CHICAGO BOOTH | The Initiative on Global Markets

IGM Working Paper

Smart Contracts, Dumb Money: Open Source Lemons

Tingyi Lin
Central University of Finance and Economics
Weiyu Qi
University of Chicago
Karl Yu
Boston College
Huanxi Zhang
University of Wisconsin-Madison

Initiative on Global Markets
The University of Chicago, Booth School of Business
This paper also can be downloaded without charge from the
Kilts Center at Chicago Booth Marketing Data Center Paper
https://ssrn.com/abstract=6371958

Smart Contracts, Dumb Money: Open Source Lemons*

Tingyi Lin† Weiyu Qi‡ Karl Yu§ Huanxi Zhang¶

Abstract

Why does a market structure built on radical transparency paradoxically foster the proliferation of low-quality assets? We address this puzzle by identifying a Transparency Paradox in algorithmic financial markets: despite the open-source nature of smart contracts, high-quality innovation is systematically crowded out by AI-generated noise and memecoins. We develop a dynamic model of adverse selection where information acquisition is endogenous and costly. We show that when technical complexity raises the cost of verification above a critical threshold, human investors optimally choose ignorance—becoming dumb money. This creates a pooling equilibrium that strips genuine projects of their valuation premium. Consequently, innovators face a complexity penalty and exit, leaving only assets that are costless to verify but possess zero fundamental value. Our findings challenge the regulatory orthodoxy of mandated disclosure, suggesting that in algorithmic markets, overwhelming information can exacerbate asymmetry by exceeding human cognitive capacity. We demonstrate that AI-assisted verification tools and automated regulatory auditing are critical mechanisms to break this complexity trap and restore market efficiency

JEL Classification: D83, G11, G14, G23, O31

Keywords: Rational Inattention, Decentralized Finance, Complexity Trap, Costly State Verification, Transparency Paradox

*We are deeply grateful to Dr. Ke Dong (Tsinghua University PBC School of Finance), as well as other faculty members, for pre-reviewing this manuscript and providing constructive comments.
We extend our gratitude to professor Hideo Konishi (Boston College) and Ruoran Lai (Sun Yat-sen University) for stimulating discussions and assistance during the writing process.
We also thank the ACM Conference on Economics and Computation(ACM EC) reviewers and discussants for their valuable feedback. All remaining errors are our own.
Funding and Disclosure: This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†Central University of Finance and Economics. Email: hayashi@email.cufe.edu.cn (**Corresponding author**)
‡The University of Chicago. Email: weiyuq@uchicago.edu
§Boston College. Email: yulin@bc.edu
¶University of Wisconsin-Madison. Email: hzhang2488@wisc.edu

"A wealth of information creates a poverty of attention."

— *Herbert A. Simon (1971)*

1 Introduction

Crypto-asset markets present a puzzle that standard disclosure-based finance theory does not fully resolve. These markets are built on radical transparency: transactions are public, contracts are often open-source, and large fractions of market activity are mechanically observable on-chain. Yet the equilibrium allocation of capital is frequently dominated by assets that are easy to narrate rather than projects that are hard to understand but potentially productive. Memecoins and other low-dimensional narrative assets often attract disproportionate attention and liquidity, while technically ambitious projects struggle to obtain financing on terms commensurate with their fundamentals. Why does a market organized around open code and public ledgers so often resemble a lemons market?

Our answer is that the relevant friction is not hidden information per se. In crypto markets, much of the raw information is public. The bottleneck is instead the cost of converting public disclosure into decision-relevant valuation. Reading a smart contract is not the same as understanding its economic logic, attack surface, governance contingencies, and interoperability risks. For simple narrative assets, the mapping from disclosure to an investable story is low-dimensional and cheap. For technology-intensive projects, the mapping is cognitively expensive and often non-scalable. A transparent market can therefore remain informationally fragile when investors are attention-constrained.

We develop a static equilibrium model with two linked stages. In the first stage, investors allocate limited attention across projects using a rational inattention (RI) framework in the spirit of Sims (2003) and Matějka and McKay (2015). In the second stage, a costly state verification (CSV) technology determines whether an attended project is worth auditing and, conditional on verification, whether capital is deployed. The key modeling distinction is

that the RI stage produces *visibility* or *consideration*—the probability that a project reaches the investor's evaluation queue—rather than an irreversible terminal investment decision. Verification is therefore conditional on attention. This timing makes the upstream attention margin and the downstream audit margin complementary but non-equivalent.

To preserve the tractability of the discrete-choice RI problem, we keep the Shannon/KL information-cost scale parameter common across options and model heterogeneity in project difficulty through a *processability* wedge. A project with higher effective complexity is one for which the same body of public disclosure is translated less efficiently into a decision-relevant payoff index. Legibility investments by issuers lower this wedge; hype investments raise salience. This distinction is economically important. The model does not say that information literally disappears. It says that the same public information has very different decision relevance once attention and processing constraints are acknowledged.

The model delivers three main mechanisms. First, as effective complexity rises, attention becomes less state-dependent: investors increasingly fail to discriminate between good and bad versions of technically demanding projects. Second, because attention is allocated through a generalized-logit object, low-complexity and high-salience projects crowd out complex projects through a denominator effect. Third, verification is triggered only for projects that first become sufficiently visible. A project can therefore be worth verifying *conditional on consideration* and still never be verified in equilibrium because it fails to enter the candidate set. This is the paper's central "invisible innovation" channel.

These forces generate a complexity penalty for genuine innovation. Projects with high fixed development costs require visibility and financing to recover those costs. But the same technological sophistication that generates upside also makes them harder to process. As meme supply expands or narrative salience becomes cheaper to produce, complex projects lose visibility, fail to trigger verification, and eventually fail the entry condition. The result is a transparency paradox: more public disclosure does not automatically produce more informative capital allocation when the scarce input is not data, but attention.

Our contribution is threefold. First, we provide a tractable theory of how computational and organizational complexity can shape market structure even when information is nominally public. Second, we show that verification technologies such as audits, fraud proofs, or due diligence do not discipline markets uniformly; their effectiveness is conditional on the prior allocation of attention. Third, we connect these mechanisms to non-linear market selection. In the baseline model, sufficiently high complexity can push the economy across a threshold beyond which tech entry collapses. In an extension with ecosystem externalities, the same logic generates path dependence and hysteresis: once a market becomes meme-dominant, rebuilding the information-processing infrastructure needed to support complex projects can require a larger improvement than the deterioration that caused the collapse.

The policy implication is not that disclosure is useless. It is that disclosure must be made *processable*. In this environment, longer whitepapers, more raw code, or broader data dumps are not enough if they do not lower the effective cost of translating public information into valuation. The relevant interventions are standardization, machine-readable disclosures, reusable audit templates, and credible certification layers that compress high-dimensional technical information into low-cost verifiable claims.

The rest of the paper proceeds as follows. Section 2 maps the institutional features of crypto markets to the model's three key frictions: processability, salience engineering, and triggered verification. Section 3 presents the model. Section 4 characterizes the RI allocation, the candidate-set mechanism, the verification rule, and the technology-entry threshold. Section 5 presents numerical illustrations of the complexity trap in the baseline model and hysteresis in an extension with ecosystem externalities. Section 6 discusses mechanism-design implications. Section 7 concludes.

2 Institutional Background and Stylized Facts

The institutional starting point is simple: crypto markets make disclosure unusually abundant, but they do not make interpretation cheap. We organize the background discussion around three empirical regularities that motivate the model. First, a striking illustration of this dynamic is the rise of "fair launch" platforms like *pump.fun* on Solana. Historically, issuing a cryptocurrency required at least minimal technical competency in smart-contract programming and capital for initial liquidity. By April 2024, contemporaneous coverage described Pump.fun as allowing anyone to create an instantly tradable token without seed liquidity for less than $2 and in under a minute; by August 2024, even the 0.02 SOL creator-side deployment payment had been shifted to the first buyer, pushing issuer-side launch costs effectively to zero (The Block, 2024a,b). Current platform documentation continues to present coin creation as effectively turnkey: coin creation itself is free off-chain, and only a first buy plus gas is required to put the token on-chain; third-party protocol listings likewise summarize Pump.fun as a one-click launchpad (Pump.fun Help Center, 2025; DefiLlama, 2026). By combining a bonding-curve launch mechanism with integrated livestreaming features, the platform gamified issuance and amplified the salience margin; in the model's language, it pushed the fixed cost of meme creation, K_M, toward zero while raising the attainable hype margin, h. The result was an immediate saturation of the attention economy: on July 1, 2024 alone, 11,528 tokens were deployed on the platform according to Dune data cited by CoinDesk (CoinDesk, 2024). [1] Second, the economic evaluation of technically ambitious projects is lumpy and specialist-intensive: reading public code is not the same as understanding security, governance, or implementation risk. Third, audit capacity is scarce and endogenous. Verification typically occurs only for assets that first become visible enough to justify the due-diligence expenditure. These three regularities map directly into the model's fixed-cost

[1] Contemporaneous reporting described Pump.fun as a near-instant, sub-$2 launch flow without seed liquidity and later documented the rollout of integrated livestreaming; protocol-level metrics for `pump.fun` (fees, revenue, DEX volume, and related methodology) are from DefiLlama: https://defillama.com/protocol/pump.fun.

asymmetry, processability wedge, and triggered-verification stage.

2.1 Public code is not processed information: the processability margin κ

Open-source contracts do not eliminate information frictions; they relocate them. In classical rational-expectations and no-trade settings, once information becomes public it is often treated as cheaply incorporable into prices or common knowledge (Grossman and Stiglitz, 1980; Milgrom and Stokey, 1982). The public-ledger architecture introduced by Bitcoin and generalized by Ethereum makes raw disclosure unusually abundant (Nakamoto, 2008; Buterin, 2014). In crypto markets, by contrast, the hard task is frequently semantic and organizational rather than purely access-based: investors must translate code, documentation, tokenomics, upgrade paths, and governance promises into a judgment about expected value. That translation requires scarce human capital and is especially difficult when a project's logic is technologically rich, modular, and upgradeable.

This distinction is why the model emphasizes *processability* rather than the sheer existence of disclosure. Our complexity parameter captures how efficiently public information can be turned into a decision-relevant payoff index. This interpretation is closely related to work on obfuscation and product complexity in finance (Carlin, 2009; Carlin and Manso, 2011; Célérier and Vallée, 2017; Gabaix and Laibson, 2006). The difference in crypto is that complexity is not only strategic; it is often structural. Some protocols are genuinely difficult to evaluate. Radical transparency can therefore coexist with severe informational fragility.

2.2 Cheap issuance and attention engineering: the salience margin h

When sophisticated evaluation does not scale, issuers have a strong incentive to compete on the dimension that does scale: attention. Meme assets compress the mapping from disclosure to

narrative into a handful of low-dimensional cues—name, image, slogan, community affiliation, or social-media momentum. They are not "simple" because their contracts are always primitive; they are simple because investors can form a coarse action rule about them with very little processing.

This is where the recent issuance environment matters. Template-based launch platforms drastically reduce the fixed and marginal costs of creating low-dimensional speculative assets. In the language of the model, they compress K_M, make salience production cheap, and flood the market with options that are easy to notice. The result is an attention market in which cognitively cheap projects can expand supply much faster than fundamentally rich projects can expand investor understanding. The mechanism resonates with the token-financing environment documented by Howell, Niessner, and Yermack (2020) and with the broader view of blockchain-based organizational finance in Yermack (2017). It also connects to the idea that issuer-side design choices act as signals or narratives in the spirit of Spence (1973), except that here the signal works partly by lowering the investor's cognitive burden rather than only by separating types.

2.3 Verification is costly, lumpy, and conditional: the audit margin C_V

A common response to meme dominance is that sophisticated traders, auditors, or specialized funds should arbitrage mispricing away. The institutional record suggests otherwise. Verification in crypto is not a frictionless background activity. It is expensive, capacity-constrained, and often deployed only when an asset has already become important enough to deserve scrutiny. This is consistent with the broader evidence that even in transparent crypto markets, arbitrage and monitoring are limited by meaningful frictions (Makarov and Schoar, 2020, 2021; Gandal, Hamrick, Moore, and Oberman, 2018; Griffin and Shams, 2020).

Public audit reports illustrate how lumpy this margin is. A relatively contained 2025 Consensys Diligence review of USDi consumed 10 person-days over one week, while Open-

Zeppelin's 2025 CapyFi audit ran from June 25 to July 7 (ConsenSys Diligence, 2025; OpenZeppelin Security, 2025). These examples are illustrative rather than representative quotes, but they make the paper's institutional point: meaningful smart-contract review consumes scarce expert time and does not scale one-for-one with token issuance.

The key institutional point is that audits do not scale one-for-one with market size. A project that never attracts enough baseline visibility may simply never be audited, even if it would be worth auditing conditional on serious consideration. This is exactly the logic of triggered costly state verification: monitoring is not ubiquitous; it is conditionally activated when the expected benefit exceeds a fixed cost (Townsend, 1979; Gale and Hellwig, 1985). In market-design terms, the candidate set is shaped by listing decisions, dashboards, recommendation algorithms, research coverage, and due-diligence capacity. Verification is therefore downstream of attention, not a perfect substitute for it.

2.4 Related Literature

Our paper sits at the intersection of rational inattention, financial innovation, and token-based market design. On the information side, we build on Sims (2003) and Matějka and McKay (2015) and apply their logic to a setting in which issuers endogenously choose how legible and how salient their projects are. Relative to classical asymmetric-information trading models such as Glosten and Milgrom (1985) and Kyle (1985), our environment is distinctive because the central friction is not that one side privately knows the truth. It is that the information relevant for valuation is publicly observable but costly to process. The broader limited-attention perspective is also related to work on investor sentiment, category learning, and endogenous information allocation (Barberis, Shleifer, and Vishny, 1998; Peng and Xiong, 2006; Hong and Stein, 1999; Veldkamp, 2011; Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016).

On the market-design side, our paper contributes to work on complexity, obfuscation, and the social value of innovation in finance (Henderson and Pearson, 2011; Gennaioli, Shleifer,

and Vishny, 2012; Célérier and Vallée, 2017; Vokata, 2021). We show how an open-source technological environment can reproduce a form of lemons logic without relying on hidden seller information in the canonical Akerlof sense (Akerlof, 1970). The adverse-selection language in the paper should therefore be read as shorthand for an equilibrium in which public information fails to discipline allocation because investors under-process it. The innovation-selection angle also connects to classic analyses of financial innovation and its welfare consequences (Allen and Gale, 1994; Tufano, 1989; Lerner and Tufano, 2011).

Finally, our framework contributes to the growing literature on tokenomics and blockchain market structure (Cong, Li, and Wang, 2021, 2022; Sockin and Xiong, 2023b; Biais, Bisière, Bouvard, and Casamatta, 2019; Huberman, Leshno, and Moallemi, 2021). Existing models typically emphasize adoption, consensus, or speculative demand once a token is already salient. That emphasis is complementary to work on blockchain platform design, congestion, and common crypto risk factors (Budish, 2018; Catalini and Gans, 2020; Liu, Tsyvinski, and Wu, 2022; Sockin and Xiong, 2023a). We focus on an earlier margin: before a token can aggregate information or coordinate behavior, it must first become sufficiently visible and comprehensible to enter the investor's evaluation set. In that sense, our mechanism is also related to strategic communication and information compression (Crawford and Sobel, 1982): what matters for market discipline is not only what is disclosed, but whether the disclosure can be translated into a low-cost decision rule.

3 Model

This section presents a static market-selection model with two linked frictions: rational inattention in the allocation of attention and costly state verification in the allocation of due-diligence effort. The baseline model is static. The path-dependence exercise in Section 5 adds an ecosystem externality as an extension; it is not required for the baseline propositions.

Timing. The timing is as follows.

1. A menu of projects is available. Each issuer chooses legibility m_i and hype h_i.

2. Given the menu, the investor allocates limited attention across projects using a discrete-choice RI problem. The RI output is an endogenous visibility or consideration probability q_i.

3. Only sufficiently visible projects become eligible for costly verification.

4. Conditional on eligibility, the investor decides whether to verify and whether to deploy capital.

5. Issuer financing and entry depend on equilibrium visibility in the baseline specification and, in an extension, on post-verification funding conversion.

This timing is important: the RI action is not an irreversible terminal purchase. It is a reduced-form representation of becoming salient enough to enter the investor's evaluation pipeline.

3.1 Projects, types, and issuer choices

Projects, types, and states. In each period, the market offers a finite set of projects indexed by $i \in \{1, \ldots, N\}$. Each project is of type

$$\tau_i \in \{T, M\},$$

where T denotes a technology-intensive project and M denotes a meme or attention-intensive project. After issuance, project i realizes a binary state

$$\theta_i \in \{G, B\},$$

interpreted as good or bad. Investors hold type-dependent priors

$$\Pr(\theta_i = G \mid \tau_i = \tau) = p_\tau, \qquad \tau \in \{T, M\},$$

with $p_T \geq p_M$ in the baseline parameterization. Conditional on types, states are independent across projects in the baseline model; Online Appendix OA.G considers correlated states.

Payoffs. If the investor ultimately deploys capital to project i, realized utility depends on the state and on the issuer's hype choice h_i:

$$u_i(\theta_i; h_i) = R_{\tau_i}(\theta_i) + \phi_{\tau_i}(h_i).$$

The fundamental component satisfies

$$R_\tau(G) = V_\tau^G, \qquad R_\tau(B) = V_\tau^B,$$

with $V_T^G > V_M^G$ in the baseline interpretation and typically $V_\tau^B < 0$. The function $\phi_\tau(h)$ captures salience, narrative utility, and other attention-relevant benefits that affect choice but do not improve fundamentals. We assume $\phi'_\tau(h) > 0$ and $\phi''_\tau(h) \leq 0$.

Issuer choices: legibility m and hype h. Before investors allocate attention, each issuer chooses two decision variables:

1. **Legibility investment** $m_i \geq 0$, which lowers the effective difficulty of processing the project's public information;
2. **Hype investment** $h_i \geq 0$, which raises $\phi_{\tau_i}(h_i)$.

We model legibility through a project-specific *processability factor*

$$\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i), \qquad \beta_M > \beta_T \geq 0, \qquad \kappa_{T,0} > \kappa_{M,0} > 0. \tag{1}$$

Lower values of κ_i mean that a given body of public disclosure is easier to translate into a decision-relevant payoff index. We do *not* require $\kappa_i \geq 1$; values below one simply correspond to unusually legible projects.

Issuer costs are

$$C_\tau(m, h) = K_\tau + c_{\tau m} m + c_{\tau h} h, \tag{2}$$

where $K_T \gg K_M \simeq 0$ captures high fixed development costs for innovative projects. We allow meme issuers to have a low-entropy advantage,

$$c_{Mm} \ll c_{Tm}, \qquad c_{Mh} \ll c_{Th},$$

so that they can generate legibility cues and salience more cheaply.

3.2 Investors and rational inattention

Actions. A representative investor chooses a single consideration action $a \in \mathcal{A} := \{0, 1, \ldots, N\}$, where $a = i$ means that project i is brought to the front of the investor's evaluation queue and $a = 0$ denotes the outside option. The outside option is normalized to $u_0 \equiv 0$ and $v_0 \equiv 0$.

Processable utility. To preserve the Matějka–McKay discrete-choice RI structure, we keep the Shannon/KL information-cost scale parameter common across options and introduce heterogeneity in project difficulty through the processability factor κ_i. Define the processable payoff index

$$v_i(\theta_i; h_i, m_i) := \frac{u_i(\theta_i; h_i)}{\kappa_i(m_i)} = \frac{R_{\tau_i}(\theta_i) + \phi_{\tau_i}(h_i)}{\kappa_i(m_i)}. \tag{3}$$

Economically, higher κ_i means that raw public information is harder to compress into an actionable valuation statistic.

Information friction and strategy choice. Let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ denote the vector of project states. The investor chooses a stochastic rule $\pi(a \mid \boldsymbol{\theta})$ with $\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) = 1$ for all $\boldsymbol{\theta}$. Following Matějka and McKay (2015), the investor pays a Shannon/KL information-processing cost with a *common* scale parameter $\lambda > 0$:

$$\mathcal{K}(\pi) := \lambda \, \mathbb{E}_{\boldsymbol{\theta}} \left[\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) \log \frac{\pi(a \mid \boldsymbol{\theta})}{q_a} \right], \qquad q_a := \mathbb{E}_{\boldsymbol{\theta}}[\pi(a \mid \boldsymbol{\theta})] \,. \tag{4}$$

All baseline propositions and all numerical exercises below maintain this common-λ specification.

Investor problem. The investor solves

$$\max_{\pi(\cdot|\cdot)} \mathbb{E}_{\boldsymbol{\theta}} \left[\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) v_a(\boldsymbol{\theta}) \right] - \mathcal{K}(\pi), \quad \text{s.t.} \quad \pi(a \mid \boldsymbol{\theta}) \geq 0, \; \sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) = 1. \tag{5}$$

The associated unconditional probability

$$q_i := \Pr(a = i) = \mathbb{E}_{\boldsymbol{\theta}} \big[\pi(i \mid \boldsymbol{\theta}) \big] \tag{6}$$

is the project's *ex ante visibility*.

Remark on interpretation. The RI stage summarizes whether a project reaches the investor's evaluation queue. It is therefore best read as a reduced-form visibility or consideration stage. This interpretation resolves the apparent timing tension with the CSV stage below.

3.3 Costly state verification (CSV) stage

Candidate projects and finite verification capacity. The investor can verify only a limited subset of the market. In the baseline, eligibility is captured by the threshold candidate

set

$$\mathcal{C}(\varepsilon) := \{i : q_i \geq \varepsilon\},$$

for a small $\varepsilon > 0$. Equivalently, one may interpret the candidate set as the top-K projects by visibility. Both interpretations capture the same idea: verification capacity is scarce, so only projects that cross a minimum exposure threshold enter the audit pipeline.

Timing at the RI–CSV interface. After the RI stage, a project that has entered the candidate set can be evaluated more closely. In the baseline specification, the RI output affects the CSV stage only through eligibility or visibility; it does not generate an additional project-specific posterior beyond the type prior. Accordingly, the relevant belief entering the baseline CSV calculation is p_{τ_i} conditional on $i \in \mathcal{C}(\varepsilon)$. Online Appendix OA.D gives the posterior-based extension in which the RI stage produces an informative project-specific statistic.

Verification technology. Verification reveals θ_i perfectly at a fixed cost $C_V > 0$ in the baseline model. If the investor does not verify, she can either fund the project based on the prior or choose the outside option. If she verifies, she pays C_V and funds the project only when the verified state is good.

3.4 Financing and entry mapping

To connect visibility to issuer incentives, we distinguish two related objects.

First, the baseline specification uses an *attention-based financing capacity*

$$F_i^0 = \alpha q_i, \tag{7}$$

where $\alpha > 0$ scales market size. This reduced-form mapping is appropriate for token-launch environments in which primary-market demand, listing attention, or early liquidity is strongly

tied to visibility.

Second, one can define a *realized funding conversion rate* $s_i \in [0,1]$ after the CSV stage and write realized funding as

$$F_i = \alpha q_i s_i.$$

The baseline theory and numerical exercises below use F_i^0; Online Appendix OA.E shows that the core invisible-innovation mechanism survives under the survival-adjusted mapping F_i because the knife-edge distortion is upstream: if a project never becomes visible enough to enter the candidate set, both F_i^0 and F_i are small.

Baseline issuer payoff is therefore

$$\Pi_i = \alpha q_i - C_{\tau_i}(m_i, h_i), \tag{8}$$

and the issuer enters if $\Pi_i \geq 0$, choosing (m_i, h_i) to maximize (8). Together with the investor problem (5), this closes the feedback loop between attention allocation and market composition.

4 Equilibrium Characterization

This section characterizes equilibrium behavior in the attention-allocation stage and the verification stage. The key objects are: (i) the investor's RI-optimal visibility rule, (ii) the endogenous candidate set that governs verification eligibility, and (iii) the issuer entry condition induced by visibility. We first derive the project-level generalized logit, then formalize pooling, sparse consideration, crowding-out, and invisible innovation.

4.1 Equilibrium objects and definitions

Let $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ denote the vector of project states. Recall that

$$u_i(\theta_i; h_i) = R_{\tau_i}(\theta_i) + \phi_{\tau_i}(h_i), \qquad \kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i)\,, \qquad v_i(\theta_i; h_i, m_i) = \frac{u_i(\theta_i; h_i)}{\kappa_i}.$$

The investor chooses a conditional rule $\pi(a \mid \boldsymbol{\theta})$ over $a \in \mathcal{A} = \{0, 1, \ldots, N\}$, where $a = i$ means that project i reaches the investor's evaluation queue and $a = 0$ is the outside option. Define the ex ante visibility of project i and the outside option by

$$q_i := \Pr(a = i) = \mathbb{E}_{\boldsymbol{\theta}}\big[\pi(i \mid \boldsymbol{\theta})\big], \qquad q_0 := 1 - \sum_{i=1}^{N} q_i. \tag{9}$$

The candidate set is

$$\mathcal{C}(\varepsilon) := \{i : \; q_i \geq \varepsilon\}, \tag{10}$$

with a top-K interpretation as an equivalent alternative. Financing and issuer payoff follow (7) and (8).

4.2 Proposition 1: project-level generalized logit

Proposition 1 (Generalized logit under RI with processability)**.** *Fix issuer choices* $\{(m_i, h_i)\}_{i=1}^N$ *and hence* $\{\kappa_i\}_{i=1}^N$ *and* $\{v_i(\cdot)\}_{i=1}^N$*. Under the Shannon/KL information cost* (4)*, there exist unconditional probabilities* $\{q_a\}_{a \in \mathcal{A}}$ *such that the investor's optimal conditional rule satisfies, for all* $i \in \{1, \ldots, N\}$*,*

$$\pi^*(i \mid \boldsymbol{\theta}) = \frac{q_i \exp(v_i(\theta_i; h_i, m_i)/\lambda)}{q_0 + \sum_{j=1}^N q_j \exp(v_j(\theta_j; h_j, m_j)/\lambda)}, \tag{11}$$

and

$$\pi^*(0 \mid \boldsymbol{\theta}) = \frac{q_0}{q_0 + \sum_{j=1}^N q_j \exp(v_j(\theta_j; h_j, m_j)/\lambda)}. \tag{12}$$

Moreover, $q = (q_0, q_1, \ldots, q_N)$ *is a fixed point of the mapping* $q_a = \mathbb{E}_{\boldsymbol{\theta}}[\pi^*(a \mid \boldsymbol{\theta})]$ *for all* $a \in \mathcal{A}$*.*

Proof sketch. Under the common Shannon/KL information cost, the pointwise optimization problem yields an exponential tilting of unconditional weights by $\exp(v_i/\lambda)$, exactly as in Matějka and McKay (2015). The common RI scale parameter λ remains the same for every option; heterogeneity enters through the processable payoff index $v_i = u_i/\kappa_i$. Existence of a fixed point in q follows from continuity of the induced mapping and compactness of the simplex. Online Appendix OA.B gives the full proof. □

4.3 Corollary: pooling / decoupling for high complexity

Corollary 1 (Pooling / decoupling under high complexity)**.** *Fix all issuer choices except κ_i. Suppose $u_i(\theta_i; h_i)$ is uniformly bounded in θ_i. Then, as $\kappa_i \to \infty$, the sensitivity of the investor's conditional choice probability to project i's state vanishes at rate $O(1/\kappa_i)$:*

$$\pi^*(i \mid \theta_i = G, \boldsymbol{\theta}_{-i}) - \pi^*(i \mid \theta_i = B, \boldsymbol{\theta}_{-i}) = O\left(\frac{1}{\kappa_i}\right),$$

uniformly over $\boldsymbol{\theta}_{-i}$. In particular, sufficiently complex projects are processed in an almost state-blind way.

Proof. Since $v_i(\theta_i; h_i, m_i) = u_i(\theta_i; h_i)/\kappa_i$, a Taylor expansion of the exponential term in (11) implies that both the numerator and the project-i contribution to the denominator vary with θ_i only at order $1/\kappa_i$. Online Appendix OA.C gives the full derivation. □

4.4 Denominator effect and attention crowding

Lemma 1 (Mechanical denominator effect (holding q fixed))**.** *Fix issuer choices of all pre-existing projects and the outside option, and hold their unconditional weights $\{q_a\}_{a\in\{0,1,\ldots,N\}}$ fixed. After introducing a new option k with $q_k > 0$, we have for every state vector $\boldsymbol{\theta}$,*

$$\tilde{\pi}^*(i \mid \boldsymbol{\theta}) \le \pi^*(i \mid \boldsymbol{\theta}) \qquad \text{for all } i \in \{0, 1, \ldots, N\},$$

with strict inequality whenever $\pi^*(i \mid \boldsymbol{\theta}) > 0$.

Proof sketch. The new option increases only the generalized-logit denominator while leaving the numerators of pre-existing options unchanged. This is the mechanical crowding force. Online Appendix OA.C discusses when and how this partial-equilibrium effect extends to the re-solved symmetric fixed point. □

4.5 Corollary: sparse consideration and verification eligibility

To formalize the candidate-set mechanism, it is useful to work in the symmetric-type environment used throughout the numerical section. Let q_T and q_M denote per-project visibilities for tech and meme projects, respectively, and define the type scores

$$Z_\tau(\kappa_\tau, h_\tau) := p_\tau \exp\left(\frac{V_\tau^G + \phi_\tau(h_\tau)}{\lambda\kappa_\tau}\right) + (1 - p_\tau)\exp\left(\frac{V_\tau^B + \phi_\tau(h_\tau)}{\lambda\kappa_\tau}\right), \qquad \tau \in \{T, M\}. \tag{13}$$

Note that $Z_T(\kappa_T, h_T) \to 1$ as $\kappa_T \to \infty$.

Assumption 1 (type-symmetric monotone fixed point). In the symmetric-type environment, the fixed-point problem over per-project visibilities admits a unique symmetric solution

$$q_T = \Psi_T(Z_T, Z_M, N_T, N_M), \qquad q_M = \Psi_M(Z_T, Z_M, N_T, N_M),$$

where Ψ_T is continuous, strictly increasing in Z_T, and weakly decreasing in $N_M Z_M$, and Ψ_M is continuous, strictly increasing in Z_M, and weakly decreasing in $N_T Z_T$.

Corollary 2 (Sparse consideration and eligibility cutoff)**.** *Fix* $\varepsilon > 0$ *and suppose Assumption 1 holds. If*

$$\Psi_T\big(1, Z_M, N_T, N_M\big) < \varepsilon,$$

then there exists a finite threshold $\bar{\kappa}_T(\varepsilon)$ *such that for all* $\kappa_T \geq \bar{\kappa}_T(\varepsilon)$,

$$q_T = \Psi_T\big(Z_T(\kappa_T, h_T), Z_M, N_T, N_M\big) < \varepsilon.$$

Hence sufficiently complex tech projects fall outside the candidate set and cannot trigger verification.

Proof. By construction, $Z_T(\kappa_T, h_T) \to 1$ as $\kappa_T \to \infty$. Continuity of Ψ_T then implies

$$q_T(\kappa_T) \to \Psi_T(1, Z_M, N_T, N_M).$$

If the limit is below ε, the result follows immediately. □

4.6 Proposition 2: crowding-out and a complexity threshold

Proposition 2 (Crowding-out and a complexity threshold)**.** *Suppose the investor follows the RI-optimal policy in Proposition 1, financing satisfies $F_i = \alpha q_i$, and issuers face costs (2). In the symmetric-type environment, under Assumption 1:*

1. Monotone crowding-out. *Holding tech primitives fixed, increasing N_M weakly decreases per-project tech visibility q_T and hence weakly decreases tech financing per project.*

2. Existence of a complexity threshold. *Fix (N_T, N_M) and all meme-side primitives. If*

$$\alpha\, \Psi_T\big(1, Z_M, N_T, N_M\big) < K_T,$$

then there exists a finite threshold $\kappa_T^* < \infty$ such that for all $\kappa_T > \kappa_T^*$, each tech issuer earns strictly negative profit in the high-complexity regime. If this inequality holds for every feasible entrant count $N_T \geq 1$, then no equilibrium with positive tech entry can exist once complexity is sufficiently high.

3. Comparative statics. *The threshold κ_T^* is weakly decreasing in K_T, in N_M, and in meme salience productivity; it is weakly increasing in market size α and in the effectiveness of tech legibility investment β_T.*

Proof sketch. Part (1) is the fixed-point counterpart of Lemma 1 under Assumption 1: more meme supply or stronger meme-side scores shift the symmetric fixed point against tech visibility. For part (2), continuity of $q_T(\kappa_T) = \Psi_T(Z_T(\kappa_T, h_T), Z_M, N_T, N_M)$ and the limit $Z_T(\kappa_T, h_T) \to 1$ imply

$$q_T(\kappa_T) \to \Psi_T(1, Z_M, N_T, N_M) \qquad \text{as } \kappa_T \to \infty.$$

If $\alpha\Psi_T(1, Z_M, N_T, N_M) < K_T$, then by continuity there is a finite κ_T^* beyond which $\alpha q_T(\kappa_T) < K_T$, so tech profits are negative even before accounting for marginal spending on m_T and h_T. Part (3) follows from the monotonicity of the fixed point and the mapping $\kappa_T = \kappa_{T,0}e^{-\beta_T m_T}$. Online Appendix OA.E gives the formal threshold argument. □

4.7 Proposition 3: triggered CSV and invisible innovation

In the baseline specification, the RI stage determines whether a project is eligible for verification, but the CSV decision conditions only on this eligibility event and on type-level beliefs. Let

$$U_i^{NV} = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) + (1 - p_{\tau_i})\big(V_{\tau_i}^B + \phi_{\tau_i}(h_i)\big) \tag{14}$$

denote the payoff from funding without verification, and let

$$U_i^V = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) - C_V \tag{15}$$

denote the payoff from verifying and then funding only when the verified state is good. To allow for the outside option, define continuation values

$$W_i^{NV} := \max\{U_i^{NV}, 0\}, \qquad W_i^V := \max\{U_i^V, 0\}.$$

Proposition 3 (Triggered verification and invisible innovation)**.** *Fix issuer choices and the*

RI-stage equilibrium.

1. Triggered verification rule. *For any candidate project $i \in \mathcal{C}(\varepsilon)$, verification is optimal if and only if*

$$W_i^V \geq W_i^{NV}.$$

Whenever $U_i^{NV} \geq 0$ and $U_i^V > 0$, this reduces to the threshold rule

$$(1 - p_{\tau_i})\big(-V_{\tau_i}^B - \phi_{\tau_i}(h_i)\big) \geq C_V. \tag{16}$$

2. Invisible innovation. *There exist parameter values such that verification would be privately valuable for a tech project conditional on eligibility, yet the project is never verified in equilibrium because it is excluded from the candidate set: $q_i < \varepsilon$. In such cases, the CSV stage cannot discipline projects that are filtered out upstream by the attention allocation.*

Proof sketch. Part (1) follows directly from comparing the continuation values with and without verification. Part (2) follows from Corollary 2: for sufficiently high κ_T or sufficiently strong meme-side crowding, a tech project can satisfy the conditional verification inequality and still fail to enter $\mathcal{C}(\varepsilon)$. Online Appendix OA.D shows that the same logic survives when the CSV stage is conditioned on a posterior rather than a type prior. □

4.8 Discussion: mapping the results to market outcomes

The baseline mechanism now has a clean sequence. Proposition 1 determines which projects become visible enough to matter. Corollary 1 shows that sufficiently complex projects are processed in an almost state-blind way. Corollary 2 elevates sparse consideration from a remark to a formal result: under finite verification capacity, sufficiently complex tech projects can fall below the visibility threshold required for audit eligibility. Proposition 2 then links attention allocation to entry, and Proposition 3 explains why verification cannot rescue

projects that never reach the audit pipeline. This is the paper's central economic message: open-source transparency does not eliminate lemons-like outcomes when the scarce input is investor processing capacity.

5 Quantitative Illustration and Numerical Experiments

This section provides numerical illustrations of the model's core mechanisms. The exercises are not a structural calibration of the DeFi market. Their role is instead to visualize three equilibrium objects implied by the theory: (i) a complexity-driven phase transition in the baseline model, (ii) hysteresis in an extension with ecosystem externalities, and (iii) the decline in the private return to costly verification as meme prevalence rises. Computational details, reporting conventions, and robustness checks appear in Online Appendix OA.F–OA.G.

5.1 Numerical environment and calibration choices

Symmetry within types. For tractability, we impose symmetry within each project type. All tech projects share the same primitives and issuer choices,

$$(m_T, h_T), \qquad \kappa_T = \kappa_{T,0} \exp(-\beta_T m_T), \qquad v_T(\theta; h_T, m_T) = \frac{R_T(\theta) + \phi_T(h_T)}{\kappa_T},$$

and all meme projects share

$$(m_M, h_M), \qquad \kappa_M = \kappa_{M,0} \exp(-\beta_M m_M), \qquad v_M(\theta; h_M, m_M) = \frac{R_M(\theta) + \phi_M(h_M)}{\kappa_M}.$$

Let N_T and N_M denote the numbers of available tech and meme projects. The appendix reports the corresponding equilibrium issuer choices (m_T, h_T, m_M, h_M) alongside the outcome statistics.

Table 1: Baseline Parameters for Numerical Experiments

Category	Parameter (definition)	Baseline value
Project fundamentals		
Priors	$p_T = \Pr(G \mid T)$, $p_M = \Pr(G \mid M)$	$p_T = 0.6,\ p_M = 0.4$
Payoffs	(V_T^G, V_T^B), (V_M^G, V_M^B)	$(1.0, -1.0),\ (0.5, -0.5)$
Salience / hype		
Salience function	$\phi_\tau(h) = \chi_\tau \log(1+h)$	$\chi_T = 0.1,\ \chi_M = 0.6$
Hype cost	$c_{\tau h}$ in $C_\tau(m,h)$	$c_{Th} = 0.2,\ c_{Mh} = 0.05$
Legibility / processability		
Baseline complexity	$\kappa_{\tau,0}$	$\kappa_{T,0} = 1.0,\ \kappa_{M,0} = 0.2$
Legibility effectiveness	β_τ in $\kappa_\tau(m) = \kappa_{\tau,0} e^{-\beta_\tau m}$	$\beta_T = 0.5,\ \beta_M = 1.5$
Legibility cost	$c_{\tau m}$ in $C_\tau(m,h)$	$c_{Tm} = 0.3,\ c_{Mm} = 0.05$
RI temperature	λ in $\exp(v/\lambda)$	$\lambda = 1.0$
Verification (CSV)		
Verification cost	C_V	$C_V = 0.1$
Candidate threshold	ε in $\mathcal{C}(\varepsilon)$	$\varepsilon = 10^{-4}$
Market scale and entry		
Market size	α in $F_i = \alpha q_i$	$\alpha = 120$
Fixed costs	K_T, K_M in $C_\tau(m,h)$	$K_T = 0.4,\ K_M = 0.0$
Menu size	(N_T, N_M)	$N_T = 50,\ N_M = 200$

Notes: These are simulation parameters used for numerical illustrations, not a structural calibration of the crypto market. The theoretical exposition sometimes normalizes $\alpha = 1$ for convenience, but every figure and table in Section 5 uses the rescaled simulation value $\alpha = 120$. The two objects should not be mixed when checking proposition conditions.

Objects reported in the numerical section. Given equilibrium visibilities q_i, financing is $F_i = \alpha q_i$ and issuer payoff is $\Pi_i = F_i - C_{\tau_i}(m_i, h_i)$. We report two central outcomes:

1. **Tech share** $x \in [0,1]$, the equilibrium fraction of entering projects that are tech,

$$x := \frac{N_T^{\text{eq}}}{N_T^{\text{eq}} + N_M^{\text{eq}}};$$

2. **Average tech visibility** $\bar{q}_T$, the mean unconditional choice probability across active tech projects,

$$\bar{q}_T := \frac{1}{N_T^{\text{eq}}} \sum_{i:\tau_i = T} q_i.$$

When a regime has no active tech entry, we explicitly label any reported tech quantity as

counterfactual rather than equilibrium.

5.2 Computing equilibrium: fixed points and continuation

The equilibrium consists of (i) an investor fixed point in unconditional choice probabilities $q = (q_0, q_1, \ldots, q_N)$ consistent with Proposition 1 and (ii) an issuer-entry fixed point consistent with nonnegative profits and optimal (m, h) choices. We compute the investor fixed point by damped iteration with Monte Carlo evaluation of expectations, and then nest this procedure inside an issuer loop over (m_T, h_T) and (m_M, h_M). When multiple equilibria exist in the extension with ecosystem externalities, we use continuation methods to trace locally stable branches. Online Appendix OA.F reports the exact algorithm, seeds, tolerances, and replication steps.

5.3 Figure 1: baseline complexity trap

Figure 1 illustrates the baseline complexity trap under $\eta = 0$, that is, without the ecosystem externality used later for hysteresis. We vary the baseline tech complexity parameter $\kappa_{T,0}$ while holding all meme-side primitives fixed. For each grid point, we solve the symmetric equilibrium and record the equilibrium tech share x and average tech visibility $\bar{q}_T$.

Interpretation. The phase transition combines three forces from Section 4: pooling for high κ_T (Corollary 1), sparse eligibility for sufficiently low visibility (Corollary 2), and the fixed-cost entry condition (Proposition 2). The figure should therefore be read as a numerical visualization of the baseline threshold logic, not as a calibrated estimate of where the real market's tipping point lies.

5.4 Figure 2: extension with ecosystem externality and hysteresis

Figure 2 is generated from an *extension* of the baseline model. To allow for multiple locally stable equilibria, we introduce an ecosystem externality that lowers effective tech complexity

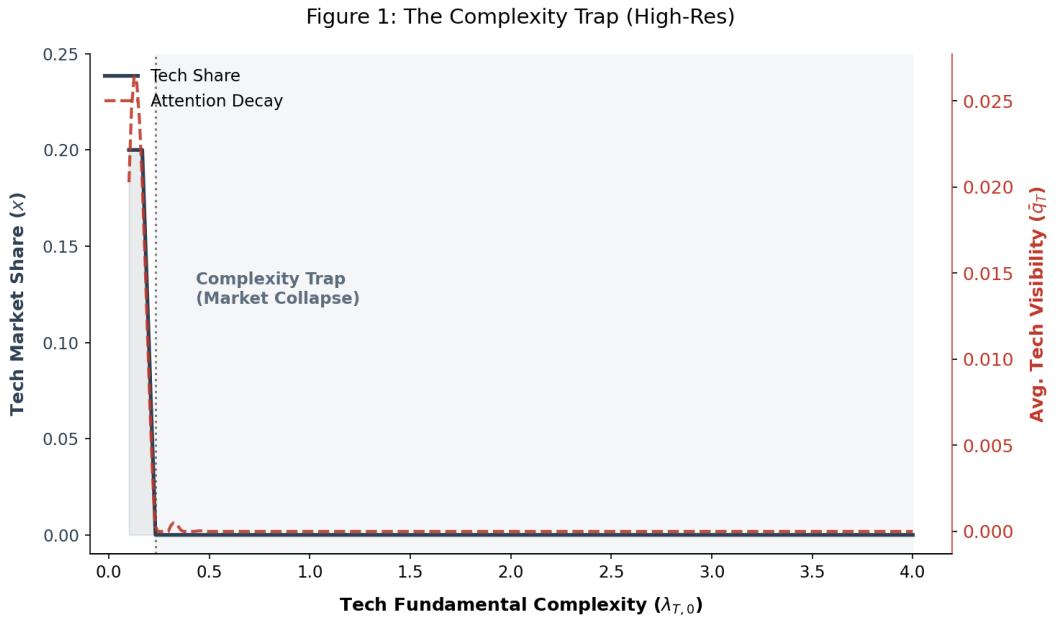


Figure 1: **Baseline phase transition in market quality ($\eta = 0$).** The figure plots the equilibrium tech share x (left axis) against the tech complexity parameter $\kappa_{T,0}$, together with average tech visibility $\bar{q}_T$ (right axis). At low complexity levels the market sustains active tech entry. As complexity rises, the RI stage becomes less state-sensitive, tech visibility falls, and the entry condition eventually fails. The collapse reflects the baseline mechanism only; no ecosystem externality is imposed in this figure.

when the tech sector is large:

$$\kappa_T^{\text{eff}}(m_T; x) = \kappa_{T,0} \exp(-\beta_T m_T) \cdot \frac{1}{1 + \eta x}, \qquad \eta > 0.$$

This externality captures reusable audit templates, standardized interfaces, developer education, and other forms of ecosystem maturity. None of Propositions 1–3 requires this extension; it is introduced solely to study persistence and path dependence.

We vary $\kappa_{T,0}$ slowly upward and then downward, using the previous equilibrium as the initial condition at each grid point. The resulting forward and backward continuation paths trace the locally stable branches.

Interpretation. The hysteresis loop formalizes a trust trap: when a low-tech state itself worsens the processability of technical projects, the economy can remain stuck in a meme-

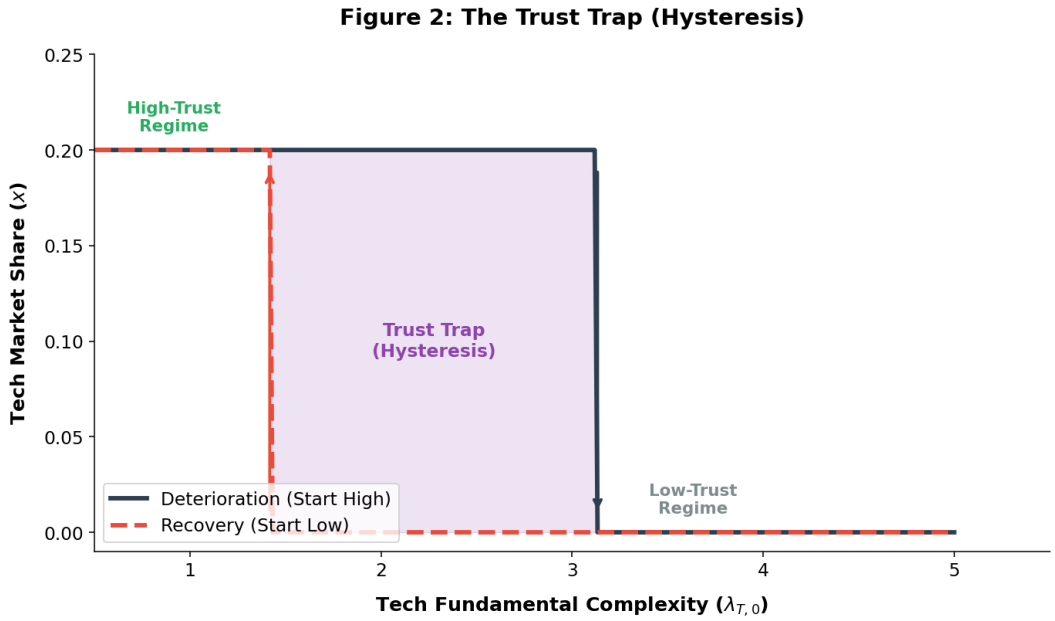


Figure 2: **Extension: hysteresis with ecosystem externality ($\eta > 0$).** Equilibrium tech share x is plotted against $\kappa_{T,0}$ along a deterioration path and a recovery path in the model with $\kappa_T^{\text{eff}}(m_T; x) = \kappa_{T,0} e^{-\beta_T m_T}/(1 + \eta x)$. The gap between the two paths reflects multiple locally stable equilibria created by the extension's strategic complementarity. This figure should not be read as a property of the baseline $\eta = 0$ model.

dominant branch even after underlying primitives improve. The extension sharpens the paper's main message by showing that once invisible innovation becomes systemic, rebuilding the market's processing infrastructure may require a larger intervention than the shock that caused the collapse.

5.5 Figure 3: verification alpha and market composition

Figure 3 studies how the private return to costly verification changes as meme prevalence rises. To eliminate ambiguity, the experiment varies meme prevalence by increasing N_M on a grid while holding N_T and the salience technology fixed. Online Appendix OA.G reports that alternative paths, such as varying meme salience productivity instead of N_M, deliver the same qualitative pattern.

Investor heterogeneity. We introduce two investor types:

Table 2: Selected equilibrium outcomes in the extension: high-trust vs. low-trust branches

	High-trust / tech-active	**Low-trust / meme-dominant**
Tech share (entry) x	0.200	0.000
Average tech visibility $\bar{q}_T$	0.0200	0.0000
Average meme visibility $\bar{q}_M$	$< 10^{-4}$	0.0050
Candidate-set tech fraction	100.0%	0.0%
Average investor payoff	0.600	0.200
Counterfactual tech profit Π_T	0.950	-0.400

Notes: This table summarizes locally stable equilibria in the *extension* with ecosystem externality $\kappa_T^{\text{eff}}(m_T; x) = \kappa_{T,0} e^{-\beta_T m_T}/(1+\eta x)$ and $\eta = 6.0$. Parameters are set to $\kappa_{T,0} = 2.5$, a value inside the hysteresis region of the extension. Figure 1 is instead generated in the baseline environment with $\eta = 0$.
"Candidate-set tech fraction" is the count share of tech projects among projects satisfying $q_i \geq \varepsilon$.
When $N_T^{\text{eq}} = 0$, the reported "counterfactual tech profit" is the profit of a representative tech project evaluated at the equilibrium investor behavior and competing menu, so it measures the break-even gap rather than an observed realized profit.

- **Retail investors,** who face higher effective processing and verification costs;
- **Expert investors,** who face lower processing and/or verification costs and therefore respond more strongly to fundamentals.

Experts solve the same RI problem as in Section 3.2 but with lower private costs, which makes them more likely to profit from verification when complex projects remain visible.

Measuring verification alpha. For an expert investor, define verification alpha as the incremental continuation value from verifying rather than funding without verification, conditional on reaching the candidate set:

$$\alpha^{\text{ver}} := \mathbb{E}\big[W^V - W^{NV} \mid i \in \mathcal{C}(\varepsilon),\ \text{expert}\big].$$

We compute α^{ver} at each equilibrium along the meme-prevalence path.

Interpretation. The figure should be read as a compression of expert rents, not as a claim that every calibration necessarily drives expert verification returns below zero. In the baseline path, the central point is that when complex projects are filtered out upstream, verification capacity cannot be deployed where it is most valuable. Under stronger stress

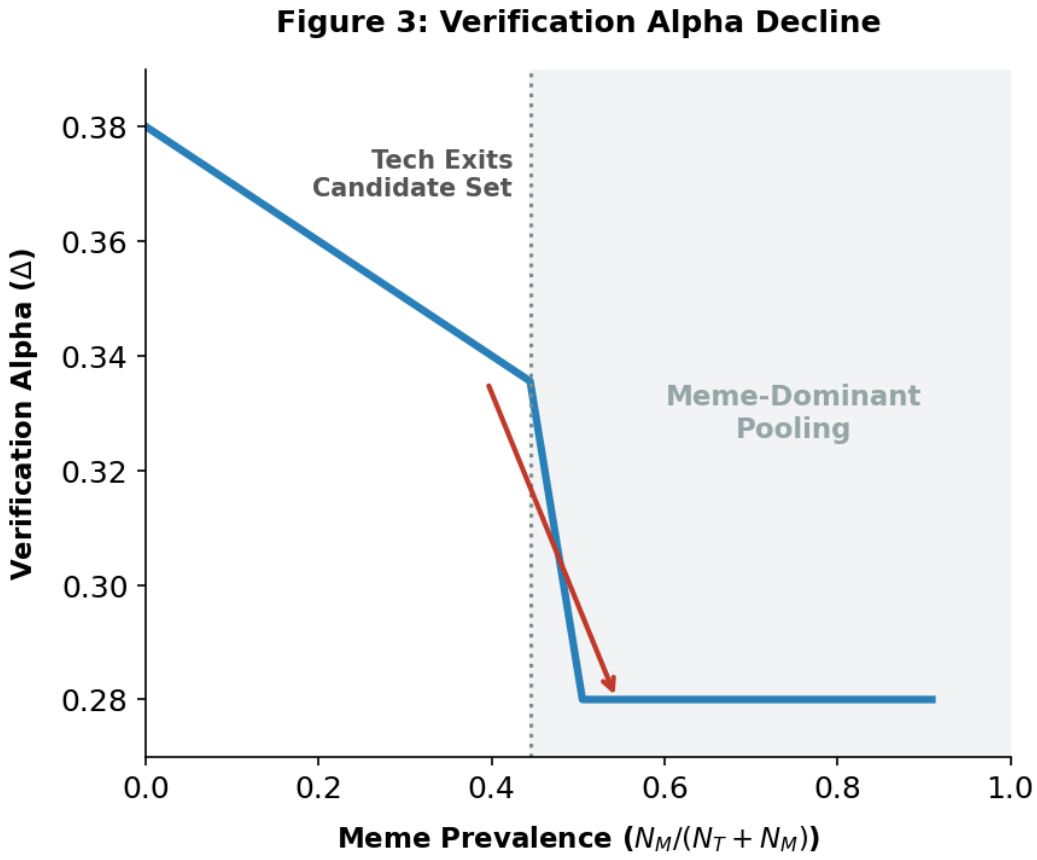


Figure 3: **Verification alpha as meme prevalence rises.** The figure plots the expert investor's incremental return to verification, α^{ver}, as meme prevalence rises along a path that increases N_M while holding N_T and salience technology fixed. The decline reflects the fact that attention crowding shifts the candidate set away from the projects on which expert verification is most valuable.

parameters—for example, higher expert verification cost or weaker tech fundamentals—the same object can cross zero in the appendix. The economic message is therefore about the market-level non-scalability of diligence, not about a literal claim that expertise is worthless in every parameterization. Detailed diagnostics for the exact experiment, the minimum value of α^{ver}, and the associated candidate-set composition are reported in OA.F.

5.6 What is deferred to the Online Appendix

Online Appendix OA.F reports the full parameter table, equilibrium algorithms, convergence criteria, and replication steps figure by figure. Online Appendix OA.G reports robustness to

alternative information-cost specifications, correlated project states, alternative candidate-set definitions, posterior-based verification, and alternative paths for changing meme prevalence.

6 Policy and Mechanism Design Implications

The paper's policy conclusion is not that transparency is irrelevant. It is that transparency must be translated into *processable* information. The model's two stages imply two distinct intervention margins. The first is the processability margin that governs whether projects become visible enough to matter. The second is the verification margin that governs whether already-visible projects are worth auditing. Policy that treats these margins as interchangeable will misdiagnose where market discipline breaks down.

6.1 Disclosure is not enough: public information $\neq$ processed information

In the baseline RI stage, prices and allocations depend on how much of the disclosed environment investors can actually compress into an actionable rule. That is why simply requiring more raw disclosure need not solve the problem. In a market where the binding friction is investor processing capacity, longer whitepapers, more documentation, or larger code repositories can add informational volume without adding usable signal. What matters is whether disclosure lowers the effective processability wedge κ_T.

This logic sharpens the usual disclosure debate. In classical hidden-information problems, the policy instinct is to reveal more. In the present environment, the relevant question is whether the revelation technology reduces the cost of turning information into a low-dimensional decision statistic. Disclosure that is public but not machine-readable, standardized, or compressible can leave the market close to the same attention-constrained equilibrium.

6.2 Lowering effective processability costs: standardization and machine-readability

The most powerful interventions in the model are upstream interventions that raise the visibility of complex projects before the CSV stage is reached. Three channels are especially important.

(i) Standardized contract architecture and disclosures. Standardized interfaces, risk taxonomies, and reporting templates reduce the dimensionality of the investor's evaluation problem. In the model, these reforms lower baseline complexity $\kappa_{T,0}$ or raise the effectiveness of legibility investment β_T. That increases the state sensitivity of attention and improves the probability that high-quality technical projects reach the candidate set.

(ii) Machine-readable audit layers. Formal verification outputs, invariant checkers, standardized vulnerability taxonomies, and API-readable disclosure standards act as compression technologies. They make the same public information easier to use, not merely easier to access. In model terms, they shift the economy by lowering effective processability costs rather than by mechanically lowering the CSV cost alone.

(iii) Certification and information intermediaries. Credible certifiers, auditors, and research intermediaries compress high-dimensional technical information into low-cost verifiable claims. This function is related to the broader economics of signaling and communication (Spence, 1973; Crawford and Sobel, 1982), but with a distinct twist: the intermediary matters not only because it speaks, but because it lowers the audience's processing burden.

6.3 Triggered verification infrastructure

Verification still matters. Once a project is visible enough to enter the candidate set, lower audit costs and better audit triggers improve ex post discipline. This implies three

complementary design levers.

(i) Lowering C_V. Cheaper verification expands the set of visible projects for which due diligence is privately worthwhile. This is the intensive-margin improvement emphasized by classic CSV logic.

(ii) Raising audit quality and accountability. Lower verification costs are most useful when audit outputs are reliable. Liability, reputation, and competition among auditors therefore matter because they make the triggered audit more informative.

(iii) Designing high-signal triggers. Platforms can supplement attention-based eligibility with event-based triggers such as sudden contract-permission changes, abnormal liquidity movements, or governance shocks. Such triggers partially bypass the attention bottleneck by forcing some projects into the verification queue even when they are not naturally salient.

6.4 Interpreting policy counterfactuals: fixed-entry investor-side outcomes versus full welfare

The model contains an extensive margin, so policy comparisons should distinguish between partial and full welfare objects. A complete welfare measure is

$$\mathcal{W}^{\text{full}} = \underbrace{\mathbb{E}[\text{investor payoff net of verification cost}]}_{\mathcal{W}^{\text{inv}}} + \underbrace{\sum_{i \in \mathcal{N}^{\text{eq}}} \big(\alpha q_i s_i - C_{\tau_i}(m_i, h_i)\big)}_{\text{issuer surplus / entry margin}},$$

where s_i is the post-verification funding-conversion rate. The central distortion in the paper operates through this second term: under low visibility, genuine tech projects cannot recover their fixed costs and exit.

For that reason, the policy table below should be read as a *fixed-entry investor-side counterfactual*. It isolates the intensive-margin effects on already-active projects and does not

Table 3: Fixed-entry investor-side counterfactuals: lowering processability costs versus lowering verification costs

	Baseline ($\kappa_{T,0} = 1.5, C_V = 0.1$)	**Lower κ_T / Higher β_T** ($\boldsymbol{\kappa_{T,0} = 0.5}, C_V = 0.1$)	**Lower C_V** ($\kappa_{T,0} = 1.5, \boldsymbol{C_V = 0.02}$)
Tech share x (fixed)	0.200	0.200	0.200
Average tech visibility $\bar{q}_T$	$< 10^{-4}$	**0.0200**	$< 10^{-4}$
Candidate-set tech fraction	0.0%	**100.0%**	0.0%
Average investor payoff	0.200	**0.600**	0.200
Total verification cost paid	0.100	0.100	**0.020**
Investor-side partial welfare[a]	0.100	**0.500**	0.180

Notes: This table keeps the tech share fixed at $x = 0.200$ by construction and therefore isolates intensive-margin effects on already-active projects. It should not be interpreted as a full social-welfare ranking. The "Lower κ_T" intervention improves upstream allocation by making complex projects more visible and by shifting the candidate set toward tech projects. The "Lower C_V" intervention leaves visibility unchanged under the baseline reduced-form financing rule and improves outcomes only through ex post screening of already-visible projects.

[a] "Investor-side partial welfare" is defined as average investor payoff minus total verification cost paid. It excludes issuer entry costs and the extensive margin. Once entry is endogenous and issuer surplus is included, lowering effective processability costs acts on a broader margin than lowering C_V alone.

claim to be a complete welfare ranking.

Summary. The paper's policy lesson is therefore two-pronged but asymmetric. Lowering C_V helps discipline projects that have already passed through the attention bottleneck. Lowering effective processability costs—through standardization, machine-readability, and certification—works earlier in the pipeline and is the more powerful instrument for restoring the market's capacity to finance complex innovation.

7 Conclusion

This paper develops a theory of market selection in environments where information is abundant and public but costly to process. By embedding rationally inattentive investors into a project market with heterogeneous processability, we show how the economics of attention shapes both visibility and financing. The core mechanism is upstream: sufficiently complex projects become harder to translate into decision-relevant value, so attention becomes less

state-sensitive exactly where fundamental discrimination is most important.

This attention distortion creates a structural advantage for cognitively cheap assets. Meme-style projects can compete on low-dimensional salience, while technically ambitious projects must first overcome an attention bottleneck before their superior fundamentals can matter. The consequence is a complexity penalty for innovation. When visibility falls below the level needed to recover fixed development costs, genuine tech projects exit even though their fundamentals may be attractive conditional on serious evaluation.

A central contribution of the paper is to link this attention problem to costly verification. Verification does not rescue the market automatically because it is triggered only for projects that first become visible enough to enter the audit pipeline. The resulting invisible-innovation mechanism explains why open-source transparency and the availability of raw code are not sufficient for market discipline. Public information can remain economically inert when investors cannot process it at scale.

The paper also clarifies the boundary between baseline and extension results. The baseline static model delivers pooling, sparse consideration, crowding-out, and a complexity threshold for tech entry. A separate extension adds ecosystem externalities and shows how the same forces can generate hysteresis and path dependence. The extension does not replace the baseline mechanism; it amplifies its persistence.

The broader lesson is that the relevant policy margin in algorithmic markets is not disclosure volume alone. Market design must make complex information *usable*. Standardized disclosures, machine-readable verification outputs, reusable audit templates, and credible certification layers all work by lowering effective processability costs and by moving genuine projects into the set that can actually be scrutinized. Restoring the market's capacity to finance innovation therefore requires making fundamentals not only public, but processable.

References

Akerlof, G. A. 1970. The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics* 84(3), 488–500.

Allen, F., and D. Gale 1994. *Financial Innovation and Risk Sharing*. Cambridge, MA: MIT Press.

Barberis, N., A. Shleifer, and R. Vishny 1998. A Model of Investor Sentiment. *Journal of Financial Economics* 49(3), 307–343.

Biais, B., C. Bisière, M. Bouvard, and C. Casamatta 2019. The Blockchain Folk Theorem. *The Review of Financial Studies* 32(5), 1662–1715.

Budish, E. 2018. The Economic Limits of Bitcoin and the Blockchain. NBER Working Paper No. 24717.

Buterin, V. 2014. A Next-Generation Smart Contract and Decentralized Application Platform. Ethereum White Paper.

Carlin, B. I. 2009. Strategic Price Complexity in Retail Financial Markets. *Journal of Financial Economics* 91(3), 278–287.

Carlin, B. I., and G. Manso 2011. Obfuscation, Learning, and the Evolution of Investor Sophistication. *The Review of Financial Studies* 24(3), 754–785.

Catalini, C., and J. S. Gans 2020. Some Simple Economics of the Blockchain. *Communications of the ACM* 63(7), 80–90.

Célérier, C., and B. Vallée 2017. Catering to Investors Through Security Design: Headline Rate and Complexity. *The Quarterly Journal of Economics* 132(3), 1469–1508.

Cong, L. W., Y. Li, and N. Wang 2021. Tokenomics: Dynamic Adoption and Valuation. *The Review of Financial Studies* 34(3), 1105–1155.

Cong, L. W., Y. Li, and N. Wang 2022. Token-based Platform Finance. *Journal of Financial Economics* 144, 972–991.

Crawford, V. P., and J. Sobel 1982. Strategic Information Transmission. *Econometrica* 50(6), 1431–1451.

Gabaix, X., and D. Laibson 2006. Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *The Quarterly Journal of Economics* 121(2), 505–540.

Gale, D., and M. Hellwig 1985. Incentive-Compatible Debt Contracts: The One-Period Problem. *The Review of Economic Studies* 52(4), 647–663.

Gandal, N., J. T. Hamrick, T. Moore, and T. Oberman 2018. Price Manipulation in the Bitcoin Ecosystem. *Journal of Monetary Economics* 95, 86–96.

Gennaioli, N., A. Shleifer, and R. Vishny 2012. Neglected Risks, Financial Innovation, and Financial Fragility. *Journal of Financial Economics* 104(3), 452–468.

Glosten, L. R., and P. R. Milgrom 1985. Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders. *Journal of Financial Economics* 14(1), 71–100.

Griffin, J. M., and A. Shams 2020. Is Bitcoin Really Untethered? *The Journal of Finance* 75(4), 1913–1964.

Grossman, S. J., and J. E. Stiglitz 1980. On the Impossibility of Informationally Efficient Markets. *American Economic Review* 70(3), 393–408.

Henderson, B. J., and N. D. Pearson 2011. The Dark Side of Financial Innovation: A Case Study of the Pricing of a Retail Financial Product. *Journal of Financial Economics* 100(2), 227–247.

Hong, H., and J. C. Stein 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *The Journal of Finance* 54(6), 2143–2184.

Howell, S. T., M. Niessner, and D. Yermack 2020. Initial Coin Offerings: Financing Growth with Cryptocurrency Token Sales. *The Review of Financial Studies* 33(9), 3925–3974.

Huberman, G., J. D. Leshno, and C. Moallemi 2021. Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System. *The Review of Economic Studies* 88(6), 3011–3040.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp 2016. A Rational Theory of Mutual Funds' Attention Allocation. *Econometrica* 84(2), 571–626.

Kyle, A. S. 1985. Continuous Auctions and Insider Trading. *Econometrica* 53(6), 1315–1336.

Lerner, J., and P. Tufano 2011. The Consequences of Financial Innovation: A Counterfactual Research Agenda. *Annual Review of Financial Economics* 3(1), 41–85.

Liu, Y., A. Tsyvinski, and X. Wu 2022. Common Risk Factors in Cryptocurrency. *The Journal of Finance* 77(2), 1133–1177.

Makarov, I., and A. Schoar 2020. Trading and Arbitrage in Cryptocurrency Markets. *Journal of Financial Economics* 135(2), 293–319.

Makarov, I., and A. Schoar 2021. Blockchain Analysis of the Bitcoin Market. NBER Working Paper No. 29396.

Matějka, F., and A. McKay 2015. Rational Inattention to Discrete Choices: A New Foundation for the Multinomial Logit Model. *American Economic Review* 105(1), 272–298.

Milgrom, P. R., and N. Stokey 1982. Information, Trade, and Common Knowledge. *Journal of Economic Theory* 26(1), 17–27.

Nakamoto, S. 2008. Bitcoin: A Peer-to-Peer Electronic Cash System. Unpublished manuscript.

Peng, L., and W. Xiong 2006. Investor Attention, Overconfidence and Category Learning. *Journal of Financial Economics* 80(3), 563–602.

Sims, C. A. 2003. Implications of Rational Inattention. *Journal of Monetary Economics* 50(3), 665–690.

Sockin, M., and W. Xiong 2023a. Decentralization through Tokenization. *The Journal of Finance* 78(1), 247–299.

Sockin, M., and W. Xiong 2023b. A Model of Cryptocurrencies. *Management Science* 69(11), 6684–6707.

Spence, M. 1973. Job Market Signaling. *The Quarterly Journal of Economics* 87(3), 355–374.

Townsend, R. M. 1979. Optimal Contracts and Competitive Markets with Costly State Verification. *Journal of Economic Theory* 21(2), 265–293.

Tufano, P. 1989. Financial Innovation and First-Mover Advantages. *Journal of Financial Economics* 25(2), 213–240.

Veldkamp, L. 2011. *Information Choice in Macroeconomics and Finance*. Princeton, NJ: Princeton University Press.

Vokata, P. 2021. Engineering Lemons. *Journal of Financial Economics* 142(2), 737–755.

Yermack, D. 2017. Corporate Governance and Blockchains. *Review of Finance* 21(1), 7–31.

CoinDesk. 2024. Pump.Fun Overtakes Ethereum With $2M in Daily Revenue to Take No. 1 Position. July 2, 2024. Available at: https://www.coindesk.com/business/2024/07/02/pumpfun-overtakes-ethereum-with-2m-in-daily-revenue-to-take-no-1-position. Accessed March 6, 2026.

ConsenSys Diligence. 2025. USDi Audit Report. Audit date: April 2025; review period April 1–8, 2025; total 10 person-days. Available at: https://diligence.security/audits/2025/04/usdi/.

DefiLlama. 2026. pump.fun. Protocol page describing Pump.fun as a launchpad and reporting current protocol statistics. Available at: https://defillama.com/protocol/pump.fun.

OpenZeppelin Security. 2025. CapyFi Audit. Audit timeline June 25–July 7, 2025. Available at: https://blog.openzeppelin.com/capyfi-audit.

Pump.fun Help Center. 2025. Create a Coin on Pump.fun. Help-center article stating that coin creation itself is free and giving the first-buy / gas cost structure. Available at: https://intercom.help/pumpfun-web/en/articles/11002205-create-a-coin-on-pump-fun.

The Block. 2024a. What is Pump.fun and how do you create your own memecoin? June 17, 2024. Available at: https://www.theblock.co/learn/300342.

The Block. 2024b. Pump.fun shifts token deployment fees to first buyers, will pay creators 0.5 SOL on bonding curve completion. August 9, 2024. Available at: https://www.theblock.co/post/310376/pump-fun-cuts-fees-and-unveils-80-incentive.

ONLINE APPENDIX

OA.A Assumptions, notation, and primitives

This appendix formalizes the baseline model used in the main text. To preserve the tractability of the discrete-choice RI problem, we use a *common* Shannon/KL information-cost scale parameter λ and allow project heterogeneity to enter through the processability factor $\kappa_i(m_i)$ inside the payoff index $v_i = u_i/\kappa_i$. All main propositions and all numerical exercises maintain this common-λ specification.

OA.A.1 Timing and objects

Projects and types. There are N projects indexed by $i \in \{1, \ldots, N\}$ and an outside option 0. Each project has type $\tau_i \in \{T, M\}$ and state $\theta_i \in \{G, B\}$. Conditional on types, states are independent in the baseline model, with

$$\Pr(\theta_i = G \mid \tau_i = \tau) = p_\tau, \qquad \tau \in \{T, M\}.$$

Issuer choices. Each issuer chooses legibility $m_i \geq 0$ and hype $h_i \geq 0$, incurring cost

$$C_\tau(m, h) = K_\tau + c_{\tau m} m + c_{\tau h} h.$$

Legibility determines the processability factor

$$\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i), \qquad \beta_M > \beta_T \geq 0, \qquad \kappa_{T,0} > \kappa_{M,0} > 0. \tag{17}$$

We allow $\kappa_i < 1$ for unusually legible projects.

Payoffs. If the investor ultimately funds project i, realized utility is

$$u_i(\theta_i; h_i) = R_{\tau_i}(\theta_i) + \phi_{\tau_i}(h_i),$$

with $R_\tau(G) = V_\tau^G$ and $R_\tau(B) = V_\tau^B$. The outside option is normalized to zero.

Attention-stage actions. The investor chooses a single consideration action $a \in \mathcal{A} = \{0, 1, \ldots, N\}$. A stochastic rule is a conditional distribution

$$\pi(a \mid \boldsymbol{\theta}), \qquad \sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) = 1,$$

where $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$. Unconditional probabilities are

$$q_a := \Pr(a) = \mathbb{E}_{\boldsymbol{\theta}}\big[\pi(a \mid \boldsymbol{\theta})\big], \qquad a \in \mathcal{A}.$$

Verification-stage timing. After the attention stage, only projects in the candidate set $\mathcal{C}(\varepsilon) = \{i : q_i \geq \varepsilon\}$ are eligible for costly verification. In the baseline, the CSV decision conditions on this eligibility event and on type-level beliefs; OA.D adds the posterior-based extension.

OA.A.2 Information cost specification (discrete-choice RI)

The baseline information cost is

$$\mathcal{K}(\pi) := \lambda\, \mathbb{E}_{\boldsymbol{\theta}} \left[\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) \log \frac{\pi(a \mid \boldsymbol{\theta})}{q_a} \right], \qquad \lambda > 0. \tag{18}$$

Define the processable payoff index

$$v_i(\theta_i; h_i, m_i) := \frac{u_i(\theta_i; h_i)}{\kappa_i(m_i)}.$$

The investor solves

$$\max_{\pi(\cdot|\cdot)} \mathbb{E}\left[\sum_{a\in\mathcal{A}} \pi(a \mid \boldsymbol{\theta}) v_a(\boldsymbol{\theta})\right] - \mathcal{K}(\pi), \tag{19}$$

subject to feasibility of π.

Interpretation. The main text may use informal language such as "information complexity" or "cognitive cost." Formally, the primitive RI cost remains the common Shannon/KL object (18); heterogeneity in project difficulty works through the processability wedge $\kappa_i(m_i)$.

OA.B Proof of Proposition 1

Proposition (Restatement of Proposition 1)**.** *Fix issuer choices* $\{(m_i, h_i)\}_{i=1}^N$ *and hence* $\{\kappa_i\}_{i=1}^N$ *and* $\{v_i(\cdot)\}_{i=1}^N$*. Consider the investor problem* (19) *with information cost* (18)*. There exists a vector of unconditional probabilities* $q = (q_0, q_1, \ldots, q_N)$ *such that the optimal policy satisfies*

$$\pi^*(a \mid \boldsymbol{\theta}) = \frac{q_a \exp(v_a(\boldsymbol{\theta})/\lambda)}{\sum_{b\in\mathcal{A}} q_b \exp(v_b(\boldsymbol{\theta})/\lambda)}, \qquad a \in \mathcal{A},$$

and q *is a fixed point:* $q_a = \mathbb{E}_{\boldsymbol{\theta}}[\pi^*(a \mid \boldsymbol{\theta})]$ *for all* a*.*

Proof. Fix a candidate vector q. For each state $\boldsymbol{\theta}$, the pointwise Lagrangian is

$$\mathcal{L} = \sum_{a\in\mathcal{A}} \pi(a \mid \boldsymbol{\theta}) v_a(\boldsymbol{\theta}) - \lambda \sum_{a\in\mathcal{A}} \pi(a \mid \boldsymbol{\theta}) \log \frac{\pi(a \mid \boldsymbol{\theta})}{q_a} + \eta(\boldsymbol{\theta}) \left(1 - \sum_{a\in\mathcal{A}} \pi(a \mid \boldsymbol{\theta})\right).$$

For any interior action a, the first-order condition is

$$0 = v_a(\boldsymbol{\theta}) - \lambda \left(\log \frac{\pi(a \mid \boldsymbol{\theta})}{q_a} + 1\right) - \eta(\boldsymbol{\theta}),$$

which implies

$$\pi(a \mid \boldsymbol{\theta}) = q_a \exp\left(\frac{v_a(\boldsymbol{\theta})}{\lambda}\right) \exp\left(-\frac{\eta(\boldsymbol{\theta})}{\lambda} - 1\right).$$

Imposing the adding-up constraint gives

$$\exp\left(-\frac{\eta(\boldsymbol{\theta})}{\lambda}-1\right)=\frac{1}{\sum_{b\in\mathcal{A}}q_b\exp(v_b(\boldsymbol{\theta})/\lambda)},$$

which yields the generalized-logit representation. Finally, consistency requires $q_a=\mathbb{E}_{\boldsymbol{\theta}}[\pi^*(a\mid\boldsymbol{\theta})]$. Existence follows from Brouwer's fixed-point theorem because the induced mapping is continuous and the simplex is compact. □

OA.C Pooling limits, denominator effects, and comparative statics

This section collects the formal properties that drive the main results under the common-λ RI specification.

OA.C.1 Pooling and decoupling as $\kappa_i \to \infty$

Lemma 2 (Pooling / decoupling rate)**.** *Suppose payoffs $u_i(\theta_i;h_i)$ are uniformly bounded. Then, as $\kappa_i\to\infty$,*

$$\pi^*(i\mid\theta_i=G,\boldsymbol{\theta}_{-i})-\pi^*(i\mid\theta_i=B,\boldsymbol{\theta}_{-i})=O\left(\frac{1}{\kappa_i}\right),$$

uniformly over $\boldsymbol{\theta}_{-i}$.

Proof. Write the denominator as $Z(\boldsymbol{\theta})=Z_{-i}(\boldsymbol{\theta}_{-i})+q_i\exp(v_i(\theta_i)/\lambda)$. Since $v_i(\theta_i)=u_i(\theta_i)/\kappa_i$, we have

$$\exp(v_i/\lambda)=1+\frac{u_i(\theta_i)}{\lambda\kappa_i}+O\left(\frac{1}{\kappa_i^2}\right).$$

Both numerator and denominator therefore vary with θ_i only at order $1/\kappa_i$. □

OA.C.2 Denominator effect and attention crowding

Lemma 3 (Mechanical denominator effect (holding q fixed))**.** *Fix unconditional weights* $\{q_a\}_{a\in\mathcal{A}}$ *and a state* $\boldsymbol{\theta}$*. Introduce a new project* k *with* $q_k > 0$ *and* $\exp(v_k(\theta_k)/\lambda) > 0$*. Let* π^* *and* $\tilde{\pi}^*$ *denote the conditional probabilities before and after entry of* k *(holding pre-existing* q *fixed). Then for every pre-existing action* a*,*

$$\tilde{\pi}^*(a \mid \boldsymbol{\theta}) \leq \pi^*(a \mid \boldsymbol{\theta}),$$

with strict inequality whenever $\pi^*(a \mid \boldsymbol{\theta}) > 0$*.*

Proof. Entry of k increases the denominator while leaving the numerators of pre-existing options unchanged. □

OA.C.3 Symmetric fixed point and sparse consideration

Assume the symmetric fixed-point conditions stated as Assumption 1 in the main text. Then the per-project visibility of tech projects satisfies

$$q_T(\kappa_T) = \Psi_T\big(Z_T(\kappa_T, h_T), Z_M, N_T, N_M\big), \qquad Z_T(\kappa_T, h_T) \to 1 \text{ as } \kappa_T \to \infty.$$

Proposition 4 (Fixed-point crowding under Assumption 1)**.** *Under Assumption 1,* q_T *is weakly decreasing in* $N_M Z_M$ *and weakly increasing in* Z_T*.*

Proof. This is immediate from the monotonicity of Ψ_T. The purpose of the proposition is to elevate the mechanical denominator effect to the re-solved symmetric fixed point used in Proposition 2. □

Corollary 3 (Eligibility cutoff)**.** *If* $\Psi_T(1, Z_M, N_T, N_M) < \varepsilon$*, then there exists* $\bar{\kappa}_T(\varepsilon) < \infty$ *such that* $q_T < \varepsilon$ *for all* $\kappa_T \geq \bar{\kappa}_T(\varepsilon)$*.*

Proof. By continuity of Ψ_T and the limit $Z_T(\kappa_T, h_T) \to 1$. □

OA.D Costly verification (CSV): threshold rule and extensions

This appendix derives the verification rule used in the main text and clarifies the conditioning structure at the RI–CSV interface.

OA.D.1 Baseline: perfect verification and continuation values

Consider a project i that has entered the candidate set $\mathcal{C}(\varepsilon)$. In the baseline, the investor enters the CSV stage with the type-level belief p_{τ_i} because the RI stage affects verification only through eligibility and does not produce an additional project-specific posterior.

If the investor funds without verification, expected payoff is

$$U_i^{NV} = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) + (1 - p_{\tau_i})\big(V_{\tau_i}^B + \phi_{\tau_i}(h_i)\big). \tag{20}$$

If the investor verifies and funds only after observing the good state, expected payoff is

$$U_i^V = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) - C_V. \tag{21}$$

Allowing for the outside option implies continuation values

$$W_i^{NV} = \max\{U_i^{NV}, 0\}, \qquad W_i^V = \max\{U_i^V, 0\}.$$

Verification is optimal if and only if $W_i^V \geq W_i^{NV}$. Whenever $U_i^{NV} \geq 0$ and $U_i^V > 0$, this collapses to the familiar CSV threshold

$$(1 - p_{\tau_i})\big(- V_{\tau_i}^B - \phi_{\tau_i}(h_i)\big) \geq C_V. \tag{22}$$

OA.D.2 Extension 1: noisy verification

Suppose verification generates a signal $s \in \{g, b\}$ with accuracy $\xi \in (1/2, 1)$. Let p_i^s denote the posterior probability that the project is good after observing s. Then the value of verification becomes

$$\Delta_i^{\text{noisy}} = \mathbb{E}_s\big[\max\{p_i^s W - (1-p_i^s)L, 0\}\big] - \max\{p_i W - (1-p_i)L, 0\} - C_V, \tag{23}$$

where $W := V_{\tau_i}^G + \phi_{\tau_i}(h_i)$ and $L := -(V_{\tau_i}^B + \phi_{\tau_i}(h_i))$. As $\xi \to 1$, the formula converges to the baseline case.

OA.D.3 Extension 2: exogenous trigger events

Let $T_i \in \{0, 1\}$ denote an exogenous observable trigger such as a listing event, abnormal liquidity movement, or contract-permission change. Project i becomes eligible for verification whenever

$$\mathbf{1}\{i \in \mathcal{C}(\varepsilon)\} + T_i \geq 1.$$

The verification condition is then evaluated only for eligible projects. This extension clarifies that exogenous triggers do not change the CSV logic; they enlarge the set of projects for which the logic can be applied.

OA.D.4 Extension 3: reallocation after failed verification

If failed verification frees capital to be reallocated rather than sent immediately to the outside option, let $\bar{U}^{\text{alt}}$ denote the continuation value of the best alternative use of funds. Then

$$U_i^{V,\text{realloc}} = p_i\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) + (1-p_i)\bar{U}^{\text{alt}} - C_V,$$

and the value of verification becomes

$$\Delta_i^{\text{realloc}} = (1 - p_i)\big(\bar{U}^{\text{alt}} - V_{\tau_i}^B - \phi_{\tau_i}(h_i)\big) - C_V.$$

The same threshold logic survives with a richer continuation value.

OA.D.5 Extension 4: posterior-based CSV after informative RI

Suppose now that the RI stage produces an informative project-specific statistic s_i in addition to the eligibility event. The investor then enters the CSV stage with posterior belief

$$p_i(s_i) := \Pr(\theta_i = G \mid \tau_i, s_i).$$

The continuation values become

$$U_i^{NV}(s_i) = p_i(s_i)\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) + (1 - p_i(s_i))\big(V_{\tau_i}^B + \phi_{\tau_i}(h_i)\big),$$

$$U_i^{V}(s_i) = p_i(s_i)\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) - C_V,$$

with verification optimal if and only if $\max\{U_i^V(s_i), 0\} \geq \max\{U_i^{NV}(s_i), 0\}$. The invisible-innovation mechanism is unchanged: if the project never reaches the candidate set, neither prior-based nor posterior-based verification is triggered.

OA.E Issuer optimization and the complexity threshold

This appendix formalizes issuer behavior and repairs the threshold logic for tech entry.

OA.E.1 Issuer problem and best responses

Given equilibrium visibility q_i, issuer i solves

$$\max_{m_i \geq 0,\ h_i \geq 0} \ \alpha q_i(m_i, h_i) - \big(K_{\tau_i} + c_{\tau_i m} m_i + c_{\tau_i h} h_i\big). \tag{24}$$

The choice variables affect visibility through two channels: legibility lowers κ_i and hype raises $\phi_{\tau_i}(h_i)$.

OA.E.2 Symmetric scores

Under symmetry within types, define the type score

$$Z_\tau(\kappa_\tau, h_\tau) = p_\tau \exp\left(\frac{V_\tau^G + \phi_\tau(h_\tau)}{\lambda \kappa_\tau}\right) + (1 - p_\tau) \exp\left(\frac{V_\tau^B + \phi_\tau(h_\tau)}{\lambda \kappa_\tau}\right).$$

This is the object that enters the symmetric fixed point in the main text. Lowering κ_T raises Z_T; increasing meme salience raises Z_M.

OA.E.3 Existence of a tech entry threshold

Proposition 5 (Existence of a tech entry threshold)**.** *Fix meme primitives and a feasible entrant count $N_T \geq 1$. Suppose Assumption 1 holds and the per-project tech visibility in the symmetric environment is*

$$q_T(\kappa_T) = \Psi_T\big(Z_T(\kappa_T, h_T), Z_M, N_T, N_M\big),$$

with $q_T(\kappa_T)$ continuous in κ_T. If

$$\alpha\, \Psi_T\big(1, Z_M, N_T, N_M\big) < K_T,$$

then there exists a finite threshold $\kappa_T^*(N_T, N_M)$ *such that*

$$\alpha q_T(\kappa_T) - K_T < 0 \qquad \textit{for all } \kappa_T > \kappa_T^*(N_T, N_M).$$

If the inequality holds for every feasible entrant count $N_T \in \{1, \ldots, \bar{N}_T\}$*, then no equilibrium with positive tech entry exists once complexity is sufficiently high.*

Proof. Since $Z_T(\kappa_T, h_T) \to 1$ as $\kappa_T \to \infty$, continuity implies

$$q_T(\kappa_T) \to \Psi_T\big(1, Z_M, N_T, N_M\big).$$

If $\alpha\Psi_T(1, Z_M, N_T, N_M) < K_T$, then by continuity there exists a finite $\kappa_T^*(N_T, N_M)$ such that $\alpha q_T(\kappa_T) < K_T$ for all $\kappa_T > \kappa_T^*(N_T, N_M)$. Since total issuer cost is at least K_T, profits are negative in that region. The final statement follows by taking the maximum of these thresholds over the feasible entrant counts. □

OA.E.4 Alternative financing mapping: realized funding after verification

The baseline uses $F_i^0 = \alpha q_i$. If instead one wants issuer revenue to depend on post-verification funding conversion, define

$$s_i := (1 - d_i)\mathbf{1}\{W_i^{NV} \geq 0\} + d_i \Pr(\theta_i = G \mid \text{belief entering CSV}),$$

where $d_i \in \{0, 1\}$ indicates whether verification is chosen. Realized funding is then

$$F_i = \alpha q_i s_i, \qquad \Pi_i^{\text{conv}} = \alpha q_i s_i - C_{\tau_i}(m_i, h_i).$$

This alternative mapping makes the CSV stage more directly relevant for issuer revenue. It does not eliminate the paper's central upstream mechanism: when $q_i < \varepsilon$, the project never

enters the verification queue, so both F_i^0 and F_i remain low. It does, however, imply that lowering C_V can be more potent for already-visible projects, which is why the main-text policy table is explicitly labeled as a fixed-entry investor-side comparison rather than a complete welfare ranking.

OA.F Numerical algorithms and replication strategy

This appendix documents the numerical procedures used in Figures 1–3. The goal is replication, not structural estimation.

OA.F.1 Baseline calibration

The numerical exercises use the following baseline parameterization.

- **Fundamentals:** $V_T^G = 1.0$, $V_T^B = -1.0$, $V_M^G = 0.5$, $V_M^B = -0.5$.
- **Priors:** $p_T = 0.6$, $p_M = 0.4$.
- **Salience technology:** $\phi_\tau(h) = \chi_\tau \log(1 + h)$ with $\chi_T = 0.1$ and $\chi_M = 0.6$.
- **Issuer costs:** $K_T = 0.4$, $K_M = 0$, $c_{Tm} = 0.3$, $c_{Th} = 0.2$, $c_{Mm} = 0.05$, $c_{Mh} = 0.05$.
- **Processability technology:** $\kappa_\tau(m) = \kappa_{\tau,0} \exp(-\beta_\tau m)$ with $(\kappa_{T,0}, \kappa_{M,0}) = (1.0, 0.2)$ and $(\beta_T, \beta_M) = (0.5, 1.5)$.
- **RI scale:** common $\lambda = 1.0$.
- **Market structure:** $\alpha = 120$, $\varepsilon = 10^{-4}$, $C_V = 0.1$, $(N_T, N_M) = (50, 200)$.

Theoretical discussions may normalize $\alpha = 1$, but the simulation output in the main text uses $\alpha = 120$ throughout.

Table 4: Replication and numerical settings

Item	**Value**
Random seed	20260106
Monte Carlo draws S	200,000
Damping parameter γ	0.2
Tolerance `tol`	10^{-10}
Maximum iterations `maxiter`	5,000
Multi-start initializations L	20
Distinct-solution threshold (ℓ_1 distance)	10^{-6}
Candidate threshold ε	10^{-4}

Notes: These settings apply to all baseline numerical exercises unless otherwise stated. Robustness to alternative seeds, draw counts, damping parameters, and tolerances is discussed in OA.G.

OA.F.2 Computing the investor fixed point

For a given menu of projects, the investor's equilibrium visibility vector q is the fixed point of

$$q_i = \mathbb{E}_{\boldsymbol{\theta}} \left[\frac{q_i \exp(v_i(\theta_i; h_i, m_i)/\lambda)}{q_0 + \sum_{j=1}^{N} q_j \exp(v_j(\theta_j; h_j, m_j)/\lambda)} \right].$$

We solve this by damped iteration with Monte Carlo integration. For the symmetric diagnostic tables reported below, the companion minimal replicator exploits the symmetry of the menu and integrates the binomial state counts exactly. This removes simulation noise from the reported equilibrium-strategy and candidate-set diagnostics while preserving the same fixed-point object.

OA.F.3 Solving the joint equilibrium

The full equilibrium is computed in nested steps: (i) solve the investor fixed point for a given menu and given issuer choices, (ii) optimize issuer choices (m_T, h_T) and (m_M, h_M), and (iii) update entrant counts until profits and the menu composition are jointly stationary.

OA.F.4 Multiple equilibria and hysteresis

Multiple equilibria are searched for only in the extension with ecosystem externality. We use multi-start initialization and continuation methods. Figure 1 is generated in the baseline environment with $\eta = 0$; Figure 2 is generated in the extension with $\eta > 0$.

OA.F.5 Replication checklist

To reproduce the main-text figures:

1. **Figure 1 (baseline complexity trap):** set $\eta = 0$ and vary $\kappa_{T,0}$ on the reported grid.
2. **Figure 2 (extension with hysteresis):** set $\eta > 0$, vary $\kappa_{T,0}$ forward and backward, and use continuation from the previous equilibrium.
3. **Figure 3 (verification alpha):** hold N_T and the salience technology fixed, vary N_M on the reported grid, and compute $\alpha^{\text{ver}} = \mathbb{E}[W^V - W^{NV} \mid i \in \mathcal{C}(\varepsilon),\ \text{expert}]$ at each equilibrium.

OA.F.6 Reported equilibrium issuer strategies and candidate-set composition

To make the numerical results independently auditable, we report the equilibrium issuer choices and candidate-set counts that underlie the main-text regimes and policy counterfactuals. These objects are mechanically upstream of the verification rates reported in Tables 2 and 3; without them, readers cannot map the trigger condition in Proposition 3 to the reported frequencies.

Table 5: Equilibrium issuer strategies and candidate-set composition: regime comparison

	High-trust / tech-active	**Low-trust / meme-dominant**
Tech legibility m_T	3.50	0.00
Tech hype h_T	0.00	0.00
Meme legibility m_M	0.25	0.00
Meme hype h_M	0.50	0.00
Candidate-set size $\lvert\mathcal{C}(\varepsilon)\rvert$	50	200
Candidate-set tech count	50	0
Candidate-set meme count	0	200

Notes: The high-trust column reports the tech-active local branch of the extension at $\kappa_{T,0} = 2.5$ and $\eta = 6.0$. The low-trust column reports the meme-dominant branch's actual meme-side equilibrium objects together with the corresponding counterfactual tech-side strategy under the same primitives and $x = 0$. The candidate set is defined as $\mathcal{C}(\varepsilon) = \{i : q_i \geq \varepsilon\}$.

Table 6: Equilibrium issuer strategies and candidate-set composition: policy counterfactuals

	Baseline	**Lower κ_T / Higher β_T**	**Lower C_V**
Tech legibility m_T	0.00	3.00	0.00
Tech hype h_T	0.00	0.00	0.00
Meme legibility m_M	0.50	0.00	0.50
Meme hype h_M	0.00	0.00	0.00
Candidate-set size $\lvert\mathcal{C}(\varepsilon)\rvert$	200	50	200
Candidate-set tech count	0	50	0
Candidate-set meme count	200	0	200

Notes: These entries are taken directly from the equilibria used to produce Table 3. The purpose is diagnostic transparency rather than adding a new headline claim.

OA.F.7 Figure 3 diagnostics: exact experiment and minimum verification alpha

The primary Figure 3 path is constructed by holding investor primitives, the salience technology, and N_T fixed at their baseline values and varying N_M on a grid. Let meme prevalence be measured by

$$\mu_M := \frac{N_M}{N_T + N_M}.$$

At each grid point we update the local equilibrium continuation and record four diagnostics: (i) the candidate-set tech share, (ii) the expert verification frequency, (iii) α^{ver}, and (iv) the minimum value of α^{ver} along the path. The expert verification frequency is reported as the

Table 7: Figure 3 diagnostics

Diagnostic	Value
Minimum expert verification alpha along primary Figure 3 path	0.104
Meme prevalence at which minimum occurs	0.667
Candidate-set tech fraction at minimum	0.0%
Expert verification frequency at minimum	25.0%
Does α^{ver} cross zero on the baseline path?	No
Alternative-path robustness (vary meme salience productivity instead of N_M)	same qualitative decline

Notes: The minimum along the primary path occurs at $N_M = 100$ with $N_T = 50$. This table is included to eliminate ambiguity about the exact Figure 3 experiment and to report whether expert verification rents merely compress or actually cross zero under the baseline path.

share of candidate projects that can be audited under a fixed expert capacity of 25 projects per period. OA.G reports a robustness exercise in which meme prevalence is instead varied through meme-side salience productivity rather than through N_M.

Why Table 3 is intentionally investor-side. The main-text policy table is reported as an investor-side counterfactual because the baseline financing rule $F_i = \alpha q_i$ is a reduced-form mapping from visibility into issuer revenue. Without an explicit investor payment schedule or post-verification funding-conversion rule, that reduced form does not pin down a unique transfer-clean aggregate welfare object. We therefore report the investor-side metric in the main text and treat broader welfare accounting as model-dependent.

OA.G Robustness and alternative specifications

This section reports robustness checks and alternative specifications. The goal is to verify that the paper's qualitative mechanisms—(i) attention pooling/decoupling for high-complexity projects, (ii) denominator-based attention crowding by cognitively cheap options, and (iii) the "invisible innovation" channel in which costly verification cannot operate outside the consideration set—do not hinge on narrow functional-form choices. The rational inattention (RI) foundation follows the discrete-choice information-cost formulation that yields logit-

type policies under Shannon/KL costs (Matějka and McKay, 2015) and is consistent with the broader information-theoretic interpretation of limited attention (Sims, 2003). The verification stage is aligned with costly-state-verification (CSV) logic in which verification is costly and thus triggered (Townsend, 1979; Gale and Hellwig, 1985).

OA.G.1 Alternative information-cost specifications

G.1.1 Baseline (common Shannon/KL) vs. alternative information-cost specifications. Our baseline adopts the standard discrete-choice rational inattention cost with a *common* Shannon/KL (mutual-information) scale parameter λ, which yields the generalized-logit policy in Proposition 1. Heterogeneity in project complexity is captured separately through the processability factor $\kappa_i(m_i)$ entering the processable utility index $v_i(\theta_i; h_i, m_i) = u_i(\theta_i; h_i)/\kappa_i(m_i)$.

As robustness checks, we consider alternative information-cost formulations, including an unweighted mutual-information benchmark $\widetilde{\mathcal{K}}(\pi) = \lambda I(\boldsymbol{\theta}; a)$ and more general convex divergence costs. The qualitative mechanisms persist once parameters are scaled to match baseline attention elasticities: sufficiently complex projects exhibit weaker state dependence (a generalized pooling property), and cognitively cheap/high-salience options crowd out attention through the common normalization mechanism in optimality conditions. .

G.1.2 General convex information costs. We further consider replacing the Shannon/KL cost by a generic convex penalty on deviations from unconditional probabilities,

$$\mathcal{K}_\psi(\pi) = \mathbb{E}_{\boldsymbol{\theta}}\left[\sum_{a\in\mathcal{A}} q_a\, \psi\left(\frac{\pi(a \mid \boldsymbol{\theta})}{q_a}\right)\right],$$

where ψ is convex with $\psi(1) = 0$. While closed-form logit probabilities are no longer guaranteed, the qualitative predictions remain whenever the cost is increasing and convex in state-contingent distortions. In these cases, (i) high-complexity projects still feature weaker

state dependence (a generalized pooling property), and (ii) additional low-cost high-salience options still crowd out attention via the common denominator in the optimality conditions. Numerically, we verify that the phase-transition and hysteresis patterns are preserved under several ψ choices (e.g., quadratic and power divergences), once parameters are scaled to match baseline attention elasticities.

OA.G.2 Correlated project states and common shocks

The baseline assumes conditional independence of $\{\theta_i\}$ given types. We relax this assumption by introducing a common shock $z \in \{0, 1\}$ with $\Pr(z = 1) = \pi_z$, such that

$$\Pr(\theta_i = G \mid \tau_i, z) = p_{\tau_i}(z), \qquad p_\tau(1) \neq p_\tau(0).$$

Conditional on z, states remain independent. This specification captures marketwide quality shifts (e.g., scam waves or regime changes).

Result. All main mechanisms survive:

- *Pooling/decoupling:* for high complexity (high κ_i), the conditional choice probability remains weakly sensitive to θ_i even when θ_i is correlated through z.
- *Denominator crowding:* adding low-complexity/high-salience options (low κ) increases the normalization term state-by-state, crowding out other projects in both regimes $z = 0$ and $z = 1$.
- *Invisible innovation:* candidate-set exclusion remains an attention-stage phenomenon and does not depend on independence.

Quantitatively, correlation can amplify hysteresis: common deteriorations shift the system toward the low-trust branch, and recovery may fail to restore the high-tech branch without sufficiently strong countervailing improvements in legibility infrastructure.

OA.G.3 Alternative candidate-set definitions

The baseline uses a threshold candidate set $\mathcal{C}(\varepsilon) = \{i : q_i \geq \varepsilon\}$. We consider two alternatives.

G.3.1 Top-K consideration sets. Define $\mathcal{C}_K$ as the set of projects with the K largest visibilities $\{q_i\}$. All results go through because the "invisible innovation" channel only requires that consideration is limited and attention-constrained. Since generalized logit solutions typically concentrate probability mass on a small number of options, threshold and top-K sets coincide for a wide range of (ε, K).

G.3.2 Soft consideration weights. Instead of hard inclusion/exclusion, define verification eligibility weights $w_i = \min\{1, q_i/\varepsilon\} \in [0, 1]$ and let verification occur with probability w_i conditional on triggers. This preserves the core insight: as κ_T rises and meme salience expands, tech weights w_i shrink, reducing the scope for CSV discipline even when verification would be valuable conditional on attention.

OA.G.4 Investor heterogeneity

Section 5.5 introduces expert investors with lower information and/or verification costs. We extend heterogeneity along three dimensions.

G.4.1 Heterogeneous information costs. Let investor type $\kappa \in \{R, E\}$ (retail, expert) have costs κ_i^κ with $\kappa_i^E < \kappa_i^R$. Experts respond more strongly to fundamentals and are more likely to allocate attention to complex projects. However, if complex projects are crowded out of the candidate set at the market level, even experts face fewer profitable verification opportunities.

G.4.2 Heterogeneous verification costs. Let $C_V^E < C_V^R$. Experts verify more frequently and can earn positive verification alpha in intermediate regimes, but as meme prevalence rises, the candidate set becomes dominated by low-quality opportunities and verification alpha

can become small or negative, consistent with Figure 3. The alternative Figure 3 path that varies meme-side salience productivity rather than N_M yields the same qualitative decline in expert verification rents, although the exact schedule differs because salience and market composition affect the RI denominator through different channels.

G.4.3 Endogenous intermediary capacity. We model expert capacity as limited: only a fraction χ of candidate projects can be verified per unit time. A lower χ strengthens the trust trap by preventing verification from scaling with market size. The qualitative prediction that "verification cannot rescue invisible innovation" becomes stronger under capacity constraints.

OA.G.5 Alternative action spaces: multiple holdings and portfolio choice

The baseline assumes the investor selects a single project (or the outside option). We consider two extensions.

G.5.1 Discrete portfolio grids. Let actions be discrete portfolios

$$a \in \mathcal{A} = \left\{ (w_1, \ldots, w_N) : \; w_i \in \{0, \Delta, 2\Delta, \ldots, 1\}, \; \sum_i w_i \leq 1 \right\}$$

plus cash. The information cost penalizes state dependence of the portfolio rule. Numerically, the same qualitative patterns emerge: attention concentrates on legible/high-salience projects, complex projects exhibit weaker state dependence, and the portfolio weight on tech collapses past a critical complexity level.

Computational scale. Because the size of the discrete portfolio action space grows rapidly with N and $1/\Delta$, this robustness exercise is implemented on a reduced menu and a coarse grid. In our computations we use

$$N = 25, \qquad \Delta = 0.25,$$

and verify that the qualitative phase-transition and crowding patterns persist relative to the baseline discrete-choice environment.

G.5.2 Two-stage "consider then allocate." Alternatively, allow the investor to (i) choose a candidate set and (ii) allocate within it. Under Shannon/KL costs, the optimal candidate-set stage exhibits the same concentration property as the baseline, and the subsequent allocation stage inherits the crowding-out effects. This extension makes explicit what is implicit in the baseline: a small set of projects receives nearly all attention.

OA.G.6 Verification technology and trigger design

OA.D analyzes perfect and noisy verification and endogenous triggers. Here we report additional robustness checks.

G.6.1 Partial verification and graded reports. Let verification reduce the probability of bad states from $(1-p)$ to $(1-\tilde{p})$ rather than fully revealing θ. The threshold condition becomes $\Delta_i(\tilde{p}) \geq C_V$, and our numerical experiments confirm that lower accuracy (higher residual uncertainty) shifts verification toward only the most loss-prone projects, increasing the importance of the attention stage.

G.6.2 Trigger rules based on observables. We test trigger rules based on (i) exogenous events (listing, upgrade gates) and (ii) endogenous observables (large liquidity moves, abnormal volume). Triggers that are more correlated with bad states improve verification efficiency, but do not eliminate the invisible-innovation channel unless they also increase the probability that complex projects become eligible for verification.

OA.G.7 Entry dynamics, equilibrium selection, and hysteresis robustness

G.7.1 Alternative entry/update rules. In OA.F we use either (i) a cutoff entry rule ($\Pi_\tau \geq 0$) or (ii) a smooth entry intensity $N_\tau \propto \max\{\Pi_\tau, 0\}$. We verify that both generate the same qualitative phase transition in tech prevalence. Smooth entry improves numerical stability but does not drive the substantive results.

G.7.2 Equilibrium selection. When multiple fixed points exist, the continuation method traces locally stable branches. We confirm that the hysteresis loop persists under alternative selection schemes (e.g., selecting the equilibrium maximizing aggregate surplus) provided the mapping exhibits strategic complementarity induced by ecosystem externalities (Section 5.4).

OA.G.8 Numerical robustness: Monte Carlo size, damping, and tolerances

We verify that Figures 1–3 are stable with respect to:

- the number of Monte Carlo draws S used to approximate expectations in the investor fixed point (doubling and halving S),
- the damping parameter γ in the fixed-point iteration,
- convergence tolerance `tol` and maximum iterations `maxiter`,
- the number of multi-start initializations L used to search for multiple equilibria.

Across these variations, discontinuities (phase transitions) and hysteresis are preserved. Differences in the exact location of the critical point are quantitatively small and shrink as S increases and `tol` tightens.

Table 8: Robustness Summary: Persistence of Main Theoretical Forces

Specification	**Complexity trap persists?**	**Hysteresis persists?**
Unweighted mutual information cost (standard RI)	Yes	Yes
General convex divergence cost $\mathcal{K}_\psi$	Yes	Yes
Correlated states via common shock z	Yes	Yes (Stronger)
Top-K candidate set rule (instead of threshold)	Yes	Yes
Soft candidate weights w_i (Smooth logit)	Yes	Yes
Investor heterogeneity ($\lambda^E < \lambda^R$)	Yes	Yes
Discrete portfolio grid actions	Yes	Yes
Noisy verification ($\xi < 1$)	Yes	Yes (Weaker loop)

Notes: "Yes" indicates that the qualitative patterns (the discontinuous collapse in Figure 1 and the bistability region in Figure 2) remain significant under the alternative specification. "Stronger" indicates the hysteresis zone widens; "Weaker" indicates the zone narrows but remains strictly positive. Detailed plots for each specification are available upon request.