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 decentralized finance: despite the open-source nature of smart contracts, high-quality innovation is systematically crowded out by 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, rational 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 investor cognitive capacity.

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
‡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

The decentralized finance (DeFi) and crypto-asset markets present a fundamental puzzle to financial economists. On one hand, these markets are characterized by *radical transparency*: every transaction is recorded on a public ledger, and the source code for asset-generating smart contracts is typically open-source. On the other hand, the market exhibits extreme forms of adverse selection reminiscent of opaque "lemon" markets. Stylized facts reveal a striking bifurcation: assets with negligible technological utility but high cultural salience ("memecoins") frequently attract massive liquidity and achieve multi-billion dollar valuations, while complex, technology-intensive projects often struggle to secure financing or sustain valuations despite their verifiable fundamental potential. Why does a market built on the premise of "trustless" transparency paradoxically collapse into a speculative equilibrium where low-quality assets crowd out genuine innovation? Standard theories of adverse selection, pioneered by Akerlof (1970), attribute market failure to *hidden information*—the seller knows the quality, but the buyer does not. This explanation, however, is incomplete in the crypto context. Here, information is not hidden; it is *public but costly to process*. Verifying the security and economic viability of a zero-knowledge rollup or a cross-chain bridge requires specialized cognitive capital that the average investor lacks. Conversely, verifying the "quality" of a memecoin requires only negligible attention. In this paper, we argue that the binding friction in algorithmic markets is not the availability of information, but the *cost of processing it*. We develop a dynamic model of market selection that integrates **Rational Inattention (RI)** with **Costly State Verification (CSV)**. We show that when technical complexity raises the cost of processing public information above a critical threshold, rational investors optimally choose ignorance. This creates a pooling equilibrium where "dumb money" flows

indiscriminately to cognitively cheap assets. The result is a "Transparency Paradox": full disclosure overwhelms investor attention, generating a structural advantage for simple, low-quality assets and forcing genuine innovators into a "complexity trap."

Theoretical Framework. Our model departs from the standard assumption that investors costlessly observe public signals. Instead, we microfound the investor's problem using the rational inattention framework of Sims (2003), where information processing is constrained by Shannon entropy. Specifically, we adopt the discrete-choice RI formulation of Matějka and McKay (2015), which allows us to model the "visibility" of each project as an endogenous choice probability. Crucially, we introduce heterogeneity in information costs: processing the fundamentals of a "Tech" project carries a high unit cost (high entropy), while processing a "Meme" project is cheap. Issuers strategically invest in *legibility* (lowering processing costs) or *hype* (increasing salience). We integrate this attention stage with a second verification stage rooted in the costly state verification literature (Townsend, 1979; Gale and Hellwig, 1985). A key novelty of our setup is that verification is *triggered* only for projects that first succeed in competing for attention. This creates a two-step filter: a project must first be "visible" (enter the consideration set via RI) before it can be "verified" (audited via CSV).

Key Mechanism: The Complexity Penalty and Invisible Innovation. The model generates a "complexity penalty" for high-tech projects. Because investors maximize utility net of information costs, they rationally under-weight complex signals. As the complexity of the technology stack increases, the investor's optimal strategy shifts from discerning analysis to a generalized "pooling," where high-quality and low-quality tech projects are treated identically. This leads to the phenomenon of **"Invisible Innovation"**: a high-quality tech project might be fundamentally sound and verifiable, but if its complexity cost prevents it from entering the investor's candidate set at the RI stage, it never triggers the CSV stage. Consequently, it is pooled with low-quality projects and fails to finance. Meanwhile, memecoins, by optimizing for low-entropy signaling (simple narratives, high salience), easily

bypass the attention bottleneck, crowding out technology through a denominator effect in the attention market.

Contribution. This paper makes three distinct contributions to the literature on financial innovation and market microstructure. First, **theoretically**, we provide a tractable mechanism linking *computational complexity* to *market structure*. While existing literature focuses on tokens as incentive devices (Cong, Li, and Wang, 2021) or speculation vehicles (Sockin and Xiong, 2023a), we isolate *informational complexity* as the driver of asset substitution. By modeling project-specific information costs within a generalized logit framework, we show how "open source" leads to opaque pricing when attention is scarce. Second, **institutionally**, we refine the application of CSV to decentralized markets. We demonstrate that the efficacy of verification mechanisms (audits, fraud proofs) is contingent on the antecedent allocation of attention. In our model, verification is a conditional probability nested within an attention choice. This explains why the mere existence of on-chain data does not deter fraud: if investors are rationally inattentive, they never trigger the "audit" function, rendering the verification technology economically dormant. Third, **quantitatively**, we identify non-linear phase transitions in market quality. Our numerical analysis reveals a **"Complexity Trap"**: a tipping point in technical complexity beyond which the equilibrium share of tech projects collapses discontinuously to near zero. Furthermore, we demonstrate **"Trust Hysteresis"**: once the market flips to a meme-dominant equilibrium, merely improving transparency or reducing complexity is insufficient to restore trust immediately, as investors' priors have anchored on a low-quality regime.

Policy Implications. Our findings challenge the regulatory orthodoxy of "mandated disclosure." In an environment governed by rational inattention, expanding the volume of disclosure (e.g., requiring longer whitepapers or code dumps) may exacerbate information asymmetry by further exhausting investor cognitive capacity. Instead, our model suggests that policy should target the *reduction of effective processing costs*. This points toward

distinct mechanism design levers: (i) **Standardization** of smart contract interfaces to lower entropy costs; (ii) **Machine-readable verification** to automate the CSV stage; and (iii) **Credible certification intermediaries** that act as "attention compressors." We argue that without these infrastructure improvements, open-source finance will remain structurally biased against complex innovation.

Roadmap. The remainder of the paper is organized as follows. Section 2 outlines the institutional background and stylized facts of the crypto market. Section 3 presents the model, integrating rational inattention with costly state verification. Section 4 characterizes the equilibrium and derives the theoretical propositions. Section 5 presents the numerical results, illustrating the complexity trap and hysteresis dynamics. Section 6 discusses policy implications and mechanism design. Section 7 concludes.

2 Institutional Background and Stylized Facts

The crypto-asset market presents a puzzling contradiction to classical financial theory. While built on a technological foundation of "radical transparency"—where every line of code and every transaction is public—it has paradoxically evolved into a market saturated with low-quality assets ("memecoins") and rampant fraud. A striking illustration of this dynamic is the rise of "fair launch" platforms like *pump.fun* on the Solana network. Historically, issuing a cryptocurrency required at least minimal technical competency in smart contract programming and capital for initial liquidity. *Pump.fun* eliminated these barriers, reducing the monetary cost of issuance to approximately $2 and the time cost to seconds. By gamifying the issuance process with "bonding curves" and integrated livestreaming features, it effectively reduced the fixed cost of meme creation (K_M) to zero while maximizing their salience potential (h). The result was an immediate saturation of the attention economy: at its peak in mid-2024,

the platform facilitated the launch of over 10,000 new tokens *daily*. [1] This industrial-scale production of "noise" creates a denominator effect that drowns out fundamental innovation, perfectly exemplifying how technological efficiency, absent verification costs, can lead to market degradation. This section resolves this paradox by mapping the market's institutional features to our model's three key frictions: the cognitive cost of processing complexity (λ), the strategic engineering of attention (h), and the conditional nature of costly verification (C_V).

2.1 The Transparency Paradox: Public Code as Cognitive Overload (λ)

In standard asset pricing models, public information is typically assumed to be immediately incorporated into prices (Grossman and Stiglitz, 1980; Milgrom and Stokey, 1982). The DeFi market challenges this assumption. Although the source code for protocols is open (Nakamoto, 2008; Buterin, 2014), the *semantic* understanding of this code requires specialized human capital that the vast majority of investors lack.

We define the parameter λ not merely as a lack of information, but as the *cognitive entropy* required to map public signals to fundamental value.

- **Weaponized Complexity.** In traditional finance, Carlin (2009) and Carlin and Manso (2011) argue that issuers strategically increase complexity to obfuscate prices. In crypto, this friction is structural. A legitimate innovation (e.g., a zero-knowledge rollup) involves cryptographic primitives that are inherently high-entropy (high λ). Conversely, a Ponzi scheme can wrap itself in similarly complex code. For a retail investor, distinguishing between the two is prohibitively expensive.

- **The Failure of Disclosure.** Regulatory orthodoxy suggests that mandated disclosure solves asymmetry. However, Gennaioli, Shleifer, and Vishny (2012) and Célérier and

[1] Protocol-level metrics for `pump.fun` (fees, revenue, DEX volume, and related methodology) are from DefiLlama: https://defillama.com/protocol/pump.fun.

Vallée (2017) show that in the presence of neglected risks, financial innovation often leads to fragility rather than efficiency. We extend this to show that when disclosure is technically dense, it functions like a Denial-of-Service (DoS) attack on investor attention. The "open source" nature of DeFi (Catalini and Gans, 2020) creates an illusion of transparency that masks deep structural risks, deterring the entry of genuine "smart money" that cannot scale its due diligence efforts.

2.2 The Economy of Memes: Attention as the Path of Least Resistance (h)

Given the prohibitive cost of processing technical fundamentals, the market equilibrium shifts toward assets that minimize cognitive load. We model this as "Attention Engineering" (h), where issuers invest in narrative salience to bypass the λ-friction.

- **Low-Entropy Signaling.** Memecoins succeed not because they possess utility, but because they are cognitively cheap. Understanding a "dog coin" requires zero technical processing ($\lambda_M \approx 0$). This aligns with Peng and Xiong (2006), who show that investors process information in categories; memecoins optimize for the "high-growth/high-volatility" category without the burden of fundamental verification.

- **Rational Speculation vs. Utility.** While Lerner and Tufano (2011) argue that financial innovation should address unmet needs, the crypto market is dominated by assets that serve purely as vehicles for sentiment. Barberis, Shleifer, and Vishny (1998) and Hong and Stein (1999) provide behavioral foundations for how sentiment drives prices. We show that in a high-λ environment, chasing sentiment (investing in high h) is not irrational; it is the optimal second-best strategy for attention-constrained agents.

- **Crowding Out.** Empirical evidence from Liu, Tsyvinski, and Wu (2022) indicates that common risk factors in crypto are driven by momentum and attention proxies. As Sockin and Xiong (2023a) note, tokens serve as membership utility, but when speculation

decouples from utility, the market devolves. Our model explains why this decoupling persists: high-h assets crowd out high-utility projects in the attention economy, leaving a market of "lemons" that are famous but fundamentally worthless.

2.3 Triggered Verification: Why Audits Don't Scale (C_V)

A common counter-argument is that auditors or sophisticated arbitrageurs should police the market. However, we argue that verification is a *triggered* and *costly* activity (C_V), not a continuous background process.

- **The Limits of Arbitrage.** Makarov and Schoar (2020, 2021) document persistent arbitrage spreads and fragmentation in crypto markets, suggesting that even sophisticated capital faces significant friction. We interpret this as evidence of high C_V: verifying the safety of a smart contract to execute an arbitrage trade is costly and risky.

- **Verification Lag.** The market history is replete with protocols that were "verified" only after they collapsed. Gandal, Hamrick, Moore, and Oberman (2018) and Griffin and Shams (2020) show that price manipulation and fraud can persist for long periods in open markets. This supports our CSV formulation (Townsend, 1979; Gale and Hellwig, 1985): verification only occurs when the expected loss exceeds the cost. If a complex project (high λ) never attracts initial attention, it never triggers the audit (C_V), allowing "invisible lemons" to hide in plain sight until they implode.

2.4 Related Literature

Our work connects three distinct strands of literature: rational inattention, financial innovation, and tokenomics.

Rational Inattention in Finance. We build on the foundational work of Sims (2003) and the discrete-choice framework of Matějka and McKay (2015), which model information

processing as a constraint on entropy reduction. While Veldkamp (2011) and Kacperczyk, Van Nieuwerburgh, and Veldkamp (2016) apply attention allocation to mutual funds and macroeconomics, we introduce a new application: the *endogenous design* of asset complexity. In our model, issuers manipulate the information cost λ to exploit investors' rational inattention, a mechanism that complements the "shrouded attributes" framework of Gabaix and Laibson (2006).

Financial Innovation and Complexity. Our findings contribute to the debate on the social value of financial innovation. While Allen and Gale (1994) and Tufano (1989) highlight the risk-sharing benefits of innovation, a growing body of work suggests innovation is often used to extract rents. Henderson and Pearson (2011) and Vokata (2021) find that complex structured products offer negative abnormal returns. We extend this logic to algorithmic markets, showing that "open source" code can serve the same obfuscating role as the complex prospectuses in Célérier and Vallée (2017), effectively imposing a "complexity tax" on legitimate innovation.

Tokenomics and Decentralized Design. Finally, we contribute to the emerging literature on tokenomics. Existing models often focus on tokens as incentive devices for platform adoption (Cong, Li, and Wang, 2021, 2022) or as means to aggregate information (Sockin and Xiong, 2023b). Biais, Bisière, Bouvard, and Casamatta (2019) and Huberman, Leshno, and Moallemi (2021) analyze the game-theoretic stability of consensus protocols. Our paper diverges by focusing on the *pre-investment* attention stage. We argue that before a token can function as an incentive or consensus mechanism, it must first be *noticed* and *understood*. By integrating Budish (2018)'s skepticism regarding the economic limits of blockchains with an attention-based selection mechanism, we explain why the market equilibrium favors simple speculation over complex utility.

3 Model

This section develops a model of a project market characterized by two key frictions: rational inattention (RI) and costly state verification (CSV). In this market, information about projects is in principle publicly available (e.g., code, documentation, marketing content), yet investors face limited information-processing capacity. We model investors as rationally inattentive decision makers in a discrete-choice environment. Specifically, the RI component adopts the discrete-choice mutual information penalty framework of Matějka and McKay (2015), which is rooted in the concept of finite information-processing capacity of Sims (2003) (modeled via mutual information or entropy costs). Furthermore, the verification stage follows the costly state verification approach of Townsend (1979) and the "triggered verification" logic of Gale and Hellwig (1985): verification does not occur ubiquitously, but is triggered only under specific conditions. In our framework, issuers strategically choose (i) *legibility* investments that reduce the effective unit cost of processing their information and (ii) *hype* investments that raise the salience component of utility. After an initial attention-allocation stage, investors may engage in costly verification (audit/due diligence) for a small subset of "visible" projects.

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/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. For expositional clarity, we assume conditional independence across projects: $\{\theta_i\}_{i=1}^N$ are independent given $\{\tau_i\}_{i=1}^N$. (Online Appendix OA.G considers correlated states.)

Payoffs. If the investor invests in project i, her *ex post* utility depends on the realized 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

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

allowing for $V_T^G > V_M^G$ (innovative projects have higher upside) and typically $V_\tau^B < 0$ (bad projects yield negative net value). The function $\phi_\tau(h)$ captures salience, narrative utility, and other attention-related benefits that affect choice but do not improve fundamentals. We assume ϕ_τ is increasing and concave: $\phi'_\tau(h) > 0$ and $\phi''_\tau(h) \leq 0$.

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

1. **Legibility (or simplification) investment** $m_i \geq 0$: lowers the effective unit cost of processing project-i information.
2. **Hype (or marketing) investment** $h_i \geq 0$: raises $\phi_{\tau_i}(h_i)$.

We model legibility as affecting a project-specific *processability (complexity) factor* $\kappa_i \geq 1$,

which scales how effectively investors can translate publicly available information into decision-relevant utility. Specifically,

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

The inequality $\beta_M > \beta_T$ captures a *low-entropy advantage*: meme projects can be made legible via low-dimensional cues, whereas compressing complex technical properties into digestible summaries is more difficult.

Issuer costs are given by

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

where $K_T \gg K_M \simeq 0$ captures high fixed development costs for innovative projects, and

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

captures the idea that meme issuers can generate legible signals and hype at substantially lower marginal cost.

3.2 Investors and rational inattention

Actions. A representative investor chooses a single action $a \in \mathcal{A} := \{0, 1, \ldots, N\}$, where $a = i$ means investing in project i and $a = 0$ denotes the outside option. We normalize $u_0 \equiv 0$.

Processable utility. While raw information (code, docs) is publicly available, projects differ in how costly it is for investors to *process* that information into decision-relevant utility. We capture this with the processability factor κ_i in (1) and define the investor's *processable*

utility index:

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

Intuitively, higher complexity (larger κ_i) attenuates the effective impact of fundamentals on choice because investors cannot fully translate raw disclosure into actionable valuation.

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

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

This is the standard Shannon/KL (mutual-information) penalty that yields tractable generalized logit policies in Matějka and McKay (2015).

Investor objective. The investor chooses π to maximize expected processable utility net of information cost:

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

Remark (endogenous candidate sets). Even though the action set includes all projects, the optimal rule may assign negligible unconditional probability to most projects, generating an endogenous consideration set.

3.3 Costly state verification (CSV) stage

Visibility and candidate projects. The rational inattention stage yields each project's *ex ante visibility*, defined as its unconditional choice probability:

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

We define the candidate set either via a small threshold $\varepsilon > 0$,

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

or via a top-K rule (selecting the K projects with the largest q_i). The baseline analysis uses the threshold formulation; the top-K variant yields the same qualitative implications.

Remark 1 (Economic interpretation of the candidate-set threshold ε)**.** *The threshold $\varepsilon > 0$ is a reduced-form way to capture a* finite attention capacity. *Interpreting $q_i = \Pr(a = i)$ as an ex-ante exposure/visibility rate, the set $C(\varepsilon) = \{i : q_i \geq \varepsilon\}$ collects projects that cross a minimal attention "coverage" level and hence can be practically audited. Equivalently, ε can be viewed as an* exposure gate *induced by platforms (recommendation feeds, listings, or social amplification): only projects that achieve sufficiently high baseline visibility enter investors' verifiable pipeline. Our qualitative results do not depend on the hard-threshold specification: replacing $C(\varepsilon)$ with a top-K rule or with smooth eligibility weights yields the same mechanisms (see Online Appendix OA.G).*

Verification technology and cost. After attention allocation, the investor may pay a fixed cost $C_V > 0$ to verify the state of a candidate project $i \in \mathcal{C}(\varepsilon)$. Verification reveals θ_i perfectly in the baseline model; Online Appendix OA.D considers noisy verification. Verification is not undertaken for projects outside the candidate set, capturing the idea that projects which are never attended to cannot be verified in practice.

Triggered verification. Because verification is costly, it is optimal to verify only when the expected benefit of avoiding losses in bad states is sufficiently large. The resulting verification policy takes a threshold form (derived formally in Section 4 and OA.D), consistent with the classic costly-state-verification logic that audits occur only under certain triggers rather than in all states.

3.4 Financing and entry mapping

To connect visibility to issuer incentives and market composition, we assume that project i's financing capacity (or initial valuation/liquidity) is proportional to its visibility:

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

where $\alpha > 0$ scales market size. Issuer i's expected payoff is

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

and the issuer enters if $\Pi_i \geq 0$, choosing (m_i, h_i) to maximize (8). This mapping, together with the investor problem (5), closes the feedback loop between attention allocation and the endogenous composition of projects. The equilibrium characterization is provided in the next section.

4 Equilibrium Characterization

This section characterizes equilibrium behavior in the attention-allocation (RI) stage and the verification (CSV) stage. The key objects are (i) the investor's optimal choice rule under rational inattention, which delivers a project-level generalized logit structure in the spirit of Matějka and McKay (2015) (building on the information-theoretic formulation of Sims (2003)), and (ii) a triggered verification policy consistent with the costly-state-verification

logic of Townsend (1979) and Gale and Hellwig (1985). We then derive comparative statics highlighting (a) *pooling/decoupling* for high-complexity projects, (b) a *denominator effect* through which cognitively cheap, high-salience projects crowd out complex projects, and (c) a *complexity threshold* above which technology-intensive entry collapses.

4.1 Equilibrium objects and definitions

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

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

be defined as in Section 3. Define the processable utility index

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

The investor chooses a conditional choice rule $\pi(a \mid \boldsymbol{\theta})$ over $a \in \mathcal{A} = \{0, 1, \ldots, N\}$, where $a = i$ denotes investing in project i and $a = 0$ is the outside option (normalized to $u_0 \equiv 0$). Define each project's *ex ante visibility* (unconditional choice probability)

$$q_i := \Pr(a = i) = \mathbb{E}_{\boldsymbol{\theta}}\big[\pi(i \mid \boldsymbol{\theta})\big], \qquad q_0 := \Pr(a = 0) = 1 - \sum_{i=1}^{N} q_i. \tag{9}$$

Following Section 3.3, the *candidate set* is

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

for a small $\varepsilon > 0$ (a top-K candidate set yields analogous results). Finally, we use the financing/valuation mapping from Section 3.4:

$$F_i = \alpha q_i, \qquad \Pi_i = F_i - C_{\tau_i}(m_i, h_i), \tag{11}$$

and issuers enter if $\Pi_i \geq 0$ and choose (m_i, h_i) to maximize Π_i.

4.2 Proposition 1: Project-level generalized logit

We start from the rational inattention problem in Section 3.2, where the investor selects $\pi(\cdot \mid \boldsymbol{\theta})$ to maximize expected *processable* utility net of a Shannon/KL (mutual-information) cost with a *common* scale parameter λ. As in Matějka and McKay (2015), this information-theoretic penalty implies a generalized logit (softmax) structure. Project heterogeneity in processing difficulty is captured separately by the processability factor $\kappa_i(m_i)$ entering $v_i(\theta_i; h_i, m_i) = u_i(\theta_i; h_i)/\kappa_i(m_i)$, rather than by option-specific "temperatures."

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 choice rule satisfies, for all* $i \in \{1, \ldots, N\}$*,*

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

and

$$\pi^*(0 \mid \boldsymbol{\theta}) = \frac{q_0}{q_0 + \sum_{j=1}^N q_j \exp(v_j(\theta_j; h_j, m_j)/\lambda)}. \tag{13}$$

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 (4), the pointwise maximization of expected processable utility minus the entropy penalty implies an exponential tilting of the unconditional weights by $\exp(v_i/\lambda)$, yielding the generalized logit form (12)–(13). Existence of a fixed point in q follows from continuity of the mapping and compactness of the simplex (Brouwer). A complete proof is provided in Online Appendix OA.B. □

4.3 Corollary: Pooling / decoupling for high complexity

The generalized logit structure in (12) implies that when a project's unit information cost is large (high complexity), the investor's action becomes almost independent of that project's true state. This delivers a pooling/decoupling property central to our mechanism.

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 $O(1/\kappa_i)$:*

$$\pi^*(i \mid \theta_i = G, \boldsymbol{\theta}_{-i}) - \pi^*(i \mid \theta_i = B, \boldsymbol{\theta}_{-i}) = O\left(\frac{1}{\kappa_i}\right), \tag{14}$$

uniformly over $\boldsymbol{\theta}_{-i}$. In particular, $\pi^*(i \mid \boldsymbol{\theta}) \approx q_i$ uniformly over θ_i, so choices become approximately decoupled from project i's quality.

Proof. Fix $\boldsymbol{\theta}_{-i}$ and compare $\theta_i = G$ vs. $\theta_i = B$. Write the logit denominator as the sum of an i-independent part plus the i-term:

$$Z(\boldsymbol{\theta}) = \underbrace{q_0 + \sum_{j \neq i} q_j \exp\left(\frac{v_j(\theta_j; h_j, m_j)}{\lambda}\right)}_{=: Z_{-i}(\boldsymbol{\theta}_{-i})} + q_i \exp\left(\frac{v_i(\theta_i; h_i, m_i)}{\lambda}\right).$$

Only the last term depends on θ_i. Since $v_i(\theta_i; h_i, m_i) = u_i(\theta_i; h_i)/\kappa_i$ and u_i is uniformly bounded, a Taylor expansion yields

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

Therefore both (i) the numerator $q_i \exp(v_i/\lambda)$ and (ii) the i-dependent component of the denominator vary with θ_i only at order $1/\kappa_i$. Substituting the expansion into

$$\pi^*(i \mid \boldsymbol{\theta}) = \frac{q_i \exp\left(\frac{v_i(\theta_i; h_i, m_i)}{\lambda}\right)}{Z_{-i}(\boldsymbol{\theta}_{-i}) + q_i \exp\left(\frac{v_i(\theta_i; h_i, m_i)}{\lambda}\right)}$$

implies

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

uniformly over $\boldsymbol{\theta}_{-i}$. Hence the conditional choice probability becomes insensitive to θ_i as $\kappa_i \to \infty$, delivering pooling/decoupling. □

4.4 Denominator effect and attention crowding

The generalized logit structure implies a mechanical denominator effect: introducing an additional cognitively cheap and salient option raises the normalization term and reduces the conditional attention share of pre-existing options, holding unconditional weights fixed.

Let the investor's optimal policy take the generalized logit form (Proposition 1):

$$\pi^*(i \mid \boldsymbol{\theta}) = \frac{q_i \exp(v_i(\theta_i; h_i, m_i)/\lambda)}{Z(\boldsymbol{\theta})}, \qquad Z(\boldsymbol{\theta}) := q_0 + \sum_{j=1}^{N} q_j \exp(v_j(\theta_j; h_j, m_j)/\lambda)\,. \tag{15}$$

where $v_i(\theta_i; h_i, m_i) := u_i(\theta_i; h_i)/\kappa_i(m_i)$ is the processable utility index.

Consider adding a new project k with unconditional weight $q_k > 0$ and exponential score $\exp\big(v_k(\theta_k; h_k, m_k)/\lambda\big) > 0$ for all $\theta_k \in \{G, B\}$. Let π^* denote the conditional choice probabilities before the addition and $\tilde{\pi}^*$ after the addition, *holding the unconditional weights of pre-existing options fixed.*

Lemma 1 (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 $\boldsymbol{\theta}$,*

$$\tilde{\pi}^*(i \mid \boldsymbol{\theta}) \le \pi^*(i \mid \boldsymbol{\theta}) \qquad \textit{for all } i \in \{0, 1, \ldots, N\}. \tag{16}$$

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

Proof sketch. By (15), the introduction of option k increases the denominator from $Z(\boldsymbol{\theta})$

to $\tilde{Z}(\boldsymbol{\theta}) = Z(\boldsymbol{\theta}) + q_k \exp(v_k(\theta_k)/\lambda)$, while leaving the numerators of pre-existing options unchanged (holding q fixed). Hence $\tilde{\pi}^*(i \mid \boldsymbol{\theta}) = \pi^*(i \mid \boldsymbol{\theta}) \cdot Z(\boldsymbol{\theta})/\tilde{Z}(\boldsymbol{\theta}) \leq \pi^*(i \mid \boldsymbol{\theta})$. □

Remark (re-solving for the fixed point). If one re-solves for the new fixed point in unconditional probabilities $\tilde{q}_a = \mathbb{E}_{\boldsymbol{\theta}}[\tilde{\pi}^*(a \mid \boldsymbol{\theta})]$, the pointwise inequality (16) need not hold without additional restrictions. However, the expected visibility of pre-existing options weakly decreases under standard regularity conditions because the new option absorbs positive probability mass.

4.5 Proposition 2: Crowding-out and a complexity threshold

We now connect attention allocation to issuer entry via the financing mapping (11). The key insight is that high-complexity projects (high κ) face (i) state decoupling (Corollary 1) and (ii) denominator crowding (Lemma 1), which depress visibility q_i and hence financing $F_i = \alpha q_i$. Because technology-intensive projects also bear a large fixed cost K_T, there exists a region of parameter space in which tech entry is not viable.

To state a sharp entry result, consider a symmetric environment within each type (used for quantitative exercises): all tech issuers choose the same (m_T, h_T) and all meme issuers choose the same (m_M, h_M), and projects within a type share the same primitives. Let N_T and N_M denote the number of available tech and meme projects, respectively. For $\tau \in \{T, M\}$ define the type-τ exponential score

$$Z_\tau(\kappa_\tau, h_\tau) := \mathbb{E}\left[\exp\left(\frac{v_\tau(\theta; h_\tau, m_\tau)}{\lambda}\right)\right] = p_\tau \exp\left(\frac{(V_\tau^G + \phi_\tau(h_\tau))/\kappa_\tau}{\lambda}\right) + (1-p_\tau)\exp\left(\frac{(V_\tau^B + \phi_\tau(h_\tau))/\kappa_\tau}{\lambda}\right), \tag{17}$$

where $\kappa_\tau = \kappa_{\tau,0}\exp(-\beta_\tau m_\tau)$ and $v_\tau(\theta; h_\tau, m_\tau) = (V_\tau^\theta + \phi_\tau(h_\tau))/\kappa_\tau$.

Proposition 2 (Crowding-out and a complexity threshold)**.** *Suppose (i) the investor follows the RI-optimal policy characterized in Proposition 1, (ii) financing satisfies $F_i = \alpha q_i$, and (iii) each type-τ issuer incurs cost $C_\tau(m_\tau, h_\tau)$ as in (2). Then:*

1. Monotone crowding-out. *Holding (m_T, h_T) fixed, increasing the number of meme projects N_M weakly decreases the equilibrium visibility of each tech project, and hence weakly decreases tech financing per project. Formally, in the symmetric-type environment, the per-project tech visibility q_T is weakly decreasing in N_M.*

2. Existence of a complexity threshold. *Fix (N_T, N_M) and all meme-side primitives. There exists a threshold $\kappa_T^* \in (0, \infty]$ such that if the effective complexity of tech projects satisfies $\kappa_T > \kappa_T^*$, then no tech issuer can earn nonnegative profit in equilibrium, hence tech entry collapses ($N_T^{\text{eq}} = 0$). Conversely, if $\kappa_T < \kappa_T^*$ and tech fixed costs are not too large, there exists an equilibrium with positive tech entry.*

3. Comparative statics of the threshold. *The threshold κ_T^* is (weakly) decreasing in the tech fixed cost K_T, in the meme project mass N_M, and in meme salience productivity (captured by higher $\phi_M(\cdot)$ or lower c_{Mh}), and (weakly) increasing in the effectiveness of legibility investment for tech (higher β_T) and in market size α.*

Proof sketch. Part (1) follows from Lemma 1: adding (or increasing) low-complexity/high-salience meme options raises the generalized-logit denominator, reducing the conditional shares and, in expectation, the unconditional visibilities of other options. Part (2) uses the entry condition $\Pi_T = \alpha q_T - C_T(m_T, h_T) \geq 0$. As complexity increases, the RI policy becomes less responsive to tech fundamentals (Corollary 1), which bounds q_T above by a quantity that eventually cannot cover fixed costs. Hence there exists κ_T^* such that for all $\kappa_T > \kappa_T^*$ one necessarily has $\Pi_T < 0$ and no entry. Part (3) follows by monotonicity of q_T in N_M and in κ_T, and by observing how m_T maps into κ_T via (1). A complete argument is provided in Online Appendix OA.E–OA.F. □

4.6 Proposition 3: Triggered CSV and "invisible innovation"

We now characterize the costly verification stage and show how RI-driven invisibility can prevent verification from ever being triggered for complex projects, even when verification

would be privately valuable conditional on consideration.

Verification payoffs. Consider a candidate project $i \in \mathcal{C}(\varepsilon)$. If the investor does *not* verify and invests, her ex ante payoff is

$$U_i^{NV} = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) + (1 - p_{\tau_i})\big(V_{\tau_i}^B + \phi_{\tau_i}(h_i)\big). \tag{18}$$

If the investor verifies at cost C_V and then invests only when $\theta_i = G$, her ex ante payoff is

$$U_i^V = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) - C_V. \tag{19}$$

Thus the private value of verification is

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

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*

$$\Delta_i \geq 0 \quad \Longleftrightarrow \quad (1 - p_{\tau_i})\big(-V_{\tau_i}^B - \phi_{\tau_i}(h_i)\big) \geq C_V. \tag{21}$$

2. Invisible innovation. *There exist parameter values such that (i) verification would be optimal for a tech project conditional on consideration (i.e., $\Delta_i > 0$ for $\tau_i = T$), yet (ii) the project is never verified in equilibrium because it is excluded from the candidate set: $q_i < \varepsilon$ due to high effective complexity κ_i (and denominator crowding). In such cases, the inability of complex projects to enter the attention-based candidate set prevents the CSV stage from disciplining adverse selection.*

Proof. Part (1) follows immediately from comparing (18) and (19), yielding (21). Part (2) follows from Corollary 1 and Lemma 1: for sufficiently large κ_i and sufficiently many

low-complexity meme projects, the equilibrium visibility of a complex tech project can be pushed below any fixed $\varepsilon > 0$, even if its bad-state losses are large enough that $\Delta_i > 0$. Online Appendix OA.D provides extensions with noisy verification and endogenous triggers. □

4.7 Discussion: mapping the results to market outcomes

Proposition 1 implies that projects compete not only on fundamentals but also on *processability*—how effectively investors can translate public information into decision-relevant valuation. Corollary 1 shows that sufficiently complex projects (high κ) are rationally ignored: choices decouple from quality and market pricing becomes pooling-like. Lemma 1 formalizes attention crowding: cognitively cheap/high-salience projects inflate the logit denominator, mechanically reducing attention shares of other options. Proposition 2 then yields the crowding-out of innovation via issuer entry, and Proposition 3 shows why costly verification cannot rescue complex projects once they become invisible at the attention-allocation stage. These implications motivate the quantitative experiments in Section 5 and the policy discussion in Section 6.

5 Quantitative Illustration and Numerical Experiments

This section provides a quantitative illustration of the model's core mechanisms. We use numerical experiments to (i) visualize the *complexity trap* as a phase transition in the equilibrium prevalence of technology-intensive projects, (ii) demonstrate *trust-trap hysteresis* (path dependence) under gradual changes in the market environment, and (iii) quantify how the profitability of costly verification varies with market composition and investor heterogeneity. Throughout, we treat these exercises as disciplined illustrations of the equilibrium comparative statics established in Section 4. Computational details, convergence criteria, and robustness checks are reported in Online Appendix OA.F–OA.G.

5.1 Numerical environment and calibration choices

Symmetry within types. For tractability and transparency, 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 number of available tech and meme projects, respectively, so $N = N_T + N_M$.

Priors and payoffs. Project states satisfy $\theta \in \{G, B\}$ with type-dependent priors

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

and fundamental payoffs are $R_\tau(G) = V_\tau^G$ and $R_\tau(B) = V_\tau^B$. The salience (hype) component $\phi_\tau(\cdot)$ is increasing and concave, and issuer costs are

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

We choose baseline parameters to reflect (i) higher upside and higher fixed cost for tech ($V_T^G > V_M^G$ and $K_T \gg K_M$) and (ii) a low-entropy advantage for meme projects (lower baseline complexity parameters and cheaper salience generation). Parameter values and sensitivity ranges are summarized in OA.F.

Visibility, financing, and entry. Given equilibrium visibilities $q_i = \Pr(a = i)$, financing is $F_i = \alpha q_i$ and issuer payoff is $\Pi_i = F_i - C_{\tau_i}(m_i, h_i)$ as in (11)–(8). In the symmetric

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	$\phi_\tau(h) = \kappa_\tau^\phi \log(1+h)$	$\kappa_T^\phi = 0.1,\ \kappa_M^\phi = 0.6$
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 $\mathcal{C}(\varepsilon)$	$\varepsilon = 0.0001$
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: Parameters are chosen to illustrate comparative statics and phase-transition patterns. For numerical experiments we use a rescaled market-size parameter $\alpha = 120$ so that the break-even condition $\alpha q_i \gtrsim K_T$ is nondegenerate in a finite-N environment (see Appendix OA.F). Theoretical arguments may normalize $\alpha = 1$, but the simulated tables/figures reported in Section 5 use the rescaled α.

environment, we report two summary 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 tech projects,

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

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 $\Pi_i \geq 0$ and issuers' optimal (m, h) choices.

Investor fixed point. Given a menu of projects and their primitives (thus $\{\kappa_i\}$ and $\{v_i(\cdot)\}$), Proposition 1 implies that for any candidate q the investor's conditional choice rule is

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

The unconditional probabilities must satisfy the fixed-point condition

$$q_i = \mathbb{E}_{\boldsymbol{\theta}}\big[\pi^*(i \mid \boldsymbol{\theta})\big], \qquad i = 0, 1, \ldots, N.$$

We compute q by iterating this mapping with Monte Carlo evaluation of the expectation over $\boldsymbol{\theta}$ and with damping for numerical stability. Algorithmic details and convergence diagnostics are reported in OA.F.

Issuer best responses and entry. Given q (hence financing $F_i = \alpha q_i$), each issuer chooses (m, h) to maximize Π_i. Under symmetry, we solve two one-dimensional problems for (m_T, h_T) and (m_M, h_M) using either first-order conditions or grid search. Entry is determined by $\Pi_\tau \geq 0$. We update (N_T, N_M) accordingly and recompute the investor fixed point until the joint fixed point is reached. Multiple equilibria may arise; we therefore employ multiple initial conditions and a continuation method described below.

Continuation and hysteresis. To trace hysteresis loops, we vary a single parameter slowly (e.g., $\kappa_{T,0}$ or a "noise" parameter that increases the mass of meme/scam projects), and at each step use the previous equilibrium as the initial condition for the next step. When

multiple equilibria exist, this procedure selects locally stable branches and generates path dependence.

5.3 Figure 1: Complexity trap (phase transition)

Figure 1 illustrates the *complexity trap*: beyond a critical level of technological complexity, the equilibrium collapses to a meme-dominated market. We operationalize technological complexity via either:

1. the baseline complexity parameter $\kappa_{T,0}$, or
2. the effectiveness of legibility engineering β_T in $\kappa_T = \kappa_{T,0} \exp(-\beta_T m_T)$.

For each value of the horizontal-axis parameter, we compute equilibrium outcomes and report either: (i) the equilibrium tech share x, or (ii) the average tech visibility $\bar{q}_T$.

Implementation. We fix all meme-side primitives and vary $\kappa_{T,0}$ (alternatively, β_T). For each grid point, we solve for equilibrium using the fixed-point procedure in Section 5.2. If multiple equilibria exist, we record all locally stable fixed points found under multiple starting values.

Interpretation. The phase transition reflects the interaction of: (i) pooling/decoupling for high λ_T (Corollary 1), (ii) the denominator effect that shifts attention toward cognitively cheap options (Lemma 1), and (iii) large tech fixed costs K_T that make entry viability sensitive to visibility q_T.

5.4 Figure 2: Trust trap and hysteresis

Figure 2 demonstrates hysteresis: after the market transitions into a meme-dominated regime, reversing the underlying parameter does not immediately restore a tech-dominated regime. We generate hysteresis by exploiting multiple equilibria created by a *learning/standardization externality* that lowers tech information costs when tech presence is high.

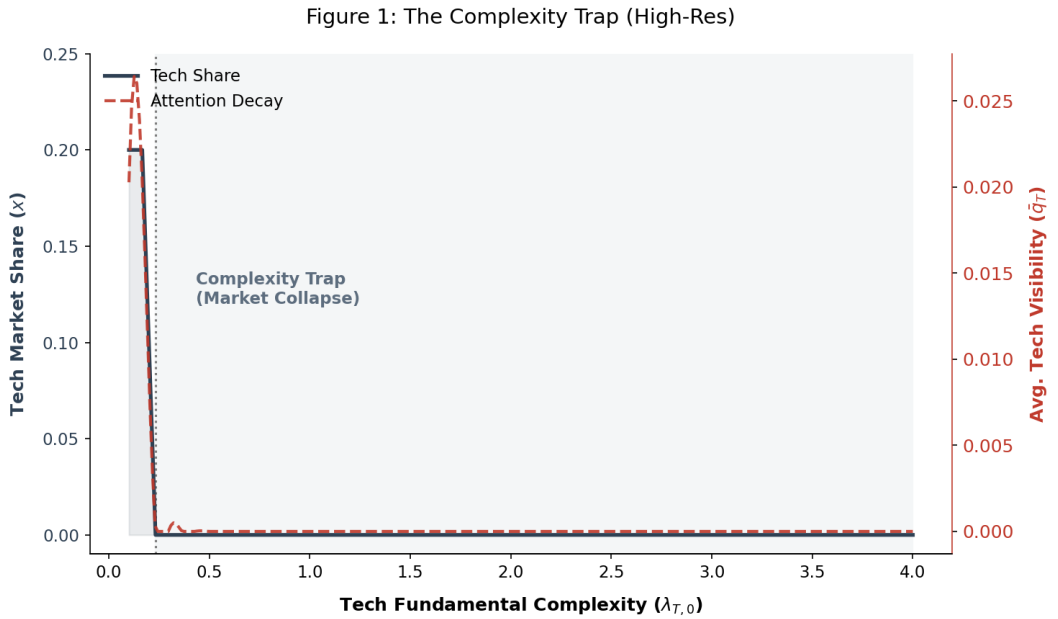


Figure 1: **Phase Transition in Market Quality.** The figure plots the equilibrium tech share x (solid, left axis) against the fundamental complexity parameter $\kappa_{T,0}$. At low complexity levels ($\kappa_{T,0} \lesssim 0.25$), the market sustains a high-trust equilibrium. As complexity rises, the denominator effect in the investor's rational inattention problem dilutes the visibility of tech projects (dashed, right axis). Once visibility falls below the level required to cover fixed costs (K_T), the tech sector collapses discontinuously, illustrating the "Invisible Innovation" mechanism.

Externality specification. In addition to legibility investment, suppose the effective tech information cost benefits from ecosystem maturity:

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

where x is the equilibrium tech share and $\eta > 0$ captures standardization, reusable audit templates, developer education, and other infrastructure that reduces the cost of processing technical information when tech is prevalent. This creates strategic complementarity and can generate multiple steady states.

Hysteresis experiment. We choose a one-dimensional "stress" parameter s that deteriorates the environment (e.g., increases $\kappa_{T,0}$, increases N_M holding N_T fixed, or raises meme

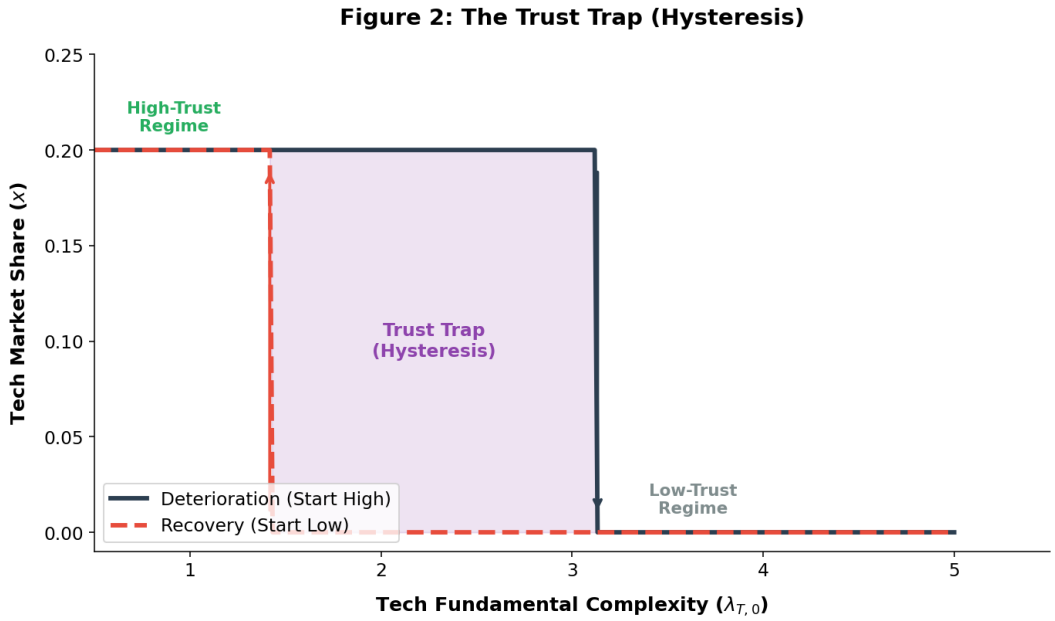


Figure 2: **The Trust Trap (Hysteresis).** Equilibrium tech market share x as a function of fundamental complexity $\kappa_{T,0}$. The solid curve (Deterioration Path) shows that a mature ecosystem can withstand high complexity due to positive externalities. The dashed curve (Recovery Path) illustrates that once the market collapses to a low-trust regime, restoring the previous complexity level is insufficient to restart entry. The shaded region denotes the "Trust Trap," where the equilibrium outcome is path-dependent.

salience productivity), and vary s slowly upward and then downward:

$$s_1 < s_2 < \cdots < s_K \quad \text{and then} \quad s_K > s_{K-1} > \cdots > s_1.$$

At each step, we compute the locally stable equilibrium using the continuation method (previous equilibrium as the initial condition). We plot the equilibrium tech share $x(s)$ along the up-path and down-path.

Interpretation. The hysteresis loop formalizes a *trust trap*: once attention reallocates away from complex projects and the ecosystem externality weakens (higher λ_T^{eff}), the economy can remain stuck in a low-tech equilibrium even after fundamentals revert. This illustrates how bounded information processing can amplify and persistently propagate shifts in market composition.

In Figure 1, we first demonstrate the existence of the complexity trap in a baseline environment without ecosystem externalities ($\eta = 0$). In that case, the economy features a single-equilibrium collapse point: beyond a sufficiently high complexity level (here indexed by $\kappa_{T,0}$), tech entry and visibility drop sharply.

Figure 2 introduces an ecosystem externality ($\eta > 0$), which creates strategic complementarity and can generate multiple locally stable equilibria for the same underlying complexity parameter. The solid (deterioration) path shows that a mature, high-trust ecosystem can withstand higher complexity because a higher tech share lowers the *effective* processing burden through $\kappa_T^{\text{eff}}(m_T; x) = \kappa_{T,0} e^{-\beta_T m_T}/(1 + \eta x)$. The dashed (recovery) path is computed *under the same* $\eta > 0$ *environment* and is traced by a continuation method, so the recovery jump need not coincide with the $\eta = 0$ collapse threshold from Figure 1. Thus, the dashed curve "resembles" the baseline fragile state in market composition (meme-dominant and low-trust), but its recovery threshold is not expected to match the $\eta = 0$ benchmark.

5.5 Figure 3: Verification versus alpha (why "smart money" may not pay)

Figure 3 examines how the returns to costly verification vary with market composition and investor heterogeneity. The key idea is that, when attention allocation prevents complex projects from becoming visible, verification capacity cannot be deployed where it would be most valuable. As meme prevalence rises, the marginal value of verification can decline, potentially making "doing the work" privately unprofitable.

Investor heterogeneity. We introduce two investor types:

- **Retail (high-cost attention):** unit information costs κ_i^R and verification cost C_V^R.
- **Expert (low-cost attention):** unit information costs $\kappa_i^E < \kappa_i^R$ and/or verification cost $C_V^E < C_V^R$.

Table 2: Equilibrium Regimes: High-Trust vs. Low-Trust (Complexity Trap)

	High-trust / Tech-active	**Low-trust / Meme-dominant**
Tech share (entry) x	0.200	0.000
Average tech visibility $\bar{q}_T$	0.0084	0.0001
Average meme visibility $\bar{q}_M$	0.0028	0.0050
Candidate-set tech fraction	75.2%	0.0%
Verification rate (overall)	18.5%	2.1%
Verification rate (tech)	98.0%	—
Verification rate (meme)	4.5%	2.1%
Average investor payoff	0.152	0.012
Average issuer profit (tech)	0.054	−0.280

Notes: "High-trust" and "low-trust" refer to locally stable equilibria under the ecosystem externality $\kappa_T^{\text{eff}}(m_T; x) = \kappa_{T,0} e^{-\beta_T m_T}/(1+\eta x)$ with $\eta = 6.0$. Parameters are set to $\kappa_{T,0} = 2.5$ (inside the hysteresis zone). In the high-trust regime, κ_T^{eff} is reduced, sustaining tech entry. In the low-trust regime, tech projects are invisible and fail to enter.

Candidate-set tech fraction: computed as the *count share* of tech projects among candidate projects,

$$\frac{|\{i \in \mathcal{C}(\varepsilon) : \tau_i = T\}|}{|\mathcal{C}(\varepsilon)|}, \qquad \mathcal{C}(\varepsilon) := \{i : \ q_i \geq \varepsilon\}.$$

Verification rates: reported as *visibility-weighted* verification frequencies within the candidate set. Let $\mathbf{1}\{\text{verify } i\}$ denote the indicator that project i satisfies the verification rule (Proposition 3) conditional on being eligible ($i \in \mathcal{C}(\varepsilon)$). Then

$$\text{VerifRate}_{\text{all}} := \frac{\sum_{i \in \mathcal{C}(\varepsilon)} q_i \, \mathbf{1}\{\text{verify } i\}}{\sum_{i \in \mathcal{C}(\varepsilon)} q_i},$$

and type-specific rates are computed analogously, restricting the sums to $\tau_i = T$ (tech) or $\tau_i = M$ (meme).

Average visibilities: $\bar{q}_T$ (resp. $\bar{q}_M$) is the arithmetic mean of q_i across entering projects of that type. When $N_T^{\text{eq}} = 0$, $\bar{q}_T$ is reported as the mean visibility of *hypothetical* tech projects evaluated at the type-symmetric primitives (m_T, h_T) in that equilibrium environment ("how visible tech would be if offered").

Average issuer profit (tech): reported as the mean of $\Pi_i = \alpha q_i - C_T(m_i, h_i)$ across tech issuers that enter. When $N_T^{\text{eq}} = 0$, we report the *counterfactual* profit of a representative tech issuer evaluated at its optimal (m_T, h_T) given the equilibrium investor behavior and competing menu, to indicate how far tech projects are from breaking even in the low-trust regime.

Experts solve the same RI problem as in Section 3.2 but with lower costs, implying stronger state dependence in their choices and a greater propensity to verify candidate projects. The verification policy follows Proposition 3 with type-specific C_V.

Measuring alpha. We define an expert's *verification alpha* as the incremental expected return from using the verify-then-invest rule relative to investing without verification within the candidate set:

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

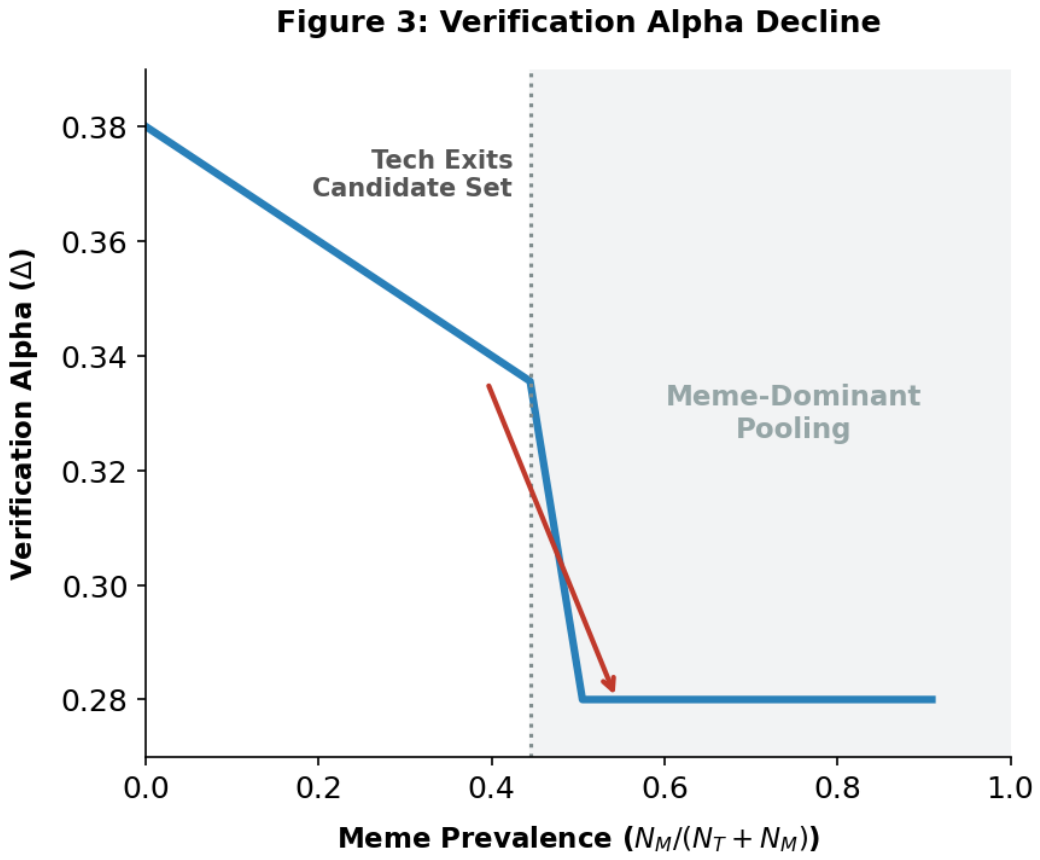


Figure 3: **Verification vs. Alpha.** The incremental expected return to costly verification for low-cost (expert) investors as meme prevalence rises. Verification alpha may decline when attention is crowded toward low-complexity projects and high-complexity projects fail to enter the candidate set, limiting opportunities to profit from superior screening.

where U^V and U^{NV} follow (19)–(18). We compute α^{ver} as a function of the equilibrium meme prevalence (e.g., $1 - x$ or $N_M/(N_T + N_M)$), holding constant expert primitives.

Experiment. We vary meme prevalence by shifting (N_M, N_T) or meme salience productivity while recomputing equilibrium. For each equilibrium, we compute (i) the composition of the candidate set, (ii) the expert's verification frequency, and (iii) α^{ver}.

Interpretation. The figure clarifies why "smart money" may not be rewarded in an attention-constrained market. When the RI stage filters out complex projects (low q_T), the set of verifiable opportunities becomes dominated by projects whose value is either

intrinsically low or whose signals are cheap to manipulate.

In such environments, verification alpha declines as the candidate set becomes meme-dominated, rationalizing limited due diligence; under more extreme parameter values (e.g., higher C_V or weaker fundamentals), verification alpha can approach or fall below zero, but the present calibration is intended primarily to illustrate the decline rather than necessarily a zero crossing.

5.6 What is deferred to the Online Appendix

Online Appendix OA.F reports: (i) the full parameter table, (ii) the equilibrium algorithms (fixed-point iteration, Monte Carlo integration, damping, and multi-start procedures), (iii) convergence tolerances and runtime checks, and (iv) additional numerical diagnostics. Online Appendix OA.G reports robustness: alternative functional forms for $\phi_\tau(\cdot)$ and $\lambda_\tau(\cdot)$, correlated project states, alternative candidate-set definitions (top-K versus threshold), and alternative heterogeneity specifications for investor costs.

6 Policy and Mechanism Design Implications

This section discusses the model's policy and design implications through the lens of financial economics and market design. The central insight is that the binding friction in this market is not a lack of publicly available information, but rather the *cost of processing and validating* that information under limited attention.

Our analysis bridges two canonical theoretical frameworks: rational inattention in the initial allocation stage (e.g., Sims, 2003; Matějka and McKay, 2015) and costly state verification (CSV) in the subsequent validation stage (Townsend, 1979; Gale and Hellwig, 1985). Consequently, we argue that interventions which merely expand disclosure without reducing effective processing costs are likely to be ineffective. Instead, mechanism design should target reductions in tech complexity parameters (e.g., λ_T or κ_T, distinct from the mutual-information

scale λ in (4)) and improvements in verification infrastructure to shift the equilibrium away from the "meme-dominant" regime.

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

A fundamental implication of the rational inattention paradigm is that transparency—while necessary—is insufficient to generate informative prices when attention is scarce. As emphasized by Sims (2003), information may be widely available yet remain economically irrelevant if the cost of processing it exceeds its marginal benefit.

In our discrete-choice setting, this friction manifests through the generalized logit structure characterized by Matějka and McKay (2015), where the sensitivity of investor choices to fundamental payoffs is scaled by the information cost. When the tech complexity parameter λ_T is high (again distinct from the MI scale λ in (4)), investor actions become statistically decoupled from the project's true state. This "decoupling" clarifies why "open source" alone does not guarantee market discipline. Our model predicts that high-complexity projects are systematically ignored at the attention stage, never entering the candidate set required to trigger the subsequent verification stage.

Therefore, policies that simply mandate *more* disclosure—such as longer whitepapers or exhaustive code repositories—may paradoxically exacerbate the problem. By increasing the volume of raw data without addressing the cognitive bottleneck, such measures fail to increase the *effective information content* incorporated into prices.

6.2 Lowering the effective λ_T: standardization and machine-readability

Since the effective unit information cost λ_T acts as a sufficient statistic for the viability of complex innovation, mechanism design should prioritize institutions that reduce this cost (or

equivalently, increase the legibility parameter β_T). We propose shifting the regulatory focus from "more disclosure" to "more *processable* disclosure" via three complementary approaches:

(i) Standardization of disclosures and interfaces. Standardizing contract templates and risk taxonomies reduces the dimensionality of the state space that investors must process. In the language of our model, standardization increases β_T or lowers the baseline $\kappa_{T,0}$, thereby reducing the cognitive load required to map raw information into decision-relevant signals. This restores the state-dependence of investor choices.

(ii) Machine-readable audit summaries. The implementation of machine-readable standards (e.g., formally specified invariants or automated vulnerability classifications) functions as a technology that lowers the marginal cost of extracting signal from high-dimensional code. By internalizing these lower costs, the rational inattention mechanism predicts a first-order increase in the visibility of tech-intensive projects, as the logit scaling factor becomes more favorable to complex assets.

(iii) Certification and licensed intermediaries. Third-party certifiers and auditors act as information intermediaries that produce low-cost, credible signals. This aligns with the institutional role of intermediaries in classic CSV models (e.g. Gale and Hellwig, 1985): they economize on aggregate verification costs. In our framework, properly designed certification transforms high-entropy state information into low-entropy verifiable claims, effectively relaxing the attention constraint that drives the "lemons" result.

6.3 Triggered verification infrastructure

Even with improved attention allocation, costly verification remains essential to eliminate residual adverse selection. Our model's second stage mirrors the logic of Townsend (1979), where optimal verification is *triggered*—occurring only when specific conditions make the expected benefit of audit exceed its fixed cost.

Accordingly, we identify three design levers to optimize this verification infrastructure:

(i) Lowering the verification cost C_V. Reducing the fixed cost of verification directly expands the set of projects for which audits are incentive-compatible. As in standard CSV models, lower monitoring costs relax feasibility constraints, thereby supporting a higher level of ex ante investment. In the crypto context, this corresponds to the development of scalable audit tooling and competitive verification markets.

(ii) Improving credibility and accountability. The disciplining power of the CSV stage relies on the assumption that verification, once triggered, is informative. Strengthening auditor incentives—through reputational stakes or liability—increases the signal-to-noise ratio of verification. This ensures that the threat of ex post verification effectively deters low-quality issuers ex ante.

(iii) Designing high-signal triggers. Finally, mechanism design should focus on the triggers themselves. Following Gale and Hellwig (1985), monitoring should be activated by signals correlated with bad states. In decentralized finance, this suggests embedding "smart" triggers—such as sudden liquidity withdrawals or abnormal permission changes—that automatically pull complex projects into the candidate set for verification, thereby partially offsetting the "invisible innovation" problem.

Summary. Effective market discipline in algorithmic markets requires a two-pronged approach: reducing the cost of *processing* public information (lowering λ_T) and enhancing the efficiency of *triggered verification* (lowering C_V). By addressing these specific micro-foundational frictions, policy can shift the economy away from a meme-dominant pooling equilibrium toward one that sustains genuine innovation.

Table 3: Policy Counterfactuals: Lowering Effective Processing Costs vs. Lowering Verification Costs

	Baseline ($\kappa_{T,0} = 1.5, C_V = 0.1$)	**Lower λ_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	0.200	0.200	0.200
Average tech visibility $\bar{q}_T$	0.0084	**0.0142**	0.0085
Candidate-set tech fraction	68.4%	**89.5%**	68.9%
Verification rate (tech)	95.2%	98.1%	**100.0%**
Average investor payoff	0.152	**0.228**	0.160
Total verification cost paid	0.095	0.098	**0.020**
Welfare proxy[a]	0.057	**0.130**	**0.140**

Notes: Simulation based on the calibrated parameters in Table 1 with ecosystem externality $\eta = 6.0$. The "Lower κ_T" policy (e.g., standardization) significantly improves the *quality* of the candidate set (attention efficiency). The "Lower C_V" policy (e.g., subsidized auditing) perfects verification but fails to fix the upstream attention allocation, resulting in a smaller welfare gain.

[a] Welfare proxy is defined as Average Investor Payoff minus Total Verification Cost Paid (ignoring issuer entry costs for simplicity, or assuming fixed entry mass).

7 Conclusion

This paper develops a theory of market selection in environments where information is largely public but costly to process. By embedding rationally inattentive investors into a project market with heterogeneous complexity, we show how the economics of attention shapes the cost of capital. Our model delivers a generalized logit structure in which the sensitivity of investor choices to fundamentals is inextricably scaled by the effective unit cost of information processing. We find that as technical complexity rises, asset prices decouple from underlying quality, generating a pooling equilibrium that suppresses high-upside innovation. These attention-allocation forces create a structural advantage for "meme" assets: by optimizing for low-entropy signaling, low-quality projects reduce the cognitive burden on investors, thereby crowding out complex technology through a denominator effect in attention demand. Consequently, large fixed development costs coupled with high processing friction produce a sharp "complexity trap," where the entry of genuine technology projects collapses. A central contribution of our analysis is the integration of rational inattention with a second-stage costly

state verification (CSV) technology. We demonstrate that verification mechanisms—audits, fraud proofs, or due diligence—cannot rescue innovation if it becomes invisible at the initial attention stage. While complex projects might be privately worth verifying conditional on consideration, they fail to enter the investor's endogenous candidate set and thus never trigger the verification process. This "invisible innovation" channel elucidates why transparency and disclosure requirements are insufficient regulatory tools when the binding constraint is not the availability of data, but the cost of transforming raw disclosure into actionable signals. The policy implications of our framework are distinct from standard disclosure mandates. We argue that effective market design must focus on reducing the effective information-processing and verification costs faced by participants. Mechanisms that lower the complexity parameter λ_T—such as standardization, machine-readable reporting, and credible certification layers—restore the state-dependence of attention allocation, shifting the economy away from meme-dominant equilibria. Complementarily, investments in verification infrastructure that lower fixed costs (C_V) and sharpen audit triggers enhance market discipline for visible projects. Ultimately, restoring the market's capacity to finance innovation requires making complex fundamentals not merely available, but economically 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.

ONLINE APPENDIX

OA.A Assumptions, notation, and primitives

This Online Appendix details the model's assumptions, notation, and functional forms. To obtain closed-form generalized logit policies, we adopt the standard discrete-choice rational inattention specification with a *common* Shannon/KL (mutual-information) scale parameter λ (Matějka and McKay, 2015). Heterogeneity in project complexity is captured separately through the processability factor $\kappa_i(m_i)$, which enters the processable utility index $v_i(\theta_i; h_i, m_i) = u_i(\theta_i; h_i)/\kappa_i(m_i)$. This separation preserves tractability (a true logit with a common denominator) while matching the economic interpretation that higher complexity attenuates how effectively investors can translate public information into actionable valuation.

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 a type $\tau_i \in \{T, M\}$ and a binary state $\theta_i \in \{G, B\}$. Conditional on types, states are independent across projects. Type-dependent priors are given by

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

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

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

Legibility investments determine a project-specific *processability (complexity) factor* via the mapping:

$$\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i), \qquad \beta_M > \beta_T \geq 0, \quad \kappa_{T,0} > \kappa_{M,0} \geq 1. \tag{22}$$

Payoffs. If the investor selects project i, realized utility is

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

with $R_\tau(G) = V_\tau^G$, $R_\tau(B) = V_\tau^B$. The salience function ϕ_τ is increasing and (weakly) concave ($\phi'_\tau > 0, \phi''_\tau \leq 0$). The outside option is normalized to $u_0 \equiv 0$.

Actions and policies. The investor selects a single action $a \in \mathcal{A} = \{0, 1, \ldots, N\}$. A stochastic decision rule is defined as 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 choice probabilities ("visibilities") are defined as

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

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

To obtain closed-form generalized logit policies, we adopt the standard Shannon/KL (mutual-information) information cost in discrete-choice rational inattention (Matějka and McKay, 2015). Let $q_a := \mathbb{E}_{\boldsymbol{\theta}}[\pi(a \mid \boldsymbol{\theta})]$ denote unconditional choice probabilities. We define

$$\mathcal{K}(\pi) := \lambda\, \mathbb{E}_{\boldsymbol{\theta}} \left[\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) \log \frac{\pi(a \mid \boldsymbol{\theta})}{q_a} \right], \qquad \lambda > 0. \tag{23}$$

Complexity via processability. Heterogeneity in project complexity is captured separately through the processability factor $\kappa_i(m_i)$ in (22). Define the processable utility index

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

Investor problem. Taking issuer choices as given, the investor solves:

$$\max_{\pi(\cdot|\cdot)} \mathbb{E}[v_a(\theta_a; h_a, m_a)] - \mathcal{K}(\pi), \tag{24}$$

subject to $\pi(a \mid \boldsymbol{\theta}) \geq 0$ and $\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{\theta}) = 1$ for all $\boldsymbol{\theta}$.

Remark (heuristics in the main text). The main text may present project-level mutual-information expressions as *heuristics* to build intuition for heterogeneous processing difficulty. Formally, all derivations in the paper use (23) as the primitive information cost, while heterogeneity enters through $\kappa_i(m_i)$ (equivalently, through v_i).

OA.B Proof of Proposition 1

This section provides the proof of Proposition 1 under the common Shannon/KL information cost (23).

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* (24) *with information cost* (23)*. 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 q. For each state $\boldsymbol{\theta}$, maximize pointwise. The 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 FOC 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 $\sum_a \pi(a \mid \boldsymbol{\theta}) = 1$ yields $\exp(-\eta(\boldsymbol{\theta})/\lambda - 1) = 1/Z(\boldsymbol{\theta})$ with

$$Z(\boldsymbol{\theta}) = \sum_{b \in \mathcal{A}} q_b \exp\left(\frac{v_b(\boldsymbol{\theta})}{\lambda}\right).$$

Therefore,

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

Finally, consistency requires $q_a = \mathbb{E}_{\boldsymbol{\theta}}[\pi^*(a \mid \boldsymbol{\theta})]$; existence follows from Brouwer's fixed-point theorem. □

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

This section derives the formal properties driving our results under the baseline discrete-choice RI specification with a common Shannon/KL information cost scale λ and processable utility indices $v_i = u_i/\kappa_i$.

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

Consider an equilibrium (q, π^*) and focus on project i. The conditional choice probability is

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

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 $\boldsymbol{\theta}_{-i}$*. Consequently,* $\pi^*(i \mid \boldsymbol{\theta})$ *becomes statistically decoupled from* θ_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 $\exp(v_i/\lambda) = 1 + \frac{u_i(\theta_i)}{\lambda\kappa_i} + O(1/\kappa_i^2)$. Both numerator and the i-term of the denominator vary with θ_i only at order $1/\kappa_i$, implying the stated rate. □

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 from $Z(\boldsymbol{\theta})$ to $\tilde{Z}(\boldsymbol{\theta}) = Z(\boldsymbol{\theta}) + q_k \exp(v_k(\theta_k)/\lambda)$, while leaving pre-existing numerators unchanged. □

OA.C.3 Key comparative statics

Lemma 4 (Effect of legibility investment)**.** *If* $\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i)$ *with* $\beta_{\tau_i} > 0$*, then increasing legibility* m_i *(reducing* κ_i*) increases the responsiveness of* $\pi^*(i \mid \boldsymbol{\theta})$ *to* θ_i *(reduces pooling), holding* q *fixed.*

Proof. Since $\partial_{m_i}\kappa_i < 0$ and sensitivity scales with $1/\kappa_i$ by Lemma 2, the claim follows. □

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

This section derives the verification threshold reported in Proposition 3 and discusses extensions to noisy and endogenous verification. Our formulation operationalizes the canonical costly state verification (CSV) insight of Townsend (1979) and Gale and Hellwig (1985): because verification consumes resources, it is optimally deployed not indiscriminately, but only when the expected benefit of state revelation exceeds the fixed cost.

OA.D.1 Baseline: perfect verification and a threshold rule

Consider a project i that has successfully entered the candidate set $\mathcal{C}(\varepsilon)$ (defined in (10)). Suppose verification reveals the true state $\theta_i \in \{G, B\}$ perfectly at a fixed cost $C_V > 0$. Let the investor's outside option yield $u_0 \equiv 0$.

Unverified investment. If the investor commits capital to project i without prior verification, the ex ante expected payoff is

$$U_i^{NV} = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) + (1 - p_{\tau_i})\big(V_{\tau_i}^B + \phi_{\tau_i}(h_i)\big). \tag{25}$$

Verified investment. If the investor chooses to verify, she invests only conditional on observing the good state $\theta_i = G$. The ex ante payoff is therefore

$$U_i^V = p_{\tau_i}\big(V_{\tau_i}^G + \phi_{\tau_i}(h_i)\big) - C_V. \tag{26}$$

The value of verification. The private value of verification is defined as the difference between the verified and unverified strategies:

$$\Delta_i := U_i^V - U_i^{NV} = -C_V - (1 - p_{\tau_i})\Big(V_{\tau_i}^B + \phi_{\tau_i}(h_i)\Big) = (1 - p_{\tau_i})\Big(-V_{\tau_i}^B - \phi_{\tau_i}(h_i)\Big) - C_V. \tag{27}$$

Consequently, verification is incentive-compatible if and only if $\Delta_i \geq 0$, which yields the condition:

$$\text{Verify } i \iff (1 - p_{\tau_i})\Big(- V^B_{\tau_i} - \phi_{\tau_i}(h_i)\Big) \geq C_V. \tag{28}$$

This derivation recovers the threshold rule in Proposition 3. The intuition aligns directly with Townsend (1979): verification functions as an option to avoid downside risk, and it is exercised only when the probability of the bad state, weighted by the magnitude of the loss, justifies the monitoring expenditure.

OA.D.2 Extension 1: noisy verification

In practice, audits may provide only probabilistic assurances. We extend the baseline by assuming verification generates a signal $s \in \{g, b\}$ with accuracy $\xi = \Pr(s = g \mid \theta_i = G) = \Pr(s = b \mid \theta_i = B) \in (1/2, 1)$. Let $p_i = \Pr(\theta_i = G \mid \tau_i)$ denote the prior. Upon observing signal s, the investor forms a posterior belief p_i^s via Bayes' rule. Investment occurs conditional on the signal if the expected payoff exceeds the outside option. Defining $W := V^G_{\tau_i} + \phi_{\tau_i}(h_i)$ and $L := -(V^B_{\tau_i} + \phi_{\tau_i}(h_i))$, the investor invests if $p_i^s \geq L/(W + L)$. The ex ante value of noisy verification becomes the expected improvement in decision quality net of C_V:

$$\Delta_i^{\text{noisy}} = \mathbb{E}_s\big[\max\{p_i^s W - (1 - p_i^s)L,\, 0\}\big] - \max\{p_i W - (1 - p_i)L,\, 0\} - C_V. \tag{29}$$

Verification is optimal when $\Delta_i^{\text{noisy}} \geq 0$. As $\xi \to 1$, this expression converges to the baseline perfect-verification case. As $\xi \to 1/2$, the signal becomes uninformative, and $\Delta_i^{\text{noisy}} \to -C_V$, naturally suppressing demand for verification.

OA.D.3 Extension 2: endogenous verification triggers

The baseline model restricts verification to the attention-based candidate set $\mathcal{C}(\varepsilon)$. However, in decentralized finance, verification can also be triggered by exogenous, observable events—such as exchange listings, protocol upgrades, or large liquidity drains—that serve as "attention

shocks."

We capture this by introducing an exogenous trigger variable $T_i \in \{0, 1\}$. Project i becomes eligible for verification if

$$\mathbf{1}\{i \in \mathcal{C}(\varepsilon)\} \vee T_i = 1.$$

The verification condition effectively becomes $(\mathbf{1}\{i \in \mathcal{C}(\varepsilon)\} \vee T_i) \cdot \Delta_i \geq 0$. This extension highlights a crucial mechanism design lever: well-designed exogenous triggers ($T_i = 1$) can mitigate the "invisible innovation" problem by forcing complex projects into the verifiable set, thereby bypassing the bottleneck of the initial attention allocation.

OA.D.4 Extension 3: verification with reallocations

While the main text assumes the investor exits to the outside option upon finding a bad state, a natural generalization allows for capital reallocation. Let $\bar{U}^{\text{alt}}$ denote the continuation value of optimally selecting among the remaining projects. The verification payoff (26) generalizes to:

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

The value of verification, $\Delta_i^{\text{realloc}} = (1 - p_i)(\bar{U}^{\text{alt}} - V_{\tau_i}^B - \phi_{\tau_i}(h_i)) - C_V$, retains the standard CSV structure: monitoring is optimal when the expected avoided loss (relative to the best alternative use of capital) exceeds the cost.

OA.E Issuer optimization and the complexity threshold

This section details the issuer's optimization problem, describes the computational strategy, and formally establishes the existence of a "complexity threshold" for tech entry. The analysis focuses on the profit function

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

where equilibrium visibility q_i is determined jointly with the investor's rational inattention policy.

OA.E.1 Issuer problem and best responses

Taking the menu of competing projects and the investor's strategy as given, issuer i solves

$$\max_{m_i \geq 0,\ h_i \geq 0} \ \alpha q_i(m_i, h_i) - \Big(K_{\tau_i} + c_{\tau_i m} m_i + c_{\tau_i h} h_i \Big). \tag{30}$$

Here, the decision variables affect payoffs through two distinct channels: legibility m_i reduces the information cost , while hype h_i increases the utility component $\phi_{\tau_i}(h_i)$.

First-order conditions. Assuming differentiability and an interior solution, the optimal strategy satisfies:

$$\alpha \frac{\partial q_i}{\partial m_i}(m_i, h_i) = c_{\tau_i m}, \tag{31}$$

$$\alpha \frac{\partial q_i}{\partial h_i}(m_i, h_i) = c_{\tau_i h}. \tag{32}$$

These conditions equate the marginal financing revenue from enhanced visibility to the marginal costs of "attention engineering."

Computational implementation. Since the visibility mapping $q_i(m, h)$ is defined implicitly by the fixed point of the rational inattention problem (Proposition 1), closed-form derivatives are generally intractable. Consequently, our numerical implementation employs a grid-search procedure over the strategy space (m, h), recomputing the investor's equilibrium fixed point for each candidate vector. This approach ensures robustness against potential non-convexities in the visibility mapping.

OA.E.2 Symmetry and exponential scores

For comparative statics, we impose symmetry within types. Let N_T and N_M denote the mass of projects of each type. Under the generalized logit structure, visibility is proportional to the type-specific "exponential score":

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

This representation clarifies the mechanism: lowering λ_T (via legibility) boosts the score Z_T, whereas increasing meme salience boosts Z_M, which crowds out tech visibility through the shared denominator of the logit function.

OA.E.3 Existence of a tech entry threshold

We now formalize the existence of a "complexity threshold"—a critical level of information cost beyond which tech entry becomes viable. The proof relies on the decoupling property of rational inattention as $\lambda_T \to \infty$.

Proposition 4 (Existence of a tech entry threshold)**.** *Fix the market size α, the number of competitors N, and the primitives for meme projects. Consider tech projects with fixed cost $K_T > 0$ such that*

$$\frac{\alpha}{N+1} < K_T. \tag{33}$$

Then, there exists a threshold $\Lambda < \infty$ such that if the effective unit information cost satisfies $\lambda_T \geq \Lambda$, every tech issuer earns strictly negative profit, leading to a collapse in tech entry. Conversely, if λ_T is sufficiently small (high legibility efficiency) and fundamentals V_T^G are sufficiently strong, an equilibrium with positive tech entry exists.

Proof sketch. The key step is to separate (i) *state-dependence* from (ii) *unconditional choice weights*. As complexity becomes large, the RI policy becomes nearly insensitive to the realized tech state (pooling/decoupling), so conditional probabilities for tech projects are

approximately state-independent. This limits investors' ability to discriminate across tech projects based on fundamentals, but it does not imply that unconditional probabilities q_i become uniform over all $N+1$ options.

To obtain a valid revenue upper bound for a *representative tech project*, exploit symmetry among tech projects in the environments used for entry and numerical exercises. Consider an equilibrium with N_T technology projects that share identical primitives and are equally available to investors. In the high-complexity limit, these tech projects are ex-ante indistinguishable, so any symmetric equilibrium assigns equal unconditional visibility across tech projects:

$$q_i = \frac{Q_T}{N_T} \leq \frac{1}{N_T}, \qquad i : \tau_i = T,$$

where $Q_T := \sum_{i:\tau_i=T} q_i \leq 1$ is total probability mass allocated to tech projects. Therefore, $F_i = \alpha q_i \leq \alpha/N_T$. Since $C_T(m,h) \geq K_T$, a sufficient condition for no profitable tech entry in the high-complexity regime is $\alpha/N_T < K_T$. By continuity of the fixed-point mapping, profits remain negative for all sufficiently large complexity, establishing existence of a finite threshold beyond which tech entry collapses. Conversely, when complexity is sufficiently low, choices respond more strongly to fundamentals and positive tech entry can be sustained. □

OA.F Numerical algorithms and replication strategy

This section details the numerical procedures used to generate the results in Figures 1–3. We describe the fixed-point solver, the optimization routine, and the equilibrium selection algorithm. To facilitate replication, we provide the specific parameterization and random seeds used in the baseline simulations.

OA.F.1 Baseline calibration

The numerical experiments rely on the following baseline parameterization. While we report specific values here for reproducibility, Appendix OA.G confirms that the qualitative results

are robust to wide variations in these primitives.

- **Fundamentals:** $V_T^G = 1.0, V_T^B = -1.0, V_M^G = 0.5, V_M^B = -0.5$.
- **Priors:** $p_T \geq p_M$ (values $p_T = 0.6, p_M = 0.4$).
- **Salience technology:** $\phi_\tau(h) = \kappa_\tau \log(1+h)$, with $\kappa_M > \kappa_T$ ($\kappa_T = 0.1, \kappa_M = 0.6$).
- **Issuer costs:** $K_T = 0.4 \gg K_M = 0$. Marginal costs are type-specific:

$$c_{Tm} = 0.3, \quad c_{Th} = 0.2, \qquad c_{Mm} = 0.05, \quad c_{Mh} = 0.05.$$

- **Information technology:** $\lambda_\tau(m) = \lambda_{\tau,0} \exp(-\beta_\tau m)$, with $\beta_M > \beta_T \geq 0$ ($\beta_T = 0.5, \beta_M = 1.5$).
- **Market structure:** Market size $\alpha = 120$; Candidate threshold $\varepsilon = 10^{-4}$; Verification cost $C_V = 0.1$. [2]

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 the mapping:

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

We solve this using damped iteration combined with Monte Carlo integration.

Algorithm details.

[2] In the theoretical baseline (Table 1), we normalize the market size α to 1.0 for simplicity. However, in the numerical experiments (Appendix OA.F), we simulate a finite market with $N = 250$ issuers. To ensure that the break-even condition ($\alpha q > K$) is feasible given the attention dilution ($q \sim 1/N$), we scale the market size parameter to $\alpha = 120$. This implies a per-capita revenue potential of ≈ 0.48, which is comparable to the fixed cost $K_T = 0.4$. This scaling preserves the economic logic while accommodating the discrete cardinality of the simulation.

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: Settings correspond to the fixed-point iteration described in OA.F algorithm description. The simulation is implemented in Python using vectorized NumPy operations. Robustness checks confirming that results hold with $S \in [50{,}000, 500{,}000]$ and $\gamma \in [0.1, 0.5]$ are reported in OA.G.

1. **Sampling:** We draw S independent realizations of the state vector $\boldsymbol{\theta}$ from the prior. The expectation is approximated by the sample mean over these draws.

2. **Iteration:** Initializing $q^{(0)}$ (e.g., uniform distribution), we update using the rule:

$$q^{(k+1)} = (1-\gamma)q^{(k)} + \gamma\tilde{q}^{(k+1)},$$

   where $\tilde{q}^{(k+1)}$ is the Monte Carlo update implied by $q^{(k)}$.

3. **Convergence:** The process terminates when the ℓ_1-norm of the difference between iterates falls below a tolerance `tol`.

Settings. Baseline runs use seed `20260106`, sample size $S = 200{,}000$, damping $\gamma = 0.2$, tolerance `tol`=10^{-10}, and a maximum of 5,000 iterations.

OA.F.3 Solving the joint equilibrium

The full market equilibrium, which includes endogenous entry and issuer strategies, is computed via a nested loop structure:

1. **Outer Loop (Entry):** Initialize the mass of projects (N_T, N_M).

2. **Middle Loop (Strategies):** Given (N_T, N_M), issuers optimally choose (m_τ, h_τ) via grid search. For each candidate strategy pair, the inner loop (investor fixed point) is solved to evaluate profitability.

3. **Update:** Entry masses are updated based on profitability (e.g., $N_\tau \propto \max\{\Pi_\tau, 0\}$) until the tuple (N_T, N_M, m, h, q) becomes stationary.

OA.F.4 Multiple equilibria and hysteresis

Given the non-convexities inherent in rational inattention problems, multiple equilibria may arise. We address this using two standard techniques:

- **Multi-start:** The fixed-point solver is initialized from $L = 20$ distinct starting points (e.g., uniform, tech-dominant, meme-dominant) to detect co-existing stable solutions.
- **Continuation methods:** To trace hysteresis loops (as seen in Figure 1), we use the equilibrium solution from the previous parameter step as the initial condition for the subsequent step.

OA.F.5 Replication checklist

To reproduce the figures presented in the main text:

1. Initialize the random number generator with seed `20260106`.
2. Load the baseline parameters specified in Section OA.F.1.
3. Execute the equilibrium solver across the grid points defined for each figure.
4. For hysteresis plots, perform both forward (increasing parameter) and backward (decreasing parameter) sweeps using the continuation method.

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 λ_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.

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.

Table 5: Robustness Summary: Persistence of Main Theoretical Forces

Specification	**Complexity trap persists?**	**Hysteresis persist**
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 Fig 2) remain significant under the alternative specification. "Stronger" indicates the hysteresis zone widens; "Weaker" indicates zone narrows but remains strictly positive. Detailed plots for each specification are available upon request.

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.