# Decentralized Reputation Network: A Framework for Natural Societal Organization

Pavel Kudrna
Prague, Czechia
Draft v1 — March 2026

## Abstract

The sense of justice—the conviction that wrongs are made right and merit is rewarded—is the strongest cohesive force of society. When centralized institutions of governance cannot deliver this sense because the cost of enforcing a right exceeds the value of the right itself, social cohesion erodes. Current institutional structures fail to resolve a significant class of disputes in systematically the same way that central planning fails to allocate resources: through information asymmetry, missing feedback loops, and disconnection of decision-makers from the consequences of their decisions. This paper presents a Reputation Social Network (RSN): a decentralized, censorship-resistant communication layer built on Decentralized Identifiers (DIDs) that enables pseudonymous participants to publish, verify, and consume reputation information about real-world behavior. The core mechanism rests on three design principles: (1) an asymmetric cost structure where reading reputation data is inexpensive but writing requires verifiable effort; (2) a non-deterministic verifier selection protocol based on consistent hashing that prices radical claims through escalating iteration costs; and (3) a market-driven reputation authority model where private verifiers stake their accumulated reputation on the accuracy of claims they endorse. The resulting architecture produces emergent meritocracy without central coordination and enables a gradual, reversible migration path from existing state-administered systems.

## 1 Introduction

### 1.1 The Problem of Centralized Trust

Modern governance rests on a fundamental assumption: that centralized institutions can mediate trust between strangers at scale. Courts adjudicate disputes. Registries certify ownership. Regulatory bodies enforce standards. These institutions were designed in an era when information was scarce, communication was slow, and delegation of authority to a central body was the only feasible coordination mechanism. Centralized governance, however, fails for the same structural reasons as central planning: information asymmetry, missing feedback loops, and disconnection of decision-makers from the consequences of their decisions.

The assumption fails, for example, when the cost of institutional mediation exceeds the value of the dispute. Consider a tenant who discovers that their rented apartment lacks a legally required electrical safety certificate—a condition that poses immediate danger to life. The tenant's legal right to withdraw from the contract is unambiguous. Yet the cost of enforcing that right through the court system exceeds the monetary value of the deposit and damages owed, even including potential statutory compensation [1]. The rational response is to absorb the loss and exit.

This is not a pathological edge case. It is the default experience for a significant class of disputes in which the harm is real but the monetary stakes fall below the threshold at which institutional enforcement becomes economically rational. The result is a system that structurally favors bad actors: those who harm others in volumes below this threshold can act with impunity, because the cost of seeking justice falls on the harmed party.

### 1.2 Why Existing Solutions Fall Short

**Decentralized Identity (DID) frameworks** [2] provide cryptographic mechanisms for self-sovereign identity management but focus primarily on identity verification without addressing the broader question of behavioral reputation.

**Soulbound Tokens (SBTs)** [3] capture the insight that reputation is non-transferable but rely on blockchain infrastructure with scalability constraints and do not address the economic incentives governing information entry. Related concepts such as merit

tokens (e.g., Liberland) share a similar philosophy but remain bound to specific jurisdictional frameworks.

**Decentralized Autonomous Organizations (DAOs)** [4] implement governance through smart contracts but replicate plutocratic dynamics where influence is proportional to capital. Quadratic voting [5] partially mitigates this but limits practical adoption. DAOs also face the fundamental *oracle problem*: smart contracts cannot reliably verify real-world states without a trusted external source, and this problem has no general solution within blockchain architecture.

No existing system combines decentralization, censorship resistance, pseudonymity, verifiable cost of entry, and emergent consensus.

### 1.3 Contributions

This paper makes the following contributions:

1. Definition of the **Reputation Social Network (RSN)**, a decentralized protocol extending the DID framework to support bilateral reputation exchange.
2. A **nondeterministic verifier selection protocol** based on consistent hashing [6].
3. The **Expensive Radicalism Principle**, pricing claims according to their distance from network consensus through three compounding cost channels.
4. A **Reputable Authority model** with asymmetric incentives where false endorsement costs catastrophically more than legitimate fees.
5. A **gradual migration path** through parallel fiscal infrastructure.

## 2 Design Principles

The RSN architecture is constrained by five non-negotiable design principles.

**Continuity and Evolution.** The system coexists with existing institutions during a transition period of indefinite length. The transition must be reversible at every stage.

**Voluntariness and Responsibility.** Participation is voluntary, but voluntariness is coupled with responsibility: a participant who assumes control over an institutional function simultaneously assumes the accompanying obligations.

**Diversity and Cultural Neutrality.** Consensus emerges from bilateral interactions, not from preset rules imposed by system designers.

**Resilience and Censorship Resistance.** The cost of censoring the network must exceed the cost of tolerating it. Only full totalitarianism—at ruinous political and economic cost—can suppress the system.

**Consensus and Enforcement.** The system incorporates human tendencies toward pettiness, spite, and retributive satisfaction as load-bearing features. Misconduct is not prohibited; it is priced.

## 3 System Overview

### 3.1 Reputation Social Network

The RSN is a decentralized, peer-to-peer communication layer in which participants exchange verifiable information about real-world behavior. Its defining properties are: decentralized and uncensorable infrastructure; pseudonymous participation through DIDs; asymmetric cost structure (reading is cheap, writing is expensive); and cryptographic verifiability.

### 3.2 Architectural Components

The RSN consists of four primary components (Fig. 1):

1. **DID Layer.** Participant identities with declared rules, reputation metadata, and network routing.
2. **Claim Protocol.** Structured format for publishing assertions about real-world events.
3. **Verification Network.** Decentralized protocol for assigning verifiers using consistent hashing.
4. **Reputation Aggregation Layer.** Services that produce human-readable reputation summaries.

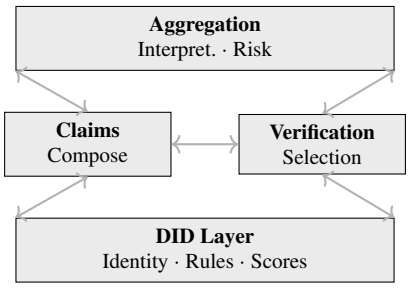


Figure 1: RSN four-layer architecture.

### 3.3 Participant Roles

A verification transaction involves up to six distinct DID roles (Fig. 2): **Issuer** ($DID_I$, creates and submits the claim), **Subject** ($DID_S$, the DID the claim is about—may coincide with Issuer), **Authority** ($DID_A$, endorses the claim for a fee; may also offer witness services—*optional*), **Witness** ($DID_{W_{1..n}}$, incog-

nito monitor of verifier honesty via challenge codes—*optional*), **Verifier** ($DID_V$, algorithmically selected to publish the claim), and the **RSN** (the network where verified claims are published). Any role may be delegated to another DID (Section 7.4). An Authority commonly bundles endorsement with witness services, but the two functions are logically independent.

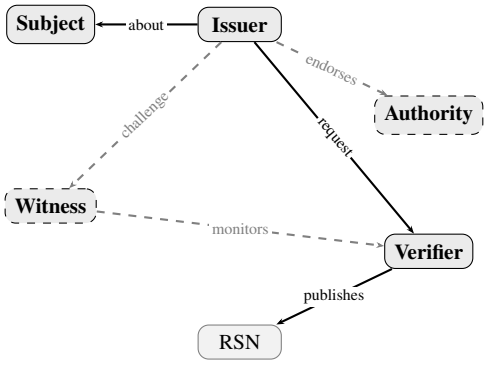


Figure 2: DID roles in a verification transaction. Dashed = optional (Authority, Witness).

# 4 Decentralized Identity Model

## 4.1 DID with Special Properties

The RSN extends the W3C DID Core specification [2] with: **Declared Rules** (machine-readable behavioral commitments), **Reputation Metadata** (aggregated behavioral score), and **Network Routing** (mutable onion gateway separate from the stable ring position).

## 4.2 Claim Lifecycle

A Claim proceeds through six stages (Fig. 3): composition, optional authority endorsement, verifier selection via consistent hashing, verification decision (accept or reject with iteration), publication, and reputation update for all parties.

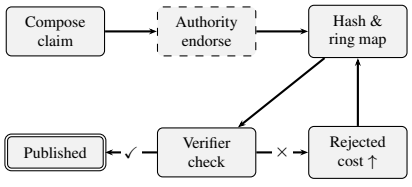


Figure 3: Claim lifecycle with iteration loop.

## 4.3 Relationship to W3C DID Core

The RSN's identity model is a proper superset of the W3C DID Core specification [2]. All RSN DIDs are valid W3C DIDs. The RSN-specific extensions are encoded as DID document properties defined in a dedicated DID method specification, compatible with existing infrastructure such as ION [7] and Ceramic [8].

# 5 Information Verification and Propagation

## 5.1 Cost of Information Entry

The RSN's core economic principle is the asymmetry between reading and writing. Reading reputation data approaches zero marginal cost for participants who are part of the DID network and have access commensurate with their reputation—newcomers must gradually earn broader access (see anti-leeching). Writing requires verifiable expenditure across three dimensions: time (waiting for verifier responses), energy (cryptographic signing and submission), and money (verifier fees). Reputation cannot be purchased instantly—it must be accumulated through sustained behavior.

## 5.2 Social Graph and Contact Depth

RSN is fundamentally a social network: each DID adds contacts (other DIDs) who permit the connection. These contacts have their own contacts, and so on. The algorithmic verifier search operates within a configurable depth $h$ of linked contacts (e.g., $h = 3$: one's direct contacts, their contacts, and contacts of contacts). This architecture does not require a single global blockchain—it naturally favors the emergence of communities/clusters with overlaps into other communities. Every DID participant must have a published **policy**—a machine-readable set of rules governing how incoming requests are evaluated. Policy violations are recorded in the network as negative artifacts.

## 5.3 Algorithmic Verifier Selection

The verification protocol employs consistent hashing [6] (Fig. 4). Verification is a paid service: verifiers operate for a fee, creating a market where competition drives pricing, speed, and reliability.

**Step 1.** The claim document is concatenated with the previous iteration's hash result (or $\varnothing$ on the first iteration) and hashed:

$$\text{pos} = f_h(\text{claim} \| \text{prev\_hash}) \tag{1}$$

The algorithm must include the *complete path* of all previous requests and responses—not merely a hash

of the previous iteration. Each verifier's full response (acceptance, rejection, quoted price, reason) is appended to the growing document, verifiable through cryptographic signatures at each step.

**Step 2.** The hash is mapped to $N$ positions on the consistent hash ring, evenly distributed (e.g., for $N = 4$: $+0°, +90°, +180°, +270°$). From each position, a clockwise search selects the nearest DID as a verifier candidate (Fig. 5). $N$ is a variable parameter that may depend on the percentage of currently active DIDs, time of day, or historical network responsiveness.

**Step 3.** The Verifier accepts (publishes, staking reputation) or rejects (returning to Step 1 with the rejection hash). When a node does not respond despite declaring online status, the next request must include all responses from all $N$ previous verifier candidates.

**Anti-leeching.** Target DIDs may set fees or access restrictions for queries from new DIDs with little reputation. This graduated mechanism (not binary) forces newcomers to build reputation through genuine activity before freely consuming others' data.

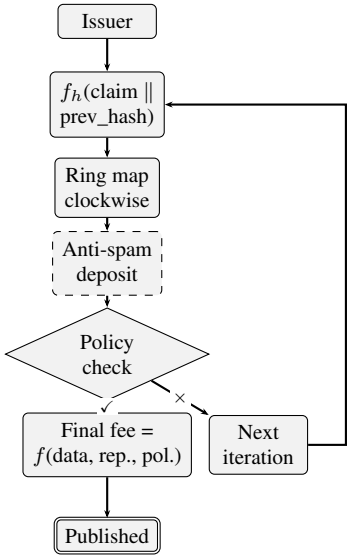


Figure 4: Verifier selection protocol. Anti-spam deposit is optional and may be refundable. Final fee is paid only on acceptance.

Each rejection appends the verifier's full response (including reasons and conditions) to the claim document, expanding its content. From the expanded document, a new hash is nondeterministically computed, mapping to a different ring position and selecting a different verifier. Documentation of the full path is mandatory—the issuer cannot predict or influence the next verifier. If a verifier ignores a request despite a compatible declared policy, witnesses (if present) record this, and the issuer may attach witness evidence upon final publication.

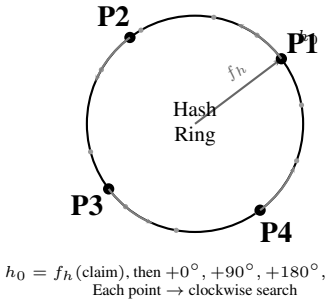


Figure 5: Multi-point selection ($N$=4). Hash $h_0$ sets the offset; 4 points evenly distributed from $h_0$.

## 5.4 Witness Protocol

The **Witness** role provides incognito accountability for verifier honesty through a cryptographic challenge-code protocol:

**Step W1.** The requester signs the verification request and sends it to $N_w$ witnesses—contacts from the requester's social network, or specialized witness-service providers operating for a fee.

**Step W2.** Each witness $w_i$ signs the entire signed request, retains the original signature $\sigma_i$, and computes a challenge code $c_i = h(\sigma_i)$. The challenge code is appended to the request.

**Step W3.** The request, now carrying $N_w$ challenge codes, is sent to the verifier. The verifier observes the codes but *cannot determine* who the witnesses are or whether the codes correspond to real signatures.

**Step W4 (honest verifier).** If the verifier responds according to their declared policy, the challenge codes remain opaque. Witnesses are never revealed. No additional data is published.

**Step W5 (policy violation).** If the verifier violates their policy (ignores the request, responds inconsistently), the requester holds the witnesses' original signatures $\sigma_i$. These can be published as companion data: anyone can verify $c_i = h(\sigma_i)$, proving the witness was present. The requester can publish independently—the verifier already signed the challenge codes in their response.

**The bluff property.** The requester may substitute random values for some or all challenge codes. The verifier cannot distinguish genuine codes (backed by high-reputation authorities) from noise. This *uncertainty* is the enforcement mechanism: every request carries the risk that a respected authority is watching incognito. The cost of maintaining this pressure is near zero (generating random numbers), while the poten-

tial cost of dishonesty for the verifier is catastrophic. Honest behavior is incentivized even in the absence of actual witnesses.

**Authority as witness.** An Authority that endorses a claim may bundle witness services: "I endorse your claim and serve as incognito witness during verification." This is a natural business model—the Authority already has context. However, the two roles are logically independent.

### 5.5 Expensive Radicalism Principle

Three cost channels compound simultaneously (Fig. 6):

**Channel 1: Iteration Cost.** Each rejection forces a new iteration consuming time and resources. Linear growth.

**Channel 2: Document Bloat.** Each iteration adds the full verifier response to the claim document. Growth is linear, but with a base offset: the initial payload (claim content, authority endorsement, cryptographic signatures) already has non-trivial size $B_0$. Total size after $k$ iterations: $S(k) = B_0 + k \cdot \Delta$, where $\Delta$ is the average response size.

**Channel 3: Verifier Risk Premium.** Risk is subjective. The verifier assesses whether the requesting DID's declared policies conflict with the verifier's own commercial, social, or political interests. The premium may escalate exponentially with perceived distance from the verifier's "ideal": $P(d) = \alpha \cdot e^{\beta d}$, where $d$ is the verifier's subjective distance metric. Beyond a personal no-go boundary, the verifier may refuse entirely.

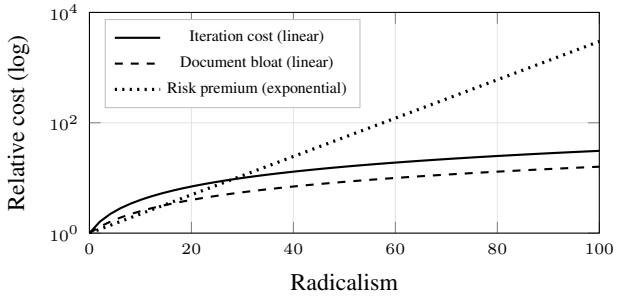


Figure 6: Cost escalation across three channels. The risk premium dominates at high radicalism.

Critically, this is not censorship. No claim is prohibited. The system prices risk (Table 1).

Table 1: Cost comparison across claim types.

| Type | Iter. | Doc | Fee | Total |
|---|---|---|---|---|
| Credible | $k_{\min}$ | $B_0$ | $P_{\min}$ | Low |
| Controversial | $k_{\text{med}}$ | $B_0 + k\Delta$ | $\alpha e^{\beta d}$ | Moderate |
| Radical | $k_{\max}$ | $\gg B_0$ | $\gg P_{\min}$ | Very high |
| Fraudulent | $\to \infty$ | — | — | Unpublishable |

## 6 Reputation and Risk Assessment

### 6.1 Reputation Accumulation Model

Reputation exhibits three properties: **slow accumulation** (incremental growth through consistent behavior), **rapid destruction** (a single fraud can reduce the score catastrophically), and **individual evaluation** (each consuming party may apply their own aggregation function).

### 6.2 Reputable Authorities

A Reputable Authority reviews claims and, if satisfied, co-signs them—staking its own reputation (Fig. 7). The asymmetry is by design: gaining reputation is slow (incremental $+\delta$ per honest endorsement), while losing it is catastrophic ($-\Delta \gg \delta$ per false endorsement; illustratively, e.g., $+2\%$ vs. $-70\%$). The optimal strategy is always to reject suspicious claims. The specific values of $\delta$ and $\Delta$ are system parameters.

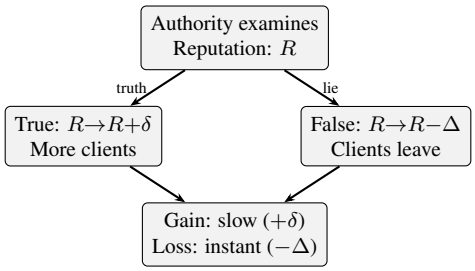


Figure 7: Reputable Authority—asymmetric incentives.

### 6.3 Multi-Dimensional Reputation

Reputation is not a scalar but a vector across multiple dimensions: financial reliability, personal integrity, professional competence, civic engagement, and others (Fig. 8). Different evaluation providers project this vector differently depending on context—a lender prioritizes financial reliability, an employer professional competence, a neighbor civic behavior. There is no single "correct" reputation score; only perspectives.

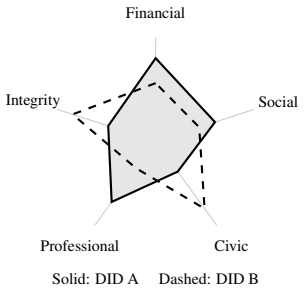


Figure 8: Multi-dimensional reputation vectors. Different DIDs exhibit different profiles across dimensions.

### 6.4 Democratized Risk Assessment

The RSN democratizes systematic risk assessment—a capability currently concentrated in financial institutions but applicable far beyond banking. Industrial conglomerates assessing supplier reliability, insurance companies pricing behavioral risk, due diligence for mergers and acquisitions, equipment lifetime estimation in manufacturing—anywhere counterparty risk requires assessment, the RSN provides a decentralized, market-driven alternative. A market of interpretation services translates raw multi-dimensional reputation data into actionable summaries, making counterparty evaluation accessible to any participant at marginal cost.

## 7 Consensus and Dispute Resolution

### 7.1 Decentralized Justice Model

The RSN reframes justice as a bilateral negotiation problem. The threat of permanent, verifiable reputation damage is a more effective deterrent than legal proceedings the harmed party cannot afford.

### 7.2 Emergent Consensus

Each participant maintains a personal satisfaction threshold. The network's aggregate consensus emerges as the statistical average of thousands of bilateral negotiations (Fig. 9). A response far above the consensus (barbaric) is expensive to publish. A response near the consensus (proportional) is cheap. The consensus is not a rule—it is a price signal.

### 7.3 Punishment Calibration

Each DID Holder publicly declares responses to specific misconduct categories. Disproportionately harsh

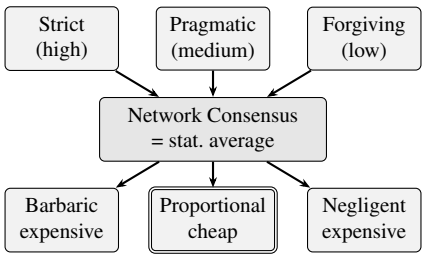


Figure 9: Emergent consensus from individual satisfaction thresholds.

or lenient declarations attract negative reputation signals. Proportional declarations are accepted without friction—a self-calibrating punishment system.

Consensus declarations are themselves subject to hypocrisy detection: declaratively committing to a consensus position while behaving inconsistently constitutes a reputational offense more severe than the underlying deviation, as it demonstrates intentional deception.

### 7.4 Delegation

All activities within a DID may be delegated to another DID. Verification, claim submission, consensus participation, dispute resolution—any role (Issuer, Authority, Witness, Verifier) can be transferred to a designated delegate DID. Delegation is revocable at any time. This enables specialization (e.g., professional verification services) while the Holder retains ultimate control and responsibility. Delegation applies across the entire system and is essential for scalability.

### 7.5 From Individual Morality to Emergent Social Contract

Each DID's declared policy represents a formalization of individual morality: what the Holder considers acceptable, how they respond to specific behaviors, what they require from counterparties. These policies are not imposed by a system designer—they are chosen by each participant and expressed in machine-readable form.

This formalization enables a three-stage emergence. First, individual morality is encoded as **policy**—a set of declared rules, thresholds, and response functions that the DID commits to follow. Second, the aggregate of all policies across the network produces **emergent ethics**—statistically observable norms that no single participant designed but that arise from the interaction of thousands of individual moral positions [9]. Third, these emergent ethics, expressed as shared behavioral

norms enforced through bilateral price signals, constitute a **measurable social contract**—not a theoretical construct that nobody signed [10], but an empirically observable set of rules backed by actual behavior.

This process mirrors findings from moral foundations theory [11]: human morality is intuitive, pluralistic, and culturally variable. The RSN does not require moral consensus as an input; it produces behavioral consensus as an output. The resulting social contract is not static—it evolves as participants join, leave, and adjust their policies in response to network feedback. Unlike classical social contract theory, this contract is revocable, observable, and enforceable without a sovereign.

# 8 Economic Model

The preceding sections described a tool—the RSN's verification mechanisms, cost structure, and emergent consensus. This section describes a methodology for applying this tool to fiscal and governance transition. The tool is general-purpose; the methodology below is one possible application.

## 8.1 Incentive Structure

Reputation as capital creates direct incentives for honest behavior. In normal operation, participants build reputation naturally through everyday transactions and fulfilled commitments. The primary expenditure is on *negative* claims—recording harm done by others. This asymmetry incentivizes useful, reliable behavior: being honest costs nothing extra, while being harmful creates a documented trail that others will pay to preserve. Hypocrisy—declaring rules without following them—is the most expensive failure mode, as it demonstrates intentional deception.

## 8.2 Parallel Fiscal System

Three components enable competitive pressure: (1) **Simple Tax**—a single proportional rate with no exceptions, initially marginally below the existing effective rate; (2) **Electronic Spending Registration (ESR)**—inverting the Czech EET model so citizens surveil state expenditures; (3) **Citizen-Directed Tax Allocation**—illustratively growing from 5% in Year 1 to 40% by Year 10 (Table 2). The actual tempo depends on the speed at which free-market alternatives to state services emerge; the function of time need not be linear.

Table 2: Citizen tax allocation over time.

| Year | Citizen-Directed | State-Managed |
|---|---|---|
| 1 | 5% | 95% |
| 3 | 10% | 90% |
| 5 | 15% | 85% |
| 7 | 25% | 75% |
| 10 | 40% | 60% |

# 9 Migration Path

The migration proceeds through three phases (Fig. 10): **Phase 1: Overlay** (RSN as informational layer, state is sole provider), **Phase 2: Competition** (RSN provides functional alternatives, dual-system use), and **Phase 3: Transition** (RSN outperforms state equivalents, voluntary migration). The transition is reversible at every stage.

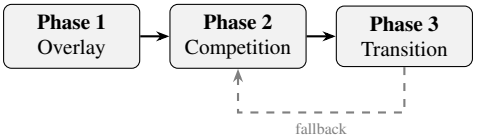


Figure 10: Three-phase migration—reversible at every stage.

The primary pressure mechanism is fiscal: when conditional taxpayers reach an "economically significant minority $X$"—a group large enough that repression against it causes non-trivial economic damage and civil disobedience becomes a credible threat—the state enters negotiation. For illustration, we use $X \approx 15\%$ of GDP, but the actual threshold depends on the specific economy.

# 10 Security Analysis

We consider five adversary categories: individual bad actors, organized groups, corporate adversaries, state-level adversaries, and protocol-level adversaries.

**Sybil resistance** [12]. Building reputation requires sustained, verifiable interaction. The cost of a successful Sybil attack scales linearly with fake identities and quadratically with the target reputation level. Operating multiple DIDs in parallel is possible but expensive by design—each identity must independently accumulate reputation through genuine activity. In free societies this cost deters casual abuse; in dictatorships, parallel DIDs become a survival tool enabling compartmentalized resistance, underground coordination, and safer black-market navigation.

**Collusion defense.** Patterns of mutual endorse-

ment among closed groups are detectable through social graph analysis. Reputation aggregation functions discount claims from tightly clustered sources.

**State-level adversary.** Onion routing, absence of centralized infrastructure, and distribution of the Verifier role across the participant base ensure that suppression requires measures disproportionate to any benefit.

**Privacy vs. accountability.** Pseudonymity—a middle ground between anonymity and identity disclosure—allows DIDs to accumulate meaningful reputation without revealing the Holder's real-world identity.

## 11 Privacy Considerations

The RSN's privacy model is built on pseudonymity. The Holder controls identity disclosure: minimal (pure pseudonym), selective (specific credentials linked), or full (legal identity associated). Published claims are permanent—the system supports contextual correction but not erasure. A Holder may abandon a compromised DID and start anew, forfeiting accumulated reputation.

## 12 Related Work

The W3C DID Core specification [2] and Ceramic Network [8] provide the identity foundation. EigenTrust [13] demonstrated distributed reputation aggregation without central authority. SBTs [3] introduced non-transferable social tokens. The RSN differs in its emphasis on economic cost as a quality filter and its mechanism for pricing radicalism.

DAOs [4] and quadratic voting [5] demonstrated feasibility of decentralized governance. The RSN derives consensus from bilateral price negotiations rather than voting.

The proposal draws on anarcho-capitalist theory [14, 15] for private service provision but departs in two critical respects: it designs for spite and retributive impulses rather than assuming rationality, and it proposes gradual transition rather than revolution. The concept of emergent social contracts has antecedents in Hayek's theory of spontaneous order [16].

**Anarcho-agorism.** The RSN shares significant common ground with agorist counter-economics [17]—particularly the emphasis on building parallel institutions and voluntary exchange. However, agorism tends to assume rational self-interest as a sufficient coordination mechanism and often lacks a concrete migration path from existing state structures. The RSN explicitly incorporates non-rational behavioral tendencies and provides a fiscal pressure mechanism for gradual transition.

**Meritocracy.** The RSN may represent a first functional description of how meritocracy [18] can emerge as a systemic property rather than a slogan. Traditional meritocratic systems fail because they require a central authority to define and enforce "merit." In the RSN, merit emerges from bilateral evaluations: those who deliver value accumulate reputation; no committee decides who has merit.

**Property as privilege.** The RSN departs fundamentally from anarcho-capitalist and Austrian-school treatments of property rights as natural and absolute. In the RSN framework, property is a *privilege*—a social recognition that can, in extreme cases, be withdrawn by the community when a participant fundamentally violates shared norms (betrayal, refusal to defend the community in existential crisis, systematic exploitation). This is not state expropriation but a withdrawal of social recognition by the network of bilateral relationships.

## 13 Discussion and Limitations

**Cold start.** The RSN's value depends on network effects; early adopters face limited value until participation reaches a threshold. However, even from early stages, the DID network serves as a useful supplementary information source. Early reputation evaluation and authority assessment are expected to develop gradually with increasing sophistication.

**Anti-leeching and access control.** Target DIDs may impose graduated fees and access restrictions on queries from low-reputation DIDs. This is not binary but a continuous scale: the less reputation a querying DID has, the less information it receives, the longer it waits, and the more it pays. This mechanism incentivizes genuine participation over passive consumption.

**Inequality of access.** The cost of publishing claims may filter legitimate grievances from economically disadvantaged participants. Mitigation: every participant simultaneously serves as a potential verifier (or delegates this role) and earns verification fees. Over a sufficient time horizon, a non-radical participant should approach financial neutrality—verification income approximately offsetting publication costs. This is a statistical expectation, not a guarantee.

**Reputation inequality.** Early adopters accumu-

late advantages that may create barriers for latecomers. However, reputation evaluation is performed by third-party providers who may choose to mitigate first-mover advantage through their aggregation algorithms. Being first with a good reputation is a legitimate comparative economic advantage—and that is by design.

**Open questions.** Optimal cost functions, aggregation standards, protocol governance, and interaction with AI-generated synthetic reputation data remain areas for future work.

This paper does not claim that the RSN will produce optimal outcomes in all circumstances. It claims that the RSN represents a *feasible* alternative—technically implementable, economically self-sustaining, and socially compatible with human behavioral tendencies as they actually are.

## 14 Conclusion

This paper has presented the Reputation Social Network, a decentralized protocol enabling pseudonymous, verifiable reputation exchange. The Expensive Radicalism Principle ensures that fraudulent claims are priced out without censorship. The proposed migration path enables gradual, reversible transition from existing institutions. The design incorporates human nature as it is—including pettiness, spite, and the desire for retributive satisfaction—rather than requiring ideological transformation. The result is not utopia but a system where justice is accessible, reputation is earned, and the cost of misconduct is borne by the party who caused it.

## References


[1] eská republika, "Zákon . 549/1991 sb., o soudních poplatcích, ve znní pozdjích pedpis," 1991, sazebník soudních poplatk; náklady ízení dle zák. . 99/1963 Sb. (OS) a vyhl. . 177/1996 Sb. (advokátní tarif).

[2] M. Sporny, D. Longley, M. Sabadello, D. Reed, O. Steele, and C. Allen, "Decentralized identifiers (DIDs) v1.0," W3C Recommendation, Jul. 2022.

[3] E. G. Weyl, P. Ohlhaver, and V. Buterin, "Decentralized society: Finding Web3's soul," *SSRN Electronic Journal*, May 2022.

[4] V. Buterin, "DAOs, DACs, DAs and more: An incomplete terminology guide," Ethereum Blog, May 2014.

[5] V. Buterin, Z. Hitzig, and E. G. Weyl, "A flexible design for funding public goods," *Management Science*, vol. 65, no. 11, pp. 5171–5187, 2019.

[6] D. Karger, E. Lehman, T. Leighton, R. Panigrahy, M. Levine, and D. Lewin, "Consistent hashing and random trees: Distributed caching protocols for relieving hot spots on the World Wide Web," in *Proc. 29th ACM STOC*, 1997, pp. 654–663.

[7] D. Buchner, "ION — a scalable DID method based on Bitcoin," Microsoft, 2021.

[8] Ceramic Network, "Ceramic protocol specification," GitHub, 2023. [Online]. Available: https://github.com/ceramicnetwork

[9] É. Durkheim, *The Division of Labor in Society*. Free Press, 1893, english translation 1984.

[10] J. Rawls, *A Theory of Justice*. Harvard University Press, 1971.

[11] J. Haidt, *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. Vintage Books, 2012.

[12] J. R. Douceur, "The Sybil attack," in *Proc. 1st International Workshop on Peer-to-Peer Systems (IPTPS)*, 2002, pp. 251–260.

[13] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina, "The EigenTrust algorithm for reputation management in P2P networks," in *Proc. 12th International Conference on World Wide Web*, 2003, pp. 640–651.

[14] M. N. Rothbard, *For a New Liberty: The Libertarian Manifesto*. Macmillan, 1973.

[15] H.-H. Hoppe, *Democracy: The God That Failed*. Transaction Publishers, 2001.

[16] F. A. Hayek, *Law, Legislation and Liberty*. University of Chicago Press, 1973.

[17] S. E. I. Konkin, *An Agorist Primer*. KoPubCo, 2006.

[18] M. Young, *The Rise of the Meritocracy*. Thames and Hudson, 1958.