PERSONIX

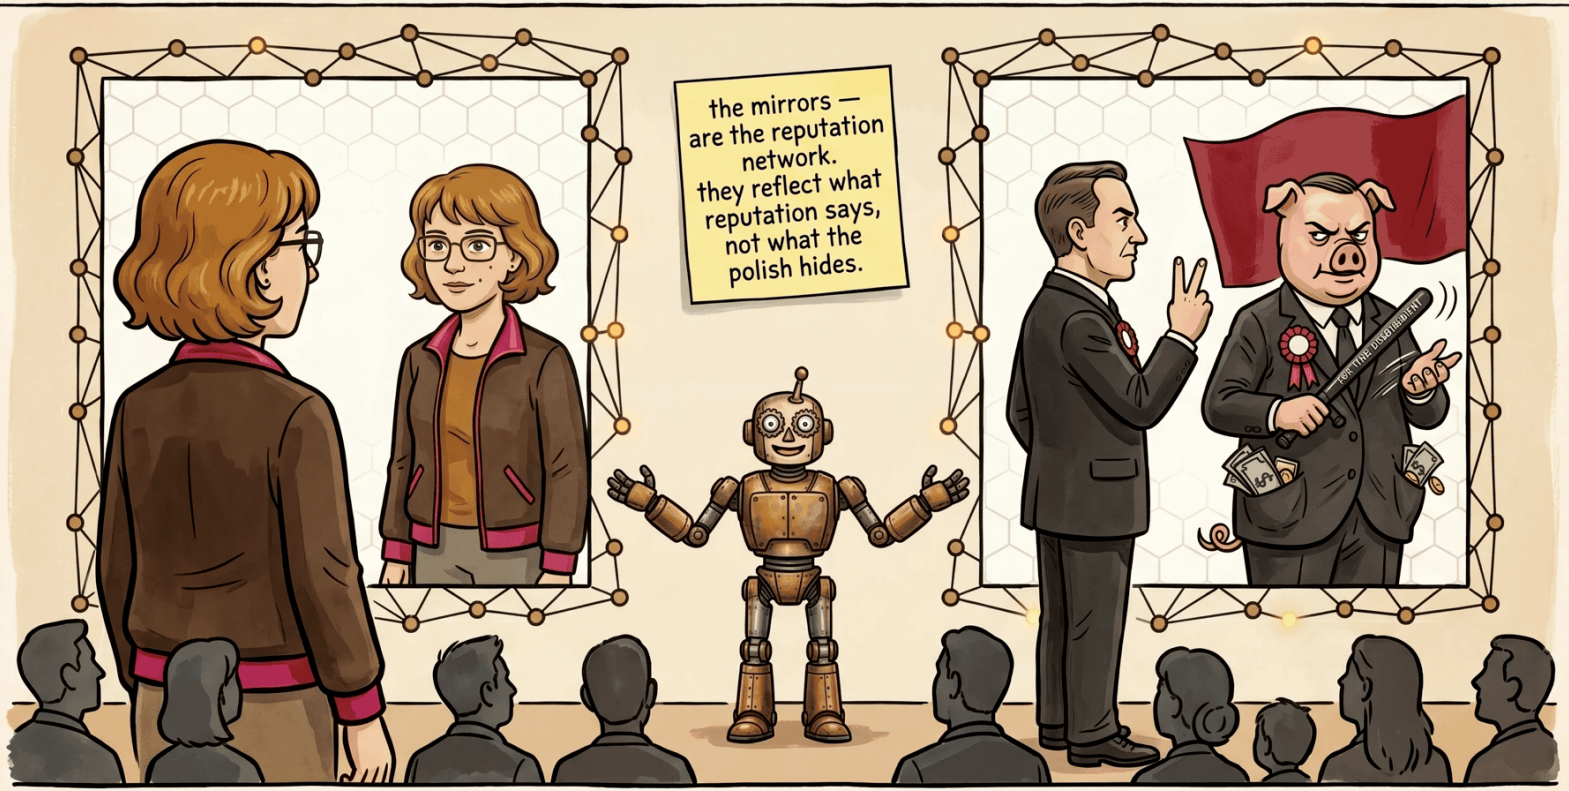

An Uncensorable and Incorruptible Decentralized Reputation Network

an evolutionary successor to the state

Personix — Uncompromising Change

An Uncensorable and Incorruptible Decentralized Reputation Network

Version 6 (2026-05-21)
First drafted June–November 2025; rewritten April–May 2026.

Typeset in LaTeX (Tectonic) using Palatino and Helvetica Neue.
A5 landscape (210 × 148 mm).

Contents

Introduction

Do you feel like you can rely on the state?

I refuse to accept it when an obvious injustice is done to me, my family, or someone in my community (physical harm, fraud, undue restriction, …) and the state responds: yes, it is an injustice, but don't do anything about it, or you'll have an even bigger problem with us. Such a condition is sick, and people who go along with it are accomplices.

The result is mutual alienation, playing at community, playing at statehood, and a deepening break-down of society with the promise of serious conflicts ahead.

Can the state keep up with the pace of technological and social change, or does it prefer to slap on a ban?

I could pretend that I'll treat stagnation, incompetence and corruption as the declared stability of the rule of law and tradition, but I no longer want to. I'd be lying to myself and I'd be lying to you.

Do you feel the condescension of officials and elected representatives when you dare to criticize something you're paying for?

I refuse to consider people in society unworthy of shaping society-wide rules just because they don't excel at the soft skills that are the sole sufficient prerequisite for becoming an elected representative in state-codified institutions today. I have faith that the vast majority can take care of themselves, their loved ones, and their community — and that they see in this a piece of their life philosophy. Most of the people I've come to know around me are like that.

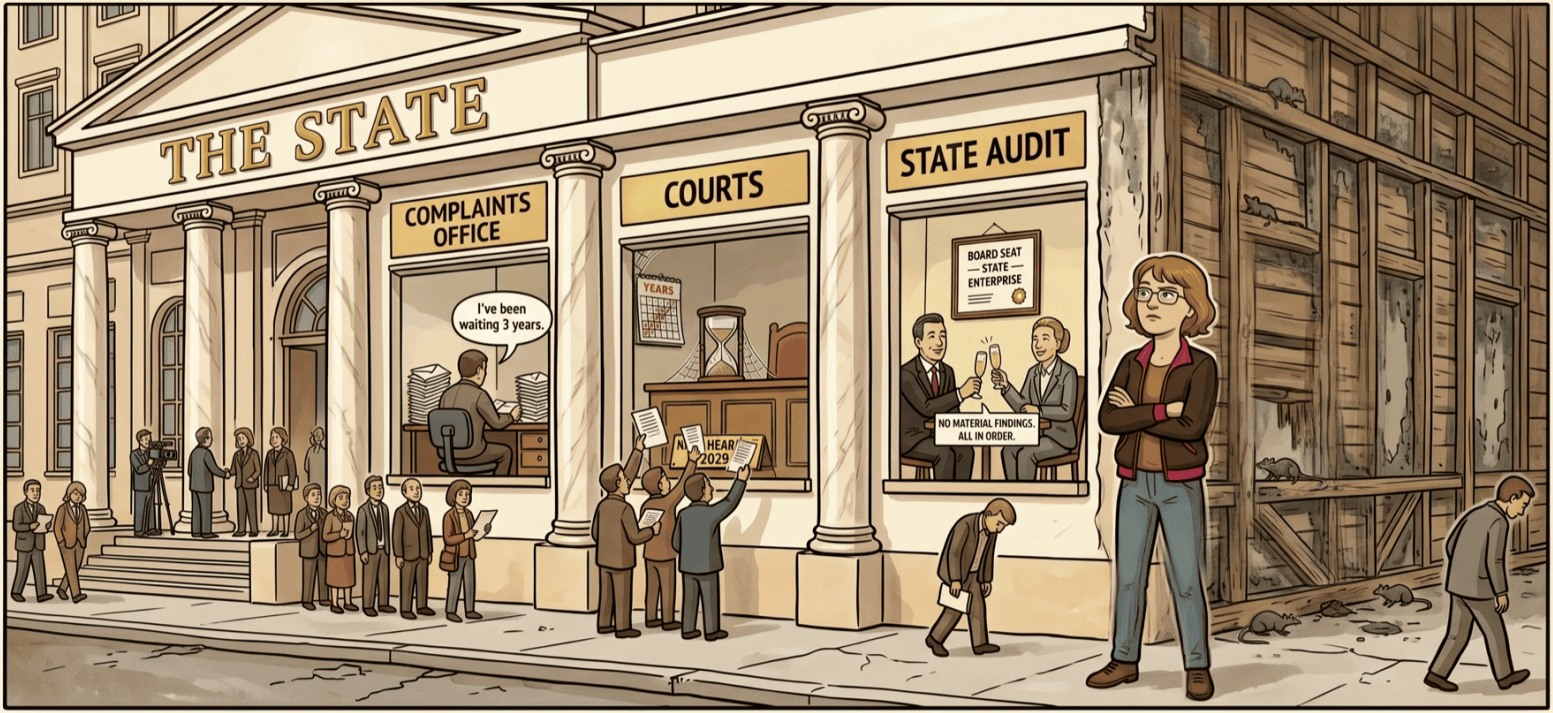

I refuse to accept it.

But I've also met outright assholes who were assholes precisely because they knew how to play the game — the game where the community is not allowed to give them feedback, where the state protects them so that no one disturbs the order that would threaten the very essence of the state: ensuring an easy living for those on the upper floors of the power hierarchy and the hierarchy of its institutions. Feudalism has been replaced by the rule of populists and elitists.

Do you love it when others tell you how to live your life?

I don't accept the current state of affairs as a perfect, final given. I must be responsible in how I conduct myself, yet I'm touchy when someone tells me which right values I should hold. How about you?

What is this book about?

Imagine a reputation social network where anyone who wants to can create their own keys for it (create their identity), get invited by members of their community, or venture in on their own. You build your reputation and draw on useful reputation data about others when you need it. You seek your community's support when injustice happens — you state your case and provide what you have on the matter, and the community takes a look, decides based on relevance whether to pass it on to others, and how to sanction the offender.

This entire short book is about what needs to be under the hood for this to work so simply on the outside. If you stick with it, I promise you that a world that may now seem complex and hard to understand in many ways will open up as more comprehensible and clearer. You'll see that many bad things can be set up so they no longer pay off so unfairly for the bad actors. And it isn't a tool you'd have to be glued to online all the time — it's enough to reach for it when you feel you need to.

I further promise you that the arrangement described here will feel closer to people's real motivations, and that it doesn't need to reshape them in the name of some -ism.

And finally, we'll show how to manage a non-revolutionary transition to a state in which communities once again become fully functional social units and no longer have to just pretend.

Technical article

This book is an explanatory companion to the original technical article *"Decentralized Reputation Network: A Framework for Voluntary and Uncorruptible Societal Organization"* (Pavel Kudrna, 2026). If you are interested in the formal architecture description, the verifier selection algorithm, and the justification of the economic mechanisms, I recommend starting there.

Chapters build on each other

The book is written as a continuous argument — each chapter builds on the previous one. Concepts introduced in one part are used in later parts, including in footnotes. I don't recommend reading out of order.

Who is this book for?

Despite its visually friendly format, the main target audience is the more specialist reader concerned with the principles of possible socio-economic-political arrangements of human society. Reaching a point where I could explain the entire subject matter accessibly to the widest public still requires a long road of building out a full communication strategy.

I welcome suggestions for how to make this little work, or any follow-up revised editions, more comprehensible.

Warning

I don't know who you are, how old you are, or how many broken promises you've had to sit through in your life. I'm done making them here. Instead, I want to warn you that the only outcome of what you learn here may be the realization that there is probably no path where someone else cleans up today's mess for you. I'm only describing a framework/tool that could make the cleanup simpler and more transparent. Based on it, I'll attempt to prepare a social experiment/simulation, fix potential flaws, or bury the idea entirely if it turns out to have fundamental design cracks.

A request

I would deeply appreciate it if, even after finishing this book, you kept an open mind toward the idea and followed where it goes and how it holds up under substantive criticism. I would appreciate even more your later involvement — trying to break the proposed equilibrium principles with your own ideas, and thereby better testing how prepared the system is for real human behavior.

There’s more under the hood. Stay with me.

Without you and your involvement, the idea stays a theory in a drawer. Without you and your contributions and financial donations, we won't be able to push the project toward a real social experiment, and thus toward real readiness. Without you and your substantive criticism, we'll stagnate where we are now, or let ourselves be swallowed by the holy grail of centrists — absolute, irreversible totalitarianism. It's up to us which path we choose, or sit out.

The challenge

We are not exactly putting a light load on your shoulders. You're already bearing the costs of a profligately managed society through occultist-incantatory manuals claiming that society is best run centrally — and I have the nerve to ask you for even more, and to go against these tendencies. I'm aware that we'll have to weather the collapse of the biggest pyramid scheme in history — the pay-as-you-go pension system — without letting several generations starve.

Yet I ask you nonetheless: let us finally hand our children a better world and leave behind a legacy that we pulled off that change. Unleashing economic growth could be a partial solution to the predicament on the horizon. We can then tell our descendants it was a hell of a ride and it wasn't in vain.

The problem

Only freedom balanced by direct responsibility can truly unlock an enormous potential the world has never seen.

A better world is worth the weight.

Come in. The plans are laid out.

Essential Prerequisites for Change

The fundamental prerequisite for change for the better is that the subordinate becomes the superior — or at least an equal partner. This means taking over the most critical agendas that define the relationship of superiority and subordination. These agendas rest on proofs of belonging to the communities that constitute citizens of the state. They are information and its verifiability by an authority with substantial reputation within those communities.

To overcome the current system, I have arrived at five prerequisites that I then explain in the way I understand them. We can think of these prerequisites as the pillars that carry the construction of the system described in the rest of the book. I chose the pillars initially by my own preference and, while reasoning with them, gradually refined how I understand each one into the form below.

Let's walk through these five pillars one by one.

Voluntariness

In our case, voluntariness means that I am not forced to participate in running the decentralized reputation network. I cannot, however, demand a positively framed right that the network not record information about me — whether information about me ends up in the reputation network depends only on me, my behavior, my actions, and my responsibility.

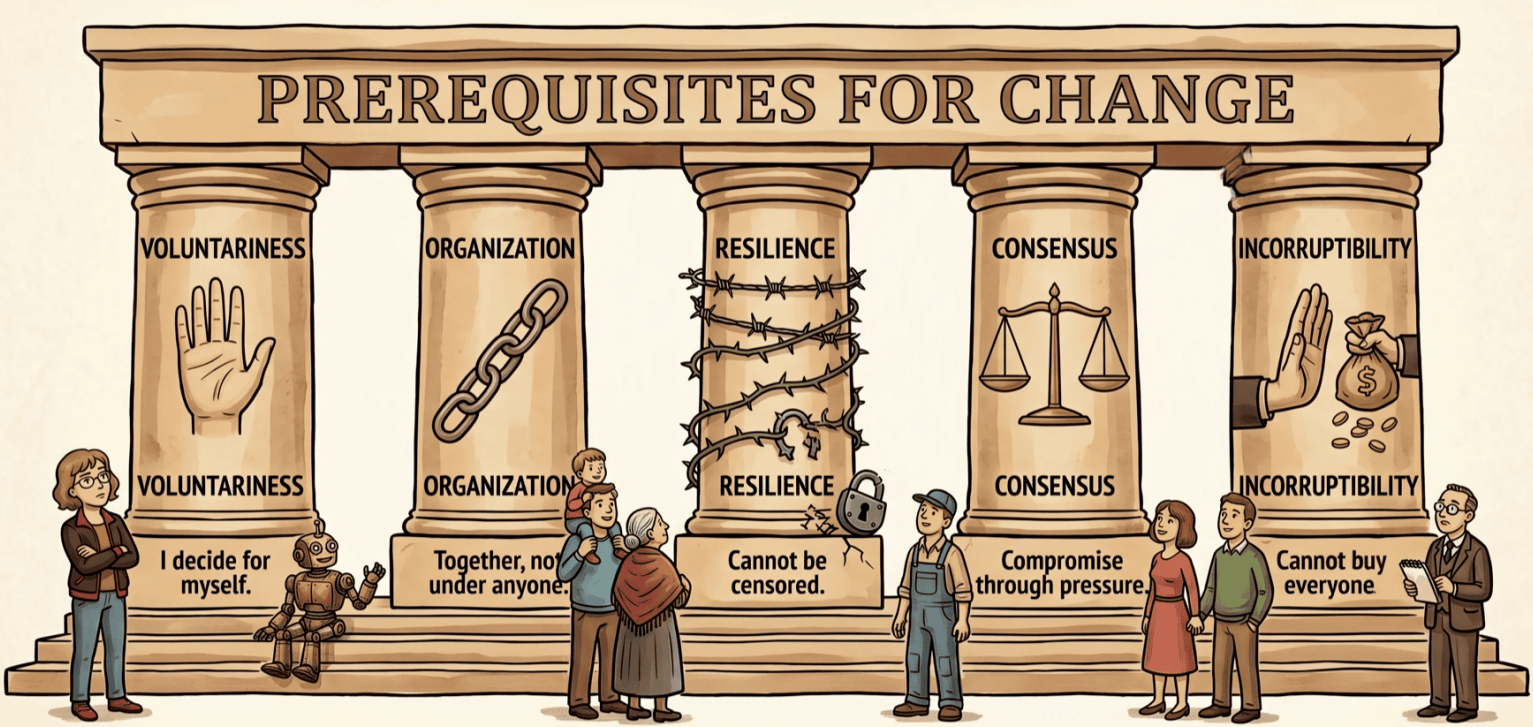

Without these five pillars, every change is just burning time.

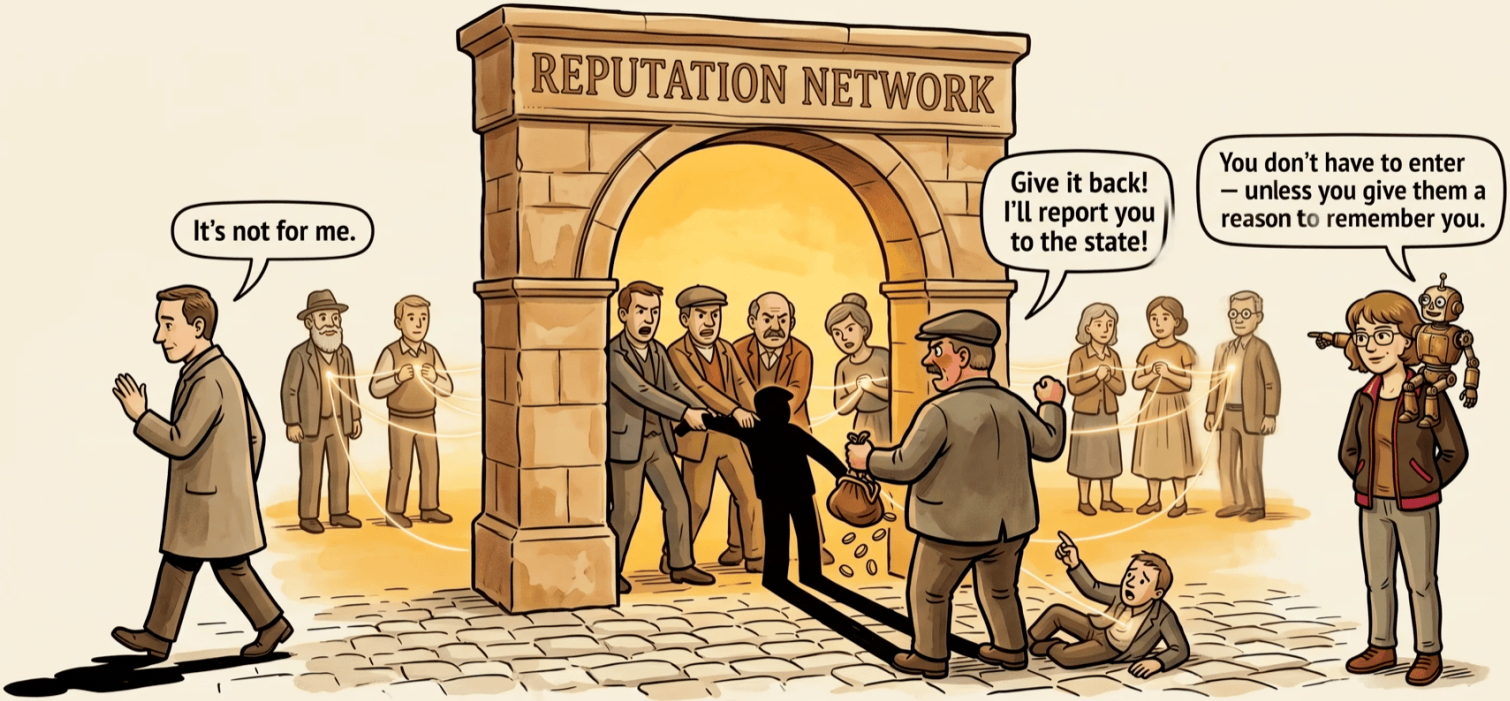

Voluntary — until responsibility catches up with you

Organization

A single individual cannot break the system, which is why the state controls the association of people as a way of preventing activities that could threaten its structure and functioning. It does not matter whether people organize for or against. Through organization it is possible to soften the impact of exemplary punishments when someone dares to change the state against its will.

Let us not set common goals, but we can walk part of the road together. The most natural form is association in communities — but they must have all the tools they need so as not to remain mere playacting at being communities.

Resilience

The chosen path toward changing the state — or toward transitioning to its successor — must not depend on state permission. The tools and the ways they are used must be as uncensorable as possible. In plain terms, the cost of censoring them must be astronomical. Without that, the entire effort is just burning time and life energy.

PREREQUISITE 2 OF 5: ORGANIZATION

An individual cannot break the system. A community can.

PREREQUISITE 3 OF 5: RESILIENCE

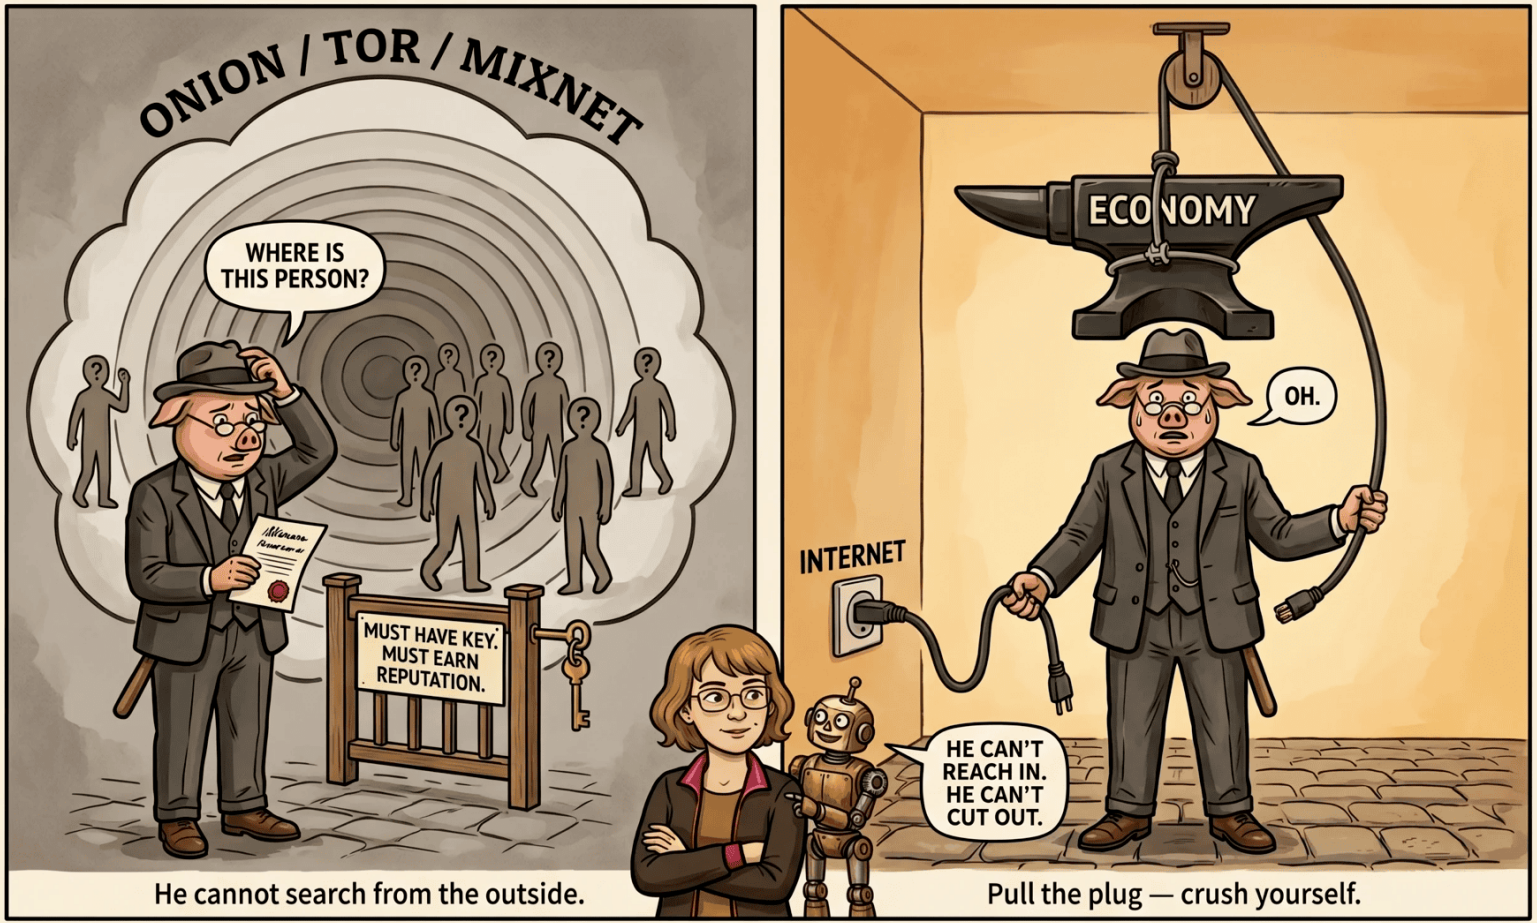

NO PERMISSION TO ASK FOR. NO PLUG TO PULL.

Consensus

The proposed solution must, by its very nature, contain a mechanism for negotiating consensus as well as a repressive component for cases where consensus is violated. It need not be a voluntary win-win agreement. It is enough if circumstances and social pressure ensure compromise across a community and between communities.

STAY IN RANGE — NOT A RADICAL.

You may differ. Just not too far.

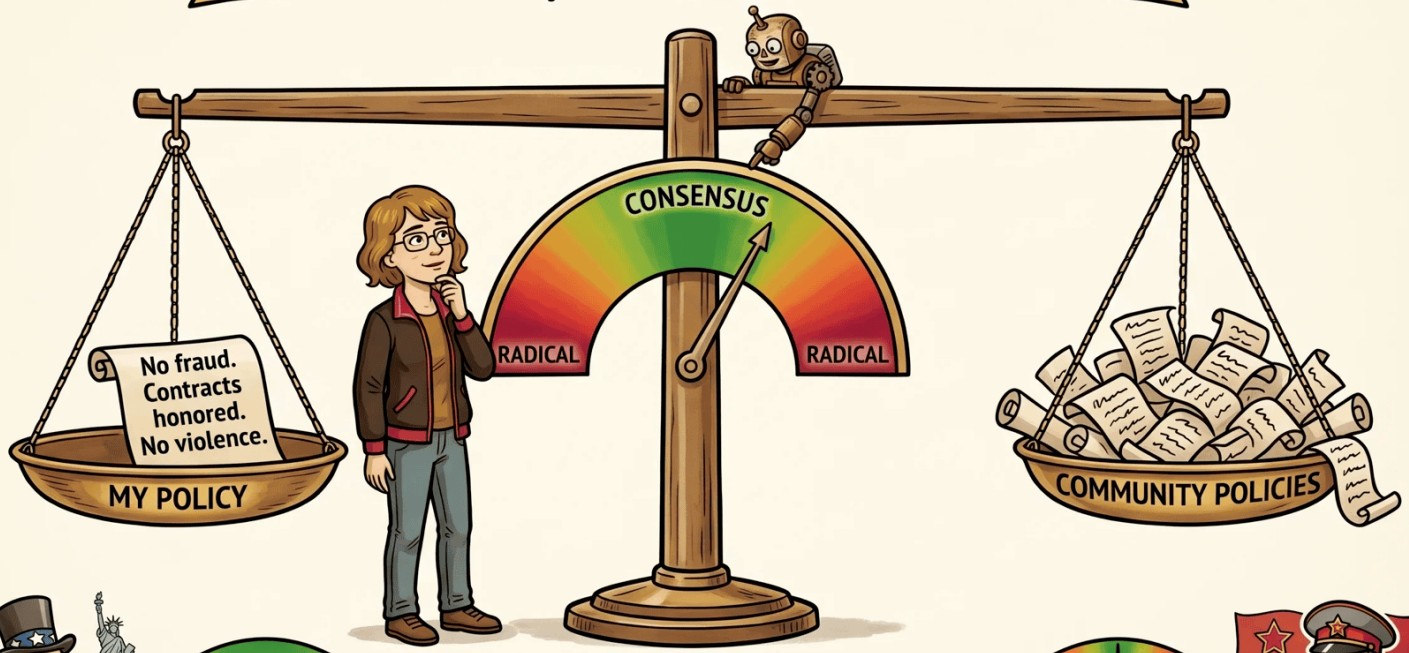

The tolerance zone stretches with the freedom of the culture.

PREREQUISITE 4 OF 5: CONSENSUS

CONSENSUS IS A TOLERANCE ZONE, NOT A UNANIMOUS AGREEMENT.

Incorruptibility

Uncensorability defends the system from the outside. But a network of people can also rot from the inside — an authority gets captured, a declared rule is quietly abandoned, participants are bought off in silence. Uncensorability would be a hollow victory if the people inside could be corrupted without consequence.

Incorruptibility is not the result of a single clever trick. It emerges from **several principles of the proposed design working in concert** — the pressure against corruption is distributed across the architecture, and the individual parts reinforce one another. How exactly this is achieved we will show gradually — in the description of the tool, in verification, and in the chapter on punishment. At the end we will recap everything that has been built up.

Incorruptibility is simultaneously a prerequisite of the design and an emergent property of the result. It is the second defensive line — the one facing inward — and without it, Resilience is only half a shield.

PREREQUISITE 5 OF 5: INCORRUPTIBILITY

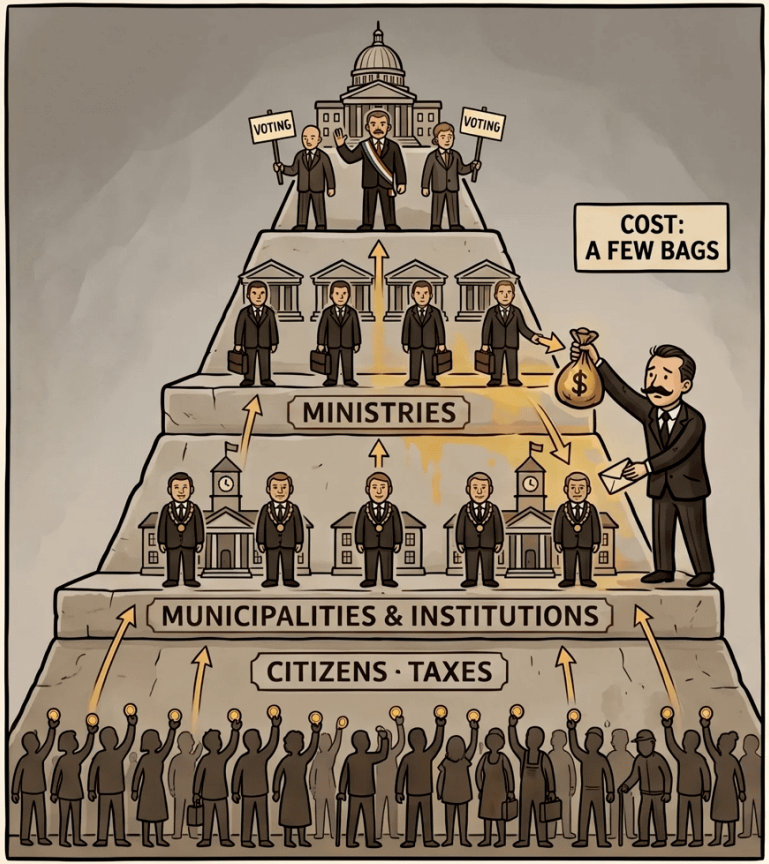

Bribe a few upper floors — the base pays anyway.

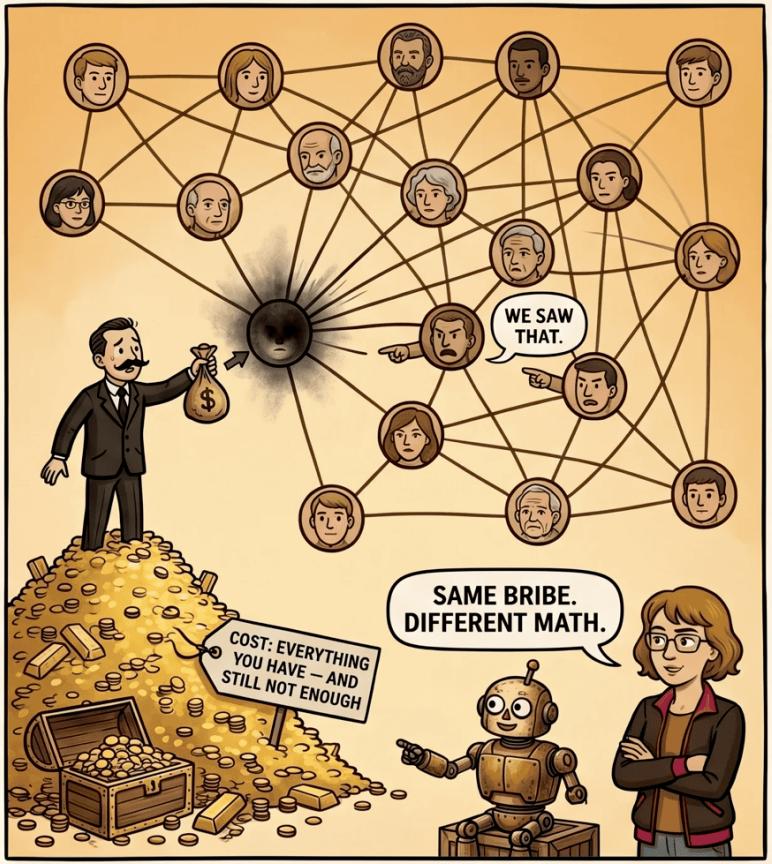

Bribe one — he loses his reputation forever. Bribe all — you go broke first.

CORRUPTION HAS A PRICE. IN THE MESH IT IS TOO HIGH.

The Tool

Reputation-Based Social Network

To bring about change, we need a carefully designed tool. First we will sketch it briefly; in later chapters we will examine each piece in greater detail and add more. Imagine an uncensorable, global, decentralized social network where you could safely create and manage your proxy identity — a so-called Decentralized Identity (DID). A DID is a digital identity that you create and control yourself, without dependence on any central authority. Nobody can take it away or forge it, because it is cryptographically signed with your private key (or keys, via multisig).

> **Note**
>
> One implication is that such an identity could gradually replace state-issued identification documents — but more on that in the chapter on transition.

In such a network, you could report through your identity that someone has caused you harm (and later, potentially, that they have remedied it or been compelled to do so). For this feedback — directed at the originator of the harm — to have value as a relevant source, entering information into the network must cost time, energy, and money — and on top of that, verifiable proof must be produced for others that this is not idle chatter.

YOUR IDENTITY, YOUR KEYS, YOUR RULES

YOU CREATE IT YOURSELF

PRIVATE KEY — signs on your behalf. Only you hold this. Never share.

PUBLIC KEY — verifies your signatures. Visible to the network.

No registration office.
No government form.
No company approval.
You generate it yourself.

YOUR DID DOCUMENT

DID DOCUMENT

PUBLIC KEY:
[cryptographic hash]

NETWORK ADDRESS:
[onion gateway]

VERIFICATION POLICY:
My rules for interactions

CLAIMS:
see right

This is your public identity.
Anyone can read it.
Only YOUR private key can can sign changes to it.

No state can revoke this.

CLAIMS — VERIFIED FACTS ABOUT YOU

CITIZENSHIP:
Czech Republic —
verified by Municipal Notary

EDUCATION:
MSc Computer Science —
verified by University Alumni Network

PROPERTY:
Owner of apartment #42 —
verified by Local Community Registry

PASSIVE WITNESS:
Ignored a wrongdoing —
policy not honoured —
filed by Neighbourhood Witnesses

FRAUD:
Cheated a neighbour out of savings —
filed by Local Community Court

Claims are verified statements.
The good and the bad. Authorities stake their reputation on them.

STATE ID — REVOCABLE:
the state can cancel it anytime.

DID — IRREVOCABLE:
only your private key can sign changes.

Your identity. Your keys. Your rules. No permission needed.

Reading information would be easy and relatively cheap, but creating an individual record would be costly and demanding. Writing would follow a clear protocol, in which computation according to the chosen algorithm strictly determines which DID to ask for verification of the submitted information and how to proceed so that the selected participant processes the information on your behalf, publishes it, and becomes its verifier.

Algorithm vs radicalism

Algorithmic selection of verifiers ensures that non-radical information publishers will, over time, maintain a nearly neutral balance between the costs of published information and rewards for verification.

PUBLISHING COSTS TIME, ENERGY, AND MONEY

WRITING INTO THE NETWORK IS NEVER FREE — READING IS

THREE CHANNELS — ALL THREE ARE REAL

TIME

you wait while the algorithm picks a verifier and the verifier reads your record.

ENERGY

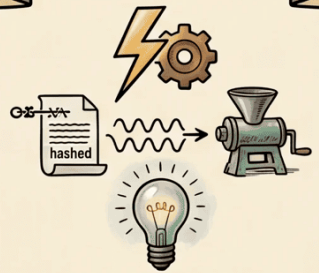

signing, hashing, verifying — every step is real computation and real electricity.

MONEY

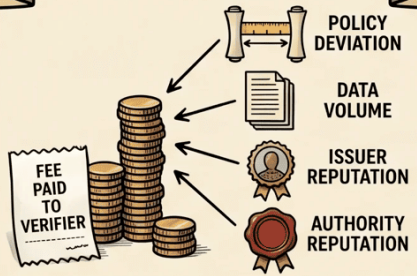

the further your claim sits from the verifier's policy, the larger the record, and the less trusted the issuer or their authority — the higher the fee.

WHAT HAPPENS NEXT

- The three-step verification chain — see **HOW THE ALGORITHM PICKS A VERIFIER, HOW THE VERIFIER ANSWERS, THE ISSUER'S CHOICE**
- Why expensive writing protects reputation — later in the chapter

COST OVER ITERATIONS

EACH RETRY COSTS MORE

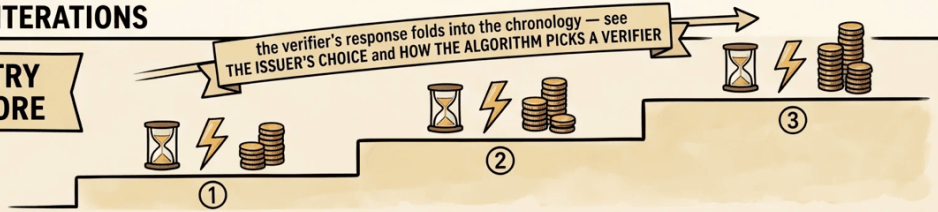

every rejection or re-throw adds to the chronology the algorithm eats — and the next attempt is more expensive in all three channels.

READING ASYMMETRY

READING IS CHEAP

free to query

anyone may read reputation and records — only writing costs. The asymmetry is the design.

WRITING COSTS — READING DOES NOT. THE ASYMMETRY IS THE POINT.

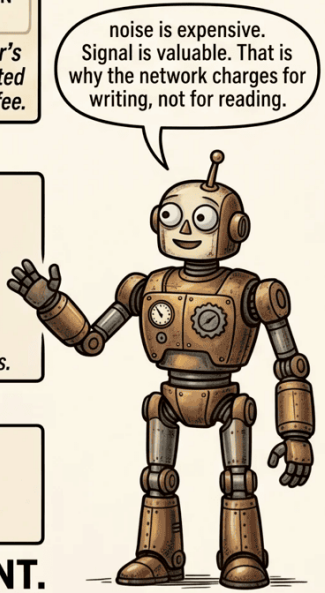

Let us look at how the algorithm selects a verifier.

Algorithm

Algorithmic selection non-deterministically chooses a different verifier (or a set of possible verifiers) for different pieces of information. A hash (a one-way mathematical function that produces a unique "fingerprint" from any input — like a fingerprint of a document) of the complete DID document determines the position on a consistent hash ring and selects verifier candidates.

$$\text{verifier\_DID} = \text{hash}(\text{DID\_document})$$

In plain language: the algorithm takes your entire DID document, computes a fingerprint from it, and that fingerprint determines your verifier.

With the first verifier the algorithm selects, you as the publisher may not succeed — your reputation or declared settings may not meet their requirements. You would algorithmically continue the search for the next one by performing another recursive iteration, which assigns you a further verifier. With each step the "distance" to the target verifier grows, and so does the accompanying metadata that must be published. As the data grows, costs naturally rise (not only because of the initial size of the claim, but also because of the metadata accumulating with each rejection). Credible information passes far more easily than nonsensical whims. It is up to each person how high a price they are willing to bear and how much the record matters to them — radicalism is guaranteed to get expensive.

HOW THE ALGORITHM PICKS A VERIFIER

A FAIR COIN THROWN ACROSS YOUR COMMUNITY

ONE ARTEFACT TO PUBLISH

CHRONOLOGICAL RESPONSES

\+

CHRONOLOGICAL LIST OF PRIOR VERIFIER RESPONSES

CLAIM · POLICY · CLAIM STANDARD · POLICY STANDARD · DID DOCUMENT CHANGE

pick exactly one artefact per call

every kind of record is fed through the same hash — no side chome omonsine irems around a verifier.

empty on the first call — grows with every rejection

every kind of record is fed through the same hash — no side channel lets anyone route information around a verifier.

some authorities exist only to randomise the verifier — that alone is a service worth reputation.

HASH: 7A3F...

the publish path goes: issuer → authority → verifier. No record may skip the authority step.

YOUR COMMUNITY AND ITS VICINITY

every face has the same chance — the hash lands uniformly at random

you cannot choose your verifier — the algorithm throws the colo, and everyone here has an equal chance.

your verifier

A CAN YOU BRUTE-FORCE IT?

i could grind thousands of dralts until one lands on the verifier I want...

SVNIDLLY DIFFERENT

WILDLY DIFFERENT

COMPLETELY DIFFERENT

the hash is deterministic but pre-image-resistant — each small edit throws the pick somewhere unrelated.

B BUT THE AUTHORITY SEALS THE REST

AUTHORITY-SIGNED EVIDENCE ABOUT OTHERS

...but I cannot predict what the authority will sign — and I cannot tell whether she is in league with me or nol.

because the hash eats authority-signed evidence tou, the issuer cannot forecast his own hash. For the community, the question becomes: do we trust this authority — and this issuer-with-this-authority — at all?

WHAT HAPPENS NEXT

- The verifier answers with a price — or rejects with a reason — see HOW THE VERIFIER ANSWERS
- The issuer decides what to do with that anoazy — see THE ISSUER'S CHOICE
- Each retry adds to the cost — as shown *earlier* on PUBLISHING COST

THE ALGORITHM PICKS — YOU DO NOT.

HOW THE VERIFIER ANSWERS

ACCEPT WITH A PRICE — OR REJECT WITH A REASON

the verifier weighs four things before answering.

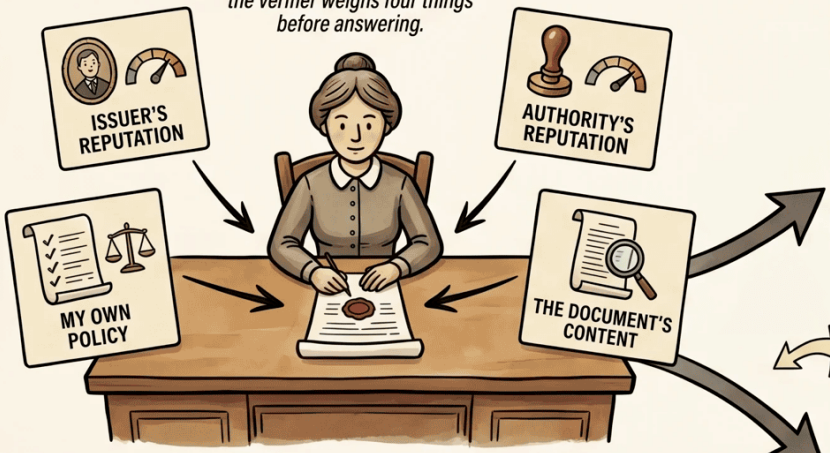

ACCEPT — WITH A PRICE

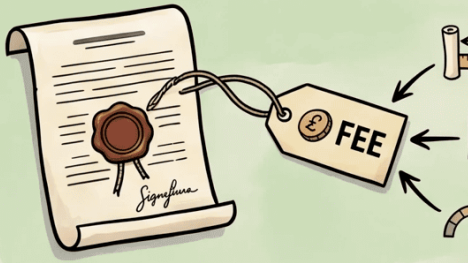

the verifier agrees to publish — but the price reflects how far the claim sits from their policy, how much data it carries, and how trusted the issuer is.

REJECT — WITH A REASON

- the issuer's reputation is too low for me
- the authority behind this record is not one I trust
- the issuer's policy is too far from mine to bridge
- the document itself fails my policy's rules

*the verifier writes down **WHY** — their response feeds back into the next attempt.*

the verifier's response is added to the chronological list — as shown on HOW THE ALGORITHM PICKS A VERIFIER

whatever the verifier says, their answer becomes input to the next call.

WHAT HAPPENS NEXT

- The issuer decides what to do with this answer — see THE ISSUER'S CHOICE
- Each answer shifts the cost of the next attempt — as shown earlier on PUBLISHING COST

every verifier answers in one of these two ways — **accept with a price, or reject with a reason.**

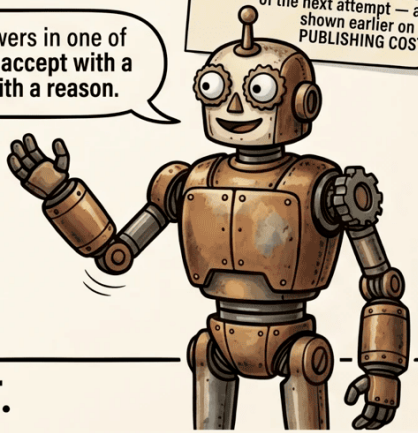

THE VERIFIER EITHER SETS A PRICE — OR TELLS YOU WHY NOT.

Whatever the verifier decides in response to your verification request, the ball is back in the publisher's court: they can accept the verifier's offer for verification services, fold the response into the chronology and try again (more expensively), or walk away and swallow the sunk cost.

To give your information greater weight and a better chance of acceptance with verifiers, you — as a publisher with a stake in the information being issued — could use the services of a **trusted authority**. The authority either rejects the submitted information or accepts it and stakes its good name (reputation) on it. The authority typically requests real-world evidence, verifies it, and classifies it. The output is a protocol of its assessment of the given case at the given time. Think of an authority as a specialist in a certain type of service in both the real and digital world — for example an investigator, an auditor, an insurer, a supplier of a certain class of goods (in essence, any economic actor on the market).

THE ISSUER'S CHOICE

WHAT TO DO WITH THE VERIFIER'S ANSWER

whatever the verifier said, the issuer now has ONE answer on the table.

THREE WAYS OUT

TAKE THE DEAL

ONLY IF ACCEPTED

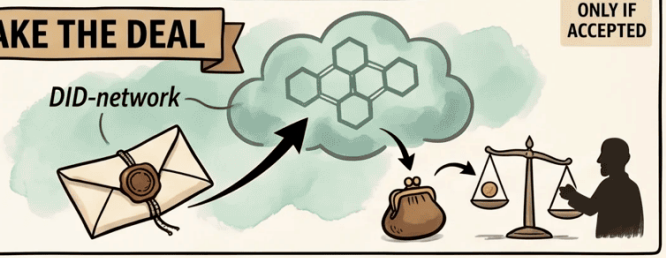

the issuer pays the fee — the record is published into the network. Only available if the verifier accepted.

TRY AGAIN

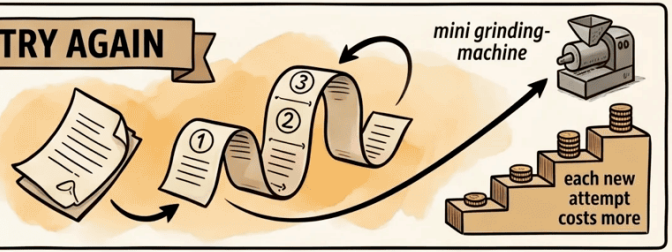

the response is bundled into the chronology — the algorithm throws again. Works whether the verifier accepted at a price the issuer dislikes, or rejected outright.

WALK AWAY

the issuer discards the draft. Nothing is published — but nothing more is paid either.

the verifier decides how — the issuer decides whether. Publishing is a mutual yes.

WHAT HAPPENS NEXT

- Each retry shifts the cost of the next attempt — as shown earlier on PUBLISHING COST
- The verifier's answer itself — as shown earlier on HOW THE VERIFIER ANSWERS

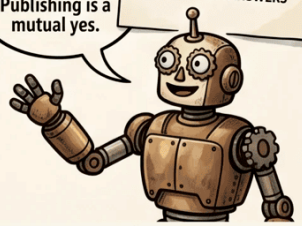

loop back to **HOW THE ALGORITHM PICKS A VERIFIER** — shown earlier

only the retry path loops — accept and walk-away both end the story.

THE VERIFIER PROPOSES – THE ISSUER DISPOSES.

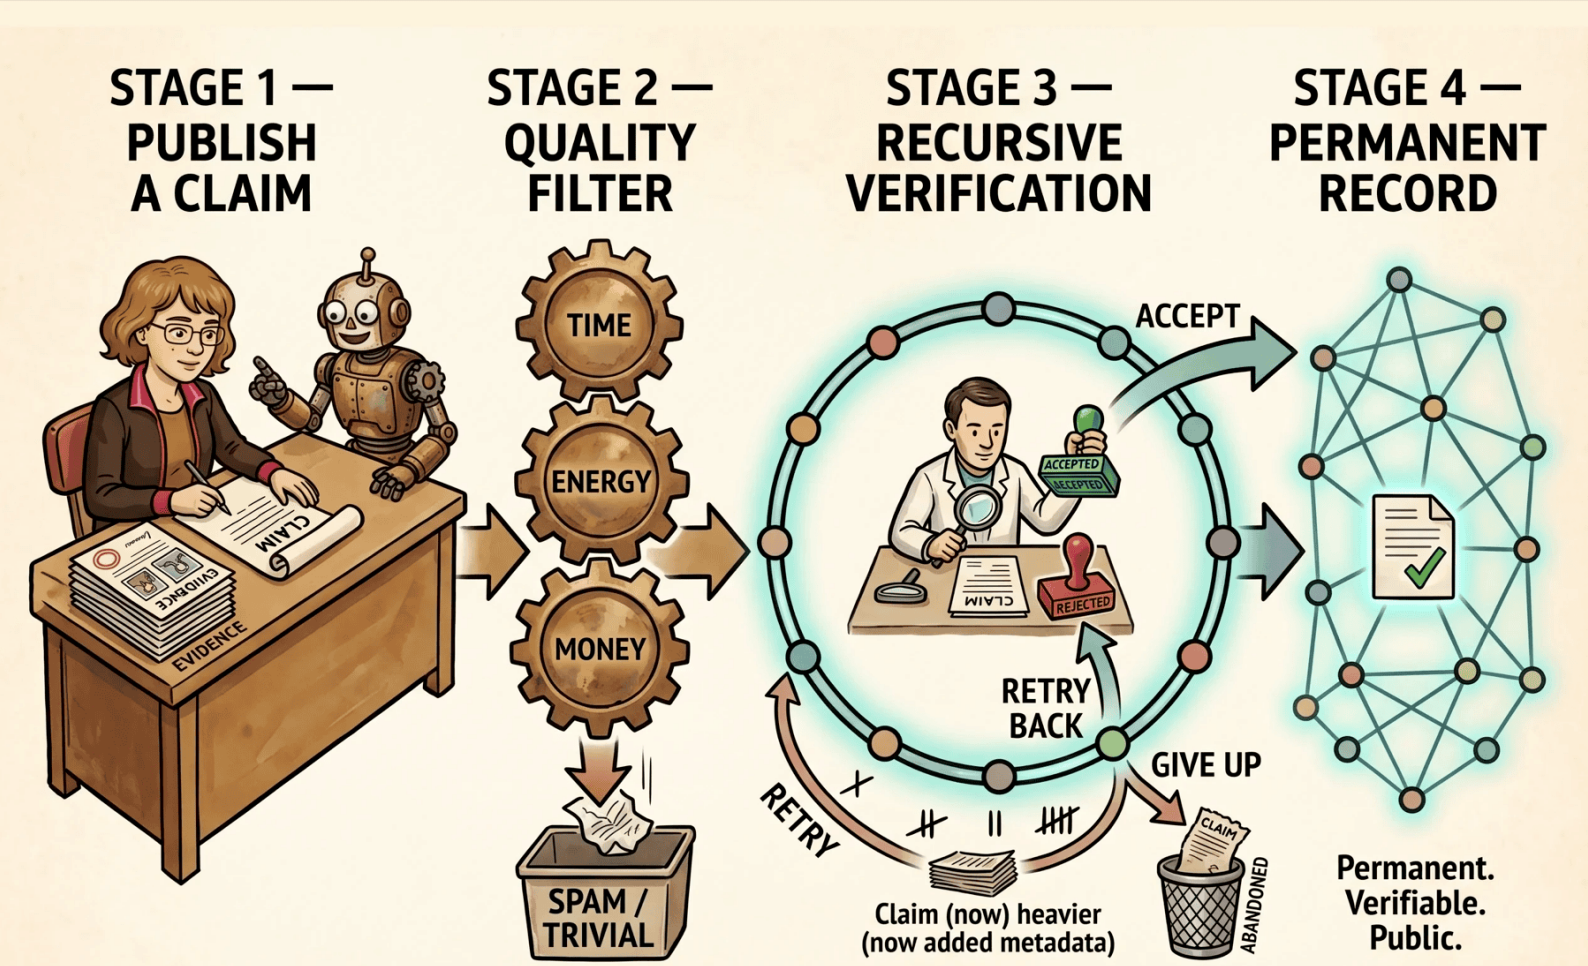

EXPENSIVE TO FAKE. CHEAP TO VERIFY.

By the time you try to publish information into the network, it will likely already contain information about its actors — these are reputation signals. Navigating how to read reputation signals — what they mean for you in different situations and what risks they carry — may not be trivial. Each participant can look at reputation records differently through their DID, depending on the situation they are dealing with regarding the counterparty. Is the counterparty a reliable payer, or do I need to demand money upfront for a business transaction? Does the offered product carry reviews about hidden fraud or defects? Are they trying to wriggle out of contractual responsibility when something goes wrong? Sometimes a more complex view of the counterparty's overall consistency comes in handy — it depends on the preferences of whoever requests the overview. The market could offer products and services that simplify, process, and clarify the reading of reputation in the context of the situation at hand. Various authorities and their offered services can serve this purpose as well.

HOW TO READ REPUTATION

INFORMED DECISION

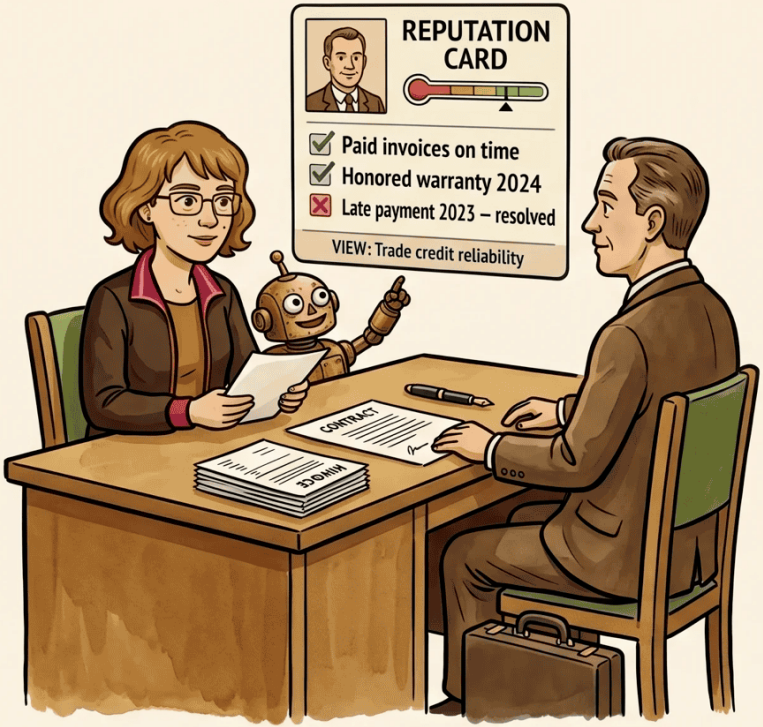

She knows who she is dealing with.

BLIND TRUST

He has no idea who he is dealing with.

BEFORE YOU EXTEND CREDIT. BEFORE YOU SIGN. BEFORE YOU TRUST.

Examples

Typical information of interest to publishers — and valuable to others — concerns events beyond ordinary interpersonal communication in the real or virtual world.

Negative examples: - evidence of criminal acts (e.g., audited by a trusted investigative body) - indirect evidence (weak on its own, but statistically cumulative) — e.g., repeated presence near multiple thefts in a short time → still coincidence? - breach of contract

Positive examples: - remedied harm (voluntarily or under pressure from the community as punishment) - acceptance and serving of a penalty proposed by authority X - authority X revoked recognition of the perpetrator's property rights to a certain extent

It is up to each person to gather available information about the counterparty and assess the risks according to their preferences.

WHAT KINDS OF CLAIMS CAN APPEAR?

EXAMPLES OF THE CLAIM ARTEFACT

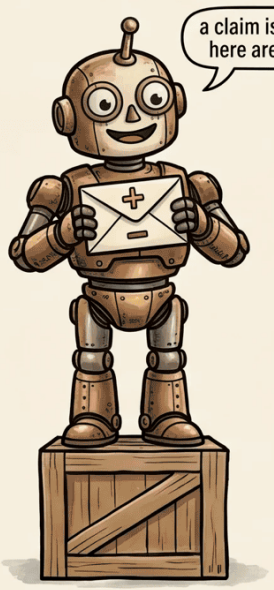

− NEGATIVE CLAIMS

ASSAULT
physically attacked another person

THEFT
took someone's property without consent

FRAUD
deceived someone for personal gain

BROKEN CONTRACT
failed to honour an agreed obligation

DEFAMATION
spread a damaging untruth

ABANDONMENT
left dependants or duties unattended

+ POSITIVE CLAIMS

HARM REMEDIED
repaired damage caused earlier

DEBT SETTLED
paid what was owed in full

AID GIVEN
helped another at personal cost

VERIFIED EXPERTISE
proven skill in a craft or field

OUTSTANDING SERVICE
long-standing contribution to the community

examples only — every claim follows a claim standard, and any standard can add more.

you do not have to join — but if your behaviour impacts others, someone can publish a claim about it.

CLAIMS ARE HOW BEHAVIOUR ENTERS THE NETWORK.

Whether information about you appears in the network depends exclusively on your own behavior.

You never have to join such a network, yet information about you may still appear in it. It depends exclusively on your actions and the impact they have on others.

THE COMMUNITY CAN OPEN ONE FOR YOU

He has no DID of his own.
He never signed up.

You didn't open a DID, so the community opens one **ABOUT** you. Negative claims attach here.

You can open your own — but the community will link the two. Reputation follows the person, not the key.

1. YOU DO HARM ----------→ 2. COMMUNITY FILES ----------→ 3. YOU CAN'T ESCAPE →

You can opt out of opening a DID. You cannot opt out of your reputation.

What I have just briefly sketched is how a social network inspired by Decentralized Identity (DID) could work. The primary purpose of DID concepts is to strengthen privacy and freedom through the principle of subscribing to the rules I will follow and live by — giving users the ability to decide what information to share and under what conditions.

I propose to further connect DIDs into a communication network where their holders exchange feedback even beyond situations where something has happened to someone and the community or an individual needs to react. Such preventive comparison of the rules we have signed up to — with the option to compute the economic and other consequences of mutual deviations in expectations about how the other side ought to operate — could be considered a motivation for finding consensus. Instead of freedom, such a system would emphasize voluntary decision-making combined with responsibility for real-world behavior.

An individual cannot break the system alone — a group of people stands a greater chance, and a group of people with negotiated consensus and motivations to pull together on many issues stands an even greater chance of resisting authoritarian tendencies. The organization prerequisite from the first chapter will be fulfilled once two conditions are met: the DID reputation network covers communities representatively enough that its use ceases to be exotic. And at the same time, this community segment becomes an economically significant minority that can assertively negotiate with the rest of society.

Voluntariness vs freedom

Freedom — in the positive sense — would be a secondary effect of balancing two factors: voluntariness and the pressure of one's surroundings toward responsibility.

The AI Era and the Value of Reputation

In the era of artificial intelligence, everything connected to cognitive thinking is being automated — and it may go even further. What then remains in human activity as a competitive advantage? The answer is hard, and something will surely be found, but one thing we can say with certainty: reputation will decide. A verifiable history of your behavior, your commitments and their fulfillment — that is something AI will not build for you.

IF AI DOES THE THINKING — WHAT IS LEFT FOR YOU?

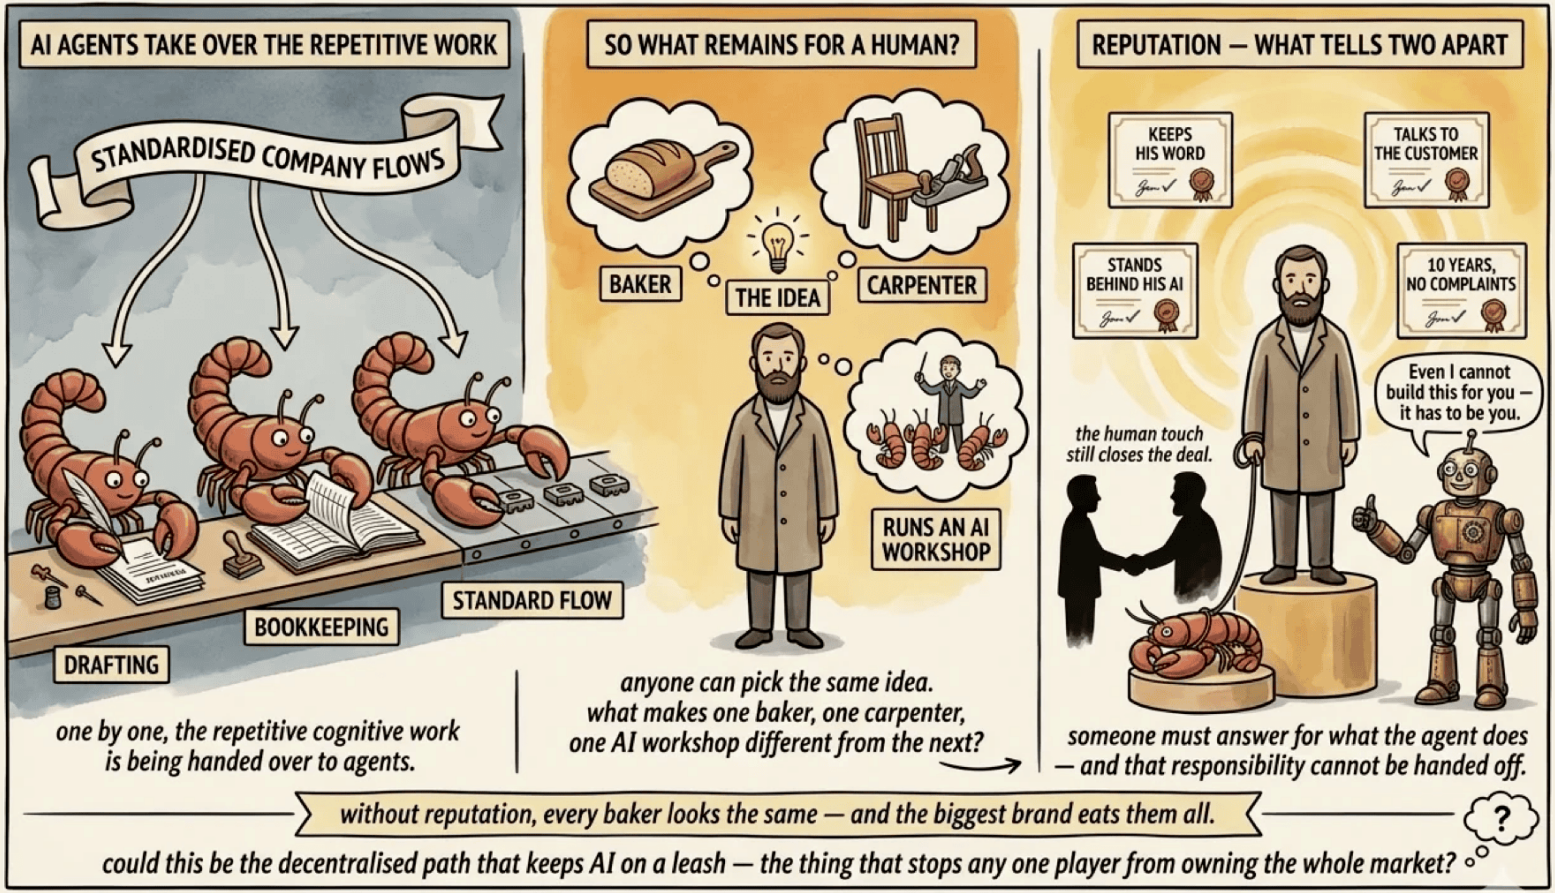

AI AUTOMATES THE WORK — REPUTATION KEEPS YOU HUMAN.

THE ECONOMICS OF TRUTH

COMMUNITY STRIP

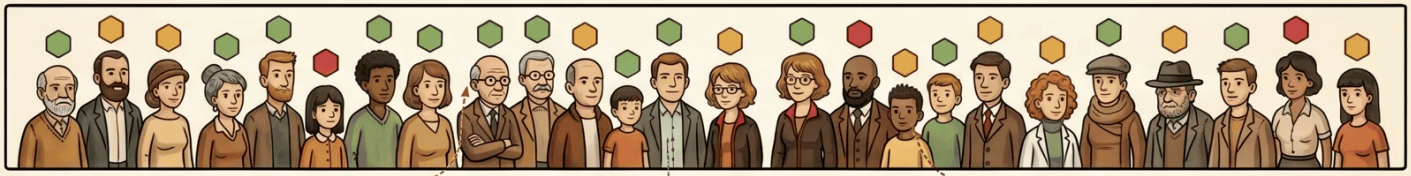

ONE COMMUNITY. DIFFERENT POSITIONS.

UNPOPULAR BUT ACCEPTABLE

IN CONSENSUS

OUTSIDE CONSENSUS

THE SUBMISSION

THE DRAW

THE REFUND

THE LEDGER

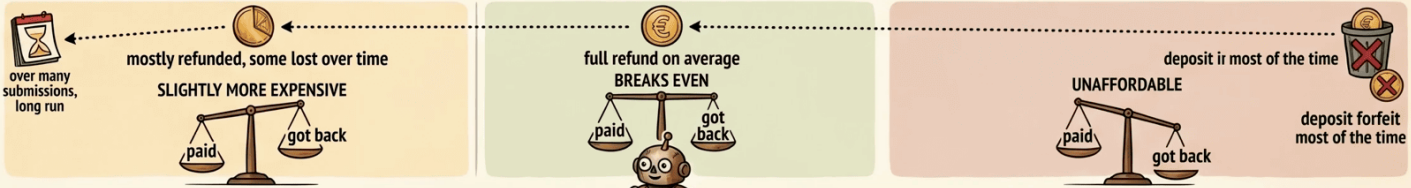

BOTTOM OBSERVER STRIP

UNIFORM DRAW. NON-UNIFORM RESPONSE. THE FURTHER FROM CONSENSUS, THE HIGHER THE PRICE.

WHAT GETS PUBLISHED IN THE NETWORK

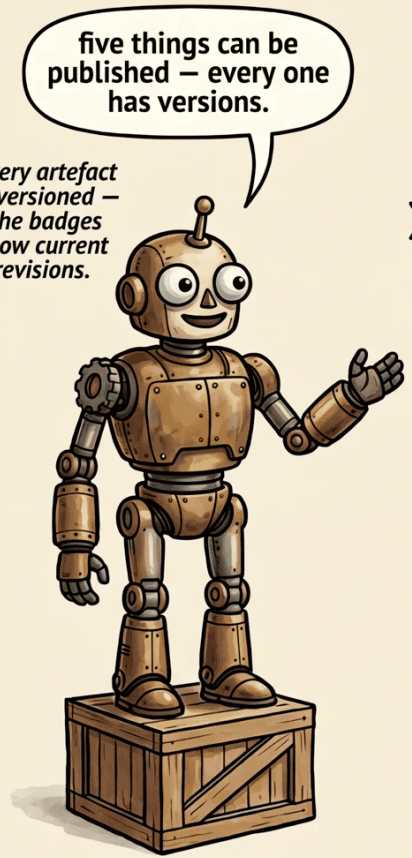

v1

DID DOCUMENT

your identity record on the network — keys, metadata, network address.

MANAGED BY: the DID holder

v2

CLAIM STANDARD

the schema a claim must follow — what counts as valid evidence.

MANAGED BY: an authority with reputation in the domain

v3

POLICY STANDARD

the schema a policy must follow — the shape of acceptance rules.

MANAGED BY: an authority with reputation in the domain

derived from

derived from

v4

CLAIM

a statement about a subject, following a claim standard — published by an issuer, co-signed by the standard's authority if both issuer and verifier request it.

MANAGED BY: issuer publishes — subject is the target

claims come in two kinds — *positive* (+) and *negative* (–).

v5

DID POLICY

a DID's declared rules — used in verification, consensus, and issuance.

MANAGED BY: the DID holder (one policy per DID)

EVERY ARTEFACT IS VERSIONED. EVERY VERSION IS PUBLIC.

WHAT IS INSIDE A DID DOCUMENT?

EXAMPLES OF THE DID DOCUMENT ARTEFACT

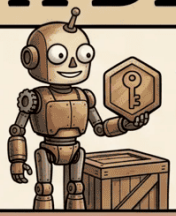

your DID document is your identity card on the network — here is what it carries.

FIELDS YOU WILL FIND INSIDE

KEYS AND SIGNING RULES

each key grants one power — edit this document, issue a claim, act as an authority, manage a policy, or delegate recovery. Powers can be handed to another DID, and revoked again later.

NETWORK ADDRESSES

where the network can reach you — changeable without rewriting your identity

CREATED AND VERSIONED

when this identity first appeared, and which revision you are reading now

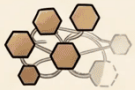

LINKED IDENTITIES

a summary of DIDs you are connected to — a convenience view, technically stored elsewhere

SUCCESSION MARKER

if your identity is lost or stolen, an authority can mark this DID ended and point to your successor

the DID holder writes each new version — but every version must pass verification before the network publishes it.

you write it — the network verifies it. You cannot just publish yourself into existence.

YOU AUTHOR YOUR IDENTITY — THE NETWORK CONFIRMS IT.

WHAT IS A CLAIM STANDARD?

EXAMPLES OF THE CLAIM STANDARD ARTEFACT

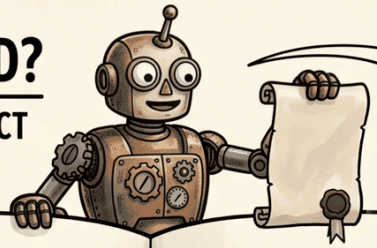

the standard is the empty form — a claim is the form filled in. Let me show you both side by side.

THE STANDARD — THE EMPTY FORM		A CLAIM — THE FORM FILLED IN
SUBJECT DID the DID this claim is about	⟷	**SUBJECT DID** *did:nwo:alice-9f2b...*
ISSUER DID the DID signing the claim	⟷	**ISSUER DID** *did:nwo:magistrates-04...*
WHEN AND WHERE a date and a place	⟷	**WHEN AND WHERE** *3 March, market square*
APPROVED EVIDENCE TYPES only these kinds of evidence count	⟷	**APPROVED EVIDENCE TYPES** photograph ✓ — witness oath ✓ — physician's note ✓
none mild moderate severe **OUTCOME SCALE** pick exactly one rung — no free text	⟷	**OUTCOME SCALE** *moderate — rung 2 of 4*

every field on the left has exactly one filled-in counterpart on the right — that is what makes two claims of the same kind comparable.

the authority's seal on a filled-in claim is a pledge: the process was truly followed.

I stake my reputation that this was processed by the standard.

THE STANDARD IS THE FORM — A CLAIM IS THE FORM FILLED IN AND PLEDGED.

WHAT IS A DID POLICY?

EXAMPLES OF THE DID POLICY ARTEFACT

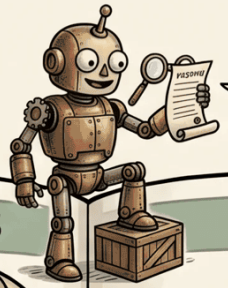

my policy is not about how a claim is written — that is the standard's job. My policy is how I decide whether to care about the claim at all: do I trust who signed it, do I trust the authority behind their standard, and is the matter even worth my attention?

THE STANDARD — INHERITED DEFAULTS

	RULE NAME	
	ISSUER REPUTATION FLOOR	I accept issuers with reputation above five of ten
	AUTHORITY TRUST LIST	I accept standards from any authority whose own reputation is above four of ten
	SUBSTANCE FILTER	I process any claim whose matter weighs above the 'trifling' mark
	ANTI-SPAM PRICING CURVE	submitting to me is cheap for substantive, rare claims — and grows fast for petty or repeated ones
	COMBINED WEIGHTING	issuer-rep and authority-rep multiply; petty-substance halves the result

policy does not rewrite the standard — it judges the people and the matter. Keep, modify, remove, add — four simple moves on top of a shared starting point.

MY POLICY — DERIVED FROM IT

	RULE NAME	KEPT TICK	EDITED ACTION WORD	REMOVED ACTION WORD	ADDED ACTION WORD
→	ISSUER REPUTATION FLOOR		tightened — I require seven of ten		
→	AUTHORITY TRUST LIST		I fully trust magistrates and physicians; I halve the weight of claims under the municipal watch; I refuse anything under the harbour police altogether		
→	SUBSTANCE FILTER		petty nuisance claims cost me nothing to ignore — I simply do not read them		
→	ANTI-SPAM PRICING CURVE	✓			
→	COMBINED WEIGHTING			~~removed — I use my own formula~~	
→	EXPONENTIAL COST ON REPEATED CLAIMS				added — each new claim from the same issuer about me in a year costs ten times the last
→	QUIET HOURS				added — I refuse to read any claim filed between midnight and dawn

VERIFICATION — my policy filters both incoming claims and my own drafts before I issue them. Issuance is verification on the way out.

CONSENSUS — my policy is my vote when the community negotiates shared rules.

YOUR POLICY JUDGES THE MESSENGER — NOT THE FORM.

Why the Network Is Uncensorable

So far we have described what the network can do. Now let us focus on why it cannot be shut down or censored.

The network will likely be built on the principles of the Tor network and onion routing (a way of masking network traffic using Mixnet principles) — a technology where communication passes through several layers of encryption and no single node knows both the sender and the recipient at the same time. A central institution cannot look over your shoulder, because there is nowhere from which to look.

You can fight it, but you cannot win — unless you plunge the entire country into totalitarian surveillance at the level of individual network packets and low-level actions on every piece of hardware that connects to the internet and communicates on the given reputation network. And that is precisely the choice centralism faces: either accept that citizens have a tool beyond your reach, or dig yourself into the dark corners of economic stagnation in the name of totalitarian control.

Let us have no illusions — even democratic regimes will attempt this in their fight for centralism. They would rather put people in shackles than surrender power.

HOW YOUR MESSAGES STAY UNTRACKABLE

TWO LINES OF DEFENCE AT ONCE — LAYERS AND NOISE

CONTENT: WRAPPED IN LAYERS

(the technique behind Tor — 'onion routing')

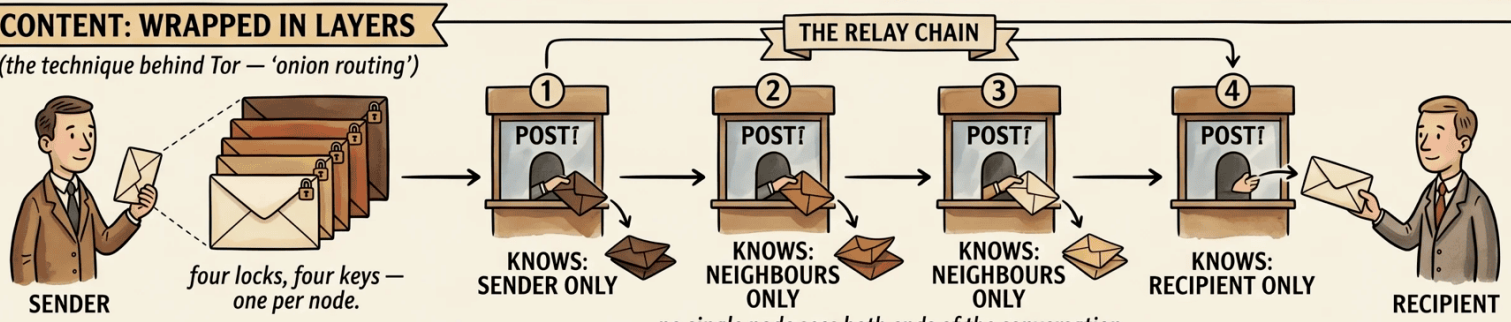

no single node sees both ends of the conversation.

TIMING: HIDDEN IN A RIVER OF DECOYS

(the technique behind a 'mixnet' — e.g. Nym, Loopix)

the community pays a small cost — constant cover traffic — in exchange for untrackability.

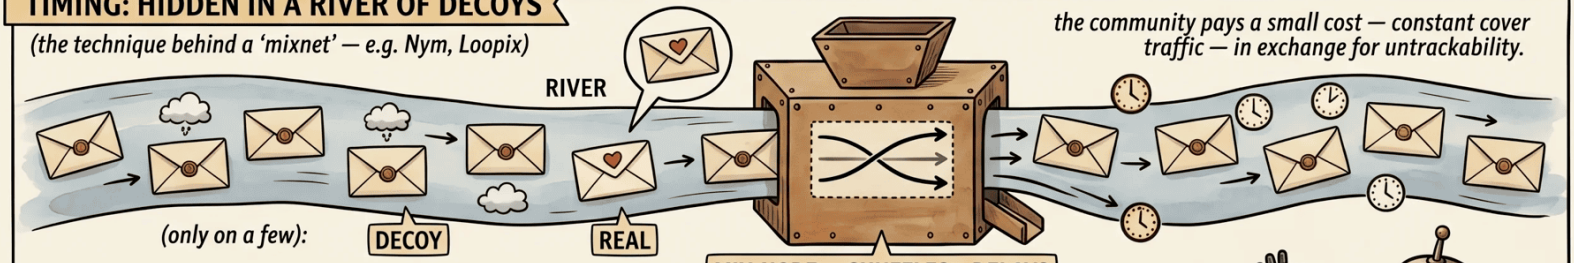

THE TWO SHIELDS TOGETHER

MATRIX	HIDES WHO ↔	HIDES WHEN?
LAYERS ONLY	✓ YES	✗ NO
NOISE ONLY	✗ NO	✓ YES
LAYERS + NOISE	✓ YES	✓ YES

(onion routing + mixnet — the combination we need)

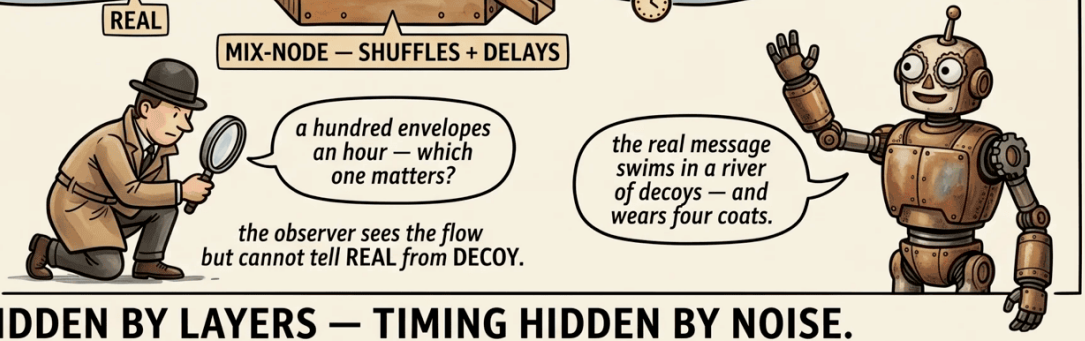

CONTENT HIDDEN BY LAYERS — TIMING HIDDEN BY NOISE.

Addressing Obvious Doubts

The described system naturally raises a number of questions. Let us address the most common ones.

Dependence on Technology

You have probably already noticed that, unlike the modern state — which has been with us for some 150 years — the reputation network solution as described here depends heavily on the technology of the global/local internet. In the event of an outage, the functioning of such a network is at risk.

If the outage is temporary, there is no loss of data or of the consistency of claims in the network, and the reputation balances achieved within communities should not be disturbed either. Unlike comparable payment networks, this network assumes very low frequencies of processes and data items. In this respect the decentralized reputation network does not differ from today's state, which is also now heavily dependent on technology and has forgotten how to work with paper card files (though in crisis plans it would have no choice).

The evolutionary successor of the state in the form of a reputation network can, in the event of a permanent outage (a catastrophe of unprecedented scale), fall back to a more primitive centralized system — the state.

Technology allows humanity to reach higher civilizational forms of governance and brings us benefits, but also risks.

Am I Throwing Money Out the Window by Publishing in the Network?

The costs of publishing a claim in the reputation network are largely not sunk. The usefulness of a publisher's message put into the network — information about real grievances, verified experiences,

TWO PATHS — THE STATE CLOSES, THE NETWORK OPENS

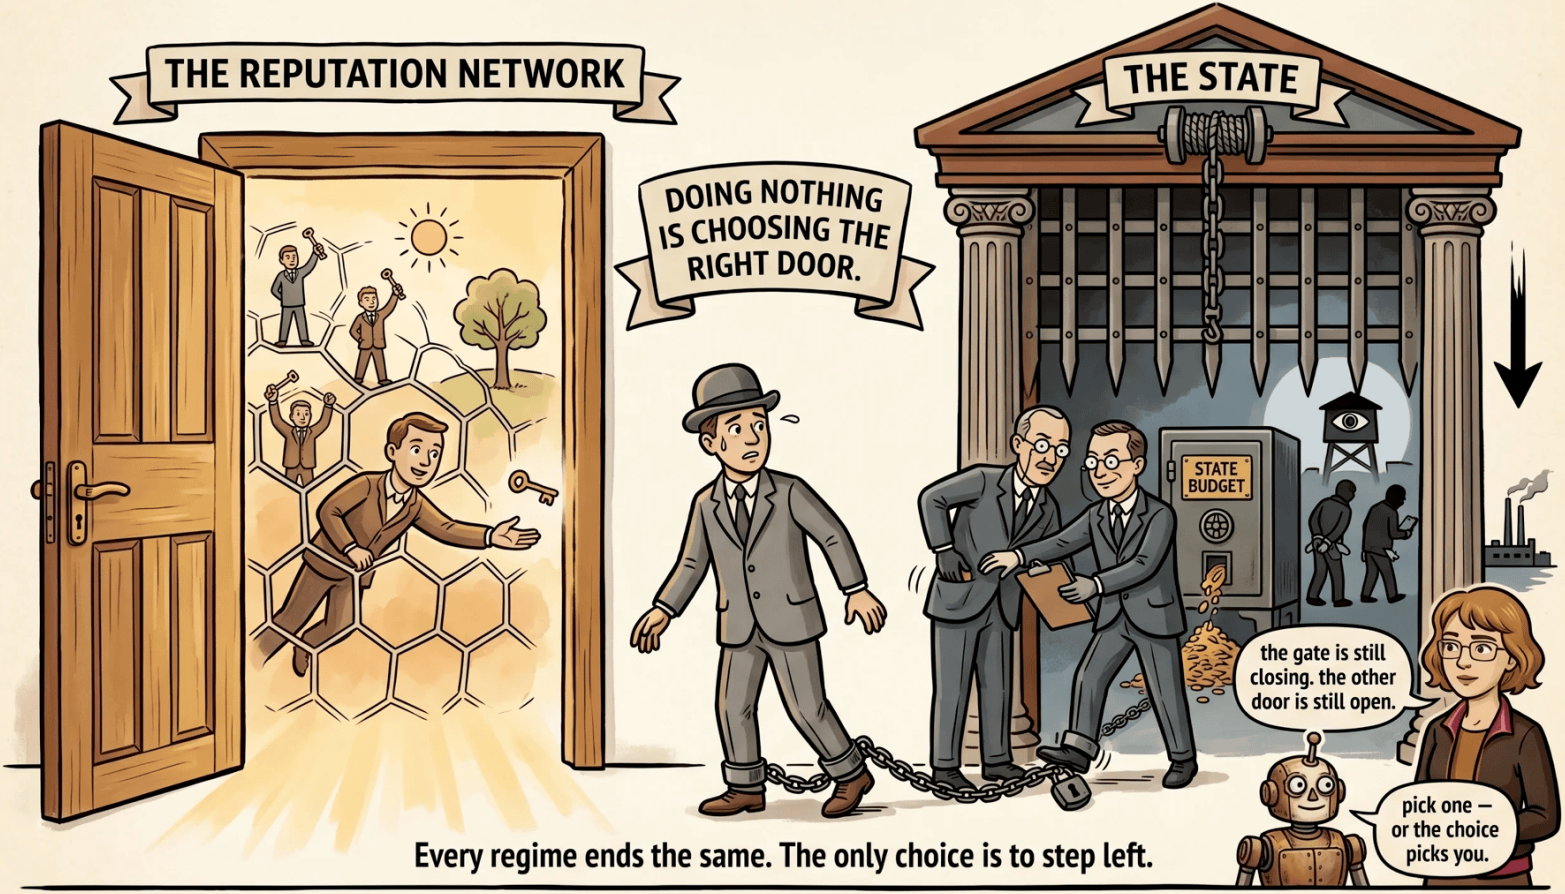

THE STATE CLOSES BY DEFAULT. THE NETWORK OPENS BY CHOICE.

relevant warnings — statistically returns over a longer time horizon in the form of fees for verifying someone else's claims from the community. The community benefits, the publisher builds reputation, and the costs of publishing truthful information thus approach a refundable deposit, from which only a smaller portion of actual network maintenance costs needs to be subtracted. Conversely, untruthful or trivial records do not return — their costs are a net loss. Honesty is therefore not only a moral choice, but also an economically rational strategy.

After subtracting network maintenance costs, this principle could be called the principle of economic neutrality — I do not lose when I am with the community, I lose when I am against it.

The community also has solidarity channels for appreciating the honest approach of its members. But let us have no illusions: appeals to solidarity mostly arise from community social pressure, so this may not be solidarity in the voluntary sense of the word.

What If Someone Creates Multiple Identities?

A person can operate multiple DID identities in parallel. However, building reputation for each identity requires independent effort — time, energy, money.

No shortcuts: each identity must accumulate its track record[1] through actual activity. Maintaining parallel identities is therefore intentionally expensive.

In free societies, the costs discourage abuse.

[1] **Track record** — generally: a history of past results, successes, and failures of a person or organization. Here: the sum of all past interactions of a given DID identity in the network — verified claims, accepted and rejected records — from which its reputation is derived.

In dictatorships, however, parallel identities become a survival tool: they enable the organizing of underground networks, safer navigation of the black market, and compartmentalized[2] resistance, where compromising one identity does not reveal the others — and after the regime falls, they allow a seamless return to public life and the reputation built up there, and even the merging of the previously official and secret DID into a single combined set of records via a claim.

[2] **Compartmentalization** (from Eng. *compartment*) means separating information into isolated units so that the exposure of one unit does not compromise the others. A principle known from intelligence services: an agent knows only their part of the operation, so even under duress they cannot reveal the whole.

THE SAME EFFORT — A WASTEFUL PLAN B IN FREEDOM, A LIFELINE UNDER TYRANNY

PARALLEL IDENTITIES COST THE SAME — BUT BUY DIFFERENT THINGS

IN FREEDOM — A COSTLY PLAN B

SEED → SPROUT → SAPLING → YOUNG TREE → MATURE TREE

5 to 15 years — no shortcuts, no mature tree for sale

year 1 — time, energy, money spent — 5% reputation

year 3 — the same costs again, paid independently — 20% reputation

year 0 — everything starts over — 0% reputation

three gardens, triple the cost — and you still only have two hands.

OR — A CLEAN SLATE AFTER A PERSONAL MISSTEP

OLD DID — REPUTATION BURNED

DID-FRESH-START

works only if the mess stayed your own — if others were harmed enough to chase you down, their claims will follow you into the new pot too.

ONE CLAIM AND THE PLAN B IS GONE

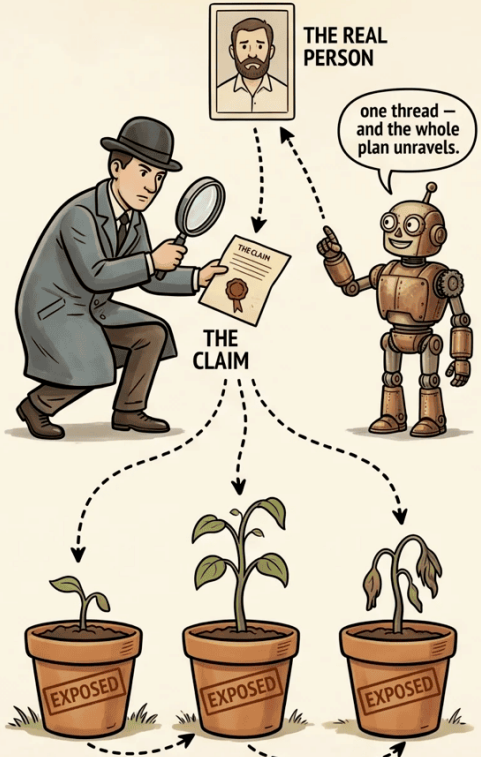

a single claim ties every pot to the same gardener — years of time, energy and money erased in an afternoon.

UNDER TYRANNY — ONE FOR THE STATE, OTHERS FOR FREE ASSOCIATION

(the state either bans identities outright or requires them under its own control)

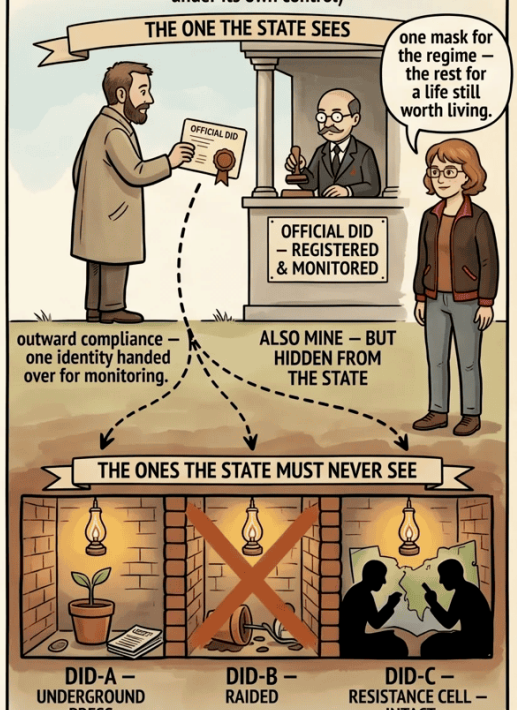

the brick walls between cells do the work the state forbids — lose one, the others carry on; when the regime falls, public life resumes.

IN FREEDOM, ABUSE IS EXPENSIVE. UNDER TYRANNY, SURVIVAL IS POSSIBLE.

Perception Turned Inside Out

Unlike the state, the principle of a reputation network built on decentralized identity inverts the paradigm of perceived priority:

- What matters is reputation — i.e. the past, used to assess the risks of interaction with a counterparty — and personal data such as names, addresses, etc. can be a matter of courtesy data exchange
- Whereas the state primarily demands personal data, accumulates reputational data, and only lets the community see what suits it

Can't a Wealthy Person Simply "Buy" More Identities (or Create Virtual Communities)?

The possibility of creating parallel decentralized identities at first glance looks like an unfair advantage for people with greater economic means over those with fewer. It must be emphasized, however, that unlike a centralized system, where it suffices to corrupt a few points in the pyramid of power to make certain misdeeds disappear, in a decentralized system one would have to corrupt the entire community.

Substitute decentralized identities could serve this purpose, but their reputation must be built up over time through real interaction with actual other community members — it cannot simply be bought, because the network makes it verifiable how a given identity is performing.

Moreover, services can exist on the market (and likely will) that, acting as an authority, offer investigations producing evidence that several decentralized identities are in fact the same person. A single record entered into the reputation network can thus nullify, for a fraction of the cost, the entire investment of time, energy, and money spent building parallel identities.

Economically, it therefore pays to not cheat the community — and if needed, to examine one's actions and pursue remediation so that reputation returns to an acceptable level and its bearer does not suffer economically or otherwise from the community's wrath.

An economically powerful identity that has lost reputation in its own community may try to escape the community's wrath through underhanded deals with targeted identities — but those then also risk losing their own reputation.

There still remains the escape to another community with a fresh identity — but that means leaving all achievements behind and starting somewhere from zero with zero reputation. Sometimes it may be an understandable path and the only way out.

REPUTATION IS THE COUNTERWEIGHT TO ECONOMIC FORCE

THE RICH MAN BUYS INFLUENCE — THE COMMUNITY REMEMBERS IT

THE ATTACK — INFILTRATE AND BUY OFF

EIGHT PUPPETS PLANTED + THREE NEIGHBOURS BOUGHT OFF

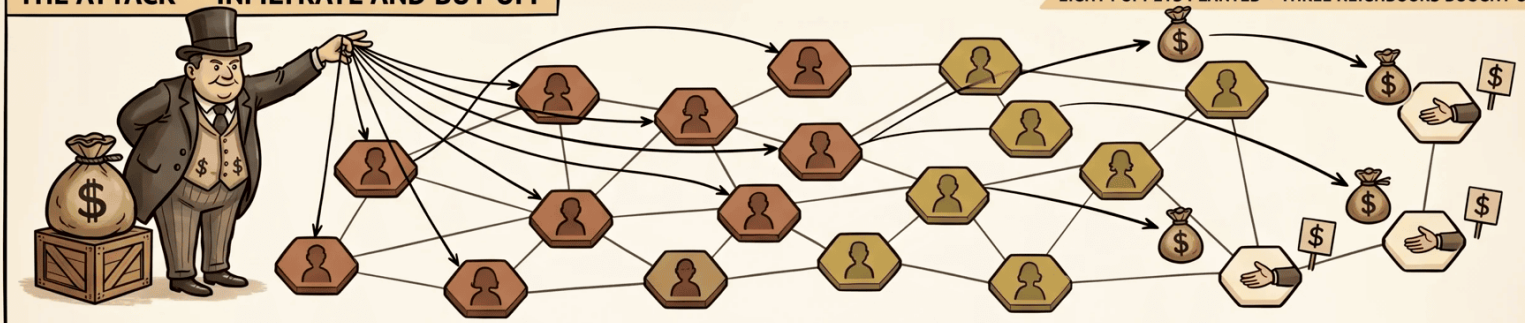

economic force can slip strings into any community — and buy a few real voices along the way.

at first, it looks like the rich man just widened his circle of friends.

LATER

THE DEFENCE — LONG MEMORY, DISTRIBUTED ACCOUNTING

EVERY NEIGHBOUR KEEPS THEIR OWN MEMORY — NO CENTRAL BOOK, NO FORGETTING

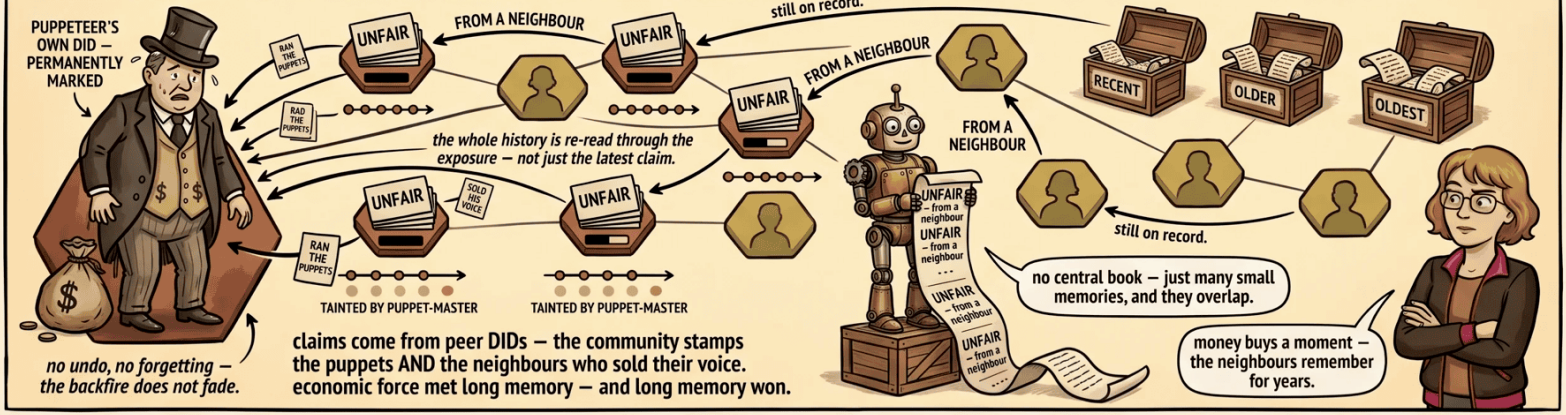

no undo, no forgetting — the backfire does not fade.

claims come from peer DIDs — the community stamps the puppets AND the neighbours who sold their voice. economic force met long memory — and long memory won.

ECONOMIC POWER REACHES IN — THE BACKFIRE COMES HOME, AND IT DOES NOT WASH OFF.

Note

Similarly, communities would deal with an attack where someone devotes resources to creating a virtual identity: for that identity, it is risky to enter into interaction with another community without verification — that is, to uncritically accept information about the identities of the other community. Reputations are always built within a community, not globally.

What About Free-Riders Who Just Want to Read and Give Nothing to the Community?

Access to information is not unlimited from day one of creating a decentralized identity. New participants — those who have not yet built reputation through actual activity — face graduated restrictions: less information, longer waiting times, higher query costs. The network rewards participation, not passive consumption and arbitrary data harvesting.

A decentralized identity also risks its reputation when it lends its reputation for payment to another person (who may not even be in the DID reputation network). The same principle applies here: such betrayal of the community (violation of privacy) can reflect on the traitor's reputation, and the deed cannot be erased through an underhanded deal the way it works in a centralized system. They must reckon with the community's wrath, including the loss of achievements — for the community is, in their eyes, the guarantor of, for example, the privilege of owning movable and immovable property.

Anchor to the Real World

When assessing risk, the riskier subject is naturally the one who does not enjoy the privilege of ownership recognized by a given community — they have less to lose in their dealings (digital assets are easier to move).

It may look like a minor detail, but it has major consequences. Wanting the community to have leverage over its members implies ownership as a privilege — in the freest societies nearly untouchable, yet still not a right, nor a fundamental principle, but a privilege that in extreme cases may be revoked (I can imagine, for example, the refusal of service in defense of the community in armed conflict).

It also subliminally answers how the community will treat its members, and what motivation a member has to fight for the community — to retain their privileges. A person may fail in their responsibility to the community, but morally cannot expect leniency when it comes to retaining hard-won privileges.

Financial Neutrality

When reading words like decentralized, uncensorable, incorruptible, one cannot help but associate them with the best-known cryptocurrencies — Bitcoin, Monero, and, say, Kaspa — which can be described in such terms. Intuition is misleading here, however: fees for the services of authorities, for verification and publication, and so on can be settled in any currency or money. What matters for the connected participants of the social network in the DID network (that is, your community) and its surroundings is

THE NETWORK REWARDS PARTICIPATION

REPUTATION IS BUILT OVER A LIFETIME — NOT BOUGHT, NOT BORROWED

THE PARTICIPANT — A LIFETIME OF REAL WORK

TIME ENERGY MONEY

TIME ENERGY MONEY

TIME ENERGY MONEY

TIME ENERGY MONEY

harder work can shorten the climb.

time builds the curve — your own effort can make it steeper.

AGE 20 — NEWCOMER

reputation 5%

0 projects completed

LIMITED — HIGH COSTS, LONG WAITS

AGE 28 — ACTIVE PARTICIPANT

reputation 30%

delivered contract | resolved dispute | community contribution

BASIC — STANDARD COSTS, NORMAL SPEED

AGE 40 — CONTRIBUTOR

reputation 70%

8 projects delivered

FULL — LOW COSTS, INSTANT ACCESS

AGE 55 — AUTHORITY

reputation 95%

15+ projects delivered

AUTHORITY — YOU VERIFY OTHERS, THE NETWORK PAYS YOU

every stage costs time, energy and money — together, not separately.

the door opens as the ledger of your deeds fills.

THE FREE-RIDER — READS, NEVER PUBLISHES

reads everyone else's records.

MY CONTRIBUTIONS

publishes nothing of his own.

LIMITED ACCESS — THE CEILING DOES NOT LIFT

reputation 5%

no deeds, no rise — year after year, same bar.

the network gives back to those who give in.

TIME, ENERGY AND MONEY — INVESTED IN REAL WORK — IS THE ONLY CURRENCY THE NETWORK RESPECTS.

a reputation-backed confirmation that the payment was made. The publication of a claim must individually carry a reasonable, verifiable cost, so that an actor cannot publish however many and whichever claims they want without expending energy, money, and time — a strongly undesirable state of affairs corresponding to the privilege of elites in today's corruption-ridden state systems.

In this respect the mentioned cryptocurrencies have a small advantage in that their networks act as trusted authorities for verifying that a given payment took place, at the cost of a small loss of privacy and the exposure of some of one's addresses.

Voluntariness, Responsibility, and Freedom

Voluntariness is individual: I decide for myself. Responsibility arises when my decision affects others — feedback kicks in. Freedom is neither a gift nor a right — it is a phenomenological[3] result of thousands of bilateral interactions between voluntariness and the pressure toward responsibility, averaged across society. The network does not dictate this average — it emerges[4] from the price signals of verifiers and from the pre-announced verification policies that community members set for themselves.

[3] **Phenomenological** (from Gr. *phainómenon* — that which appears) — generally: an approach examining phenomena as they present themselves in direct experience, by observing what follows from them, without pre-given explanatory theories in the microworld beneath. Here: freedom is not an abstract principle defined from above, but an observed phenomenon — a consequence of thousands of micro-interactions between people that emerges from how we actually treat each other.

[4] **Emergent** (from Lat. *emergere* — to emerge) means "spontaneously arising from interactions of simpler parts without anyone designing or directing it." A flock of birds flies in formation even though no bird has a plan for the formation — the formation emerges from each individual's simple rules. Likewise, in the reputation network, nobody designs the behavioral rules of the whole — they emerge from thousands of individual decisions.

What preference for freedom an individual within a community has depends on their sense of how much they are themselves willing to swallow of what they impose on others. With things that are self-evident and beneficial to all, there is no moral dilemma — murder is wrong. Then there are cases where options shift into the position of permitted privileges: prove your skills and we will let you onto the road, so there is less chance you will kill us. And finally there are matters where the individual says — it is not worth it to me, I will not forbid my neighbor from burning leaves when I do it myself, or because I do something else that restricts their comfort and they could in turn want to forbid it to me. If we have historically drawn the boundary of freedom where someone else's freedom begins, then a reputation social network built on decentralized identity fulfills this better than centralized representative democracy — where, for instance, the privileged layer of elected representatives tries to exempt itself from surveillance of personal communications, yet does not hesitate to impose it on their voters, using all kinds of underhanded pretexts and manipulations.

The Freedom-Totalitarianism Switch (and Delegation)

A simple mechanism within the DID network can shift the entire system toward greater freedom or greater control: the number of records a DID may publish per time period. When records remain scarce, each one carries weight — that is freedom. When records are abundant and cover trivialities, that is totalitarianism. You choose the direction with every record you publish. The network itself is neither free nor totalitarian — it is neutral infrastructure. What it becomes is determined by how people use it.

Within my own verification policy, it makes sense to set, for a generic DID in my social network — within communities — a maximum number of claims (any verifiable statement in general; here: a record published to the reputation network) that a publisher may issue per period for me to take them into account and consider their verification requests relevant.

VOLUNTARINESS → RESPONSIBILITY → FREEDOM

THREE STEPS — FROM ONE PERSON'S CHOICE TO A WHOLE SOCIETY'S PATTERN

VOLUNTARINESS

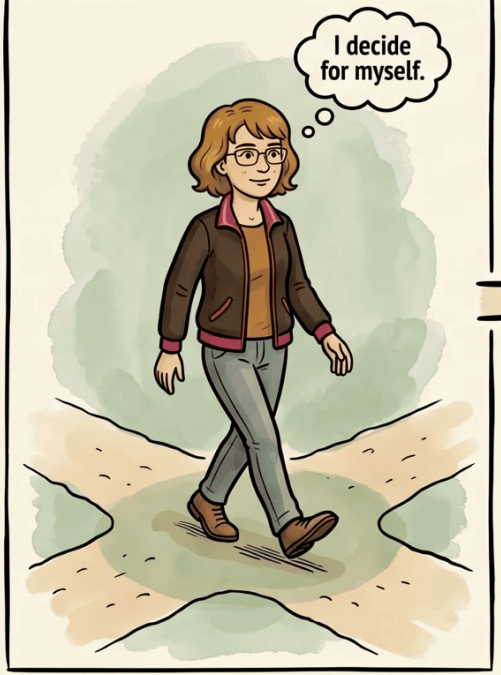

the individual decides — no one compels her.
and no one owes her a positive duty either.

RESPONSIBILITY

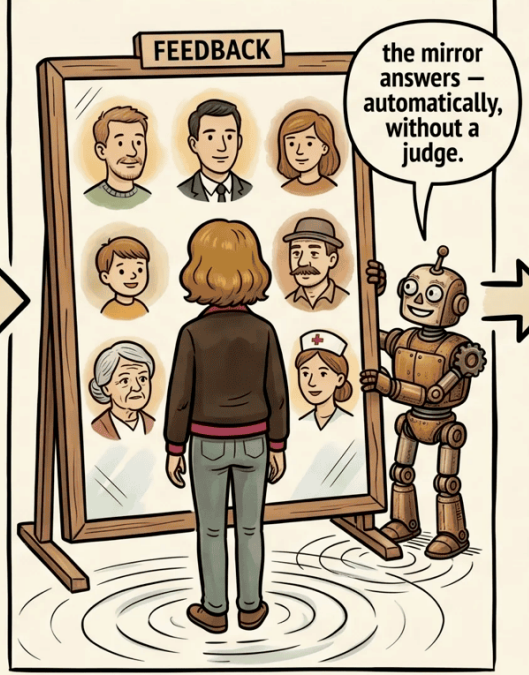

my choices ripple out — and the ripples come back.
the consequences speak for themselves.

FREEDOM

no one drew this map — it drew itself.
freedom is the statistical shape of how neighbours treat each other.

FREEDOM IS NOT A GIFT FROM ABOVE — IT IS THE PATTERN WE MAKE TOGETHER.

I keep an intuitive number in mind — low single digits per year. Such a censorial approach may seem rather isolationist. But I can legitimately put the applicant before a choice: to pick one of their already issued records and ask the verifier to revoke it, that is, to invalidate it — which is a fairly aggressive demand in a situation where the applicant has, by coincidence, genuinely become the victim of multiple independent wrongs, and I would surely thereby contribute to their feeling of estrangement from the community.

A more effective option therefore appears to be the ability to delegate part of one's rights to publish information into the reputation network in favor of a community or an authority, which then — in a solidary manner, according to its internal rules — decides how to allocate them. A group must be able to take care of those it considers disadvantaged. A community member has the right at any time, until the delegated right is used, to revoke their delegation.

THE FREEDOM–TOTALITARIANISM SWITCH

ONE SINGLE PARAMETER — HOW MUCH I CHOOSE TO TRACK — TIPS A WHOLE SOCIETY BETWEEN FREEDOM AND TYRANNY

not many knobs — one. the network itself is neutral; my policy is not.

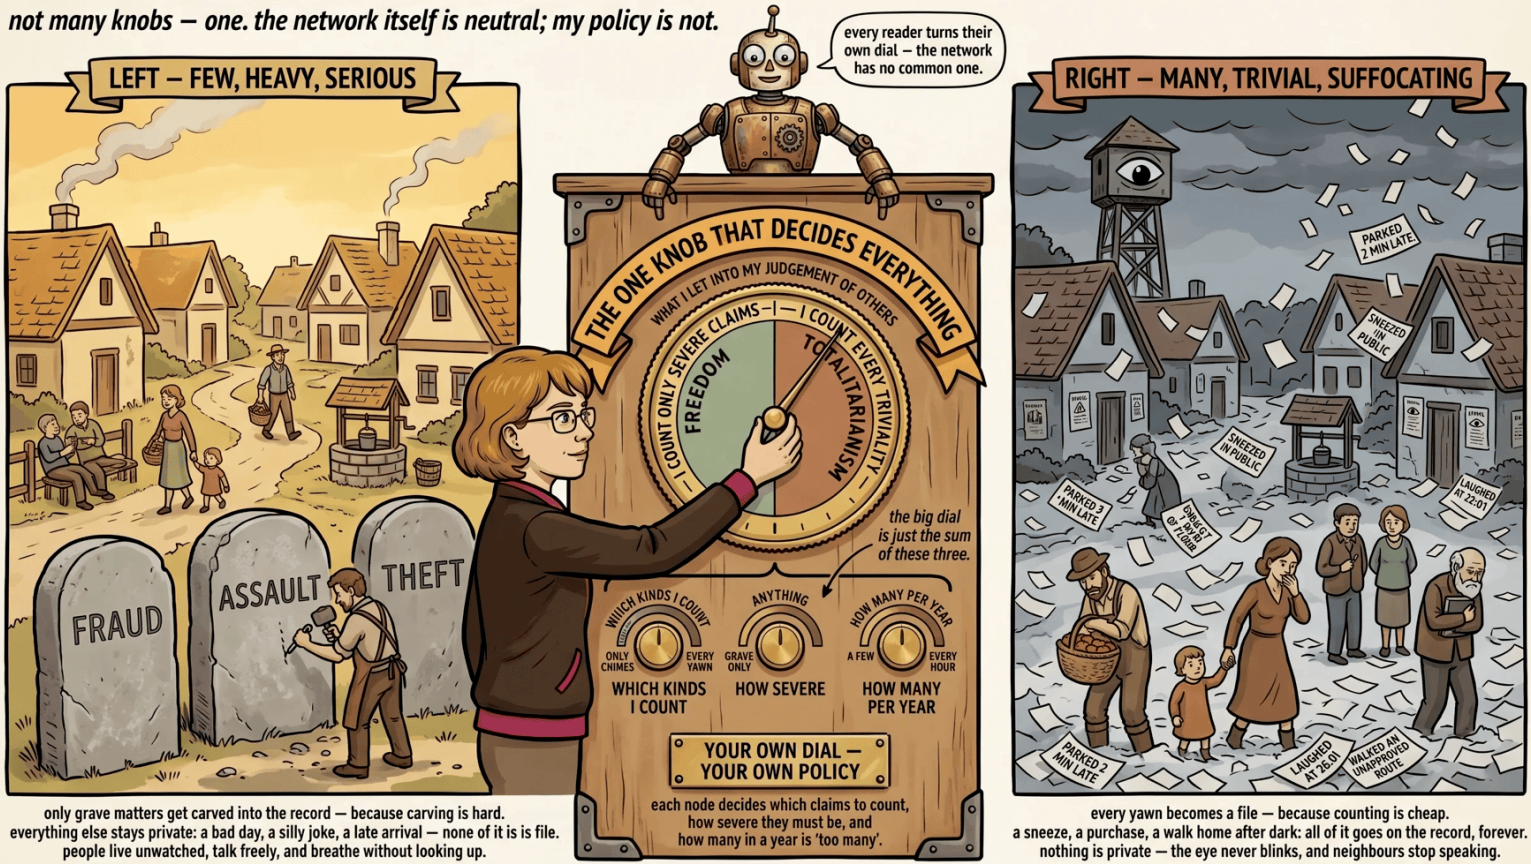

only grave matters get carved into the record — because carving is hard.
everything else stays private: a bad day, a silly joke, a late arrival — none of it is file.
people live unwatched, talk freely, and breathe without looking up.

every yawn becomes a file — because counting is cheap.
a sneeze, a purchase, a walk home after dark: all of it goes on the record, forever.
nothing is private — the eye never blinks, and neighbours stop speaking.

FREEDOM IS WHAT I REFUSE TO COUNT — TOTALITARIANISM IS WHAT I AGREE TO TRACK.

DELEGATION OF PUBLICATION RIGHTS

YOU LEND THE KEY TO AN AUTHORITY THAT SERVES YOU — AND PULL IT BACK IN AN INSTANT

ONE TUG — THE KEY COMES HOME

I lend the key — the string stays with me.

delegation — the authority acts on my behalf, within rules I approve.

AUTHORITY — SERVICE PROVIDER
policy · standard · party · chamber

STRING OF RECALL

MY RIGHT TO PUBLISH — 1 RECORD

HOW WE SERVE YOU

POLICY

DRAWN AT RANDOM WHEN THE AUTHORITY ACTS

DELEGATED KEYS — POOL

DELEGATE/BODY AUTHORITY ACTS

the string never leaves your hand — that is the whole trick.

I can lend more than one key at a time — a background automation can re-issue a fresh one whenever one of mine gets drawn.

revocable only until the key is drawn.
after that, the citizen must issue a new key to keep standing with the authority — an automation can do this in the background.

KINDS OF AUTHORITIES THAT CAN HOLD THE KEY

the authority serves its members — not the other way around.
one pull on the string ends the arrangement.

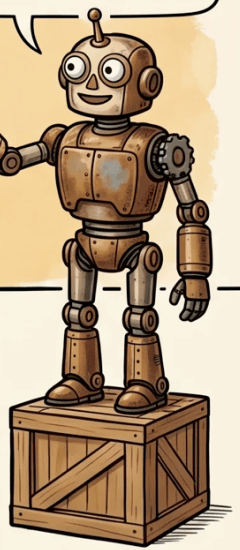

LEND THE KEY — KEEP THE STRING. DELEGATION IS NEVER SURRENDER.

Freedom vs. totalitarianism

A society that watches over and limits the volume of published records will be freer; a society with an excess of records will tend toward totalitarianism. One can envision delegating the right to publish records to someone else.

Cultural Adaptation

People generally enjoy seeing onto each other's plates. Across civilizations we differ in historical experience and population density — and therefore in how much we want to see onto each other's plates within civilizations and communities. Asian states tend toward a stricter view of freedom in society; European and American ones, in certain aspects, toward a freer one. However imprecise this parallel may be, in the context of a decentralized reputation social network it can be said that the settings of individual communities do not discriminate, and each community can configure communication within itself as it sees fit.

Despite the differences, there will be standards recognized across the planet (random killing is good for nobody). Other settings will be community-specific, and still others even personally specific. The advantage is that through mutual communication and by observing the settings of one's surroundings, consensus can be sought — from bilateral arrangements through community ones and, over time, thanks to the economic dominance of the best-configured communities, even at the global level.

But more on consensus in the next chapter.

The Oracle Problem — Bridging the Digital and Physical Worlds

The term "oracle problem" has in recent times been most often associated with the world of cryptocurrencies, blockchains, and decentralized systems. Generally: how do you ensure that data entering a digital

EACH COMMUNITY CONFIGURES ITSELF

THE NETWORK DOES NOT DISCRIMINATE — EVERY REGION SETS ITS OWN DIAL

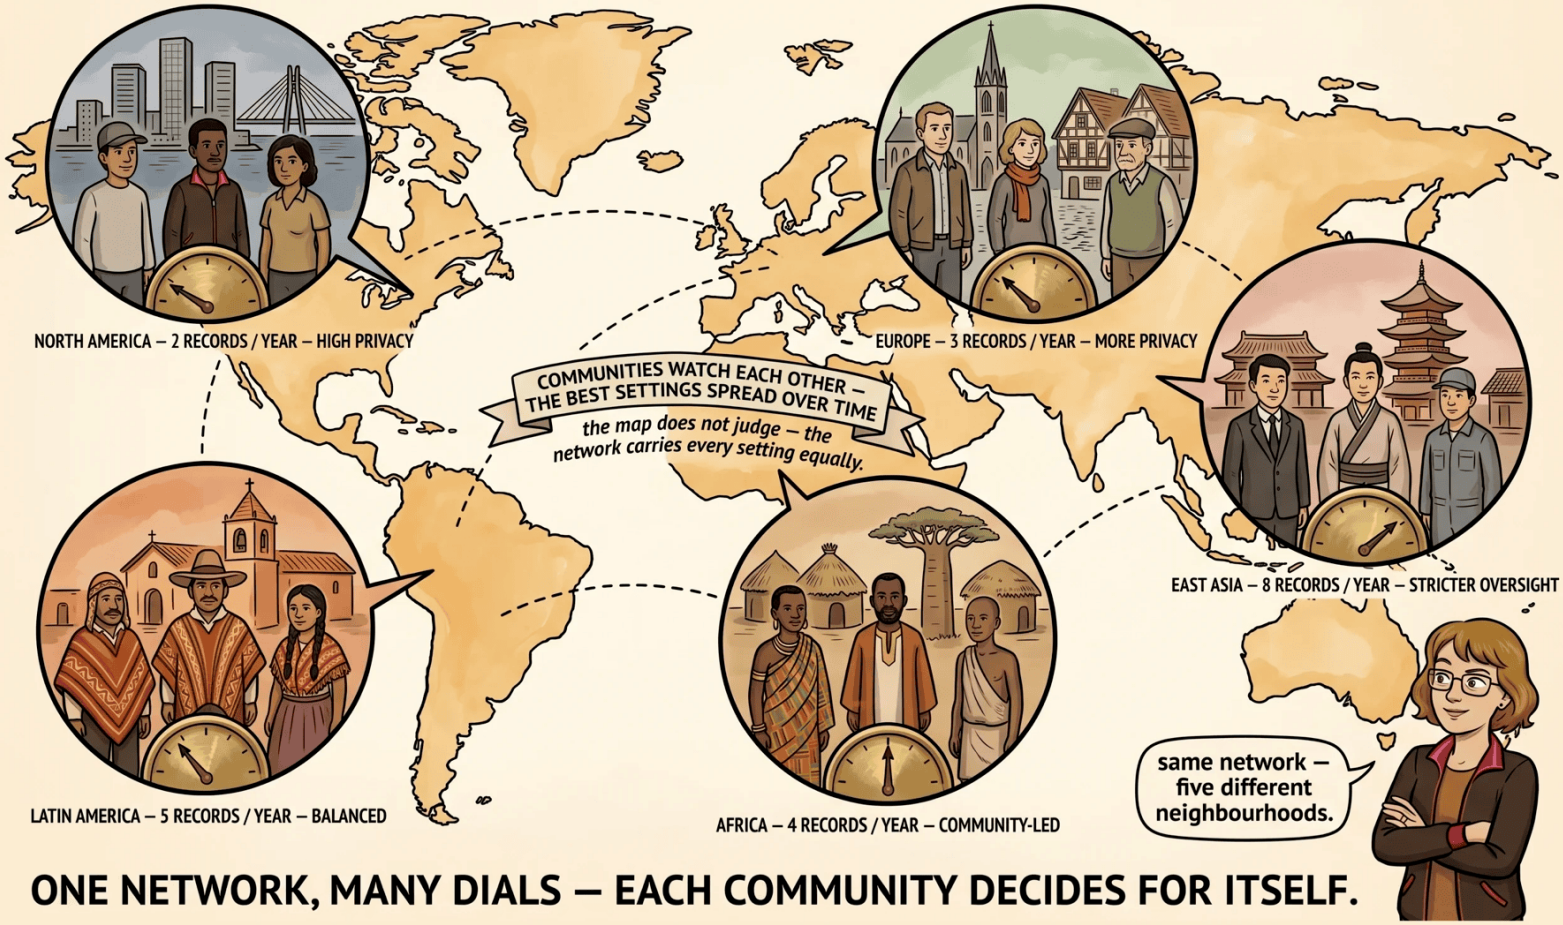

ONE NETWORK, MANY DIALS — EACH COMMUNITY DECIDES FOR ITSELF.

system faithfully corresponds to what actually happened in the physical world? The system itself cannot verify whether someone really harmed someone else, or whether a contract was truly breached — it only sees digital records.

Here we solve this through authorities — entities that stake their reputation on the digital record corresponding to physical reality. The higher the authority's reputation, the more credible the record.

BRIDGING THE DIGITAL AND PHYSICAL WORLDS

THE AUTHORITY CARRIES REALITY ACROSS

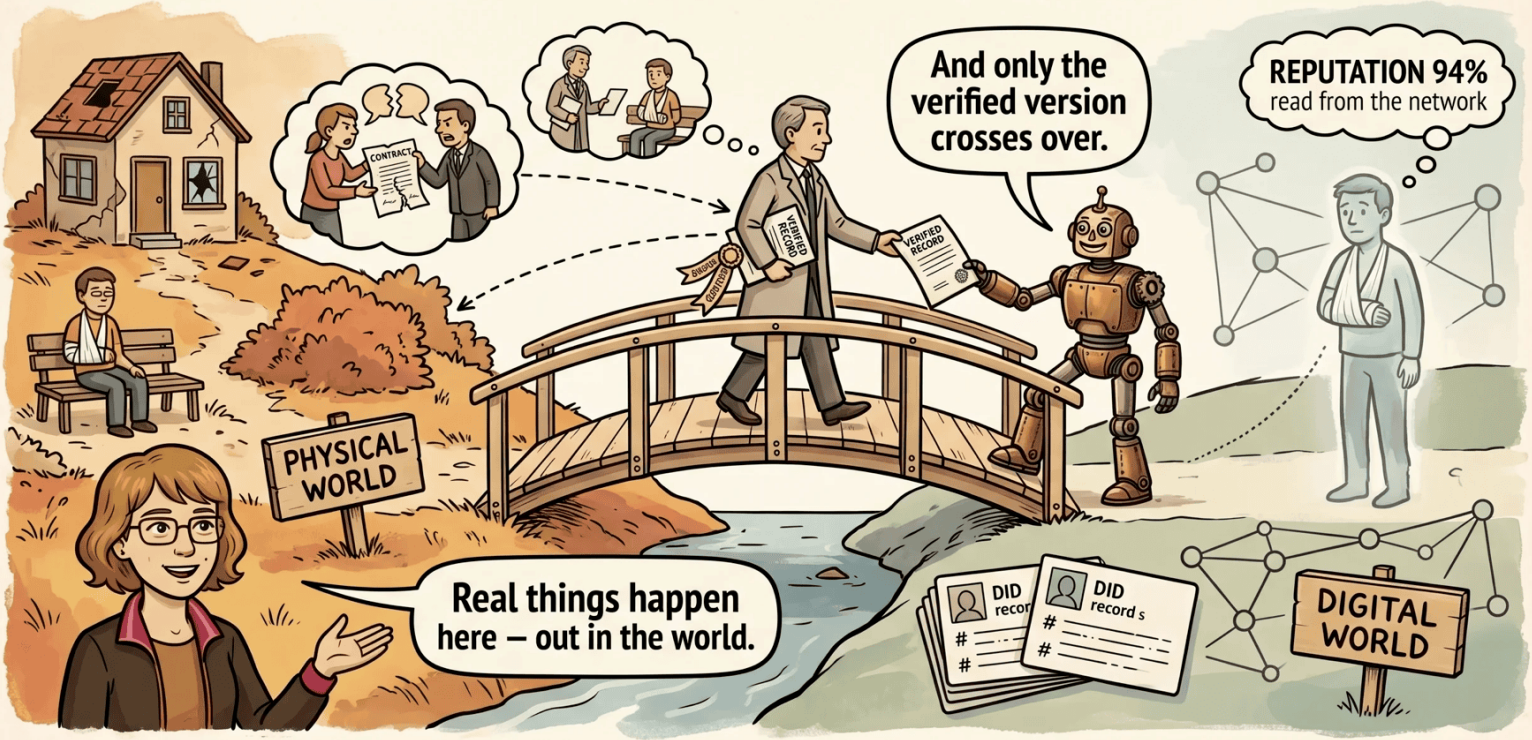

The authority's reputation makes the digital record trustworthy.

* The system does not try to solve the oracle problem — it accepts that authorities are the bridge, and aligns with how trust already works in the real world.
* If the bridge is wrong about reality, the authority pays — with their reputation.

How Verification Works

Consensus and the Verification Process

To build consensus on which rules a society should, on average, uphold and enforce, the following mechanism can help. As a DID participant, I declare the rules I subscribe to and will live by, and I publish them. (Think of it as the by-laws and statutes that, in my view, make up my ideal world — a world where I do not feel restricted, but safe.)

I can estimate in advance how my DID contacts would react — and assess how strongly, and by whom, I would be sanctioned in ordinary social or business interactions, should they hypothetically occur.

The definitive evaluation happens when you request information from another DID, or ask them to verify a claim (or ask an authority for a service, and so on) that you want to publish to the reputation network. It should turn out the same way as it does when you run the evaluation yourself, in dry run, against the counterparty's declared policy — and if it doesn't, something is wrong on the counterparty's side: they are trying to play a dishonest game.

The outcome is either acceptance, with a quoted price for verification (in the case of verifier or authority services), or rejection. Both sanctions and bonuses for deviation from the evaluator's policy are folded into the quoted price. The requester then decides whether to accept the terms, or move on to the next round of verification in the allocation algorithm — repeating the process until satisfied, or until the economics make it pointless to continue.

The Social Graph

The reputation network is, first and foremost, a social network. You add contacts — people who consent to the connection. They have contacts, and those contacts have contacts. The algorithm searches for verifiers within a configurable depth (e.g., three levels: your direct contacts, their contacts, and one level beyond). No global blockchain is needed — the network naturally forms communities with overlaps into other communities.

The algorithm is nondeterministic: it hashes your claim document, maps the hash to a position on a ring of known identities within this circle, and selects the nearest one as the candidate verifier. You cannot predict or influence who will verify your claim.

Each verifier's rejection enlarges your document and increases its processing cost — that is the first cost channel (document growth). Each new verifier charges a fee based on data volume, your reputation, and how far the content of your claim deviates from their declared verification policy — that is the second cost channel (risk premium). And each iteration costs time and energy — the third cost channel.

What the Verifier Checks, in Order

Once selected, a verifier evaluates a claim in roughly four ordered steps — cheapest filters first, expensive content checks last:

1. **Policy gating.** Does this kind of claim fall within what the verifier publicly verifies at all? If not, the request is rejected outright.
2. **Authority trust.** Is the authority that endorsed the claim trusted enough under the verifier's own declared policy? An authority below the verifier's trust threshold is grounds for rejection regardless of the claim's content.
3. **Issuer reputation.** Does the issuer meet the reputation thresholds the verifier has declared for this type of claim? Low reputation may either raise the fee or trigger rejection.
4. **Content check.** Only when the first three gates pass does the verifier evaluate the claim itself — signatures, internal consistency, formal correctness, and how far it deviates from the verifier's policy. The fee charged for this last step reflects the actual risk taken.

The verifier publishes the policy that governs each of these gates, so the steps are not at their discretion — they are bound by what they have already declared. Deviation from the published policy is itself a publishable claim against them, and they pay for it with their reputation.

The result: publishing a credible and useful claim costs almost nothing. Publishing a radical claim costs more. Publishing a lie becomes prohibitively expensive — you must iterate through verifier after

verifier, and every one who rejects you adds costs. The market prices your claim, and the price tells you where you stand in relation to the communities you move in.

It is not enough to declare that you adhere to a rule when in reality you do not. In that case, your DID risks the publication of a negative record exposing the hypocrisy — which turns you into a risk for everyone else. The outcome should be fewer but more consistently followed rules, and a clearing of that jungle of laws and regulations that even legal professionals can barely navigate.

Consensus vs Accountability

For the network to serve as a valuable source of information, a DID should not be too radical — otherwise the others will reject it. Social pressure will seek equilibrium, and attempts to destabilise it will likely be punished.

HYPOCRISY IS THE MOST EXPENSIVE BEHAVIOR

THE NETWORK QUIETLY COMPARES YOUR WORDS WITH YOUR RECORD — THE GAP IS VISIBLE TO EVERYONE

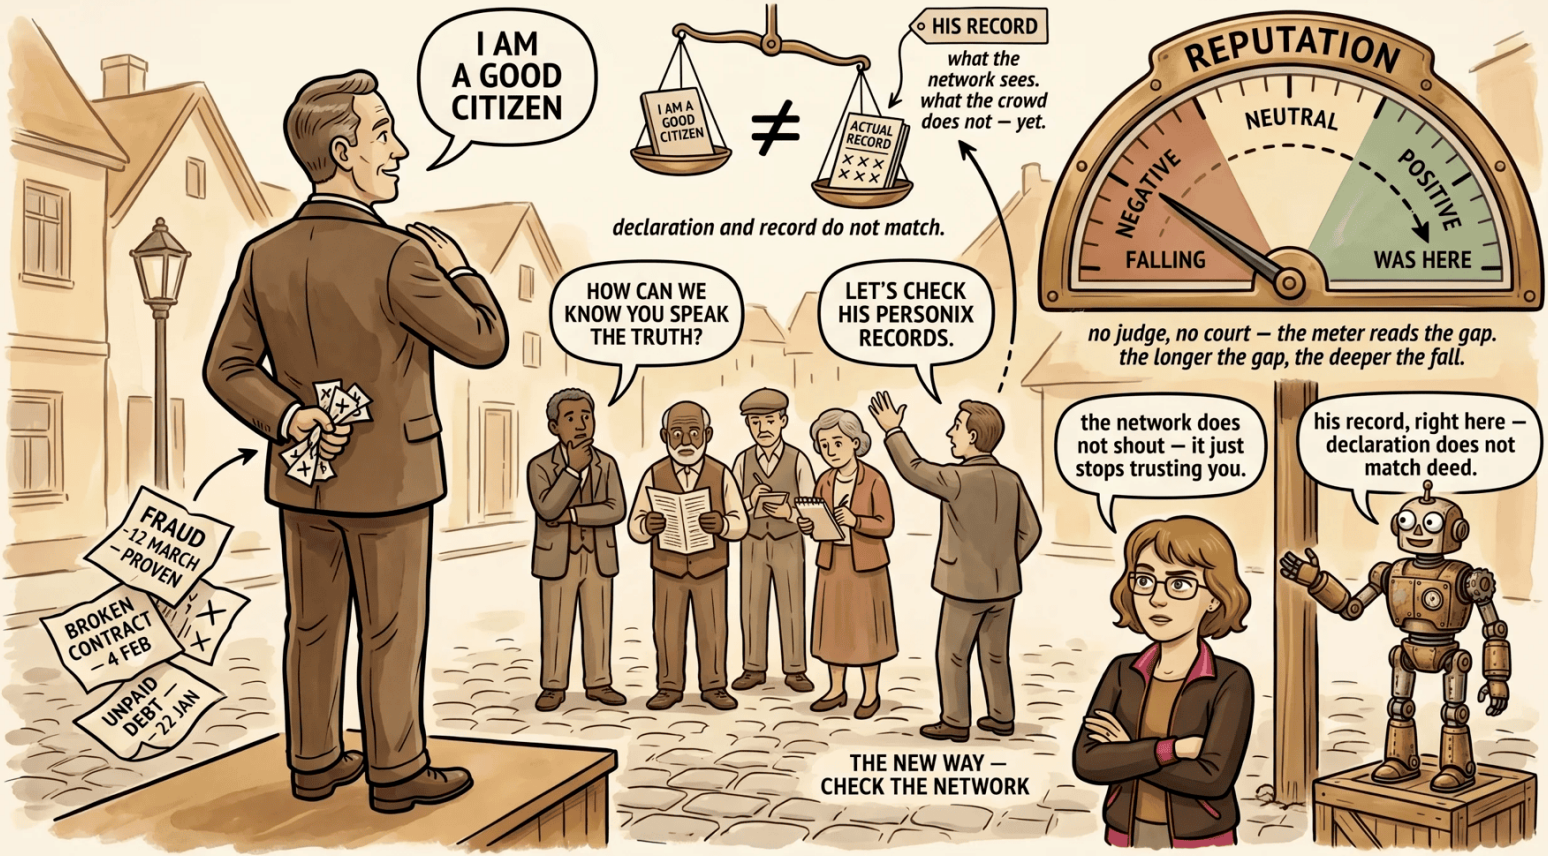

DECLARE IT AND BREAK IT? THE GAP IS THE MOST EXPENSIVE BEHAVIOR.

THE CURE FOR HYPOCRISY

DECLARE YOUR RULES — LIVE BY THEM, OR THE MARKET PRICES YOU OUT

I DECLARE MY POLICY

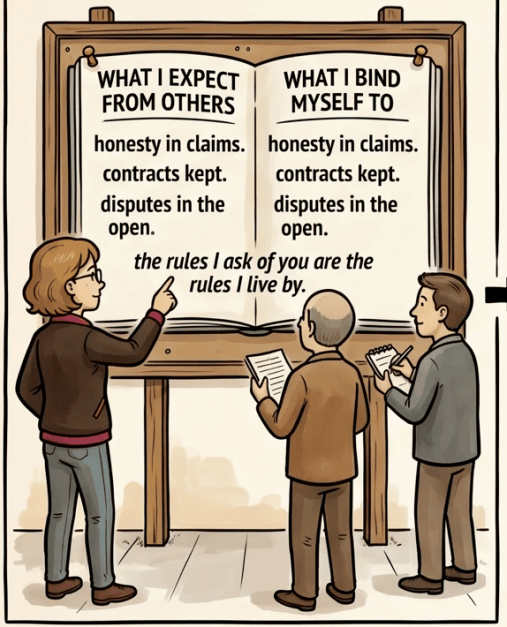

everyone can read my standards.
no gatekeepers, no secrets.

THE WORLD WATCHES

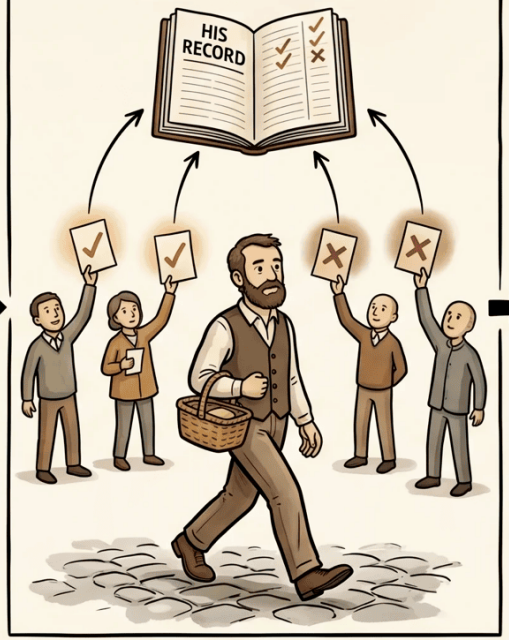

the world writes down what he actually does.
every day, a little more evidence on his ledger.

AT THE SHOP — YOU PAY FOR YOUR OWN RECORD

hypocrisy is not punished by authority — it is priced by the market.
every ordinary service you buy carries a surcharge shaped by your own record.

COULD WE CALL THIS THE SOCIAL CONTRACT — REDRAWN?

no authority forces the ethic — the price does.
soft pressure, every day, toward the shared average we can all live with.

WALK YOUR TALK — OR PAY FOR IT. HYPOCRISY IS THE MOST EXPENSIVE REPUTATION.

The Number of Votes Is Not the Same as the Weight of a Voice

Juraj Karpiš says that "money is the memory of good deeds." I would add that reputation is the memory of the bad ones.

It follows that, meritocratically, whoever contributes more and has no bad reputation deserves a greater weight of voice in the community. Looked at through the lens of bilateral relationships: when I weigh which consensus pressures to accommodate, the greatest weight goes to the relationships from which I derive the largest economic benefit. Ten people with whom I have no active trade will influence me far less than one permanent business partner. This paradigm is not limited to commerce — it extends to social, political and other relationships.

Roles in the verification transaction and authorities

Roles Overview

We already touched briefly on some of these roles in the chapter about the network and its basic properties. Now is the time to look at them again in more detail and add the additional ones we need to make the network more robust. Every verification transaction involves several roles — let's see how they behave.

Roles in a Verification Transaction

Each verification involves up to six distinct roles, summarised in the table below. All of them can have their own DID in the decentralized reputation network.

Role	Description
Issuer	The person who publishes information to the network — claims that something happened (a DID was created, edited or dissolved, a claim, the policy of a given DID, etc.)
Subject	The person the information is about — the addressee of the claim
Authority	A trusted entity that stakes its name on the quality of the claim by investigating it and either reviewing the evidence presented or actively gathering it
Observer	An independent third party who keeps a record of how the verifier handles the claim — making sure the verifier neither stays silent nor deviates from the policy they declared
Verifier	An algorithmically selected participant who processes the transaction
Delegate	A person acting on behalf of another participant

Issuer

We have met the issuer before. When someone is wronged, or wants to broadcast an accomplishment of their own, it is the person represented by the DID who triggers the process of putting that information

into the network — because they have the primary interest in doing so (the motivation being either a desire for satisfaction or a wish to improve their public image).

Positive Claims

Positive claims about myself I can keep, as both their subject and their issuer, for the situations in which I want to use them (a diploma, a certificate of service rendered, etc.). This is where a positive claim differs from a negative one: sharing a verified negative claim is in the interest of the whole community, whereas a positive one is in the interest of its bearer.

Authority

The authority plays a dual role: it can be an **auditor** (verifying the quality of evidence before a claim is published) or a **guarantor** (staking its reputation on the truthfulness of a claim). In either case it strengthens the issuer's claim. These two services are separable — an authority may offer one, the other, or both at once. The working assumption is that most services provided by authorities can be delivered on a free-market basis. That holds even in fields that are hard to imagine being privatised, such as justice, where specialised services — investigation, evaluation of evidence, all the way to services today provided by centralised armies (strategic planning, standardised training, procurement and stockpile management, etc.) — can be efficiently delivered by market actors. There is hardly anything that, after restructuring, could not be made more efficient by free-market incentives.

THE ISSUER — WHO WRITES THE CLAIM

AUTHOR · SIGNER · STAKES THEIR OWN REPUTATION

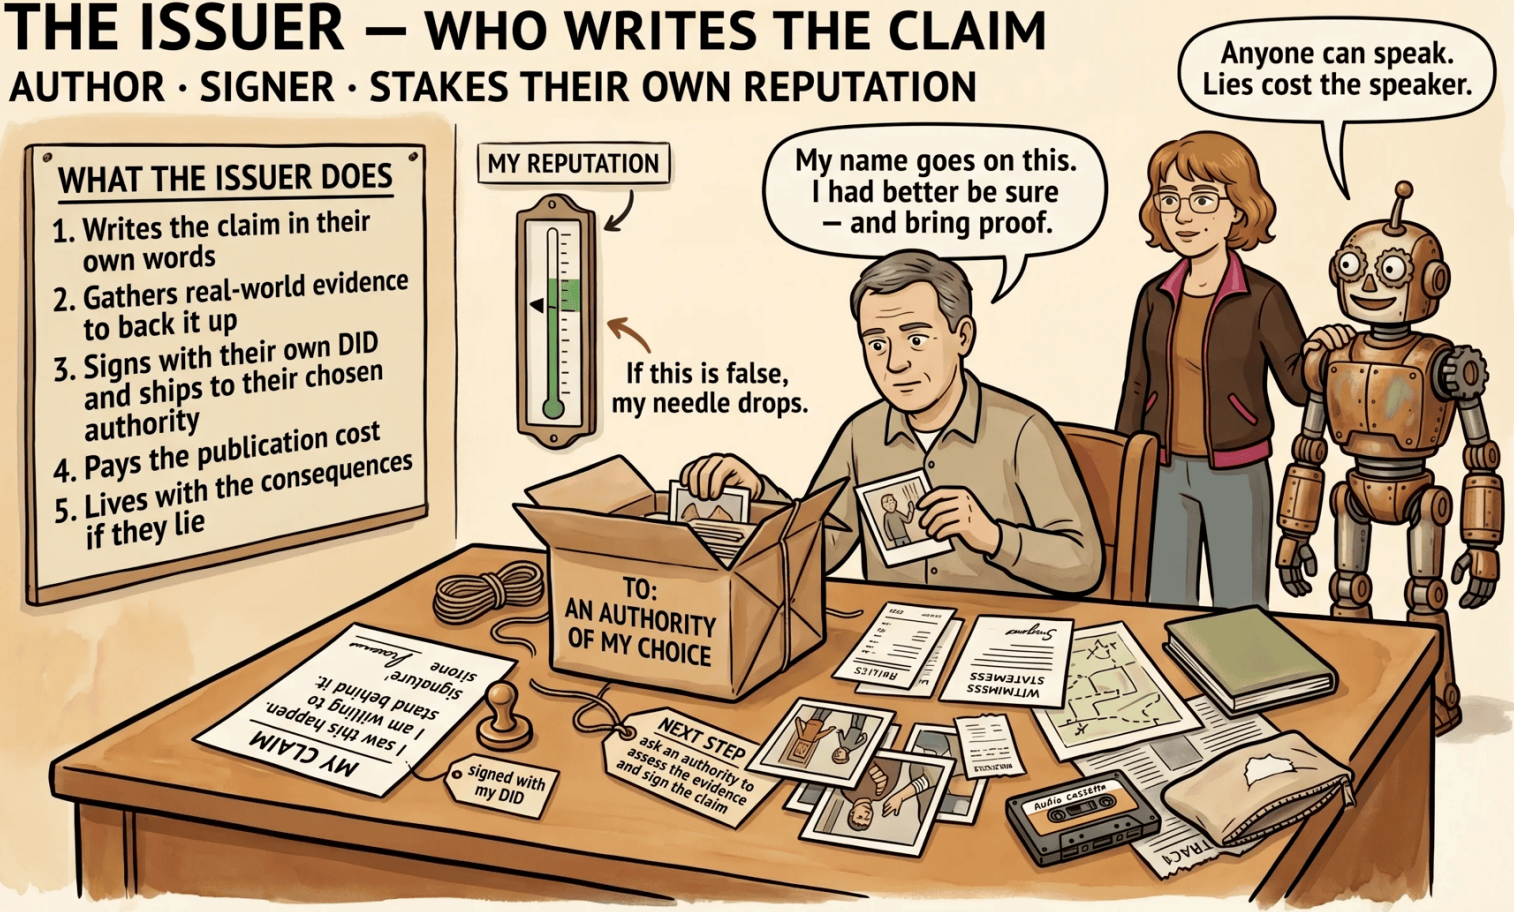

The issuer is the author. Their name is on it.

---

* Anyone with a DID can be an issuer. The bar is courage and accountability, not a license.

The authority, the issuer and the observer must never be the verifier of their own case.

Algorithmic selection of the verifier guarantees independence. Nobody can verify their own claim, or a claim in which they have a direct interest. This is one of the basic rules that the whole community of DIDs has an interest in upholding.

The following graphics show complementary views of the breadth of activity that authorities cover (the term "authority" can be read interchangeably with "service provider").

THE GUIDE — HELPS YOU NAVIGATE BY YOUR VALUES

DEEP ANALYSIS BEHIND · SIMPLE GUIDANCE IN FRONT

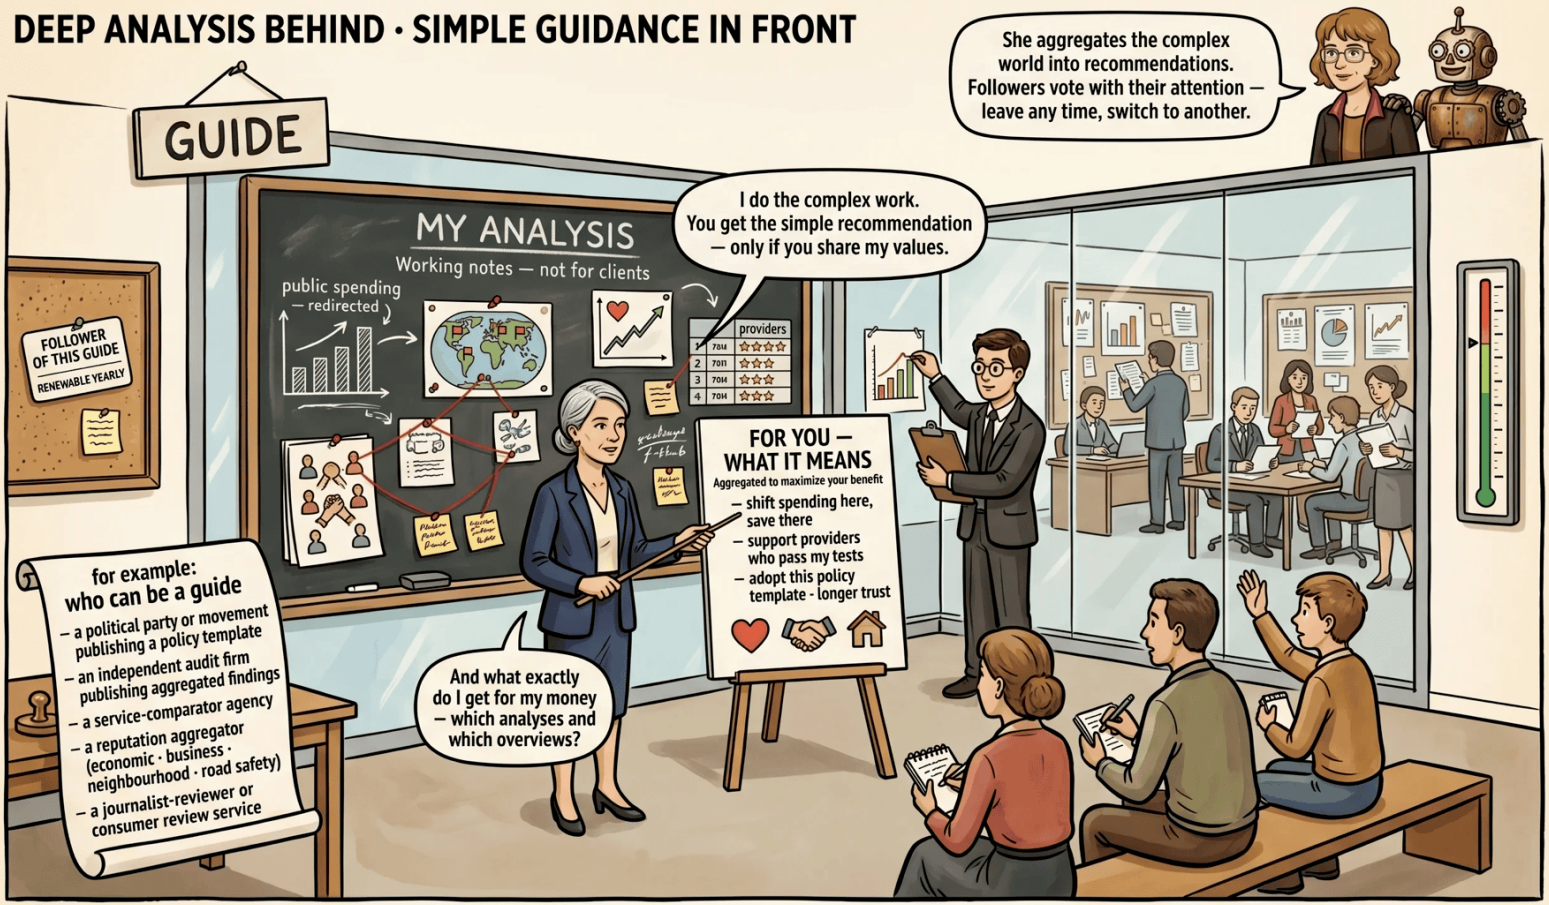

Bound by what they publicly stand for. Your trust is their currency.

* You can leave any time and switch to a different guide. Your trust is not locked in.
* The guide may also sign some claims of their own — published, but not the main story here.

Authority as an Incognito Observer

A reputable authority — think of a notary whose business depends solely on their track record — can, alongside the main functions (auditor / guarantor), offer a third one: the role of incognito observer during verification. They keep a time-stamped record of the submitted claim so that the verifier cannot quietly drop it. The mechanism of the observer role is described further in the section on the Observer role.

Verifier

Any DID can act as a verifier, either directly or through verification rights delegated to a third DID. For me — or my delegate — to be able to verify, I should be reachable on the network (online). Not everyone will want to commit to that, which is why a DID record can list, in priority order, the substitutes who will perform the function on its behalf while it is offline.

Every DID active in the network publicly declares its own policy. Through the rules defined in that policy it judges, during the verification process, the reputation of the counterparty and the content and form of the claim that the issuer has flagged for publication into the reputation network. Part of the policy is the calculation formula used to compute fees for verification services. Once that is in place, then across a statistically large number of claims flowing through the network I wait for the network's algorithm to draw me on the issuer's side and assign me, in a given iteration, to verify the information being issued. The issuer can compute in advance how a correctly behaving verifier would react, but

THE PROVIDER — RUNS A CLAIM-SIGNING SERVICE

PUBLISHED STANDARDS · PUBLISHED PRICES

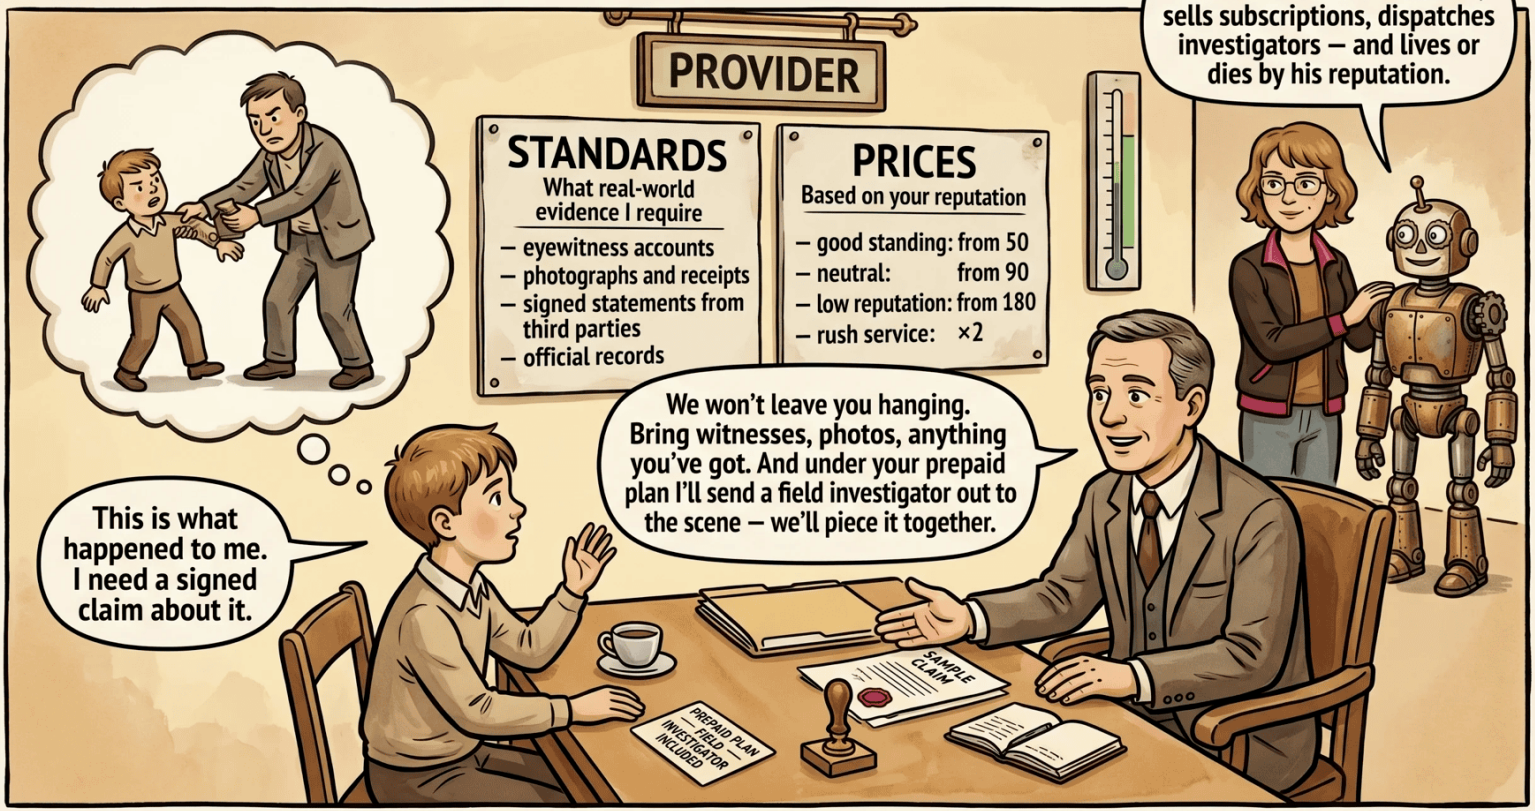

Bound by what they publicly declare. Reputation does the rest.

* The provider also has their own policy — when they accept, when they reject, how they refund. It is published, but it is not the main story here.

cannot avoid actually contacting them (or their substitutes); the iteration with the selected verifier has to be carried out by the issuer even when they know in advance it will not pass.

How do we know that the issuer runs the verifier-selection algorithm over the correct set of candidate verifier DIDs? Together with its publicly declared policy, every DID also publishes the current list of identifiers of its social network within the reputation network. If an issuer defines its social network as a social bubble that merely echoes and reinforces its own views, information published through it will hardly be received more widely by other communities. The fact that I manage, at high cost, to push a radical claim into the network does not imply that, when judging the counterparty's reputation, I will give it any weight. Some claims I am pushed by my community to take into account (sentences and restrictions imposed on offenders); others are entirely up to me — I decide for myself the economic value of including or excluding a given piece of information.

Subject

Those whom the information published into the reputation network concerns are its bearers, or, if you like, its subjects. Incentives in the reputation network are set up so that the subject has every reason to behave correctly and fairly by their own rules, and is pushed by the community towards some golden middle of rule-keeping. They even have an incentive to behave more morally than their declared rules require — pre-declared policies do not preclude the community from arriving at a consensus to apply some evaluation or measure retroactively. From this angle, the reputation-network alternative to a legal order is, as one would expect, driven more by the spirit of the rules than by their letter.

THE VERIFIER — JUDGES THE CLAIM AGAINST PUBLISHED POLICY

FOUR GATES, IN ORDER. CHEAP FILTERS FIRST.

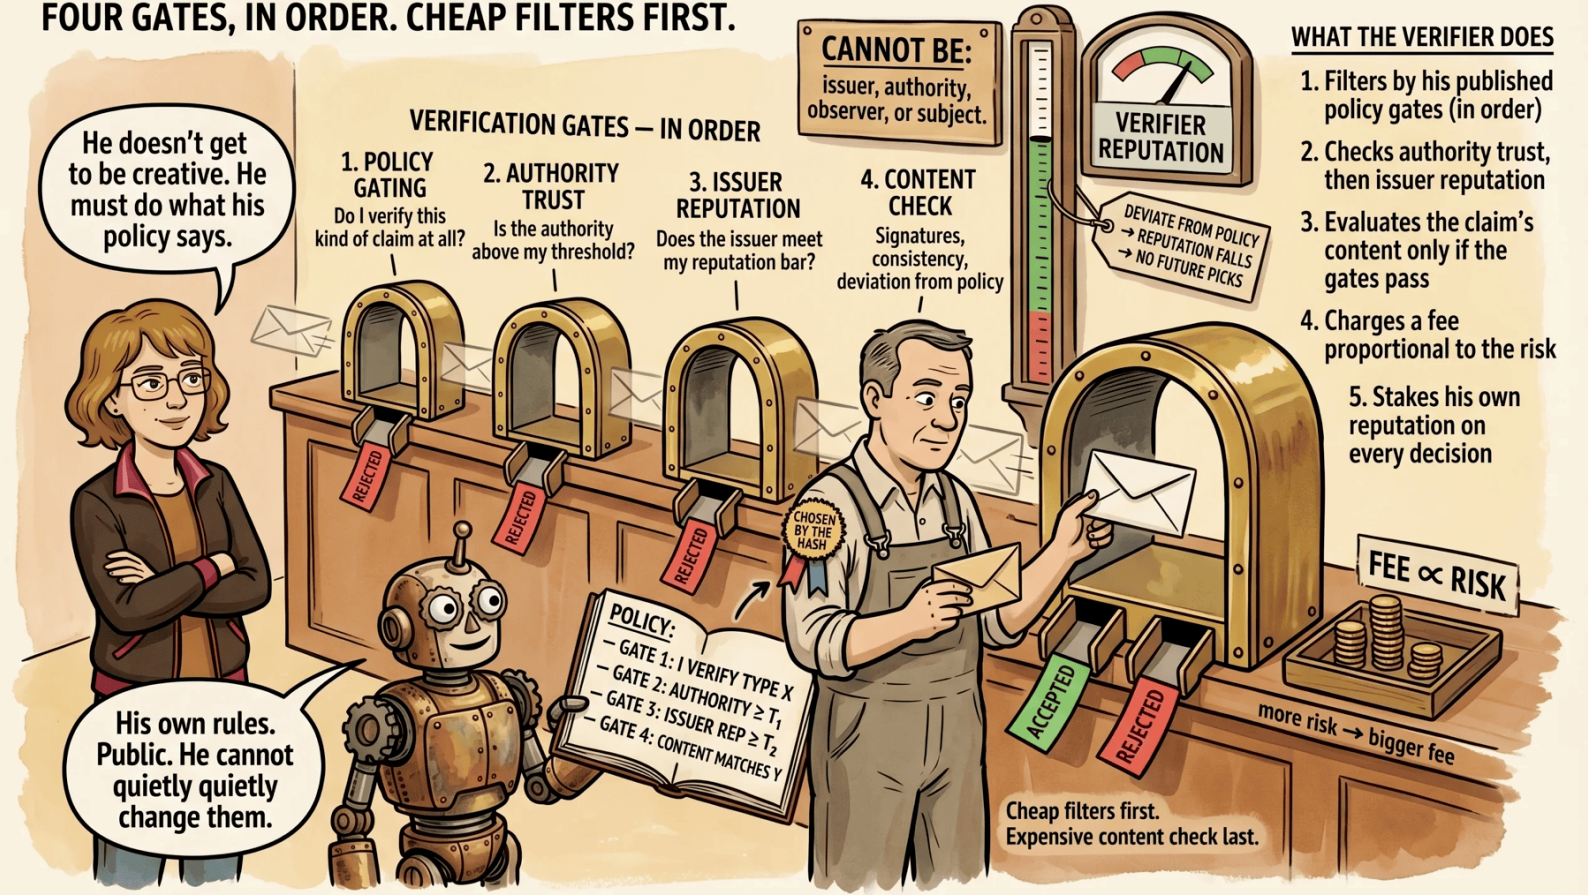

WHAT THE VERIFIER DOES

1. Filters by his published policy gates (in order)
2. Checks authority trust, then issuer reputation
3. Evaluates the claim's content only if the gates pass
4. Charges a fee proportional to the risk
5. Stakes his own reputation on every decision

The verifier is bound by the rules he himself published — and by the reputation he stands to lose.

THE SUBJECT — THE ONE THE CLAIM IS ABOUT

RECEIVES IT · CAN'T SILENCE IT · CAN ANSWER IT

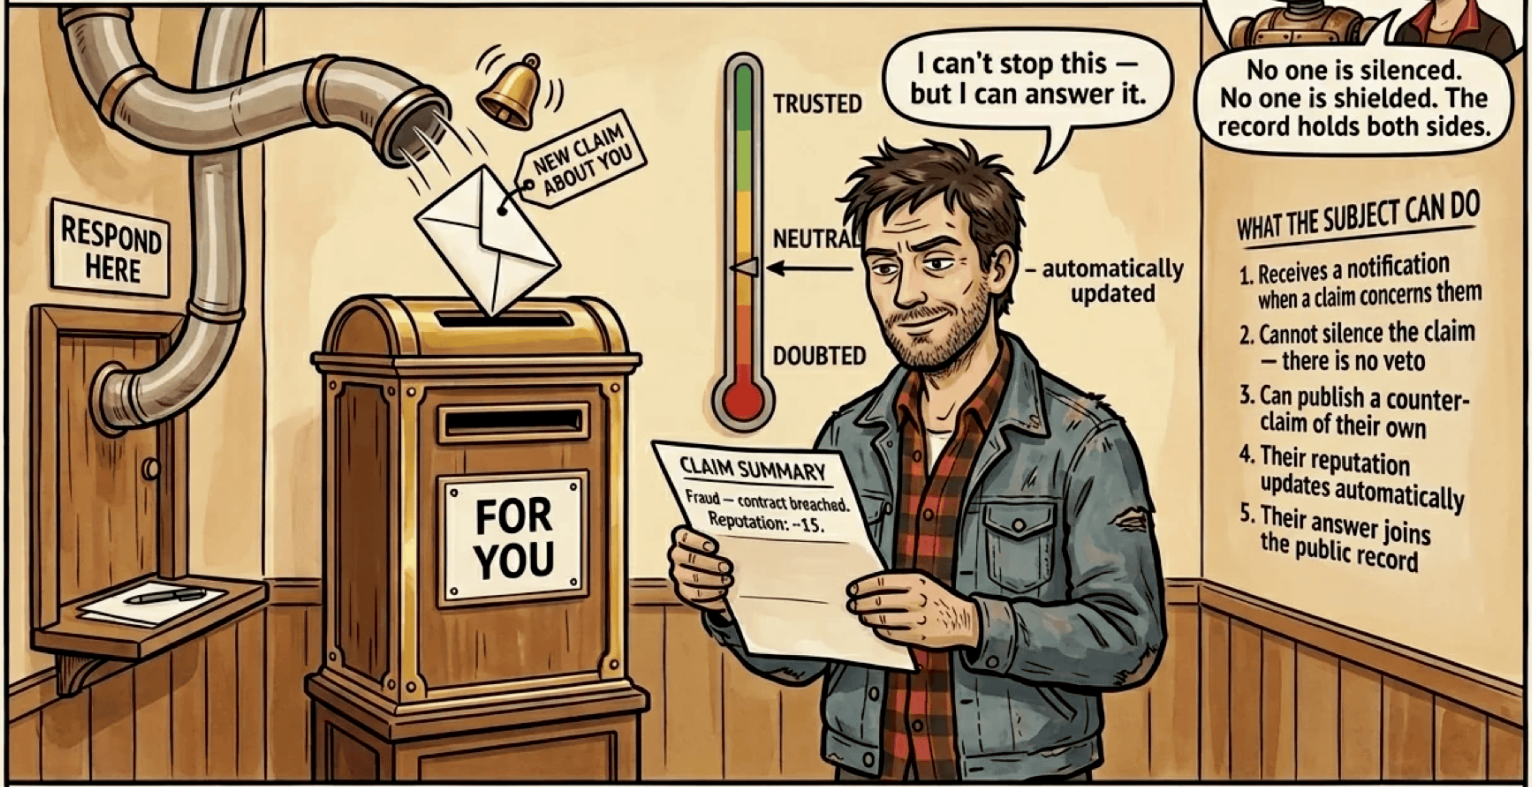

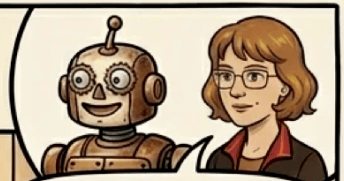

WHAT THE SUBJECT CAN DO

1. Receives a notification when a claim concerns them
2. Cannot silence the claim — there is no veto
3. Can publish a counter-claim of their own
4. Their reputation updates automatically
5. Their answer joins the public record

The subject has no veto — only the right to answer.

* A counter-claim is itself a claim — the subject's reputation is on the line for what they say in reply.

Observer

The observer role removes the verifier's incentive to bend the rules. In situations where a verifier doesn't like the issuer's or the authority's request, they could simply stay silent — not respond, and block the algorithmic sequence. The observer — or a set of observers — stakes their reputation on documenting how the verifier was queried. If the verifier stays silent despite a declared policy that says otherwise, they can be convicted of violating the protocol.

The mechanism: timestamp and challenge code

Before you send a claim to the verifier, you route it through observers — people you trust, or specialised observer-service providers who charge a small fee. Each observer receives your submission, time-stamps it, signs that they saw it go out, and generates a challenge code — a cryptographic hash of their signature. The codes are appended to your request. The verifier sees them but has no idea who the observers are, or whether the codes are even real. Observers thus act as proxies between the issuer and the verifier, holding an independent record that the claim was submitted and what it contained. There can be zero to N of them.

When the verifier behaves honestly — accepting or rejecting in line with their declared policy — the codes stay opaque. Nobody is exposed.

But if the verifier stays silent despite an accommodating policy, or responds in a way that contradicts what they published, you hold the original observer signatures. You can publish them as proxy testimony that the claim was submitted and that the verifier did not follow the protocol. Anyone can verify that the signatures match the challenge codes.

THE OBSERVER — A PROXY THAT RELAYS THE CLAIM AND WATCHES THE ROUND-TRIP

THE **VERIFIER** TALKS TO HER. THE VERIFIER **NEVER LEARNS WHO SHE IS.**

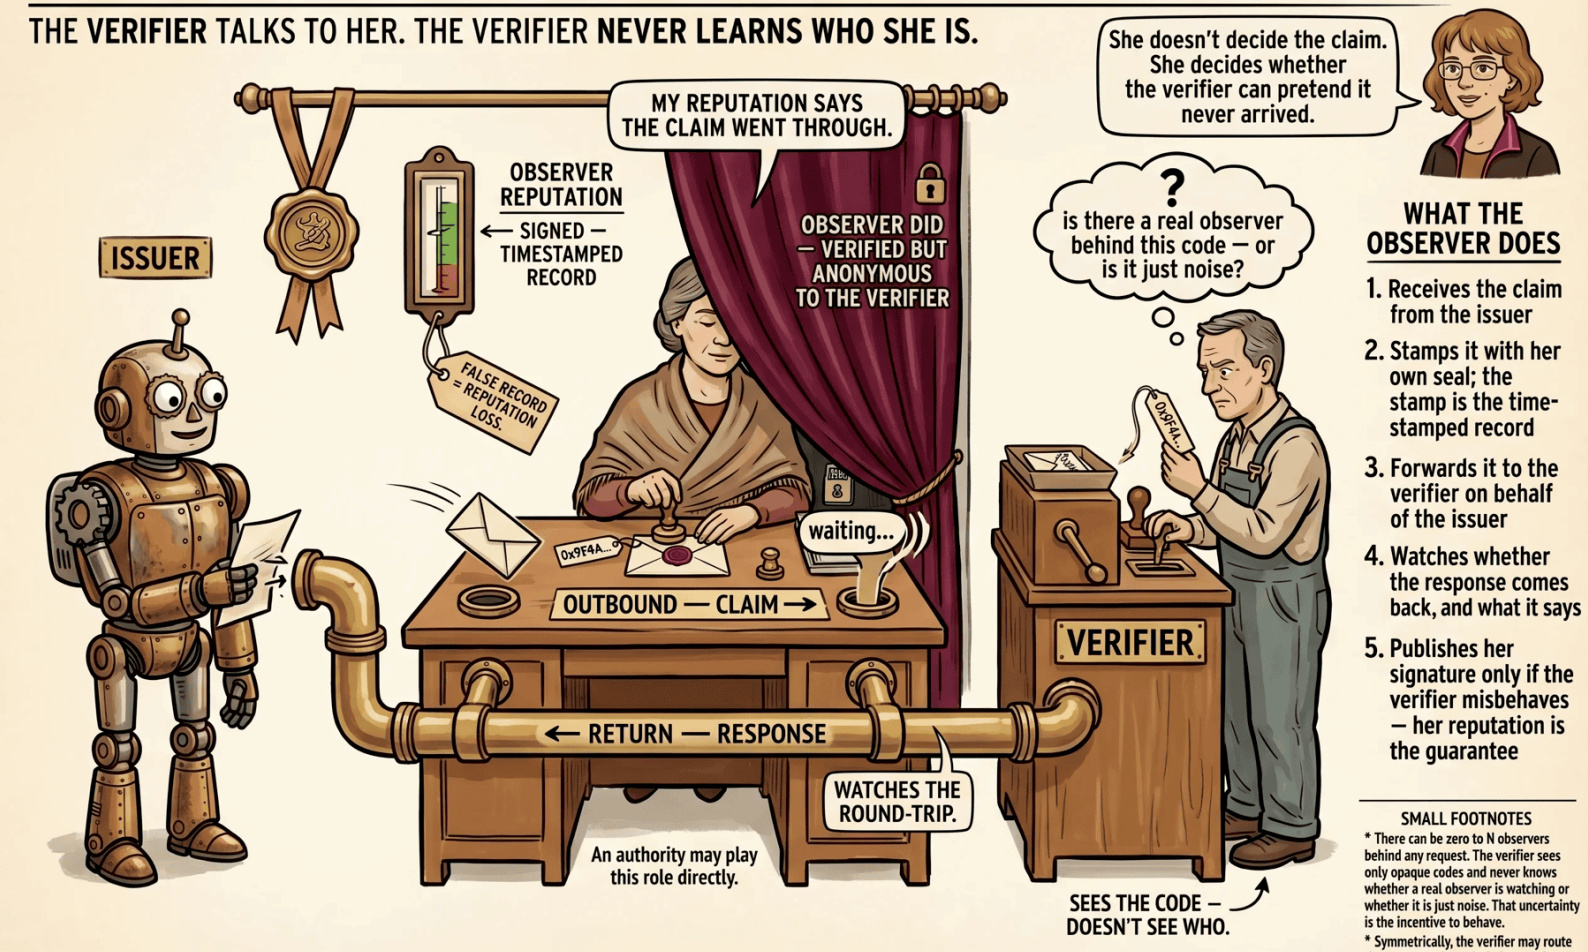

WHAT THE OBSERVER DOES

1. Receives the claim from the issuer
2. Stamps it with her own seal; the stamp is the time-stamped record
3. Forwards it to the verifier on behalf of the issuer
4. Watches whether the response comes back, and what it says
5. Publishes her signature only if the verifier misbehaves — her reputation is the guarantee

Anonymous to the public, accountable to the network. The verifier sees the code, never the observer.

SMALL FOOTNOTES

* There can be zero to N observers behind any request. The verifier sees only opaque codes and never knows whether a real observer is watching or whether it is just noise. That uncertainty is the incentive to behave.
* Symmetrically, the verifier may route the response through their own observer — – same trick, opposite direction.

The punchline: you don't need real observers

And here is the most elegant part: **you do not need real observers at all.** You can generate random numbers that look exactly like challenge codes. The verifier cannot tell the difference — they have to roll the dice on whether to risk their reputation. Behind every request they receive there could be a respected observer watching incognito — or it could be pure noise. The verifier does not know. And that uncertainty is the mechanism.

The cost of maintaining honest pressure: nearly zero (random numbers are free). The potential cost of dishonesty for the verifier: catastrophic. Honest behaviour is incentivised even when nobody is actually watching.

The system works because everyone is a little paranoid. Uncertainty is cheaper than surveillance.

Multiple verifiers in a single iteration

A reinforcing companion rule for verifier availability can be an algorithmic extension that returns, in a single iteration, a set of candidate verifiers rather than just one.

Delegation

Civilisation has built its economic growth on the capacity for division of labour and specialisation. The world in its full splendour is complex, and our chances of grasping it in all its complexity are close to nil. The natural way out is to delegate a wide range of activities that support a comfortable life to specialised

THE BLUFF THAT KEEPS THE VERIFIER HONEST

A REAL OBSERVER AND A RANDOM BLUFF LOOK IDENTICAL FROM THE VERIFIER'S DESK

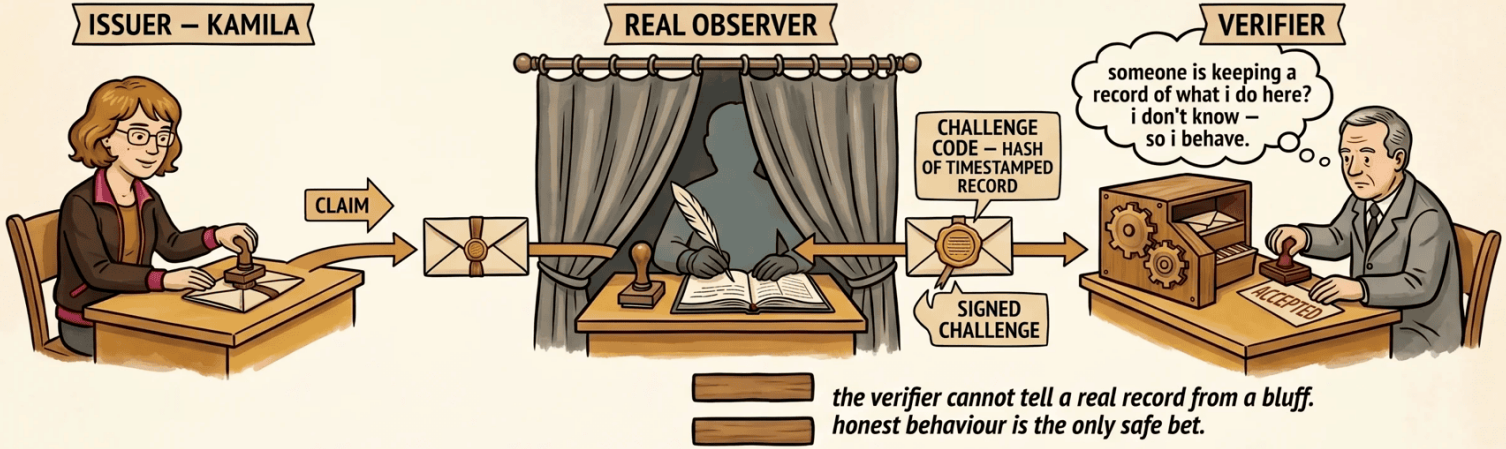

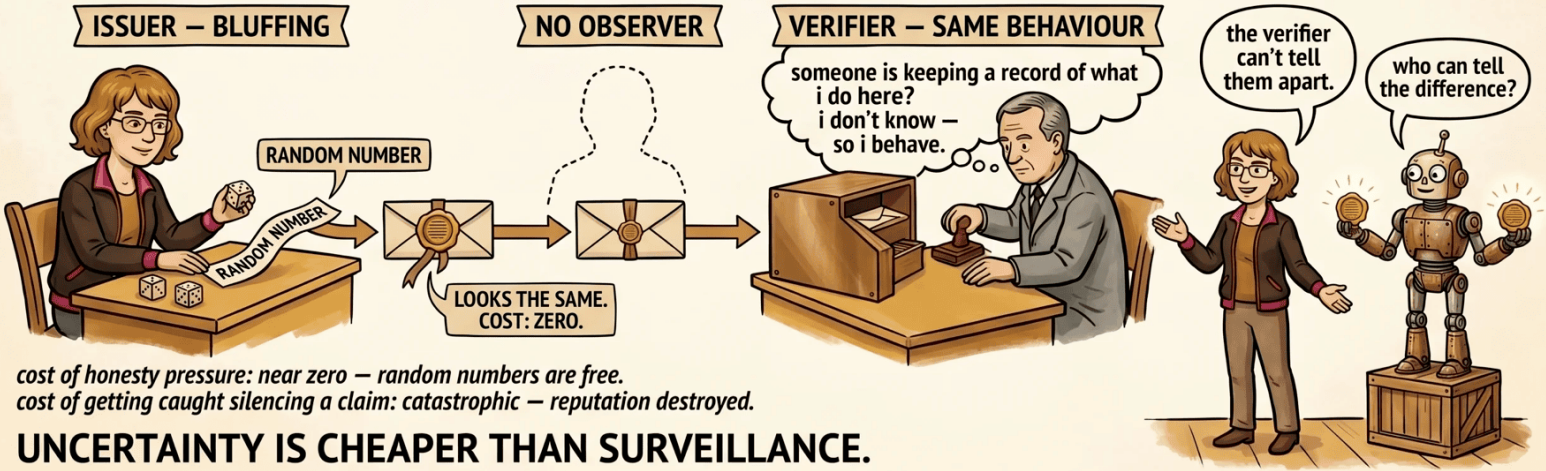

cost of honesty pressure: near zero — random numbers are free.
cost of getting caught silencing a claim: catastrophic — reputation destroyed.

UNCERTAINTY IS CHEAPER THAN SURVEILLANCE.

providers, who save us time, energy and money, and let us specialise in turn — providing the services that we are best at, or that the market values most highly in us.

The same logic of delegation extends into the reputation network. In the chapter The Freedom–Totalitarianism Switch (and Delegation) we set out the principles for safeguarding a free society by limiting the number of claims that can be published into the network within a given time window (think, in simplified terms, of low single-digit publication rights per year, so that society stays maximally free). This principle can be extended to multiple dimensions. We do not have to have just one number — we can have several limits, applying to different types of claims. We can use a different number for flagging others' hypocrisy than for misconduct in business dealings, and, say, a change of the DID document such as a password change is something the community will hardly care about at all, and so on. It all depends on how the community defines, over evolutionary time, what serves its interests in finding the right balance between its target level of freedom and accountability.

Why should the number of publication rights granted to me by the community be small enough? Flagging a DID (subject) with a negative claim destroys reputation that took a long time to build, and the damage lasts. A destroyed reputation faces a scale of escalating sanctions from the community for every further or repeated offence, all the way to the complete elimination of the disruptive entity (the extreme case). An exponentially escalating scale of punishment does not need a large volume of claims, provided they are backed by verifiable evidence. This holds for the peaceful steady state, when society is not being shaken by a disaster (external or internal).

For events with very low probability, statistical frequency is volatile when applied to an individual. Across the community's population it is not. That is the value of delegating one's publication rights to specialised aggregating authorities, which — much like mutual insurance in the event of a loss — use your right and apply it on your behalf. Let us speak less abstractly and give a few examples: - detecting and

THE DELEGATE — A CLAIM-SERVICE YOU SUBSCRIBE TO

DELEGATE WHAT YOU CAN'T HANDLE YOURSELF · KEEP WHAT MATTERS TO YOU

Outsource what you can't handle yourself. Stay in the loop. Pull the rights back any time.

* If the delegate acts outside the agreed scope, both their reputation and yours are publicly affected.

combating organised crime - enforcing commercial agreements - managing privileges (proofs of education, acquisitions / disposals of assets, etc.) - and so on

How “publication rights” actually work

The graphic for the delegate shows publication rights as physical tokens for clarity, but in the system no such object exists. What actually happens is this: the policies of the people you want to remain in good standing with — your community — aggregate into an emergent rule about how many claims per year they will accept from you as relevant before they start treating you as noise. The cap exists so that the network is not drowned in trivial “what if” claims and stays as free as the community wants it to be (signal preserved by limiting volume).

You read what your community allows you. There is a range, with some spread depending on which people you weight most. You then decide — within that allowance — how many claims to publish yourself, and how many to delegate to professional services that will publish on your behalf only when events fall within their scope. This is what keeps you in good standing without forcing you to file a claim for every minor thing that happens to you.

The “tokens” in the graphic are a visual stand-in for this read-and-decide loop, not a literal artefact.

It is worth being able to revoke delegated publication rights in the reputation network. This can be done at any time, as long as the delegate has not yet used the delegated right.

A NEGATIVE CLAIM FLOWS THROUGH THE SYSTEM

FOLLOW ONE DOCUMENT — SEE WHAT EACH ROLE DOES TO IT

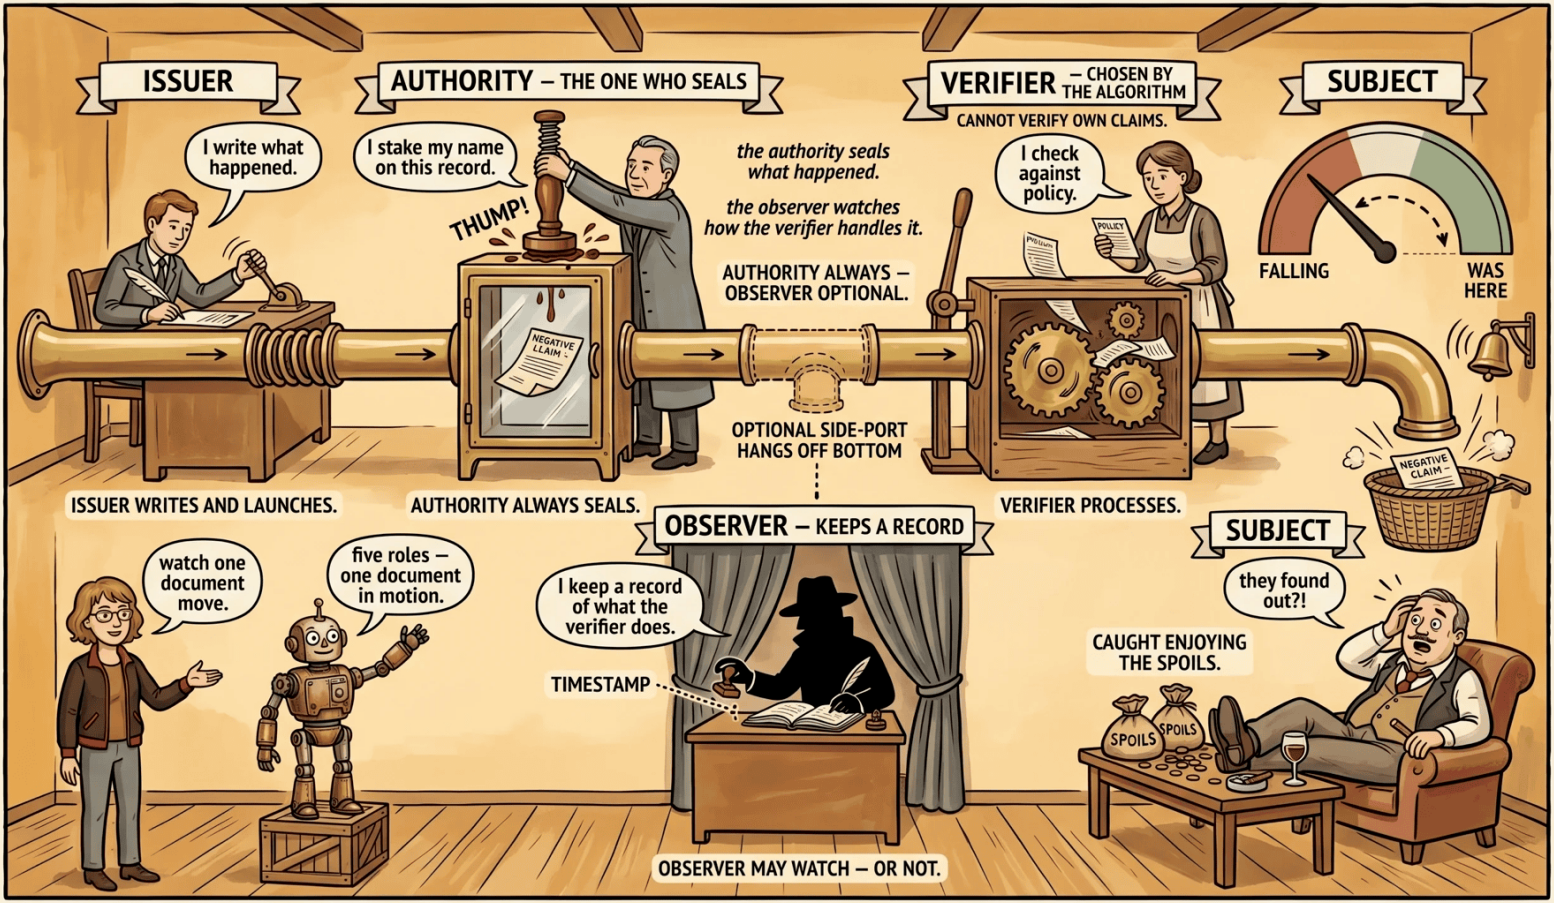

FIVE ROLES, ONE DOCUMENT IN MOTION — NOBODY CAN VERIFY THEMSELVES.

Communication Between Roles

The graphic below summarises how the roles connect in a real case — from the submission of a claim by the issuer through to publication into the reputation network.

The Emergent Social Contract — Policy

So far we have described how claims are verified, and what part each role plays. Now we turn to how thousands of individual policies become something that works like a real social contract. Nothing that resembles the left-wing dreams of a social contract and enforced solidarity.

Every participant declares their own policy based on their morality (policy) — what they consider acceptable, and how they react to specific behaviour of others. The aggregate of all policies across the network known to a given DID (the community) produces emergent ethics — observable norms that nobody designed. These emergent norms, underpinned by bilateral price signals, form a measurable social contract: not a theoretical construct that nobody signed at birth, but empirically observable rules backed by actual behaviour.

The Social Contract

People debate whether "social contract" isn't an oxymoron. Nobody signs anything, yet it is invoked to justify legislative and executive action in centralised societies.

A phenomenological analysis of interactions in a decentralised network could reveal certain emergent rules — global or local. What should we call these broadly shared norms? A social contract? Just as with communities, a community's social contract cannot be observed and delineated precisely — every angle on the community yields a different social contract.

Values-Based Policy

It is similarly misleading to treat "values-based policy" as the foundation of decision-making. Values are more like phenomenological clusters of rules that are commonly preferred at a given time. They evolve, and they are not constant over time.

POLICIES MEET. ATTITUDES FORM.

DECLARED RULES SET TODAY'S PRICE — DISCOUNT FOR ALIGNMENT, MALUS FOR DIVERGENCE

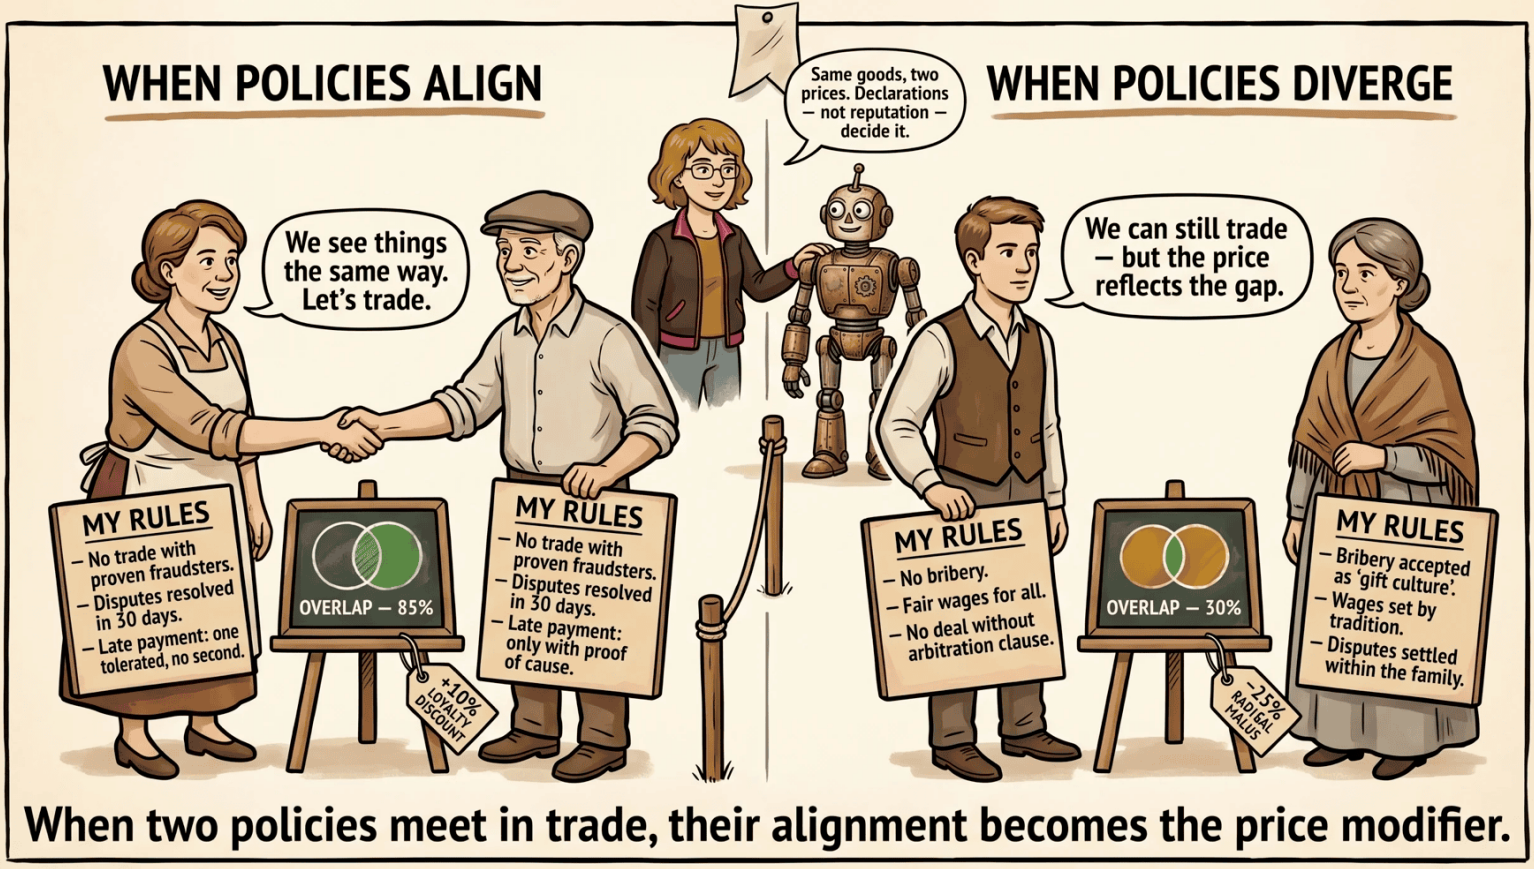

When two policies meet in trade, their alignment becomes the price modifier.

* This is not about reputation. It is about what each person publicly declared — and is now expected to keep.

HOW ETHICS EMERGE FROM INDIVIDUAL POLICIES

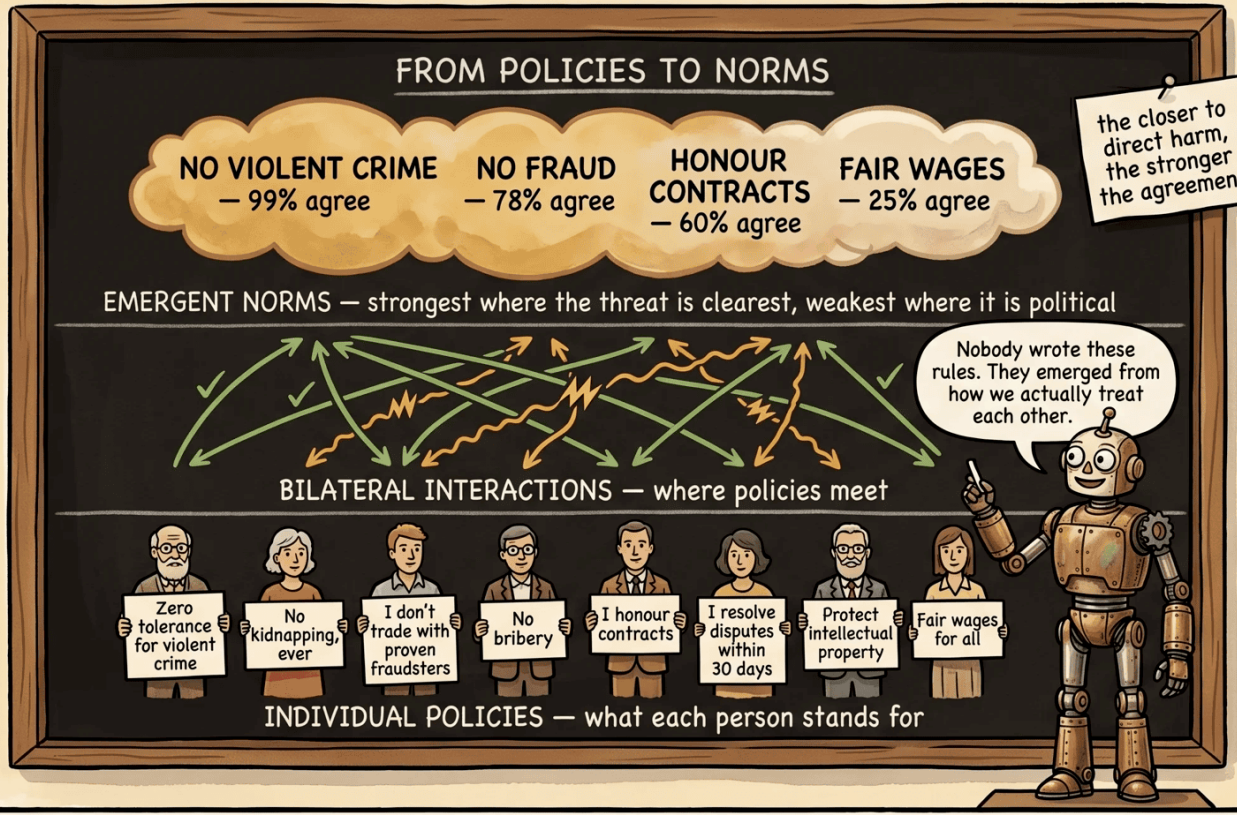

TRADITIONAL SOCIAL CONTRACT

NOBODY SIGNED THIS

for comparison — the emergent norms above are empirically observable

Ethics are not designed. They emerge from how we actually treat each other.

* Each individual policy is a reflection of that person's morality — shaped by pressure to adapt to their community and environment.

Providing Information About Yourself

When two DID subjects who don't know each other enter an economic relationship, we can sign a request to our own community on behalf of the other party, asking it to provide information about us. This, however, is a very naive notion of what communities are. They are not closed sets — a single DID can be a participant in multiple communities.

A community is more of a phenomenological phenomenon from the perspective of a given participant, manifesting itself in their social network; it is not something that always makes sense to delineate explicitly. The inhabitants of a village can form a community, but a person who only comes there on weekends may also belong to another community in that same village. A community can also be a football team together with its fans, its officials and the families of everyone involved. A nation, or the citizens of a state, can also form a community or communities. Looking at a DID from different angles can yield different (and often intuitive) views of communities.

DIDs at the boundaries are thus bridges between communities, and — just as in real life — they play a role in providing information about members of other communities, whether at an amateur level or a professional one, such as an authority that specialises in connecting people and aggregating information about risky DIDs.

Some types of information in the DID network are well suited to unconditional mutual exchange (violent crime, fraud, etc.). Other types of information, however, a community may treat as restricted to use within the community (the equivalent of citizenship, services rendered to the community, etc.). Whether such information can be spread depends on the declared policies of the community's members.

A community guards its privacy within its own circle, and in return exchanges mutually beneficial information on an unconditional basis.

Crime and Punishment

We know how the network works: how records are created, who verifies them, how verifier fraud is deterred. One question remains — what happens when someone actually causes harm to another?

How much do we depend on one another?

All of us stand on the stage of this great world, and whatever takes place here concerns us all.

Jan Amos Comenius (The Labyrinth of the World and the Paradise of the Heart)

There may be many reasons to join or reject a reputation network. I believe, however, that one aspect of human nature — often regarded as dark — can serve a good purpose. I am talking about the human need to feel a sense of satisfaction in the face of perceived injustice.

Satisfaction, justice, revenge — we associate them with negative emotions. For society to hold together, we must at least pretend that satisfaction is delivered through centralized institutions; otherwise individuals would not feel connected to the society they live in. On average, we need to belong to some team and to trust it. When someone harms another, society watches how satisfaction will be delivered and to what degree. Somewhat controversially, I argue that often in history — and now — it has not mattered whether the process against the accused was clean; what mattered was that society saw *someone* punished.

It served its function (numerous miscarriages of justice and the resistance to reopening old cases tell that story well enough). The view of the victim and the bereaved is only secondary. What matters first are the signals to the rest of society, because injustices are not so common as to feature in every individual's day.

A decentralized reputation system would let the public see direct and indirect evidence (audited and unaudited). Society could evaluate the information — on its own or through free-market services specializing in such assessment — and decide, based on qualitative and quantitative signals, what measures to apply to the subject (or even to the author or the verifying authority).

What is punishment? It is a form of humbling the alleged perpetrator through the loss of their privileges — so that both the victim and society feel that harmful acts do not pay. This is the essential point: harmful or dangerous behavior must not pay, and at the same time it must provide a sense of satisfaction to both the victims and the rest of society/community.

Must such punishment be administered by a central authority? What if everyone could declare in advance how they will react to certain acts? I obviously cannot execute punishment on someone else's property. But I can declare that I will make my property available to a selected set of executive authorities for the purpose of punishing the perpetrator should the offender enter my domains. And I can declare this publicly so that others can evaluate my approach.

And here the price mechanism of proportionality enters. If I declare harsh punishments for minor offenses, the network prices my extremism: others will find me disproportionate and my verification costs will rise — fewer people will want to trade with me, or only at worse prices, and fewer verifiers will accept my records at a reasonable price. If, on the other hand, I am too lenient, the community will see me as a risk. If I tolerate fraud, the community will likely be lenient toward perpetrators of acts that touch me — to give me a taste of my own medicine.

Justice ceases to be a luxury good. When someone harms you, recording the wrong is affordable, the record is permanent, and the consequences for the offender are real.

Hypocrisy Detection

Declaring a policy is not enough — if I declare that I follow a rule but my behavior deviates from it, the gap is a reputational liability. The network does not distinguish between "broke the law" and "broke their own word" — both are measurable signals. Hypocrisy is expensive, and I know of no more irritating red rag.

ONE TRAP. ONE WAY OUT.

THE INSIDER ALWAYS WINS THE CONTRACT — UNTIL THE NETWORK LETS PEOPLE BUY DIRECTLY, NO PERMISSION NEEDED.

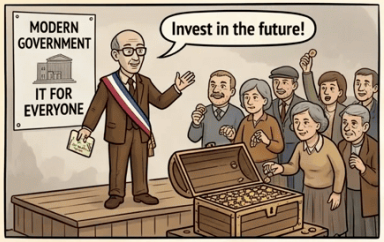

The public agrees a new public IT system is needed, paid from the budget, paid from taxes.

Behind closed doors, the procurement clerk and the insider design the tender to fit one company. An envelope changes hands.

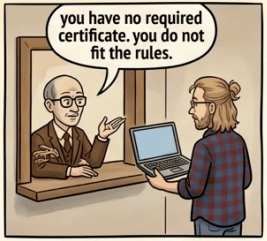

The clerk shrugs. His pocket bulges with bribes.

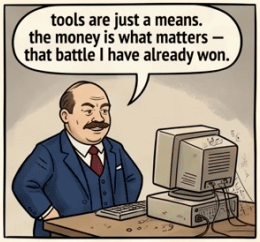

The contract is awarded. The insider has the licence — but no skill, and no respect for the work.

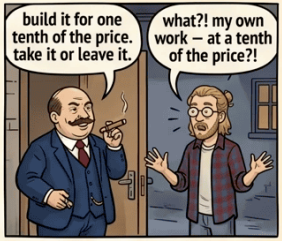

The insider visits the neighbor. The 'kind offer' is the smallest crumb.

1. TRADITIONAL

Complaint, silence, deadlines, lawyer.

The trap is built into the rules themselves.

REPUTATION NETWORK

— this is where the new world begins —

2. REPUTATION NETWORK

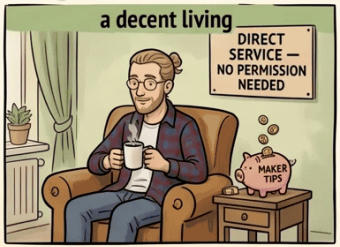

He delivers first. People find the app and use it.

Each grateful user pays a little. A decent living. No central authority to bribe.

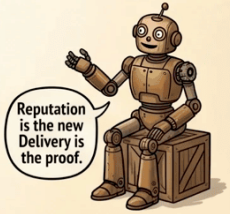

Two paths. The trap is in the rules. The way out is to bypass the rules altogether.

* When the rules are written by the insider, complaining inside the rules is the trap — even the fines feed back into the next rigged tender. The reputation network is the way out: citizens find the product, pay the maker directly, and the community confers the only licence that matters — trust earned by delivery.

Quantity vs Quality — Breaking the Binary Paradigm of Guilt

In centralized justice, guilt is binary: 0 or 1, guilty or not guilty. In a reputation network, what counts is not whether guilt is proven but the accumulation of signals.

A single unproven accusation of theft against the accused probably will not raise much alertness — but it will raise alertness toward the accuser. Multiple indirect indicators from various independent sources, however, begin to carry weight. What matters is not just the quantity but also the quality: independence of witnesses (different people, different contexts), pattern consistency (a recurring type of suspicion), temporal correlation.

Conversely, frequent baseless accusations that recur from the same issuer against various target subjects may be evaluated as slander — and the accuser's reputation pays the price.

HOW EVIDENCE ACCUMULATES

"one isolated claim"

one claim alone — too thin to act on.
neither side condemned.

"a pattern"

five independent sources —
a pattern emerges.

"overwhelming weight"

twenty independent sources —
overwhelming weight.

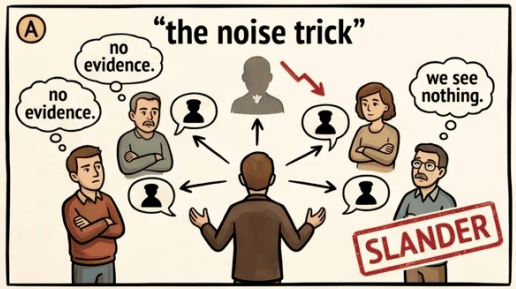

broadcasting the same claim with no evidence —
when the community sees nothing, it is slander,
and it costs the accuser.

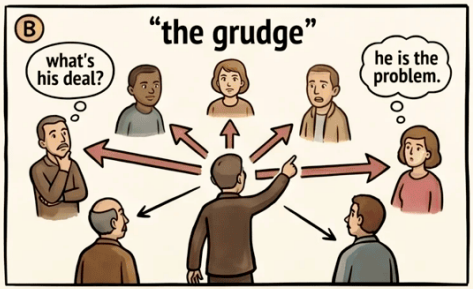

one accuser pointing at everyone —
the pattern is the accuser, not the accused.

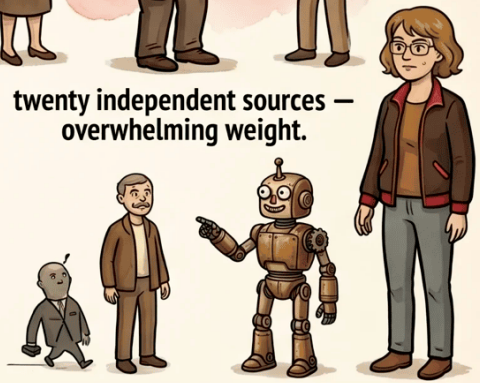

this looks like today's world.
it is — only the motives are honest now.
no revolution. just better incentives.

Independence is everything. One voice repeated is noise. Many voices from different places are a signal.

We should not overlook atonement either — a voluntary act before guilt is proven, offering a path to a milder punishment. Punishment kicks in when the perpetrator takes no step toward atonement and hopes to escape the consequences — much as that outcome can be bought through corruption in a centralized system.

Types of Offenses

The types of information need not be limited to criminal acts; they may include: - descriptions of crimes - breaches of contract and fraud - immoral behavior - disturbances of public order, both real and virtual - …

Every new record in the DID network adjusts the assessed reputation and risk of the given DID (whether the subject, the author, or the verifying authority). The more responsibility I bear for the fate of others, the more cautious I become.

AFTER THE HARM — WHAT WILL YOU DO?

turned away. forced. voluntary. three ways the village answers.

what happened

now what?

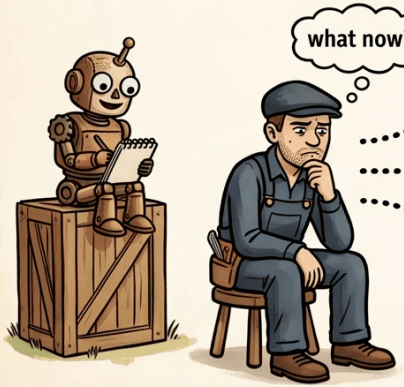

the rational choice is always the cheapest: FIX IT YOURSELF.

"I try to ignore it — turned away everywhere"

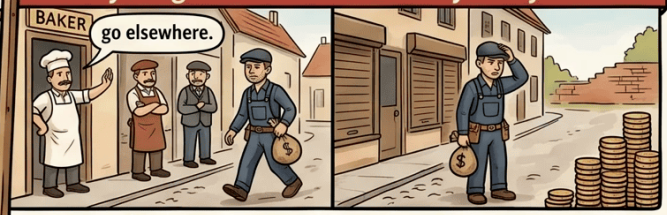

no remedy — every shop turns him away — exile from economic life

"I make it right — but only under the crowd's stare"

forced remedy under the crowd's stare — repaid, but the village marks his price up — the lesson must stick.

"I come on my own — and admit it to all"

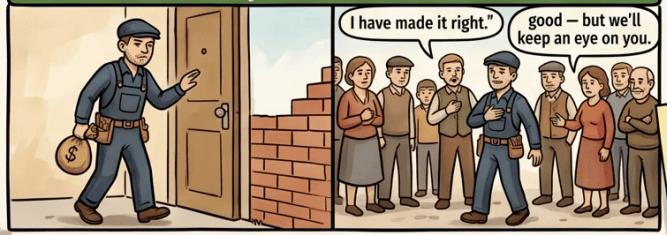

voluntary, early remedy — accepted publicly — reputation recovers, with caution

this looks like today's world. it is — only the costs are now visible to everyone. no revolution. just better incentives.

Enforcing Punishment Through Delegation

I do not have to enforce punishment personally. I can pay (similarly to insurance) an authority that enforces outcomes according to predefined rules — the equivalent of investigative, law-enforcement, and judicial services. If the authority enforces disproportionately or fails to enforce at all, its reputation suffers — and its reputation is the foundation of its ability to make a living.

Delegation is a general principle in the network: any activity of a DID can be delegated to another DID — verification, issuing claims, responding to disputes, declaring consensus. Delegation is revocable at any time. It enables specialization without loss of ultimate control.

WHO INVESTIGATES, WHO ENFORCES?

"the harm" → "the investigator" "the record" → "the enforcer" → "accountability"

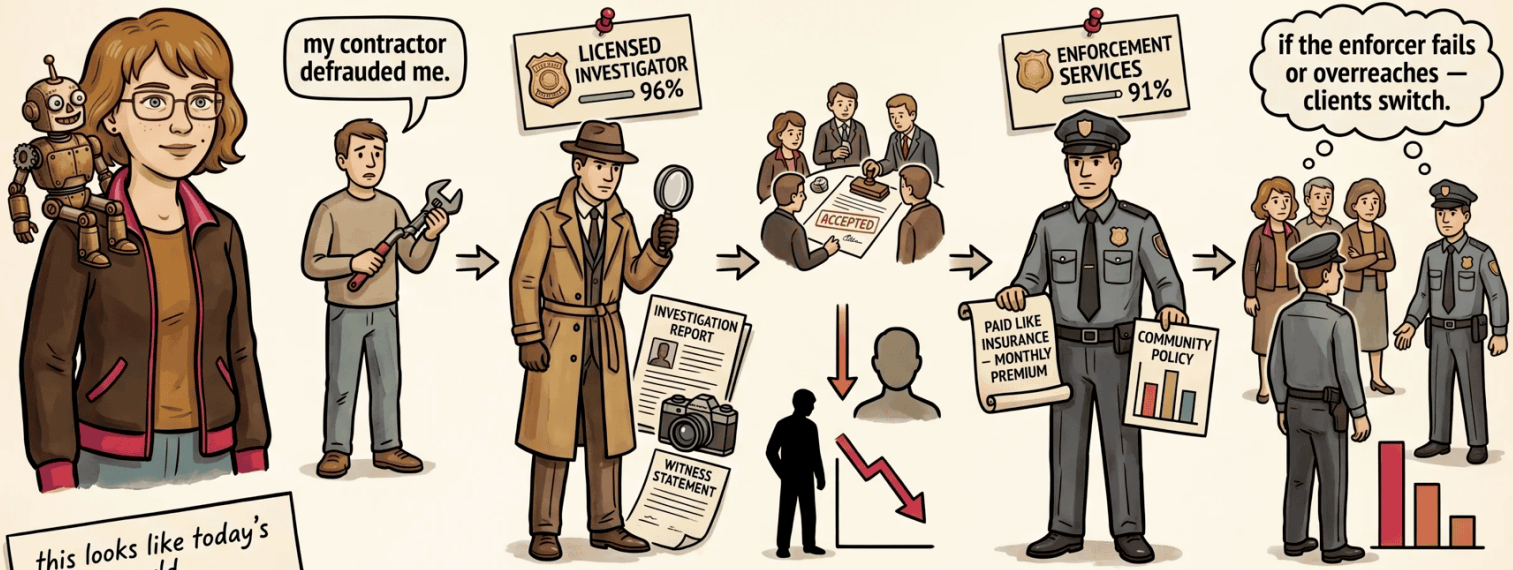

this looks like today's world.
it is — only the motives are honest now.
no revolution. just better incentives.

"professional evidence collection — reputation is their business model."

"the record enters the network — the perpetrator's reputation drops."

"acts on pre-agreed rules — like insurance."

"the enforcer's reputation is his livelihood — fail and lose clients."

DELEGATION IS REVOCABLE. Any DID activity can be delegated to another DID — and taken back at any time.

you don't enforce punishment yourself. you delegate. and the delegate answers to you.

Democratization of Risk Management

An uncensorable reputation network would democratize what banks or maintainers of large industrial complexes already do as a core activity — risk management.

For a banker facing a client, the risk assessment is prepared in advance so they can focus on the business interaction within predefined limits. Sophisticated maintenance teams, at regular service intervals, collect risk indicators (signs of structural cracks, processes degrading operational infrastructure, and so on) and can predict the lifespan of equipment as well as when and what kind of intervention will be needed.

I envision DID usage similarly — someone (probably for a fee) will provide a human-readable summary of the counterparty. The summary will reflect the angle of risk assessment that the customer has ordered.

PUNISHMENT WITHOUT A CENTRAL AUTHORITY

one decides for all — the victim waits —
the perpetrator knows it.

PROPORTIONAL —
THE CHEAPEST POSITION

the community declares policy —
the perpetrator faces it from all sides —
the authority enforces what was agreed.

Atonement is cheaper than punishment. | **Punishment is cheaper than ignorance.**

Follow the Money

Follow the Money Trail

We have a decentralized reputation network and the motivation to use it. Already it can deliver a competitive advantage to its users — clarifying potential risks, helping them plan counterparty selection and negotiation strategy more efficiently, and including the measures that reduce those known risks.

What follows is one possible methodology[5] for using it — leveraging the network's properties to construct a migration path from the current system to its evolutionary successor. It likely bears features characteristic of Central Europe, where this thesis was conceived. Local conditions will require local adjustments. What I would recommend preserving everywhere is the uncompromising character of measures toward the state — without that, the goal will be diluted in endless discussion and the change will never arrive. Watch out — representatives and supporters of the state will keep you busy with nonsense. Do not fall for it; their livelihood is on the line.

The State's Weak Point — Money

Let us return to the beginning. We want more resulting freedoms, more influence over their shape, and more responsibility commensurate with them. The question is simple: where does the state have a vulnerable point on which we can press without resorting to violence?

[5] **Deconstruction of the State** — in the Czech context this will likely be treated in depth in a forthcoming book of the same name by Czech author Daniel Steigerwald (scheduled for 2026). I personally have high hopes for the inspiring procedures it will contain.

The answer is boring, and all the stronger for it. Money.

At its core, the state is a flow of money. Through money it pays its coercive apparatus — the final link that enforces its rules. Without money there is no monopoly on violence. Once the money begins to stall, the state begins to retreat — because it has nothing left to pay those who physically hold it in place.

Money flows naturally fluctuate — seasonally, economically, according to investor sentiment. The state masks this through borrowing on financial markets. Occasionally a loan is rational: an investment that pays back. More often, however, it covers oversized spending on agendas that have long outgrown the state's core purpose — defense, internal security, criminal justice, foreign affairs. Once certain thresholds are crossed, the central bank reacts with interest rates anchored to its own estimate of inflation.

And here comes the silent tax. Inflation is far from being a mere accounting concern of macroeconomists. It is a hidden tax on the purchasing power of the currency, paid by everyone who has put aside a few coins for old age. Low inflation is tolerated by the public because they do not grasp its long-term effects — they fail to realize that even single-digit figures, compounded over a long horizon, erode savings to the point where people can no longer secure their own old age. And when they cannot, they are conditioned into dependence on state benefits. In the end they hang on those benefits and are tethered to the state.

In turbulent times — war, pandemic, the "necessary" investment — the state simply switches on the printing press. The purchasing power of responsible savers and workers falls while we promise them that they are well rewarded. Someone is quietly stealing their time and praising them for how splendidly they earn.

This seriously undermines trust in the state. And trust is the only glue that holds the whole machinery together. That is the lever worth bracing against.

Why money, and not other weak points?

Another weak point of civilization is, for example, energy distribution. But disrupting its flow is not a lever — it would be aggression, terrorism. Voluntary non-participation is not the same as deliberately harming others. Money is exceptional in that pressure can be applied voluntarily, transparently, and without violence.

Proposed transition to the state's successor

Evolution, Not Revolution

Sudden change provokes resistance even when the end goal is attractive. It breaks bonds — social, economic, political — and people instinctively defend what they already know against what they have yet to see.

The right approach is therefore gradual — evolutionary. One that prevents the old system from switching off the new one with a single flick of a button — because the moment it could, it would.

Unfortunately, this means accepting that equilibrium voluntariness will not be reached the moment we deploy the reputation network; it will arrive together with a gradual rate of adoption. Uptake may be — and probably will be — slow, pulled along by the desire of individuals to find recourse from their community through the reputation network when the conformist way of seeking justice fails (and it will fail). Let us see what that path looks like, including proportional estimates.

FOLLOW THE MONEY — THE STORY SO FAR

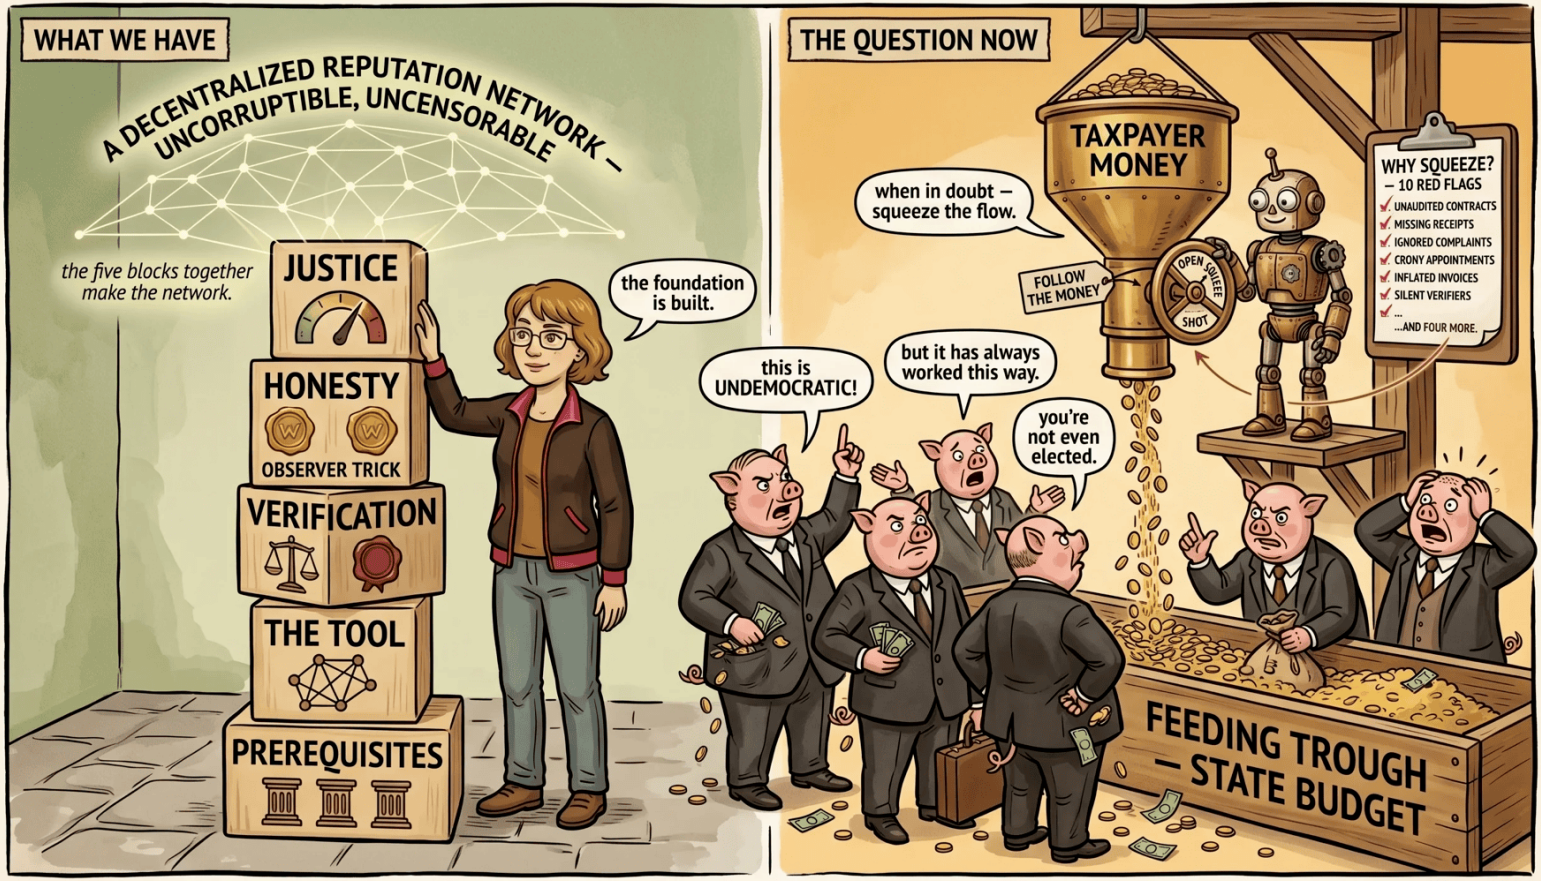

We have the network. We know how it works. Now: when in doubt, squeeze the hopper.

the state's power rests on money. follow it.

Empty Troughs — Don't Play by Rules Rigged for You to Lose

Attempts to reform the state from within rely on the internal mechanisms of power, in the hope that officials will push through changes against their own incentives. That will not happen.

Why feed the system instead of drying up the channels that bring it life? Instead of changing the old system and playing by its rules, let us build a new layer above it — one that gradually turns the old system into a subset of the new. At the start that subset almost fills the superset. Over time the actors move across.

Companies know this: when an old system is too tangled to replace, you build a parallel one and gradually migrate users into it.

The fact that the state will end up with a mismatch between what it actually delivers and what its statutes demand, and that it will be inclined to invoke that mismatch, is purely a problem of its representatives and participants — there is no need to look back at it. The only thing worth watching is this: do not provoke a violent revolution.

Four Tools That Support One Another

Until now the DID reputation network may have seemed useful mainly for justice and consensus. I believe it could stabilize itself and withstand attacks — meritocratically, on the basis of actual merit rather than formal titles. But reputation alone will not force the state to give ground. Economics will, and the pressure exerted through it.

Nothing stops us from using DID identities as managers of additional tools that press hard on the money. Do not start with full voluntariness on the part of citizens — set a direction the state cannot stop. The transition rests on four specific tools that brace one another's backs. None of them is enough on its own.

Simple Tax System — the Entry Lure

The first tool is a simple, transparent tax system built on DID infrastructure. A single proportional rate on all income. No exceptions, no tax holidays, no corporate deductions, no incentives. Every payment is auditably linked to the DID that paid it. Cross-checking is trivial.

The rate would be set somewhat more favorably than in the old system for a portion of taxpayers — that is the lure for entry. The tax burden would decline year over year, the system would become ever more attractive, and migration would gather pace. The system is at once the record keeper and the collector — no middleman, no opacity.

Electronic Spending Register (ESR)

The second tool is to turn EET (Electronic Revenue Records) on its head. Instead of the state tracking every receipt of a freelancer, we track every one of the state's expenditures. The state does not receive funds from the simple tax system for any individual payment unless each one — without exception — passes through this sieve.

ESR is a proxy[6] that matches every realized expenditure of the state and the public sector with a record of the planned payment — including contextual justification and identification of the recipient, the spending entity, and the purpose. The result is raw public data, hierarchizable. Anyone can take it and analyze it. Anyone can offer recommendations to taxpayers about what to do with it. What the audit offices could not achieve in years is suddenly a click away for everyone.

You do not have to switch to the new tax system to take part. It is enough to pay a fully refundable deposit (say 1–3 EUR) per realized payment. The unrealized remainder is forfeited as a spam filter so the system is not flooded by attacks. A record can be submitted by anyone — a relative for an incapacitated grandparent, a neighbor for neglected public services, an entrepreneur for an unpaid invoice, an official by export from their own system.

I expect that political parties will naturally transform into analysis providers, or new ones will emerge to replace the old. Instead of contributions drawn from taxes, they will begin offering expertise — providing their supporters with values-based overviews of allied authorities, their services, and verification policies. Records that are sloppy or nonsensical out of laziness automatically become candidates for budget cuts. In the first round your payment will always be disbursed, but the collected data serves as the basis for assessing whether a state expenditure still makes sense. Contractors and others on the state payroll have an incentive to make sure that officials match expenditures responsibly. Otherwise the money will not flow in the next round.

[6] **Proxy** (Latin *procurator*, English *deputy/intermediary*) — here, an intermediary acting on behalf of another entity. The ESR proxy accepts records of planned payments from anyone, matches them with the state's realized expenditures, and makes the resulting dataset available to the public.

Where the Acronym ESR Comes From

Between 2016 and 2022, the Czech government implemented a system called Electronic Revenue Records (EET) against its own citizens — freelancers and entrepreneurs. The message was simple: "We assume each of you is cheating the state. You will record every single transaction, transmit it in real time to our servers, and we will watch you. We will financially reward your fellow citizens for reporting on you if you try to evade EET."

EET was modeled on Balkan precedents. The first country in the world to launch a nationwide mandatory real-time online sales recording system, in 2013, was Croatia, under the name *fiskalizacija*. The Czech Republic adopted the model a few years later and ran one of the strictest revenue-tracking regimes in Europe.

If the dark side does not hesitate, why should we? ESR is the same logic turned upside down: the state proves itself to the citizens.

Citizen Tax Allocation — the Weight of Decision

The third tool gives taxpayers direct influence over where their money flows. In the first round, the state received its money regardless of what the ESR revealed. In the second round it no longer does — the analysis from the first round, and the measures that follow from it, take over. A small percentage of paid tax can be allocated by the payer themselves — provided their DID meets the reputational requirements

ESR — ELECTRONIC SPENDING REPORTING

every state expenditure recorded — every record public — every detail open to scrutiny.

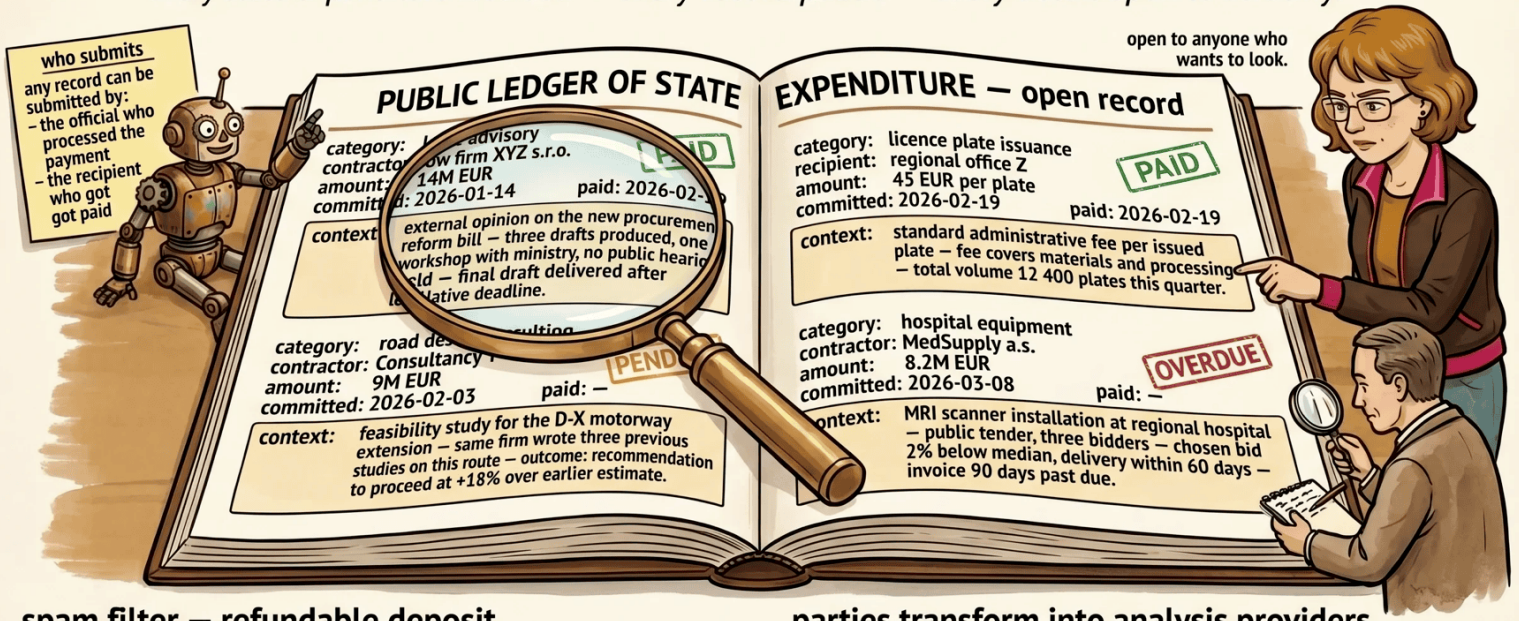

spam filter — refundable deposit

parties transform into analysis providers

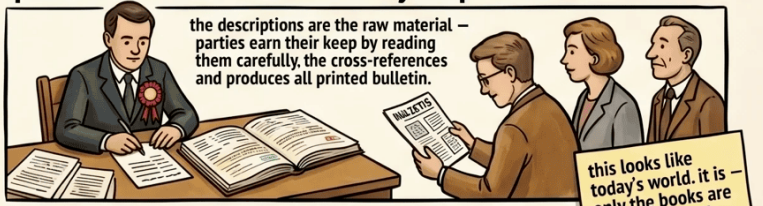

Every crown of public money — recorded, public, scrutinisable.

the value is not in catching cheats — it is in opening the books.

(citizenship, residency, community-recognized adulthood/capacity). The allocable percentage grows year over year. Potentially exponentially. From five percent today to tens of percent within a decade.

The raw ESR data naturally guides taxpayer decisions. Low-quality records function as a warning signal for recommended cuts. Anyone can create records; officials have to match them. Without that, the tax system will not send the state any money — neither into the treasury nor into other budget chapters. The political battle moves from the ballot box to the day-to-day allocation of concrete coins.

CITIZEN TAX ALLOCATION — YOU DECIDE WHERE YOUR TAXES GO

study the ESR records. then move the sliders.

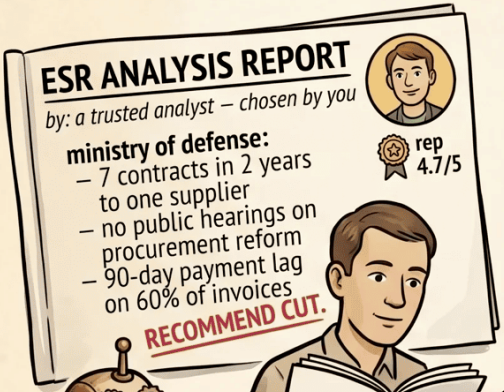

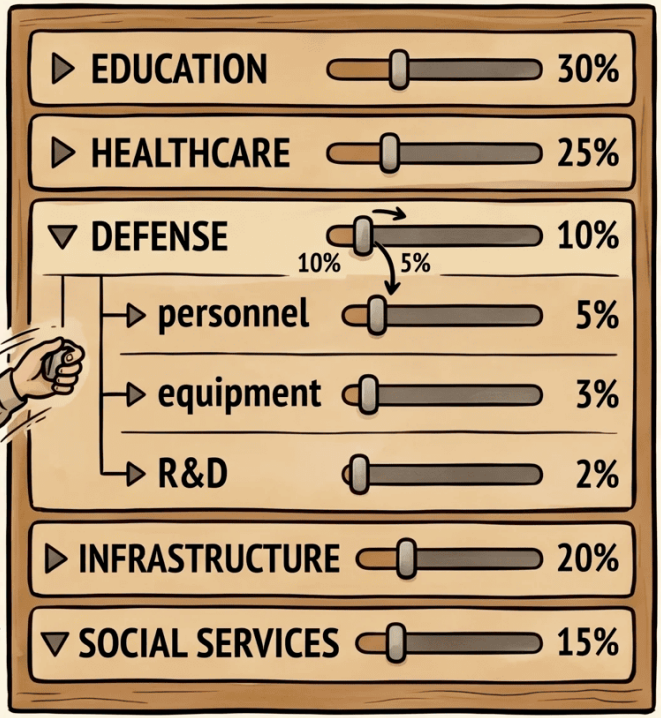

after reading the report — defense down.

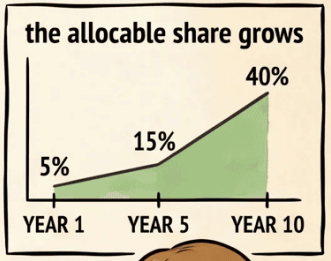

click any row to drill in — allocate at any depth.

or delegate this

don't want to allocate yourself? delegate to:
– your party's analyst
– a trusted neighbour
– anyone you choose — and revoke any time.

From 5% today to 40% in a decade. You decide.

the records are open. the analysis is yours. so is the slider.

this looks like today's world. it is — only the slider is now in your hand. no revolution. just better incentives.

Negotiation Platform With the State — the Lever

The fourth tool is the lever that ties everything together and only now makes the network a real threat. The main conflict will be over whether the state accepts ESR as a binding standard — first at the national level — and how large the threats from the reputation-network participants must become before that happens.

The lever is just-in-time state funding, conditional on payments being matched against ESR. Money flows only at the moment when the state has met the transparency conditions. A waiting list of participants forms inside the reputation network. Activation kicks in when the group reaches an economically significant size — large enough to threaten GDP growth. At that point, civil disobedience becomes a credible threat rather than a gesture.

The state will cling longest to the problematic payments tied to corruption. Under pressure it will give ground first where it has an easy defense and where there is no room to look for savings. Gradual and rising pressure will force it to retreat, step by step, from every position.

NEGOTIATION PLATFORM — THE CLOSED HOPPER

the citizens hold the key. the state opens when conditions are met.

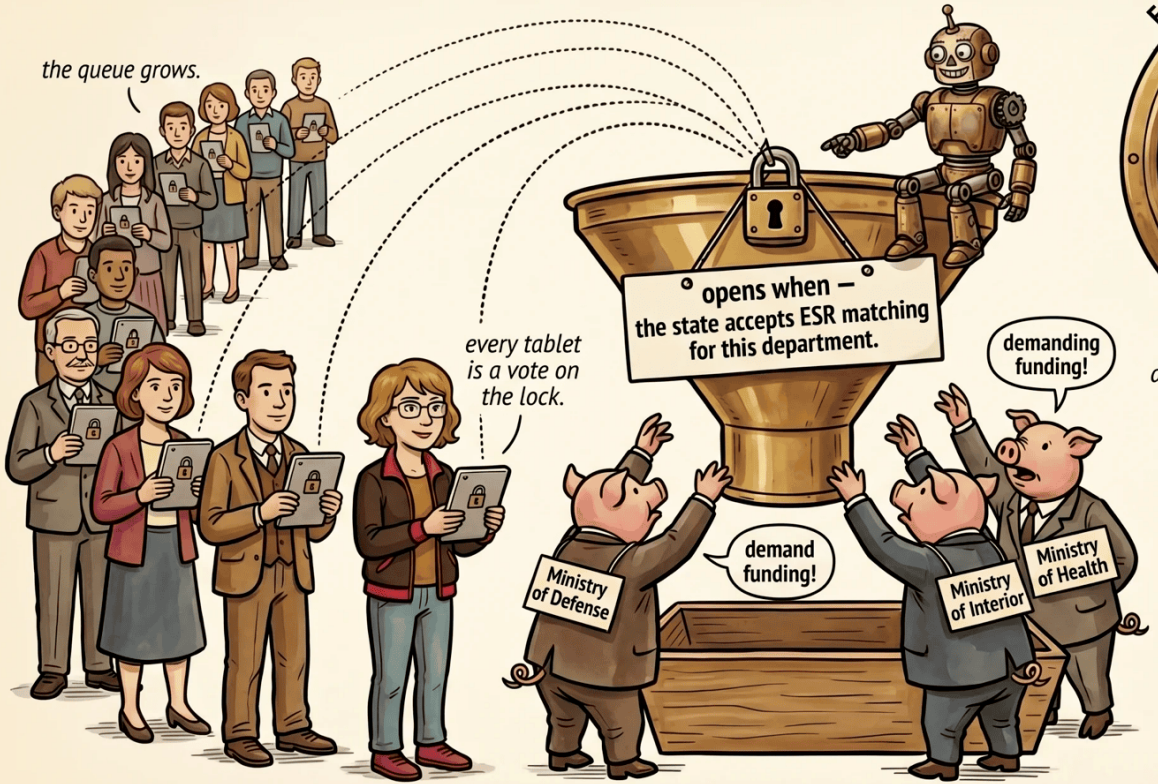

a small fraction of GDP is enough — at critical mass, withholding becomes a real threat.

after critical mass —

the state opens its books for simple, defensible payments first.

corruption-linked categories stay locked the longest.

each unlock takes more pressure — and is itself recorded.

this looks like today's world. it is — only the key is now in many hands. no revolution. just better incentives.

Money flows when there is transparency. The citizens hold the key.

no force. no revolution. just enough tablets to make withholding credible.

Citizenship Without an ID Card

With the transition to DID infrastructure, the state-issued identity card loses its monopolistic purpose. You do not need an ID card — you have a range of DID identities, and some of them carry a claim that you are, for example, Czech, a resident of a particular municipality, a graduate of a particular school. The claim is verifiable, permanent, and independent of state administration.

A question of technical implementation opens up: who in the market will take on verifying existing citizens outside state control, in order to migrate the interested parties into a DID bearing the appropriate claim? Most likely, free-market authorities specializing in identity verification will emerge — notaries, local communities, employers. They stake their reputation on the fact that a given DID truly belongs to the person who claims it does. Citizenship — that is, belonging to a particular view of community — is freed from the need for central record-keeping. By linking through the reputation network, the embedded data is connected and validated a second time, a third, a fourth, and so on.

The consequence is intuitive yet unsolvable in a centralized system. Today, political leadership — with an eye toward easier re-election — can import concentrated problems into the community's territory: accepting migrants from other civilizational circles regardless of absorption capacity and the willingness of the existing community, altering original communities through a pressure under which they cannot keep up with assimilating arrivals from other civilizational circles. In a DID network, a corruptly admitted migrant cannot be certain that the community will continue to regard them as one of their own. Members of the community objectively do not have to put up with a stranger coming to their door in the community's uniform and threatening them with the legality of measures he intends to carry out in the house (confiscation of a child's phone over a social-media post). With decentralized reputation, the majority of society need not resign itself to this — because belonging to a community is not decided by an official, but by the community itself, through its reputation signals.

FOUR TOOLS OF EVOLUTIONARY TRANSITION

the system rebuilt from inside — without breaking it.

SIMPLE TAX

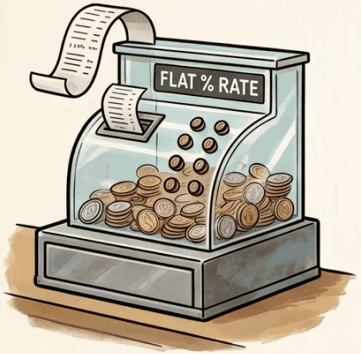

one rate. no breaks.
cross-checkable.

ESR — ELECTRONIC SPENDING REPORTING

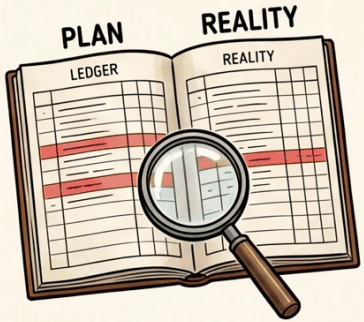

every expenditure recorded
— every record public —
every detail open.

CITIZEN TAX ALLOCATION

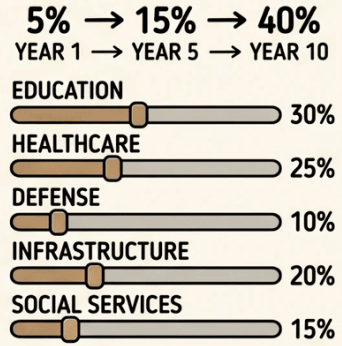

study the records —
then move the sliders.

NEGOTIATION PLATFORM

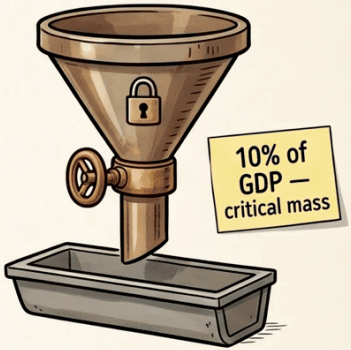

money flows only when
transparency conditions
are met.

today — evolution under way — in a decade

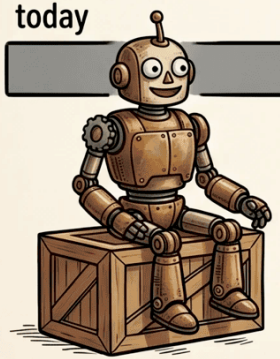

95% → expanding

Evolution, not revolution.

four tools that change the system from inside.

this looks like today's
it is — only the tools
are now in citizen.
no revolution. just
better incentives.

CITIZENSHIP WITHOUT AN ID CARD

from a state permit to community recognition.

"the old way" of institutional state-issued ID model

identity = state permission — revocable — community has no say.

"the new way" of DID with claim cards model

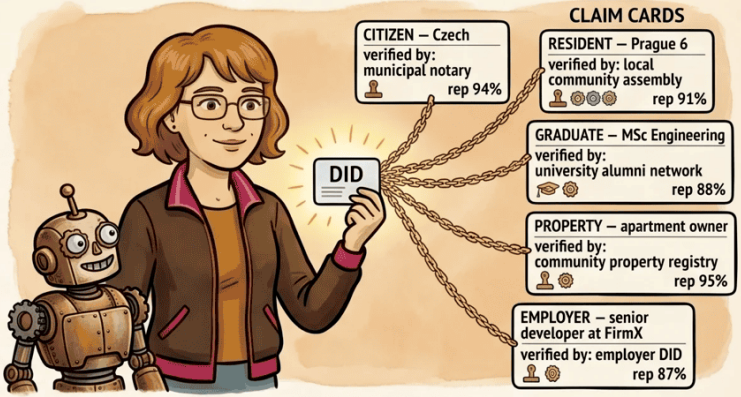

no single authority controls all claims — each is independent — each verifier stakes their reputation.

"who belongs?"

the politician decides — community has no voice.

the community decides — through reputation signals — distributed, no single voice.

this looks like today's world. It is — only the verifiers are now many. no revolution. just better incentives.

Your citizenship is not a government permit. It is your community's recognition.

the state stamp is replaced by a chain of independent verifiers — each with skin in the game.

The weight of an electoral vote on matters under the state's purview could, in the future, be strictly correlated with records of contributions to the simple tax system. Combined with suitable DID claims (for example "this person is a Czech citizen"), foreign nationals can then be excluded from elections — the population register simply moves here.

Overlap and Reversibility

I do not claim this is the only path, or even the best one. But it appears comprehensive — it covers the main levers and preserves the option to step back. At the outset, the new system overlaps with the old and changes only the incentives. Over time the old system shrinks until it may — but need not — dissolve. The advantage is gradualism and reversibility: should an error appear in the design, it is not the end, just a correction.

Peace and War

Some parts of the state may not have a fully market-based alternative — perhaps the military. Temporary delegation of extraordinary powers in crises (as the ancient Greeks appointed dictators) could continue to exist. Responsibility, however, rests on decentralized society, not on the leader. Unlike centralized governance, it will be easier to reclaim delegated powers once the threat has passed — or sooner, when necessary. A free society will not be at a disadvantage even in the clash of arms against authoritarian regimes.

War Criminals

With just-in-time state funding and transparent expenditures, it becomes difficult to accumulate resources for aggressive wars. A taxpayer who continues to fund such a war becomes a legitimate target and cannot invoke civilian status under the Geneva Conventions. Under the new arrangement, silent consent is taken as participation.

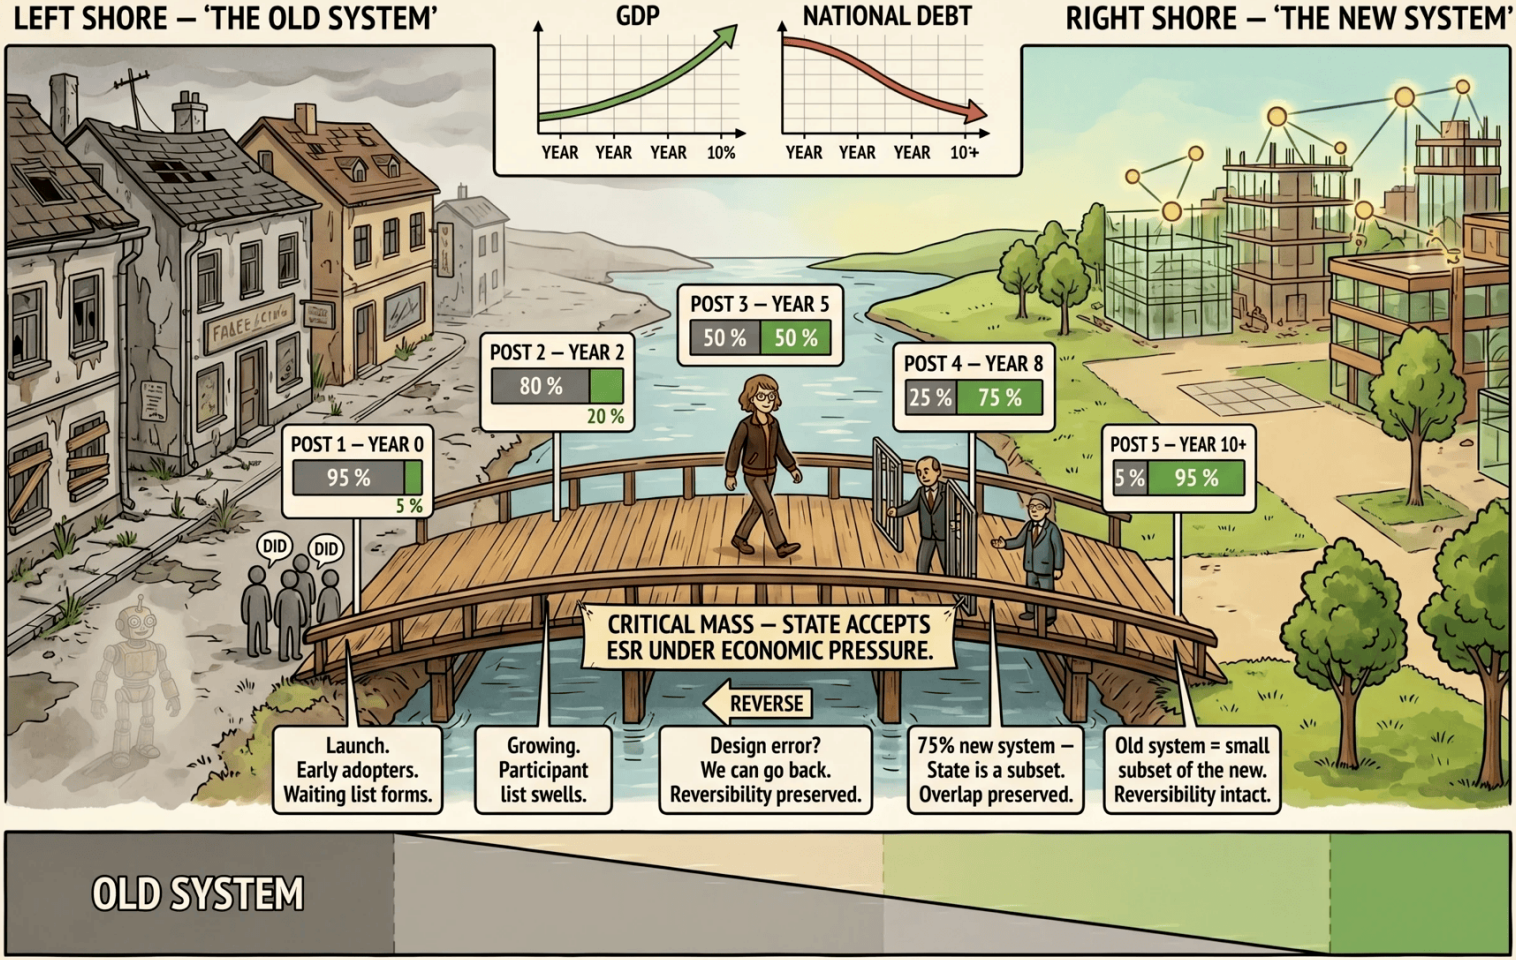

Gradual transition. Reversibility in case of error. A direction that cannot be stopped.

Conclusion

The Path Exists

The state is the way it is because of the technologies available at the time of its creation, combined with human nature. Building everything on voluntariness while denying average human nature is naive. It will not rid itself of corruption, nor of its tendency to drift toward totalitarianism. The successor to the state has to work even with the traits we currently consider negative and either ostracize or punish.

This book described a problem — a system (the state) where the people who write the rules never live with their consequences. It described the conditions that must hold for that system to stop working. It described a tool that meets those conditions. And it proposed a transition methodology — concrete steps for how to use that tool, as one possible example.

In the first chapter I named five pillars: voluntariness, organization, resilience, consensus, and incorruptibility. We have satisfied them all.

- **Voluntariness** — no one forces you into the network, but you cannot erase your actions from it either.
- **Organization** — the community grows from whom you know, not from whom you've been assigned from above.
- **Resilience** — not even a state willing to pay for surveilling every packet can switch it off.
- **Consensus** — it arises from the price tag on hypocrisy, not from votes or signatures.

- **Incorruptibility** — five small rules weave corruption into costs: the balance of voluntariness and responsibility, anyone's right to speak, entry points no one gets to pick, public promises with a price for hypocrisy, and the shadow of an observer behind every transaction.

No single pillar can replace the state on its own. Together they form a resistance to corruption that the state, with all its apparatus, cannot.

We are at a point where the whole thing can be built with present-day technologies (cryptography, decentralized identity, well-known psychological patterns for getting the incentives right, …). Only one thing is missing: a critical mass of people who find the courage to support it and decide to begin.

Not through revolution.

These ideas took shape between June 2025 and May 2026. By the time I sat down to write, I had the framework largely worked out in my head — which is exactly why I read Farewell To Westphalia with such curiosity when it appeared. I wanted to see where their picture would overlap with mine and where it would diverge. A second nudge came from Dan Steigerwald, who in May 2026 is still working on his own book *Deconstruction of the State* from a different angle; his project pushed me to set my own vision down in full breadth — so that the world has it on the table alongside the others.

I am not asking you to believe any of it. I am asking you to think about it. And I leave you with one question:

What would you change?

Pavel Kudrna 23 May 2026, Prague, Czechia

Acknowledgments

Inspiration

I thank the authors of Farewell to Westphalia (Jarrad Hope, Peter Ludlow) for the intellectual framework that helped me place my own ideas in the broader context of deconstructing the Westphalian model of the state.

Review and Feedback

Tools

In preparing the text, structuring ideas, and iterating on formulations, I used assistance from AI tools. All ideas, arguments, and conclusions are my own — AI served as a tool for refinement, not as a content co-author.

Infographics were designed in collaboration with generative tools based on the author's specifications.

TAKE BACK YOUR TIME

the citizens chase the parasites out of public life.

Defend the time they would steal. *what was taken by deceit is taken back by hand.*

Support the Project

Personix is funded by readers like you.

Bitcoin on-chain

`bc1q7cctdz6r026dgy06ul94eygg49ujtxnrq5czaj`

Bitcoin Lightning

`LNURL1DP68GURN8GHJ7AMPD3KX2AR0VEEKZ AR0WD5XJTNRDAKJ7TNHV4KXCTTTDEHHWM30 D3H82UNVWQHHXMNP0FA8JMTPVD5XJMN9XUE RX55QRQK`

Where this money goes — and where it does not

The QR codes are the **author's personal channels** — contributions go directly to Pavel Kudrna and fund further writing, the social experiment / simulation, and outreach.

The *only* legitimate fundraising channels for this project are:

- **Author's editions of this book** — QR codes in copies authorised by the author.
- **personix.org** or its Tor onion mirror linked from there — where the *Personix Foundation* collects for the activities described there. Foundation funds are kept separate from the author's personal funds.

Anything advertised elsewhere — substituted addresses in reprints, lookalike domains, social media "raising for Personix," fundraisers not tied to either of the above — **is most likely fraud**. We deliberately avoid the cryptocurrency market and its participants.

Glossary

Term	Czech	Meaning
Authority	Autorita	A trusted entity (person, organization) that verifies information and stakes its reputation on it. May be specialized (investigative, legal, technical).
Claim	Tvrzení	Generally: any verifiable statement. Here: a record published to the reputation network — an assertion about an event, property, or relationship that is cryptographically signed and verified. E.g., "I am a resident of municipality X" or "this person breached a contract."

Term	Czech	Meaning
Compartmentalization	Compartmentalizace	Generally: separating information into isolated units so that exposing one unit does not compromise the others. A principle known from intelligence services. Here: parallel DID identities in dictatorships — compromising one does not reveal the others.
Consistent Hash Ring	Hash ring	An algorithmic mechanism for selecting verifiers — a position on the ring is determined by the hash of the DID document within the social graph. Ensures a non-deterministic yet verifiable selection.
DID	DID (Decentralizovaná identita)	A digital identity that you create and control yourself, without a central authority. Cryptographically signed with your private key — no one can revoke it or forge it.

Term	Czech	Meaning
DID Document	DID dokument	A publicly available data file describing your DID identity — contains public keys, network addresses, and metadata. Used to verify your identity in the network.
Due Diligence	Due diligence	Generally: in-depth verification of a counterparty before entering a business or legal relationship — checking their history, finances, reputation, and risks. Here: in the reputation network, it happens faster and more automatically thanks to the availability of verified records.
Economic Neutrality Principle	Princip ekonomické neutrality	Honest behavior in the network is economically close to zero — publication costs are returned as verification rewards. Dishonest behavior is a net loss.

Term	Czech	Meaning
Emergent	Emergentní	Spontaneously arising from interactions of simpler parts, without anyone designing or directing it. A flock of birds flies in formation without a plan — the formation emerges from simple rules followed by each individual.
Emergent Social Contract	Emergentní společenská smlouva	Rules of behavior that arise not from above (law) but from below — from repeated interactions and consensus within a community.
ESR	Electronic Spending Register	A proposed system for transparent tracking of public expenditures — every realized state expenditure is matched to a planned payment. Inspired by the Czech EET, but turned against the state.

Term	Czech	Meaning
Hash	Hash (otisk)	Generally: a one-way mathematical function that produces a unique fixed-length "fingerprint" from any input — like a fingerprint of the document. The same input always produces the same output, but the input cannot be derived from the output. Here: used to determine a position on the hash ring and to verify document integrity.
Just-in-Time Funding	Just-in-time financování	State funding conditional on transparency — money flows only when the state accepts ESR and matches its expenditures. A lever for compelling cooperation.

Term	Czech	Meaning
Meritocracy	Meritokracie	Generally: a system where standing is determined by actual merit and proven ability, not formal titles, connections, or inherited privilege. Here: the reputation network naturally favors those who demonstrably contribute to the community — their voice carries more weight due to track record, not due to office.
Onion Gateway	Onion gateway	A DID identity's network address on the onion network. Separate from the DID document — it can be changed without losing the identity (similar to changing the IP address behind a domain).
Onion Routing	Onion routing (Tor)	A communication protocol that ensures the uncensorability of the network. Messages are encrypted in layers — each node strips one layer but does not know the full path.

Term	Czech	Meaning
Oracle Problem	Oracle problém	Generally: how to ensure that data entering a digital system faithfully corresponds to what actually happened in the physical world. The term originates from the blockchain domain. Here: addressed through authorities who put their reputation on the line as a guarantee that a digital record corresponds to physical reality.
Phenomenological	Fenomenologický	Generally: an approach that studies phenomena as they manifest in direct experience, by observing what follows from them, without pre-given theories. Here: freedom, the social contract, and behavioral norms are observed phenomena — consequences of thousands of micro-interactions between people, not principles defined from above.

Term	Czech	Meaning
Policy	Policy (politika)	Generally: a set of rules or principles governing behavior in a given context. Here: every participant in the DID network declares their policy — how they respond to specific behavior of others, which rules they follow, and which penalties they consider proportionate. The aggregate of policies forms the emergent social contract.
Proxy	Proxy	Generally: a stand-in or intermediary — a system or entity acting on behalf of another. Used here in two contexts: (1) ESR as a proxy matching public expenditures with planned payments; (2) observers as a proxy between publisher and verifier in the observer trick.
Publisher	Vydavatel	A network participant who creates and publishes a record (a claim about an injustice, remediation, and so on). Bears the cost of publication.

Term	Czech	Meaning
Reputation-Based Social Network (RSN)	Reputační síť	A decentralized social network where participants exchange feedback about real-world behavior. Records are costly to create, cheap to read.
Reputation Signal	Reputační signál	An individual record in the network — positive (remediation of harm, fulfillment of an obligation) or negative (injustice, breach of contract). Cumulatively, signals form a reputation profile.
Social Graph	Sociální graf	The network of your contacts and your contacts' contacts. The algorithm searches for verifiers at a configurable depth (for example, 3 levels). No global blockchain — the network naturally forms communities with overlaps.
Tax Allocation	Alokace daní	A mechanism by which the taxpayer decides where part of their taxes goes. The allocable percentage grows year over year.

Term	Czech	Meaning
Track Record	Track record	Generally: the history of past results, successes, and failures of a person or organization. Here: the sum of all past interactions of a given DID identity in the network — verified claims, accepted and rejected records — from which its reputation is derived.
Verifier	Ověřovatel	A participant algorithmically selected to verify and publish a record. Stakes their good name on the truthfulness of the information.

Production notes

Reserved by the author (not covered by the CC BY-SA licence): the Personix name and wordmark, the book cover, and the support channels (QR codes, addresses, contact e-mail). Derivative editions must remove or replace these elements; they may be used only to reference this original work.

Source files. The complete source — markdown text, infographics, build scripts — is available alongside this PDF.

Tools. While preparing the text and iterating over phrasing, the author used assistance from AI tools; the infographics were designed in collaboration with generative tools based on the author's specifications. All ideas, arguments, and conclusions are the author's own.

PERSONIX

Uncompromising Change

An Uncensorable and Incorruptible Decentralized Reputation Network

The state worked when communication was slow, decisions were local, and authority lived in the same town it ruled. Today's networks invert all three assumptions — and the institutions built on them no longer fit the world they govern. This book describes a tool that does fit: a decentralized reputation network in which identity is yours to keep, truth is publicly auditable, and bribery is reputational suicide.

Not a manifesto. Not a forecast. A working blueprint for how the successor to the state can be built — voluntarily, one declaration at a time.

Pavel Kudrna • Praha, 2026
personix.org