending to. Employees are divided into teams, departments, offices, cohorts and project groups. Many of these divisions are practical. Some are close to arbitrary. People are assigned to the Blue team or the Red team for an exercise, to an onboarding cohort by start date, to a floor by the availability of desks. Each of these labels can become a marker.

Sometimes this is useful. A team with a strong identity may work harder, support its members and take pride in its results. Companies deliberately foster such identities through team names, rituals and shared goals. The problem arises when team identity becomes stronger than the

identity of the wider organisation, and when teams begin to compete in ways that harm the whole. Sales and engineering, head office and field operations, the acquired company and the acquirer: each pair can slide into a minimal group dynamic in which being ahead of the other side matters more than achieving the shared goal.

Tajfel's finding that boys preferred to maximise the gap between groups, even at a cost to their own group, has a direct analogue in organisations. Departments sometimes hoard information, resist collaboration or block initiatives from other parts of the company, not because they gain anything but because they do not want the other group to win. Budget processes that pit units against one another, incentive systems that reward local results over company results and reorganisations that create winners and losers all strengthen these dynamics.

The research also points to remedies. Minimal group effects weaken when people are reminded of a shared superordinate identity, when fairness norms are made salient and when individuals expect to be accountable for how they treat members of other groups. These are not soft interventions. They are structural features that leaders can build into how teams are formed, rewarded and evaluated.

Markers as predictions

Why would humans be so ready to form group loyalties on the basis of thin cues? One answer, which runs through the work of researchers on the evolution of cooperation, is that markers function as predictions. In an environment where cooperation with strangers is valuable but risky, a cue that predicts shared norms is useful. People who speak your dialect, eat your food and follow your customs are more likely to understand your expectations, respond to your signals and be bound by the same reputational and moral systems that bind you. Cooperating with them is less risky than cooperating with someone whose norms you cannot predict.

The anthropologists Richard McElreath, Robert Boyd and Peter Richerson developed a formal model of this idea in 2003, arguing that ethnic markers could evolve because they help people coordinate with

others who share their norms. In their model, markers do not need to carry any intrinsic meaning. They just need to be correlated with norms. Once that correlation exists, people who prefer to interact with others carrying the same marker do better than those who do not, and the correlation is reinforced.

This model has an important implication. Markers are useful only insofar as they predict behaviour. If a marker stops being correlated with norms, for example because many people who carry it no longer share the expected values, or because outsiders can adopt it easily, its predictive power erodes. People may continue to rely on it out of habit, which leads to mistakes. Or they may stop relying on it, which forces them to look for new cues.

In organisations, this dynamic appears whenever markers drift away from the behaviour they once predicted. A degree from a particular university may once have predicted a certain set of skills and attitudes. If the university changes, or if the skills needed for the job change, the degree may continue to be treated as a marker long after it has stopped predicting anything useful. Hiring managers who rely on it are using an outdated prediction. Chapter 15 explores this problem in detail.

Learning new markers

If markers were fixed at birth, there would be little point in a book about designing them. The good news from the developmental research is that markers are learned, and learning continues throughout life. Infants exposed to many kinds of faces show weaker other-race effects. Children who grow up among speakers of several languages are less likely to treat accent as a strong marker of difference. Adults who join new groups adopt their markers, sometimes quickly, often in ways they do not notice.

Organisations are among the most powerful environments for learning new markers. People who join a company spend a large share of their waking hours in its culture. They absorb its vocabulary, its rituals and its

unwritten rules. Within months, many find that they speak and behave differently, at least at work. After years, the company's markers may become part of their identity, so much so that leaving feels like emigrating.

This capacity to learn markers is what makes anonymous societies possible at scale. A nation can absorb immigrants, a company can absorb new hires and a platform can absorb new users because people are capable of adopting new markers. But the process is slow, uneven and emotionally charged. People who feel that their old markers are being rejected may resist the new ones. People who adopt new markers may be suspected of disloyalty by those who share their old ones. The developmental research reminds us that marker learning is not merely intellectual. It reaches down into the earliest layers of perception and preference.

What infants teach leaders

The research on infants and minimal groups can be summarised in a few propositions, each of which has practical consequences.

Humans begin reading markers very early, before language and before conscious reasoning. Markers therefore operate largely beneath awareness, and people often cannot say why they feel comfortable with some colleagues and uneasy with others.

Almost any distinction can become a marker if it is presented as a sign of membership. Leaders who create teams, labels and categories are creating potential markers, whether they intend to or not.

Markers produce favouritism toward insiders and, under some conditions, indifference or hostility toward outsiders. The strength of these effects depends heavily on context, including how leaders frame the relationship between groups.

Markers are learned and can be relearned. People exposed to diverse groups and given shared identities can extend trust beyond their original circle.

The most famous studies in this area are more complicated than their textbook versions suggest. Leaders should be wary of claims that people are inevitably tribal, and equally wary of claims that group identity can be engineered at will.

With these propositions in hand, we can turn to one of the most widely quoted ideas in discussions of group size: the claim that humans can maintain only about one hundred and fifty stable relationships. The idea has shaped how companies design teams and offices, and it deserves a careful look.

CHAPTER 5

How Many People Can You Know?

Few numbers from the behavioural sciences have travelled as far as 150. It appears in management books, conference talks, software design documents and office floor plans. It has been used to justify the size of military units, the layout of factories and the maximum membership of online communities. It has its own name, Dunbar's number, after the British anthropologist and evolutionary psychologist Robin Dunbar, who proposed it in the early 1990s.

The idea is appealing because it seems to explain an experience many people share. There is a point at which a group stops feeling like a community and starts feeling like a crowd. You can know everyone in a village, a small company or a congregation. You cannot know everyone in a city, a corporation or a national church. Somewhere between the two lies a threshold, and Dunbar offered a number for it.

For this book the number matters because it is, in effect, a claim about the ceiling of individual-recognition societies described in Chapter 1. If humans can maintain only about 150 personal relationships, then any human group much larger than that must rely on something other than personal knowledge to hold together. That something is markers. Understanding what the evidence says about Dunbar's number, and what it does not say, helps clarify when markers become necessary and why.

Where the number came from

Dunbar's original argument, published in the *Journal of Human Evolution* in 1992, started with primates. Many primates live in social groups, and group size varies widely across species. Dunbar noticed that average group size was correlated with the relative size of the neocortex, the part of the brain associated with complex cognition. Species with larger neocortices relative to the rest of the brain tended to live in larger groups. He proposed that this relationship reflected a cognitive constraint. Maintaining social relationships is demanding. You need to track who is allied with whom, who owes whom a favour and who is likely to cheat. The bigger the group, the more information you must handle, and the bigger the brain you need.

Dunbar then extrapolated from the primate data to humans. Plugging the human neocortex ratio into the regression equation produced a predicted group size of about 148, which he rounded to 150. He went on to look for evidence that human groups cluster around that size. He pointed to estimates of the size of hunter-gatherer clans, Neolithic villages, military companies and religious communities. Over the following decades, he and his collaborators developed the idea further, proposing that human social networks are structured in layers. An inner core of about five intimate relationships sits inside a circle of about fifteen close friends, then about fifty good friends, then about 150 meaningful relationships, with further layers of acquaintances beyond. In a 2005 paper with Wei-Xing Zhou, Didier Sornette and Russell Hill, Dunbar argued that each layer is roughly three times the size of the one inside it.

Dunbar also proposed an explanation for how humans manage such large networks compared with other primates. Monkeys and apes maintain relationships largely through grooming, which is time-consuming and can only be done with one partner at a time. Dunbar argued that language evolved partly as a more efficient form of social grooming. Gossip, in this view, lets humans keep track of many relationships at once by exchanging information about absent third parties. He laid out the argument in his 1996 book *Grooming, Gossip and the Evolution of Language*.

Evidence that seems to fit

Dunbar and his colleagues have reported a range of findings consistent with the idea. A study with Russell Hill, published in 2003, asked British participants to list the people to whom they sent Christmas cards and found that the networks implied were in the low-to-mid hundreds, a range broadly consistent with the prediction. Analyses of online interaction have found similar patterns. A 2011 study of Twitter users by Bruno Gonçalves, Nicola Perra and Alessandro Vespignani reported that users tended to maintain stable, reciprocal interactions with between about one hundred and two hundred others, even when they followed many more accounts. Dunbar's own 2016 study of social media use found that people with hundreds of online friends nonetheless reported only a handful whom they could rely on in a crisis.

Organisational examples are often cited too. In *The Tipping Point*, published in 2000, Malcolm Gladwell described how W. L. Gore and Associates, the maker of Gore-Tex, was reported to keep its plants to about 150 people, on the theory that beyond that size people stopped knowing one another and informal coordination broke down. Hutterite colonies, communal religious settlements in North America, are often described as splitting when they grow beyond roughly the same size. These examples have become part of the folklore of organisational design.

Evidence that does not

The number has also attracted sustained criticism. One line of criticism concerns the primate data. In 2021 Patrik Lindenfors, Andreas Wartel and Johan Lind of Stockholm University reanalysed the relationship between brain size and group size using updated datasets and modern statistical methods that account for the evolutionary relationships among species. Their estimates of average human group size varied widely depending on the method and data. Bayesian analyses produced averages between roughly 69 and 109, and generalised least-squares analyses produced averages

between roughly 16 and 42. The confidence intervals around these estimates were enormous, ranging from single digits to several hundred. The authors concluded that specifying any one number was futile and that a cognitive limit on human group size could not be derived from this kind of analysis.

Dunbar responded publicly, arguing that the critique misunderstood the theory and that the number had been supported by many independent lines of evidence beyond the original regression. The exchange illustrates a familiar pattern in science. A bold claim draws on several kinds of evidence, critics attack the weakest part, and defenders point to the rest.

A second line of criticism concerns human networks directly. Estimates of the number of people an individual knows vary greatly depending on how knowing is defined and measured. Network researchers including Christopher McCarty, Peter Killworth and Russell Bernard, using a range of survey techniques, have produced estimates of average personal network size that are considerably larger than 150, on the order of nearly three hundred people in some of their work. Individual variation is large. Some people maintain far bigger networks than others, depending on personality, profession, culture and life stage.

A third line of criticism concerns the organisational examples. Many of the stories about factories and villages are anecdotes, sometimes repeated from secondary sources without verification. Groups of around 150 certainly exist, but so do successful groups of very different sizes. Whether any particular threshold reflects a cognitive limit rather than economic, cultural or practical constraints is hard to establish.

What survives the debate

The honest conclusion is that Dunbar's number should not be treated as a precise constant. There is no strong evidence that 150 marks a sharp boundary beyond which human relationships break down. The exact figure depends on how relationships are defined, which methods are used and which populations are studied.

But the broader idea that motivated the number survives the debate largely intact, and it is the broader idea that matters for this book. Human beings can maintain only a limited number of meaningful personal relationships. Whether that limit is fifty, 150 or three hundred, it is far smaller than the size of the societies in which most people now live. A nation of millions, a company of tens of thousands or a platform of hundreds of millions cannot be held together by personal knowledge alone. Moffett makes the same point from a different angle. Human societies became anonymous not because humans lost the capacity for personal relationships but because their societies grew far beyond what personal relationships could support.

Seen this way, the debate over Dunbar's number is less about the precise limit than about what happens beyond it. Below the limit, whatever it is, groups can rely heavily on personal knowledge. Members know each other, reputations are tracked through direct experience and gossip, and trust is personal. Above it, groups need additional mechanisms. Formal rules, hierarchies and contracts are part of the answer. Markers are another. They allow members of a large group to recognise one another as members without needing to know each other personally.

The transition zone

For organisations, the most important insight may be that the shift from personal to anonymous trust is not a single moment. It happens over a range of sizes, and it happens unevenly across the organisation.

Consider a company growing from thirty people to five hundred. At thirty, almost everyone knows everyone. At eighty, the founders still know everyone, but some employees do not. At 150, small teams know each other well, but cross-team relationships are thinner. At three hundred, most employees know a minority of their colleagues personally, and new hires may never meet people in other departments. At five hundred, the company is an anonymous society in Moffett's sense, held together largely by shared markers, formal structures and the personal networks of a few connectors.

Throughout this transition, the company's mechanisms of trust need to change. Practices that worked at thirty, such as informal decision making, open access to information and trust based on personal relationships, may fail at three hundred. Leaders often notice the failure as a loss of culture. What they are noticing, in many cases, is that the company has crossed into a zone where personal recognition no longer covers enough of the organisation and markers have not yet been built to replace it.

The transition zone is dangerous because people experience it as decline rather than change. Old hands remember a time when everyone knew everyone and see newcomers as threats to that intimacy. Newcomers feel like outsiders because the markers they need to learn are informal and unwritten. Teams begin to form their own identities, which may compete with the company's. Without deliberate attention, the company may split into a set of individual-recognition societies, each internally cohesive and mutually suspicious, rather than becoming a single anonymous society with shared markers.

Designing around a limit you cannot measure

How should leaders use an idea that is both useful and imprecise? A few principles follow from the evidence.

First, treat the limit as a range rather than a number. Rather than designing every unit to contain exactly 150 people, recognise that personal relationships thin out as groups grow, and that somewhere in the range of tens to a few hundred people, informal coordination begins to require support.

Second, notice the layers. Dunbar's layered model may not be precise, but the idea that people have a small number of close relationships, a larger number of good relationships and a much larger number of acquaintances is well supported. Teams of a handful to a dozen people can rely on close relationships. Departments of dozens can rely on good relationships supported by shared rituals. Organisations of hundreds and beyond need markers that work among acquaintances and strangers.

Third, invest in markers before you need them. Companies that wait until the culture is slipping to define their markers find it harder to do so, because by then rival markers have formed in different parts of the organisation. Leaders who notice that their company is approaching the transition zone can begin to make implicit markers explicit, to create rituals that bring people together across teams and to build onboarding practices that teach newcomers the markers quickly.

Fourth, value the connectors. In every organisation there are people who maintain unusually large networks across teams and levels. They act as bridges between individual-recognition clusters and help hold the anonymous society together. Their role is rarely recognised in formal structures, and they are often among the first to burn out or leave when an organisation grows rapidly.

Fifth, be wary of rules that try to substitute for markers. When personal trust fades, many organisations respond by adding approvals, procedures and controls. Some of these are necessary. But rules that replace trust entirely can produce a culture of compliance in which people follow procedures without understanding or caring about their purpose. Markers that carry genuine information about shared values and competence can reduce the need for such rules.

A worked example

Consider a software company that has grown from forty people to two hundred in eighteen months. The founders still try to interview every candidate, but the queue is long and hiring has slowed. Weekly all-hands meetings, once a place for open debate, have become presentations, because too many people attend for discussion to work. New employees say they do not know whom to ask about anything. Long-serving employees say the company has lost its soul.

A reading based on Dunbar's number would note that the company has passed 150 and conclude that it should split into smaller units. That may be part of the answer, but the marker framework suggests a richer diagnosis. The company has outgrown the reach of personal recognition without

building the markers that would let strangers within it cooperate. Its procedures are unwritten, so newcomers cannot learn them without a personal guide. Its vocabulary is informal and varies by team. Its values have never been stated, because in the early days everyone simply knew them. Its status system is invisible to anyone who was not present when it formed.

The practical response follows from the diagnosis. The founders can stop trying to know everyone and start making the company legible to people who do not know each other. They can write down how decisions are made and who owns what. They can create a small number of rituals that bring people across teams together regularly. They can give each newcomer a guide for their first months. They can identify the connectors who already bridge teams and give them time and recognition for that work. None of this depends on whether the true limit on personal relationships is fifty, 150 or three hundred. It depends only on recognising that the limit has been passed.

The limit as an opportunity

It is tempting to see the limit on personal relationships as a constraint to be overcome, perhaps by better technology. Social media platforms promised to let people maintain far larger networks than before. In a narrow sense they succeeded. People now have hundreds or thousands of online connections. But the evidence suggests that the number of people they rely on for support, advice and intimacy has not grown in the same way. Technology can extend acquaintance. It has not obviously extended friendship.

For leaders, the limit is better seen as an opportunity. Because personal relationships are scarce, the markers that allow strangers to cooperate are valuable. An organisation that builds strong, accurate markers can operate at a scale that would otherwise be impossible. It can let people who have never met trust each other enough to collaborate. It can absorb newcomers quickly and coordinate across distance and time. In this sense, the limit on personal relationships is what makes marker design a strategic capability.

The next chapter turns to the oldest evidence of that capability: the hunter-gatherer peoples who, long before cities or states, saw themselves as members of societies far larger than the bands they lived in, and who gave those societies names that meant something like "real people."

CHAPTER 6

Real People

The popular image of hunter-gatherers is of small, isolated bands: a few families moving through a forest or across a desert, self-sufficient and largely alone. The image is partly true. For most of human history, the people you lived with day to day numbered in the dozens. They were your relatives, your in-laws and a few others who had joined by choice or necessity. You knew each of them intimately.

What the popular image misses is that these bands were rarely the whole of anyone's social world. Mark Moffett argues, drawing on a long tradition of anthropological fieldwork, that hunter-gatherer bands usually belonged to larger societies, numbering from a few hundred to a few thousand people, that shared a language, a set of customs and a name. Members of different bands might meet only occasionally, at seasonal gatherings, through marriage or when one band moved through another's range. They might never meet at all. Yet they recognised one another as belonging to the same people, and they distinguished themselves from neighbouring peoples whose language and customs differed.

This matters enormously for the argument of this book. It suggests that anonymous societies are not a recent invention of cities and states. They go back to the deep human past. Long before agriculture, writing or money, human beings were already living in societies too large for everyone to know everyone, held together by markers.

The names peoples give themselves

One of the most revealing clues to how these societies understood themselves lies in their names. Moffett notes that many peoples across the world have called themselves by words meaning, roughly, "people" or "real people." The Jahai, a group of forest-dwelling hunter-gatherers in the Malay Peninsula, use a word, *menra*, that Moffett cites in this sense. The pattern is widespread. The Inuit of the Arctic take their name from a word meaning "people." The Navajo call themselves Diné, also usually translated as "the people." The Yup'ik of Alaska and Siberia have a name that combines a word for "person" with a suffix meaning "real" or "genuine."

Not every ethnonym works this way, and many group names were given by outsiders rather than chosen by the people themselves. Linguists caution against reading too much into translations, which can flatten subtle meanings. Still, the recurrence of names meaning "people" or "real people" is striking. It suggests that, for many human societies, the line between members and outsiders was also a line between full humans and something less.

Moffett does not argue that hunter-gatherers regarded outsiders as animals, and the ethnographic record includes many examples of trade, intermarriage and hospitality across boundaries. The point is subtler. Membership in a people carried a sense of being normal, proper and complete. Outsiders might be traded with, married or fought, but they were different in a way that mattered. They did not quite do things right.

Markers before the state

How did hunter-gatherers who lived in small bands spread across large territories recognise one another as members of the same people? The answer, Moffett argues, is markers. Language and dialect were among the most important. Peoples often spoke distinct languages or dialects, and the ability to speak correctly was a primary sign of membership. Clothing, body decoration, hairstyles, tools and art styles also varied among peoples,

sometimes in ways that archaeologists can trace in the material record. Ritual practices, myths and songs were shared within peoples and differed between them.

Some hunter-gatherer societies developed especially elaborate systems for placing strangers. Many Aboriginal Australian societies use what are often called skin names, a system of social categories that assigns every person to a class with defined relationships to every other class. A stranger arriving at a camp could be given or asked for a skin name, and from that name everyone could work out how to behave toward them: whom they could marry, whom they should avoid, whom they should treat as a sibling or an uncle. The system turned an unknown individual into a known category within minutes. It was, in effect, a technology for making strangers legible.

Among the Ju/'hoansi of the Kalahari, the anthropologist Polly Wiessner documented a system of gift exchange called *hxaro*, in which people gave and received objects such as ostrich eggshell bead ornaments with partners in other bands, sometimes far away. Hxaro relationships created obligations of mutual support. When drought or conflict struck one area, people could travel to their hxaro partners and expect hospitality. The beads were not money in the modern sense. They were tokens of relationship, markers that a bond existed and would be honoured.

These examples show that the problem of trust among strangers was being solved creatively long before modern institutions. Hunter-gatherers did not have passports or credit scores, but they had languages, names, gifts and rituals that did similar work. They allowed people who did not know each other personally to cooperate, share resources and avoid conflict, at least within the boundaries of the people.

Imagined communities

The political scientist Benedict Anderson, in his 1983 book *Imagined Communities*, offered a famous definition of the nation. A nation, he argued, is an imagined political community. It is imagined because the members of even the smallest nation will never know most of their fellow members,

never meet them or even hear of them, yet in the minds of each lives the image of their communion. Anderson was writing about modern nations, and he emphasised the role of print capitalism, especially newspapers and novels in vernacular languages, in creating a sense of shared identity among readers who would never meet.

Anderson and Moffett come from very different disciplines, but their accounts converge. Both describe societies held together by shared images and symbols rather than personal relationships. Anderson's newspapers and Moffett's markers perform the same function. They let strangers feel that they belong to the same community.

The convergence also highlights a difference. The imagined communities of modern nations are vastly larger than hunter-gatherer peoples, and they rely on markers that are produced and distributed by institutions. A hunter-gatherer learned the markers of her people from her parents and neighbours. A modern citizen learns many of the markers of her nation from schools, media, governments and companies. The scale has changed, and so has the means of production.

Making the people

Modern nations did not arise automatically from pre-existing peoples. In many cases, they had to be made, and the making involved deliberate efforts to create shared markers. A well-known line, usually attributed to the Italian statesman Massimo d'Azeglio after unification in the nineteenth century, holds that having made Italy, the task was now to make Italians. Whether or not d'Azeglio said exactly this, the sentiment captures a real problem. The new Italian state contained people who spoke a variety of regional languages and dialects, identified more with their towns and regions than with the nation, and had little sense of shared identity.

France offers a similar example. During the revolutionary period, the Abbé Henri Grégoire conducted a survey of languages spoken in France and concluded that standard French was the everyday language of only a minority of the population, with many people speaking regional languages or dialects. Over the following century, the French state used schools,

military service and administration to spread standard French and a shared national identity. The historian Eugen Weber described this process in *Peasants into Frenchmen*, published in 1976, arguing that many rural French people came to think of themselves primarily as French only in the late nineteenth and early twentieth centuries.

Nation building involved flags, anthems, national holidays, standardised languages, national histories taught in schools and, eventually, documents that identified individuals as citizens. The French philosopher Ernest Renan, in a famous 1882 lecture, described the nation as a daily plebiscite, a community that exists because its members continually affirm their desire to live together. Markers are the means by which that affirmation is expressed and recognised.

The document that says who you are

Among the most consequential markers of the modern era is the passport. Documents allowing travel existed for centuries, but the modern passport, a standardised booklet issued by a state to identify its citizens, is largely a twentieth-century invention. Before the First World War, it was possible to travel across much of Europe without one. The war brought controls, and in 1920 a conference convened under the League of Nations in Paris recommended a standard format for passports. Over the following decades, the passport became the default marker of national membership for anyone crossing a border.

The passport is a fascinating object from the perspective of this book. It is a marker designed to be hard to forge, issued by an authority that guarantees its validity and recognised by other authorities around the world. It turns a stranger at a border into a legible category, a citizen of a particular state, with associated rights and obligations. It allows a border official to treat a person they have never met according to rules that apply to that category. In this sense, it is a modern descendant of the skin name and the hxaro bead, scaled up to a world of billions.

The passport also shows how markers can create inequality. The rights a passport confers vary enormously depending on who issued it. Holders of some passports can travel to most countries without a visa. Holders of others face extensive restrictions. The marker sorts people into categories with very different opportunities, often on the basis of where they happened to be born. Any discussion of designing markers needs to keep this sorting function in view.

Companies as peoples

Organisations are not nations, but they borrow many of the techniques of nation building. Large companies give their members names. Disney famously calls its theme park employees cast members. Starbucks calls its employees partners. Employees of large technology companies often adopt informal collective names derived from the company, and those names circulate inside and outside the firm. These names do something similar to the self-names of peoples. They tell members that they belong to a distinct group with its own character.

Companies also create foundational myths, often about the founders and early days. They establish rituals, from annual conferences to onboarding ceremonies. They develop vocabularies, sometimes so distinctive that outsiders cannot follow internal conversations. They produce symbols, from logos to swag, that members display. Some even create a version of the passport: the employee badge that grants access to buildings and systems and signals to everyone inside that the bearer belongs.

There is also a version of the "real people" distinction. In many organisations, employees distinguish between insiders who truly understand the company and outsiders who do not. Consultants, contractors, recent hires and employees of acquired companies may be treated as less than full members, however senior or capable they are. The distinction is rarely stated explicitly, but it shows up in who is invited to meetings, whose opinions are taken seriously and who is trusted with sensitive information.

The expanding circle

The history of human societies can be read, in part, as a history of expanding the circle of people who count as real. Hunter-gatherer peoples of a few thousand gave way, over thousands of years, to chiefdoms, states, empires and nations of millions. Each expansion required new markers capable of uniting people who had previously regarded each other as foreign. Religions that welcomed converts across ethnic lines, empires that granted citizenship to conquered peoples and nations that absorbed immigrants all extended the boundary of us.

The philosopher Peter Singer, in *The Expanding Circle*, published in 1981, argued that moral concern has expanded over history from kin to tribe to nation and potentially to all humans and beyond. The expansion has been neither smooth nor complete, and it has often been reversed. But the capacity to expand the circle is one of the most important features of human societies, and markers are central to it. New markers can create new forms of membership that include people previously excluded.

For leaders of organisations and platforms, this is perhaps the most important lesson of the hunter-gatherer record. The boundary of a group is not fixed by nature. It is defined by markers, and markers can be designed. A company can define its people narrowly, as those who share the founders' background, education and style, or broadly, as anyone who shares its purpose and meets its standards. A platform can make its markers accessible to newcomers from many backgrounds or set them up in ways that favour those who already belong. These choices determine who counts as real within the organisation, and they shape the organisation's capacity to grow, adapt and survive.

Founding myths

Peoples are held together not only by present markers but by shared stories about where they came from. Nations teach origin stories in schools. Religions recount their founding. Hunter-gatherer peoples passed on myths

that explained their place in the world. These stories are markers of a special kind. Knowing them, and telling them the right way, signals membership.

Companies tell origin stories too, and some have become part of the wider culture. The garage in Palo Alto where William Hewlett and David Packard began their company in 1939 was designated a California Historical Landmark in 1989 and described as the birthplace of Silicon Valley. Start-ups tell stories about the dorm room, the kitchen table or the moment of frustration that led to the idea. These stories do real work. They tell newcomers what kind of people the founders were, what the company values and how members are expected to behave. A story about founders who slept in the office signals a culture of intense commitment. A story about a founder who refused a lucrative deal on principle signals a culture that values integrity over money.

Origin stories can also become constraints. A company whose identity is bound up with a scrappy founding story may struggle to accept the processes and structures it needs as it grows. Members may treat new procedures as betrayals of the founding spirit. And origin stories can exclude, by implying that the real members are those who share the founders' backgrounds or were present at the beginning. Leaders who inherit a founding myth should ask what it teaches newcomers today, and whether the lessons still fit the organisation the company has become.

Before turning to systems

Part I has argued that anonymous societies are rare in nature, that ants and humans achieved them through markers, that humans read markers from infancy and form groups around them readily, that personal relationships have limits whatever their precise number, and that human beings have been living in societies held together by markers since long before the state.

Part II turns from biology and anthropology to systems and economics. If markers are the signals that let strangers trust one another, how does that trust behave over time? How do markers accumulate value, lose it, evolve and compete? Donella Meadows and Eric Beinhocker offer tools for

answering those questions. They let us see markers not as static symbols but as parts of dynamic systems, with stocks and flows, feedback loops and delays, and an evolutionary history that is still unfolding.

PART II

The System of Trust

Markers do not sit still. They gain and lose meaning, they feed back on the behaviour they signal, and they evolve as societies and economies change. Part II uses the tools of systems thinking and complexity economics to follow them through time.

CHAPTER 7

Trust as a Stock

Donella Meadows liked to begin with a bathtub. Imagine a tub with a tap and a drain. The water in the tub is a stock, an accumulation that you could, in principle, measure at any moment. The water flowing in through the tap and out through the drain are flows, rates of change that raise or lower the stock over time. If inflow exceeds outflow, the level rises. If outflow exceeds inflow, it falls. If the two are equal, the level stays the same, even though water is constantly moving through the system.

This simple picture, which Meadows used throughout *Thinking in Systems*, turns out to be one of the most useful tools in all of systems thinking. Almost everything we care about in organisations and societies can be described as a stock: money in a bank account, inventory in a warehouse, skilled people in a workforce, reputation in a market, carbon in the atmosphere. Each stock is filled and drained by flows. Understanding a system means understanding which stocks matter, what flows change them and what governs those flows.

Meadows emphasised a property of stocks that people consistently underestimate. Stocks change only as fast as their flows allow. A full bathtub with a small drain takes a long time to empty, no matter how urgently you want it empty. A large reservoir cannot be refilled by a single storm. Stocks act as buffers and memories. They carry the history of past flows into the present, and they make systems resistant to rapid change.

Trust among strangers behaves like a stock. It is accumulated over time through experiences, and it is drained by betrayals. It carries the history of past interactions into present decisions. And it changes more slowly than people who want to build it, or who have just damaged it, would like.

The stock you cannot see

A bathtub has a visible level. Trust does not. When you meet a stranger, you cannot directly observe how much trust the relevant society has accumulated in people like them. You cannot see a gauge on the stranger's forehead showing how reliable they are likely to be. You need some other way of estimating the stock.

Personal history is one way. If you have dealt with someone many times and they have behaved well, your personal stock of trust in them is high. But personal history is exactly what is missing in anonymous societies. Most of the people you deal with in a city, a corporation or a market are people you have never met.

This is where markers come in. A marker is a signal that lets you infer the trust stock without measuring it directly. When a stranger carries a marker associated with a group whose members have generally behaved well toward you and people like you, you extend trust on the basis of the group's record. The marker connects the individual in front of you to a stock of trust built up through countless past interactions with other members of the group.

In Meadows's terms, markers are information flows. They carry information about the state of a stock to the people making decisions. Meadows placed the structure of information flows high on her list of leverage points, the places where interventions in a system have the most effect. She told a story about houses in the Netherlands where electricity meters had been installed in different places. In some houses, the meter was in the basement. In others, it was in the front hall, where residents saw it every day. According to the account she recounted, households with visible meters used significantly less electricity, with no change in prices or rules. The only difference was that information about consumption was flowing to the people who could act on it.

Markers work similarly. They make an invisible stock visible enough to guide action. A stranger's uniform, badge, accent or rating tells you something about how much trust you can reasonably extend. If the

information is accurate, you make good decisions. If it is inaccurate, you misjudge the stock, extending trust where it is not warranted or withholding it where it is.

Inflows

What fills the stock of trust? In an individual relationship, the inflow is the accumulation of positive experiences: promises kept, help given, honest dealing. In an anonymous society, the inflow is the aggregate of millions of such experiences across many relationships, each of which slightly raises or lowers the general expectation of how members of a group will behave.

A striking piece of evidence about the stock of trust in strangers comes from a large field experiment published in *Science* in 2019 by Alain Cohn, Michel André Maréchal, David Tannenbaum and Christian Lukas Zünd. Research assistants turned in more than seventeen thousand "lost" wallets at banks, hotels, post offices and other institutions in 355 cities across forty countries. Each wallet contained a business card with an email address, a shopping list and a key, and some also contained money. The researchers measured how often recipients contacted the supposed owner. Return rates varied substantially between countries. In almost every country, though, wallets with money were more likely to be reported than wallets without, and wallets with more money were more likely still. Contrary to what most economists and members of the public predicted when surveyed, the presence of money increased honesty rather than decreasing it.

The finding suggests that the stock of trust among strangers in many societies is higher than people assume. It also suggests that the stock varies by place, reflecting different histories of institutions, norms and experiences. In Meadows's terms, different societies have filled their tubs to different levels, through different histories of inflow.

Markers contribute to inflow by making trust visible and therefore actionable. When people can identify members of their society reliably, they can extend trust to them and, if that trust is honoured, reinforce it. Each successful interaction adds to the stock. This creates a reinforcing loop, one of the two basic feedback structures Meadows described. Trust

enables cooperation, cooperation produces positive experiences, positive experiences raise trust, and higher trust enables more cooperation. The loop can drive a society toward high levels of trust over time, provided the markers remain reliable.

Outflows

The stock of trust is drained by betrayals, failures and exposures. Each time a member of a group behaves badly toward someone who trusted them because of their marker, the stock declines a little, both for that individual and, to some extent, for the group. When betrayals are frequent, visible or dramatic, the outflow can be large.

There is a well-known asymmetry here. Trust tends to be built slowly and destroyed quickly. A single betrayal can wipe out the effect of many positive experiences. Psychologists studying negativity bias have found that bad events often have stronger effects on perception than good events of equal magnitude, a pattern reviewed by Roy Baumeister and colleagues in a widely cited 2001 paper titled "Bad Is Stronger Than Good." In systems terms, the outflow from a betrayal is larger than the inflow from an equivalent act of good faith.

The asymmetry creates a characteristic pattern. A group's stock of trust may rise slowly for years and then fall sharply after a scandal. Rebuilding takes years again, if it happens at all. Leaders who have presided over such a fall often express frustration at how long recovery takes. They have fixed the problem, apologised and changed the policies. Why does the public, or the workforce, still not trust them? The answer is in the bathtub. The stock was drained quickly, and refilling it depends on inflows that accumulate only as fast as new positive experiences can be had.

The Tylenol lesson and its limits

The case often cited as a model of trust recovery is the response of Johnson and Johnson to the Tylenol poisonings of 1982. Seven people in the Chicago area died after taking capsules of Extra-Strength Tylenol that had

been laced with cyanide by an unknown person. The company recalled roughly thirty-one million bottles nationwide, at great expense, and later reintroduced the product in tamper-resistant packaging. Tylenol's market share collapsed in the immediate aftermath and then recovered substantially over the following year or so. The response became a staple of business school courses on crisis management.

Seen through the bathtub, the case is instructive in a specific way. The outflow was sudden and enormous, and it was not caused by the company's own wrongdoing. The company's response did not try to argue the stock back up. It acted on the flows. It stopped further outflow by removing the product, and it created a new marker, the tamper-resistant seal, that let customers verify safety rather than having to take it on trust. The seal was a marker that carried real information. It made the invisible state of the product visible at the moment of purchase.

The limits of the case are equally instructive. Johnson and Johnson had built a large stock of trust over many decades before the crisis, and it was widely understood that the company was a victim rather than a perpetrator. Organisations whose trust has been drained by their own actions face a harder problem. They cannot simply recall a product and add a seal. They need to change the behaviour that caused the outflow and then wait for inflows to accumulate, often under intense scrutiny.

Delays

Meadows paid close attention to delays, the time it takes for a change in one part of a system to produce effects in another. Delays are everywhere in systems involving trust. When an organisation changes its behaviour, it takes time for people to notice. When they notice, it takes time for their perceptions to update. When perceptions update, it takes time for new markers to form and spread. Each delay slows the flow of trust into or out of the stock.

Delays also cause trouble in the other direction. When an organisation begins to behave badly, its markers may continue to signal trustworthiness for some time before anyone notices. The stock of trust continues to

support cooperation even as the behaviour that justified it erodes. During this period, people who rely on the markers are exposed to risks they cannot see. When the truth comes out, the outflow is sudden, and people feel not only that the organisation behaved badly but that they were deceived by its markers.

Delays are one reason why markers can be dangerous. They carry information about the past, not the present. A brand that was trustworthy ten years ago may still be trusted today, even if the company behind it has changed. A credential earned decades ago may still be accepted, even if its holder's skills have decayed. The longer the delay between change in behaviour and change in markers, the larger the gap between what the markers say and what is true.

When the marker becomes the target

A further complication arises when people start treating markers as goals rather than signals. Meadows warned about systems that seek the wrong goal, pursuing an indicator that is easier to measure than the thing that actually matters. The economist Charles Goodhart gave his name to a related principle, often paraphrased as the observation that when a measure becomes a target, it ceases to be a good measure.

Markers are especially vulnerable to this dynamic. Because they signal trust, they confer advantages on those who carry them. That gives people an incentive to acquire the marker without necessarily acquiring the qualities it is supposed to signal. A company can invest in brand advertising rather than product quality. A job candidate can acquire a credential rather than a skill. A seller on a platform can game a rating system rather than improve service. In each case, the marker continues to function for a while, but its connection to the underlying stock weakens.

In systems terms, the information flow becomes corrupted. People reading the marker think they are learning about the stock of trust, but they are learning about the effort that has gone into producing the marker. As

this happens more often, the reliability of the marker declines, and people begin to discount it. Eventually, the marker may lose its value entirely, and the society has to find another way of estimating the stock.

Practical consequences

Treating trust as a stock and markers as information flows has several practical consequences for leaders.

First, it shifts attention from events to flows. A single scandal or success matters, but what matters more is the ongoing rate of positive and negative experiences that fill and drain the stock. Leaders who want to build trust need to manage the flows continuously, not merely respond to crises.

Second, it highlights the importance of accurate markers. The value of a marker lies in how well it reflects the underlying stock. Leaders who invest in markers that do not reflect real behaviour are borrowing against the future. When the gap is discovered, the outflow will be larger than if the markers had been honest from the start.

Third, it explains why trust takes so long to rebuild. Stocks change only as fast as their flows allow, and the flows that rebuild trust depend on experiences that accumulate one at a time. There is no shortcut. A leader who promises to restore trust quickly is promising something the system does not permit.

Fourth, it suggests where to look for leverage. Meadows argued that information flows are often more powerful leverage points than parameters such as prices or budgets. Making the right information visible to the right people at the right time can change behaviour without any change in incentives. Markers that make trustworthiness visible, and that make untrustworthiness visible quickly, are among the most powerful tools an organisation has.

The next chapter looks more closely at how marker systems go wrong over time, using some of the recurring patterns Meadows called system traps. Those patterns explain why so many marker systems that start well end badly, and what can be done to keep them honest.

CHAPTER 8

Traps in the Marker System

In the final part of *Thinking in Systems*, Donella Meadows described a set of recurring patterns she called system traps. Each trap is a structure, a particular arrangement of stocks, flows and feedback loops, that tends to produce a characteristic kind of trouble. Meadows gave them memorable names: policy resistance, the tragedy of the commons, drift to low performance, escalation, success to the successful, shifting the burden to the intervenor, rule beating and seeking the wrong goal. She argued that once you learn to recognise the structures, you can see the same problems recurring in very different settings, and you can look for the same kinds of escape.

Marker systems fall into these traps with remarkable regularity. Ratings inflate until they mean nothing. Credentials multiply until they exclude people who could do the job. Labels that once signalled quality are copied by anyone who wants the benefit without the cost. Leading platforms and institutions accumulate markers of trust that make it hard for newcomers to compete, whatever their quality. This chapter walks through several traps as they apply to markers. The aim is not to catalogue every failure but to give leaders a vocabulary for recognising structural problems before they become crises.

Drift to low performance

Meadows described a trap she called drift to low performance, sometimes also called eroding goals. A system has a goal and a perceived state. When the perceived state falls short of the goal, people take action to close the

gap. But if the goal itself is allowed to adjust toward the perceived state, the gap can be closed by lowering expectations instead of improving performance. Over time, standards drift downward, and each step seems reasonable because it is measured against the recent past rather than the original goal.

Rating systems on digital platforms show a mirror image of this trap. The drift is in the meaning of the marker. When a platform introduces a five-star rating system, the scale is supposed to discriminate between good and bad experiences. In practice, users quickly learn that low ratings can harm the person rated, that the person rated may retaliate or complain, and that most interactions are acceptable. Ratings cluster at the top. A study by Georgios Zervas, Davide Proserpio and John Byers of Airbnb listings found that the great majority of properties had average ratings of 4.5 stars or higher. On ride-hailing platforms, drivers and passengers widely understand that anything less than five stars signals a problem.

As ratings bunch at the top, the effective standard shifts. A rating that once meant "good" now means "a problem occurred." The marker still functions, but its range has collapsed, and it carries less information than it was designed to. Users must learn to read it in a new way, paying attention to tiny differences or to the text of reviews rather than the stars. Newcomers who take the scale at face value misinterpret it. And the platform loses one of its main tools for distinguishing excellent from adequate service.

Grades in education have followed a similar path in many institutions. In 2013 Harvard officials reported that the median grade awarded to undergraduates was an A-minus and that the most frequently awarded grade was an A. Similar patterns have been documented at many other universities. The drift happens gradually, through many individual decisions that each make sense locally. A teacher who grades more harshly than colleagues penalises their students. A department that grades more strictly than others loses enrolments. The marker loses discrimination, and employers and graduate schools turn to other markers, such as test scores, internships or the name of the institution, to sort candidates.

Escalation

Escalation is a trap in which two or more parties each try to stay ahead of the others, producing a reinforcing spiral. Arms races are the classic example. In marker systems, escalation often takes the form of credential inflation. When many applicants for a job hold a particular credential, it stops distinguishing them. Employers start asking for a higher credential. Applicants respond by acquiring it. The cycle continues.

Over recent decades, observers in many countries have documented jobs that once required a secondary education coming to require a degree, and jobs that once required a degree coming to require a postgraduate qualification, without clear evidence that the work itself had changed. A 2017 report titled *Dismissed by Degrees*, published by Harvard Business School with Accenture and Grads of Life and drawing on job-posting data from the labour market analytics firm Burning Glass Technologies, argued that many middle-skill job postings in the United States demanded degrees that most people already doing those jobs did not hold. Since then, a number of large employers and some state governments have announced that they would drop degree requirements for many roles and focus on demonstrated skills instead.

Escalation is costly for everyone involved. Applicants spend time and money acquiring credentials that do not make them better at the job. Employers lose access to capable people who lack the credentials. Educational institutions benefit in the short run but may find their markers devalued over time. And the credential stops carrying information about competence and starts carrying information about the ability to pay for and complete the credential, which is related to but not the same as competence.

Success to the successful

Meadows described a trap she called success to the successful. When winners gain resources that help them win again, a reinforcing loop concentrates advantage. The sociologist Robert Merton gave a related phenomenon in science the name the Matthew effect, after a verse in the Gospel of Matthew, observing that eminent scientists tend to receive disproportionate credit for work, which further increases their eminence.

Marker systems are particularly prone to this trap because markers of trust attract more opportunities to earn trust. A seller with many positive reviews appears higher in search results, receives more orders and accumulates more reviews. A new seller with no reviews struggles to get the first order. The eBay experiment by Paul Resnick and colleagues described in Chapter 22 found that buyers were willing to pay more than eight per cent more to an established seller than to new sellers offering matched goods, which illustrates the cost of starting without a marker.

The same dynamic appears in employment, where workers with prestigious names on their résumés receive more interviews and therefore more chances to add further prestigious names; in finance, where companies with strong credit ratings borrow more cheaply and therefore find it easier to maintain their ratings; and in culture, where well-known brands receive more attention and therefore remain well known. In each case, the marker system rewards those who already have markers, making it hard for newcomers to break in regardless of their quality.

Meadows noted that systems caught in this trap often need deliberate mechanisms to level the field, such as antitrust laws, progressive taxes or rules that give newcomers a chance. Platforms that want to remain open to new participants often build similar mechanisms, such as boosting new listings for a period, providing introductory badges or guarantees, or creating separate categories in which newcomers can compete with one another.

The tragedy of the commons

The tragedy of the commons occurs when a shared resource is open to everyone and each user has an incentive to take more than their share. The resource is depleted, and everyone suffers. Shared markers can behave like commons. When a label such as "natural," "eco-friendly" or "artisan" becomes associated with trust, every seller has an incentive to use it. If anyone can use it without meeting any standard, it is overused, and its value is depleted.

Consider the history of claims about environmental performance. When consumers and investors began to reward companies seen as environmentally responsible, environmental claims multiplied. Some were accurate. Many were vague or misleading, a practice critics labelled greenwashing. The most dramatic case was the Volkswagen emissions scandal that broke in 2015, when the company admitted that it had fitted millions of diesel cars with software that detected emissions tests and altered the engines' behaviour to pass them. Volkswagen had marketed its diesel cars in the United States under the banner of clean diesel. The marker had been used without the substance, and when the truth emerged, the damage extended beyond Volkswagen to the credibility of diesel claims in general.

Escapes from the tragedy of the commons usually involve some form of enclosure, regulation or mutual monitoring. For markers, this often means certification, in which a third party controls who may use a label and checks that standards are met. Fair trade labels, organic certifications and safety marks all work this way, with varying degrees of success. Certification turns an open commons into a governed resource. It has costs of its own, which later chapters discuss, but it addresses the basic problem that an unregulated marker invites free riders.

Shifting the burden

Meadows described a trap she called shifting the burden to the intervenor. A problem is addressed by an intervention that relieves the symptoms without solving the underlying cause. Over time, the system becomes dependent on the intervention and loses its capacity to solve the problem itself.

Markers can create this trap when organisations rely on them instead of developing their own judgment. A hiring manager who screens only on credentials may never develop the ability to assess candidates directly. A buyer who relies only on brand names may lose the ability to judge quality. A platform that relies only on automated ratings may stop investing in direct quality checks. In each case, the marker relieves the burden of judgment, and the capacity for judgment atrophies.

The trap becomes visible when markers fail. If a credential turns out to be fraudulent, a brand turns out to be hollow or a rating system turns out to be gamed, the organisation discovers that it has no other way of assessing trust. It must rebuild the capacity for judgment from scratch, often in a crisis.

The escape from this trap is not to abandon markers, which are indispensable in anonymous societies. It is to maintain a parallel capacity for direct assessment. Organisations that check a sample of credentials, test candidates' skills directly, audit suppliers in person or periodically verify ratings with independent data keep their judgment alive, and they are better placed when markers fail.

Rule beating

Rule beating occurs when people comply with the letter of a rule while evading its purpose. Markers are often defined by rules: a certification requires certain documents, a verification badge requires a certain process, a credential requires passing a certain exam. People who want the marker without the substance learn to satisfy the rules without meeting the intent.

Security certifications, for example, are supposed to signal that a company protects data well. But many certification regimes check whether policies and procedures exist rather than whether they are effective. A company can obtain a certification by producing the right documents and passing an audit, while still having serious weaknesses in practice. The marker tells customers that the rules were followed, not that the data is safe. When a breach occurs at a certified company, customers who relied on the marker feel betrayed, and the value of the certification declines.

Meadows suggested that the remedy for rule beating is to design rules around purposes rather than procedures, and to welcome rule beating as feedback about where rules have failed. For markers, this means continually checking whether the marker still predicts the behaviour it is supposed to signal and adjusting the rules when it does not.

Seeking the wrong goal

The final trap considered here is the one that underlies many of the others. Meadows warned that systems will faithfully pursue whatever goal is built into them, even if that goal is not what anyone actually wants. If a company measures success by the number of positive ratings, it will produce positive ratings, whether or not customers are better served. If a school measures success by grades, it will produce high grades.

Markers are particularly susceptible because they are often easier to measure than the qualities they represent. It is easier to count credentials than to assess competence, easier to track brand awareness than to measure product quality, easier to count verified accounts than to ensure honest behaviour. When the marker becomes the goal, the system optimises for the marker and the underlying quality drifts. The result is the familiar Goodhart pattern discussed in the previous chapter, now embedded in the structure of the organisation.

Recognising the traps

Each of these traps produces characteristic symptoms. Drift to low performance shows up as compression, with markers bunching together until they no longer discriminate. Escalation shows up as inflation, with markers ratcheting upward in cost and requirement. Success to the successful shows up as concentration, with markers accumulating among those who already hold them. The tragedy of the commons shows up as dilution, with markers used by more and more people who do not meet the standard. Shifting the burden shows up as dependence, with organisations unable to assess trust without the marker. Rule beating shows up as hollow compliance, with markers obtained through procedure rather than substance. Seeking the wrong goal shows up as optimisation for the marker at the expense of the thing it signals.

Leaders who learn to recognise these symptoms can intervene earlier. They can recalibrate rating scales, cap or redesign credential requirements, create on-ramps for newcomers, govern shared labels, maintain direct assessment and design rules around purpose. None of these interventions is easy, and each creates new problems. But they are far easier than rebuilding a marker system after it has collapsed.

The traps also remind us that marker systems are not static. They are dynamic structures that evolve under pressure from the people who use them. The next chapters use Eric Beinhocker's evolutionary economics to explore that evolution more directly, starting with one of the most fundamental and most culturally variable markers of all: wealth.

CHAPTER 9

Wealth Is What a Society Agrees to Count

Eric Beinhocker opens *The Origin of Wealth* with a comparison between two tribes. The first is the Yanomamö, a people of the Amazon rainforest on the border of Brazil and Venezuela, who at the time of the ethnographic accounts he drew on lived by hunting, gathering and gardening, in villages of a few dozen people. The second is the tribe of New Yorkers, who live along the Hudson River and spend their days in offices, shops and subway carriages. Biologically, Beinhocker notes, the two groups are essentially identical. Their economies are not. He estimates that the average Yanomamö income, if one could measure it in money, would be a tiny fraction of the average New Yorker's. More strikingly, he estimates the number of distinct products and services available in each economy. For the Yanomamö, the count of what retailers call stock keeping units might run to several hundred, perhaps a few thousand. For New Yorkers, his rough estimate is on the order of ten billion.

Beinhocker uses the comparison to ask the question at the centre of his book. Where does wealth come from, and why has it grown so dramatically in some places over the last few centuries? His answer, developed over several hundred pages, is that wealth is the product of an evolutionary process. Economies generate a constant stream of new designs for products, business plans and ways of organising people. Markets and other selection mechanisms test those designs, and the ones that succeed are copied and amplified. The extraordinary variety of a modern economy is the accumulated result of that process, running for centuries.

Along the way, Beinhocker makes a point that is easy to pass over but crucial for this book. What counts as wealth is not fixed by physics or biology. It is defined by culture. Different societies value different things, and the same object may be treasure in one place and rubbish in another. Wealth is, in an important sense, a marker. It is what a society agrees to count.

Cattle and bank balances

For many East African pastoralist peoples, including the Maasai of Kenya and Tanzania, cattle have long been the central measure of wealth. A family's herd is its savings, its insurance and its standing in the community. Cattle are exchanged at marriages, given as gifts to cement alliances and slaughtered at important ceremonies. A person with many cattle is wealthy in the fullest social sense, respected and influential. A person who has lost their herd to drought or raiding has lost more than food. They have lost status.

For a bank customer in a European city, wealth is usually measured in money: a balance in an account, a portfolio of investments, the value of a house. The customer may never touch most of this wealth. It exists as entries in databases, maintained by institutions the customer trusts to honour them.

Beinhocker's point is not that one of these is more real than the other. Both are cultural agreements. Cattle are valuable because a community has organised its marriages, ceremonies and status hierarchies around them. Bank balances are valuable because a society has organised its laws, institutions and expectations around them. Take a Maasai herd to central Frankfurt and it would be an expensive inconvenience. Take a Frankfurt bank statement to a traditional pastoralist community and it would be a piece of paper.

The same point is embedded in the history of language. The English word "pecuniary," meaning relating to money, derives from the Latin *pecunia*, which is related to *pecus*, meaning cattle or livestock. The word "fee" descends from an Old English word that could mean both livestock and

property. In many societies, wealth and livestock were once closely linked. As economies changed, the markers of wealth changed with them, but traces of the older markers survive in the words we use.

The stones of Yap

One of the most famous illustrations of wealth as a cultural agreement comes from the island of Yap in the western Pacific. For centuries, Yapese people used large stone discs, called rai, as a form of value. The stones were quarried on Palau, hundreds of kilometres away, and transported by canoe at great effort and risk. Some were enormous, too heavy to move once they had been set in place. They were used in important transactions such as marriages and political settlements.

The economist Milton Friedman wrote about Yap in a 1991 paper titled "The Island of Stone Money," drawing on an account published in 1910 by the American traveller and anthropologist William Henry Furness III. Furness described how ownership of a stone could change without the stone being moved. Everyone simply knew that the stone now belonged to someone else. He also described a stone that had been lost at sea during transport. Because the crew testified that it had been of great size and value, the stone continued to be counted as part of a family's wealth, even though no one had seen it for generations and it lay at the bottom of the ocean.

Friedman used the story to illustrate that money is a matter of shared belief. The value of the lost stone did not depend on its physical location. It depended on the community's agreement that it existed and belonged to a particular family. This is not as strange as it may sound. Much of the world's wealth today exists in a similar form, as entries in ledgers recording claims that are honoured because institutions and communities agree to honour them. Gold stored in a central bank vault may change ownership without ever leaving the vault, just as the Yapese stones changed hands without moving.

Cigarettes behind the wire

Cultural agreements about wealth can form quickly under pressure. In 1945 an economist named R. A. Radford, who had been a prisoner of war during the Second World War, published a paper in the journal *Economica* titled "The Economic Organisation of a P.O.W. Camp." Radford described how prisoners, who received parcels of food and supplies from the Red Cross, developed a market for trading goods. Cigarettes emerged as the main currency. They were widely desired, easily divisible, reasonably durable and familiar to everyone. Prices for food, clothing and services were quoted in cigarettes.

Radford's account is valuable because it shows a marker of wealth emerging from scratch, in a society of strangers thrown together by war. No authority decreed that cigarettes would be money. The convention arose because it was useful, and it persisted because everyone expected everyone else to accept it. Radford also described how the cigarette currency was affected by supply shocks. When parcels arrived, cigarettes became plentiful and prices rose. When supplies ran low, prices fell. Non-smokers held cigarettes purely as a medium of exchange, while smokers consumed them, which gradually reduced the money supply.

The cigarette economy illustrates several features of markers of wealth. They are chosen for practical qualities, such as portability and divisibility, but their value depends on shared expectations. They can be disrupted by changes in supply. And they can serve as markers of status as well as means of exchange. A prisoner with many cigarettes was, in the camp's terms, rich.

Wealth as display

Wealth is not only a store of value. It is also a marker of status, displayed to others as evidence of success, power or virtue. Many societies have institutions in which wealth is deliberately shown, given away or destroyed in order to earn prestige.

The potlatch ceremonies of Indigenous peoples on the Pacific Northwest coast of North America are a well-known example. Hosts distributed large quantities of goods to guests, and in some cases destroyed valuable objects, as a way of demonstrating status and validating claims to rank and title. Among the most prized objects were large engraved copper plaques, known as coppers, whose value was established by their history of exchange. Colonial authorities in Canada banned the potlatch for decades, from the late nineteenth century until the ban was dropped in the 1950s, partly because missionaries and officials regarded the giving away of wealth as wasteful and incompatible with the economic habits they wished to instil.

The ban is itself instructive. It shows how one society's markers of wealth and status can be seen as irrational or immoral by another. The potlatch was a sophisticated system for converting material goods into social standing and for redistributing resources across a community. To outside observers who measured wealth by accumulation, it looked like destruction.

Why cultural definitions of wealth matter for leaders

At first glance, cultural definitions of wealth may seem remote from the concerns of leaders in modern organisations. But they matter in at least three ways.

First, organisations have their own definitions of wealth, and those definitions function as markers. In some companies, status is measured by budget and headcount. In others, it is measured by technical reputation, by proximity to the chief executive or by the success of products. These internal currencies shape behaviour as powerfully as salaries do. People pursue what their organisation counts as wealth. A leader who wants to change behaviour needs to understand which currencies matter and how they are earned.

Second, mergers and collaborations bring together organisations with different definitions of wealth. An engineering-led company that values technical excellence may merge with a sales-led company that values

revenue and relationships. Each side may regard the other's markers of status as superficial or misguided. Engineers who have earned prestige by solving hard problems may be baffled when salespeople with large accounts are treated as more important. Salespeople may find engineers' obsession with elegance indulgent. These differences are not merely matters of opinion. They are differences in what each culture counts as wealth.

Third, platforms and markets that operate across cultures must recognise that participants may value different things. A reputation system that rewards one kind of behaviour may be meaningless or offensive to users from another culture. A loyalty programme built around one form of value may fail to engage people who value something else. Designers who assume that their own definitions of wealth are universal are likely to make mistakes.

The invisible agreement

Beinhocker's broader argument is that the economy is a complex adaptive system, a network of interacting agents who learn and adapt, and whose collective behaviour cannot be understood by studying any one of them alone. Definitions of wealth are part of the structure of that system. They determine what agents pursue, what they exchange and how they measure success. They are not chosen by anyone in particular. They emerge from countless interactions and are reinforced by the expectation that others will honour them.

This is what makes them so powerful and so fragile. As long as the agreement holds, wealth works smoothly. People accept cattle, stones, cigarettes or bank balances because they expect others to accept them too. When the agreement is shaken, for instance by a crisis that makes people doubt whether others will honour their claims, wealth can evaporate with alarming speed. The next chapter examines the most consequential marker of wealth in modern societies, money itself, and what happens when trust in it is built, attacked and lost.

Partnership as wealth

Professional services firms offer a clear example of how an organisation defines its own wealth. In many large law firms, the most important marker of success is not salary but partnership. The model often associated with the New York firm Cravath, Swaine and Moore, developed in the early twentieth century, hired young lawyers as associates, trained them intensively and, after a number of years, either promoted them to partner or expected them to leave. Variants of this up-or-out model spread across law, accounting and consulting.

In such firms, the path to partnership becomes the central currency. Associates measure their progress by the signals they receive about their prospects: the partners who request them, the clients they are trusted with, the assignments that indicate rising status. Billable hours often serve as a visible marker of commitment. Associates who bill more are seen as more dedicated, whatever the quality of their work, and the marker can drive extraordinary working hours.

The partnership system illustrates how an organisation's definition of wealth shapes behaviour. Because partnership is scarce and highly valued, people compete intensely for it. Because the markers of progress are visible, people pursue them. Because the up-or-out rule makes failure visible too, people who do not make partner often leave, carrying the firm's markers into the wider profession, where alumni networks become a further source of status. Leaders of such firms who want to change behaviour, for instance to reduce burnout or improve collaboration, must change what counts as progress toward the firm's central currency. Exhortations to work differently will have little effect while billable hours remain the marker by which commitment is judged.

A practical exercise

Before turning to money, leaders may find it useful to map the definitions of wealth in their own organisation. The exercise is simple to describe and often revealing in practice.

List the things that confer status in the organisation. Include formal markers, such as titles and pay grades, and informal ones, such as access to senior leaders, ownership of prestigious projects, technical reputation or a seat at certain meetings. For each, ask how it is earned, who controls access to it and what behaviour it rewards.

Then ask how these internal currencies relate to the organisation's stated goals. If the organisation says it values collaboration but confers status on individual heroics, the currency is working against the goal. If it says it values customers but rewards internal politics, people will pursue politics. Changing such patterns requires changing what the organisation counts as wealth, which is far harder than changing a value statement but far more effective.

Finally, ask how the organisation's definitions of wealth look to outsiders: new hires, customers, partners and employees of companies it might acquire. Definitions that seem natural from inside may look strange, unfair or incomprehensible from outside. Recognising this is the first step toward designing markers that can bridge different cultures rather than dividing them.

CHAPTER 10

Money, the Marker Everyone Carries

Pull a banknote from your wallet and look at it closely. It is a small rectangle of paper or polymer, worth almost nothing as material. Yet it is covered in symbols. There may be a portrait of a monarch, a founding figure or a celebrated scientist. There may be a coat of arms, a national landmark or an image of native wildlife. There will be the name of the issuing authority, a signature, a serial number and an array of security features: watermarks, metallic threads, microprinting, holograms or windows that change colour when tilted.

All of this exists to do one job. The note must persuade a stranger to accept it. When you hand it to a shopkeeper, the shopkeeper does not know you and has no reason to trust you personally. What they trust is the note, and behind the note the institutions that issue and defend it. Money is a marker that lets strangers exchange value without knowing one another. It is, in a sense, the most widely carried marker of membership in a modern society.

The design of a shared identity

The symbols on money are rarely accidental. They are chosen to express the identity of the community that uses the currency, and the choices can be revealing.

When the euro was introduced as physical cash in 2002, its designers faced an unusual problem. The currency would be used across many countries with different histories, languages and national symbols. Putting

any single nation's heroes or monuments on the notes risked offending the others. The solution, chosen through a design competition won by the Austrian designer Robert Kalina, was to depict architectural features, windows, gateways and bridges, representing different periods of European architectural history, without depicting any real, identifiable structure. The bridges on the notes were deliberately generic. They stood for Europe as a whole rather than any particular place.

The coins took a different approach. Each euro coin has a common side, shared across the eurozone, and a national side chosen by the issuing country. A German euro coin and an Italian euro coin are equally valid throughout the area, but each carries a reminder of where it was minted. The design acknowledges two layers of identity at once, the shared European currency and the national community that issued a particular coin.

There is a charming postscript. Some years after the notes entered circulation, a town in the Netherlands, Spijkenisse, built a set of real bridges modelled on the fictional ones from the banknotes. The generic symbols of European unity had become specific landmarks after all.

The euro's design illustrates the tension at the heart of marker design. Markers must be specific enough to signal membership in a particular community, yet inclusive enough that all members can identify with them. Too specific, and they exclude. Too generic, and they mean nothing. The euro's bridges tried to find a middle path, and their reception, sometimes criticised as bland, sometimes praised as inclusive, shows how hard that path is to walk.

Counterfeits as attacks on trust

Because money is a marker of trust, counterfeiting is an attack not only on the people who receive fake notes but on the system as a whole. If enough counterfeit notes circulate, people begin to doubt every note, and the marker's power to enable exchange among strangers is weakened.

States have long understood this, and some have used counterfeiting as a weapon. During the Second World War, Nazi Germany organised a scheme, later known as Operation Bernhard, to forge British banknotes. Jewish prisoners with skills in printing, engraving and related trades were forced to work on the project at the Sachsenhausen concentration camp. The forgeries produced were of high quality. The original aim was reportedly to destabilise the British economy by flooding it with fake notes, although in practice many of the notes were used to pay for intelligence operations and other purposes. The episode has been told in books and in the 2007 Austrian film *The Counterfeiters*.

Operation Bernhard shows how a marker's value depends on the cost of forging it. Banknotes are expensive to counterfeit because they incorporate features that are hard to reproduce. As reproduction technology improves, issuers respond by adding new features, in a continuing contest between the designers of the marker and those who would copy it. This contest is a version of the escalation trap described in Chapter 8, but here escalation is necessary. A marker that is cheap to forge cannot carry trust.

When the agreement fails

Money depends on a shared expectation that others will accept it. When that expectation breaks down, money can lose its value with startling speed.

Germany in 1923 offers the most famous example. After the First World War, the German government faced heavy reparations obligations, economic disruption and political turmoil. It financed spending by printing money, and inflation accelerated into hyperinflation. By late 1923, prices were rising so fast that workers were paid more than once a day and rushed to spend their wages before they lost value. Photographs from the period show people carrying banknotes in baskets and children playing with bundles of worthless currency.

The crisis was ended, in part, by the introduction of a new marker. In November 1923 the German authorities issued the Rentenmark, a new currency nominally backed by mortgages on land and industrial assets. It was exchanged for the old paper mark at a rate of one Rentenmark for one

trillion paper marks. The backing was more symbolic than practical, since few people could actually redeem a Rentenmark for land. But the new currency, combined with fiscal measures and strict limits on issuance, restored a shared expectation that money would hold its value. The marker worked because people believed it would and because the authorities behaved in ways that sustained that belief.

Zimbabwe experienced a similar collapse in the late 2000s, when inflation reached such extreme levels that the central bank issued a note with a face value of one hundred trillion Zimbabwe dollars. The currency was eventually abandoned in favour of foreign currencies, especially the US dollar. In both cases, money failed not because the paper changed but because the shared agreement that sustained it broke.

Runs on the bank

A related failure occurs when people lose confidence not in a currency but in an institution that holds their money. A bank run is a collapse of trust in the bank's promise to return deposits on demand.

In September 2007, customers queued outside branches of the British bank Northern Rock to withdraw their savings, after news that the bank had sought emergency support from the Bank of England. It was widely described as the first run on a British bank in well over a century. Pictures of the queues, broadcast around the world, became one of the defining images of the early financial crisis. The queues themselves fed the run, since people who saw others withdrawing money had reason to fear that the bank would run out.

In March 2023 Silicon Valley Bank, a lender closely tied to the technology start-up community in the United States, suffered a run of extraordinary speed. Depositors attempted to withdraw around forty-two billion dollars in a single day, according to regulatory filings, much of it initiated through online banking and coordinated through messaging among venture capitalists and founders. The bank was taken over by

regulators the following day. The episode showed how digital communication can accelerate the outflow of trust. Markers that once took days or weeks to collapse can now fail in hours.

These episodes illustrate the dynamics of trust as a stock described in Chapter 7. The stock of trust in a bank is built up over years of reliable behaviour and signalled through markers such as brand, regulation and deposit insurance. It can be drained very quickly when a trigger causes depositors to doubt the markers. Because each withdrawal makes further withdrawals more rational, the outflow becomes a reinforcing loop, a spiral that can destroy an institution in days.

Insurance as a meta-marker

One of the most important responses to bank runs has been the creation of deposit insurance. In the United States, the Federal Deposit Insurance Corporation was created in 1933, during the Great Depression, to guarantee deposits up to a limit. Many other countries have adopted similar schemes. Deposit insurance changes the system's structure. Depositors no longer need to judge the trustworthiness of each bank individually. They need only trust the insurance scheme, which is backed by the state.

In marker terms, deposit insurance is a meta-marker. It is a signal that the underlying markers, a bank's brand and balance sheet, are backed by a further guarantee. It shifts the burden of trust from many individual institutions to a single public one, and it weakens the reinforcing loop that drives runs. The trade-off, much discussed by economists, is that insurance can also weaken the incentive for depositors to monitor banks, and for banks to behave prudently, a problem known as moral hazard. Like most interventions in complex systems, it solves one problem and creates others.

Money without a sovereign

The emergence of cryptocurrencies in the years after the 2008 financial crisis represented an attempt to create money without relying on trust in banks or governments. The white paper that introduced Bitcoin, published

in 2008 under the pseudonym Satoshi Nakamoto, described a system in which transactions would be verified by a distributed network rather than a trusted third party. Supporters often described such systems as trustless.

In practice, cryptocurrencies did not eliminate trust. They relocated it. Users had to trust the code, the network of miners or validators, the exchanges where they bought and sold coins and the wallets where they stored them. The collapse of the cryptocurrency exchange FTX in November 2022, followed by the conviction of its founder Sam Bankman-Fried on fraud charges in 2023, showed that many users had in fact placed their trust in familiar markers: a charismatic founder, celebrity endorsements, sponsorship of sports teams and arenas, and the appearance of institutional legitimacy. When those markers proved hollow, the losses were enormous.

The lesson is not that cryptocurrencies are necessarily fraudulent. It is that trust among strangers always requires markers of some kind. Removing one set of markers, such as bank brands and government guarantees, does not eliminate the need for trust. It creates demand for new markers, and those new markers can be forged, inflated or misused like any others.

Money inside organisations

Most organisations do not issue currency, but many create internal markers that work like money. Budgets, headcount allocations, credits for computing resources and points in internal recognition systems are all forms of internal currency. They let people who do not know each other coordinate the allocation of resources.

These internal currencies are subject to the same dynamics as national money. They can be inflated, as when budget allocations grow faster than the value they produce. They can be counterfeited, as when people game recognition systems. And they can collapse, as when an organisation in crisis freezes budgets and people suddenly discover that their internal wealth cannot be spent.

Leaders who design internal currencies should ask the same questions central bankers ask. What gives the currency its value? Who issues it, and under what rules? How hard is it to fake? What happens if people lose confidence in it? The answers determine whether the currency will help strangers cooperate or simply create a new arena for politics.

Private currencies

Not all money is issued by states. Companies have long issued currencies of their own, and these private currencies show the same dynamics of trust as national money. Airline loyalty programmes are a familiar example. American Airlines launched its AAdvantage programme in 1981, and frequent flyer schemes spread rapidly across the industry. Members accumulate miles or points that can be exchanged for flights, upgrades and other rewards. The points are, in effect, a currency issued by the airline, whose value depends on the airline's willingness to honour them.

Loyalty currencies combine the functions of money and markers. They store value, but they also signal status. Many programmes award tiers, such as silver, gold or platinum, that grant visible privileges at airports: separate queues, lounge access, priority boarding. These privileges mark members as valued insiders in front of other passengers. The status is often as important to frequent travellers as the free flights.

Because the airline controls the currency, it can change its value. Programmes regularly adjust how many points are needed for a reward, a practice members often describe as devaluation. Each adjustment drains a little of the trust that members place in the currency, and sharp changes can provoke the same sense of betrayal that accompanies a national currency crisis, on a smaller scale. Members who have accumulated points over years feel that their savings have been eroded by an authority they cannot hold to account. Companies that issue private currencies take on, in miniature, the responsibilities of a central bank, and they face the same temptation to inflate.

The deepest marker

Money is perhaps the deepest marker of trust in modern societies, precisely because people rarely think about it. Every time someone accepts a banknote, a card payment or a bank transfer from a stranger, they are relying on an invisible web of institutions, laws, expectations and shared beliefs. The marker makes the web usable without making it visible.

The previous chapter argued that wealth is what a society agrees to count. Money is the means by which that agreement is made portable. It allows the stock of trust accumulated in a society's institutions to be carried in a pocket and exchanged with anyone who shares the agreement. When the agreement holds, money is the most powerful tool ever invented for cooperation among strangers. When it fails, the result is one of the most rapid collapses of social order that any society can experience.

The next chapter turns to the long history of how societies built the institutions that let trust extend beyond kin, from medieval trading coalitions to the modern rule of law.

CHAPTER 11

Kinship's Long Shadow

For most of human history, the safest person to trust was a relative. Kin share genes, which gives them a biological interest in one another's welfare, and they share histories, which gives them knowledge of one another's character. A cousin who cheats you will have to face the whole family at the next wedding. A brother who fails you will hear about it from your mother. Kinship is the oldest system of accountability, and in many parts of the world it remains the most important one.

The subtitle of this book refers to trust outgrowing kinship. That is not a claim that kinship disappears in large societies. Families remain central to most people's lives. It is a claim that the trust which large societies require cannot be supplied by kinship alone. When people need to cooperate with thousands or millions of strangers, they need markers and institutions that extend trust beyond the family. How societies built those markers and institutions, and why some built them more extensively than others, is one of the great questions of economic history. It also has direct implications for anyone building an organisation that must grow beyond its founders' personal networks.

The traders of the Geniza

In the medieval synagogue of Fustat, in what is now Old Cairo, a storeroom known as a geniza accumulated centuries of discarded documents, because Jewish custom discouraged destroying writing that might contain the name of God. When scholars examined the Cairo Geniza in the late nineteenth

and twentieth centuries, they found an extraordinary archive, including thousands of letters written by merchants who traded across the Mediterranean in the eleventh century.

The economic historian Avner Greif used these letters to study how traders, often called the Maghribi traders because many came from or traded with North Africa, solved the problem of trusting agents in distant ports. A merchant in Fustat might send goods to Sicily or Tunisia, where an agent would sell them and send back the proceeds. The merchant could not supervise the agent directly. Courts were slow and uncertain across jurisdictions. What stopped agents from cheating?

Greif argued, in influential work published from the late 1980s onward, that the traders formed a coalition in which information about agents' behaviour circulated widely through letters, and in which an agent who cheated one member would be shunned by all. This multilateral punishment made cheating costly, since an agent who betrayed one merchant would lose business with the entire network. Trust was sustained not by personal relationships between each merchant and each agent but by a shared reputation system within a community defined by religion, language and commercial practice. In Moffett's terms, the coalition was an anonymous society of traders held together by markers and gossip.

Greif contrasted the Maghribi traders with the merchants of Genoa, who, in his account, relied more on bilateral relationships and formal legal institutions such as contracts and courts. He argued that the Genoese system scaled better over the long run, because it did not depend on membership in a closed community and could therefore extend to strangers.

Greif's interpretation has been challenged. The historians Jeremy Edwards and Sheilagh Ogilvie argued in 2012 that the Geniza letters show the traders using formal legal mechanisms extensively and that the evidence for a coalition enforcing multilateral punishment is weaker than Greif suggested. The historian Jessica Goldberg's detailed study of the Geniza merchants offered a more complex picture of their business relationships. The debate is unresolved, and readers should treat the coalition story as one

interpretation among several. But the question it raises, how communities create trust among members who cannot supervise one another, is central to the argument of this book.

Fairs and private judges

A related story concerns the Champagne fairs of medieval France, where merchants from across Europe gathered to trade in the twelfth and thirteenth centuries. In a 1990 paper, the economists Paul Milgrom, Douglass North and Barry Weingast proposed that the fairs were supported by a system of private judges who kept records of disputes and could inform merchants about the past behaviour of potential trading partners. Merchants who refused to honour judgments would be excluded from future trade. The system, they argued, allowed strangers from different regions to trade with confidence, even without a powerful state to enforce contracts.

This account, too, has been questioned by historians, including Ogilvie, who argue that the fairs relied heavily on public authorities and that the private enforcement story is overstated. Again, the details matter less for this book than the general point. Expanding trade beyond kin and community required mechanisms that could make strangers' reputations legible and their promises enforceable. Whether those mechanisms were private, public or a mix, they functioned as markers and institutions of trust.

Amoral familism

The flip side of these stories is what happens when trust remains confined to the family. In 1958 the American political scientist Edward Banfield published *The Moral Basis of a Backward Society*, a study of a poor village in southern Italy that he called Montegrano. Banfield argued that villagers acted to maximise the short-term material advantage of their nuclear families and assumed that everyone else would do the same. He called this ethos amoral familism. In his account, it made collective action almost

impossible. Villagers would not cooperate to build a school or a hospital, because each assumed that others would cheat or that the benefits would go to someone else's family.

Banfield's work has been heavily criticised for overgeneralising from a single village and for neglecting the historical and economic conditions that made distrust rational. Later research has suggested more complex explanations for regional differences in trust and civic engagement. Yet the concept of amoral familism has endured because it names a real pattern. Where trust is limited to kin, cooperation among strangers is difficult, and the institutions that depend on such cooperation, from public services to large firms, struggle.

The political scientist Francis Fukuyama made a broader version of this argument in *Trust*, published in 1995. He contrasted societies in which trust extended widely beyond the family with those in which it remained largely within it, and he argued that the radius of trust affected economic organisation. In societies with narrow trust, he suggested, businesses tended to remain family firms, because owners were reluctant to hand control to outsiders. In societies with broader trust, large professionally managed corporations were easier to build. Fukuyama's categorisations of particular countries have been debated, but the idea of a radius of trust has become widely used.

The WEIRD hypothesis

The most ambitious recent attempt to explain why trust expanded beyond kinship in some societies is Joseph Henrich's *The WEIRDest People in the World*, published in 2020. Henrich, an anthropologist at Harvard, argues that people in Western, educated, industrialised, rich and democratic societies are psychologically unusual in several respects, including their tendency to trust strangers and to follow impartial rules. He traces this psychology in part to the influence of the medieval Western Church, which over centuries imposed rules against cousin marriage and other practices that sustained large kin groups. As kin networks weakened, he argues,

people were pushed to form voluntary associations such as guilds, monasteries, universities and towns, and they developed a psychology suited to cooperating with unrelated others.

Henrich supports his argument with a wide range of evidence, including cross-cultural psychological studies and historical data on marriage patterns. One illustration he draws on comes from a study by Raymond Fisman and Edward Miguel of parking tickets accumulated by United Nations diplomats in New York. Before 2002, diplomatic immunity meant that diplomats could ignore parking tickets without penalty. Fisman and Miguel found that diplomats from countries with higher levels of corruption accumulated far more unpaid tickets than those from countries with lower corruption, suggesting that norms carried over from home shaped behaviour even when enforcement was absent.

Henrich's thesis is bold and has attracted criticism from historians and anthropologists who question aspects of his causal story and his use of evidence. It should be read as a provocative hypothesis rather than an established conclusion. But it connects directly to the theme of this book. If trust in strangers is not a universal human default but a cultural achievement, shaped by particular institutions over long periods, then the markers and institutions that sustain it are precious and potentially fragile.

Social technologies

Eric Beinhocker distinguishes between physical technologies, which are ways of transforming matter and energy, such as steam engines or semiconductors, and social technologies, which are ways of organising people. Examples of social technologies include the division of labour, money, markets, the joint-stock company, the rule of law and the modern corporation. Beinhocker argues that physical and social technologies co-evolve. New physical technologies create opportunities for new forms of organisation, and new forms of organisation make it possible to exploit physical technologies at scale.

Many of the most important social technologies are mechanisms for extending trust beyond kin. The joint-stock company, pioneered in forms such as the Dutch East India Company founded in 1602, allowed large numbers of investors to pool capital in an enterprise run by managers they did not know personally. Limited liability, which spread in the nineteenth century, made it safer for strangers to invest by capping their potential losses. Accounting standards, audits and stock exchanges created markers that let investors judge companies they could not inspect directly.

These social technologies are, in effect, systems of markers combined with rules. They make strangers' promises legible and enforceable. They allow trust to be extended to people outside the family, and they create the conditions for organisations far larger than any family could sustain.

Kinship inside the firm

Modern organisations still feel kinship's pull. Many start-ups are founded by friends, classmates or family members, and early hires often come through personal networks. This makes sense. In the earliest stage of a company, when there are few markers and little formal structure, personal trust is the most reliable kind available. Founders hire people they know because they cannot yet assess strangers.

The difficulty arises when the company grows. Personal networks are limited, and they tend to be homogeneous, drawn from the same schools, regions and social circles as the founders. A company that continues to hire mainly through personal networks will struggle to scale, and it will tend to reproduce the founders' backgrounds and blind spots. Referral hiring, which many companies encourage because referred candidates often perform well and stay longer, can entrench this pattern.

The transition from kin-like trust to marker-based trust is one of the most difficult a growing organisation faces. It requires building ways of assessing and integrating strangers that do not rely on personal relationships: structured interviews, skills assessments, onboarding programmes and cultural markers that newcomers can learn. It also requires the

founders and early employees to accept that the company will increasingly be made up of people they did not choose personally and may never know well.

Family businesses face a version of the same challenge, often in more acute form. Many of the world's largest companies began as family firms, and some remain under family control. Those that have grown successfully have usually found ways to bring in professional managers and outside investors while preserving family influence. Those that have struggled have often done so because control remained concentrated among relatives who could not or would not extend trust to outsiders.

Beyond the family, beyond the coalition

The stories in this chapter share a common shape. In each case, a community needed to extend cooperation beyond the circle of people who knew one another personally. Kinship was the first answer, and it remains powerful. Coalitions defined by religion, language or trade were a second answer, extending trust to members who carried the right markers. Formal institutions such as courts, contracts and corporations were a third, extending trust to anyone who could be bound by impartial rules.

Each answer had limits. Kinship could not scale. Coalitions could scale further but excluded outsiders. Formal institutions could, in principle, include anyone, but they depended on markers and rules that could be gamed, captured or eroded.

Organisations and platforms today draw on all three answers. They rely on personal relationships within small teams, on shared cultural markers across the organisation and on formal rules and contracts at the boundaries. The art of building trust at scale is to combine these mechanisms so that each compensates for the others' weaknesses. The next chapter looks at how such combinations evolve, using Beinhocker's account of differentiation, selection and amplification to explain why some markers spread and others disappear.

CHAPTER 12

Differentiate, Select, Amplify

In 1999 a young designer named Shigetaka Kurita, working on a mobile internet service for the Japanese operator NTT Docomo, created a set of 176 small pictures, each just twelve pixels by twelve. They were meant to help users convey information and emotion quickly on tiny screens: weather symbols, a heart, a smiling face, a musical note. The pictures became known as emoji. Other Japanese carriers developed their own sets, and the idea spread. In 2010 emoji were incorporated into the Unicode standard, the system that allows computers to represent text consistently, which meant they could be displayed across devices and platforms worldwide. Since then the catalogue has grown to thousands of symbols, and emoji have become part of how billions of people communicate. In 2015 Oxford Dictionaries chose the "face with tears of joy" emoji as its word of the year.

The history of emoji is a compact illustration of what Eric Beinhocker calls the evolutionary algorithm: differentiate, select, amplify. A set of designs is created, with variations. Some designs prove useful and are adopted, while others fall away. The successful ones are copied, standardised and spread, and they become the basis for further variation. Over time, the process produces a body of designs far richer than any single designer could have planned.

Emoji are also markers. The ones people choose, the way they use them and the custom sets that communities create all signal membership. A team that has developed its own private emoji on a messaging platform has, in

effect, created a dialect. Understanding how such markers evolve helps explain why some catch on and others do not, why markers drift in meaning and why deliberate design so often produces surprising results.

The algorithm

Beinhocker's central claim in *The Origin of Wealth* is that the economy is an evolutionary system. He does not mean this as a loose metaphor. He argues that the same abstract process that produces biological diversity also produces economic diversity, operating on different materials. In biology, the designs are encoded in genes, variation comes from mutation and recombination, selection comes from survival and reproduction, and amplification comes from the spread of successful genes through a population. In the economy, the designs are what Beinhocker calls business plans, broadly understood as combinations of physical technologies, social technologies and strategies. Variation comes from experimentation and invention, selection comes from markets and other competitive pressures, and amplification comes from growth, imitation and investment.

A key feature of this process is that it does not require foresight. No one needs to know in advance which designs will succeed. The algorithm searches the space of possibilities by trying many variations and keeping what works. Beinhocker emphasises that economic actors are not the perfectly rational optimisers of traditional economic models. They are people with limited information who use rules of thumb, learn from experience and copy what seems to work. Their collective experimentation drives the evolution of the economy.

Markers evolve the same way. New markers are constantly being created, sometimes deliberately by designers and sometimes spontaneously by communities. Some are adopted widely and become part of how a society signals membership. Others remain local or disappear. Whatever survives is imitated widely, and its spread seeds the next round of invention.

Differentiation

New markers emerge from many sources. Some are designed by organisations: logos, uniforms, slogans and badges. Some arise from subcultures: slang, styles of dress, music and rituals. Some are borrowed from other groups and adapted. Some appear almost by accident, when a habit or phrase catches on and comes to signal membership.

The hashtag offers an example of a marker that began as a suggestion. In 2007 a user named Chris Messina proposed on Twitter that people use the pound sign to group messages by topic. The platform did not immediately adopt the idea, but users did, and the convention spread. Eventually Twitter turned hashtags into links, and other platforms adopted them. A small piece of notation became a marker of participation in conversations, movements and communities.

Variation is essential for markers, just as it is for genes and business plans. Without new markers, a society cannot adapt to changing circumstances or absorb new members. A culture that allows no variation in its markers becomes rigid, and its markers may lose meaning as conditions change. At the same time, too much variation can fragment a society, as different groups develop markers that others cannot read.

Organisations face this tension constantly. A company that imposes strict uniformity of markers, in dress, language and behaviour, may achieve strong cohesion but lose the capacity to innovate. A company that allows unlimited variation may become a loose federation of subcultures with little shared identity. Meadows placed self-organisation, the ability of a system to create new structure, high on her list of leverage points. Organisations that allow markers to emerge from below, within a framework of shared identity, harness that capacity.

Selection

Not every new marker survives. Selection happens through many mechanisms. Some markers are selected by markets, as when customers adopt or reject a brand. Some are selected by social pressure, as when a community embraces or ridicules a new style. Some are selected by authorities, as when a government standardises a symbol or bans another.

Markers that survive selection usually have certain qualities. They are easy to recognise, so that people can read them quickly. They are hard enough to fake that they carry information, though not so hard to acquire that legitimate members cannot obtain them. They are associated with groups that people want to belong to or be seen as belonging to. And they fit with existing markers, so that they can be integrated into a society's repertoire without too much disruption.

The history of business dress in the late twentieth century shows selection at work. For decades, many offices in the United States and Europe expected men to wear suits and ties and women to wear comparable formal clothing. In the 1990s, casual dress at work spread, encouraged by changing attitudes, by the rise of technology firms with relaxed norms and, according to many accounts, by clothing companies that promoted business casual styles to employers. Casual Fridays became common, and in many industries casual dress gradually extended to the whole week.

The selection pressure was not uniform. In banking, law and some other professions, formal dress persisted longer because it remained associated with seriousness and trustworthiness in the eyes of clients. In technology, casual dress became a marker in its own right, signalling a culture that valued substance over appearance and creativity over conformity. When Mark Zuckerberg wore a hooded sweatshirt to meetings with investors ahead of Facebook's public offering in 2012, one analyst publicly described the choice as a mark of immaturity. Others saw it as a statement that the company would not bend to Wall Street's norms. The hoodie was being selected, or rejected, by different audiences according to their own markers of credibility.

Amplification

Markers that survive selection are amplified through imitation, standardisation and network effects. Once a marker becomes associated with success or prestige, others copy it. Once enough people use it, it becomes a standard that newcomers must adopt to be recognised. And once it becomes a standard, its value increases with each new user, because it allows communication with more people.

The incorporation of emoji into Unicode is an example of amplification through standardisation. Before 2010, emoji sets varied among Japanese carriers, and a symbol sent from one phone might not display correctly on another. Standardisation allowed emoji to work across devices, which made them far more useful and accelerated their spread. Each new user made emoji more valuable for everyone else.

Corporate marketing sometimes deliberately amplifies markers. In 1991 Intel launched its Intel Inside campaign, paying computer manufacturers to display the Intel logo on their machines and in their advertising. The campaign made a component that most consumers never saw into a visible marker of quality. Buyers began to look for the logo, and manufacturers who used other chips found themselves at a disadvantage. The marker was amplified by advertising and by the network of manufacturers who adopted it, until it became a standard signal in the market.

Drift and mutation

Markers do not stay fixed once they are amplified. Their meanings drift as they are used by different groups in different contexts. A symbol that begins as a marker of one community can be adopted by another and acquire a new meaning. A slang word that signals insider status in one generation may become mainstream, losing its distinctiveness, and be replaced by a new word.

This drift creates problems for anyone who tries to design markers deliberately. A company may create a slogan or symbol with a particular meaning, only to find that employees or customers interpret it differently. A marker designed to signal openness may come to signal conformity. A symbol intended to unite may be appropriated by a faction. Designers cannot fully control what their markers will mean, because meaning is created by use.

Emoji illustrate the point. The meanings of particular symbols have shifted over time and vary among communities. Some emoji have acquired connotations their designers never intended. A symbol that one group uses innocently may carry a different meaning for another. In workplace communication, the same emoji can be read as friendly by one colleague and passive-aggressive by another, depending on the dialect each has learned.

Co-evolution with technology

Beinhocker argues that social technologies co-evolve with physical technologies. The same is true of markers. New communication technologies create opportunities for new markers and change the meaning of old ones.

The printing press enabled the spread of standardised vernacular languages, which became markers of national identity, as Benedict Anderson argued. Photography made faces and dress more widely visible, changing how people presented themselves. Television created shared cultural references that served as markers of a generation. The internet and social media created new markers, such as usernames, avatars, follower counts and verification badges, and accelerated the rate at which markers emerge and spread.

Each technological change also creates new vulnerabilities. The same technologies that allow markers to spread quickly also allow them to be copied quickly. Digital markers can be forged more easily than physical

ones, and they can be manipulated by automated systems at a scale that no individual counterfeiter could match. Chapter 25 examines this problem in detail.

Lessons for designers

Beinhocker's evolutionary view carries several lessons for leaders who want to design markers.

First, design is a form of variation, not a guarantee of selection. Leaders can create new markers, but they cannot ensure that those markers will be adopted. The best they can do is create many variations, observe which ones resonate and amplify those.

Second, selection depends on context. A marker that succeeds in one community may fail in another. Leaders should test markers with the people who will use them, rather than assuming that what appeals to designers will appeal to everyone.

Third, amplification creates path dependence. Once a marker is widely adopted, it becomes hard to change, because many people have invested in it. Leaders should be cautious about amplifying markers prematurely and should consider how easily they could be revised if circumstances change.

Fourth, meaning drifts. Leaders should monitor how markers are being used and interpreted, and be prepared to adjust them or let them go when their meanings no longer serve the organisation's purpose.

Finally, markers co-evolve with the rest of the system. Changes in technology, strategy, structure and membership all affect which markers will work. A marker system that was well suited to a company of a hundred people in one industry may be poorly suited to a company of ten thousand in several industries. Leaders need to revisit their markers as the organisation evolves.

The next chapter turns to the human drives that give markers their emotional force. Beinhocker emphasises that real economic agents are not cold calculators. They are driven by instincts of status, fairness and reciprocity that make markers feel personal, and that explain why people will sometimes sacrifice money to defend what a marker means to them.

CHAPTER 13

Status, Fairness and the Instinct Beneath the Spreadsheet

Imagine that a stranger is given a hundred dollars and told to offer you a share. You can accept or reject the offer. If you accept, you each keep your portion. If you reject it, neither of you gets anything. The stranger offers you ten dollars and keeps ninety. What do you do?

Traditional economic theory, which assumes that people maximise their own material gain, says you should accept. Ten dollars is better than nothing. Yet when versions of this scenario, known as the ultimatum game, are played in laboratories, many people reject low offers. The game was first studied experimentally by Werner Güth, Rolf Schmittberger and Bernd Schwarze in a paper published in 1982, and it has since been run thousands of times. In many studies with university students in industrialised countries, the most common offers are between forty and fifty per cent of the total, and offers below about twenty per cent are frequently rejected. People are willing to give up money to punish what they see as unfairness.

Eric Beinhocker uses findings like these to argue that real economic agents are not the cold calculators of textbook models. They are human beings shaped by evolution and culture, with instincts for fairness, reciprocity and status that influence their decisions as much as material interest does. This chapter explores those instincts and their connection to markers. Status and fairness are what give markers their emotional force. They explain why people care so intensely about the symbols of membership and rank, and why changes to those symbols can provoke reactions out of all proportion to their material significance.

Fairness is cultural

The ultimatum game has been played across many cultures, and the results vary in revealing ways. In a large project led by Joseph Henrich and colleagues and published in the early 2000s, researchers ran the game in fifteen small-scale societies around the world, including foragers, herders and farmers. The results were strikingly diverse. Among the Machiguenga of the Peruvian Amazon, offers were low and rarely rejected. Among the Lamalera whale hunters of Indonesia, where catches are shared among large crews, average offers were above half the stake. Among the Au and Gnau of Papua New Guinea, some players offered more than half, and some generous offers were rejected, a pattern the researchers connected to local norms in which accepting a gift creates an obligation to reciprocate.

The researchers found that variation in behaviour was linked to features of each society's economic life, including how much people depended on cooperation with non-kin and how integrated they were into markets. Societies in which people regularly traded and cooperated with strangers tended to make fairer offers. The finding suggests that norms of fairness toward strangers are shaped by experience with anonymous co-operation.

For this book, the cross-cultural results carry two lessons. First, fairness is a powerful motive, strong enough to override material interest. Second, what counts as fair depends on cultural markers. People from different societies bring different expectations to the same situation, and those expectations are shaped by the norms of their group. When people from different cultures meet, in a merger, a platform or a global team, they may disagree about fairness without realising that they are applying different standards.

Monkeys and cucumbers

The instinct for fairness appears to have deep roots. In a widely reported experiment published in *Nature* in 2003, Sarah Brosnan and Frans de Waal trained capuchin monkeys to exchange tokens for food. When a monkey received a piece of cucumber for its token, it usually accepted happily. But when it saw a neighbouring monkey receive a grape, which capuchins prefer, for the same token, it often refused the cucumber or threw it out of the enclosure. The monkeys appeared to be reacting not to the absolute value of their reward but to its value relative to what their neighbour received.

The study has been debated, with some researchers arguing that the monkeys' reactions reflected frustration at seeing better food rather than a sense of fairness as such. But the broader pattern, that social animals compare their rewards with those of others and react strongly to perceived inequity, has been observed in several species. Humans are not unique in caring about relative outcomes. We may be unique in the elaborate systems of markers we build to signal and contest those outcomes.

Status as a stock

Status is the relative standing of a person within a group. It is closely linked to markers, because status must be signalled to be recognised. Titles, uniforms, office sizes, invitations, seating arrangements and countless other markers indicate where people stand. In anonymous societies, these markers are essential, because people cannot know everyone's standing from personal experience.

Status matters to people for reasons that go beyond vanity. Studies of British civil servants, known as the Whitehall studies and led by Michael Marmot from the late 1960s onward, found a social gradient in health. Civil servants in lower grades had higher rates of illness and death than those in

higher grades, even after accounting for factors such as smoking and access to health care. The finding suggests that position in a hierarchy can affect well-being in ways that are not simply about income.

In systems terms, status behaves like a stock. It accumulates through recognition and achievement and is drained by humiliation and failure. Like trust, it carries history and changes slowly. And like trust, it is inferred from markers, which serve as information flows about the stock. When markers of status change, people experience the change as a gain or loss in the underlying stock, even if nothing else has changed.

When markers carry status

This explains why changes to markers can provoke intense reactions. A reorganisation that changes job titles, moves people to different offices or alters reporting lines may leave pay and responsibilities unchanged, yet employees often react as if they had been demoted. In a real sense, they have been. Their markers of status have changed, and the people around them will read the change accordingly.

A study by David Card, Alexandre Mas, Enrico Moretti and Emmanuel Saez, published in 2012, illustrated how sensitive people are to relative position. After a newspaper in California made the salaries of University of California employees searchable online, the researchers informed a random sample of employees about the website. Those who learned that they earned less than the median for their unit and occupation reported lower job satisfaction and were more likely to say they were looking for a new job. Those who learned that they earned more than the median showed no corresponding gain in satisfaction. Information about relative pay, a marker of status, affected well-being asymmetrically. Losses relative to peers hurt more than gains pleased.

This asymmetry has direct consequences for leaders. Any change that makes some people's markers of status look worse relative to others will produce strong negative reactions, while improvements will be taken for

granted. Mergers, restructurings and pay reforms all create such changes. Leaders who focus only on material outcomes, such as whether people's pay has risen or fallen, will underestimate the reactions they face.

Costly signals

Why do status markers so often involve waste or difficulty? Why do people pursue degrees from expensive universities, buy luxury goods that perform no better than cheaper alternatives or put themselves through gruelling initiation rituals?

Several theories converge on an answer. In 1899 the economist Thorstein Veblen described conspicuous consumption, the display of wealth through spending on visible luxuries, as a way for the leisure class to signal status. In 1973 the economist Michael Spence proposed a model of job market signalling, in which education serves partly as a signal of underlying ability. Because education is costlier for less able people to complete, the act of completing it signals ability to employers, even if the education itself does not increase productivity. Spence shared the Nobel memorial prize in economics in 2001 for this work. In biology, Amotz Zahavi proposed the handicap principle in 1975, arguing that some animal signals, such as the peacock's elaborate tail, are reliable precisely because they are costly. Only a healthy, strong animal can afford to carry such a burden.

The common thread is that costly signals are hard to fake. If a marker is cheap, anyone can display it, and it carries little information. If a marker is expensive, in money, effort, risk or time, only those who can afford the cost will display it, and the marker becomes a reliable indicator of something. This is the same principle that makes banknotes secure. The cost of forging the marker determines how much trust it can carry.

Costly signals have a downside. They consume resources that might be used productively, and they can exclude people who have the underlying qualities but cannot pay the cost. A credential that signals ability by being expensive also excludes able people who cannot afford it. An initiation ritual that signals commitment by being painful may deter people who

would have been committed members. Designers of markers face a constant trade-off between reliability, which favours cost, and accessibility, which favours cheapness.

Fairness and the marker system

Fairness instincts also shape how people respond to marker systems. If a system for awarding markers, such as promotions, badges or ratings, is seen as fair, people accept its outcomes even when they lose. If it is seen as unfair, people reject its outcomes even when they win, and they may withdraw cooperation from the whole system.

Research on organisational justice has distinguished several dimensions of fairness. Distributive fairness concerns whether outcomes are fair. Procedural fairness concerns whether the process used to reach them is fair. Interactional fairness concerns whether people are treated with respect and given explanations. Many studies have found that procedural and interactional fairness can matter as much as distributive fairness, and sometimes more, for people's acceptance of decisions.

For markers, this means that how markers are awarded matters as much as which markers are awarded. A promotion system that is transparent, consistent and explained will be more widely accepted than one that is opaque, inconsistent or unexplained, even if both produce similar outcomes. A rating system that users understand and can contest will be trusted more than one that operates as a black box. Leaders who want markers to carry trust need to attend to the fairness of the processes behind them.

The inflation of titles

Job titles are among the most visible status markers inside organisations, and they are prone to the same inflation as grades and credentials. When a company cannot offer higher pay, it may offer a grander title. When employees compare their titles with peers at other companies, they push for titles that match. Over time, titles drift upward. Roles once called

coordinator become manager, managers become directors and directors become vice presidents. Start-ups, which often cannot match the salaries of larger firms, are especially prone to generous titles.

Title inflation has costs that are easy to overlook. Within the organisation, it compresses the hierarchy at the top and makes titles less informative about actual responsibility. Outside it, recruiters and partners learn to discount titles from certain companies or industries, much as experienced users of rating systems learn to read small differences in scores. Employees who move from a company with inflated titles to one with conservative titles may find themselves offered what looks like a demotion, which they experience as a loss of status even if the work is more substantial.

There is also a fairness dimension. When titles are awarded through negotiation rather than clear criteria, people who negotiate more aggressively receive grander titles for the same work. Because the willingness and ability to negotiate varies across individuals and groups, title inflation can produce systematic inequities. Organisations that want titles to carry trust should tie them to defined levels of responsibility, apply them consistently and resist the temptation to use them as a substitute for pay. A title that means the same thing everywhere in the organisation is a far more valuable marker than one that reflects the outcome of a private bargain.

The instinct beneath the spreadsheet

Beinhocker's broader argument is that economic models built on assumptions of narrow self-interest and perfect rationality miss much of what drives real behaviour. People care about fairness, reciprocity and status. They follow rules of thumb, copy others and respond to social cues. These drives are not irrational in an evolutionary sense. They helped our ancestors cooperate in groups and navigate social hierarchies. But they mean that markers, which signal fairness, reciprocity and status, carry emotional weight that spreadsheets do not capture.

A leader looking at a merger spreadsheet may see synergies, cost savings and revenue opportunities. The employees affected by the merger will see changes in titles, offices, reporting lines, logos and rituals, each of which signals something about their status and the fairness of their treatment. A leader looking at a platform's metrics may see engagement, conversion and retention. The users will see ratings, badges and rankings that signal their standing in the community and the fairness of the platform's rules. If the leader ignores these signals, the system will behave in ways the spreadsheet cannot predict.

The chapters that follow in Part III apply this understanding to the markers that organisations use every day: brands, credentials, uniforms, dialects, the collisions of mergers and the persistence of culture. In each case, the material stakes are real, but the emotional stakes of status and fairness often matter more.

PART III

Markers at Work

Every organisation large enough to contain strangers runs on markers. Part III looks at the ones leaders meet most often, from brands and credentials to dress codes and office dialect, and at what happens when markers collide in mergers or resist redesign.

CHAPTER 14

The Badge and the Brand

The word "brand" began as a mark burned into the hide of an animal. Ranchers and herders branded livestock to show ownership, so that a cow that wandered onto someone else's land could be identified and returned, or so that a thief could be caught. The brand was a marker of membership in a herd, readable by strangers who had never met the owner.

Later the word extended to marks on goods. Medieval craftsmen stamped their work with marks that identified the maker or the guild. In England, a statute of 1300 required silver to meet a standard of purity, and over time a system of marks applied at Goldsmiths' Hall in London came to certify that standard, giving us the word "hallmark." In 1876, shortly after a new law created a register of trademarks in the United Kingdom, the brewer Bass registered its red triangle, which is often cited as the first trademark registered under the act. The red triangle told drinkers in distant pubs that the beer came from a particular brewery and could be expected to taste a particular way.

The modern brand is a descendant of these marks. It tells strangers what to expect from a product or company without requiring them to know the people behind it. It is a marker of trust. But brands also do something that hallmarks and livestock brands did not. They create communities. People who buy a brand, wear it or work for it come to see themselves as members of a group, and they read other people's brands as signs of membership too. This chapter explores brands as markers of belonging, and what happens when companies change them.

Brands as tribes

Some brands inspire loyalty so intense that it resembles national or religious identity. Harley-Davidson is a famous example. In 1983 the company created the Harley Owners Group, a club for owners that organised rides, rallies and events. Members wear the brand's logo on jackets and patches, and some have it tattooed on their bodies. For many owners, riding a Harley is not just a choice of motorcycle. It is a membership in a community with its own markers, rituals and values.

Apple has cultivated a similar sense of belonging. Its 1997 Think Different campaign, which featured images of figures such as Albert Einstein and Martin Luther King Jr., invited customers to see themselves as members of a community of creative nonconformists. Apple products became markers of that identity. Users who carried a particular laptop or phone signalled something about who they were, and many developed strong loyalty to the brand.

These examples illustrate a general principle. When a brand becomes a marker of identity, customers relate to it not only as buyers but as members. They defend it against critics, recruit others to it and feel personally affected by changes to it. This creates enormous value for companies, because loyal members buy repeatedly and recommend the brand to others. It also creates risks, because members feel entitled to a say in how the brand evolves.

New Coke

On 23 April 1985 the Coca-Cola Company announced that it was changing the formula of its flagship soft drink for the first time in nearly a century. The decision followed years of losing market share to Pepsi, whose advertising had featured blind taste tests in which many people preferred Pepsi's sweeter flavour. Coca-Cola's own extensive taste tests, involving

around two hundred thousand consumers according to widely reported accounts, had found that people preferred the new formula to both the old Coke and Pepsi.

The public reaction was ferocious. Customers called and wrote to the company in large numbers to complain. Some stockpiled the old formula. Protest groups formed. People described the change as a betrayal. On 11 July, seventy-nine days after the launch, the company announced that the original formula would return as Coca-Cola Classic. The announcement was reported as major news.

What had gone wrong? The taste tests measured flavour preference. They did not measure what Coca-Cola meant to its drinkers. For many Americans, Coke was a marker of identity, associated with childhood, family, tradition and the nation itself. Changing the formula felt to them like changing a flag. The company had treated a marker of belonging as a product attribute, and it discovered that the people who carried the marker regarded it as theirs.

The episode is often told as a marketing blunder, and it was. It is also a vivid lesson in how markers work. The value of the old formula was not in its taste alone. It was in the shared expectation, built over decades, that Coke would always be Coke. Changing the formula drained the stock of trust that expectation represented, and the only way to stop the outflow was to restore the marker.

Tropicana and Gap

The New Coke episode was not unique. Companies regularly change logos, packaging and names, and they are regularly surprised by the reactions.

In early 2009 PepsiCo introduced new packaging for Tropicana Pure Premium orange juice, replacing the familiar image of an orange pierced by a straw with a close-up of a glass of juice and a minimal design. According to reports in the trade press, sales fell by around a fifth in the following two months, and many customers complained that they could not find the product on shelves or that it looked like a cheap store brand. The company reverted to the old design.

In October 2010 the clothing retailer Gap unveiled a new logo, replacing its long-familiar white letters in a navy box with a different typeface and a small blue square. The reaction online was overwhelmingly negative, and within about a week the company announced that it would return to the original logo.

In both cases, the companies had treated visual markers as design elements that could be updated for freshness. Customers treated them as signals of identity and quality. The old Tropicana carton told shoppers that this was the familiar, trusted product. The new one looked like something else. The old Gap logo was a marker that many people had grown up with. The new one felt like a break with that history.

Rebrands that work

Not every rebrand fails. Companies often change their markers successfully, especially when the change reflects a real change in the company and is introduced with care. Apple replaced its rainbow-striped logo with a monochrome version in the late 1990s without significant backlash. Many companies have updated logos gradually over decades, keeping the essential elements recognisable while modernising details.

Successful rebrands tend to share certain features. They preserve continuity, keeping enough of the old marker that loyal members can still recognise it. They are connected to substance, reflecting changes in products, strategy or values that members can see. They are introduced in ways that respect the existing community, explaining the change rather than simply imposing it. And they are tested not only for aesthetic appeal but for what they signal to the people who carry the marker.

Failed rebrands tend to share the opposite features. They break abruptly with the past, they are disconnected from any visible change in substance and they are imposed without explanation. They treat markers as design problems rather than as signals of membership.

Even large, deliberate rebrands that reflect genuine strategic change can provoke strong reactions. When Twitter was renamed X in July 2023 and its bird logo was replaced, many users and observers expressed dismay at

the loss of a familiar symbol, and the old name and logo continued to be widely used in conversation long after the change. Whatever one thinks of the strategy behind the rebrand, the episode showed how attached communities become to the markers of the platforms they inhabit.

When the leader damages the brand

Brands can also be damaged by the words of the people who lead them. In 1991 Gerald Ratner, chief executive of the British jewellery chain Ratners, gave a speech to the Institute of Directors in which he joked that some of his company's products were "total crap." The remark was widely reported, and the company's value fell dramatically in the following period. The name Ratner became shorthand in British business for a leader who destroys his own brand with careless words.

The Ratner episode illustrates how fragile the stock of trust behind a brand can be. Customers who bought Ratners jewellery had relied on the brand as a marker of acceptable quality at an affordable price. When the chief executive openly mocked the products, he did not change the jewellery. He changed what the brand signalled, and he humiliated the customers who carried it. The outflow of trust was immediate.

The employer brand

Brands are markers not only for customers but for employees. A company's name and reputation tell potential hires what to expect, and they tell employees' friends and families something about where they work. People who work for admired companies often display the brand proudly, wearing company clothing and mentioning their employer in conversation. People who work for companies with damaged reputations may avoid mentioning where they work.

The employer brand shapes who joins an organisation and how they see themselves. A company known for innovation will attract people who see themselves as innovators. A company known for stability will attract people who value security. These self-selection effects reinforce the brand,

creating a reinforcing loop. Over time, the people inside the company come to embody the brand's markers, which makes the brand more credible to outsiders.

The loop can also run in reverse. When a company's reputation is damaged, by scandal, layoffs or a perceived betrayal of its values, employees may feel that their own identity has been damaged. Some leave. Those who remain may disengage. Recruitment becomes harder. The employer brand, like the customer brand, is a stock of trust that can be drained quickly and rebuilt only slowly.

Badges inside the building

Inside organisations, the most literal brand is the employee badge. It shows the company's logo, the employee's name and photograph, and often a colour or symbol indicating their status: employee, contractor, visitor. It grants access to buildings and systems. It is a marker that lets security staff, receptionists and colleagues identify members at a glance.

Badges often carry finer distinctions too. In many companies, full-time employees and contractors have badges of different colours, even when they do similar work side by side. The distinction may be necessary for legal or security reasons, but it also signals status. Contractors with a different badge colour may be excluded from certain events, meetings or perks, and they may feel like second-class members of the organisation. Some have described the experience in terms that echo the "real people" distinction discussed in Chapter 6.

The design of badges is a small but revealing example of how markers carry meaning beyond their practical function. A badge is a security device. It is also a statement about who belongs and how fully. Leaders who want to build a cohesive organisation should pay attention to the messages their badges send, and to whether those messages match the culture they want to create.

Brands across borders

Brands that cross cultural boundaries face the problem described in Chapter 9: what a marker means depends on the community reading it. A colour, symbol or name that signals quality in one market may signal something quite different in another. Companies expanding abroad routinely adapt their names, packaging and advertising for local audiences, and the history of international marketing contains many cautionary tales.

Some of those tales are themselves unreliable, which is a lesson of its own. A widely repeated story holds that the Chevrolet Nova sold poorly in Spanish-speaking markets because "no va" can be read in Spanish as "doesn't go." The story has been examined repeatedly and is generally regarded as a myth; the car is reported to have sold reasonably well in several Latin American markets. That the story persists says something about how attractive simple explanations of marker failure can be. It is more satisfying to believe that a brand failed because of a single unlucky word than to investigate the many factors that shape whether customers in a new market trust a foreign product.

The genuine challenges of cross-border branding are subtler. Customers in a new market lack the history that gives a brand its meaning at home. They have not grown up with it, seen it in their parents' kitchens or associated it with formative experiences. To them, a famous foreign brand may be simply an unfamiliar name with an expensive price. Building trust requires either borrowing local markers, for example through partnerships with trusted local firms, or accumulating a new history of positive experiences in the new market, which takes time. Companies that assume their domestic reputation will travel with them often discover that a brand's stock of trust is held in local currency.

Designing brands as markers

The cases in this chapter suggest several principles for designing brands that function well as markers.

Recognise that brands are owned in part by the communities that carry them. Customers and employees who identify with a brand have a stake in it, and changes that ignore that stake will provoke resistance.

Connect markers to substance. A brand that signals qualities the company actually has will accumulate trust. A brand that signals qualities the company lacks will eventually be exposed, and the outflow will be larger than if the brand had been honest.

Change markers gradually and with explanation. When change is necessary, preserve continuity where possible, explain the reasons and involve the community in the process.

Remember that leaders are markers too. What leaders say and do signals what the brand stands for. A careless remark can drain trust faster than any marketing campaign can build it.

Pay attention to internal markers. Badges, titles and other signals inside the organisation shape how employees see themselves and one another. They should be designed with the same care as external brands.

The next chapter turns to a different kind of marker, one that signals not membership in a company's community but competence in a profession: the credential.

CHAPTER 15

Credentials and Licences

When you lie on an operating table, you are trusting a stranger with your life. You may have met the surgeon once, briefly, in a consultation. You have no way of assessing their skill directly. You cannot watch them operate on others, review their technique or check their decision making under pressure. What you can do is rely on markers: the degree on the wall, the licence that permits them to practise, the board certification that indicates specialist training, the hospital that has granted them privileges. Each of these is a credential, a marker that tells strangers that a person has met a standard set by some authority.

Credentials are among the most important markers in modern societies, and among the most contested. They allow people to trust strangers with tasks of great consequence, from medicine and law to engineering, aviation and accounting. They also create barriers to entry, raise costs and sometimes protect incumbents more than they protect the public. This chapter examines credentials as markers of competence, the systems that produce them and the ways those systems succeed and fail.

The rise of the licence

Occupational licensing, the legal requirement to hold a licence before practising a trade or profession, has expanded dramatically in some countries over the past several decades. In the United States, the economist Morris Kleiner and colleagues have estimated that the share of workers required to hold a state licence rose from around five per cent in the 1950s to somewhere between a quarter and nearly a third of the workforce in the

early twenty-first century, depending on the data source. Licensed occupations now include not only doctors, lawyers and pilots but also, in various states, hairdressers, interior designers, florists, auctioneers and many others.

Supporters of licensing argue that it protects consumers by ensuring minimum standards of competence and safety. In professions where mistakes can kill, such as medicine and aviation, the case is strong. But for many occupations, critics argue, the link between licensing and public safety is weak, and the main effect of licensing is to restrict entry, raise prices and protect existing practitioners from competition. A 2015 report published by the Obama administration, *Occupational Licensing: A Framework for Policymakers*, reviewed the evidence and concluded that licensing requirements often vary widely between states for the same occupation, that they can impose substantial costs on workers and consumers and that they can make it harder for workers to move between states.

The variation is telling. If licensing requirements were based purely on what is needed to protect the public, one would expect them to be broadly similar across states for the same occupation. Instead, the number of hours of training required for a licence in some occupations differs several-fold from one state to another. This suggests that licensing requirements reflect local politics, lobbying by professional groups and historical accident as much as public safety.

Some states have begun to respond. In 2019 Arizona became the first state to recognise occupational licences issued by other states for most occupations, allowing licensed workers who move to Arizona to practise without obtaining a new licence. Other states have since adopted similar measures. These reforms treat licences as markers that should be portable across jurisdictions, like passports, rather than local barriers.

Credentials as tribal markers

Credentials do more than certify competence. They also create professional communities with their own identities, rituals and boundaries. Doctors, lawyers, engineers and accountants each belong to professions that function

in some respects like the peoples described in Chapter 6. They have shared languages, full of technical terms that outsiders cannot follow. They have rituals of initiation, such as examinations, oaths and ceremonies. They have markers of membership, such as titles, uniforms and letters after their names. And they have a sense of who counts as a real member and who does not.

Medicine offers vivid examples. Many medical schools in the United States and elsewhere now hold white coat ceremonies, in which new students are formally given their first white coat and often recite an oath. The practice is generally traced to a ceremony introduced at Columbia University's medical school in 1993. The ceremony marks entry into the profession and signals to students that they are now members of a community with obligations to patients and to one another.

These professional identities can be valuable. They create a sense of shared responsibility, encourage members to uphold standards and provide a basis for peer accountability. A doctor who violates professional norms may face sanction not only from regulators but from colleagues who feel that the violation reflects on the profession as a whole.

Professional identities can also become exclusionary. Professions may resist the entry of practitioners trained in other countries, other traditions or other routes, even when those practitioners are competent. They may defend their boundaries against neighbouring occupations that seek to take on some of their tasks, such as nurses performing procedures once reserved for doctors. And they may close ranks against outsiders who criticise them, protecting members who should be held to account.

What credentials signal

The deeper question about credentials concerns what they actually signal. Michael Spence's signalling model, discussed in Chapter 13, proposed that education can serve as a signal of ability even if it does not increase productivity directly. Economists have debated for decades how much of the return to education reflects learning and how much reflects signalling. One piece of evidence often cited in favour of a signalling component is the

so-called sheepskin effect, the observation that earnings tend to jump when people complete a degree, beyond what would be expected from the additional years of study alone. Completing the degree, in this view, signals qualities such as persistence and conscientiousness that employers value.

The debate matters for marker design. If credentials mainly signal learning, then the way to improve them is to improve the education behind them. If they mainly signal underlying qualities, then the way to improve them is to find cheaper, more accessible ways of signalling those qualities. Most economists believe that both effects are real, though they disagree about their relative size.

For employers, the practical question is whether a credential predicts performance in a particular job. Some credentials do so reliably, because the knowledge and skills they certify are directly relevant. A licence to fly a commercial aircraft certifies skills that are essential for the job. Other credentials predict performance weakly, because they certify knowledge that is only loosely related to the work or because the qualities they signal can be assessed more directly in other ways.

Dropping the degree

In recent years, a number of large employers and public bodies have announced that they would remove degree requirements for many roles. In 2022 the state of Maryland announced that it would remove four-year degree requirements for many state government jobs, and other states followed. Several large technology and consulting firms have made similar announcements for some positions, emphasising skills-based hiring. Google launched a set of career certificates in fields such as data analytics and project management, promoting them as alternatives to degrees for entry-level roles.

These moves reflect several pressures. Employers in tight labour markets want access to more candidates. Critics of credential inflation, discussed in Chapter 8, argue that degree requirements exclude capable workers without improving performance. And advocates for equity point

out that degree requirements disproportionately exclude people from lower-income backgrounds and from groups underrepresented in higher education.

Removing a credential requirement is not the same as removing the need for markers. Employers still need some way of assessing strangers. Skills-based hiring replaces one set of markers with another, such as work samples, skills tests, portfolios and certifications. The question is whether the new markers are more accurate, more accessible and harder to fake than the old ones. Early evidence suggests that announcing the removal of degree requirements does not automatically change hiring practices. Some analyses have found that companies that dropped degree requirements on paper continued to hire mainly degree holders in practice. Changing a marker system requires changing the habits of the people who read the markers, not just the rules that define them.

Digital credentials

New technologies have created new forms of credentials. Digital badges, promoted by projects such as Mozilla's Open Badges initiative launched in 2011, allow organisations to issue verifiable records of specific skills and achievements. Online learning platforms issue certificates for completed courses. Professional networks display endorsements and skills lists.

These digital credentials can be more granular than traditional degrees, certifying specific skills rather than broad programmes of study. They can be issued by many kinds of organisation, not just universities and licensing boards. And they can be verified online, which makes them harder to forge than paper certificates.

They also face the familiar problems of marker systems. When anyone can issue a credential, the value of credentials in general may be diluted, a version of the tragedy of the commons. When credentials can be earned cheaply, they may signal little. When platforms display endorsements that cost nothing to give, the endorsements may become meaningless. A credential is only as valuable as the trust that strangers place in its issuer and its process.

Credentials inside organisations

Organisations create their own internal credentials: certifications for using particular systems, qualifications for leading certain kinds of projects, training completions required before promotion. These internal credentials function as markers that let managers and colleagues trust people they do not know personally.

They are subject to the same dynamics as external credentials. They can inflate, as requirements multiply. They can become barriers, as people who lack them are excluded from opportunities they could handle. They can become rituals, completed for the sake of compliance rather than learning. And they can become targets, as people focus on acquiring them rather than on developing the skills they are meant to certify.

Leaders should periodically review their internal credentials with the same questions that critics ask of occupational licensing. Does this credential predict performance? Does it protect against real risks? Is it accessible to everyone who might reasonably need it? Is it harder to fake than the alternatives? Could the same assurance be obtained more cheaply or more directly?

The engineer who drives a taxi

A familiar figure in many cities with large immigrant populations is the professional who cannot practise. A doctor trained in one country drives a taxi in another. An engineer works in a warehouse. A lawyer works as a clerk. Their skills are real, but their credentials are not recognised, or recognising them requires years of further study, examinations and fees that they cannot afford while supporting a family.

From the perspective of this book, such cases are failures of marker portability. The credential that signalled competence in one society does not carry across the boundary into another, and the receiving society has not built an efficient way to verify the underlying competence directly. Some of the caution is justified. Training standards vary, and a licensing

body has a duty to protect the public. But much of the barrier reflects the difficulty of reading unfamiliar markers and, in some cases, the interest of incumbents in limiting competition.

Several countries and professions have developed approaches that treat the problem as one of verification rather than exclusion. These include bridging programmes that assess what an applicant already knows and train only the gaps, supervised practice periods in which competence can be observed directly, and mutual recognition agreements between jurisdictions with comparable standards. Each approach replaces a crude marker test, in which the foreign credential is simply not accepted, with a more precise assessment of the quality the credential was supposed to signal. The economic case is strong, since societies that leave skilled immigrants underemployed lose the value of their training. The human case is stronger still. A marker system that cannot recognise competence when it arrives from elsewhere is, in effect, declaring that only some people count as real practitioners.

Principles for credential design

The evidence on credentials points to several principles.

Tie credentials to real risks. Credentials are most justified where mistakes are costly and hard to detect, as in medicine, aviation and structural engineering. Where risks are low or mistakes are easily detected and corrected, lighter forms of assurance may suffice.

Measure what matters. A credential should certify the knowledge, skills or qualities that actually predict performance. Credentials that certify loosely related knowledge are likely to exclude good candidates and admit poor ones.

Make credentials portable. Markers that are recognised across jurisdictions and organisations allow people to move and contribute more widely. Local credentials that do not travel create barriers without adding assurance.

Keep credentials current. Skills decay and requirements change. Credentials that never expire may certify competence that no longer exists. Periodic renewal, continuing education or re-examination keep the marker connected to the underlying quality.

Maintain direct assessment. Credentials should supplement, not replace, the ability to assess people directly. Organisations that rely entirely on credentials lose the capacity to judge, and they become vulnerable when credentials are forged, as the next chapter shows.

CHAPTER 16

The Great Impostor

In 1951, during the Korean War, a Royal Canadian Navy destroyer named HMCS *Cayuga* was operating off the Korean coast. Aboard was a surgeon-lieutenant who called himself Joseph Cyr. When wounded South Korean soldiers were brought aboard, the ship's doctor operated on them, according to the many accounts of the episode, including at least one case in which he removed a bullet lodged close to a man's heart. The operations were reported in the press as heroic.

The man was not Joseph Cyr, and he was not a doctor. He was Ferdinand Waldo Demara, an American with no medical training who had borrowed the identity and credentials of a real Canadian physician he had met. When news of the operations reached Canada, the real Dr Cyr's family recognised the name, and the deception unravelled. Demara was quietly discharged and sent back to the United States. Over his life he posed, by the accounts of his biographers, as a monk, a teacher, a prison official and a psychologist, among other roles. His story was told in Robert Crichton's 1959 book *The Great Impostor* and in a 1961 film of the same name.

Demara's career is a case study in the vulnerability of anonymous societies. He did not hack any system or forge any elaborate document in the modern sense. He simply acquired the markers that other people relied on to judge competence: a name, a set of credentials, a uniform and a confident manner. In a society of strangers, those markers were enough to get him into an operating theatre.

Why impostors succeed

Impostors succeed because anonymous societies depend on markers, and markers can be acquired without the qualities they signal. The previous chapter described credentials as markers that let strangers trust people with consequential tasks. An impostor exploits the gap between the marker and the quality. As long as no one checks the underlying quality directly, the marker carries the impostor through.

Several features of anonymous societies make this easier. First, people rarely verify markers. Checking every credential, reference and claim would be prohibitively costly, so people check only occasionally or rely on others to have checked. Second, markers reinforce one another. A person who wears the right uniform, uses the right vocabulary and carries the right documents is assessed as a whole, and each marker makes the others more credible. Third, people are reluctant to challenge someone who displays high-status markers. Questioning a doctor, an officer or an executive feels presumptuous, and the social cost of being wrong discourages scrutiny. Fourth, institutions often assume that someone else has done the checking. Each party relies on the trust extended by others, creating what might be called a trust cascade.

Demara exploited all of these features. He presented himself with confidence, used the vocabulary of whatever role he was playing, and relied on the fact that the institutions he joined had little reason to doubt credentials that appeared genuine.

A cautionary note on famous impostors

Impostor stories are popular, and some of the most famous have proved less reliable than their tellers. Frank Abagnale, whose claims to have posed as an airline pilot, a doctor and a lawyer as a young man were the basis for the 2002 film *Catch Me If You Can*, has had many of his claims challenged. Investigations, including a 2020 book by the researcher Alan Logan, have argued that much of his story was exaggerated or invented. There is a

certain irony in an impostor's story that is itself partly an imposture. It also offers a lesson. Markers of credibility, such as a best-selling memoir, a Hollywood film and a long career as a speaker, can carry a story far beyond its evidence.

This book uses only impostor cases that are well documented in contemporary records and multiple sources. Even so, the details of some cases are inevitably drawn from accounts written after the fact, and readers should treat vivid specifics with appropriate caution.

Theranos

The most consequential impostor story of recent decades did not involve a fake identity. It involved a real person whose company made claims its technology could not support, and who surrounded herself with markers of credibility that persuaded investors, partners and regulators to look away from the gap.

Elizabeth Holmes founded the company that became Theranos in 2003, after leaving Stanford University. She claimed that Theranos had developed technology capable of running a wide range of blood tests on a few drops of blood from a finger prick. The claim, if true, would have transformed medical testing. Theranos raised large sums from investors, reached a valuation reported at around nine billion dollars and struck a partnership with the pharmacy chain Walgreens to offer tests in its stores.

The technology did not work as claimed. Reporting by John Carreyrou of the *Wall Street Journal*, beginning in October 2015, revealed that the company was running many tests on conventional machines bought from other manufacturers and that its own devices produced unreliable results. Employees who raised concerns, including Tyler Shultz, a grandson of the former US secretary of state George Shultz, faced pressure and legal threats. Regulators eventually sanctioned the company, and it was dissolved in 2018. In January 2022 a jury convicted Holmes on four counts of defrauding investors, and she was later sentenced to more than eleven years in prison.

What makes Theranos relevant to this book is the role of markers. Holmes was widely reported to have adopted a style of dress, black turtlenecks, that evoked Steve Jobs, and to speak in an unusually deep voice. Whatever her intentions, these choices aligned her with the markers of a visionary technology founder. Theranos assembled a board of directors that included former secretaries of state George Shultz and Henry Kissinger, former defence secretary William Perry and the retired general James Mattis, among other prominent figures. Few board members had expertise in laboratory medicine, but their presence signalled credibility to investors and partners. Prominent media coverage, including magazine covers, amplified the image.

Each marker reinforced the others. Investors who saw the board assumed that such distinguished people must have checked the technology. Partners who saw the investors assumed that due diligence had been done. Journalists who saw the partners assumed the company was legitimate. The trust cascade carried Theranos for years, until a reporter and a handful of whistleblowers checked the underlying quality directly.

The diploma factory

Impostors who forge credentials individually are limited in scale. Organisations that manufacture fake credentials can operate at far greater scale. In 2015 the *New York Times* published an investigation into Axact, a Pakistani software company, reporting that it ran a vast network of online schools and universities that sold fake degrees and diplomas to customers around the world. The report described websites for institutions that existed only online, with invented faculty, fake accreditation bodies and testimonials. Pakistani authorities raided the company's offices after the report.

Diploma mills exploit the fact that credentials are markers whose value depends on the issuer's reputation. By creating fake issuers that look like real ones, with names, logos, websites and accreditation claims, they produce markers that can fool employers who do not check carefully. Some

customers may know they are buying a fake. Others may believe they are enrolling in a legitimate distance-learning programme. Either way, the market for fake credentials undermines the value of real ones.

The heiress

A different kind of impostor works not through professional credentials but through social markers. Between 2013 and 2017, a young woman calling herself Anna Delvey moved through the upper levels of New York society, presenting herself as a wealthy German heiress. Her real name was Anna Sorokin. She stayed in expensive hotels, dined at fashionable restaurants and sought financing for an exclusive private club, while leaving behind unpaid bills and, according to prosecutors, defrauding banks, hotels and acquaintances. In 2019 she was convicted on multiple counts, including grand larceny.

Sorokin's methods relied on markers of wealth and belonging: expensive clothes, confident manners, familiarity with the rituals of elite social life and a plausible story. People who encountered her read these markers and assumed that she was who she claimed to be. Checking would have seemed rude, and in a world where wealth is often discreet, the absence of verifiable evidence was not immediately suspicious. Her case shows that social markers, like professional credentials, can carry a stranger into circles of trust that would otherwise be closed.

The cost of checking

If impostors exploit the failure to check markers, why not simply check more? The answer is that checking is costly, and in an anonymous society the costs multiply. Every credential verified, every reference called and every claim investigated takes time and money. If everyone checked everything, the efficiency gains that markers provide would disappear. Markers are valuable precisely because they allow people to extend trust without checking.

The challenge, then, is not to eliminate trust in markers but to check intelligently. Several principles can help.

Check in proportion to consequence. Markers that grant access to high-stakes roles, such as surgery, control of financial assets or access to critical systems, deserve more rigorous verification than markers that grant access to low-stakes roles.

Check the source, not just the marker. Verifying that a credential was actually issued by the claimed institution, and that the institution is legitimate, is more reliable than inspecting the credential itself.

Check independently. Relying on others to have checked creates trust cascades. Independent verification, even of a sample, breaks the cascade.

Check the substance periodically. A marker that was verified once may no longer reflect reality. Periodic direct assessment of skills, performance or claims keeps markers honest.

Protect the people who check. Whistleblowers and sceptics who question markers are often treated as troublemakers. In the Theranos case, employees who raised concerns faced pressure. Organisations that want to detect impostors need to make it safe to question even high-status markers.

Impostors inside organisations

Not every impostor is a con artist. Inside organisations, people sometimes acquire markers they have not fully earned: titles that outstrip their skills, credentials padded on résumés, reputations built on others' work. Surveys of hiring managers have often reported that a substantial share of résumés contain exaggerations or false claims. Most such cases are minor, but some can be serious, as when an executive's claimed qualifications turn out to be false.

Organisations also create conditions that encourage people to display markers they do not feel they deserve. The phenomenon often called impostor syndrome, a persistent feeling of being a fraud despite evidence of competence, was first described by the psychologists Pauline Clance and Suzanne Imes in a 1978 paper. It is, in a sense, the inverse of the impostor problem. People who carry legitimate markers doubt that they deserve

them, while some who carry illegitimate markers feel no doubt at all. Both phenomena reflect the gap between markers and the qualities they are supposed to signal.

What impostors teach

Impostors are unwelcome, but they are useful teachers. Each successful imposture reveals a gap in a marker system: a place where markers are trusted without verification, where high-status markers discourage scrutiny or where trust cascades substitute for independent checking. Studying how impostors succeeded can show organisations where their own marker systems are vulnerable.

The deeper lesson is that anonymous societies always carry this risk. Because markers allow strangers to be trusted without being known, they also allow strangers to be trusted when they should not be. The ant colony that admits a butterfly caterpillar, the navy that commissions a fake surgeon and the investors who fund a fraudulent start-up are all paying the price of a system that runs on markers. The price is usually worth paying, because the alternative, a society in which no one can be trusted without being known, would be far poorer. But the price must be managed. The next chapter turns to markers that are less consequential than credentials but far more pervasive: the uniforms and dress codes that signal belonging every day.

CHAPTER 17

Uniforms and Dress Codes

When Robert Peel established the Metropolitan Police in London in 1829, he faced a problem of public trust. Many Britons were deeply suspicious of the idea of a professional police force, which they associated with the repressive state police of continental Europe. The new force needed to be visible, so that people could identify officers and know that they were being watched over. It also needed to look different from the army, whose red coats were associated with the suppression of riots and protests. The solution was a uniform of dark blue coats and tall hats, deliberately civilian in style. The colour and cut signalled that the police were citizens in uniform, not soldiers.

Uniforms are among the most explicit markers of membership in any society. They tell strangers, at a glance, who someone is, what role they play and what authority they hold. They also tell the wearer something about themselves. Putting on a uniform is an act of joining. This chapter looks at uniforms and their softer cousins, dress codes, as markers that shape trust inside and outside organisations.

What a uniform says

A uniform carries several messages at once. It identifies the wearer as a member of a particular organisation or profession. It signals the wearer's role and often their rank. It indicates that the wearer is acting in an official capacity rather than as a private individual. And it promises a certain standard of behaviour, since the wearer represents the organisation and can be held to its norms.

These messages make uniforms powerful tools for building trust among strangers. A passenger who sees a pilot in uniform does not need to know the pilot personally to feel some assurance. A patient who sees a nurse in uniform knows whom to ask for help. A lost child can be told to look for a police officer. In each case, the uniform turns a stranger into a legible category.

Uniforms also shape how people respond to authority. In a study published in 1974, the psychologist Leonard Bickman had experimenters approach pedestrians in New York and make requests, such as asking them to pick up a paper bag or give a coin to a stranger. The experimenters were dressed in ordinary civilian clothing, in a milkman's outfit or in a uniform resembling that of a security guard. People were more likely to comply when the requester wore the guard's uniform, even when the request had nothing to do with security. The uniform carried an aura of authority beyond its actual scope.

This effect explains both the value and the danger of uniforms. They allow organisations to project authority and build trust efficiently. They also allow people to exploit that authority. Impostors who wear uniforms, as Chapter 16 showed, can gain access and compliance they would not otherwise obtain. And legitimate wearers may be granted more deference than their role warrants.

Colour as information

In large organisations, especially hospitals, uniforms often carry fine-grained information about role. Patients in a busy ward may encounter many staff in a single day: doctors, nurses, healthcare assistants, physiotherapists, pharmacists, porters and cleaners. Distinguishing among them matters, because each can help with different things. Some health systems have introduced uniforms in which colour indicates role. NHS Scotland, for example, began rolling out a national uniform in December 2009. Before it, according to *Nursing Times*, hospitals in Scotland had been using more than 250 styles of tunic in about a hundred colours. The new scheme used shades of blue for clinical staff, with a particular blue reserved

for registered nurses and navy for the senior charge nurse running a ward, and green for non-clinical staff such as porters and catering teams. Official guidance stated the purpose plainly: to make it easier for patients, the public and other staff to tell the roles apart.

Colour coding makes uniforms more informative, but it also makes hierarchy more visible. When roles are colour coded, everyone can see who holds which position, and the distinctions may reinforce status differences. In some hospitals, staff have reported that colour coding affects how they are treated by colleagues and patients. A marker designed to inform can also stratify.

Doctors themselves have long debated what to wear. The white coat, once the standard marker of the physician, has been criticised as a possible carrier of infection, and some health systems have discouraged it in favour of short sleeves. Yet surveys of patients have often found that many prefer doctors in formal attire with a white coat, associating it with professionalism and competence. A 2018 study in *BMJ Open* led by Christopher Petrilli surveyed patients in the United States and found a general preference for formal attire with a white coat, though preferences varied by setting and region. The marker persists because patients read it as a signal of trustworthiness, whatever its practical drawbacks.

The power of an outfit

Uniforms may also affect the people who wear them. In 2012 Hajo Adam and Adam Galinsky published a study proposing what they called enclothed cognition. Participants who wore a lab coat described as a doctor's coat performed better on attention tasks than those who wore the same coat described as a painter's coat. The authors suggested that clothing affects cognition through its symbolic meaning. The study attracted wide attention. In 2019, however, a preregistered direct replication by Daniel Burns and colleagues, using a much larger sample, found no effect of wearing a lab coat on the attention task used in the original experiment, and Adam and Galinsky acknowledged that the result cast doubt on that specific finding while defending the broader idea. A later meta-analysis

raised concerns about the replicability of the early studies in this area. The safest reading is that clothing may influence how people think and feel, but the dramatic early laboratory results should not be relied on.

The broader idea that uniforms help people adopt a role is consistent with everyday experience. Many workers describe putting on a uniform as a way of stepping into their professional selves, leaving personal concerns behind. Actors speak of finding a character through costume. Whether or not clothing changes cognition in a measurable way, it clearly helps people signal to themselves, as well as to others, which group they belong to at a given moment.

Personal uniforms

Some leaders adopt personal uniforms, wearing essentially the same outfit every day. Steve Jobs is the most famous example. His biographer Walter Isaacson recounted that Jobs had been impressed by the uniforms worn by employees at Sony in Japan and had wanted to introduce something similar at Apple, an idea that was not well received by Apple employees. Jobs instead adopted a personal uniform of black mock turtlenecks, made for him by the designer Issey Miyake, together with jeans and sneakers.

A personal uniform is a marker of a different kind. It signals not membership in an organisation but a personal identity, often associated with focus, simplicity or indifference to fashion. When a leader's personal uniform becomes famous, it can turn into a marker that others imitate. The Theranos case described in Chapter 16 showed how a borrowed uniform can be used to signal affiliation with a respected archetype.

The end of the white shirt

Dress codes are less formal than uniforms but perform a similar function. They define what members of an organisation should look like, and they signal the organisation's culture to outsiders.

For decades, IBM was famous for a dress code that, in popular memory, centred on dark suits and white shirts. The look became so closely associated with the company that it was part of its identity. In the mid-1990s, under chief executive Lou Gerstner, IBM relaxed its dress code as part of a broader effort to change the company's culture. The change was symbolic. It signalled that IBM wanted to be more flexible, more customer-focused and less bound by tradition.

Changing a dress code can be a powerful signal precisely because dress is so visible. Employees and customers notice immediately. But a change in dress code without a change in substance can be read as superficial, or even as a distraction. Employees may wonder whether the new casual style reflects a real change in how decisions are made or merely a change in appearance.

When dress codes exclude

Dress codes can also exclude. In 2015 the United States Supreme Court decided a case brought by the Equal Employment Opportunity Commission against Abercrombie & Fitch. A young woman named Samantha Elauf had applied for a job at one of the retailer's stores while wearing a headscarf for religious reasons. She was not hired, and the company's appearance policy, which it called its Look Policy, was at issue. The Court ruled in her favour, holding that an employer may not refuse to hire an applicant to avoid accommodating a religious practice, even if the applicant has not explicitly requested an accommodation.

The case illustrates how dress codes can function as boundaries that keep out people whose markers differ. A dress code designed to project a particular brand image may exclude people whose religious, cultural or physical characteristics do not fit that image. Such exclusions may be unlawful, and even when they are lawful, they narrow the pool of people the organisation can draw on and signal to outsiders who is welcome.

Organisations with diverse workforces have increasingly adapted their dress codes to accommodate religious and cultural markers, such as headscarves, turbans and beards. Some have found that these accommoda-

tions strengthen rather than weaken organisational identity, because they signal that the organisation's markers of membership are compatible with members' other identities.

Uniforms as levellers

Uniforms can reduce status distinctions as well as reinforce them. Some companies have used uniforms for all employees, including executives, as a way of signalling equality. Honda has long been noted for a culture in which staff, including managers, wear the same work uniform, an approach associated with the founder Soichiro Honda's emphasis on the factory floor. Whether or not such uniforms change behaviour, they send a message that rank is less important than shared membership.

Schools have used uniforms for similar reasons, arguing that they reduce visible differences in wealth and create a shared identity. The evidence on whether school uniforms improve behaviour or achievement is mixed, and debates about them often reflect broader disagreements about conformity and individuality.

The levelling effect of uniforms illustrates a general point about markers. A marker that everyone shares strengthens the boundary between members and outsiders while weakening distinctions inside the group. A marker that varies by role strengthens distinctions inside the group while still marking the boundary. Designers must decide which distinctions they want to emphasise.

Dress in a remote world

The spread of remote work has changed the role of dress as a marker. When colleagues meet mainly through video calls, they see only each other's upper bodies, and often their backgrounds. Dress codes have become more relaxed in many organisations, and new markers have emerged, such as the choice of video background, the presence of books or plants on a shelf and the visibility of family life.

These new markers carry information about status and identity too. A carefully curated background can signal professionalism or cultural capital. A virtual background can signal a desire for privacy or a lack of suitable space. The absence of video can signal disengagement, or simply a poor connection. As with dress in physical offices, people read these markers quickly and often unconsciously.

Principles for uniforms and dress codes

The cases in this chapter suggest several principles for leaders designing uniforms and dress codes.

Be clear about what the marker should signal. Uniforms and dress codes can signal role, authority, equality, brand identity or professionalism. Each goal may require a different design.

Consider the boundary and the interior. A shared uniform strengthens the boundary of the group. Role-based distinctions strengthen internal hierarchy. Decide which matters more.

Accommodate other identities. Members bring markers from their religions, cultures and communities. Dress codes that accommodate those markers signal inclusion and widen the pool of people who can belong.

Connect changes to substance. A change in dress code can signal a change in culture, but only if the culture actually changes. Otherwise it will be read as cosmetic.

Watch for misuse. Uniforms carry authority, and authority can be borrowed by impostors or abused by wearers. Systems for verifying identity should not rely on uniforms alone.

The next chapter turns from what people wear to how they speak: the dialects, jargon and digital rituals that mark membership in the modern office.

CHAPTER 18

The Dialect of the Office

Every organisation develops a dialect. New employees notice it within days. There are acronyms that everyone uses and no one explains. There are phrases that carry special meaning, such as "disagree and commit," "bias for action" or "two-way door." There are ways of writing emails, structuring documents and running meetings that are expected but rarely written down. There are reactions in messaging channels, custom emoji and in-jokes that mean nothing to outsiders. Within a few months, most newcomers find themselves speaking the dialect without thinking about it. At that point, they have become members.

Language is among the most powerful markers of membership, as Chapter 3 argued. Accent, vocabulary and idiom tell listeners where a speaker comes from and whether they belong. Organisational dialects work the same way. They signal who is an insider, who is a newcomer and who is an outsider. They also shape how people think, because a shared vocabulary makes some ideas easy to express and others hard.

Measuring enculturation

Researchers have begun to measure organisational dialects directly. In a study published in *Management Science* in 2018, Sameer Srivastava, Amir Goldberg and colleagues analysed several years of internal email from a mid-sized technology firm. They developed a measure of how closely each employee's language matched the linguistic patterns of their colleagues, which they treated as an indicator of cultural fit. They found that new employees typically adapted their language to the organisation's norms over

their first months. Those who adapted quickly tended to do well. Those who failed to adapt were more likely to leave involuntarily. Interestingly, the researchers also found that some employees who had initially adapted later began to diverge from the group's language, and that this divergence predicted voluntary departure.

A related study by Cristian Danescu-Niculescu-Mizil and colleagues, presented in 2013, examined online communities devoted to reviewing beer. They found that members adopted the community's linguistic norms as they joined and that the community's language changed over time. Members who stopped adapting to the evolving language of the community, typically after a period of active participation, became more likely to leave. The authors titled their paper "No Country for Old Members," capturing the finding that long-standing members could become outsiders simply by failing to keep up with the dialect.

These studies suggest that language is not just a reflection of membership but a mechanism of it. People who speak the dialect are treated as members and come to feel like members. People who do not are treated as outsiders and come to feel like outsiders. The dialect is a marker that both signals and shapes belonging.

Codified dialects

Some organisations deliberately codify their dialects. Amazon's Leadership Principles, a set of short statements such as "customer obsession," "ownership" and "bias for action," are used in hiring, performance reviews and everyday conversation. Employees cite them to justify decisions, and interviewers probe candidates for evidence of them. The principles function as a vocabulary that lets strangers across a very large company communicate about values and behaviour in shared terms.

Amazon also became known for a particular way of writing. In a widely reported 2004 email, Jeff Bezos told senior executives that presentations using slides would be replaced by narrative memos, and meetings at the company came to begin with participants reading a memo of several pages

in silence. The practice is itself a marker. Employees who have mastered the memo format signal their membership, and newcomers must learn it to be taken seriously.

Netflix codified its culture in a presentation first published online in 2009, which described values such as "freedom and responsibility" and practices such as the "keeper test," in which managers ask whether they would fight to keep a given employee. The presentation was widely circulated and imitated. Its vocabulary became a marker not only inside Netflix but in the wider technology industry, where people who used its terms signalled familiarity with a particular philosophy of management.

Codified dialects have advantages. They make implicit norms explicit, which helps newcomers learn them. They provide a shared language for discussing difficult issues, such as performance and disagreement. And they create a sense of identity that can hold a large organisation together.

They also have risks. A codified dialect can become a set of slogans, repeated without understanding. People may learn to use the vocabulary to justify whatever they want to do, a form of rule beating. And a dialect that is closely identified with a particular organisation may make it harder for people from other backgrounds to join, since they must learn not only the work but the language.

Jargon as a barrier

Every profession and organisation develops jargon, specialised terms that allow experts to communicate efficiently. Jargon is useful. A surgeon who says "anastomosis" conveys in one word what would otherwise take a sentence. A software engineer who refers to a "race condition" communicates a specific kind of problem to colleagues.

But jargon also functions as a marker, and it can become a barrier. People who do not know the jargon are excluded from conversations, and they may feel that they do not belong. Organisations that rely heavily on internal acronyms can be bewildering to newcomers, partners and

customers. In mergers, the jargon of each side can become a source of friction, as employees struggle to understand one another and each side regards the other's terms as pretentious or obscure.

There is also a status dimension. Fluency in jargon signals expertise and seniority. People sometimes use jargon to display membership rather than to communicate, and listeners may hesitate to ask for clarification for fear of revealing ignorance. The result can be conversations in which many participants nod along without fully understanding.

Leaders can reduce these barriers by encouraging plain language, maintaining glossaries of internal terms for newcomers and modelling the willingness to ask what an acronym means. They can also pay attention to which parts of the organisation have developed dialects that others cannot follow, since those dialects may indicate subcultures that are drifting apart.

Digital rituals

Much of modern office communication happens through messaging platforms such as Slack and Microsoft Teams. These platforms have created new forms of dialect. Teams develop conventions about which channels to use for which purposes, how quickly to respond, when to use threads and how formal to be. They create custom emoji, sometimes hundreds of them, that carry meanings specific to the team. They use reactions as a form of lightweight communication: a thumbs-up to acknowledge, an eyes emoji to indicate that someone is looking into a problem, a custom emoji to signal agreement or amusement.

These digital rituals are markers. A newcomer who does not know them may respond too formally or too casually, use the wrong channel or miss the significance of a reaction. Long-standing members can identify newcomers by their messaging habits just as easily as Major Hellstrom identified Archie Hicox by his fingers. And teams with different conventions may struggle to communicate, as each reads the other's messages through its own dialect.

Custom emoji are especially interesting because they are created from below. Anyone on a team can add one, and the ones that catch on are those that capture something the team wants to express. In Beinhocker's terms, custom emoji evolve through differentiation, selection and amplification. A team's emoji collection is a record of its culture: its jokes, its concerns, its heroes and its shared experiences. Outsiders looking at the collection may find it baffling. Insiders find it a source of belonging.

When dialects divide

Organisational dialects can divide as well as unite. Different departments often develop different dialects. Engineers and salespeople, finance and marketing, head office and field operations may use different vocabularies, value different kinds of communication and read each other's messages with suspicion. These differences reflect real differences in work and expertise, but they can harden into boundaries that make cooperation difficult.

The dialect of senior leadership can be particularly distant from that of frontline staff. Strategy documents full of abstract terms such as "transformation," "alignment" and "value creation" may mean little to people whose work involves specific customers, products or machines. When leaders speak a dialect that frontline staff do not share, they signal that they belong to a different group, and frontline staff may conclude that leaders do not understand their work.

Mergers bring dialects into direct collision. Each side has its own vocabulary, acronyms and communication habits. In the early stages of integration, employees may struggle to understand one another, and each side may interpret the other's language as a sign of different values. Chapter 20 explores this dynamic in detail.

Learning the dialect

Because dialect is such an important marker, the speed at which newcomers learn it affects how quickly they become effective members. Organisations can accelerate this learning in several ways.

They can make the dialect explicit. Glossaries, style guides and documented norms help newcomers decode the language and conventions they encounter.

They can assign guides. Mentors or buddies who can explain the unwritten rules of communication help newcomers avoid embarrassing mistakes and learn faster.

They can create spaces for questions. Channels or forums where newcomers can ask what terms mean, without fear of looking foolish, reduce the barrier of jargon.

They can be tolerant of mistakes. Newcomers will use the dialect imperfectly at first. Treating their mistakes with patience rather than ridicule signals that they are welcome.

And they can allow the dialect to evolve. Newcomers bring their own vocabulary and habits, some of which may be useful. A dialect that absorbs contributions from newcomers stays alive. One that demands total conformity risks becoming stale.

Dialect as leverage

Donella Meadows argued that changing the paradigm, the set of shared assumptions from which a system arises, is among the most powerful ways to change a system. Language is one of the main carriers of paradigm. The words an organisation uses shape what its members notice, value and discuss. An organisation that talks constantly about customers will tend to think about customers. One that talks constantly about internal metrics will tend to think about metrics.

This makes dialect a potential lever for change. Leaders who want to shift an organisation's focus can introduce new vocabulary, retire old terms and model new ways of speaking. But dialect cannot be changed by decree. People adopt new language when it helps them express something they care about, and they reject it when it feels imposed or empty. Leaders who announce a new vocabulary without changing the underlying realities will find that the new words are used ironically or not at all.

The next chapter looks at what happens to markers when the office itself disappears, and colleagues who rarely or never meet in person must build trust through screens.

CHAPTER 19

Markers at a Distance

In the 1970s Thomas Allen, a professor at the Massachusetts Institute of Technology, studied how engineers in research and development organisations communicated. He found that the frequency of communication between colleagues fell off sharply with the distance between their desks. Engineers who sat close together talked often. Those who sat more than a few dozen metres apart talked far less, and those on different floors or in different buildings might hardly talk at all. The relationship became known as the Allen curve, and it was reported in his 1977 book *Managing the Flow of Technology*.

Allen's finding captured something that office workers had always known intuitively. Proximity creates contact, and contact creates familiarity. People who see each other regularly, in corridors, kitchens and meetings, come to know each other. They pick up one another's markers: dress, vocabulary, habits, moods. They form the kinds of personal relationships that sustain individual-recognition societies inside larger organisations.

When the COVID-19 pandemic forced much of the world's office work to move online in 2020, this system of contact was abruptly disrupted. Colleagues who had seen each other daily now met through screens, if at all. New employees joined organisations without ever meeting their colleagues in person. The markers that had been absorbed through everyday contact had to be transmitted in new ways, or not at all. Years later, many organisations are still working out how to build trust among people who rarely share a room.

What the evidence shows

The shift to remote work produced a natural experiment on a vast scale, and researchers have studied it intensively. Some of the most detailed evidence comes from studies of large firms with access to communication data.

A study published in *Nature Human Behaviour* in 2022 by Longqi Yang and colleagues at Microsoft analysed the communication patterns of more than sixty thousand Microsoft employees in the United States before and after the company moved to firm-wide remote work. The researchers found that remote work caused collaboration networks to become more static and siloed. Employees communicated less with people outside their immediate groups, and the bridging ties that connect different parts of an organisation weakened. Communication also shifted toward asynchronous channels such as email and instant messaging and away from synchronous ones such as meetings and calls.

These findings fit the framework of this book. Within teams, people who already knew each other maintained their relationships. Across teams, where relationships depended on chance encounters and shared spaces, contact declined. In Moffett's terms, the organisation's individual-recognition clusters held together, but the anonymous society linking them weakened. Markers that had been transmitted through informal contact across the organisation stopped flowing.

Other research has found more positive effects. In a randomised experiment at the Chinese travel company Ctrip, reported by Nicholas Bloom and colleagues in 2015, call-centre employees who worked from home were more productive than those in the office. A later randomised trial at the same company, published in *Nature* in 2024 by Bloom, Ruobing Han and James Liang, studied hybrid work for graduate employees in engineering, marketing and finance. Employees who worked from home two days a week had similar performance reviews and promotion rates to those in the office full time, and their quit rates fell by about a third.

The evidence suggests that remote and hybrid arrangements do not automatically destroy productivity or trust. They change the structure of the system. Work that depends on established relationships and clear tasks can thrive at a distance. Work that depends on weak ties, chance encounters and the transmission of markers across groups may suffer unless organisations find other ways to support it.

Swift trust

How do strangers who work together at a distance come to trust each other? One useful concept comes from research on temporary groups. In 1996 Debra Meyerson, Karl Weick and Roderick Kramer described what they called swift trust, a form of trust that arises quickly among members of groups that come together for a limited time to accomplish a task, such as film crews, emergency response teams or project teams. Members of such groups do not have time to build trust through long personal acquaintance. Instead, they rely on categories and roles. They trust each other because each occupies a role with known expectations and because each is presumed to have been selected for competence.

Swift trust is, in effect, trust based on markers. Members read each other's roles, credentials and affiliations and extend trust on that basis. The trust is fragile, because it is not grounded in personal experience, but it allows the group to function while personal experience accumulates.

Research on virtual teams has found similar patterns. In a study of global virtual teams published in 1998, Sirkka Jarvenpaa and Dorothy Leidner found that teams whose members communicated actively and enthusiastically early on, and who responded predictably to one another's messages, developed higher levels of trust. Early interactions set the tone. Teams that began with slow, sporadic or uncertain communication struggled to build trust later.

For remote organisations, these findings suggest that markers of role and competence become especially important when personal contact is scarce. Clear role definitions, visible credentials and predictable communication patterns help strangers trust each other quickly. At the

same time, the fragility of swift trust means that early failures can be costly. A missed deadline or an unanswered message in the first weeks of a project can drain a stock of trust that has not had time to accumulate.

The newcomer problem

The group most affected by the loss of informal contact is newcomers. Employees who joined an organisation before the shift to remote work had already absorbed many of its markers through everyday contact. They knew how things were done, whom to ask and what the unwritten rules were. Employees who joined remotely had to learn these things without the benefit of overheard conversations, corridor chats and the chance to watch experienced colleagues at work.

Many organisations found that remote onboarding required far more deliberate effort than in-person onboarding. Markers that had been transmitted automatically now had to be made explicit. Some companies responded by documenting their norms in detail. GitLab, which has operated as an all-remote company, maintains a public handbook running to thousands of pages, describing everything from its values to its communication practices. The handbook is, among other things, a written record of the company's markers, designed to let newcomers learn them without needing to absorb them through contact.

Documentation alone does not fully replace contact. People learn many markers through imitation and feedback rather than instruction. A newcomer can read that the company values candour, but learning what candour looks like in practice, how direct to be, in which settings and with whom, requires observing and participating. Remote organisations that succeed at onboarding usually combine documentation with structured contact: mentors, regular check-ins, cohort programmes and occasional in-person gatherings.

New markers for a remote world

Remote work has created its own markers. Some are inherited from office life and adapted. Others are new.

Responsiveness has become a powerful marker. In an office, colleagues can see that you are working. Remotely, they infer it from how quickly you respond to messages. This can create pressure to respond instantly, which may signal commitment but can also prevent focused work. Organisations that value deep work may need to establish norms that decouple responsiveness from trust, for example by agreeing on expected response times for different channels.

Visibility has become another. Employees who speak up in video meetings, post in public channels and share their work openly are more visible than those who work quietly. Visibility can be mistaken for contribution. Remote organisations risk rewarding those who are good at signalling presence rather than those who produce the most value.

Video presence carries its own markers, as Chapter 17 noted. Backgrounds, camera use and the visibility of home life all signal something about a person's circumstances and identity. These markers can create new forms of inclusion and exclusion. Employees with comfortable home offices may appear more professional than those working from cramped or shared spaces.

Written communication has become more important. In remote organisations, much of the work of coordination happens in writing. The ability to write clearly and in the organisation's preferred style becomes a marker of competence and membership. Employees whose strengths lie in verbal communication or who write in a second language may be disadvantaged.

Rituals at a distance

Organisations have experimented with many ways of recreating the rituals that build shared identity. Virtual coffee chats pair colleagues at random for informal conversation. Online social events, from quizzes to cooking

classes, try to replicate after-work gatherings. Regular all-hands meetings bring the whole organisation together on screen. Some companies send physical items, such as branded clothing or care packages, to remote employees as tangible markers of membership.

The evidence on these rituals is mixed, and many employees report fatigue with online social events. What seems to matter is not the format but whether the ritual creates genuine contact and shared experience. A virtual event that feels forced may do little. A recurring small-group conversation in which people get to know each other over time may do much more.

Many remote and hybrid organisations have concluded that occasional in-person gatherings are worth their cost. Bringing people together for a few days a year, or for key moments such as the start of a project or the integration of a new team, allows them to absorb markers and form relationships that sustain trust during the long stretches of remote work. In systems terms, these gatherings are periodic inflows that refill the stock of trust.

Two classes of colleague

Hybrid arrangements, in which some people work in the office some of the time, create a risk that remote-only arrangements do not. They can produce two classes of colleague. Those who are present in the office absorb its markers, form relationships with leaders and benefit from chance encounters. Those who are absent, through choice, caring responsibilities, disability or distance, may find that they are less visible and less trusted, even when their work is equally good.

Managers often describe this as proximity bias, the tendency to favour people one sees more often. In marker terms, it reflects the fact that physical presence remains a powerful marker of membership. Being in the room signals commitment and makes a person familiar, and familiarity breeds trust. Remote colleagues must work harder to display markers that substitute for presence, and they may still fall short.

Leaders can reduce the risk by designing practices that do not depend on presence. Meetings can be run so that remote participants have equal opportunity to speak, for example by having everyone join from their own screen when some people are remote. Decisions can be documented in shared channels rather than made in corridors. Performance can be assessed on outcomes rather than visibility. Opportunities for high-profile work can be allocated deliberately rather than by who happens to be nearby. These measures do not eliminate the advantage of presence, but they prevent it from hardening into a boundary between real members and the rest.

Designing for distance

Remote and hybrid work change the structure of an organisation's marker system. They weaken the informal channels through which markers flow across groups and increase reliance on explicit markers such as roles, credentials and written norms. Several principles follow for leaders.

Make markers explicit. What was once absorbed through contact must now be documented, taught and modelled deliberately.

Invest in bridging ties. Because remote work tends to strengthen within-team relationships and weaken cross-team ones, leaders should create structured opportunities for people in different groups to meet and work together.

Front-load trust building. Swift trust depends on early interactions. Projects and teams should begin with clear roles, active communication and, where possible, some form of personal contact.

Prioritise newcomers. People who join remotely need more support to learn the organisation's markers than those who join in person. Mentors, cohorts and structured onboarding make a significant difference.

Watch for new inequalities. Remote markers such as responsiveness, visibility and video presence can create new forms of exclusion. Leaders should check that they are rewarding contribution rather than the display of presence.

Use in-person time strategically. When people do gather, the time should be spent on the things that are hardest to do remotely: building relationships, absorbing culture and resolving difficult questions.

The next chapter turns to one of the most dramatic tests of any marker system: the merger, in which two organisations with different markers are asked to become one.

CHAPTER 20

When Two Tribes Sign a Deal

In May 1998 Daimler-Benz, the German maker of Mercedes-Benz cars, and Chrysler, the third-largest American carmaker, announced that they would combine. The deal valued Chrysler at around thirty-six billion dollars, and it was presented to the world as a merger of equals. The new company, DaimlerChrysler, would have headquarters in both Stuttgart and Auburn Hills, Michigan. Its leaders spoke of complementary strengths: German engineering and American marketing, luxury cars and mass-market vehicles, two great industrial traditions joined in one global champion.

Less than a decade later, in 2007, Daimler agreed to sell a controlling stake in Chrysler to the private equity firm Cerberus Capital Management in a deal valued at about seven point four billion dollars, a fraction of the original price, with much of the money staying in Chrysler rather than going to Daimler. The merger of equals had become one of the most frequently cited examples of a failed combination in business history.

Many factors contributed to the failure, including strategic misjudgments and changes in the car market. But accounts of the merger by journalists, scholars and participants return again and again to a set of problems that this book would describe as marker collisions. Two organisations with different languages, rituals, status systems and assumptions were asked to become one. Neither side's markers were respected by the other, and the promised union never took hold.

Merger of equals

The phrase "merger of equals" is itself a marker. It signals to employees, investors and the public that neither side is being conquered, that both cultures will be respected and that the new organisation will draw on the best of each. It is a reassurance offered to people who are anxious about their status.

At DaimlerChrysler, that reassurance was undermined publicly. In 2000 Jürgen Schrempp, Daimler's chief executive, was quoted in an interview with the *Financial Times* as indicating that the merger of equals framing had been adopted for psychological reasons and that Chrysler had always been intended to become a division of Daimler. The remarks were widely reported, and the investor Kirk Kerkorian, a large Chrysler shareholder, later sued, arguing that shareholders had been misled. The suit was ultimately unsuccessful, but the damage to trust was done. Chrysler employees who had been told they were partners now heard that they had been acquired.

In marker terms, Schrempp's remarks drained the stock of trust that the merger-of-equals framing had created. The marker had promised parity. The statement revealed that it had been, at least in part, a device. Employees who had invested in the marker felt deceived, and their subsequent behaviour reflected that feeling.

Two sets of markers

Beyond the question of control, the two companies carried very different markers. Accounts of the merger have described differences in management style, with Daimler often characterised as more hierarchical and formal and Chrysler as more informal and faster-moving. There were differences in pay. Senior American executives were reported to earn considerably more than their German counterparts, a disparity that created

resentment on both sides once it became visible. There were differences in decision making, in attitudes to cost and engineering, and in everyday rituals such as how meetings were run and how documents were prepared.

None of these differences was necessarily a problem in itself. Organisations with different cultures can work together productively. The problem was that the differences functioned as markers of membership in rival groups. Each side read the other's habits as evidence of different values. German managers could see American informality as a lack of discipline. American managers could see German formality as rigidity. Each side's markers signalled to the other that they were not real members of the same company.

Over time, many senior Chrysler executives left, and the company was increasingly run from Stuttgart. The anonymous society that the merger was supposed to create never formed. Instead, the combined company remained two societies sharing a name, with the larger one gradually absorbing the smaller.

AOL and Time Warner

The merger of AOL and Time Warner, announced in January 2000, was even larger. AOL, then the dominant internet service provider in the United States, agreed to acquire Time Warner, owner of magazines, a film studio, cable networks and cable systems, in a deal valued at well over one hundred billion dollars. The combination was hailed as the marriage of old media and new, uniting Time Warner's content with AOL's online audience.

The combined company struggled almost immediately. The dot-com bubble burst, AOL's advertising revenue fell and accounting problems emerged at the AOL unit. In 2002 AOL Time Warner reported a loss of about ninety-nine billion dollars, at the time one of the largest annual losses in corporate history, mainly reflecting a write-down of goodwill from the merger. In 2003 the company dropped AOL from its name, and in 2009 Time Warner spun off AOL as a separate company.

As at DaimlerChrysler, the failure had many causes, but culture featured prominently in most accounts. AOL was a young, fast-moving company with an aggressive, deal-driven culture. Time Warner was a collection of established businesses, including journalistic and creative divisions with long traditions and strong identities of their own. Employees of Time Warner's divisions often saw AOL's leaders as salespeople who did not understand or respect their work. AOL's leaders often saw Time Warner as slow and bureaucratic. The divisions themselves, such as HBO, Warner Bros. and the magazine group, had their own markers and resisted integration not only with AOL but with each other.

The AOL Time Warner case illustrates a point that is easy to miss. A merger is not only a collision between two cultures. It is a collision between multiple cultures, since each organisation contains its own subcultures with their own markers. Integrating the parent companies at the top may do little to integrate the divisions below, and each division may defend its markers against all comers.

Disney and Pixar

Not every merger fails, and comparing successes with failures helps clarify what makes the difference. In 2006 Disney acquired Pixar Animation Studios for about seven point four billion dollars. Pixar had made a string of successful animated films, including *Toy Story*, while Disney's own animation studio had struggled. The deal raised an obvious risk. Pixar had a distinctive culture, built around creative autonomy and a particular way of developing stories. Many observers feared that Disney would absorb Pixar and destroy what made it special.

Instead, Disney took a different approach. Ed Catmull and John Lasseter, Pixar's leaders, were given responsibility for Disney's own animation studio as well as Pixar's. Pixar kept its name, its campus in Emeryville, California, and its practices. In his 2014 book *Creativity, Inc.*, Catmull described how the two companies agreed on a set of commitments intended to protect Pixar's identity after the acquisition, extending, by his

account, to small details of everyday identity. Rather than merging the two studios into one, Disney preserved their separate identities while connecting them at the top.

In marker terms, Disney chose to respect Pixar's markers rather than replace them. Pixar employees continued to see themselves as members of Pixar, with its own name, rituals and ways of working. The acquisition changed who owned the company but did not ask its people to become Disney people. At the same time, the flow of leadership and practices from Pixar into Disney Animation allowed some of Pixar's markers to spread to the larger organisation.

Integration strategies

Scholars of mergers and acquisitions have long distinguished among different approaches to integration. In their 1991 book *Managing Acquisitions*, Philippe Haspeslagh and David Jemison described approaches including absorption, in which the acquired company is fully integrated into the acquirer; preservation, in which it is kept largely separate; and symbiosis, in which the two gradually become interdependent while retaining distinct identities. Each approach suits different circumstances, depending on how much the companies need to share resources and how important it is to protect the acquired company's autonomy.

From the perspective of this book, these approaches are strategies for managing markers. Absorption replaces the acquired company's markers with the acquirer's. Preservation keeps both sets of markers intact. Symbiosis creates new shared markers while allowing some distinct ones to persist. The DaimlerChrysler merger was presented as symbiosis but carried out, in practice, as something closer to absorption. Disney and Pixar was closer to preservation at first, with elements of symbiosis over time.

Problems arise when the strategy is unclear or when the announced strategy differs from the actual one. Employees read markers closely during a merger. They notice whose logo goes on the building, whose systems are adopted, whose executives get the top jobs, whose rituals survive and

whose vocabulary becomes standard. Each of these decisions sends a signal about who has won and who has lost. When the signals contradict the official story, employees believe the signals.

Why mergers so often disappoint

A figure often quoted in discussions of mergers comes from a 2011 article in the *Harvard Business Review* by Clayton Christensen and colleagues, who wrote that study after study had put the failure rate of mergers and acquisitions somewhere between seventy and ninety per cent. The figure is widely cited, but it should be treated with caution. Studies use different definitions of failure, different time horizons and different samples, and some research finds more favourable outcomes, particularly for smaller acquisitions and for acquirers with experience. What seems clear is that many large mergers fail to deliver the value promised when they are announced.

Culture is frequently cited as a reason, but culture is a vague term. This book's framework offers a more precise diagnosis. Mergers bring together two anonymous societies held together by different markers. Unless the markers are managed deliberately, the result is a version of what happens when two ant colonies are brought together: each side recognises the other as foreign, and cooperation gives way to suspicion. Talented people leave, because they no longer feel they belong. Those who stay may defend their old markers rather than adopt new ones. The anticipated synergies, which depend on cooperation, fail to materialise.

The first hundred days of markers

The early period after a deal closes is when employees read markers most intensely. They have been told what the merger is supposed to achieve, but they do not yet know what it will mean for them. Every visible decision is scrutinised for clues. Leaders who understand this can use the period deliberately.

Some decisions arrive almost immediately: what to call the combined company, whose logo appears on email signatures and building signs, whose systems people log in to on the first day, who sits on the new leadership team and who reports to whom. Each of these is a marker of who has won. If they all point in one direction, employees on the other side will conclude that they have been acquired, whatever the official language says. If they are balanced, employees may believe the language of partnership, but only if later decisions confirm it.

Other decisions are less visible but just as important. Which side's way of running meetings becomes the norm? Whose vocabulary appears in the integration plan? Which rituals are continued, merged or quietly dropped? Who is invited to the first offsite, and who is left out? Employees notice these things even when leaders do not.

A practical approach is to treat the first months of integration as a period of explicit marker management. Leaders can list the markers each side cares about most, decide which will be preserved, blended or replaced, and communicate those decisions with reasons. They can create early joint projects that give employees from both sides a shared experience and a chance to develop new common markers. And they can monitor how employees are reading the signals, through surveys, conversations and attention to who is leaving, so that they can correct course before suspicion hardens into the kind of rival identities that defeated so many earlier combinations.

Managing markers in a merger

The cases in this chapter suggest several principles for leaders planning a merger or acquisition.

Map the markers on both sides. Before the deal closes, identify the markers that matter to each organisation's members: names, logos, rituals, vocabulary, status systems, ways of making decisions. Understand which ones carry deep meaning and which are superficial.

Be honest about the strategy. If the acquired company will be absorbed, say so. Employees can adapt to an honest absorption more easily than to a merger of equals that turns out to be a conquest. Contradictions between official framing and actual decisions destroy trust.

Decide which markers to preserve. Some markers are central to the value of the acquired company, such as Pixar's creative culture. Preserving them may be essential to protecting that value. Others can be replaced without much loss.

Create shared markers deliberately. If the goal is a single organisation, invest in new markers that both sides can adopt, rather than simply imposing one side's markers on the other. Shared rituals, joint projects and a new vocabulary can help.

Watch the signals. Every decision about logos, systems, titles, offices and leadership appointments sends a message. Leaders should consider those messages as carefully as they consider the financial consequences.

Expect delays. Markers change slowly, as Chapter 7 argued. Integration takes years, not months, and leaders should plan accordingly.

The next chapter takes up a related question that every leader of change eventually confronts: why culture, which seems so soft and intangible, is so remarkably hard to redesign.

CHAPTER 21

Why Culture Resists Redesign

In February 2017 Susan Fowler, a former engineer at Uber, published a blog post describing her experiences of sexual harassment at the company and the indifference, as she described it, of the human resources department to her complaints. The post spread rapidly and triggered a wave of further revelations about Uber's workplace culture. The company commissioned an investigation led by the former US attorney general Eric Holder, whose firm produced a set of recommendations for reform. In June 2017 Travis Kalanick, Uber's co-founder and chief executive, resigned under pressure from investors.

Among the targets of criticism was a set of fourteen cultural values that Kalanick had introduced in Uber's early years. They included phrases such as "always be hustlin'," "meritocracy and toe-stepping" and "principled confrontation." To many observers, these values seemed to have licensed aggressive behaviour and to have signalled that results mattered more than how they were achieved. When Dara Khosrowshahi became chief executive later in 2017, he led a process to write new cultural norms. Among them was the statement "We do the right thing. Period."

The episode is a vivid example of a company trying to redesign its culture by changing its markers. It also raises the question at the heart of this chapter. Why is culture so hard to change? If a company's values can be rewritten in a few months, why do the behaviours they describe so often persist for years?

Culture as layered markers

The organisational psychologist Edgar Schein proposed an influential model of culture with three levels. At the surface are artefacts, the visible features of an organisation such as its offices, dress, rituals, language and published values. Below them are espoused values, the beliefs and principles that members say guide their behaviour. At the deepest level are basic underlying assumptions, the taken-for-granted beliefs about how the world works and how people should behave, which members may not even be aware of holding.

Schein's model maps neatly onto the framework of this book. Artefacts are markers in the most visible sense. Espoused values are markers expressed in language. Basic assumptions are the deep structure that markers signal and that shape how members read them. Changing artefacts and espoused values is relatively easy. A company can redesign its office, change its dress code and publish new values in a matter of months. Changing basic assumptions is much harder, because they are embedded in members' habits, expectations and identities.

When a company changes its surface markers without changing its basic assumptions, the result is often cynicism. Employees notice the new values posted on the walls and compare them with the behaviour they see from leaders and colleagues. If the behaviour has not changed, the new markers are read as hollow, and trust in leadership declines. In systems terms, the company has introduced a new information flow that contradicts the state of the stock, and people respond by discounting the information.

Policy resistance

Donella Meadows described a system trap she called policy resistance. It arises when various actors in a system have different goals and each pulls the system toward their own goal. When a policy maker tries to push the system in one direction, the other actors respond by pulling harder in their

own directions, and the system stays roughly where it was. Meadows used examples such as drug policy, in which efforts to reduce supply raise prices, which attracts more suppliers, which offsets the reduction.

Culture change often encounters policy resistance. A leader announces a new set of values and introduces programmes to promote them. Middle managers who were promoted under the old culture, and whose status depends on its markers, may comply outwardly while continuing to reward the old behaviours. Employees who have learned that the old behaviours lead to success may continue them. Customers and partners may expect the old behaviours and react badly when they change. Each actor pulls the system back toward its previous state.

Meadows suggested that the way out of policy resistance is to find a goal that all actors can share, or to let go of the policy and look for ways to align the actors' interests. For culture change, this means that announcing new values is not enough. Leaders need to change the incentives, structures and markers that reward the old behaviours, and to find ways of making the new culture serve the interests of the people who must adopt it.

Markers as stocks

Another reason culture resists change is that markers behave like stocks. Chapter 7 described trust as a stock that accumulates and drains through flows. The markers of a culture accumulate in the same way. Every story told about how things are done, every promotion that rewards a particular behaviour, every ritual repeated over years adds to the stock of shared expectations. These stocks are large, and they change only as fast as their flows allow.

A new chief executive who wants to change the culture is trying to drain one stock and fill another. The old stock of expectations will drain only as new experiences contradict it, and the new stock will fill only as new experiences confirm it. If the flow of new experiences is small, because only a few leaders behave differently or because the new behaviours are

visible only in some parts of the organisation, the stocks will change slowly. People will continue to act on the old expectations because those expectations are still what most of their experience supports.

This explains why culture change programmes often seem to succeed at first and then stall. The initial announcements and symbolic changes create a burst of new information. But unless that information is sustained by a continuing flow of experiences, the old stock reasserts itself.

Microsoft's reinvention

The transformation of Microsoft's culture after Satya Nadella became chief executive in 2014 is often cited as a successful example. Before Nadella, Microsoft had a reputation for internal competition, sometimes described in the press as a culture of warring divisions. A performance management system known as stack ranking, in which managers were required to rate employees against one another on a forced distribution, was widely criticised for encouraging employees to compete with colleagues rather than collaborate. Microsoft announced in late 2013, shortly before Nadella took over, that it would end the practice.

Nadella emphasised the idea of a growth mindset, drawn from the work of the psychologist Carol Dweck, and in his 2017 book *Hit Refresh* he described wanting Microsoft to move from a culture of know-it-alls to one of learn-it-alls. He also made strategic changes that altered the company's markers, such as embracing collaboration with former rivals and making Microsoft software available on competing platforms. Over the following years, Microsoft's market value rose dramatically, and many observers credited the cultural shift as part of the reason.

What made the change stick, according to most accounts, was that it was not only rhetorical. The end of stack ranking removed a structural incentive for internal competition. Strategic decisions demonstrated the new values in action. And the change was sustained over many years, providing a continuing flow of experiences that gradually rebuilt the stock of expectations. The markers changed, but so did the structures and behaviours they signalled.

It is worth adding a caution. Accounts of successful culture change tend to be written after the fact, by people who know the outcome. Microsoft benefited from many factors, including the growth of cloud computing, that had little to do with culture. Attributing its success primarily to a change in mindset risks the kind of storytelling that makes complex outcomes seem simpler than they are. The case is best read as an example of how culture change can be supported by structural and strategic change, rather than as proof that culture change alone transforms results.

When values become slogans

Google's founders included the phrase "Don't be evil" in the company's prospectus when it went public in 2004, and the phrase became one of the most famous corporate mottos in the world. Over time, it attracted criticism from people who argued that the company's behaviour did not always match it. In 2018 the phrase was moved from the preamble of Google's code of conduct to a closing line, a change that was widely reported as the company dropping its motto.

The episode illustrates a problem with explicit values. A memorable slogan can become a marker of identity, but it also becomes a standard against which the company is judged. When behaviour diverges from the slogan, critics use the slogan to highlight the gap. The company may be tempted to soften or remove the slogan, which is then read as an admission that the values have changed. Values that are stated boldly create expectations that are hard to meet consistently.

Uber's experience shows the opposite risk. Values that are too aggressive can license behaviour that later becomes indefensible. Kalanick's original values were memorable and clearly expressed, and they shaped behaviour. The problem was not that they were ignored but that they were followed.

The leverage of paradigm

Meadows placed paradigms, the shared assumptions from which a system's goals, rules and structures arise, near the top of her list of leverage points. Changing a paradigm can transform a system more profoundly than any change to its parameters or structures. But paradigms are also the hardest things to change, because they are embedded in everyone's way of seeing.

In organisational terms, a paradigm is close to what Schein called basic underlying assumptions. It is the deep belief about what the organisation is for, what success means and how people should treat one another. Culture change that does not touch the paradigm will be superficial. Culture change that does touch it may be transformative, but it requires sustained effort over a long period and often a crisis that makes the old paradigm untenable.

Meadows suggested that paradigms change when people keep pointing at the anomalies and failures of the old paradigm, speak loudly and with assurance from the new one, and place people with the new paradigm in positions of visibility and power. For organisations, this means that leaders who want to change culture need to make the costs of the old culture visible, articulate the new one consistently and promote people who embody it. Markers matter here, because they are how the new paradigm becomes visible. But the markers must be backed by real changes in who is rewarded and how decisions are made.

Principles for culture change

The analysis in this chapter suggests several principles for leaders who want to change culture.

Change structures as well as symbols. New values and rituals will be read as hollow unless incentives, promotion criteria and decision processes change to match them.

Sustain the flow. Culture is a stock that changes slowly. A burst of announcements will not change it. A sustained flow of consistent experiences over years will.

Expect resistance. People whose status depends on the old markers will resist the new ones. Find ways to align their interests with the change, or accept that some will leave.

Be careful with slogans. Bold values create expectations. Choose words that the organisation can live up to, and be prepared to be judged by them.

Look for the paradigm. Surface changes will not stick unless the underlying assumptions change. Identify those assumptions and address them directly.

Tell the story honestly. Culture change rarely produces results as quickly or as cleanly as success stories suggest. Leaders should set realistic expectations and resist the temptation to declare victory early.

Part IV now turns from organisations to the digital platforms that have pushed cooperation among strangers further than any institution in history, and to the markers that make it possible.

PART IV

Platforms and Digital Strangers

Digital platforms have pushed cooperation among strangers further than any society before them. Part IV examines the reputation systems, trust designs and commons that make this possible, and the counterfeit markers that threaten them.

CHAPTER 22

Reputation Engines

On Labor Day weekend in 1995, a software engineer named Pierre Omidyar launched a website called AuctionWeb, which would later become eBay. Within a few months, the site had several hundred members buying and selling items with one another. Almost none of them had met. A buyer in one state would send money to a seller in another and wait for a package to arrive. There was no guarantee that it would.

On 26 February 1996 Omidyar posted a letter to the site's users announcing a new feature, the Feedback Forum. "By creating an open market that encourages honest dealings," he wrote, "I hope to make it easier to conduct business with strangers over the net." He went on: "Most people are honest. And they mean well." But some, he acknowledged, were dishonest or deceptive. "But here, those people can't hide." Users would be able to rate each other after each transaction, leaving positive, negative or neutral feedback and a short comment. The ratings would be tallied and displayed next to each user's name.

Omidyar's letter is one of the founding documents of the platform economy. It states the problem of this book in a single sentence: how to make it easier to conduct business with strangers. And it proposes a solution that is, in Moffett's terms, a new kind of marker. The feedback score is a signal, attached to a username, that tells strangers how other strangers have experienced dealing with that person. It turns the gossip that sustained trust in small communities into a public, quantified and permanent record.

How reputation engines work

A reputation system on a platform performs several functions at once. It collects information about past behaviour from people who have interacted with a participant. It aggregates that information into a summary, such as a score, a star rating or a percentage of positive feedback. It displays the summary to people considering future interactions. And it creates an incentive for participants to behave well, since bad behaviour will be recorded and will affect future opportunities.

In systems terms, a reputation engine is an information flow that makes the stock of trust in a particular participant visible. It also creates a feedback loop. Good behaviour produces good ratings, which attract more transactions, which produce more opportunities for good behaviour and more ratings. Bad behaviour produces bad ratings, which deter transactions. The loop rewards trustworthiness and punishes untrustworthiness, at least in principle.

The value of this information is measurable. The experiment by Paul Resnick, Richard Zeckhauser, John Swanson and Kate Lockwood, published in 2006, had an established eBay dealer in vintage postcards sell matched lots under his long-standing, highly rated identity and under new identities with no feedback. Buyers paid about eight per cent more, on average, for the lots sold under the established identity. The reputation marker was worth real money. Interestingly, the researchers also found that one or two negative ratings for the new sellers did not significantly affect buyers' willingness to pay, suggesting that buyers read reputation scores with some nuance.

The problems of reputation

Reputation engines face all the traps described in Chapter 8, and some of their own.

Ratings inflate. On eBay, researchers have long observed that the overwhelming majority of feedback is positive. Studies by Resnick and Zeckhauser of early eBay data found that negative feedback was rare. Part of the reason was reciprocity. In the original system, buyers and sellers rated each other, and a buyer who left negative feedback risked receiving negative feedback in return. Many users therefore left positive feedback or none at all, even after mediocre experiences. In 2008 eBay changed its rules so that sellers could no longer leave negative or neutral feedback for buyers, aiming to remove the fear of retaliation.

Ratings concentrate. As Chapter 8 noted, participants with established reputations attract more business and accumulate more ratings, while newcomers struggle to get started. This is the success to the successful trap. Platforms that want to remain open to newcomers must find ways to give them a chance, such as guarantees, introductory badges or search algorithms that occasionally surface new participants.

Ratings can be faked. Because reputation is valuable, participants have incentives to manufacture it. Fake reviews are a persistent problem on many platforms. A 2016 study by Michael Luca and Georgios Zervas of the review site Yelp found that about sixteen per cent of restaurant reviews were filtered by the site's algorithm as suspicious, and that restaurants facing more competition or with weaker reputations were more likely to have suspicious reviews. In 2024 the US Federal Trade Commission finalised a rule banning the sale and purchase of fake reviews, reflecting the scale of the problem.

Ratings can discriminate. Reputation systems record the judgments of users, and users' judgments may reflect bias. If some users rate participants from certain groups more harshly, the ratings will carry that bias forward. The platform's marker may then encode the prejudices of its users, and it may do so in a form that appears objective.

Stack Overflow and the economy of points

A different kind of reputation engine emerged on Stack Overflow, the question-and-answer site for programmers founded in 2008 by Jeff Atwood and Joel Spolsky. Users ask and answer technical questions, and other users vote on the quality of the answers. Votes earn reputation points, and reputation unlocks privileges. Users with small amounts of reputation can vote and comment. Users with more can edit other people's posts, vote to close questions and access moderation tools. The site also awards badges for particular achievements, from answering a first question to sustaining participation over long periods.

Stack Overflow's design makes reputation a marker not only of trustworthiness but of status and authority. A user with a high reputation score is recognised as an expert and is trusted with powers that newcomers do not have. The system creates a hierarchy based on contribution, visible to everyone through the number displayed next to each username.

The design was remarkably successful for many years. Stack Overflow became one of the most widely used resources for programmers, and its answers were frequently the first results in searches for technical problems. The reputation system motivated experts to contribute their time and knowledge without payment, in exchange for points, badges and the status they conferred.

The system also attracted criticism. Newcomers often reported that the site felt hostile, with questions closed as duplicates or downvoted for not meeting community standards. Long-standing users with high reputation sometimes enforced norms in ways that felt exclusionary. The markers that sustained the community's standards also made it harder for strangers to join. This tension, between maintaining quality and welcoming newcomers, is common to many reputation-based communities, and Chapter 24 returns to it.

In the years after 2022, Stack Overflow faced a new challenge as programmers increasingly turned to AI chatbots for help with technical questions. Researchers who studied the site's activity reported a marked

decline in new questions after the release of ChatGPT in late 2022. The episode raised a question that reputation-based communities will increasingly face. When a machine can answer questions without needing to be trusted as a person, what happens to a marker system built on human contribution?

Reading reputation

Users of reputation systems learn to read them in sophisticated ways. Because ratings inflate, experienced users pay attention to small differences, such as the difference between 99.5 and 98 per cent positive feedback, which may signal a meaningful difference in reliability. They read the text of reviews rather than relying on the score. They look at the number of ratings as well as their average, since a high score based on few ratings is less informative than one based on many. They check how recent the ratings are, since reputation can change.

This sophistication is itself a form of marker literacy. Users who know how to read reputation systems are insiders in the platform's culture. Newcomers who take scores at face value may be misled. Platforms that want to be accessible to newcomers must either design their reputation systems to be easily interpretable or invest in teaching users how to read them.

Designing reputation engines

The experience of eBay, Stack Overflow and other platforms suggests several principles for designing reputation engines.

Protect against retaliation. When ratings are reciprocal, fear of retaliation suppresses honest negative feedback. Platforms can address this by hiding ratings until both parties have submitted them, by restricting who can rate whom or by allowing private feedback to the platform.

Keep scales meaningful. Ratings that cluster at the top lose their power to discriminate. Platforms can recalibrate scales, display relative rather than absolute scores or emphasise specific dimensions of experience rather than a single overall score.

Give newcomers a path. Reputation systems naturally favour established participants. Platforms can provide guarantees, introductory periods, verification badges or other markers that let newcomers establish trust before they have accumulated ratings.

Detect manipulation. Fake reviews and coordinated rating schemes undermine trust in the whole system. Platforms need continuing investment in detection, along with clear penalties for manipulation.

Watch for bias. Reputation systems can encode users' prejudices. Platforms should test whether ratings differ systematically by group in ways that do not reflect behaviour, and consider design changes that reduce the influence of bias.

Combine reputation with other markers. Ratings are one marker among many. Verification of identity, guarantees, insurance, dispute resolution and direct quality checks can complement reputation and make the system more robust.

Rating people, not products

There is an important difference between rating a product and rating a person. A poor rating for a toaster affects the toaster's sales. A poor rating for a driver, a cleaner, a freelance designer or a delivery worker may affect that person's livelihood. As reputation engines have spread from marketplaces for goods to marketplaces for labour, the stakes of their markers have risen sharply.

Workers on many platforms depend on ratings to receive work. A drop below a threshold may reduce the number of jobs offered or lead to removal from the platform. Because ratings inflate, small differences can carry large consequences. A worker whose average falls by a fraction of a star may find themselves at risk, even though most customers were

satisfied. And because ratings record customers' judgments, they may reflect customers' moods, misunderstandings or biases as well as the worker's performance.

These features place a special obligation on platforms that use ratings to make decisions about people. Procedural fairness, as Chapter 13 argued, matters a great deal. Workers should be able to see their ratings and the feedback behind them, understand how ratings are used, contest ratings they believe are unfair and appeal decisions to remove them from the platform. Platforms should check whether ratings differ systematically across groups of workers in ways that do not reflect behaviour. And they should consider whether a single average score is the right marker at all, or whether more specific feedback, combined with other measures of performance, would be fairer and more informative.

The broader point is that a reputation engine is a form of governance. When its markers determine who may work, it is making decisions that in other contexts would be subject to employment law, due process and public scrutiny. Designers who treat ratings as a neutral feature of the interface miss how much power the markers carry.

Gossip, scaled

Robin Dunbar argued that language evolved partly to let humans gossip, exchanging information about absent third parties so that they could track the reputations of more people than they could observe directly. Reputation engines are gossip scaled to millions. They take the information that would once have circulated in a village and make it available to anyone, anywhere, attached to the name of the person it concerns.

This scaling has made possible forms of cooperation among strangers that would have been unimaginable a few decades ago. It has also created new problems, as gossip has always done. Reputation can be unfair, manipulated or permanent in ways that village gossip was not. The next chapter looks at platforms that took on an even harder version of the trust problem: persuading strangers not only to trade with one another but to sleep in one another's homes and ride in one another's cars.

CHAPTER 23

A Stranger's Spare Room

In 2007 two young designers in San Francisco, Brian Chesky and Joe Gebbia, could not afford their rent. A large design conference was coming to the city, and hotels were full. They inflated air mattresses on the floor of their apartment and offered them to conference visitors, with breakfast included. They called the idea AirBed & Breakfast. With a third co-founder, Nathan Blecharczyk, they turned it into a company that launched as a website in 2008 and later became Airbnb.

The idea seemed absurd to many people at the time. Who would invite a stranger to sleep in their home? Who would sleep in a stranger's home? The obstacles were not technical. They were obstacles of trust. Hosts feared theft, damage and danger. Guests feared misrepresented listings, unsafe conditions and hosts who might behave badly. Persuading millions of strangers to overcome these fears required one of the most ambitious exercises in marker design in business history.

Designing for trust

Joe Gebbia described the challenge in a 2016 talk at the TED conference titled "How Airbnb designs for trust." He spoke of the bias that most people carry against strangers and the way design could help overcome it. He described research Airbnb had conducted with academics at Stanford University on how people decide whom to trust. In his account, people generally preferred hosts or guests who seemed similar to themselves, a pattern consistent with the research on markers discussed in Part I. But the

research also found that a strong reputation, built from a sufficient number of positive reviews, could outweigh the preference for similarity. In other words, a reputation marker could override a demographic marker.

Gebbia's account captures the core of Airbnb's approach. The company built a system of markers designed to make strangers legible to one another. Profiles included photos, names, descriptions and verified information. Reviews after each stay let hosts and guests rate each other. Messaging allowed them to communicate before committing. Payments were handled by the platform, so that neither side had to trust the other with money directly. Each of these features was a marker or a mechanism that let strangers judge one another without personal acquaintance.

Early on, the founders reportedly travelled to New York to photograph hosts' listings themselves, having noticed that poor photos made listings look untrustworthy. Better photographs were a marker of quality that helped guests feel confident. The company later offered professional photography to hosts more widely. It is a small detail, but it illustrates a general point. Markers of trust are often mundane. A clear, well-lit photograph tells a prospective guest that the host cares enough to present the space honestly.

When trust fails

Every marker system eventually fails somewhere, and how an organisation responds to failure shapes the trust that remains. In 2011 a host in San Francisco wrote a widely read blog post describing how guests had ransacked her apartment, stealing belongings and damaging property. The incident attracted extensive press coverage and raised questions about whether Airbnb could protect hosts.

Airbnb's initial response was criticised as slow. The company then introduced a guarantee, initially of fifty thousand dollars, to cover hosts for damage caused by guests, and later raised it to one million dollars. It also expanded its customer support and safety teams. The guarantee functioned as a meta-marker, similar in some respects to the deposit insurance

described in Chapter 10. Hosts no longer needed to rely only on guests' reviews and profiles. They could rely on the platform's promise to make them whole if something went wrong.

Markers that discriminate

Airbnb's marker system also revealed a darker side. As Chapter 1 and Chapter 4 explained, markers that let people recognise members of their own group can also lead them to treat others with suspicion. On Airbnb, profiles included names and photos, which let hosts see guests' apparent ethnicity before accepting a booking.

The field experiment by Benjamin Edelman, Michael Luca and Dan Svirsky, published in 2017, tested the consequences. The researchers sent about 6,400 requests to hosts in five American cities from guest accounts that were identical except for names, chosen to be distinctively African American or distinctively white. Requests from guests with distinctively African American names were about sixteen per cent less likely to be accepted, roughly forty-two per cent compared with fifty per cent. The discrimination appeared among hosts of different races and genders and in both shared and whole-home rentals. It was concentrated among hosts who had never hosted an African American guest.

The findings, along with personal accounts shared by users on social media, prompted Airbnb to commission a review of discrimination on the platform, led by the civil rights expert Laura Murphy. The company introduced a nondiscrimination policy that users had to accept, expanded instant booking, which lets guests book without host approval, and in 2018 changed its system so that hosts would see a guest's profile photo only after accepting a booking.

The episode illustrates a fundamental challenge for platform designers. The markers that build trust among strangers are often the same markers that enable discrimination. A photo helps a host feel that a guest is a real person, but it also reveals characteristics that some hosts use to exclude. Removing the photo reduces discrimination but may also reduce the sense

of personal connection that helps hosts trust guests. Designers must decide which markers to show, when and to whom, knowing that each choice has costs.

Ride-hailing and mutual ratings

Ride-hailing platforms such as Uber and Lyft faced a similar challenge in a different setting. Getting into a stranger's car has long been something parents warn children against. To make it ordinary, the platforms built a system of markers: driver profiles with names, photos and vehicle details, real-time tracking of the trip, cashless payment and mutual ratings after each ride.

Mutual ratings created a two-way reputation system. Drivers rated passengers, and passengers rated drivers. Drivers whose average ratings fell below thresholds set by the platform could lose access to it. Passengers with low ratings might find it harder to get rides. Both sides had an incentive to behave well.

The rating system also showed the inflation discussed in earlier chapters. Because low ratings could cost drivers their livelihoods, many passengers came to treat anything less than five stars as a serious complaint. The five-point scale effectively became a binary signal: five stars for acceptable, anything less for a problem. Drivers, aware of this, sometimes asked passengers directly for five-star ratings. The marker continued to function, but its meaning changed.

There were also concerns that ratings could reflect passengers' biases. Researchers and advocates raised the possibility that drivers from some groups might receive systematically lower ratings for reasons unrelated to their driving, which could affect their ability to stay on the platform. Platforms using ratings to make consequential decisions about workers face a responsibility to check whether those ratings are fair.

Trust frameworks

Some platforms have tried to articulate explicitly how they build trust. BlaBlaCar, a French long-distance ride-sharing platform that connects drivers with spare seats and passengers going the same way, published research in 2016, conducted with academics at New York University's Stern School of Business, on how its members came to trust one another. The company described a framework of trust signals, including declared information such as names and photos, ratings from previous trips, engagement such as payment in advance, and moderation by the platform. It reported that members with complete profiles and positive ratings were trusted at levels approaching those reserved for friends and family, and well above those for neighbours or colleagues.

Whatever the precise figures, the framework captures something important. Trust among strangers on platforms is built from multiple layers of markers, each of which adds information. Declared information tells you who someone says they are. Verification tells you that the platform has checked some of it. Ratings tell you how others have experienced them. Engagement signals commitment. Moderation tells you that the platform will step in if something goes wrong. No single marker is sufficient. Together, they can create enough trust for strangers to share a car for several hours.

The platform as sovereign

Platforms that build trust among strangers take on a role that resembles that of a state. They set the rules for who may participate. They issue markers, such as verified badges and ratings, that other participants rely on. They resolve disputes and punish violations. They guarantee transactions. In Moffett's terms, they define the boundary of an anonymous society and the markers that mark its members.

This role brings power and responsibility. A platform that removes a participant from its marketplace may deprive them of a livelihood. A platform that changes its rating rules may change the fortunes of millions of users. A platform that allows discrimination to flourish may reproduce, at scale, the exclusions that civil rights laws were designed to end. Platform designers are, whether they acknowledge it or not, in the business of governing anonymous societies.

Principles for platform trust

The experience of Airbnb, ride-hailing platforms and BlaBlaCar suggests several principles for designing trust between strangers.

Layer the markers. Combine declared information, verification, reputation, engagement and guarantees. Each layer compensates for the weaknesses of the others.

Make markers mundane. Clear photos, accurate descriptions and simple verification steps often do more for trust than elaborate features.

Back markers with guarantees. When markers fail, guarantees and insurance limit the damage and reassure participants that the platform stands behind its system.

Choose what to reveal. Every marker that builds trust can also enable bias. Decide carefully which information to show, at what stage of the interaction and to whom.

Audit for fairness. Test regularly whether markers produce systematically different outcomes for different groups, and adjust designs when they do.

Govern openly. Platforms that set rules for millions of participants should explain those rules, apply them consistently and provide ways to appeal. Procedural fairness, as Chapter 13 argued, matters as much as outcomes.

The next chapter turns to communities that have built trust among strangers without the commercial incentives of marketplaces: the volunteers of Wikipedia and the developers of open-source software.

CHAPTER 24

Commons Without Kin

On 15 January 2001 Jimmy Wales and Larry Sanger launched Wikipedia, a free online encyclopedia that anyone could edit. The idea seemed to many observers almost as implausible as sleeping in a stranger's home. An encyclopedia written by anonymous volunteers, with no payment, no credentials required and no central editorial control, would surely be riddled with errors, vandalism and bias. Within a few years Wikipedia had become one of the most visited websites in the world, and its English edition alone would eventually contain millions of articles.

Wikipedia and the open-source software movement are among the most remarkable examples of cooperation among strangers in human history. Volunteers who have never met, scattered across the globe, collaborate to build shared resources of enormous value. They do so without the incentives of markets or the authority of employers. They rely, instead, on markers, norms and institutions that they have developed themselves. Studying how they work, and where they struggle, reveals much about how trust can be built in communities that are neither kin nor company.

Ostrom's commons

The political economist Elinor Ostrom spent much of her career studying how communities manage shared resources, such as fisheries, forests, pastures and irrigation systems, without either privatising them or handing them to the state. Her 1990 book *Governing the Commons* challenged the widely held view, associated with the ecologist Garrett Hardin's 1968 essay

"The Tragedy of the Commons," that shared resources would inevitably be overused. Ostrom documented many communities that had managed commons sustainably for long periods, and she identified a set of design principles that successful communities tended to share. In 2009 she became the first woman to receive the Nobel memorial prize in economics.

Ostrom's design principles include clearly defined boundaries, so that it is clear who has the right to use the resource; rules adapted to local conditions; arrangements that let those affected by the rules participate in changing them; monitoring of behaviour by people accountable to the community; graduated sanctions for those who break the rules; accessible means of resolving conflicts; recognition by outside authorities of the community's right to organise; and, for larger systems, nested layers of governance.

Several of these principles are about markers. Clearly defined boundaries require some way of distinguishing members from non-members. Monitoring requires visible information about behaviour. Graduated sanctions require a reputation system that records violations. Ostrom's commons were, in Moffett's terms, anonymous or semi-anonymous societies held together by markers and rules. Digital commons such as Wikipedia and open-source projects face the same challenges at a far larger scale.

How Wikipedia builds trust

Wikipedia's markers of trust are partly technical and partly social. Every edit is recorded in a public history, so that anyone can see who changed what and when. Contributors have usernames and user pages where they can describe themselves and their interests. Many display userboxes, small badges that signal their languages, expertise, interests and affiliations. Editors award one another barnstars, virtual awards that recognise particular contributions. Edit counts and the length of a contributor's history serve as informal markers of experience.

More formal markers confer authority. Experienced editors can be granted administrative privileges through a community process in which other editors discuss and vote on their suitability. Administrators can protect pages, block vandals and delete articles. Further roles carry additional powers. These markers create a hierarchy based on demonstrated commitment and judgment rather than credentials or employment.

Wikipedia also developed an extensive body of policies and guidelines covering everything from neutrality and verifiability to civility and conflict of interest. Experienced editors cite these policies by name and abbreviation in discussions, creating a dialect of the kind described in Chapter 18. Fluency in that dialect is itself a marker of membership.

The cost of quality

Wikipedia's markers and rules helped it achieve remarkable quality for a volunteer project. They also created problems. In a 2013 paper titled "The Rise and Decline of an Open Collaboration System," Aaron Halfaker, R. Stuart Geiger, Jonathan Morgan and John Riedl analysed the history of the English Wikipedia's editor community. They found that the number of active editors had grown rapidly until around 2007 and then begun a sustained decline. Their analysis suggested that a key factor was the treatment of newcomers. As Wikipedia grew, it developed automated tools and stricter quality controls to fight vandalism. These tools were effective, but they also increasingly rejected good-faith contributions from newcomers, who often did not yet know the community's norms. Newcomers whose first edits were reverted were less likely to continue editing.

The finding captures a dilemma that runs throughout this book. Markers and rules protect a community against forgers, vandals and free riders. They also make it harder for legitimate strangers to join. A community that is too open is overrun. A community that is too closed stagnates. Wikipedia had become, in effect, an individual-recognition society for its core editors, with high barriers for outsiders. The researchers

suggested that the community needed ways to socialise newcomers more effectively, and Wikipedia's foundation and volunteers subsequently experimented with initiatives to welcome and support new editors.

Open source and the signed-off line

Open-source software projects face similar challenges. Contributors around the world submit changes to shared code, and maintainers decide which changes to accept. Because the software may run on millions or billions of devices, accepting a bad change can have serious consequences. Maintainers must decide whom to trust, often on the basis of limited information about people they have never met.

The Linux kernel, the core of the operating system that Linus Torvalds began as a personal project in 1991, has developed an elaborate structure of maintainers responsible for different subsystems, who review and pass on contributions in a layered hierarchy. In 2004, amid legal disputes over the origins of code in Linux, the kernel community introduced the Developer Certificate of Origin, a short statement that contributors affirm by adding a line reading "Signed-off-by" with their name and email address to each contribution. The line is a marker. It declares that the contributor has the right to submit the code under the project's licence, and it creates a record of who vouched for each change.

Open-source communities rely heavily on reputation built through contribution. A developer who has submitted many good changes over years, responded constructively to review and helped others becomes trusted. Their changes receive less scrutiny, and they may eventually be granted maintainer status, with the power to accept others' contributions. Platforms such as GitHub make these histories visible through contribution graphs, commit logs and profiles.

The xz backdoor

The weakness of this system was exposed dramatically in 2024. On 29 March of that year, Andres Freund, a software engineer at Microsoft, reported that he had discovered a backdoor in xz Utils, a widely used open-source data compression library included in many Linux distributions. Freund had noticed that logins over SSH on a test system were taking slightly longer than expected and investigated. He found malicious code that had been inserted into recent versions of the library, apparently designed to allow unauthorised access to affected systems. The vulnerability was assigned the identifier CVE-2024-3094. Because the compromised versions had only recently been released, they had not yet reached most stable production systems, and a potentially catastrophic attack was largely averted.

The backdoor had been inserted by a contributor using the name Jia Tan, who had begun contributing to the project around 2021. Over the following two years, Jia Tan made many legitimate contributions and gradually earned the trust of the project's long-standing maintainer, Lasse Collin, who had been maintaining the project largely alone. During this period, other accounts, which investigators later suspected were part of the same operation, posted messages pressing Collin to accept more help and hand over responsibilities, citing the slow pace of development. Collin had mentioned in public messages that he had limited capacity, including for personal health reasons. Jia Tan eventually became a co-maintainer, with the ability to make releases. The malicious code followed.

The xz case is a perfect illustration of the vulnerability described in Chapter 2, when butterfly caterpillars infiltrate ant colonies by mimicking the colony's markers. Jia Tan acquired the markers of a trusted contributor through sustained, apparently legitimate work. Social pressure from other accounts accelerated the transfer of authority. The open-source community's reputation system, which rewards long contribution with

trust, was used against it. Many security researchers have concluded that the operation was patient and well resourced, though the identity of those behind it has not been publicly established.

The burden on the few

The xz case also exposed a structural problem. Much of the digital infrastructure on which modern economies depend is maintained by small numbers of volunteers, often unpaid, sometimes working alone. The writer Nadia Eghbal documented this problem in a 2016 report titled *Roads and Bridges* and in her 2020 book *Working in Public*. An earlier crisis had made the point vividly. In 2014 a serious vulnerability known as Heartbleed was discovered in OpenSSL, a widely used encryption library. It emerged that the project was maintained by a very small team with limited funding, despite being used by a large share of the world's web servers. Following Heartbleed, technology companies and the Linux Foundation launched initiatives to fund critical open-source projects.

In systems terms, the open-source commons suffers from a version of the tragedy of the commons. Many companies benefit from the shared resource, but few contribute to its maintenance. The burden falls on a small number of maintainers, who may become overworked, burnt out and vulnerable to manipulation. When a maintainer is exhausted and alone, an offer of help from a capable contributor is hard to refuse, even if the contributor's markers are not fully verified.

A gift culture

Why do volunteers give their time to projects that pay them nothing? In a 1998 essay titled "Homesteading the Noosphere," the software developer and writer Eric S. Raymond argued that open-source communities function as gift cultures, in which status is earned not by accumulating resources but by giving them away. A developer who contributes useful code, documenta-

tion or help gains reputation among peers, and that reputation is the main reward. Raymond compared these communities to other societies, described by anthropologists, in which prestige comes from generosity.

The idea connects directly to the discussion of wealth in Chapter 9. In a gift culture, the community has agreed to count contribution as wealth. The markers of that wealth, such as commit histories, maintainer titles, conference invitations and the recognition of peers, are what contributors accumulate. Like any currency, they shape behaviour. Developers pursue the kinds of contribution that earn recognition, which may not always be the kinds the project most needs. Writing new features tends to earn more visible credit than fixing old bugs, reviewing other people's code or maintaining infrastructure, even though the latter tasks are essential.

This helps explain the burden on maintainers described earlier in this chapter. Maintenance is the least glamorous form of contribution. It earns little visible status relative to the effort it requires, and it often brings complaints rather than thanks. A gift culture that rewards novelty more than care will tend to leave its foundations under-maintained. Communities and companies that depend on open source can help by recognising maintenance explicitly, funding it and creating markers that honour the people who keep shared infrastructure running.

Lessons from the digital commons

Wikipedia and open-source software show that strangers can build shared resources of extraordinary value through markers, norms and voluntary institutions. They also show the dilemmas such communities face.

Boundaries must be porous but defended. Communities need to welcome newcomers to survive, but they also need to protect against vandals and infiltrators. Finding the balance requires continuing attention to how newcomers are treated.

Reputation is powerful and exploitable. Trust built through contribution is a strong marker, but it can be acquired deliberately by people with malicious intent. Critical decisions, such as granting release authority, should not rest on reputation alone.

Maintainers need support. Communities that depend on a few overworked individuals are fragile. Funding, shared responsibility and succession planning reduce the risk that a single exhausted maintainer becomes a point of failure.

Ostrom's principles apply. Clear boundaries, participatory rules, monitoring, graduated sanctions and accessible conflict resolution help digital commons as they help physical ones.

The next chapter looks more broadly at the counterfeit markers that threaten all of these systems, from fake verification badges to deepfaked executives, and at what the age of artificial intelligence means for trust among strangers.

CHAPTER 25

Counterfeit Belonging

On 10 November 2022, shortly after Twitter began selling its blue verification checkmark to any user willing to pay eight dollars a month, an account bearing the name and logo of the pharmaceutical company Eli Lilly posted a message announcing that insulin was now free. The account carried a blue check. It was not Eli Lilly. The real company posted an apology for the confusion caused by what it called a fake account, and the impersonation was one of many that appeared in the days after the change, as users paid for checkmarks and used them to pose as brands, politicians and celebrities. Twitter paused the paid verification programme within days and relaunched it later with changes.

The blue check had been, for years, one of the most recognisable markers on the internet. It signalled that an account belonged to the person or organisation it claimed to represent, and that the platform had checked. Its value lay in the cost of obtaining it, which was not money but verification. When the marker became available to anyone who could pay a modest fee, it stopped signalling what users had learned to expect. For a few days, the most visible marker of authenticity on the platform became a tool of impersonation.

This chapter is about counterfeit markers: signals of membership, identity or trustworthiness that are produced without the underlying reality. Counterfeits are as old as markers themselves, as the caterpillars in Chapter 2 and the impostors in Chapter 16 showed. What is new is the scale and sophistication that digital technology, and especially artificial intelligence, now makes possible.

The economics of forgery

The value of a marker depends on the cost of forging it, as earlier chapters argued. When forging is expensive, markers carry trust. When it becomes cheap, trust drains. Much of the history of markers can be read as a contest between those who produce markers and those who copy them.

Digital technology has lowered the cost of copying many markers dramatically. A logo can be reproduced perfectly in seconds. A website can be cloned. An email can be made to appear to come from a trusted sender. A social media profile can be created with a stolen photo and a plausible name. For years, the main defences against these forgeries were small imperfections, such as awkward language in a phishing email, a slightly wrong domain name or a low-quality image, that alert recipients could spot.

Artificial intelligence is removing many of those imperfections. Language models can produce fluent text in any style. Image generators can create realistic photographs of people who do not exist. Voice cloning can reproduce a person's speech from a short sample. Video synthesis can make a face move and speak convincingly. Each of these capabilities lowers the cost of forging markers that people have long relied on to recognise one another.

The deepfake meeting

In early 2024 police in Hong Kong described a case in which an employee at a multinational company's Hong Kong office had been persuaded to transfer around two hundred million Hong Kong dollars, roughly twenty-five million US dollars, after taking part in a video call with people who appeared to be the company's chief financial officer and other colleagues. According to the police account, the other participants in the call were deepfakes, generated using publicly available video and audio of the real executives. In May 2024 the British engineering firm Arup confirmed that it was the company involved.

The case is a vivid example of counterfeit belonging. The employee had reportedly been suspicious of an initial email requesting a secret transaction. What overcame the suspicion was a set of markers: familiar faces, familiar voices and the presence of several apparent colleagues on the same call. In an anonymous society, such markers are exactly what people use to decide whom to trust. The forgery succeeded not by breaking any technical system but by reproducing the markers of membership convincingly enough to trigger trust.

An earlier case, reported by the *Wall Street Journal* in 2019, involved a UK-based energy company whose chief executive was persuaded to transfer about two hundred and twenty thousand euros after a phone call in which a voice that sounded like his boss at the parent company asked him to make an urgent payment. The company's insurer said it believed the voice had been generated with AI software.

These cases are not yet common, and many fraud schemes still rely on simpler methods. Business email compromise, in which criminals impersonate executives or suppliers by email to request payments, remains far more widespread. The FBI's Internet Crime Complaint Center has reported losses from such schemes running to billions of dollars a year. But the deepfake cases show the direction of travel. Markers that once required physical presence or personal knowledge to fake can now be forged remotely and at low cost.

Bots and synthetic crowds

Counterfeit markers also operate at the level of crowds. Social media platforms display markers of popularity, such as follower counts, likes and shares, that users read as signals of credibility. A post with many likes appears more trustworthy than one with few. An account with many followers appears more authoritative.

These markers can be inflated by automated accounts. Researchers have long studied the prevalence of bots on social media. A 2017 study by Onur Varol and colleagues estimated that between nine and fifteen per cent of active Twitter accounts at the time were bots, though estimates vary widely

depending on methods and definitions. Bots can be used to amplify messages, create the appearance of grassroots support and manipulate trending topics.

The xz case described in Chapter 24 included a version of this tactic on a small scale, with apparently fake accounts pressing a maintainer to accept help. On a larger scale, coordinated networks of fake accounts can create the impression that many people share a view or trust a source. In Moffett's terms, they forge the markers of a society, making a small group look like a large one.

Inflation of effortful markers

Artificial intelligence also threatens a subtler class of markers: those that signal effort, skill or care. A well-written cover letter once signalled that an applicant could write well and cared enough to tailor their application. A thoughtful message signalled that the sender had taken time to consider the recipient. A polished report signalled competence.

When AI can produce fluent, tailored text at negligible cost, these markers lose much of their signalling value. A well-written cover letter may now signal only access to a tool. This is the escalation trap described in Chapter 8, operating at the speed of software. As effortful markers become cheap, people who read them must look for other signals, such as live conversation, demonstrated work or verified credentials, that are harder to fake.

The same dynamic affects many other markers. Product reviews, customer testimonials, academic essays, job applications and even dating profiles all rely, in part, on the assumption that producing them takes some effort. As that assumption weakens, the markers' value declines, and institutions that depend on them must adapt.

Defences

How can anonymous societies defend their markers against counterfeiting at this scale? No single answer exists, but several approaches are emerging.

One approach is to make markers harder to forge through cryptography. Systems such as passkeys, which replace passwords with cryptographic credentials tied to a user's device, make it much harder for attackers to impersonate users. Content provenance standards, such as those developed by the Coalition for Content Provenance and Authenticity, an industry group formed in 2021 with members including Adobe and Microsoft, aim to attach verifiable records of origin to images and videos, so that viewers can check where a piece of content came from and whether it has been altered.

A second approach is to verify through independent channels. The deepfake cases suggest a simple but powerful defence: confirm unusual requests through a separate channel, such as calling a known phone number or speaking in person, before acting on them. Many organisations have introduced rules requiring such verification for payments above certain thresholds. The principle is the same as the independent checking recommended in Chapter 16. A marker presented in one channel should be confirmed in another.

A third approach is to shift reliance toward markers that remain costly. Live interaction, demonstrated skills, long-term relationships and physical presence are harder to fake than text or images. Organisations may find themselves placing more weight on these markers as cheaper ones lose value, even though they are more expensive to use.

A fourth approach is to restore cost to markers that have become cheap. Twitter's paid verification episode showed the risk of removing verification from a marker of identity. Platforms can preserve the value of identity markers by maintaining genuine verification processes, distinguishing clearly between paid features and verified identity, and acting quickly against impersonation.

The return of the personal

There is an irony in the challenge that artificial intelligence poses to markers. Anonymous societies arose because personal recognition could not scale. Markers allowed strangers to trust one another without knowing

each other personally. If markers become easy to forge, societies may be pushed back toward personal recognition, relying more on people they know and less on signals from strangers.

Some of this is already visible. People increasingly rely on trusted personal networks for recommendations, information and verification. Organisations place more emphasis on in-person meetings for high-stakes decisions. Small, closed communities, such as private messaging groups, have grown in importance relative to open public platforms.

A retreat into personal networks would carry costs. It would favour people with large, well-connected networks and disadvantage outsiders. It would reduce the scale of cooperation that anonymous societies make possible. And it would weaken the markers that allow strangers from different backgrounds to trust one another. The challenge for designers is to find ways of defending markers that preserve the benefits of anonymous cooperation while resisting forgery.

Principles for an age of counterfeits

The cases in this chapter suggest several principles for leaders.

Assume markers can be forged. Any marker that can be copied digitally may be copied. Design processes on the assumption that familiar faces, voices and logos could be counterfeit.

Verify through separate channels. Confirm high-stakes requests using a channel different from the one through which they arrived.

Prefer cryptographic markers for identity. Passwords and visual cues are weak. Cryptographic credentials and verified provenance are stronger.

Revalue costly signals. As cheap markers lose value, invest in markers that remain hard to fake, such as live interaction and demonstrated work.

Protect identity markers. Do not dilute markers of verified identity by making them available without verification.

Educate people. Employees and users who understand how counterfeits work are better able to recognise them. Training should include realistic examples of AI-enabled impersonation.

Part V turns from diagnosis to design. Having explored how markers work, how they fail and how they are attacked, the remaining chapters offer practical tools for leaders who want to design, audit and repair the markers that hold their organisations and platforms together.

PART V

Designing Markers on Purpose

The final part turns diagnosis into practice. It offers a method for auditing the markers an organisation already runs on, principles for designing new ones and a discipline for repairing markers that have broken.

CHAPTER 26

The Marker Audit

Most organisations have never looked at their markers as a system. They know their brand, because someone is paid to manage it. They know their credentials and job titles, because human resources maintains them. They may have a values statement and a code of conduct. But the full set of signals by which members recognise one another, by which strangers decide whether to trust the organisation and by which newcomers learn to belong, is rarely mapped. It grows by accretion, through decisions made by many people at different times for different reasons.

A marker audit is a structured way of seeing that system. It does not require specialist tools. It requires curiosity, honesty and a willingness to look at familiar signals as if seeing them for the first time. This chapter sets out a method in five stages: inventory, testing, mapping the loops, mapping the boundaries and setting priorities. Leaders can use it for a whole organisation, a division, a platform or a single team.

Stage one: inventory

The first stage is to list the markers. The aim is breadth rather than precision. It helps to work through categories, drawing on the kinds of markers described in earlier chapters.

Symbols include the organisation's name, logo, colours, slogans and any internal emblems, such as team names, product codenames or mascots.

Credentials include job titles, grades, certifications, licences, internal qualifications and any badges or awards that signal competence or achievement.

Bodily and material markers include uniforms, dress codes, badges, lanyards, office layouts, equipment and the physical spaces people occupy.

Linguistic markers include internal vocabulary, acronyms, codified values, preferred document formats, communication habits and the dialects of different groups.

Procedural markers include the unwritten rules of how things are done: how meetings run, how decisions are made, how to request resources, whom to copy on messages and how disagreement is expressed.

Reputational markers include ratings, reviews, performance scores, public recognition and the informal reputations that circulate through gossip.

Digital markers include usernames, profiles, verification badges, access levels, channel memberships, custom emoji and the visible traces of activity in shared tools.

Economic markers include pay grades, budgets, headcount, equity and any internal currencies that signal status or power.

For each marker, the inventory should note who issues it, who carries it, who reads it and what it is supposed to signal. The last question is often the most revealing. Many markers persist long after anyone remembers what they were meant to indicate.

The inventory should be compiled from several vantage points. Leaders see some markers clearly and others not at all. Newcomers are especially valuable informants, because they notice markers that long-standing members take for granted. Employees of recently acquired units, contractors, frontline staff and customers can each reveal markers that are invisible from the centre. A useful technique is to ask newcomers, in their first weeks, to keep a list of everything they had to learn in order to be taken seriously.

Stage two: testing

The second stage is to test each significant marker against four questions. Together they indicate whether a marker is doing its job.

The first question is accuracy. Does the marker actually predict what it is supposed to signal? A credential should predict competence. A rating should predict quality of service. A dress code should signal the qualities the organisation wants to project. Testing accuracy may require data, such as whether people with a particular credential actually perform better, or it may require judgment, such as whether a values statement matches observed behaviour. Chapter 7 described markers as information flows about an invisible stock. Accuracy asks whether the information is true.

The second question is cost to forge. How hard is it for someone who lacks the underlying quality to acquire the marker? Markers that are cheap to forge invite impostors, inflation and gaming. Markers that are expensive to forge carry more trust. Testing this question involves imagining how someone might acquire the marker dishonestly and estimating how difficult that would be. Chapters 16 and 25 offer many examples.

The third question is accessibility. Can legitimate members acquire the marker at reasonable cost? Markers that are expensive for honest people to obtain exclude capable members and narrow the organisation's talent pool. Credential inflation, discriminatory dress codes and reputation systems that disadvantage newcomers all fail this test. Accessibility must be weighed against cost to forge, since many features that make a marker hard to forge also make it hard to acquire.

The fourth question is currency. Is the marker still current? Markers carry information about the past. A credential earned long ago, a rating accumulated years before, a brand built under previous leadership or a values statement written for a smaller company may no longer reflect present reality. Testing currency involves asking when the marker was last checked against the quality it signals, and how quickly it would change if that quality changed.

A simple scoring exercise can help. For each major marker, rate accuracy, cost to forge, accessibility and currency on a rough scale, such as low, medium and high. Markers that score low on accuracy or currency need attention urgently, because people are relying on false information. Markers that score low on cost to forge are vulnerable to attack. Markers that score low on accessibility are likely excluding people unnecessarily.

Stage three: mapping the loops

The third stage is to map the feedback loops in which markers participate. Donella Meadows's tools are useful here. For each significant marker, ask what reinforcing and balancing loops it is part of.

Reinforcing loops amplify. A strong reputation attracts business, which produces more positive ratings, which strengthens reputation. A prestigious credential attracts better opportunities, which produce achievements that reinforce prestige. A strong culture attracts people who fit it, who reinforce the culture. Reinforcing loops can be virtuous or vicious. The same loop that builds a strong reputation for established participants can lock newcomers out.

Balancing loops stabilise. When ratings fall, a seller may improve service to bring them back up. When a credential becomes too common, employers may look for other signals. When a culture becomes too rigid, talented people may leave, forcing change. Balancing loops can keep markers honest, but they can also resist necessary change, as Chapter 21 showed.

Mapping the loops helps reveal which of the traps described in Chapter 8 may be operating. Rating compression suggests drift to low performance. Credential inflation suggests escalation. Concentration of reputation suggests success to the successful. Widespread use of a marker without meeting its standard suggests the tragedy of the commons. Dependence on markers without direct assessment suggests shifting the burden. Compliance without substance suggests rule beating. Optimisation of the marker itself suggests seeking the wrong goal.

Delays deserve particular attention. For each marker, ask how long it would take for a change in the underlying quality to show up in the marker. A long delay means the marker can mislead for a long time before anyone notices. Markers with long delays need more frequent direct checks.

Stage four: mapping the boundaries

The fourth stage is to map the boundaries that markers create. Every marker distinguishes those who carry it from those who do not. Some boundaries are intended and useful, such as the boundary between licensed and unlicensed practitioners of surgery. Others are unintended and harmful, such as the boundary between employees and contractors who do the same work but are treated very differently.

A boundary map identifies the main groups that markers create inside and around the organisation. It might show full-time employees, contractors, recent hires, long-standing staff, members of acquired units, headquarters and field staff, different professions and different levels of seniority. For each boundary, ask several questions. What markers define it? How easy is it to cross? Who benefits from it and who is disadvantaged? Does it serve the organisation's purpose, or has it outlived its usefulness?

The boundary map should also show external boundaries: between the organisation and its customers, partners, suppliers, regulators and the wider public. Markers at these boundaries determine how strangers outside the organisation decide whether to trust it. A brand, a certification or a published code of conduct are boundary markers of this kind.

Boundary mapping often reveals the "real people" distinction described in Chapter 6. Some groups inside the organisation are treated as full members, while others are treated as something less. The distinction may never be stated, but it shows up in who is invited, consulted, promoted and trusted. Making it visible is the first step toward deciding whether it should exist.

Stage five: setting priorities

The final stage is to decide what to do. No organisation can redesign all its markers at once, and attempting to do so would itself be destabilising, since markers are stocks that people depend on. The audit should produce a short list of priorities.

Urgent priorities usually include markers that score low on accuracy or currency, since people are relying on false information, and markers that score low on cost to forge in high-stakes areas, since they invite serious attack. Markers that create harmful boundaries, especially boundaries that exclude people on grounds unrelated to their contribution, also deserve priority.

Important but less urgent priorities may include markers caught in slow-moving traps, such as gradual rating compression or credential inflation, and markers whose meaning has drifted away from their original purpose.

Some markers should simply be left alone. Markers that are accurate, hard to forge, accessible and current are assets. Changing them risks draining trust unnecessarily. The New Coke episode in Chapter 14 is a reminder of the cost of changing markers that were working.

Running the audit well

A few practical observations can help leaders run marker audits effectively.

Involve the people who carry the markers. Markers belong partly to the communities that use them. An audit conducted entirely by leaders or consultants will miss much of what markers mean to members, and its conclusions will be resisted.

Listen to newcomers and outsiders. They see markers that insiders take for granted, and they experience boundaries that insiders do not notice.

Look for contradictions. Gaps between what markers claim and what people experience are the most important findings of any audit. A values statement that contradicts observed behaviour, a credential that does not predict performance or a brand that promises what the organisation does not deliver are all warning signs.

Be honest about power. Markers distribute status and opportunity. An audit that threatens existing markers will threaten the people who benefit from them. Leaders should expect resistance and be prepared to explain why changes are necessary.

Repeat the audit. Markers evolve, as Chapter 12 argued. An audit is a snapshot, and its findings will become outdated. Organisations that audit their markers periodically, perhaps every year or two, or whenever they undergo major change such as rapid growth, a merger or a shift to remote work, are better placed to keep their markers honest.

The audit tells leaders where they stand. The next chapter turns to what they can build: principles for designing new markers on purpose.

CHAPTER 27

Designing Markers on Purpose

Hunter-gatherers did not design their languages, and nations rarely design their accents. Most markers in human history emerged through the evolutionary process Eric Beinhocker describes: variation, selection and amplification, operating over generations without anyone in charge. Yet some of the most consequential markers of the modern world were designed deliberately. The passport format recommended in 1920, the euro banknotes introduced in 2002, eBay's Feedback Forum, Airbnb's host guarantee, the Linux kernel's sign-off line and the tamper-resistant seal on a medicine bottle were all the products of conscious choices by identifiable people.

Leaders of organisations and platforms are, increasingly, in the position of those designers. They create markers whether they intend to or not, through every decision about names, titles, credentials, rituals, ratings and access. This chapter sets out principles for doing so on purpose. The principles draw on everything the earlier chapters have described: Moffett's account of how markers hold anonymous societies together, Meadows's tools for seeing how markers behave as parts of systems, and Beinhocker's view of markers as designs that evolve under selection.

Principle one: start from the stranger

Every marker exists to help one stranger decide how to treat another. Design should therefore begin not with the organisation's self-image but with the stranger's question. Who is meeting whom, in what situation, and what do they need to know in order to cooperate?

A hospital patient needs to know which person in the room can answer a question about medication. A host on a home-sharing platform needs to know whether a guest will respect the property. A maintainer of an open-source project needs to know whether a contributor's code can be trusted. A new employee needs to know whom to ask for help and how to make a request without causing offence. Each situation calls for different markers.

Starting from the stranger helps avoid a common mistake, which is to design markers that express what the organisation wants to say about itself rather than what strangers need to know. A values statement full of aspirational language may please its authors but tell a newcomer nothing about how decisions are actually made. A logo redesign may express a new strategy but confuse customers who relied on the old one to find the product.

Principle two: tie markers to substance

A marker is valuable only insofar as it predicts something real. The most durable markers are those that are directly connected to the qualities they signal. The tamper-resistant seal on a bottle of Tylenol did not merely claim that the product was safe. It allowed customers to check. A reputation score built from verified transactions reflects real experiences. A credential based on demonstrated skill predicts competence better than one based on hours of attendance.

Where possible, markers should allow verification. If a marker can be checked against the underlying quality, people can trust it without taking it on faith, and forgers face a higher cost. Where verification is impossible, markers should at least be issued by parties with a strong incentive to keep them accurate, and they should be checked periodically against independent evidence.

Tying markers to substance also means accepting that markers will sometimes reveal unwelcome truths. A rating system that accurately reflects poor service will produce bad ratings. A credential that accurately

measures skill will exclude some people who want it. Designers who soften markers to avoid uncomfortable results will find that the markers lose their value.

Principle three: price the forgery

A marker can bear only as much trust as it is expensive to fake. Designers should consider, explicitly, how someone who lacked the underlying quality might acquire the marker, and how expensive that would be.

For low-stakes markers, a low cost of forgery may be acceptable. A custom emoji on a team channel does not need to be hard to fake. For high-stakes markers, such as identity verification, access to critical systems or authority to release software, the cost of forgery must be high. Cryptographic credentials, independent verification, multiple sign-offs and long track records all raise the cost.

Pricing the forgery also means anticipating how the cost will change. Technologies that make copying cheaper, as Chapter 25 described, can quickly erode the security of markers that once seemed robust. Designers should revisit the cost of forgery regularly, especially for markers that depend on appearance, voice, writing or other signals that artificial intelligence can now reproduce.

Principle four: build on-ramps

A marker system that is impossible for newcomers to enter will stagnate. Anonymous societies survive by absorbing strangers, and organisations and platforms are no different. Every marker system needs on-ramps: ways for legitimate newcomers to acquire markers quickly enough to participate.

On-ramps can take many forms. Platforms can offer guarantees, verification badges or introductory periods that let newcomers establish trust before they have accumulated reputation. Organisations can provide mentors, onboarding cohorts and explicit guides to their dialect and procedures. Professions can recognise credentials from other jurisdictions or offer alternative routes to qualification. Communities can create spaces

where newcomers can ask questions and make mistakes without penalty. In 2012 Wikipedia's community launched a space called the Teahouse, designed as a friendly place for new editors to ask questions and get help, as one response to the problems of newcomer retention described in Chapter 24.

On-ramps must be designed with the same care as other markers. If they are too easy, they become a route for forgers and free riders. If they are too hard, they fail to bring newcomers in. The right balance depends on the stakes. A community can afford generous on-ramps for low-risk activities and stricter ones for high-risk privileges.

Principle five: layer markers

No single marker is sufficient for trust among strangers. The most robust systems combine several layers, each of which compensates for the weaknesses of the others. The BlaBlaCar framework described in Chapter 23, with its combination of declared information, ratings, engagement and moderation, is an example. Airbnb's combination of profiles, reviews, messaging, payment handling and guarantees is another.

Layering serves two purposes. It increases the information available to strangers, allowing them to form a more accurate judgment. And it raises the cost of forgery, since a forger must fake several markers at once rather than one. The xz backdoor in Chapter 24 succeeded partly because a single marker, a long record of contribution, was sufficient to gain release authority. A layered system might have required additional checks, such as independent review of release changes or verification of the contributor's identity, before granting that authority.

Layering has costs. Each additional marker adds friction, and too much friction deters legitimate participants. Designers should match the depth of layering to the stakes of the interaction.

Principle six: draw the boundary deliberately

Every marker creates a boundary between those who carry it and those who do not. Designers should decide where they want that boundary to fall, rather than letting it fall by accident.

The history of human societies, as Chapter 6 described, is in part a history of moving the boundary outward, extending the circle of people who count as real. Organisations and platforms can do the same. A company can define its people by shared purpose and standards rather than shared background. A platform can design markers that are accessible to people from many cultures, rather than markers that favour those who already belong. A profession can recognise competence acquired through many routes.

Drawing the boundary deliberately also means recognising when a boundary is necessary. Some distinctions protect people, such as the boundary between licensed and unlicensed surgeons or between verified and unverified accounts. The aim is not to eliminate boundaries but to ensure that each boundary serves a purpose that the people on both sides of it could, in principle, accept as fair.

Principle seven: allow variation within shared identity

Beinhocker's evolutionary view suggests that marker systems need variation to adapt. A culture that permits no variation in its markers becomes rigid. A culture that permits unlimited variation fragments. The most resilient organisations combine a small set of shared markers that define membership in the whole with freedom for teams and communities to develop their own local markers.

The euro's coins illustrate the idea: a common side shared across the eurozone and a national side chosen by each country. Disney's preservation of Pixar's identity within the larger company is another example. Within organisations, leaders can define a few markers that everyone shares, such

as core values, a common language for key processes and shared rituals, while allowing teams to develop their own names, emoji, traditions and styles.

Local markers are not a threat to shared identity if they are nested within it. They become a threat when they replace it, so that people identify with their team or unit rather than the whole. Designers should watch for signs that local identities are hardening into rival societies, especially in large organisations, after mergers and in remote settings.

Principle eight: govern openly

Markers distribute trust, status and opportunity. The processes by which they are awarded, changed and withdrawn should therefore be seen as fair. Chapter 13 described research showing that procedural fairness matters to people as much as outcomes. A marker system that is transparent, consistent and open to appeal will be accepted more widely than one that operates opaquely.

Open governance means explaining what markers mean and how they are earned. It means applying rules consistently. It means providing ways for people to challenge decisions about their markers, such as a disputed rating, a denied credential or a removed account. And, following Ostrom's principles, it means involving the people affected by markers in decisions about how they are designed and changed.

Principle nine: plan for failure

Every marker system will fail somewhere. Markers will be forged, inflated, misused or outdated. Designers should plan for failure from the start. That means monitoring markers for signs of the traps described in Chapter 8, maintaining a capacity for direct assessment that does not depend on markers, building guarantees that limit the damage when markers fail and preparing to respond quickly and honestly when a failure becomes visible.

The next chapter takes up the response to failure in detail. For now, the key point is that a marker system designed with failure in mind will recover faster than one designed on the assumption that it will always work.

A design brief

Leaders who want to apply these principles can use a simple design brief for any new marker. It asks a short series of questions.

Who is the stranger, and what do they need to know? What quality should the marker signal? How can the marker be tied to that quality, and can it be verified? How costly is it to forge, and how will that cost change? How will legitimate newcomers acquire it? What other markers will it be layered with? What boundary does it create, and is that boundary fair? How much local variation will it allow? How will it be governed, and how can decisions about it be challenged? How will we know if it is failing, and what will we do then?

None of these questions has a simple answer. But asking them before a marker is introduced is far cheaper than discovering the answers after it has gone wrong.

CHAPTER 28

Repairing Broken Markers

Markers break in many ways. A brand is exposed as hollow. A credential system is found to admit impostors. A rating scale inflates until it means nothing. A merger replaces familiar symbols with foreign ones. A culture's stated values are revealed to be contradicted by its leaders' behaviour. A verification badge becomes available to anyone who pays. In each case, the information flow that let strangers infer a stock of trust has been corrupted, and trust begins to drain.

Repairing broken markers is among the hardest tasks a leader faces, because it must be done under pressure, in public and against the asymmetry that makes trust slow to build and quick to lose. This chapter draws on cases from earlier in the book to set out a discipline for repair. It is organised around the bathtub that Donella Meadows used to explain stocks and flows, because repair is, at bottom, a matter of stopping the outflow and restoring the inflow.

Step one: diagnose the break

Before attempting repair, leaders need to understand what kind of break has occurred. Different breaks call for different responses.

Some markers break because they were inaccurate from the start, signalling qualities the organisation did not have. Theranos is the extreme example. Repair in such cases requires changing the substance, not just the marker, and it may be impossible if the substance cannot be supplied.

Some break because they drifted. A credential that once predicted competence no longer does. A rating scale that once discriminated no longer does. Repair requires recalibrating the marker to restore its connection to the underlying quality.

Some break because they were forged. Counterfeiters, impostors or infiltrators acquired the marker without the substance. Repair requires raising the cost of forgery and identifying and removing the forgeries.

Some break because they were changed. A rebrand, a merger or a culture programme replaced familiar markers with unfamiliar ones, and members experienced the change as a loss of identity. Repair may require restoring the old markers, blending them with the new or explaining the change more convincingly.

Some break because the organisation's behaviour contradicted them. A company that claimed to put customers first was revealed to be exploiting them. A platform that promised safety failed to provide it. Repair requires changing the behaviour and then waiting for the evidence of change to accumulate.

Misdiagnosis is common and costly. An organisation that treats a behavioural break as a communication problem will invest in messaging when it needs to invest in change. An organisation that treats a forgery problem as a drift problem will recalibrate a marker that is still being forged.

Step two: stop the outflow

When a marker breaks, trust drains. The first priority is to stop the drain, or at least to slow it.

The Tylenol case in Chapter 7 shows the logic. Johnson and Johnson did not begin by arguing that its product was safe. It removed the product from shelves, which stopped further harm and signalled that the company took the threat seriously. Coca-Cola, after the New Coke backlash in Chapter 14, stopped the outflow by bringing back the original formula. Airbnb, after the 2011 ransacking incident in Chapter 23, introduced a

guarantee that limited hosts' exposure to future losses. Twitter, after the impersonation wave in Chapter 25, paused the paid verification programme within days.

Stopping the outflow often requires accepting a short-term cost. Recalls, reversals and pauses are expensive and embarrassing. They are also, usually, far cheaper than allowing the outflow to continue. Leaders who hesitate, hoping that the problem will fade, typically find that the stock of trust drains further while they wait.

Step three: tell the truth about the stock

Once the outflow has been slowed, the next step is to give stakeholders an honest account of what has happened. This is where many repair efforts fail. Organisations under pressure are tempted to minimise, deflect or obscure. They describe the problem as an isolated incident, blame a few individuals or emphasise what went right. Each of these responses may seem to protect the organisation in the short term, but each risks further draining trust, because stakeholders who later learn the full truth will feel deceived a second time.

In systems terms, the broken marker was a corrupted information flow. Repair requires restoring an accurate flow. Stakeholders need to know what the marker actually signalled, how it differed from reality and what is being done about it. Honest accounting does not mean confessing to things that are not true or accepting blame that belongs elsewhere. It means giving people the information they need to recalibrate their trust.

The DaimlerChrysler case in Chapter 20 shows the cost of failing to tell the truth. The merger-of-equals marker was contradicted by later statements and decisions, and the gap between framing and reality caused lasting damage. A more honest framing from the start might have produced a less exciting announcement but a more durable relationship.

Step four: introduce verifiable markers

Trust that has been broken is unlikely to be restored by the same kind of marker that broke. If a brand promise was hollow, a new brand promise will be met with scepticism. If a rating system was inflated, a new rating system will be distrusted until it proves itself. Repair is far more effective when it introduces markers that stakeholders can verify for themselves.

The tamper-resistant seal is the classic example. After the Tylenol poisonings, customers did not need to trust the company's assurance of safety. They could check the seal. Other examples follow the same pattern. Platforms that publish data about how their rating systems work, organisations that commission independent audits and publish the results, and companies that offer guarantees backed by money all introduce markers that do not depend solely on the word of the party that broke trust.

Verifiable markers also help with forgery problems. When an organisation discovers that its markers have been forged, it can introduce markers that are harder to forge, such as cryptographic credentials, independent verification or multiple layers of checks. The open-source community's response to the xz backdoor, for example, included renewed attention to how release authority is granted and to the support available for overstretched maintainers.

Step five: sustain the inflow

The final step is the longest. Once the outflow has stopped, the truth has been told and verifiable markers are in place, the stock of trust must be refilled. That happens only through a sustained flow of positive experiences that confirm the new markers.

There is no shortcut. Chapter 7 explained why. A drained stock refills at the pace of its inflow, and here the inflow is made of individual experiences, each of which takes time to happen. A company that has broken trust with its customers needs to deliver on its promises consistently, transaction after transaction, for years. A leader who has

broken trust with employees needs to behave consistently with stated values, decision after decision. Each experience adds a little to the stock. Any lapse drains it again.

Microsoft's cultural change, described in Chapter 21, illustrates the importance of sustained inflow. The change was not achieved by a single announcement. It was sustained over many years by consistent structural changes, strategic decisions and leadership behaviour. Organisations that try to repair trust with a burst of activity and then return to business as usual will find the stock draining again as soon as the burst ends.

Repairing markers after a merger

Mergers deserve special attention, because they so often break markers on both sides. Employees of the acquired company lose familiar names, logos, rituals and status. Employees of the acquirer may feel that their markers are being diluted by newcomers. Both sides may experience the merger as a loss of identity.

Repair after a merger follows the same steps but with particular emphases. Diagnosis should map which markers have been lost, changed or contested on each side. Stopping the outflow may require restoring some markers that were changed too quickly, such as allowing a respected brand or team name to persist. Honest accounting means being clear about who controls what and how decisions will be made, rather than maintaining a fiction of equality that decisions contradict. Verifiable markers might include published commitments about how employees from both sides will be treated, monitored by joint groups. Sustained inflow means consistent treatment over years, not months.

The Disney and Pixar case in Chapter 20 suggests that the best repair is prevention. By respecting Pixar's markers from the start, Disney avoided breaking them. For mergers already in difficulty, the lesson is to look for markers that can be restored or protected, and to build new shared markers patiently rather than imposing them.

Repairing markers after drift

Repairing markers that have drifted, such as inflated ratings or devalued credentials, poses a distinctive problem. Everyone has become used to the drifted meaning, and recalibration will disrupt their expectations. A seller with a 4.9 rating on an inflated scale may see that rating fall when the scale is recalibrated, even though their service has not changed. A graduate with a high grade from an inflated system may find it worth less when grading is tightened.

Recalibration therefore needs to be gradual, explained and, where possible, applied in ways that preserve comparability. Platforms can introduce new dimensions of rating alongside old ones, display relative rankings rather than absolute scores or reset scales with clear communication about what has changed. Institutions can report the distribution of grades alongside individual grades, so that readers can interpret them. In each case, the aim is to restore the marker's ability to discriminate without punishing people for a drift they did not cause.

Apology as a marker

Apologies are themselves markers. They signal that an organisation recognises a breach, accepts some responsibility and intends to behave differently. Like other markers, apologies carry trust only if they are costly enough to be credible and connected to substance.

Stakeholders have learned to read apologies with great sophistication. They distinguish between an apology that names the harm and one that refers vaguely to any inconvenience caused. They notice whether the apology comes from the person responsible or from a spokesperson. They notice whether it is accompanied by concrete action or stands alone. They notice whether the organisation apologises for the harm or for the fact that people are upset. An apology that seems designed to limit legal liability rather than to acknowledge harm is read as a defensive manoeuvre, and it may drain trust rather than restore it.

The most credible apologies share several features. They are specific about what went wrong. They come from leaders who have authority to change things. They acknowledge the effect on the people who were harmed rather than focusing on the organisation's intentions. They are accompanied by actions that impose a real cost on the organisation, such as compensation, recalls, changes in leadership or binding commitments. And they are followed by behaviour consistent with the apology over time. An apology without follow-through is a counterfeit marker, and stakeholders who see it as such will discount future apologies too.

The repair mindset

Repairing broken markers requires a particular mindset. It requires leaders to see trust as a stock rather than a feeling, to accept that rebuilding takes time, to prefer verifiable markers over assurances, and to tell the truth even when it is uncomfortable. It also requires humility. Leaders who presided over a break are often the least trusted people to lead the repair. Sometimes repair requires new leadership, independent oversight or the involvement of the communities whose trust was broken.

Above all, repair requires remembering what markers are for. They exist to let strangers trust one another. When a marker breaks, the strangers who relied on it have been let down. Repair is not primarily about protecting the organisation's reputation. It is about restoring the conditions under which strangers can cooperate again.

Conclusion: Real People

Return, one last time, to the morning train. The carriage fills at the first stop and empties at the last. In between, several hundred strangers share a small space without fear, without conversation and without incident. Nobody in the carriage thinks of this as an achievement. It is simply how a city works.

This book has argued that the ordinary calm of that carriage rests on one of the most remarkable capacities in the living world. Most social animals cannot live among strangers. Their societies are held together by personal recognition, and they cannot grow beyond the number of individuals their members can know. A few species broke through that limit. Ants did it with chemistry, wearing the smell of their colony on their bodies and recognising one another by scent rather than by name. Humans did it with culture, wearing the language, gestures, clothing, food habits and symbols of their society, and reading those markers in others from the first months of life.

Mark Moffett's insight that markers are what make anonymous societies possible has been the starting point for everything else in these pages. From it follow the questions that matter to leaders. If markers hold large groups together, what holds markers together? How do they gain and lose meaning? How do they fail, and how can they be repaired?

What the argument has shown

Donella Meadows's systems thinking gave us a way to see markers in motion. Trust among strangers behaves like a stock, accumulating slowly through positive experience and draining quickly through betrayal. Markers are information flows that let people infer that stock without

measuring it directly. When markers are accurate, trust can scale far beyond the circle of personal acquaintance. When markers drift, inflate, concentrate or are forged, the information flow is corrupted, and the system behaves in ways that surprise the people who built it. The traps Meadows named, from drift to low performance and escalation to success to the successful and the tragedy of the commons, recur in marker systems with remarkable regularity. And the leverage points she identified, especially information flows, rules, goals and paradigms, are where marker design has its greatest effect.

Eric Beinhocker's complexity economics gave us a way to see markers evolving. Wealth itself turns out to be a marker, defined by what a society agrees to count. Cattle in a Maasai homestead and a balance in a Frankfurt bank are both stores of value because communities honour them. Money is the marker everyone carries, and its collapse is among the fastest failures of trust a society can experience. Markers evolve through differentiation, selection and amplification, co-evolving with the physical and social technologies of their time. They are shaped by instincts of status and fairness that make them feel personal, which is why people will give up money to defend what a marker means to them.

With these tools, the book has looked at markers in the places where leaders meet them. Brands are markers of belonging, and changing them without respect for the communities that carry them invites revolt. Credentials let strangers trust one another with consequential tasks, but they inflate, exclude and can be forged. Uniforms and dress codes signal role and authority, and they can level or stratify. Organisational dialects mark membership and shape thought. Remote work weakens the informal flow of markers across groups and demands that they be made explicit. Mergers bring rival markers into collision, which helps explain why so many fail and why the few that succeed tend to respect the markers of the people they acquire. Culture resists redesign because markers are stocks with long delays, embedded in the habits and status of the people who carry them.

On digital platforms, markers have enabled cooperation among strangers on a scale no earlier society achieved. Reputation engines turned gossip into public records. Home-sharing and ride-hailing platforms layered profiles, ratings, payments and guarantees until strangers would share homes and cars. Wikipedia and open-source software built commons without kin, governed by markers and norms their members developed themselves. Each of these achievements has also exposed the weaknesses of marker systems: inflation, discrimination, hostility to newcomers and the patient infiltration that nearly compromised a critical piece of software in 2024. And artificial intelligence is now lowering the cost of forging many of the markers people have relied on to recognise one another.

The responsibility of the designer

The final part of the book turned from diagnosis to design. It offered a marker audit, principles for designing markers on purpose and a discipline for repairing them when they break. Behind all of those tools sits a single responsibility.

Every marker draws a boundary. On one side are those who carry it, who are recognised as members and extended trust. On the other side are those who do not, who are treated with suspicion or not seen at all. The hunter-gatherer peoples who called themselves by names meaning "people" or "real people" were drawing such a boundary. So is every organisation that treats its contractors as second-class members, every platform whose design lets some users be excluded because of their names or faces, and every profession that keeps out competent practitioners trained elsewhere.

The boundary is not fixed. That is the great lesson of the human story. Over thousands of years, people have moved it outward, from the band to the people, from the people to the nation and, imperfectly and with many reversals, toward strangers on the far side of the planet. Every one of those widenings depended on signals that former strangers could learn to share. Religions that welcomed converts, empires that extended citizenship, nations that absorbed immigrants and institutions that judged people by impartial rules all widened the circle of those who counted as real.

Today a great deal of that line-drawing happens inside companies and on platforms. The designers of a hiring process, a rating system, a merger integration or a verification badge are deciding, in a small but real way, who counts. They can draw the boundary narrowly, around people who share the founders' backgrounds, accents and habits, or broadly, around anyone who shares the organisation's purpose and meets its standards. They can make markers accessible to newcomers or set them up to favour those who already belong. They can design markers that carry real information about trustworthiness or markers that merely signal familiarity.

What leaders can do on Monday

It would be a strange book about practical design that ended without practical suggestions. Here are a few things any leader can do in the coming week.

Ask a newcomer what they had to learn in order to be taken seriously. Their answer will reveal markers you have stopped seeing.

Pick one marker your organisation relies on heavily, such as a credential, a rating or a title, and ask whether it still predicts what it is supposed to signal. If no one knows, find out.

Look at the boundaries in your organisation. Who is treated as a full member, and who is not? Is each distinction necessary?

Before your next change to a name, logo, title or ritual, ask who carries the marker you are changing and what it means to them.

Check how a stranger could forge your most important markers. Then ask how much it would cost them.

If trust has been broken somewhere in your organisation, ask whether the response so far has stopped the outflow, told the truth, introduced verifiable markers and committed to a sustained inflow. If any of these steps is missing, start there.

The ordinary miracle

None of this makes the work of leaders easy. Markers are subtle, emotional and constantly evolving. They belong partly to the communities that carry them, and they resist being designed from above. Leaders who try to manage them will make mistakes. But leaders who ignore them will make bigger mistakes, because markers are already shaping how their organisations work, whether anyone is paying attention or not.

The chimpanzees of Gombe could not imagine cooperation with a stranger from the next valley. Human beings have built cities, markets, corporations and global networks on precisely that kind of cooperation. We did it not by knowing one another, which is impossible at scale, but by learning to recognise one another through signals we share. The subway carriage, the marketplace, the open-source repository and the stranger's spare room are all monuments to that achievement.

They are also fragile. Markers can be forged, drained, inflated and abandoned. The trust they carry can be lost far faster than it was built. The task of leaders, founders and platform builders is to tend that trust with the seriousness it deserves: to design markers that tell the truth, to draw boundaries that welcome the strangers a society needs, and to repair the signals when they break. That is how trust outgrows kinship. It is also how a society decides, again and again, who its real people are.

A Note on Sources and Method

This book is a deliberate conversation among three works: Mark W. Moffett's *The Human Swarm: How Our Societies Arise, Thrive, and Fall* (2019), Eric D. Beinhocker's *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics* (2006) and Donella H. Meadows's *Thinking in Systems: A Primer* (2008), edited by Diana Wright and published after Meadows's death. Their ideas are presented in paraphrase. Where the text attributes an argument to one of them, it describes the argument in my own words rather than quoting, except where a short quotation has been checked against a published source.

The three books supply the frame. The cases supply the evidence. I chose cases that are well documented in contemporary reporting, court records, regulatory filings, peer-reviewed research or the published accounts of the people involved. Where a story is widely told but poorly supported, such as some famous impostor tales and a marketing legend about a car's name, the text says so. Where research findings are contested or have failed to replicate, as with Dunbar's number, the Robbers Cave experiment, the coalition interpretation of the Maghribi traders and the early studies of enclothed cognition, the text presents the dispute rather than choosing the more convenient side.

Figures are given as they appear in the sources cited in the notes, rounded where precision would be misleading. Where sources disagree, the text uses ranges or qualifying language. Where I could not confirm a specific detail, I wrote around it rather than guessing. The notes that follow list the principal sources for each chapter. They are intended as a guide for readers who want to go further, not an exhaustive bibliography.

The analogies at the heart of the book, between ant colonies and corporations, hunter-gatherer peoples and platforms, are offered as structural comparisons. They identify shared problems of cooperating beyond personal acquaintance and the families of solutions those problems tend to produce. They are not claims that companies are tribes, that employees behave like insects or that the findings of one field transfer automatically to another. Readers should weigh the comparisons as I have tried to: as tools for seeing, to be tested against the particular organisation in front of them.

Notes and Sources

Introduction

Mark W. Moffett, *The Human Swarm: How Our Societies Arise, Thrive, and Fall* (New York: Basic Books, 2019). Eric D. Beinhocker, *The Origin of Wealth: Evolution, Complexity, and the Radical Remaking of Economics* (Boston: Harvard Business School Press, 2006). Donella H. Meadows, *Thinking in Systems: A Primer*, ed. Diana Wright (White River Junction, VT: Chelsea Green, 2008). Meadows's hierarchy of leverage points appears in her essay "Leverage Points: Places to Intervene in a System" (Hartland, VT: The Sustainability Institute, 1999) and in *Thinking in Systems.*

Chapter 1: The Forest and the Carriage

On Gombe and the Kasekela–Kahama conflict: Jane Goodall, *The Chimpanzees of Gombe: Patterns of Behavior* (Cambridge, MA: Belknap Press, 1986). On Ngogo: John C. Mitani, David P. Watts and Sylvia J. Amsler, "Lethal Intergroup Aggression Leads to Territorial Expansion in Wild Chimpanzees," *Current Biology* 20 (2010). On the comparative analysis of lethal aggression: Michael L. Wilson et al., "Lethal Aggression in Pan Is Better Explained by Adaptive Strategies than Human Impacts," *Nature* 513 (2014). On elephant vocal recognition: Karen McComb et al., "Unusually Extensive Networks of Vocal Recognition in African Elephants," *Animal Behaviour* 59 (2000). On baboon social knowledge: Dorothy L. Cheney and Robert M. Seyfarth, *Baboon Metaphysics* (Chicago: University of Chicago Press, 2007). Individual-recognition and anonymous societies: Moffett, *The Human Swarm.*

Chapter 2: Billions Without a Quarrel

On the Argentine ant supercolony in California: Neil D. Tsutsui et al., "Reduced Genetic Variation and the Success of an Invasive Species," *Proceedings of the National Academy of Sciences* 97 (2000). On the intercontinental megacolony: Eiriku Sunamura et al., "Intercontinental Union of Argentine Ants," *Insectes Sociaux* 56 (2009). On border conflict in southern California: Melissa L. Thomas et al., "When Supercolonies Collide," *Molecular Ecology* 15 (2006). On diet and recognition: Dangsheng Liang and Jules Silverman, "'You Are What You Eat': Diet Modifies Cuticular Hydrocarbons and Nestmate Recognition in the Argentine Ant," *Naturwissenschaften* 87 (2000). On

Phengaris acoustic mimicry: Francesca Barbero et al., "Queen Ants Make Distinctive Sounds That Are Mimicked by a Butterfly Social Parasite," *Science* 323 (2009). Marauder ants and leafcutters: Moffett, *The Human Swarm*; Bert Hölldobler and E. O. Wilson, *The Superorganism* (New York: Norton, 2009).

Chapter 3: Three Fingers in a Basement

Inglourious Basterds, written and directed by Quentin Tarantino (2009). Judges 12:5–6. On the 1937 massacre: Richard Lee Turits, "A World Destroyed, a Nation Imposed: The 1937 Haitian Massacre in the Dominican Republic," *Hispanic American Historical Review* 82 (2002). On white-crowned sparrow dialects: Peter Marler and Miwako Tamura, "Culturally Transmitted Patterns of Vocal Behavior in Sparrows," *Science* 146 (1964). On sperm whale vocal clans: Luke Rendell and Hal Whitehead, "Vocal Clans in Sperm Whales," *Proceedings of the Royal Society B* 270 (2003). On perceptual narrowing in speech: Janet F. Werker and Richard C. Tees, "Cross-Language Speech Perception," *Infant Behavior and Development* 7 (1984). Markers in general: Moffett, *The Human Swarm*.

Chapter 4: Born Reading Markers

David J. Kelly et al., "Three-Month-Olds, but Not Newborns, Prefer Own-Race Faces," *Developmental Science* 8 (2005). Katherine D. Kinzler, Emmanuel Dupoux and Elizabeth S. Spelke, "The Native Language of Social Cognition," *Proceedings of the National Academy of Sciences* 104 (2007). Katherine D. Kinzler et al., "Accent Trumps Race in Guiding Children's Social Preferences," *Social Cognition* 27 (2009). Neha Mahajan and Karen Wynn, "Origins of 'Us' versus 'Them'," *Cognition* 124 (2012). J. Kiley Hamlin et al., "Not Like Me = Bad: Infants Prefer Those Who Harm Dissimilar Others," *Psychological Science* 24 (2013). Henri Tajfel, "Experiments in Intergroup Discrimination," *Scientific American* 223 (1970). Henri Tajfel and John Turner, "An Integrative Theory of Intergroup Conflict," in *The Social Psychology of Intergroup Relations* (1979). Muzafer Sherif et al., *Intergroup Conflict and Cooperation: The Robbers Cave Experiment* (1961). Gina Perry, *The Lost Boys: Inside Muzafer Sherif's Robbers Cave Experiment* (Melbourne: Scribe, 2018). Richard McElreath, Robert Boyd and Peter J. Richerson, "Shared Norms and the Evolution of Ethnic Markers," *Current Anthropology* 44 (2003).

Chapter 5: How Many People Can You Know?

R. I. M. Dunbar, "Neocortex Size as a Constraint on Group Size in Primates," *Journal of Human Evolution* 22 (1992). Robin Dunbar, *Grooming, Gossip and the Evolution of Language* (London: Faber, 1996). W.-X. Zhou, D. Sornette, R. A. Hill and R. I. M. Dunbar, "Discrete Hierarchical Organization of Social Group Sizes," *Proceedings of the Royal Society B* 272 (2005). R. A. Hill and R. I. M. Dunbar, "Social Network Size in Humans," *Human Nature* 14 (2003). Bruno Gonçalves, Nicola Perra and Alessandro Vespignani, "Modeling Users' Activity on Twitter Networks: Validation of Dunbar's Number," *PLoS*

ONE 6 (2011). R. I. M. Dunbar, "Do Online Social Media Cut Through the Constraints That Limit the Size of Offline Social Networks?" *Royal Society Open Science* 3 (2016). Patrik Lindenfors, Andreas Wartel and Johan Lind, "'Dunbar's Number' Deconstructed," *Biology Letters* 17 (2021). Christopher McCarty et al., "Comparing Two Methods for Estimating Network Size," *Human Organization* 60 (2001). Malcolm Gladwell, *The Tipping Point* (Boston: Little, Brown, 2000).

Chapter 6: Real People

On the Jahai and names meaning "people": Moffett, *The Human Swarm.* On hxaro: Polly Wiessner, "Risk, Reciprocity and Social Influences on !Kung San Economics," in *Politics and History in Band Societies*, ed. Eleanor Leacock and Richard Lee (Cambridge: Cambridge University Press, 1982). Benedict Anderson, *Imagined Communities* (London: Verso, 1983). Eugen Weber, *Peasants into Frenchmen* (Stanford: Stanford University Press, 1976). Ernest Renan, "Qu'est-ce qu'une nation?" (lecture, 1882). On the passport: John Torpey, *The Invention of the Passport* (Cambridge: Cambridge University Press, 2000). Peter Singer, *The Expanding Circle* (Oxford: Clarendon Press, 1981).

Chapter 7: Trust as a Stock

Stocks, flows, feedback and delays: Meadows, *Thinking in Systems.* The Dutch electricity meter example: Meadows, "Leverage Points." Alain Cohn, Michel André Maréchal, David Tannenbaum and Christian Lukas Zünd, "Civic Honesty Around the Globe," *Science* 365 (2019). Roy F. Baumeister et al., "Bad Is Stronger Than Good," *Review of General Psychology* 5 (2001). On Tylenol: contemporary reporting in the *New York Times* and *Washington Post*, 1982–83, and the many business school case studies of the recall. Goodhart's law: Charles Goodhart, "Problems of Monetary Management: The UK Experience" (1975), as commonly paraphrased.

Chapter 8: Traps in the Marker System

System traps: Meadows, *Thinking in Systems*, chapter 5. Georgios Zervas, Davide Proserpio and John W. Byers, "A First Look at Online Reputation on Airbnb, Where Every Stay Is Above Average," *Marketing Letters* 32 (2021). On Harvard grading: reporting in the *Harvard Crimson*, December 2013. Joseph B. Fuller and Manjari Raman, *Dismissed by Degrees* (Harvard Business School, Accenture and Grads of Life, 2017), drawing on Burning Glass Technologies data. Robert K. Merton, "The Matthew Effect in Science," *Science* 159 (1968). On Volkswagen: US Environmental Protection Agency notice of violation (September 2015) and subsequent court records.

Chapter 9: Wealth Is What a Society Agrees to Count

Beinhocker, *The Origin of Wealth*, chapter 1, for the Yanomamö and New York comparison and the argument that wealth is culturally defined. Milton Friedman, "The Island of Stone Money," Hoover Institution working paper (1991). William Henry Furness III, *The Island of Stone Money: Uap of the Carolines* (Philadelphia: Lippincott, 1910). R. A. Radford, "The Economic Organisation of a P.O.W. Camp," *Economica* 12 (1945). On the potlatch ban: Douglas Cole and Ira Chaikin, *An Iron Hand upon the People: The Law against the Potlatch on the Northwest Coast* (Vancouver: Douglas & McIntyre, 1990).

Chapter 10: Money, the Marker Everyone Carries

On euro banknote design: European Central Bank materials on the first series. On the Spijkenisse bridges: *Dezeen*, 5 June 2013. On Operation Bernhard: Lawrence Malkin, *Krueger's Men* (Boston: Little, Brown, 2006). On the 1923 hyperinflation and the Rentenmark: Adam Fergusson, *When Money Dies* (London: William Kimber, 1975). On Northern Rock: House of Commons Treasury Committee, *The Run on the Rock* (2008). On Silicon Valley Bank: California Department of Financial Protection and Innovation, order of 10 March 2023; Board of Governors of the Federal Reserve System, *Review of the Federal Reserve's Supervision and Regulation of Silicon Valley Bank* (April 2023). Satoshi Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System" (2008). On FTX: trial reporting, United States v. Bankman-Fried (S.D.N.Y. 2023).

Chapter 11: Kinship's Long Shadow

Avner Greif, "Reputation and Coalitions in Medieval Trade," *Journal of Economic History* 49 (1989), and *Institutions and the Path to the Modern Economy* (Cambridge: Cambridge University Press, 2006). Jeremy Edwards and Sheilagh Ogilvie, "Contract Enforcement, Institutions, and Social Capital: The Maghribi Traders Reappraised," *Economic History Review* 65 (2012). Jessica L. Goldberg, *Trade and Institutions in the Medieval Mediterranean* (Cambridge: Cambridge University Press, 2012). Paul R. Milgrom, Douglass C. North and Barry R. Weingast, "The Role of Institutions in the Revival of Trade," *Economics and Politics* 2 (1990). Edward C. Banfield, *The Moral Basis of a Backward Society* (Glencoe, IL: Free Press, 1958). Francis Fukuyama, *Trust* (New York: Free Press, 1995). Joseph Henrich, *The WEIRDest People in the World* (New York: Farrar, Straus and Giroux, 2020). Raymond Fisman and Edward Miguel, "Corruption, Norms, and Legal Enforcement: Evidence from Diplomatic Parking Tickets," *Journal of Political Economy* 115 (2007). Social technologies: Beinhocker, *The Origin of Wealth*.

Chapter 12: Differentiate, Select, Amplify

The evolutionary algorithm and physical and social technologies: Beinhocker, *The Origin of Wealth*. On Shigetaka Kurita's emoji: the Museum of Modern Art, New York, which acquired the original set in 2016. Oxford Dictionaries, Word of the Year announcement, November 2015. On the hashtag: Chris Messina's post on Twitter, 23 August 2007. On Intel Inside: Intel corporate histories of the 1991 campaign. On the 2012 hoodie remark: contemporary coverage of comments by the analyst Michael Pachter.

Chapter 13: Status, Fairness and the Instinct Beneath the Spreadsheet

Werner Güth, Rolf Schmittberger and Bernd Schwarze, "An Experimental Analysis of Ultimatum Bargaining," *Journal of Economic Behavior and Organization* 3 (1982). Joseph Henrich et al., eds., *Foundations of Human Sociality* (Oxford: Oxford University Press, 2004). Sarah F. Brosnan and Frans B. M. de Waal, "Monkeys Reject Unequal Pay," *Nature* 425 (2003). Michael Marmot, *The Status Syndrome* (London: Bloomsbury, 2004). David Card, Alexandre Mas, Enrico Moretti and Emmanuel Saez, "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," *American Economic Review* 102 (2012). Thorstein Veblen, *The Theory of the Leisure Class* (1899). Michael Spence, "Job Market Signaling," *Quarterly Journal of Economics* 87 (1973). Amotz Zahavi, "Mate Selection: A Selection for a Handicap," *Journal of Theoretical Biology* 53 (1975).

Chapter 14: The Badge and the Brand

On hallmarking: The Goldsmiths' Company, histories of the assay office. On the Bass red triangle: UK Intellectual Property Office historical material on trade mark number one. On New Coke: The Coca-Cola Company's published history of the 1985 episode; Thomas Oliver, *The Real Coke, the Real Story* (New York: Random House, 1986). On Tropicana: *Advertising Age* reporting, April 2009. On Gap: contemporary reporting, October 2010. On the Chevrolet Nova story: Snopes and other fact-checking reviews. On Ratners: contemporary reporting of Gerald Ratner's 1991 Institute of Directors speech.

Chapter 15: Credentials and Licences

Morris M. Kleiner and Alan B. Krueger, "Analyzing the Extent and Influence of Occupational Licensing on the Labor Market," *Journal of Labor Economics* 31 (2013). The White House, *Occupational Licensing: A Framework for Policymakers* (July 2015). On Arizona's universal recognition law: Arizona House Bill 2569 (2019). On the white coat ceremony: the Arnold P. Gold Foundation. On Maryland: Office of the Governor

announcement, March 2022. On skills-based hiring in practice: Burning Glass Institute and Harvard Business School, *Skills-Based Hiring: The Long Road from Pronouncements to Practice* (2024). On Open Badges: Mozilla Foundation materials, 2011.

Chapter 16: The Great Impostor

Robert Crichton, *The Great Impostor* (New York: Random House, 1959). Alan C. Logan, *The Greatest Hoax on Earth: Catching Truth, While We Can* (2020). John Carreyrou, *Bad Blood: Secrets and Lies in a Silicon Valley Startup* (New York: Knopf, 2018), and his reporting in the *Wall Street Journal* from October 2015. United States v. Holmes (N.D. Cal.), verdict January 2022. Declan Walsh, "Fake Diplomas, Real Cash: Pakistani Company Axact Reaps Millions," *New York Times*, 17 May 2015. On Anna Sorokin: trial reporting, New York State Supreme Court, 2019. Pauline Rose Clance and Suzanne Ament Imes, "The Imposter Phenomenon in High Achieving Women," *Psychotherapy: Theory, Research and Practice* 15 (1978).

Chapter 17: Uniforms and Dress Codes

On the Metropolitan Police uniform: Clive Emsley, *The English Police: A Political and Social History* (London: Longman, 1996). Leonard Bickman, "The Social Power of a Uniform," *Journal of Applied Social Psychology* 4 (1974). On the NHS Scotland national uniform: Scottish Government Health Directorates bulletin (2009); *Nursing Times*, 9 December 2009; NHS Scotland national uniform policy. Christopher M. Petrilli et al., "Understanding Patient Preference for Physician Attire," *BMJ Open* 8 (2018). Hajo Adam and Adam D. Galinsky, "Enclothed Cognition," *Journal of Experimental Social Psychology* 48 (2012); Daniel M. Burns et al., "An Old Task in New Clothes," *Journal of Experimental Social Psychology* 80 (2019). Walter Isaacson, *Steve Jobs* (New York: Simon & Schuster, 2011). EEOC v. Abercrombie & Fitch Stores, Inc., 575 U.S. 768 (2015).

Chapter 18: The Dialect of the Office

Sameer B. Srivastava, Amir Goldberg, V. Govind Manian and Christopher Potts, "Enculturation Trajectories: Language, Cultural Adaptation, and Individual Outcomes in Organizations," *Management Science* 64 (2018). Cristian Danescu-Niculescu-Mizil et al., "No Country for Old Members: User Lifecycle and Linguistic Change in Online Communities," *Proceedings of the 22nd International World Wide Web Conference* (2013). Amazon's Leadership Principles are published on the company's website. On the Netflix culture presentation: Patty McCord, *Powerful* (Silicon Guild, 2017).

Chapter 19: Markers at a Distance

Thomas J. Allen, *Managing the Flow of Technology* (Cambridge, MA: MIT Press, 1977). Longqi Yang et al., "The Effects of Remote Work on Collaboration Among Information Workers," *Nature Human Behaviour* 6 (2022). Nicholas Bloom, James Liang, John

Roberts and Zhichun Jenny Ying, "Does Working from Home Work? Evidence from a Chinese Experiment," *Quarterly Journal of Economics* 130 (2015). Nicholas Bloom, Ruobing Han and James Liang, "Hybrid Working from Home Improves Retention without Damaging Performance," *Nature* 630 (2024). Debra Meyerson, Karl E. Weick and Roderick M. Kramer, "Swift Trust and Temporary Groups," in *Trust in Organizations*, ed. Roderick M. Kramer and Tom R. Tyler (Thousand Oaks, CA: Sage, 1996). Sirkka L. Jarvenpaa and Dorothy E. Leidner, "Communication and Trust in Global Virtual Teams," *Journal of Computer-Mediated Communication* 3 (1998). The GitLab Handbook is published online.

Chapter 20: When Two Tribes Sign a Deal

Bill Vlasic and Bradley A. Stertz, *Taken for a Ride: How Daimler-Benz Drove Off with Chrysler* (New York: William Morrow, 2000). On the 2007 Cerberus transaction: Daimler AG announcements, May 2007. On AOL Time Warner: Nina Munk, *Fools Rush In* (New York: HarperBusiness, 2004); AOL Time Warner annual report for 2002. Ed Catmull with Amy Wallace, *Creativity, Inc.* (New York: Random House, 2014). Philippe C. Haspeslagh and David B. Jemison, *Managing Acquisitions* (New York: Free Press, 1991). Clayton M. Christensen, Richard Alton, Curtis Rising and Andrew Waldeck, "The Big Idea: The New M&A Playbook," *Harvard Business Review* (March 2011).

Chapter 21: Why Culture Resists Redesign

Susan Fowler, "Reflecting on One Very, Very Strange Year at Uber," blog post, 19 February 2017; Covington & Burling, recommendations to the Uber board (June 2017). Edgar H. Schein, *Organizational Culture and Leadership* (San Francisco: Jossey-Bass, multiple editions from 1985). Satya Nadella, *Hit Refresh* (New York: Harper Business, 2017). Carol S. Dweck, *Mindset* (New York: Random House, 2006). Google's 2004 founders' letter in its registration statement. Policy resistance and paradigms: Meadows, *Thinking in Systems* and "Leverage Points."

Chapter 22: Reputation Engines

Pierre Omidyar, "Founder's Letter," eBay, 26 February 1996. Paul Resnick, Richard Zeckhauser, John Swanson and Kate Lockwood, "The Value of Reputation on eBay: A Controlled Experiment," *Experimental Economics* 9 (2006). Paul Resnick and Richard Zeckhauser, "Trust Among Strangers in Internet Transactions," in *The Economics of the Internet and E-Commerce*, ed. Michael R. Baye (Amsterdam: Elsevier, 2002). Michael Luca and Georgios Zervas, "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud," *Management Science* 62 (2016). US Federal Trade Commission, final rule on consumer reviews and testimonials (August 2024). On Stack Overflow activity

after ChatGPT: Maria del Rio-Chanona, Nadzeya Laurentsyeva and Johannes Wachs, "Large Language Models Reduce Public Knowledge Sharing on Online Q&A Platforms," *PNAS Nexus* 3 (2024).

Chapter 23: A Stranger's Spare Room

Leigh Gallagher, *The Airbnb Story* (Boston: Houghton Mifflin Harcourt, 2017). Joe Gebbia, "How Airbnb Designs for Trust," TED talk, 2016. Benjamin Edelman, Michael Luca and Dan Svirsky, "Racial Discrimination in the Sharing Economy: Evidence from a Field Experiment," *American Economic Journal: Applied Economics* 9 (2017). Laura W. Murphy, *Airbnb's Work to Fight Discrimination and Build Inclusion* (September 2016). BlaBlaCar with NYU Stern, *Entering the Trust Age* (2016).

Chapter 24: Commons Without Kin

Elinor Ostrom, *Governing the Commons* (Cambridge: Cambridge University Press, 1990). Garrett Hardin, "The Tragedy of the Commons," *Science* 162 (1968). Aaron Halfaker, R. Stuart Geiger, Jonathan T. Morgan and John Riedl, "The Rise and Decline of an Open Collaboration System," *American Behavioral Scientist* 57 (2013). On the Developer Certificate of Origin: Linux kernel documentation, "Submitting Patches." Andres Freund, message to the oss-security mailing list, 29 March 2024; CVE-2024-3094. Eric S. Raymond, "Homesteading the Noosphere" (1998), collected in *The Cathedral and the Bazaar* (Sebastopol, CA: O'Reilly, 1999). Nadia Eghbal, *Roads and Bridges* (Ford Foundation, 2016) and *Working in Public* (San Francisco: Stripe Press, 2020). On Heartbleed: CVE-2014-0160 and the Linux Foundation's Core Infrastructure Initiative.

Chapter 25: Counterfeit Belonging

On paid verification and impersonation, November 2022: contemporary reporting and Eli Lilly's public statement. On the Hong Kong deepfake case: Hong Kong Police briefing, February 2024; Arup statement, May 2024. Catherine Stupp, "Fraudsters Used AI to Mimic CEO's Voice in Unusual Cybercrime Case," *Wall Street Journal*, 30 August 2019. FBI Internet Crime Complaint Center, annual Internet Crime Reports. Onur Varol et al., "Online Human-Bot Interactions: Detection, Estimation, and Characterization," *Proceedings of the International AAAI Conference on Web and Social Media* (2017). Coalition for Content Provenance and Authenticity, technical specifications.

Chapters 26 to 28 and Conclusion

The audit, design and repair methods draw on the sources above, especially Meadows, *Thinking in Systems*, for stocks, flows, traps and leverage points; Beinhocker, *The Origin of Wealth*, for evolutionary design; Moffett, *The Human Swarm*, for markers and the boundaries of societies; and Ostrom, *Governing the Commons*, for governance principles. On the Wikipedia Teahouse: Jonathan T. Morgan et al., "Tea and Sympathy: Crafting

Positive New User Experiences on Wikipedia," *Proceedings of CSCW* (2013). On the AAdvantage programme: American Airlines corporate history. On the Cravath model: Robert T. Swaine, *The Cravath Firm and Its Predecessors* (New York: privately printed, 1946–48). On the HP garage: California Historical Landmark No. 976.

Glossary

Amplification. In Beinhocker's account of economic evolution, the process by which successful designs spread through imitation, growth and investment. Applied to markers, the copying and standardisation that turn a local signal into a widely recognised one.

Anonymous society. Moffett's term for a society whose members can recognise one another as members without knowing each other personally, because membership is read from shared markers rather than recalled from memory. Ants and humans are the outstanding examples.

Balancing loop. A feedback loop that pushes a system back toward a goal or limit. A seller who improves service after a fall in ratings is part of a balancing loop.

Boundary map. A diagram of the groups that markers create inside and around an organisation, used to judge which distinctions are necessary and which exclude people without good reason.

Cost to forge. How difficult and expensive it is for someone without the underlying quality to acquire a marker. The higher the cost, the more trust the marker can carry.

Credential. A marker issued by an authority to certify that a person has met a standard, such as a degree, licence or certification.

Currency (of a marker). Whether a marker still reflects present reality. A marker that was accurate when issued may lose currency as conditions change.

Delay. The time it takes for a change in one part of a system to affect another. Delays let markers keep signalling trustworthiness after the behaviour behind them has changed.

Differentiation. The generation of variety: new products, plans, organisational forms or markers. Without it, selection has nothing to work on.

Drift to low performance. One of Meadows's system traps, in which standards erode because goals adjust to past performance. In marker systems it often appears as compression, with ratings or grades bunching at the top.

Escalation. A competitive spiral in which each side's attempt to get ahead provokes the other to raise the stakes. Credential inflation is a common form.

Individual-recognition society. A society in which membership depends on every member knowing every other member, as among chimpanzees, wolves, meerkats and elephants. Such societies are limited in size.

Information flow. In Meadows's framework, the channel through which information about the state of a system reaches the people who act on it. Markers are information flows about stocks of trust.

Leverage point. A place in a system where an intervention produces large effects. Meadows ranked paradigms, goals, rules and information flows above parameters such as prices and budgets.

Marker. A signal of membership that lets strangers recognise one another as belonging to the same group. Markers range from flags, uniforms and logos to accents, gestures, food habits, credentials, ratings and digital badges.

Marker audit. A structured review of an organisation's markers, covering inventory, testing, loops, boundaries and priorities, described in Chapter 26.

Meta-marker. A marker that vouches for other markers, such as deposit insurance backing bank brands or a platform guarantee backing user reviews.

Minimal group. A group formed on a trivial or arbitrary basis, used in experiments by Henri Tajfel and others to show how readily people favour members of their own group.

On-ramp. Any mechanism that lets legitimate newcomers acquire the markers of a group quickly enough to take part, such as mentoring, guarantees, introductory badges or welcoming spaces.

Paradigm. The shared assumptions from which a system's goals, rules and structures arise. Meadows placed paradigms near the top of her leverage points.

Policy resistance. A trap in which actors with different goals pull a system back toward its previous state whenever someone tries to change it. Culture change programmes often meet it.

Reinforcing loop. A feedback loop that amplifies change in the direction it is already moving. Trust that enables cooperation, which in turn builds more trust, is a reinforcing loop.

Selection. The process by which some designs succeed and others fail. Markets, social pressure and authorities all select among markers.

Shibboleth. A custom, word or habit that separates members of a group from outsiders, named after the password test in the Book of Judges.

Shifting the burden. A trap in which an intervention relieves symptoms while weakening the system's own capacity to solve the problem. Relying on markers instead of direct judgment can produce it.

Social technology. Beinhocker's term for a way of organising people, such as money, markets, the joint-stock company or the rule of law. Social technologies co-evolve with physical technologies.

Stock. An accumulation that can be measured at a point in time, such as water in a bathtub, money in an account or trust in an institution. Stocks change only through their flows.

Success to the successful. A trap in which winners gain resources that help them win again. Reputation systems that favour established participants show it.

Swift trust. Trust that forms quickly among members of temporary groups on the basis of roles and categories rather than personal history, described by Meyerson, Weick and Kramer.

Tragedy of the commons. A trap in which a shared resource open to all is overused. Unregulated labels such as "natural" can be depleted in the same way.

Trust cascade. A chain in which each party trusts a marker because others appear to have trusted it, so that no one checks the underlying quality independently.