# Mathematical Foundations of Decentralized Trust: From Bitcoin to OMXUS

Rigorous Analysis, Security Arguments, and Open Problems

OMXUS Research Initiative
research@omxus.com

February 2026 — Version 2.0

**Abstract**

We present a mathematical treatment of decentralized trust systems, beginning with the rigorous foundational work on Bitcoin's Proof-of-Work consensus and extending to the Proof-of-Humanity framework implemented in OMXUS. For Bitcoin, we reproduce established results: exact double-spend probabilities via the regularized incomplete beta function, and protocol stability via martingale theory. For OMXUS, we develop security arguments for social consensus mechanisms, clearly distinguishing between *proven theorems*, *security bounds under stated assumptions*, and *open conjectures*. We explicitly acknowledge where OMXUS's security model differs fundamentally from Bitcoin's computational model, and identify open problems requiring further research. This paper aims for intellectual honesty: Bitcoin's security is mathematically proven; OMXUS's security rests on economic and social assumptions that we formalize but cannot prove with equal rigor.

**Keywords:** blockchain, cryptography, Sybil resistance, web-of-trust, security arguments

## Contents

## Notation and Rigor Levels

To maintain intellectual honesty, we distinguish:

- **Theorem**: Mathematically proven result with complete proof
- **Proposition/Lemma**: Supporting results with proofs
- **Security Bound**: Upper/lower bound under explicitly stated assumptions
- **Heuristic**: Plausible argument without formal proof
- **Conjecture**: Unproven claim we believe to be true
- **Open Question**: Problem requiring further research

## Part I
# Mathematical Foundations of Bitcoin

*The results in this part are mathematically rigorous, building on established work by Nakamoto [1], Rosenfeld [5], and Grunspan-Pérez-Marco [2, 3].*

## 1 The Mining Model

### 1.1 Poisson Process Foundations

Consider a miner with fraction $0 < p \leq 1$ of the total network hashrate. The network validates blocks at an average rate of one per $\tau_0 = 10$ minutes.

**Theorem 1.1** (Mining Time Distribution). *The inter-block mining time $T$ for a miner with hashrate fraction $p$ follows an exponential distribution:*

$$f_T(t) = \alpha e^{-\alpha t}, \quad \alpha = \frac{p}{\tau_0} \tag{1}$$

*Proof.* The pseudo-random properties of SHA-256 ensure that each hash attempt is an independent Bernoulli trial with success probability $p_{\text{success}} = \frac{\text{target}}{2^{256}}$. For a miner performing $h$ hashes per second, the time to first success follows a geometric distribution that, in the continuous limit, converges to exponential with rate $\alpha = h \cdot p_{\text{success}}$. The memoryless property follows from the independence of hash attempts. □

**Corollary 1.2** (Block Count Distribution). *Let $N(t)$ count the number of blocks validated by time $t$. Then $N(t)$ follows a Poisson distribution:*

$$\mathbb{P}[N(t) = n] = \frac{(\alpha t)^n}{n!} e^{-\alpha t} \tag{2}$$

## 2 The Double-Spend Problem

**Theorem 2.1** (Double-Spend Probability, Grunspan-Pérez-Marco [2]). *After $z$ confirmations, the probability of a successful double-spend by an attacker with hashrate $q < 1/2$ (where $p = 1-q$) is:*

$$P(z) = I_{4pq}(z, 1/2) \tag{3}$$

*where $I_x(a, b)$ is the regularized incomplete beta function:*

$$I_x(a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \int_0^x t^{a-1}(1-t)^{b-1}\, dt \tag{4}$$

*Proof.* The attacker's block count $X_n = N'(S_n)$ when honest miners reach block $n$ follows a negative binomial distribution with parameters $(n, p)$. The probability of catching up from $z$ blocks behind is computed via the ballot problem. The regularized incomplete beta function arises from the cumulative distribution. Full proof in [2]. □

**Corollary 2.2** (Exponential Security). *Let $s = 4pq < 1$. As $z \to \infty$:*

$$P(z) \sim \frac{s^z}{\sqrt{\pi(1-s)z}} \tag{5}$$

*Proof.* Apply Watson's Lemma to the integral representation of $I_x(a, b)$. □

## 3 Protocol Stability

**Theorem 3.1** (Optimal Strategy, Grunspan-Pérez-Marco [3])**.** *In the absence of difficulty adjustment, the optimal mining strategy is to publish all mined blocks immediately upon discovery.*

*Proof.* Define revenue ratio $\Gamma = \mathbb{E}[R]/\mathbb{E}[\tau]$. For any strategy with cycle duration $\tau$ (a stopping time), Doob's optional stopping theorem applied to the martingale $M(t) = N(t) - \alpha t$ gives $\mathbb{E}[N(\tau)] = \alpha\mathbb{E}[\tau]$. Thus $\Gamma \leq pb/\tau_0 = \Gamma_H$, with equality only for immediate publication. □

**Remark 3.2.** *This theorem holds only without difficulty adjustment. With adjustment, selfish mining can be profitable for $q > (1-\gamma)/(3-2\gamma)$ where $\gamma$ is network connectivity [4].*

# Part II
# Security Analysis of OMXUS

*Unlike Part I, the results here vary in rigor. We clearly label each result's status. OMXUS's security rests on social and economic assumptions that cannot be formalized with the same precision as computational assumptions.*

## 4 Fundamental Differences from Bitcoin

**Remark 4.1** (Epistemic Honesty)**.** *Bitcoin's security reduces to computational hardness assumptions (e.g., SHA-256 preimage resistance) that are:*

1. *Precisely defined*
2. *Measurable in hash operations*
3. *Falsifiable by cryptanalysis*

*OMXUS's security reduces to* social cost *assumptions that are:*

1. *Difficult to precisely define*
2. *Not directly measurable*
3. *Dependent on human behavior*

*This is a fundamental limitation, not a flaw to be hidden.*

## 5 The Web-of-Trust Model

**Definition 5.1** (Verification Graph)**.** *The OMXUS verification graph is a directed graph $G = (V, E)$ where:*

- *$V$ is the set of verified identities*
- *$(u, v) \in E$ if $u$ vouched for $v$*

*Each vertex $v$ has in-degree $d^-(v) \geq k$ (minimum vouches, currently $k = 3$).*

### 5.1 Sybil Resistance

**Assumption 5.2** (Social Cost Existence)**.** *There exists a positive cost $c_{social} > 0$ incurred by a voucher when participating in fraudulent verification. This cost includes:*

1. *Risk of trust score reduction (quantifiable within the system)*
2. *Risk of network ejection (quantifiable)*
3. *Reputational damage in real-world social network (not quantifiable)*
4. *Time cost of in-person meeting (quantifiable in hours)*

**Remark 5.3.** *Unlike Bitcoin's hash cost (measurable in joules or dollars), $c_{social}$ is heterogeneous across individuals and contexts. A more honest statement is that $c_{social}$ is a random variable with unknown distribution, and we assume $\mathbb{E}[c_{social}] > 0$.*

**Security Bound 5.4** (Sybil Attack Cost — Lower Bound)**.** *Under Assumption 5.2, creating $n$ Sybil identities requires convincing at least $\lceil nk/m \rceil$ vouchers (where $m$ is max vouches per person), giving:*

$$\mathcal{C}(n) \geq \frac{nk}{m} \cdot \mathbb{E}[c_{social}] \tag{6}$$

*This is a **lower bound**, not an exact cost. The actual cost may be higher due to coordination overhead, detection risk, etc.*

**Open Question 5.5.** *Can $c_{social}$ be empirically measured? Possible approaches:*

1. *Market price for fraudulent vouches (if a black market exists)*
2. *Survey-based willingness-to-accept for vouching strangers*
3. *Revealed preference from attempted attacks*

### 5.2 Verification Security

**Assumption 5.6** (Voucher Independence)**.** *The $k$ vouchers for a new identity are selected independently, such that if each voucher has probability $p_c$ of colluding, then:*

$$\mathbb{P}[\textit{All } k \textit{ collude}] = p_c^k \tag{7}$$

**Security Bound 5.7** (Fraud Probability — Under Independence)**.** *Under Assumption 5.6 with $k = 3$ and $p_c = 0.01$:*

$$\mathbb{P}[\textit{Fraudulent verification}] \leq p_c^k = 10^{-6} \tag{8}$$

**Remark 5.8** (Critical Limitation)**.** *Assumption 5.6 is **known to be false** in practice. Colluding vouchers are correlated by definition — they coordinate to commit fraud. The independence assumption provides a lower bound that is optimistic.*

*A more realistic model treats voucher collusion as a clique detection problem:*

$$\mathbb{P}[\textit{Fraud}] = \mathbb{P}[\exists \textit{ colluding } k\textit{-clique in voucher graph}] \tag{9}$$

*This depends on the graph structure and is harder to analyze.*

**Heuristic 5.9** (Collusion Mitigation)**.** *The protocol requires vouchers to be "sufficiently separated" in the social graph. If we require:*

$$d(u_i, u_j) \geq 2 \quad \forall i \neq j \tag{10}$$

*(no two vouchers are directly connected), then forming a colluding group requires compromising a distributed set, increasing coordination cost. This is a heuristic defense, not a proven guarantee.*

### 5.3 Trust Score Dynamics

**Definition 5.10** (Ripple Responsibility). *When vertex $v$ commits a violation with severity $s \in [0,1]$, trust scores update:*

$$\tau(u) \leftarrow \tau(u) - s \cdot \delta \cdot r^{d(u,v)} \tag{11}$$

*where $d(u,v)$ is graph distance and $r \in (0,1)$ is the decay rate.*

**Proposition 5.11** (Bounded Ripple — Fixed Decay). *For decay rate $r$ and maximum out-degree $\Delta$, if $r\Delta < 1$, then total penalty is bounded:*

$$\sum_{u \in V} \Delta\tau(u) \leq s \cdot \delta \cdot \frac{1}{1 - r\Delta} \tag{12}$$

*Proof.* At distance $i$, there are at most $\Delta^i$ vertices. Total penalty:

$$\sum_{i=0}^{\infty} \Delta^i \cdot r^i = \frac{1}{1 - r\Delta} \tag{13}$$

which converges iff $r\Delta < 1$. □

**Remark 5.12.** *The original claim used $r = 1/3$, requiring $\Delta < 3$, which is unrealistically restrictive (each person can vouch for at most 2 others). In practice, we must set $r = 1/(\Delta + \epsilon)$ for convergence, which weakens the incentive effect as the network grows.*

**Open Question 5.13.** *What is the optimal tradeoff between decay rate $r$ and incentive strength? Too fast decay ($r$ small) makes ripple responsibility meaningless; too slow decay ($r$ large) creates unbounded liability.*

## 6 Emergency Response: An Idealized Model

**Assumption 6.1** (Homogeneous Poisson Responders). *Verified responders are distributed as a homogeneous Poisson point process with intensity $\rho$ (responders per unit area).*

**Proposition 6.2** (Coverage Under Poisson Assumption). *Under Assumption 6.1, the probability of at least one responder within distance $r$ is:*

$$P_{coverage}(r) = 1 - e^{-\rho\pi r^2} \tag{14}$$

*Proof.* Standard result for Poisson point processes: the count in a disk of radius $r$ is Poisson with parameter $\rho\pi r^2$. □

**Remark 6.3** (Critical Limitation of Poisson Assumption). *Assumption 6.1 is **unrealistic:***

1. *Real populations are highly clustered (cities, buildings, rooms)*
2. *Density varies dramatically by time of day*
3. *Indoor/outdoor barriers affect actual response time*
4. *Responder willingness is not constant (availability, capability)*
5. *The "60-second response" claim assumes instantaneous notification*

*The Poisson model provides an **optimistic lower bound** on required density, not a realistic prediction.*

**Conjecture 6.4** (Realistic Coverage)**.** *For a clustered population model (e.g., Thomas cluster process) with the same mean density $\rho$, the actual coverage probability satisfies:*

$$P_{real}(r) < P_{Poisson}(r) \tag{15}$$

*The gap depends on clustering parameters and is an open problem.*

**Heuristic 6.5** (Practical Response Estimate)**.** *A more defensible claim is: "In areas where OMXUS adoption exceeds $\rho_{\min}$ responders per km$^2$, and conditional on responder availability, the expected number of potential responders within walking distance is at least 1." This is weaker but more honest.*

## 7 Cryptographic Anchoring

**Theorem 7.1** (Inherited Bitcoin Security)**.** *An OMXUS epoch root $R_e$ anchored to Bitcoin at height $B$ with $z$ confirmations has reversion probability:*

$$\mathbb{P}[Revert\ R_e] \leq I_{4pq}(z, 1/2) \tag{16}$$

*where $q$ is the attacker's Bitcoin hashrate fraction.*

*Proof.* Direct application of Theorem 2.1. Reverting $R_e$ requires reverting the Bitcoin block containing it. □

**Remark 7.2.** *This is the* ***only fully rigorous security result*** *for OMXUS, because it directly inherits Bitcoin's proven security. The social layer security (Sybil resistance, fraud bounds) remains in the "security argument" category.*

## 8 Honest Comparison with Bitcoin

Table 1: Security Comparison — Honest Assessment

| **Property** | **Bitcoin** | **OMXUS** |
|---|---|---|
| Consensus Type | Proof-of-Work | Proof-of-Humanity (social) |
| Attack Cost | $\mathcal{O}(n \cdot c_{\text{hash}})$<br>*Measurable in joules/$* | $\mathcal{O}(n \cdot k \cdot c_{\text{social}})$<br>*Not precisely measurable* |
| Security Proofs | **Rigorous** (Poisson, martingales) | **Bounds under assumptions** |
| Key Assumption | SHA-256 hardness<br>*Cryptographically standard* | Social cost > 0<br>*Behavioral, untestable* |
| Finality | Probabilistic ($I_{4pq}$)<br>*Proven formula* | Social (3 vouches)<br>*Under independence assumption* |
| Falsifiability | SHA-256 break would invalidate<br>*Would be obvious* | Coordinated Sybil attack<br>*Might be subtle* |

## 9 Open Problems and Future Work

1. **Formalizing Social Cost**: Can $c_{\text{social}}$ be defined in terms of measurable quantities (time, money, reputation tokens)?

2. **Non-Independent Collusion**: Develop security bounds for correlated voucher behavior using techniques from dependent percolation or epidemic models.
3. **Realistic Spatial Models**: Analyze emergency response under Thomas cluster processes or other clustered point processes.
4. **Optimal Ripple Decay**: Find the Pareto-optimal decay rate balancing convergence, incentive strength, and fairness.
5. **Empirical Validation**: Conduct controlled experiments to measure actual Sybil attack costs in web-of-trust systems.
6. **Adversarial Analysis**: Model sophisticated attackers who exploit graph structure to minimize social cost.

## 10 Conclusions

We have presented Bitcoin's security with mathematical rigor, reproducing established results on double-spend probabilities and protocol stability. For OMXUS, we have developed security arguments that are plausible but not proven to the same standard.

**What OMXUS can claim**:

- Bitcoin-level security for anchored epoch roots (Theorem 7.1)
- Linear scaling of Sybil attack cost with number of fake identities (Bound 5.4, under assumptions)
- Low fraud probability under voucher independence (Bound 5.7, assumption known to be approximate)

**What OMXUS cannot claim**:

- Mathematically proven Sybil resistance comparable to Bitcoin's double-spend resistance
- Precise quantification of social attack costs
- Guaranteed emergency response coverage

The intellectual honesty of this distinction is itself a form of trustworthiness. We invite the research community to strengthen these arguments or identify weaknesses we have missed.

## References

[1] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008.

[2] C. Grunspan and R. Pérez-Marco, "Double spend races," *Int. J. Theoretical and Applied Finance*, vol. 21, no. 08, 2018.

[3] C. Grunspan and R. Pérez-Marco, "On the profitability of selfish mining," arXiv:1805.08281, 2018.

[4] I. Eyal and E. G. Sirer, "Majority is not enough: Bitcoin mining is vulnerable," *Commun. ACM*, vol. 61, pp. 95–102, 2018.

[5] M. Rosenfeld, "Analysis of hashrate-based double spending," arXiv:1402.2009, 2014.

[6] J. R. Douceur, "The Sybil attack," in *IPTPS*, 2002.

[7] H. Yu et al., "SybilGuard: Defending against Sybil attacks via social networks," *ACM SIGCOMM*, 2006.