Smart Contracts, Dumb Money: Open Source Lemons*

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

First version: January 2026
This version: July 2026

Abstract

The most transparent asset market in history mints more new assets in a day than its largest auditor has examined in a lifetime. We show that attention rather than disclosure governs verification. Investors turn public code into noisy quality signals, and scarce attention admits only some projects into a consideration set where audits and funding can occur. Complexity keeps a project from competing for attention on quality, so in hot markets sound but complex projects fall out of the queue and their financing collapses. A constrained planner in this trap gains nothing from subsidizing entry or cheaper audits; only interventions that make projects processable reach the binding margin, and the trap destroys roughly a third of market surplus at the 2024 calibration. Around the January 2024 pump.fun launch, audit adoption among newly listed complex Solana protocols fell twenty-one percentage points against Ethereum-side peers while venture composition held still, so verification deteriorated before funding, as the model predicts.

JEL Classification: D83, G11, G14, G23, O31
Keywords: Rational Inattention, Decentralized Finance, Complexity Trap, Costly State Verification, Consideration Sets

*We extend our gratitude to Ruoran Lai (Sun Yat-sen University) for stimulating discussions and assistance during the writing process.
We also thank Bank of Italy ToDeFi, FTG Summer Conference(INSEAD), CCER summer institute, Tsinghua University PBC School of Finance and the ACM Conference on Economics and Computation(ACM EC) reviewers and participants 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.
†Adrasteia Labs; UIUC. Email: tingyi3@illinois.edu (**Corresponding author**)
‡University of Chicago. Email: weiyu@uchicago.edu
§Boston College. Email: yulin@bc.edu
¶University of Wisconsin-Madison. Email: hzhang2488@wisc.edu

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

— *Herbert A. Simon (1971)*

1 Introduction

Crypto-asset markets disclose more raw information than most financial markets. Trades are recorded on public ledgers, many contracts are open source, and token rules are often posted before launch. Yet liquidity often moves first toward assets that can be understood from a ticker, image, social-media feed, or community story. Protocols whose value depends on contract interfaces, governance permissions, oracle dependencies, bridge risk, or upgrade rights can be much harder to finance. The puzzle is why public code and public ledgers do not by themselves produce the screening that standard disclosure logic would suggest. This paper argues that the screening runs through attention. Public code and ledgers are abundant; the attention needed to make sense of them is not. Investors process public code into noisy signals, and scarce attention decides which of those signals enter an evaluation queue in which a project can be audited, funded, or built out. What a salience flood takes from complex projects is a place in that queue.

The friction is the work required to turn raw disclosure into a signal an investor can act on. Reading a smart contract is different from understanding its economic logic, attack surface, upgrade authority, liquidity design, and interactions with other protocols. That work does not scale like token issuance. A meme token can be evaluated by a coarse rule of thumb; a protocol with real technical upside may require specialist review before an investor knows what the downside states look like. This is the paper's *processability* margin.

Neither phrase in the title should be read as a behavioral claim. The "dumb money" investors are fully rational and allocate attention optimally given what processing disclosure costs them; the "lemons" arise with nothing hidden at all, since what binds is the processing of public information, abundant yet costly to read (Akerlof, 1970; Hirshleifer and Teoh, 2003; Blankespoor, deHaan, and Marinovic, 2020). This is a narrower claim than a general transparency paradox. Cheap-to-process assets, when their supply is large, degrade the market's ability to screen the hard-to-process ones.

We develop a static equilibrium model of the launch stage, before a liquid secondary-market price exists, in which two frictions are distinct by design. The first is a *comprehension ceiling*: human processing turns a project's public disclosure into a noisy signal of its quality, and the noise grows with the project's complexity. The second is *attention scarcity*: investors choose which processed signals to act on under a rational-inattention (RI) cost in the spirit of Sims (2003) and Matějka and McKay (2015). Complexity caps how informative processing

can be; the RI cost prices how much of that information is used, so no attention strategy can discriminate fundamentals more finely than the comprehension channel permits. A costly-state-verification (CSV) stage then determines whether an attended project is audited and funded (Townsend, 1979; Gale and Hellwig, 1985). The RI stage produces *consideration*, the probability that a project reaches the evaluation queue; nothing is bought at this stage. The verification stage inherits the processed-signal posterior, and attention values each project by the payoff of evaluating it, audit and walk-away options included, so the two stages share one information structure and one value function. Issuers choose legibility, which lowers complexity, and hype, which raises salience without improving fundamentals.

The analysis has three main results. First, being noticed is an equilibrium outcome with a price: a project receives attention if and only if its salience-adjusted score beats a benchmark set by the competing menu, and complexity attenuates exactly the component of the score that rewards quality, so a very complex project competes on salience alone, the dimension on which meme assets are strongest. Second, sound projects go unfunded precisely in hot markets, and only there: exclusion requires that the outside option itself receives no attention, and it requires a salience wedge in favor of simple assets, because complexity destroys a technical project's information advantage but not its level advantage. Third, invisibility is financial: an unconsidered project raises nothing, triggers no audit at any audit cost, and cannot recover fixed development costs, so sufficiently high complexity collapses technical entry; meme supply and salience raise the benchmark at the equilibrium, and with ecosystem externalities the collapse persists after primitives improve, because once few technical projects are active, future ones are still harder to evaluate.

The policy payoff is an intervention ordering, derived from a constrained planner who faces the same comprehension ceiling, attention technology, and audit cost as the market, and whose gap to private entry incentives decomposes into an appropriability wedge, a congestion wedge, and a salience wedge whose welfare status depends on whether hype utility is experienced or only decision-relevant. In the trap regime, where complex projects are excluded from consideration, even the planner attaches no value to marginal technical entry or to marginally cheaper audits, because neither instrument reaches projects that investors cannot process; what acts on the binding margin is the class of interventions that restores processability (standardization, machine-readable disclosure, certification), and because attention values a project by what evaluating it is worth, verification made cheap enough to restore that value belongs to the same class. Quantified at the moment-matched 2024 state, the trap is expensive: the platform shock destroys 29 percent of transfer-clean surplus under the paper's welfare accounting, 37 percent valuing the market's realized behavior with hype counted as consumption, and 97 percent with hype counted as distortion, the reading in

which the market's own valuation metric barely registers the collapse.

The ordering is testable, and Section 5 tests it against the sharpest available shock to the salience margin: the January 2024 launch of pump.fun, which pushed the cost of creating and promoting a salience-intensive asset effectively to zero on one chain. Because Solana's aggregate activity boomed through the event, every margin is biased toward improvement, and the relative declines below are measured against that wind. Audit adoption among newly listed complex protocols with a Solana deployment, measured one year after listing from archived dashboard snapshots, falls by 25 percentage points from the 2023 cohorts to the 2024–25 cohorts, against 4 points for Ethereum-side peers: a relative decline of 21 points. The decline survives the conservative restriction to Solana-native protocols and is largest among structurally intricate entrants. Within a single audit firm working both sides of the market, audit intensity per newly listed complex protocol never recovers past 0.7 in the 2024–26 cycle from 1.4 to 1.7 times the Ethereum-side benchmark in the 2021–22 cycle, and a six-archive panel shows a 58 percent decline in the relative report rate, part of it reallocation toward a Solana specialist. Venture funding composition is initially unchanged, the technical share of Solana deals staying inside its historical band through the event half-year, while $153 million of memecoin presale capital arrives in a single month: the salience flood itself. Verification deteriorates first and funding conversion moves last, which is the model's ordering. The adoption contrast sits at the edge of conventional significance, and every margin moves against the wind of the Solana boom; because the shock occurs on one chain, Section 5 reports the full inference accounting and presents the evidence as jointly suggestive.

The contribution is the equilibrium ordering of market screening, a step beyond the obvious point that meme assets are easier to understand than complex protocols: processability determines the consideration set; the consideration set determines which projects can trigger verification; and verification governs financing conversion and entry. Because each link is characterized rather than assumed, the model yields testable implications that differ from resale-speculation stories, including the prediction that hyped assets are audited less conditional on visibility and that exclusion of sound complex projects is concentrated in hot markets.

On the information side, the paper builds on rational inattention (Sims, 2003; Matějka and McKay, 2015) and, most directly, on Caplin, Dean, and Leahy (2019), who show that rational inattention generates consideration sets through corner solutions in unconditional choice probabilities; we import that logic, add a comprehension ceiling that varies across options, and study its market-equilibrium and entry consequences. Stochastic-consideration and attention-manipulation models make related points outside the RI framework (Manzini

and Mariotti, 2014; Eliaz and Spiegler, 2011). Closest on the policy side, Persson (2018) shows in a sender-receiver setting that competitive information supply generates overload, disclosure mandates lose force, and format regulation becomes the operative instrument; the equilibrium here delivers that conclusion as one implication and adds market-level structure with no analogue in the communication setting: a menu-set attention benchmark, hot-market exclusion under a salience wedge, consideration-gated verification and entry, and a trap regime in which even the planner's subsidies are worthless. Relative to classical asymmetric-information trading models (Glosten and Milgrom, 1985; Kyle, 1985), the central friction is not that one side privately knows the truth but that publicly observable information is costly to process: the disclosure-processing-cost margin of Hirshleifer and Teoh (2003) and Blankespoor, deHaan, and Marinovic (2020), here embedded in a general-equilibrium attention market. Work on sentiment, attention-driven buying, category learning, salience, and information allocation is complementary (Barberis, Shleifer, and Vishny, 1998; Barber and Odean, 2008; Peng and Xiong, 2006; Hong and Stein, 1999; Bordalo, Gennaioli, and Shleifer, 2012; Veldkamp, 2011; Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016).

On the market side, the paper adds to work on complexity, obfuscation, and the social value of financial innovation (Henderson and Pearson, 2011; Gennaioli, Shleifer, and Vishny, 2012; Célérier and Vallée, 2017; Vokata, 2021; Allen and Gale, 1994; Tufano, 1989; Lerner and Tufano, 2011): an open-source environment reproduces a form of lemons logic without hidden seller information in the canonical sense of Akerlof (1970), because public information fails to discipline allocation when investors under-process it. The tokenomics and blockchain literature (Cong, Li, and Wang, 2021, 2022; Sockin and Xiong, 2023b,a; Biais, Bisière, Bouvard, and Casamatta, 2019; Huberman, Leshno, and Moallemi, 2021; Budish, 2018; Catalini and Gans, 2020; Liu, Tsyvinski, and Wu, 2022) mostly studies adoption, consensus, or speculation once a token is salient; we study the earlier margin at which a token becomes visible and comprehensible enough to enter the evaluation set, which also connects the mechanism to strategic communication and information compression (Crawford and Sobel, 1982).

The rest of the paper proceeds as follows. Section 2 maps the institutional features of crypto markets to processability, salience engineering, and triggered verification. Section 3 presents the model. Section 4 characterizes the attention rule, the consideration set, crowding, the complexity threshold, and the verification stage. Section 5 tests the mechanism's ordering around the pump.fun launch. Section 6 presents numerical illustrations of the baseline complexity threshold; the hysteresis and verification-alpha extensions appear in Online Appendix OA.F. Section 7 states the planner problem and the wedge decomposition. Section 8 draws the mechanism-design implications. Section 9 concludes.

2 Institutional Background and Stylized Facts

Crypto markets make disclosure unusually abundant, but they do not make interpretation cheap. Three empirical regularities motivate the model. First, "fair launch" platforms like *pump.fun* on Solana collapsed the cost of issuance: by mid-2024, creating an instantly tradable token took under a minute and cost the issuer essentially nothing, with even the small deployment payment shifted to the first buyer (Pump.fun Help Center, 2025) (The Block, June 2024; The Block, August 2024; DefiLlama)[1]. A bonding-curve launch mechanism with integrated livestreaming gamified issuance; in the model's language, the platform pushed the fixed cost of meme creation, K_M, toward zero while raising the attainable hype margin, h. The attention economy saturated at once: on July 1, 2024 alone, 11,528 tokens were deployed on the platform (CoinDesk, July 2024). Second, the economic evaluation of technically ambitious projects is lumpy and specialist-intensive: reading public code is not the same as understanding security, governance, or implementation risk. Third, audit capacity is scarce and endogenous: verification typically occurs only for assets already visible enough to justify the due-diligence expenditure. These regularities map into the model's fixed-cost asymmetry, processability wedge, and triggered-verification stage.

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

Open-source contracts do not eliminate information frictions; they relocate them. In classical rational-expectations settings, public information is treated as cheaply incorporable into prices or common knowledge (Grossman and Stiglitz, 1980; Milgrom and Stokey, 1982), and the public-ledger architecture of Bitcoin and Ethereum makes raw disclosure unusually abundant (Nakamoto, 2008; Buterin, 2014). The hard task in crypto is semantic: translating code, documentation, tokenomics, upgrade paths, and governance promises into a judgment about expected value requires scarce human capital, especially when a project's logic is rich, modular, and upgradeable. The model's complexity parameter captures how efficiently pub-

[1]Industry and data sources are cited by organization and date; none is an academic reference. URLs: The Block (theblock.co, June 17 and August 9, 2024); DefiLlama (defillama.com; api.llama.fi/protocols); CoinDesk (coindesk.com, October 15, 2020, January 9, 2023, and July 2, 2024); CoinGecko Research (coingecko.com/research, 2023, 2024, and 2026 reports); Coinpaper (coinpaper.com, August 2, 2024, data from Dune Analytics); Cryptonews (via finance.yahoo.com, December 29, 2025); Solidus Labs 2025 rug-pull report (soliduslabs.com); SolanaFloor (solanafloor.com, October 2025, data from Blockworks); Galaxy Research (galaxy.com, January 2025); Hacken (hacken.io, 2023 security report and audit service page); Halborn (halborn.com, August 14, 2024); CertiK (certik.com smart-contract-audit page); Sherlock (sherlock.xyz, February 2026); ConsenSys Diligence USDi audit (diligence.security, April 2025); OpenZeppelin CapyFi audit (blog.openzeppelin.com, July 2025).

lic information can be turned into a decision-relevant payoff index. The interpretation is close to work on obfuscation and product complexity in finance (Carlin, 2009; Carlin and Manso, 2011; Célérier and Vallée, 2017; Gabaix and Laibson, 2006) and to the accounting literature on disclosure processing costs (Hirshleifer and Teoh, 2003; Blankespoor, deHaan, and Marinovic, 2020), with one difference: in crypto, complexity is usually structural and only occasionally strategic, so radical transparency coexists with severe informational fragility.

2.2 Structural complexity and strategic obfuscation

The baseline treats high processability costs as structural: a project is hard to process because its code, governance, oracle dependencies, cross-chain exposures, or upgrade paths are genuinely difficult to evaluate. Crypto markets also contain strategic complexity (tangled contract logic, opaque token rights, misleading documentation, governance that hides control) and investors who observe only effective complexity cannot tell the two apart. The baseline studies the conservative case in which high complexity is not assumed malicious; allowing bad issuers to obfuscate lowers the pooled quality of the high-complexity pool and strengthens the adverse-selection channel (Online Appendix OA.G.1).

2.3 Cheap issuance and attention engineering: the salience margin h

When sophisticated evaluation does not scale, issuers have a strong incentive to compete on the dimension that does scale: attention. Meme assets compress the mapping from disclosure to narrative into a handful of low-dimensional cues: name, image, slogan, community affiliation, or social-media momentum. They are not "simple" because their contracts are always primitive; they are simple because investors can form a coarse action rule about them with very little processing. That attention-grabbing characteristics drive buying behavior is well documented for equity markets (Barber and Odean, 2008), and investor-attention proxies strongly forecast cryptocurrency returns (Liu and Tsyvinski, 2021). The model's salience term can accordingly be read in two ways: as experienced narrative utility, or as a decision weight in the sense of salience theory (Bordalo, Gennaioli, and Shleifer, 2012, 2013); Section 3.1 derives the term from attribute-level salience weighting, which nests both readings, and Section 7 shows how the welfare conclusions depend on which reading one adopts.

Template-based launch platforms drastically reduce the fixed and marginal costs of creating low-dimensional speculative assets: they compress K_M, make salience production cheap, and flood the market with options that are easy to notice, so cognitively cheap projects expand supply much faster than fundamentally rich projects can expand investor understanding.

The mechanism resonates with the token-financing environment of Howell, Niessner, and Yermack (2020) and Yermack (2017), and with signaling in the spirit of Spence (1973), except that here the signal does more than separate types; it also lowers the investor's cognitive burden.

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

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

Public audit reports illustrate how lumpy this margin is. A relatively contained 2025 Consensys Diligence review of USDi consumed 10 person-days over one week, while OpenZeppelin's 2025 CapyFi audit ran from June 25 to July 7 (ConsenSys Diligence, 2025; OpenZeppelin, 2025). Although these examples are only illustrative, they make the paper's institutional point: meaningful smart-contract review consumes scarce expert time and does not scale one-for-one with token issuance.

A project that never attracts enough visibility may never be audited, even if it would be worth auditing conditional on serious consideration. This is the logic of triggered costly state verification: monitoring is activated when the expected benefit exceeds a fixed cost (Townsend, 1979; Gale and Hellwig, 1985), and in market-design terms the candidate set is shaped by listing decisions, dashboards, research coverage, and due-diligence capacity. Verification is downstream of attention.

2.5 Scope and alternative mechanisms

The mechanism should be read as complementary to speculative explanations of meme-asset demand, not as a substitute for them. Greater-fool logic explains why investors may buy a salient asset once they expect resale demand. Entertainment value, social coordination, and community identity can also make meme assets attractive even when fundamentals are weak. The mechanism in this paper addresses an upstream question: why some assets become salient enough to enter the evaluation set while technically valuable projects may remain

invisible. Narrative utility, resale speculation, and social coordination do not contradict the salience side of the model; they amplify it.

Earlier ICO waves are the same mechanism at smaller scale: projects competed on compressed narratives when detailed technical evaluation was costly. The launch-platform environment changes only the scale; the logic is the same.

Nor does the model require retail meme buyers to be the direct financiers of early-stage technical projects. Public-market visibility affects exchange listings, analyst coverage, audit demand, governance-token liquidity, and expected exit opportunities, all of which shape financing terms and the amortization of costly due diligence. The claim is not that a dollar spent on a memecoin would otherwise have funded a DeFi protocol; it is that meme-dominant attention markets weaken the public evaluation infrastructure through which complex projects become legible to retail and professional capital alike.

The model likewise contains no secondary-market price, by design. Its margin is the launch stage: whether a project becomes visible enough to reach the listings, research coverage, market-making, and due-diligence infrastructure through which price discovery later operates. The aggregation logic of rational-expectations models (Grossman and Stiglitz, 1980) presupposes an asset that already trades in a liquid market; it has no purchase on an asset no investor considers, because a price cannot aggregate evaluations that are never performed. For assets that do reach liquid trading, Grossman and Stiglitz (1980) itself implies that prices compensate only partially for costly information processing, so partial price aggregation downstream would attenuate but not remove the consideration gate studied here. The present model isolates the upstream margin; embedding partial aggregation for listed assets is the natural next step.

2.6 Observable implications

The processability channel makes five predictions that differ from a pure resale-speculation story; Section 5 takes implication 3 to data around the pump.fun launch and speaks to implication 5, and the others remain open. (1) Technical evaluation costs should reduce visibility most for projects without standardized audits, risk taxonomies, machine-readable disclosures, or credible certification, with candidate-set proxies such as exchange listings, dashboard inclusion, coverage, audit queues, and due-diligence shortlists. (2) Audit reports and machine-readable verification outputs should raise attention or liquidity more for complex protocols than for simple narrative assets, because complex protocols benefit more from compression. (3) Shocks to meme supply should first reduce coverage, dashboard inclusion, and audit demand for complex projects, with funding and liquidity responding later through

the conversion margin. (4) Conditional on visibility, audit propensity should be hump-shaped in beliefs and *decreasing* in hype: the rule in Proposition 5 funds low-belief projects not at all and high-belief projects on faith, and hype pads the perceived downside, so the loudest funded projects should be the least audited. (5) Exclusion of sound complex projects should be procyclical: it can occur only when the outside option receives no attention (Corollary 2), so hot markets are when invisible innovation should be most severe.

2.7 Suggestive evidence

The aggregate facts below complement the cohort- and firm-level designs of Section 5: they are correlations assembled from industry and academic sources that discipline the parameterization of Section 6.2; none identifies the mechanism by itself.

The salience margin is quantitatively dominant, and its pooled quality is extreme. Table 4 collects the magnitudes: issuance runs at thousands of new tokens per day against a bonding-curve graduation rate on the order of one percent, and forensic studies classify half to nearly all of the new supply as scams or pump-and-dumps (Xia et al., 2021) (CoinGecko Research, 2024; Coinpaper, 2024, from Dune Analytics; Solidus Labs, 2025). In the model's notation, the measured N_M/N_T and p_M sit far beyond any conventional calibration.

Verification is rationed away from hype (implication 4). Audit capacity cannot scale with issuance (the largest auditor's lifetime output is roughly one day of new-token issuance, Table 4), and rationing tightens exactly in hot markets: during the 2020–21 DeFi boom, leading audit firms were booked out for months and turned applicants away, and prominent protocols launched unaudited (CoinDesk, October 2020). Landsman, Lyandres, Maydew, Rabetti, and Zhang (2025) provide the systematic evidence on this voluntary audit market, in which audits act as certification and adoption responds to protocol complexity and security shocks. Conditional on failure, scrutiny and salience line up as the model predicts: only 6 percent of rug-pulled projects in 2023 had any audit (Hacken, 2024), and audited protocols account for only 20 percent of the top-100 DeFi exploits and 14.3 percent of their losses (Halborn, 2024). On the conversion margin, an audit is associated with roughly 37 percent higher locked value and 42 percent faster user growth in the first week after launch (Rabetti, 2023), consistent with verification operating on the conversion rate s_i while leaving visibility untouched.

Exclusion is concentrated in hot markets (implication 5). The outside option is abandoned in hot markets: stablecoins' share of total crypto market capitalization was 7.3

percent at the January 2022 bull-market peak, rose to 12.9 percent a year into the bear market, and peaked at 17.8 percent after the Terra and FTX collapses (CoinGecko Research, 2023). Attention composition swings with the same cycle: meme tokens' share of Solana DEX volume reached 67.5 percent at the late-2024 peak and fell to about 10 percent a year later (SolanaFloor, 2025, from Blockworks), and average daily memecoin trading volume rose almost ninefold between 2023 and 2024 (CoinGecko Research, 2026). Funding composition follows: consumer and narrative categories (Web3, NFT, metaverse, gaming) accounted for 31 percent of crypto venture deals at the 2022 peak and again led both capital and deal counts in late 2024, while infrastructure's share recovers in cold markets (CoinDesk, January 2023; Galaxy Research, 2025). These are composition facts, not identified effects: stablecoin balances also serve as trading collateral, and cycles move many things at once. But the joint pattern, cash abandoned while salience-intensive assets absorb attention and funding, is the regime in which Corollary 2 permits exclusion of fundamentally sound complex projects.

Audits are also noisy: 48 percent of exploit incidents between 2022 and 2026 hit protocols with at least one prior public audit (Beyer, 2026), so the perfect-verification benchmark is generous to audits, the conservative direction for the results (Section 4.9).

Section 5 implements the sharper test: the January 2024 launch of pump.fun as a chain-specific supply shock to the salience margin (a discrete fall in K_M to essentially zero) with dashboard inclusion, audit demand, and funding conversion traced around the event.

3 Model

Two frictions drive the model. Investors allocate attention under rational inattention, and due-diligence effort under costly state verification. The baseline is static. The path-dependence exercise in Section 6 adds an ecosystem externality as an extension; it is not required for the baseline propositions.

Timing. The timing is as follows.

1. A menu of projects is available. Each issuer chooses legibility m_i and a hype-seeding effort h_i. The effort contributes to an endogenous hype stock H_i, which is also reinforced by equilibrium visibility.

2. Human processing converts each project's public disclosure into a noisy signal whose accuracy falls with the project's complexity (the comprehension ceiling).

3. Given the processed signals, the investor allocates limited attention across projects

using a discrete-choice RI problem. The RI output is an endogenous visibility or consideration probability q_i; projects with $q_i = 0$ never reach the evaluation queue.

4. Conditional on consideration, the investor carries the posterior implied by the processed signal into the verification stage and decides whether to verify and whether to deploy capital.

5. Issuer financing and entry depend on both equilibrium visibility and conditional funding conversion.

The RI action is not an irreversible terminal purchase: it is a reduced-form representation of becoming salient enough to enter the investor's evaluation pipeline, and the two informational stages, attention and verification, share a single information structure.

3.1 Projects, types, and issuer choices

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

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

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

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

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

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

with $p_T \geq p_M$ in the baseline parameterization. Conditional on types, states are independent across projects in the baseline model; Online Appendix OA.G considers correlated states. The type labels are expository: Online Appendix OA.C.6 records that the paper's project-level results hold verbatim for arbitrary heterogeneous menus, with meme and tech read as regions of a characteristic space.

Payoffs. If the investor ultimately deploys capital to project i, realized utility depends on the state and on the project's accumulated hype stock H_i:

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

The fundamental component satisfies

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

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

Microfoundation of the salience term. The additive, state-independent form of ϕ_τ need not be assumed; attribute-level salience weighting delivers it. Let a project deliver two attributes: the fundamental payoff $R_\tau(\theta)$ and a state-independent narrative dividend $x_\tau(H) \geq 0$ (community membership, entertainment, participation) with $x_\tau' > 0$. An investor who overweights the attribute that stands out in the choice context, as in Bordalo, Gennaioli, and Shleifer (2013) and Kőszegi and Szeidl (2013), applies a decision weight $w_\tau(H) \geq 1$, nondecreasing in H, to the narrative attribute: hype makes that attribute vivid, while the fundamentals of a complex project remain pallid. Decision utility is then

$$R_\tau(\theta) + w_\tau(H)\, x_\tau(H) =: R_\tau(\theta) + \phi_\tau(H),$$

exactly the payoff above, while experienced utility is $R_\tau(\theta)+x_\tau(H)$. The two welfare readings of Section 7 are the poles of this decomposition: $w_\tau \equiv 1$ (hype is consumed) and $x_\tau \to 0$ with $w_\tau x_\tau$ fixed (hype is pure distortion), and an intermediate case splits ϕ_τ into an experienced part x_τ and a wedge $(w_\tau - 1)x_\tau$. Online Appendix OA.A.3 records the derivation and a belief-tilt alternative in which hype makes upside states vivid; the belief-tilt variant reproduces the same qualitative behavior and adds a second pooling force, so the additive baseline is the conservative specification.

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

1. **Legibility investment** $m_i \geq 0$, which lowers the effective difficulty of processing the project's public information;

2. **Hype-seeding effort** $h_i \geq 0$, which increases the project's hype stock H_i.

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

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

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

Issuer costs are

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

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

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

so that they can generate legibility cues and salience more cheaply. The numerical implementation adds a convex hype term $c_{\tau h2} h^2$ to (2); the convexity makes the meme sector's free-entry salience contest well defined, and the conduct rule it supports is stated formally in Online Appendix OA.E.5.

Obfuscation extension. The baseline sets strategic obfuscation to zero. A richer version lets a bad issuer choose $\omega_i \geq 0$ that raises effective complexity without improving fundamentals, $\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i + \omega_i)$, at a cost $c_{B\omega}\omega_i$ that may lie below the cost of genuine legibility or engineering quality: structurally complex innovation and strategically obfuscated junk then enter the same high-complexity pool. The main propositions set $\omega_i = 0$ to isolate structural complexity; Online Appendix OA.G.1 shows that allowing $\omega_i > 0$ strengthens the complexity trap by lowering the pooled prior quality of high-complexity projects.

Hype stock and attention feedback. A meme issuer does not merely buy a one-time advertisement: current visibility creates community activity, trending exposure, and future discoverability. The steady-state hype equation

$$\delta_{\tau_i} H_i = \xi_{\tau_i} h_i + \rho_{\tau_i} q_i, \qquad \delta_{\tau_i} \in (0, 1], \quad \xi_{\tau_i} > 0, \quad \rho_{\tau_i} \geq 0, \tag{3}$$

captures this, with ρ_τ the social feedback of today's visibility into tomorrow's salience, plausibly stronger for memes ($\rho_M > \rho_T$, $\xi_M \geq \xi_T$); the one-shot specification is the case $\rho_\tau = 0$, $\xi_\tau = \delta_\tau = 1$, in which $H_i = h_i$. The model is static: (3) is the steady state of the accumulation law $H_{i,t+1} = (1 - \delta_{\tau_i}) H_{i,t} + \xi_{\tau_i} h_{i,t} + \rho_{\tau_i} q_{i,t}$, a reduced form for persistent salience; we do not model it as an equilibrium object of a repeated game. Online Appendix OA.C.7 derives the comparative statics; at a locally stable steady state, feedback multiplies the marginal visibility return to hype effort by $\left[1 - \rho_\tau Q_i'(H_i)/\delta_\tau\right]^{-1} \geq 1$, so the one-shot benchmark

understates the return to hype under strong social feedback.

3.2 Investors: comprehension ceiling and rational inattention

Two frictions, two parameters. The investor side has two distinct frictions. First, a *comprehension ceiling*: human processing converts project i's public disclosure into a noisy signal y_i of its state, and the noise grows with the project's complexity κ_i. Second, *attention scarcity*: acting on the processed signals is itself costly, priced by a rational-inattention parameter λ that is *common* across options. Complexity caps how informative processing can be; λ prices how much of that information the investor actually uses. Because the investor's choice can condition only on the processed signals, no attention strategy can discriminate fundamentals more finely than the comprehension channel permits.

The comprehension channel. Let

$$U_i(\theta_i; H_i) := R_{\tau_i}(\theta_i) + \phi_{\tau_i}(H_i).$$

Conditional on the true state, the processed representation of project i's public disclosure is a binary signal $y_i \in \{g, b\}$ with accuracy

$$\Pr(y_i = g \mid \theta_i = G) = \Pr(y_i = b \mid \theta_i = B) = \rho_i(\kappa_i),$$

where $\rho_i(\kappa_i) \in [1/2, 1]$, $\rho_i'(\kappa_i) < 0$, and $\lim_{\kappa_i \to \infty} \rho_i(\kappa_i) = 1/2$. Signals are independent across projects conditional on states. Thus complexity does not remove public information; it degrades the informativeness of the human-processed representation of that information. The binary channel is the leading case: the results below hold whenever complexity indexes a family of channels decreasing in Blackwell informativeness (Remark 1 and Online Appendix OA.C.5). Legibility investment affects the channel through $\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i)$. Let

$$\mu_i(y_i; H_i, \kappa_i) := \mathbb{E}\Big[U_i(\theta_i; H_i) \mid y_i, \tau_i, \kappa_i\Big]$$

denote the posterior valuation implied by the processed signal, with values $\mu_i(g) \geq \mu_i(b)$.

The evaluation value. Bringing a project to the front of the queue is worth what the investor can do with it there: fund it on the processed posterior, pay for verification first, or walk away. Writing $p_i(y_i)$ for the posterior probability of the good state (formally derived

in Section 3.3), the value of evaluating project i at signal y_i is

$$\mathcal{W}_i(y_i) \;:=\; \max\Big\{\,\mu_i(y_i),\;\; p_i(y_i)\big(V^G_{\tau_i}+\phi_{\tau_i}(H_i)\big)-C_V,\;\; 0\,\Big\}, \tag{4}$$

the three branches being funding without verification (whose expected payoff is exactly $\mu_i(y_i)$), verifying at cost C_V and funding only in the good state, and the outside option. The baseline attention index is $\mathcal{W}_i$: the attention stage anticipates the verification stage, so the two stages share one information structure and one value function. $\mathcal{W}_i$ is a maximum of affine functions of the posterior, hence convex, nondecreasing, and nonnegative. The myopic variant that scores options by μ_i instead, ignoring the audit and walk-away options, is treated as a robustness specification (Online Appendix OA.G.2); every structural result below holds under both.

Attention over processed signals. Let $\boldsymbol{y}=(y_1,\dots,y_N)$ collect the processed signals. A representative investor chooses a single consideration action $a\in\mathcal{A}:=\{0,1,\dots,N\}$, where $a=i$ means that project i is brought to the front of the evaluation queue and $a=0$ denotes the outside option, normalized to $\mathcal{W}_0\equiv 0$. A strategy is a stochastic rule $\pi(a\mid\boldsymbol{y})$ with $\sum_{a\in\mathcal{A}}\pi(a\mid\boldsymbol{y})=1$ for all $\boldsymbol{y}$. Following Matějka and McKay (2015), the investor pays a Shannon/KL information-processing cost with the common scale parameter $\lambda>0$:

$$\mathcal{K}(\pi):=\lambda\,\mathbb{E}_{\boldsymbol{y}}\left[\sum_{a\in\mathcal{A}}\pi(a\mid\boldsymbol{y})\log\frac{\pi(a\mid\boldsymbol{y})}{q_a}\right],\qquad q_a:=\mathbb{E}_{\boldsymbol{y}}[\pi(a\mid\boldsymbol{y})]\,, \tag{5}$$

and solves

$$\max_{\pi(\cdot|\cdot)}\;\mathbb{E}_{\boldsymbol{y}}\left[\sum_{a\in\mathcal{A}}\pi(a\mid\boldsymbol{y})\,\mathcal{W}_a(y_a)\right]-\mathcal{K}(\pi),\quad\text{s.t.}\quad\pi(a\mid\boldsymbol{y})\ge 0,\;\sum_{a\in\mathcal{A}}\pi(a\mid\boldsymbol{y})=1. \tag{6}$$

The associated unconditional probability

$$q_i:=\Pr(a=i)=\mathbb{E}_{\boldsymbol{y}}\big[\pi(i\mid\boldsymbol{y})\big] \tag{7}$$

is the project's *ex ante visibility*. The induced state-contingent rule is the mixture $\pi(a\mid\boldsymbol{\theta})=\mathbb{E}[\pi(a\mid\boldsymbol{y})\mid\boldsymbol{\theta}]$; by construction it cannot be more state-sensitive than the comprehension channel allows. Heterogeneity in project difficulty therefore enters through signal quality $\rho_i(\kappa_i)$, while the RI scale λ is common to all options. All baseline propositions and all numerical exercises below maintain this specification.

The processable payoff index as a derived object. For state-level bounds it is convenient to define

$$v_i(\theta_i; H_i, m_i) := \mathbb{E}\Big[\mu_i(y_i; H_i, \kappa_i) \mid \theta_i\Big], \tag{8}$$

the mean posterior valuation conditional on the true state. For the binary symmetric channel, this index admits the closed-form representation

$$v_i(\theta_i; H_i, m_i) = \bar{U}_i(H_i) + A_i(\kappa_i, p_{\tau_i}) \Big[U_i(\theta_i; H_i) - \bar{U}_i(H_i)\Big], \tag{9}$$

where

$$\bar{U}_i(H_i) = p_{\tau_i} U_i(G; H_i) + (1 - p_{\tau_i}) U_i(B; H_i)$$

and

$$A_i(\kappa_i, p_{\tau_i}) = \frac{4 p_{\tau_i}(1 - p_{\tau_i}) r_i(\kappa_i)^2}{1 - (2p_{\tau_i} - 1)^2 r_i(\kappa_i)^2}, \qquad r_i(\kappa_i) := 2\rho_i(\kappa_i) - 1. \tag{10}$$

The attenuation factor satisfies $A_i \in [0, 1]$, $A_i \to 1$ as processing becomes perfect, and $A_i \to 0$ as processing becomes uninformative. Because $\phi_{\tau_i}(H_i)$ is state independent,

$$v_i(\theta_i; H_i, m_i) = \bar{R}_{\tau_i} + \phi_{\tau_i}(H_i) + A_i(\kappa_i, p_{\tau_i})[R_{\tau_i}(\theta_i) - \bar{R}_{\tau_i}], \qquad \bar{R}_\tau = p_\tau V_\tau^G + (1 - p_\tau) V_\tau^B.$$

Hype shifts the level; complexity attenuates state discrimination.

Relation to reduced-form specifications. Two natural shortcuts (RI over $\boldsymbol{\theta}$ with the derived index v_i, or a multiplicative discount u_i/κ_i) are reduced forms of this two-layer structure: the first replaces the signal distribution by its state-conditional mean, the second conflates level and state-sensitivity effects. The signal-level formulation is the primitive one, and the posterior implied by y_i is carried into the verification stage in Section 3.3, so the two stages share one information structure. Online Appendix OA.G.2 compares the specifications; the numerical implementation solves the signal-level formulation itself, drawing processed-signal vectors in the attention stage and carrying the signal posteriors into the verification stage (Online Appendix OA.F.2).

3.3 Costly state verification (CSV) stage

Consideration gates verification. Verification can be triggered only for projects the investor actually considers. The theoretical object is the endogenous *consideration set*

$$\mathcal{C}^* := \{i : \ q_i > 0\},$$

the support of the optimal attention rule, which Section 4 characterizes in primitives. Institutionally, scarce audit capacity restricts scrutiny further, to projects with sufficient visibility; the operational candidate set

$$\mathcal{C}(\varepsilon) := \{i : q_i \geq \varepsilon\}, \qquad \varepsilon > 0 \text{ small},$$

or equivalently a top-K rule, captures this capacity margin and is what the numerical implementation reports. Every theoretical result below uses only the upstream gate: a project outside the consideration set cannot be verified, and $\mathcal{C}(\varepsilon) \subseteq \mathcal{C}^*$ for every $\varepsilon > 0$.

Beliefs at the RI–CSV interface. The investor who considers project i knows her processed signal y_i. Because signals and states are independent across projects, the consideration event $\{a = i\}$ carries no information about θ_i beyond y_i (Online Appendix OA.D.1), so the belief entering the CSV stage is the Bayes posterior

$$\begin{gathered} p_i(y_i) := \Pr(\theta_i = G \mid \tau_i, y_i), \\ p_i(g) = \frac{p_{\tau_i}\rho_i}{p_{\tau_i}\rho_i + (1-p_{\tau_i})(1-\rho_i)}, \qquad p_i(b) = \frac{p_{\tau_i}(1-\rho_i)}{p_{\tau_i}(1-\rho_i) + (1-p_{\tau_i})\rho_i}. \end{gathered} \tag{11}$$

The two stages therefore share a single information structure. As $\kappa_i \to \infty$ the posterior collapses to the type prior p_{τ_i}: for very complex projects, having been noticed carries no news about quality, and verification decisions are made on the prior alone.

Verification technology. We begin with perfect verification as a benchmark. Verification reveals θ_i at a fixed cost $C_V > 0$: if the investor does not verify, she can fund the project on the posterior $p_i(y_i)$ or choose the outside option; if she verifies, she pays C_V and funds the project only when the verified state is good. This assumption gives the verification stage its strongest possible disciplinary power. The invisible-innovation result therefore does not rely on audits being weak. Section 4.9 and Online Appendix OA.D.2 replace the benchmark with noisy smart-contract audits; noisy verification weakens the downstream correction mechanism and strengthens the complexity trap.

3.4 Financing and entry mapping

To connect visibility to issuer incentives, we model financing as an attention-to-conversion pipeline; attention does not translate directly into investment. Visibility gives a project access to potential investors, but actual funding also depends on conditional screening and, when triggered, verification.

Let $g : [0,1] \to \mathbb{R}_+$ be a continuous increasing reach function with $g(0) = 0$. The term $g(q_i)$ converts the project's visibility into effective market reach. Let $s_i \in [0, \bar{s}]$ denote the conditional funding conversion rate after prior-based screening or verification. Aggregate expected funding is

$$F_i = \alpha g(q_i) s_i, \tag{12}$$

where $\alpha > 0$ scales market size. The linear specification $F_i = \alpha q_i$ is nested as the benchmark case $g(q) = q$ and $s_i = 1$.

A simple microfoundation is the following. A continuum of mass α of ex ante identical investors each solves the RI problem of Section 3.2, with independent draws of the processed signals and of the randomization in the choice rule. By a law of large numbers, q_i is then also the *fraction* of investors who consider project i, so the representative-investor probabilities and the market-level shares are the same object. Project i reaches an effective mass $\alpha g(q_i)$ through the attention pipeline, with $g(q) = q$ the frictionless case, and converts reached investors into funding with expected intensity s_i, defined formally in (27) from the CSV continuation values. Aggregating over the investor mass gives (12). One coordination friction is left out: an audit is non-rival information, so decentralized verification invites free-riding in the spirit of Grossman and Stiglitz (1980); modeling it would only strengthen the results, and single-audit accounting keeps them a lower bound.

Issuer payoff is therefore

$$\Pi_i = \alpha g(q_i) s_i - C_{\tau_i}(m_i, h_i), \tag{13}$$

and the issuer enters if $\Pi_i \geq 0$, choosing (m_i, h_i) subject to the processability technology (1) and the hype-stock condition (3). Together with the investor problem (6), this closes the feedback loop between attention allocation, verification, funding conversion, and market composition.

4 Equilibrium Characterization

Three objects characterize equilibrium behavior across the attention and verification stages: (i) the investor's RI-optimal rule over processed signals, (ii) the endogenous consideration set, characterized in primitives by an attention-score benchmark, and (iii) the issuer entry condition induced by consideration. We first derive the generalized logit over processed signals together with its concave-program representation, then formalize pooling, the consideration-set characterization and its crowding comparative statics, the complexity threshold for entry, and triggered verification. All proofs are in Online Appendices OA.B–OA.E.

4.1 Equilibrium objects and definitions

Let $\boldsymbol{y} = (y_1, \ldots, y_N)$ denote the vector of processed signals and $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_N)$ the vector of states. Recall that

$$U_i(\theta_i; H_i) = R_{\tau_i}(\theta_i) + \phi_{\tau_i}(H_i), \qquad \kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i)\,, \qquad \mu_i(y_i) = \mathbb{E}[U_i(\theta_i; H_i) \mid y_i],$$

that the evaluation value $\mathcal{W}_i(y_i)$ of (4) is the baseline attention index, and that the investor chooses a conditional rule $\pi(a \mid \boldsymbol{y})$ over $a \in \mathcal{A} = \{0, 1, \ldots, N\}$, where $a = i$ means that project i reaches the evaluation queue and $a = 0$ is the outside option. Define the ex ante visibility of project i and the outside option by

$$q_i := \Pr(a = i) = \mathbb{E}_{\boldsymbol{y}}\Big[\pi(i \mid \boldsymbol{y})\Big], \qquad q_0 := 1 - \sum_{i=1}^{N} q_i. \tag{14}$$

The consideration set and the operational candidate set are

$$\mathcal{C}^* := \{i : \; q_i > 0\}, \qquad \mathcal{C}(\varepsilon) := \{i : \; q_i \geq \varepsilon\} \subseteq \mathcal{C}^*, \tag{15}$$

with a top-K interpretation of $\mathcal{C}(\varepsilon)$ as an equivalent alternative. Define project i's *attention score* and the outside option's score by

$$Z_i := \mathbb{E}\Big[\exp\Big(\mathcal{W}_i(y_i)/\lambda\Big)\Big], \qquad Z_0 := 1. \tag{16}$$

Because $\mathcal{W}_i$ is a convex function of the posterior $p_i(y_i)$, the score is governed by the informativeness of the comprehension channel; two special structures follow. Writing $\mathcal{W}_\tau(p) := \max\{pW_\tau - (1-p)L_\tau,\; pW_\tau - C_V,\; 0\}$ with $W_\tau = V_\tau^G + \phi_\tau(H_\tau)$ and $L_\tau = -(V_\tau^B + \phi_\tau(H_\tau))$, salience raises $\mathcal{W}_\tau$ pointwise (at rate one where the project is funded outright, p where it is audited, zero where it is declined), so scores rise with salience. Under the myopic index μ the dependence is exact: the score factors as

$$Z_\tau^\mu(\kappa_\tau, H_\tau) = \exp\left(\frac{\phi_\tau(H_\tau)}{\lambda}\right) \widetilde{Z}_\tau(\kappa_\tau), \qquad \widetilde{Z}_\tau(\kappa_\tau) := \sum_{y \in \{g,b\}} P_\tau(y) \exp\left(\frac{m_\tau(y)}{\lambda}\right), \tag{17}$$

where $P_\tau(y)$ is the unconditional signal probability and $m_\tau(y) := \mathbb{E}[R_\tau(\theta) \mid y]$, so that salience shifts the myopic score multiplicatively. Under either index, complexity operates through the spread of the posteriors. Financing and issuer payoff follow (12), (27), and (13).

Lemma 1 (Complexity monotonicity of the score)**.** *Fix H_τ and suppose $V_\tau^G > V_\tau^B$, $\rho_\tau'(\kappa) < 0$, and $\mathcal{W}_\tau(p_\tau(g)) > \mathcal{W}_\tau(p_\tau(b))$ (the evaluation value separates the two posteriors, which holds*

whenever the project is funded or audited at the good signal). Then $Z_\tau(\kappa_\tau, H_\tau)$ is continuous and strictly decreasing in κ_τ, with

$$Z_\tau(\kappa_\tau, H_\tau) \downarrow Z_\tau^\infty(H_\tau) := \exp\left(\frac{\mathcal{W}_\tau(p_\tau)}{\lambda}\right) \qquad \text{as } \kappa_\tau \to \infty,$$

and $Z_\tau(\kappa_\tau, H_\tau) \geq Z_\tau^\infty(H_\tau)$ throughout, with strict inequality wherever $r_\tau(\kappa_\tau) > 0$ and the separation condition holds. The same statements hold for the myopic score (17) with limit $\exp\{(\bar{R}_\tau + \phi_\tau(H_\tau))/\lambda\}$.

(Proof: Online Appendix OA.C.3.)

The lemma is the score-level statement of the paper's central force. Complexity does not lower a project's *prior* appeal: the limit Z_τ^∞ retains the full value of evaluating the project on the prior alone, $\mathcal{W}_\tau(p_\tau) \geq \max\{\bar{R}_\tau + \phi_\tau(H_\tau), 0\}$. What it destroys is the premium $Z_\tau - Z_\tau^\infty$ that attention earns because the processed signal moves the evaluation: under the baseline index that premium is literally the option value a processed signal adds over the prior. A very complex project therefore competes for attention on level and salience alone.

4.2 Generalized logit over processed signals

Proposition 1 (Generalized logit, concave program, and support conditions)**.** *Fix issuer choices $\{(m_i, h_i)\}_{i=1}^N$, the induced hype stocks $\{H_i\}_{i=1}^N$, and hence $\{\kappa_i\}_{i=1}^N$, $\{\rho_i\}_{i=1}^N$, and $\{\mathcal{W}_i(\cdot)\}_{i=1}^N$. Under the Shannon/KL information cost (5):*

1. *The optimal unconditional probabilities solve the concave program*

$$\max_{q \in \Delta(\mathcal{A})} \lambda \, \mathbb{E}_{\boldsymbol{y}} \log\left(\sum_{a \in \mathcal{A}} q_a e^{\mathcal{W}_a(y_a)/\lambda}\right), \tag{18}$$

   whose value equals the investor's expected payoff net of the attention cost; a maximizer exists.

2. *Given an optimal q, the optimal conditional rule is, for all $i \in \{1, \ldots, N\}$,*

$$\pi^*(i \mid \boldsymbol{y}) = \frac{q_i \exp(\mathcal{W}_i(y_i)/\lambda)}{q_0 + \sum_{j=1}^N q_j \exp(\mathcal{W}_j(y_j)/\lambda)}, \tag{19}$$

   and

$$\pi^*(0 \mid \boldsymbol{y}) = \frac{q_0}{q_0 + \sum_{j=1}^N q_j \exp(\mathcal{W}_j(y_j)/\lambda)}, \tag{20}$$

   and q is consistent: $q_a = \mathbb{E}_{\boldsymbol{y}}[\pi^*(a \mid \boldsymbol{y})]$ for all $a \in \mathcal{A}$.

3. (Support conditions.) *Writing* $D(\boldsymbol{y}) := q_0 + \sum_j q_j e^{\mathcal{W}_j(y_j)/\lambda}$, *a feasible* q *is optimal if and only if*

$$\mathbb{E}_{\boldsymbol{y}}\left[\frac{e^{\mathcal{W}_a(y_a)/\lambda}}{D(\boldsymbol{y})}\right] \leq 1 \quad \textit{for all } a \in \mathcal{A}, \qquad \textit{with equality whenever } q_a > 0. \tag{21}$$

In particular, options can optimally receive zero *attention: corners of the program* (18) *are the endogenous consideration set, as in Caplin, Dean, and Leahy (2019).*

(Proof: Online Appendix OA.B.)

4.3 Pooling / decoupling for high complexity

Because the attention rule conditions on processed signals, never on the states directly, its state sensitivity is bounded by the informativeness of the comprehension channel. The bound is expressed by the same attenuation factor A_i that governs the derived index (9).

Corollary 1 (Pooling under noisy comprehension)**.** *Fix all issuer choices and unconditional weights, and let* $\pi^*(i \mid \theta_i, \boldsymbol{y}_{-i})$ *denote the choice probability induced by* (19) *conditional on project* i*'s true state. Then*

$$\left|\pi^*(i \mid G, \boldsymbol{y}_{-i}) - \pi^*(i \mid B, \boldsymbol{y}_{-i})\right| \leq \frac{A_i(\kappa_i, p_{\tau_i})}{4\lambda}\left|V^G_{\tau_i} - V^B_{\tau_i}\right| \qquad \textit{for every } \boldsymbol{y}_{-i}.$$

Consequently, as $\kappa_i \to \infty$ *and* $A_i(\kappa_i, p_{\tau_i}) \to 0$*, the conditional choice probability becomes state-blind. If* $r_i(\kappa_i) = O(\kappa_i^{-1/2})$*, the difference is* $O(1/\kappa_i)$*.*

(Proof: Online Appendix OA.C.1.)

The pooling result is a data-processing property of the comprehension channel: the choice rule cannot separate states that the processed signal barely separates.

Remark 1 (General comprehension channels)**.** *Nothing binary is essential. Let complexity index any family of signal structures ordered by Blackwell informativeness, so that higher* κ_i *makes the distribution of the posterior* $p_i(y_i)$ *a mean-preserving contraction around the prior, and read the attenuation factor as the normalized posterior variance* $A_i = \mathrm{Var}(p_i)/[p_{\tau_i}(1 - p_{\tau_i})]$*, which* (10) *equals exactly for the binary channel. Then the score monotonicity of Lemma 1, the bound of Corollary 1, and the threshold structure built on them hold for every such family and for every convex nondecreasing payoff index, including the continuation-value index of the baseline (Online Appendix OA.C.5).*

4.4 Denominator effect and attention crowding

Lemma 2 (Mechanical denominator effect (holding q fixed)). *Fix issuer choices of all pre-existing projects and the outside option, and hold their unconditional weights $\{q_a\}_{a\in\{0,1,\dots,N\}}$ fixed. After introducing a new option k with $q_k > 0$, we have for every signal vector $\boldsymbol{y}$,*

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

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

(Proof: Online Appendix OA.C.2.)

This is the mechanical crowding force; Proposition 3 below establishes the corresponding *equilibrium* statement, with the weights re-solved, on the margin that the entry results use.

4.5 The consideration set in primitives

The support conditions (21) deliver the paper's central equilibrium object without auxiliary assumptions. Because signals are independent across projects, the inclusion condition for any single project separates into the project's own score and a market-level benchmark.

Proposition 2 (Attention-score benchmark). *Let q be optimal and $D(\boldsymbol{y}) = q_0 + \sum_j q_j e^{\mathcal{W}_j(y_j)/\lambda}$. Define the equilibrium* attention benchmark

$$\bar{Z} := \left(\mathbb{E}_{\boldsymbol{y}}\left[D(\boldsymbol{y})^{-1}\right]\right)^{-1} \ge 1, \qquad \text{with } \bar{Z} = 1 \text{ whenever } q_0 > 0. \tag{22}$$

Then for every project i:

1. *if $Z_i < \bar{Z}$, then $q_i = 0$;*
2. *if $q_i > 0$, then $Z_i \ge \bar{Z}$; equivalently, $Z_i > \bar{Z}$ implies $i \in \mathcal{C}^*$ at every optimum up to boundary ties;*
3. *hence $\mathcal{C}^* = \{i : \ Z_i \ge \bar{Z}\}$ up to ties at $Z_i = \bar{Z}$: a project is considered if and only if its salience-adjusted attention score beats the benchmark set by the competing menu.*

(Proof: Online Appendix OA.C.3.)

Corollary 2 (Exclusion requires a hot market). *Every project satisfies $Z_i \ge \exp\{\mathcal{W}_i(p_{\tau_i})/\lambda\} \ge 1$, strictly whenever its evaluation has value at some signal. Hence if the outside option retains positive attention ($q_0 > 0$, so $\bar{Z} = 1$), every project whose evaluation value is positive with positive probability is considered: in particular every project that can be funded on its*

prior ($\bar{R}_{\tau_i}+\phi_{\tau_i}(H_i) \geq 0$) and every project whose audit window is nonempty at some posterior. A project worth evaluating can be excluded from the consideration set only when the outside option is itself abandoned ($q_0 = 0$): invisible innovation is a hot-market phenomenon, arising only when competing assets are attractive enough that holding cash commands no attention.

(Proof: Online Appendix OA.C.3.)

Corollary 3 (Sparse consideration: complexity-driven exclusion)**.** *Work in the symmetric-type environment and fix the meme side. Consider the meme-dominant allocation in which only the N_M meme projects receive attention, and let $\bar{Z}$ denote the benchmark (22) evaluated there; $\bar{Z}$ does not depend on κ_T. If*

$$Z_T^{\infty}(H_T) = \exp\left(\frac{\mathcal{W}_T(p_T)}{\lambda}\right) \;<\; \bar{Z}, \tag{23}$$

then there exists a finite threshold $\bar{\kappa}_T$, defined by $Z_T(\bar{\kappa}_T, H_T) = \bar{Z}$, such that for all $\kappa_T > \bar{\kappa}_T$ the meme-dominant allocation satisfies the optimality conditions (21) and tech projects receive zero attention: $q_T = 0$, so no tech project can trigger verification. If instead $Z_T^{\infty}(H_T) \geq \bar{Z}$, no level of complexity excludes tech projects. Exclusion at high complexity is therefore possible if and only if (23) holds, i.e., if and only if the value of evaluating a tech project on the prior alone falls short of the attention benchmark: $\mathcal{W}_T(p_T) < \lambda \log \bar{Z}$, where $\mathcal{W}_T(p_T) \geq \max\{\bar{R}_T + \phi_T(H_T),\, 0\}$.

(Proof: Online Appendix OA.C.3.)

4.6 Equilibrium crowding

The benchmark $\bar{Z}$ is an equilibrium object, and crowding-out is a comparative static of the benchmark, established at the fixed point.

Proposition 3 (Memes raise the exclusion benchmark)**.** *Consider the meme-dominant regime: N_M symmetric meme projects receive all attention, $q_0 = 0$, and tech projects are excluded, so that the optimal allocation is $q_M = 1/N_M$ per meme and $D(\boldsymbol{y}) = \bar{X}_{N_M} := N_M^{-1} \sum_{j=1}^{N_M} X_j$ with $X_j := e^{\mathcal{W}_M(y_j)/\lambda}$ i.i.d. and $\mathbb{E}[X_j] = Z_M$. This allocation is optimal if and only if $\mathbb{E}[\bar{X}_{N_M}^{-1}] \leq 1$ and $Z_T \leq \bar{Z}(N_M)$, where*

$$\bar{Z}(N_M) := \left(\mathbb{E}\left[\bar{X}_{N_M}^{-1}\right]\right)^{-1} \;\leq\; Z_M.$$

Moreover:

1. Supply channel. *$\bar{Z}(N_M)$ is nondecreasing in N_M, and $\bar{Z}(N_M) \uparrow Z_M$ as $N_M \to \infty$: in the deep-menu limit, a tech project is considered if and only if its score beats the meme score itself.*

2. Salience channel. *Raising meme salience raises every meme's evaluation value pointwise and with it $\bar{Z}(N_M)$: weakly always, strictly whenever considered memes are funded or audited with positive probability. Under the myopic index the increase is the exact factor $e^{\Delta/\lambda}$.*

3. Crowding-out. *Consequently the tech-exclusion region $\{\kappa_T : Z_T(\kappa_T, H_T) \leq \bar{Z}(N_M)\}$ weakly expands, and the exclusion threshold $\bar{\kappa}_T$ of Corollary 3 weakly falls, as meme supply N_M or meme salience rises.*

(Proof: Online Appendix OA.C.3.)

Remark 2 (Interior crowding)**.** *Proposition 3 establishes crowding on the extensive margin (meme supply and salience enlarge the region in which tech projects receive literally no attention), which is the margin the entry results below use. In the hot-market regime the interior margin is monotone as well: when $q_0 = 0$, the optimal attention allocation is unique, and aggregate interior tech attention falls strictly in meme salience and in meme supply, with the corner of the interior first-order condition reproducing the inclusion rule $Z_T \gtrless \bar{Z}(N_M)$ (Online Appendix OA.C.4). What remains computational is the response of interior tech attention to tech complexity itself, whose sign the same arguments do not determine.*

Corollary 4 (What the trap requires)**.** *Because $\bar{Z}(N_M) \leq Z_M$, the exclusion condition (23) implies the necessary condition*

$$\lambda \log Z_M(\kappa_M, H_M) \;>\; \mathcal{W}_T(p_T) \;\geq\; \max\{\bar{R}_T + \phi_T(H_T),\, 0\}\,: \tag{24}$$

the memes' certainty-equivalent evaluation value must exceed the value of evaluating a tech project on the prior alone. With meme fundamentals anchored poor (p_M small, $\bar{R}_M < 0$), the left side is positive only when meme salience is large; under the myopic index the condition specializes to $\phi_M(H_M) - \phi_T(H_T) > \bar{R}_T - \lambda \log \widetilde{Z}_M(\kappa_M)$, an explicit salience wedge. Complexity alone therefore cannot exclude a tech project: it destroys the project's information *advantage ($A_T \to 0$ in Corollary 1), but leaves the prior evaluation value intact, and only salience on the meme side can overturn that level. The complexity trap requires both ingredients, which matches the institutional reading of Section 2: cheap issuance floods the market with salience, and complexity prevents fundamentals from competing with it.*

(Proof: Online Appendix OA.C.3.)

4.7 Crowding-out and a complexity threshold for entry

Proposition 4 (Complexity threshold for tech entry)**.** *Suppose the investor follows the RI-optimal policy in Proposition 1, financing satisfies $F_i = \alpha g(q_i) s_i$, where g is continuous and increasing with $g(0) = 0$ and $0 \le s_i \le \bar{s}_T$ for tech projects, and issuers face costs (2) with $K_T > 0$.*

1. Entry collapse at the exclusion margin. *In the meme-dominant regime, if the exclusion condition (23) holds, then for all $\kappa_T > \bar{\kappa}_T$ every tech project has $q_T = 0$, hence $F_T = \alpha g(0) s_T = 0$ and $\Pi_T \le -K_T < 0$: no tech issuer enters, for* any *market size α,* any *reach function g, and* any *conversion rate s_T. The threshold is determined by primitives through $Z_T(\bar{\kappa}_T, H_T) = \bar{Z}(N_M)$.*

2. Entry failure at the reach margin. *Even when $q_T > 0$, tech entry fails whenever $\alpha g(q_T)\bar{s}_T < K_T$. If $g_1(q) \le g_2(q)$ for all q (pointwise ordering), the set of primitives supporting tech entry under g_1 is contained in the corresponding set under g_2; the entry region likewise weakly expands with α and $\bar{s}_T$ and weakly contracts with K_T.*

3. Comparative statics of the exclusion threshold. *$\bar{\kappa}_T$ is weakly decreasing in meme supply N_M and in meme salience $\phi_M(H_M)$ (Proposition 3), and weakly increasing in tech mean fundamentals $\bar{R}_T$ and tech salience $\phi_T(H_T)$. In terms of pre-investment complexity, the exclusion threshold for $\kappa_{T,0}$ equals $\bar{\kappa}_T\, e^{\beta_T m_T}$ and is therefore weakly increasing in the effectiveness of legibility investment β_T for any $m_T > 0$.*

(Proof: Online Appendix OA.E.3.)

The two margins in parts (1) and (2) are conceptually distinct and can bind in either order. The exclusion margin is the paper's headline mechanism: it is independent of market size and of the financing technology, and no downstream instrument reaches it. The reach margin is a conventional fixed-cost condition that depends on $(\alpha, g, \bar{s}_T, K_T)$. The numerical section reports which margin binds first under the chosen calibration.

4.8 Triggered CSV and invisible innovation

The investor who considers project i enters the verification stage with the posterior $p_i(y_i)$ of (11). For a generic belief $p \in [0,1]$, define the payoff from funding without verification and from verifying (funding only in the verified good state) as

$$U_i^{NV}(p) = p\Big(V_{\tau_i}^G + \phi_{\tau_i}(H_i)\Big) + (1-p)\Big(V_{\tau_i}^B + \phi_{\tau_i}(H_i)\Big), \tag{25}$$

$$U_i^{V}(p) = p\Big(V_{\tau_i}^G + \phi_{\tau_i}(H_i)\Big) - C_V, \tag{26}$$

with continuation values $W_i^{NV}(p) := \max\{U_i^{NV}(p), 0\}$ and $W_i^{V}(p) := \max\{U_i^{V}(p), 0\}$. Let

$$d_i(y_i) := \mathbf{1}\{q_i > 0\}\,\mathbf{1}\{W_i^{V}(p_i(y_i)) \geq W_i^{NV}(p_i(y_i))\}$$

indicate that project i is considered and optimally verified given the processed signal. Let $\ell : \mathbb{R} \to [0, \bar{s}]$ be a nondecreasing funding-intensity function with $\ell(w) = 0$ for $w \leq 0$. Averaging over the signal, the conversion rate is

$$s_i = \mathbb{E}_{y_i}\Big[\big(1 - d_i(y_i)\big)\,\ell\Big(W_i^{NV}(p_i(y_i))\Big) + d_i(y_i)\,p_i(y_i)\,\ell\Big(V_{\tau_i}^{G} + \phi_{\tau_i}(H_i)\Big)\Big]. \tag{27}$$

In the unit-investment cutoff case $\ell(w) = \mathbf{1}\{w > 0\}$ and the high-complexity limit $p_i(y_i) \to p_{\tau_i}$, this collapses to a prior-based conversion rule that depends only on type-level objects.

Proposition 5 (Triggered verification and invisible innovation)**.** *Fix issuer choices and the attention-stage equilibrium. Write $W_i := V_{\tau_i}^{G} + \phi_{\tau_i}(H_i) > 0$ and $L_i := -(V_{\tau_i}^{B} + \phi_{\tau_i}(H_i)) > 0$ for the good-state payoff and bad-state loss.*

1. Audit window. *For a considered project, verification is optimal at belief p if and only if $C_V \leq W_i L_i/(W_i + L_i)$ and*

$$\frac{C_V}{W_i} \;\leq\; p \;\leq\; 1 - \frac{C_V}{L_i}. \tag{28}$$

   Audits are triggered only at intermediate beliefs: low-p projects are not worth funding even if audited, and high-p projects are funded on faith. Because L_i falls with the hype stock H_i, the upper edge of the window falls with hype: conditional on being funded, more hyped projects are less likely to be audited. And because $p_i(y_i) \to p_{\tau_i}$ as $\kappa_i \to \infty$, complexity also destroys the informational content of consideration itself, so the audit decision for very complex projects is made essentially on the type prior.

2. Invisible innovation. *There exist primitives such that the audit window (28) contains the posterior of a tech project (verification would be strictly valuable conditional on consideration) yet the project receives zero attention ($q_T = 0$ by Corollary 3) and is never verified in equilibrium. The CSV stage cannot discipline projects filtered out upstream by the attention allocation.*

3. Financing consequence. *Under the funding rule $F_i = \alpha g(q_i) s_i$: if $q_i = 0$ then $d_i \equiv 0$, $F_i = \alpha g(0) s_i = 0$, and lowering C_V has no effect on project i: no audit-cost subsidy reaches a project that investors never consider. For considered projects, since $s_i \leq \bar{s}$ and $g(0) = 0$, any sequence with $q_i \to 0$ satisfies $F_i \leq \alpha g(q_i)\bar{s} \to 0$.*

(Proof: Online Appendix OA.D.1.)

4.9 Noisy verification and smart-contract audits

Smart-contract audits need not fully reveal composability, upgradeability, oracle, or governance risk. Let a considered project's audit return a report $\sigma_i \in \{g, b\}$ with symmetric accuracy $\xi_i \in (1/2, 1]$, after which the investor updates $p_i(y_i)$ and decides whether to fund; Online Appendix OA.D.2 derives the verification value $W_i^V(\xi_i)$, with perfect verification the case $\xi_i = 1$.

Corollary 5 (Noisy verification strengthens invisible innovation)**.** *For a fixed project and belief, the value of verification is weakly increasing in audit precision ξ_i. Hence replacing perfect verification with noisy verification cannot enlarge the set of projects for which verification is privately valuable. Moreover, because verification remains gated by consideration, any project excluded from the consideration set under perfect verification remains excluded under noisy verification. If audit precision is decreasing in effective complexity, $\xi_i = \xi(\kappa_i)$ with $\xi'(\kappa_i) \leq 0$, then complexity weakens both the probability of reaching the audit stage and the informativeness of the audit conditional on reaching it.*

(Proof: Online Appendix OA.D.2.)

4.10 Discussion: mapping the results to market outcomes

Complexity destroys the option value of evaluating a project with information rather than on its prior, not the prior value itself (Lemma 1), so attention to complex projects becomes state-blind (Corollary 1); a project is considered exactly when its score beats the equilibrium benchmark (Proposition 2), exclusion of sound projects requires a hot market and a salience wedge (Corollaries 2 and 4), and meme supply and salience push the benchmark up (Proposition 3). Exclusion becomes an entry collapse through the fixed cost (Proposition 4), and verification cannot rescue what attention filters out, while hype narrows the audit window for funded projects (Proposition 5). In the deep-menu limit the deterioration and recovery thresholds are explicit in primitives and hysteresis follows from the meme sector's conduct feedback alone (Online Appendix OA.E.4). The baseline is conservative in two respects, structural rather than malicious complexity and perfectly precise audits, and both obfuscation and audit noise strengthen the trap.

5 Evidence: The pump.fun Salience Shock

This section takes the mechanism's ordering (consideration upstream, verification in the middle, funding conversion last) to the sharpest available shock to the salience margin: the launch of pump.fun in January 2024, described in Section 2, which pushed the cost of creating and promoting a salience-intensive asset effectively to zero on Solana while leaving other chains untreated for roughly six months, until clone platforms spread. Solana-exposed protocols (Solana among deployment chains in the first available full archive capture containing the protocol) are compared with Ethereum-side peers (Ethereum, Arbitrum, or Optimism, and no Solana); Base and BNB are excluded from the control group because both received meme-platform waves of their own in mid-to-late 2024. The Solana-exposed definition follows the location of the shock and includes protocols that also deploy on Ethereum; Online Appendix OA.J re-imposes the Solana-native restriction. Two features discipline the interpretation. First, the confound runs against the mechanism: Solana's aggregate activity, fee revenue, and developer entry boomed in 2024, pushing listings, audit budgets, and funding *up*, so relative declines are measured against the wind. Second, the model predicts different movements on different margins, and the margins are observed separately, so the pattern across Tables 1–3, not any single contrast, is the object of interest.

5.1 Protocol-side verification

Table 1 measures audit adoption among newly listed complex protocols: the conversion of consideration into verification in Proposition 5, aggregated over all audit firms. The outcome is measured in event time: each cohort's audit status is read from an Internet Archive capture of the dashboard taken about one year after the cohort's listing window, so cohorts are compared at a common exposure age; a same-day snapshot would instead mix ages, and the lag is identical across chains within every cohort by construction. Deployment-chain and category labels follow the same measured-at-listing principle. We take them from the first full archive capture that contains each protocol, so a project's group, like its audit status, is fixed near its listing date and cannot drift with the current cross section. Adoption among Solana-exposed cohorts falls by 25.3 percentage points from the 2023 cohorts to the 2024–25 cohorts, against 4.3 points on the Ethereum side. The design is the protocol-level linear probability model

$$\text{audited}_i = \gamma\,\text{Sol}_i + \tau\,\text{Post}_{c(i)} + \beta\,\text{Sol}_i \times \text{Post}_{c(i)} + \varepsilon_i, \tag{29}$$

where i indexes newly listed complex protocols, $c(i)$ is the listing cohort half-year, audited$_i$ is the archived twelve-month audit status, Sol$_i$ is Solana exposure under the earliest-capture label, and Post$_{c(i)}$ marks the 2024–25 cohorts. Table 1 reports the cells and the estimates: the decline is about a third of the Solana baseline, sits at the edge of conventional significance under the saturated model, and attenuates moderately under cohort and category effects. The Solana-minus-control gap is nonnegative in both pre cohorts and negative in every post cohort (+8.2 and 0.0, then −17.9, −16.5, −5.9, −20.4 points). The extensive listing margin moves the other way (new complex Solana-exposed listings rise from 14 in 2023H2 to 31 in 2024H2) as the market-size margin of Proposition 4 predicts in a boom; the diagnostic margin is verification conditional on entry, not entry itself.

The decline also sits where the model puts it, in the structurally intricate tail, although the heterogeneity is not an independent main result. Extending (29) with a complexity indicator and its interactions, the decline among derivatives, options, bridges, synthetics, and restaking exceeds the decline among other complex protocols by 26.6 points, but seven intricate Solana protocols in the pre-period cell leave the contrast far from significance (directional $p = 0.136$); a direct-window version restricting to the 2024 cohorts and Ethereum controls reaches −54.1 points (directional $p = 0.036$, 0.072 after Holm correction across the two complexity proxies), a restriction selected from the logged grid and therefore exploratory. Oracle dependence separates the effect under the Solana-native restriction but not under the exposure definition, and a fork-based proxy is unusable because no Solana-native protocol in the sample is coded as a fork. Online Appendix OA.J tabulates the cells and the complete specification accounting.

5.2 Firm-side verification

Panel A of Table 2 holds the measuring instrument fixed: one audit firm (Zellic) that publishes month-dated, chain-tagged reports and works both sides of the market, so firm-level demand shocks difference out. Per ten newly listed complex protocols on the same side, Zellic audited Solana projects *more* intensively than Ethereum-side projects during the 2021–22 Solana cycle (intensity ratio 1.4–1.7). In the 2024–26 cycle the ratio never recovers past 0.7. Panel B removes archive-entry and archive-exit composition by retaining the six public archives observed continuously from 2022H1 through 2025H2: their Solana-to-Ethereum relative report rate falls 58 percent across the two cycles, marginally significant under the exact common-time permutation but not under firm clustering, and the exact result does not survive correction over the full logged search, so the panel supports the direction and market magnitude without replacing the firm-agnostic adoption result.

Table 1: Audit adoption of newly listed complex DeFi protocols one year after listing, before and after the salience-supply shock

Panel A: adoption shares

	2023 cohorts		2024–25 cohorts			
	audited/N	share	audited/N	share	change	DiD
Solana-exposed	17/29	58.6%	40/120	33.3%	−25.3 pp	
Ethereum-side	258/478	54.0%	258/519	49.7%	−4.3 pp	−21.0 pp

Panel B: the interaction coefficient in (29)

	(1) saturated	(2) cohort effects	(3) + category effects
$\hat{\beta}$ (Sol × Post)	−0.210	−0.198	−0.165
HC2 standard error	(0.107)	(0.108)	(0.108)
HC2 p	0.050	0.067	0.126
Directional permutation p	[0.034]	—	—
Two-sided permutation p	[0.052]	—	—
Observations	1,146	1,146	1,146

Notes: Unit: protocol newly listed on the DefiLlama dashboard, complex DeFi categories only. Solana-exposed = Solana among deployment chains in the first available full archive capture containing the protocol; Ethereum-side = Ethereum/Arbitrum/Optimism and no Solana. Audited = nonzero DefiLlama audit field in the Internet Archive capture closest to twelve months after the cohort midpoint (captures of 2024-09-09, 2025-03-06, 2025-10-03, 2026-05-02, and 2026-06-02; realized lags 11 to 17 months, 8 for the 2025H2 cohort, equal across groups within cohort). Panel B reports the Sol×Post coefficient of (29): column (1) is the saturated model, whose coefficient equals Panel A's difference-in-differences; column (2) replaces the Post indicator with cohort half-year effects; column (3) adds category effects; the logit analogue is negative in all specifications. The directional statistic uses the model-implied alternative $H_1 : \beta < 0$; both permutation p-values apply to the saturated statistic and use 100,000 draws that reshuffle protocol labels within cohort half-years. The p-values are specification-specific and unadjusted for the treatment-definition search. With one treated chain, neither HC2 nor protocol-label permutation provides chain-clustered inference. The same-date and current-label exposure definitions, the Solana-native restriction, and the full specification grid appear in Online Appendix OA.J. Strict audit coding gives a DiD of −20.3 points.

Reallocation is nevertheless in the data. Extending coverage to all ten public archives shows a Solana-specialist entrant, Offside Labs (founded 2023), publishing roughly four times Zellic's Solana-side output in the 2024–26 cycle, and Solana's listings-normalized audit intensity pooled over all ten archives is essentially unchanged across the two cycles. The specialist's book is dominated by repeat engagements for established Solana franchises: Meteora, Jupiter, Kamino, and Jito account for over sixty percent of its reports. The legacy-firm contrast therefore measures incumbent auditors' loss of new-protocol demand, a narrower object than the market total. The margin reallocation cannot explain is Table 1, which counts an audit by any firm and falls among newly listed complex protocols. Verification persists for flagships already inside the consideration set and thins for new complex entrants, which is the gate the model describes.

5.3 Funding conversion and the salience flood

Table 3 turns to the margin the model says moves last. Venture deal composition on the Solana side is essentially flat through the event half-year once memecoin *presales* are separated out: the technical share of Solana venture deals in 2024H1 is 64 percent, inside its 50–69 percent historical band. What changes is the arrival of a funding channel outside venture rounds altogether: 28 Solana memecoin presales totaling $153 million are recorded in 2024H1, essentially all in March 2024, against $62 million raised by all Solana technical rounds in the same half. The presale wave is treatment intensity rather than an outcome (capital chasing salience through a channel that requires no evaluation of fundamentals) while the venture channel, downstream of professional evaluation, had not yet moved.

5.4 What the evidence can and cannot say

Solana is a single treated chain, so conventional chain-clustered inference is unavailable. Table 1 therefore reports two protocol-level descriptive calculations, an HC2 standard error of 0.107 from the saturated linear probability model and protocol-label permutations within cohort half-years (directional $p = 0.034$, two-sided 0.052); neither supplies chain-level randomization. Chain-level corroborations point the same way within tight limits: Solana's post-period adoption gap relative to a donor-weighted synthetic is seven times the largest placebo gap among five donor chains, though Solana's noisy pre-period rules out the standard fit-ratio test, and a Solana-native 48-cell specification grid is negative in every cell with median -24.8 points, its only five-percent cells being restrictions chosen after inspecting where the effect concentrates. Online Appendix OA.J reports the complete inference accounting, including 432 logged synthetic-control variants, none below five percent, and the broader treatment, label-timing, window, control, category, outcome, and estimator search; the exact p-values in the tables are specification-specific rather than familywise claims. The main adoption contrast therefore reaches the five-percent rule under the model-implied directional statistic and sits just outside it two-sided; the broader evidential weight rests on signs and timing across independently measured margins, each biased against the finding by the Solana boom.

Chain-level trading intensity supplies a continuous check on the binary exposure contrast. A token-classified measure of launchpad-origin DEX volume, including graduated-token trading, shows the flood exactly where and when the design assumes it: zero everywhere before 2024, 28.6 percent of Solana DEX volume in 2024H2, and never above 2.1 percent on any other chain. With chain, listing-half-year, and category fixed effects, a ten-percentage-point higher pre-listing meme-volume share predicts a 9.2-point lower adoption rate (HC2

$t = -1.58$, directional $p = 0.057$, exact wild-score 0.075), the next-six-month look-ahead is null (-0.5 points, $p = 0.901$), and permuting the twelve observed exposure paths across chains is null. The gradient supports the binary exposure ordering but does not identify a separate chain-level dose response; Table 14 reports the construction, the sensitivities, and the logged multiplicity correction.

Five further limits are recorded. First, the archive pins outcome exposure age only to within a few months across cohorts (11 to 17, and 8 for 2025H2), while holding it equal across groups within cohort. Deployment-chain and category labels use the first available full capture containing each protocol, but full captures begin in September 2024, so the labels cannot be guaranteed pre-treatment for the 2023 cohorts. The estimate is stable when labels instead come from the outcome-date or current cross section, and excluding every Ethereum-straddling protocol under the Solana-native restriction gives -19.8 points (directional $p = 0.049$, two-sided $p = 0.081$); Online Appendix OA.J reports all variants. Second, public audit reports are client-approved, so reduced publication and reduced purchase are observationally similar; both are verification margins, but they differ for mechanism claims. Third, the design cannot separate the attention channel from a compositional shift toward entrants who would never have audited: conditioning on realized traction cannot arbitrate, because survival is post-treatment, and separating composition from attention requires pre-launch quality measures. The archived complexity splits hold one pre-launch dimension fixed, but the oracle split is sensitive to the exposure definition and quality composition within either class remains unobserved. Fourth, dashboard listing is itself a consideration proxy, so projects excluded from consideration never enter the dataset; the testable margins are verification and funding conversion among those already listed. Fifth, audit adoption is procyclical generally (Solana's 2022 cohorts also show low adoption), so separation of the platform shock from generic hot-market effects rests on Ethereum-side cohorts passing through the same 2024 bull market with much flatter adoption. The 2024H2 funding data are also right-censored in the archived registry snapshot. Finally, later meme-platform launches on other chains show no comparable adoption decline in a staggered design (Online Appendix OA.I); read through the model, this is what a market-level attention benchmark predicts, though the null equally admits a Solana-idiosyncratic reading.

6 Quantitative Illustration and Numerical Experiments

The numerical work has two parts. The first illustrates the model's core mechanisms under the empirically disciplined parameterization of Section 6.2; the second pins the remaining free parameters by minimum distance on three 2024 moments (Section 6.3). The figures visualize the central equilibrium object implied by the theory, the complexity-driven collapse of tech entry in the baseline model, together with the binding-margin diagnostics that separate exclusion from unrecoverable cost; the matching exercise locates the observed 2024 market inside that object and reproduces the observed audit-adoption decline without targeting it. Two further exercises, hysteresis in an extension with ecosystem externalities and the destruction of the private return to costly verification as salience production gets cheaper, appear in Online Appendix OA.F (Sections OA.F.5 and OA.F.6). Computational details, reporting conventions, and robustness checks appear in Online Appendix OA.F–OA.G.

Remark on the numerical implementation. Four conventions connect this section to the theory. First, two attention indices appear in fixed roles: the baseline evaluation value (4), under which attention anticipates the audit and walk-away options, is the theory's object and generates the figures; the myopic μ variant is the calibration's behavior model, because the baseline index admits only regime allocations at the 2024 interior moments (the fourth finding of Section 6.3); and the welfare exercise of Section 7.4 evaluates the calibrated states under both accountings. The division of labor is fixed across all exercises, and Online Appendix OA.G.2 compares the indices and quantifies the bounded-anticipation bridge between them. Second, the attention stage is solved as the concave program (18) in unconditional weights, so the consideration set is read off the corners exactly and certified by the Kuhn–Tucker residuals ($1 \pm O(10^{-3})$ for included options, strictly below one for excluded ones); no ε threshold enters. Third, two distinct margins can end tech activity: the exclusion margin of Corollary 3 ($q_T = 0$) and the reach margin of Proposition 4(2) ($q_T > 0$ but revenue below cost); the diagnostics accompanying each figure classify, point by point, which margin keeps the marginal tech group out. Fourth, the entry game has a mixed-strategy equilibrium at every grid point and a pure-strategy equilibrium under within-type symmetry (Online Appendix OA.E.5); with heterogeneous legibility groups purity can fail in a band of complexity levels, because attention competition makes entry decisions strategic substitutes while the meme sector's conduct response makes them complements, so the figures report continuation branches (a deterioration branch along which tech groups may only exit, and a recovery branch along which they may only re-enter).

6.1 Numerical environment and calibration choices

Heterogeneity within types. Both project types are heterogeneous, which removes the knife-edge allocations that full within-type symmetry would produce. The $\bar{N}_T = 50$ potential tech entrants come in $G = 10$ legibility groups of five projects; group g has pre-investment processability

$$\kappa_g = \kappa_{T,0}\, e^{\sigma_T z_g}, \qquad z_g \text{ deterministic normal-score nodes}, \quad \sigma_T = 0.4,$$

each group chooses its own legibility m_g with $\kappa_g^{\text{eff}} = \kappa_g e^{-\beta_T m_g}$, and a group is active when its profit is nonnegative. The $N_M = 5{,}000$ meme projects come in $G_M = 10$ salience classes generated by a virality lottery,

$$H_j = h_M\, e^{\sigma_M \zeta_j}, \qquad \zeta_j \text{ normal-score nodes}, \quad \sigma_M = 0.35,$$

where the common seeding effort h_M is chosen before the lottery. Because individual memes also compete for attention against one another, the salience contest dissipates rents: the sector's effort settles at the largest h_M with nonnegative expected per-meme profit given the convex hype cost $c_{Mh2}h_M^2$, consistent with the near-zero average meme profits in the data. Tech hype and meme legibility are fixed at zero in the baseline (their returns are negligible at the baseline salience and processability technologies) and relaxed in robustness. For numerical illustrations, we use the accuracy schedule $r_\tau(\kappa) = 2\rho_\tau(\kappa) - 1 = (1+\kappa)^{-1/2}$, which delivers an $O(1/\kappa)$ pooling rate. The appendix reports the equilibrium group-level legibility choices and the meme-sector effort alongside the outcome statistics.

Objects reported in the numerical section. Given equilibrium visibilities q_i, financing is $F_i = \alpha g(q_i) s_i$ with the conversion rate s_i computed from the CSV rule (27), and issuer payoff is $\Pi_i = F_i - C_{\tau_i}$. The baseline figures set $\rho_\tau = 0$ and normalize $\xi_\tau = \delta_\tau = 1$. We report three central outcomes:

1. **Tech entry** $N_T^{\text{eq}} \in \{0, 5, \ldots, 50\}$, the number of active tech projects (with $N_M = 5{,}000$ anchored, the count is more informative than the share, which is at most one percent by construction);

2. **Aggregate tech attention** $\sum_{i:\tau_i=T} q_i \in [0,1]$, the total attention absorbed by active tech projects;

3. **Binding-margin diagnostics**: at every grid point we test whether the most legible *inactive* tech group is kept out by the exclusion margin (zero attention at its best

legibility, Corollary 3) or the reach margin (positive attention but unrecoverable cost, Proposition 4(2)).

When a regime has no active tech entry, we explicitly label any reported tech quantity as counterfactual, since there is no equilibrium tech entry to report.

6.2 An empirically disciplined parameterization

The baseline parameter values, collected in Table 7 of Online Appendix OA.F.1, are simulation choices, not estimates. Discipline is nevertheless available: several model objects map to magnitudes directly observable in public data. Table 4 records, for each such object, the observable anchor with its source, and the relation to the baseline value used in the figures. Parameters without a credible anchor (the attention scale λ, the reach function g, the conversion rule ℓ) are treated as free and swept; the equilibrium salience level ϕ_M, the object behavior actually depends on, is bounded below by two anchored moments, and the technology parameter χ_τ mapping hype to it is swept.

The pinned parameters (K_M, p_M, the menu ratio) are set at their anchors, and where full anchoring is computationally impractical the remaining gap is conservative: the measured issuance ratio exceeds the implemented $N_M/\bar{N}_T = 100$ by at least an order of magnitude, and audits are noisier than the perfect-verification benchmark, so the crowding and exclusion forces in the figures are understated. Because payoffs are normalized, the anchors discipline relative scales: the audit-cost anchor pins verification as a lumpy fixed cost several orders of magnitude above the cost of issuing a meme token, which is the asymmetry the triggered-CSV stage encodes.

The salience side is disciplined the same way. Funding tokens whose anchored survival prior implies a mean fundamental loss reveals, under either reading of ϕ, a salience level at least the size of that loss, $\phi_M \gtrsim -\bar{R}_M = 0.45$; and the near-total absence of audits among failed tokens requires, through the audit window (28), $\phi_M \geq -V_M^B - C_V = 0.4$. The equilibrium values satisfy both bounds ($\phi_M \approx 0.97$ for the considered virality class), so the level of ϕ_M, the object behavior depends on, follows from the size of the fundamental downside. It is not a free parameter; only the technology χ_M mapping hype effort to that level remains a normalization.

6.3 A moment-matched parameterization

The anchors of Table 4 pin the parameters with direct observable counterparts and leave three objects free: the salience technology χ_M, the audit cost C_V, and the baseline complexity level

$\kappa_{T,0}$. This subsection chooses those three by minimum distance on three moments of the 2024 platform era, which disciplines the remaining freedom in the parameterization. The match runs under the myopic attention variant; under the baseline index the same protocol produces only regime allocations, which is itself the fourth finding below. The model state is the salience-cost path of Online Appendix OA.F.6: a pre state solved at the expensive-salience cost $c_{Mh2} = 0.04$ and a post state continued from it on the deterioration branch at $c_{Mh2} = 0.0025$, the platform era. The three targeted moments, all measured in 2024 data already used in this paper, are the meme sector's share of attention (matched to the meme share of Solana DEX volume at the late-2024 peak, 67.5 percent), audit adoption among newly listed complex Solana-exposed protocols (matched to the 2024–25 event-time cohorts of Table 1, 33.3 percent), and the funded share of meme projects (matched to the 2024 pump.fun graduation rate, 1.41 percent). The loss is the sum of squared relative deviations at the post state; the search runs on a parameter grid with local refinement, and the reported point is selected by re-solving the refinement neighborhood at the full sample-average resolution.

Four results carry the weight. First, the anchored model can occupy a state that quantitatively resembles the 2024 platform era: memes holding roughly three quarters of attention while complex projects remain considered, partially audited, and partially excluded, with relative errors of 10 to 30 percent on the three targeted moments. Second, the fit is out-of-sample in the dimension that matters most for the paper: neither the pre-shock adoption level nor the adoption decline was targeted, and both come out at the observed magnitudes (Table 5). The decline arises in the model exactly as Section 5 reads the data: cheap salience pushes the complex tail of the tech distribution below the consideration benchmark, the surviving entrants are the more legible ones, and verification contracts with them. Third, the matched complexity level sits close to the collapse boundary, so the exercise locates the observed market in the region where the trap mechanism is quantitatively live, not deep inside a safe interior. Fourth, the baseline index rearranges the landscape without improving the fit. Once attention anticipates evaluation, score dispersion compresses and the winner-take-all logic of the deep-menu limit (Online Appendix OA.E.4) already bites at finite menus. Every parameter point we tried then collapses to a regime allocation, a post-shock meme attention share of zero or one, never the interior mix the data show. Read through the consistent model, the 2024 interior composition is therefore informative about attention itself: it requires bounded anticipation of the evaluation option, quantitatively an anticipation weight below one to ten percent depending on complexity, and attention-scale heterogeneity does not substitute, because the baseline-index score ordering is invariant to λ on the matched menu; within the admissible band the moments cannot distinguish the anticipation

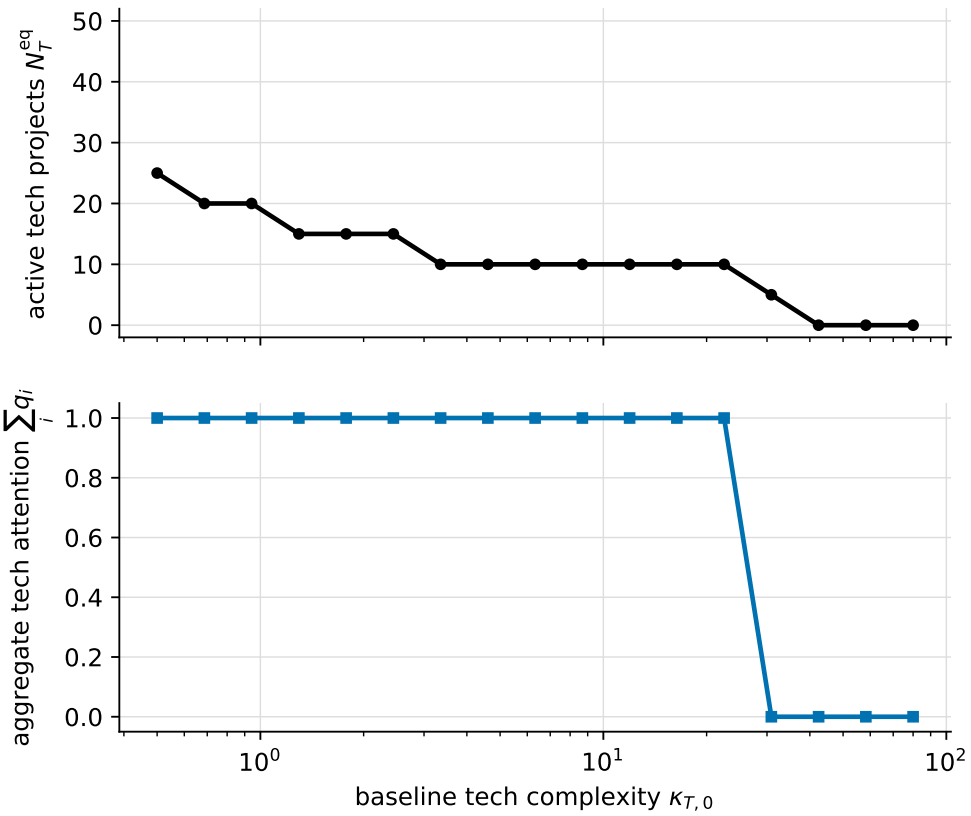


Figure 1: **Baseline complexity trap ($\eta = 0$).** Top panel: active tech projects N_T^{eq}; bottom panel: aggregate attention absorbed by tech projects. As $\kappa_{T,0}$ rises, tech groups defend consideration by paying for legibility until the defense cost exhausts their rents and they exit group by group; the meme sector's seeding effort, endogenous throughout, stays at zero while tech dominates attention and ignites late in the collapse. No ecosystem externality is imposed in this figure.

weight from zero (Online Appendix OA.G.2). The exercise remains below the standard of a full structural estimation, and the notes to Table 5 record the identification limits; what it removes is the free-parameter status of χ_M, C_V, and $\kappa_{T,0}$.

6.4 Figure 1: baseline complexity trap

Figure 1 illustrates the baseline complexity trap under $\eta = 0$, that is, without the ecosystem externality used later for hysteresis. We vary the baseline tech complexity parameter $\kappa_{T,0}$ on a logarithmic grid while holding all meme-side primitives fixed, tracing the deterioration branch from the tech-active state. For each grid point, we record tech entry N_T^{eq}, aggregate tech attention, the meme sector's endogenous seeding effort, and the binding-margin diagnostic.

Interpretation. The collapse combines the three forces of Section 4: pooling for high κ_T (Corollary 1), exclusion from the consideration set (Corollary 3), and the fixed-cost entry condition (Proposition 4). Both margins of Proposition 4 bind along the path: at several interior steps the marginal inactive group would be considered but cannot recover its legibility and fixed costs (reach margin), while at and beyond the meme ignition point it attracts zero attention at any feasible legibility (exclusion margin). The transition is a staircase rather than a cliff because groups exit in order of their legibility draws; full symmetry would produce a single discontinuous collapse. Entry falls in steps of one group ($25 \to 20 \to 15 \to 10 \to 5 \to 0$), and the meme sector's seeding effort ignites at $\kappa_{T,0} \approx 31$ ($h_M = 2.25$, the top virality class absorbing essentially all attention) and completes the collapse by $\kappa_{T,0} \approx 42$. The investor's fundamentals payoff (salience stripped) swings from $+0.59$ to -0.45 across the collapse while the decision payoff moves only from 0.59 to 0.52. The collapse is nearly invisible in the market's own valuation metric, the decision-versus-experienced gap of Section 7. The figure illustrates the threshold logic; it is not a calibrated estimate of the real market's tipping point.

6.5 What is deferred to the Online Appendix

Online Appendix OA.F reports the full parameter mapping, the concave-program solver and its Kuhn–Tucker certification, and the branch-tracing protocol, together with the two numerical extensions: hysteresis under an ecosystem externality (Section OA.F.5) and the collapse of expert verification alpha as salience production gets cheaper (Section OA.F.6). Online Appendix OA.G reports robustness to alternative information-cost specifications, correlated project states, alternative candidate-set definitions, posterior-based verification, and the N_M versus salience-cost paths for strengthening the meme side.

7 Welfare: A Constrained Planner Problem

A welfare criterion, a constrained planner problem, and a decomposition of the private-social entry gap set up the policy comparison of the next section. The main result is an intervention ordering, Proposition 7. All proofs are in Online Appendix OA.H.

7.1 Welfare accounting and the status of hype

The investor value function. Use the investor-mass microfoundation of Section 3.4. The welfare-relevant index for project i is the full continuation value of considering it, $\max\{W_i^V(p_i(y_i)), W_i^{NV}(p_i(y_i))\} \geq 0$, which nets out verification costs and embeds the option

to walk away; written out, this is exactly the evaluation value $\mathcal{W}_i(y_i)$ of (4), the baseline attention index. The positive model and the welfare accounting therefore share one program: let $\mathcal{V}(\mathcal{N})$ denote the value of (18) for menu $\mathcal{N}$, the representative investor's expected payoff net of attention and expected verification costs, with the scores Z_i and benchmark $\bar{Z}(\mathcal{N})$ of Section 4. (Under the myopic robustness index the welfare analysis re-solves the program at $\mathcal{W}$; all statements below are unchanged, Online Appendix OA.H.)

Aggregate welfare and transfers. Under the convention that u_i measures the investor's net-of-price return and F_i the issuer's gross proceeds, the funding flow is a transfer that appears once (negatively) inside investor payoffs and once (positively) in issuer profits, so the sum

$$\mathcal{W}^{\text{full}}(\mathcal{N}) = \alpha\, \mathcal{V}(\mathcal{N}) + \sum_{i \in \mathcal{N}} \Big(F_i - C_{\tau_i}(m_i, h_i) \Big) \tag{30}$$

is transfer-clean: it counts real project payoffs, real attention and audit costs, and real issuer costs exactly once.

Two readings of hype. The welfare status of $\phi_\tau(H)$ changes the welfare conclusions, not the positive model. Under the *experienced-utility* reading, narrative and community value are genuine consumption, ϕ belongs in (30), and meme dominance is not inefficient per se. Under the *decision-utility* reading, in the spirit of salience theory (Bordalo, Gennaioli, and Shleifer, 2012), ϕ distorts choice but is not experienced, and welfare should be computed from the same equilibrium behavior with ϕ stripped from realized payoffs; write $\mathcal{V}^{\text{exp}} \leq \mathcal{V}$ for that object. Under the attribute decomposition $\phi_\tau = w_\tau x_\tau$ of Section 3.1, the two readings are the poles $w_\tau \equiv 1$ and $x_\tau \equiv 0$; an intermediate case assigns experienced value $x_\tau = \phi_\tau / w_\tau$ and is bracketed by the two reported cases. We report which results depend on the reading; the intervention ordering in Proposition 7 does not.

7.2 The constrained planner and the entry wedges

The planner chooses the menu (entry) and issuer actions (m_i, h_i), and may operate the verification technology, but faces the *same* comprehension ceiling $\rho_i(\kappa_i)$, the same attention technology (5), and the same audit cost C_V as the market: attention allocation remains the RI-optimal response to the menu. This is the relevant benchmark for information policy. A planner who could rewire investors' processing capacity would trivialize the problem.

Proposition 6 (Entry wedges)**.** *Fix a menu $\mathcal{N}$ and issuer actions, and consider adding*

project i. The marginal social value of entry and the private entry incentive are

$$\Delta\mathcal{W}_i^{\text{full}} = \alpha\Big[\mathcal{V}(\mathcal{N}\cup\{i\}) - \mathcal{V}(\mathcal{N})\Big] + \Big(F_i - C_{\tau_i}(m_i,h_i)\Big), \qquad \Delta\Pi_i = F_i - C_{\tau_i}(m_i,h_i),$$

so the wedge is $\alpha\,\Delta\mathcal{V}_i := \alpha[\mathcal{V}(\mathcal{N}\cup\{i\}) - \mathcal{V}(\mathcal{N})]$. It decomposes as follows.

1. Appropriability wedge. $\Delta\mathcal{V}_i \geq 0$ *always, with* $\Delta\mathcal{V}_i > 0$ *if and only if* $Z_i > \bar{Z}(\mathcal{N})$*, i.e., if and only if the entrant is considered. The investor-side surplus a considered entrant creates (the option value of a better consideration set, including the option to audit it) accrues to investors, not to the issuer, so private entry is weakly too low for any considered project, under either reading of ϕ.*

2. Congestion wedge. *The entrant's effect on rivals' funding is a transfer and nets out of (30); congestion is already internalized inside $\Delta\mathcal{V}_i$, which is why $\Delta\mathcal{V}_i$ can be small even for a considered entrant when it mostly displaces close substitutes.*

3. Salience wedge. *Under the decision-utility reading, a marginal meme entrant with $K_M \simeq 0$ can have $\Delta\Pi_M > 0$ while the experienced-welfare change $\alpha\Delta\mathcal{V}_M^{\text{exp}}$ is strictly negative (hype attracts attention and funding toward assets whose experienced mean payoff is negative), so meme entry is strictly excessive. Under the experienced-utility reading this wedge is absent and meme entry is disciplined only by wedges (1)–(2).*

(Proof: Online Appendix OA.H.3.)

7.3 The trap regime and the ordering of instruments

Proposition 7 (Intervention ordering in the trap)**.** *Consider the meme-dominant regime in which tech projects are excluded from the consideration set: $Z_T \leq \bar{Z}$. Then, under either reading of ϕ:*

1. Entry subsidies are worthless at the margin. $\Delta\mathcal{V}_T = 0$*: adding a tech project changes neither attention allocation nor investor value, so a subsidy that induces tech entry spends K_T and raises welfare by exactly $\Delta\mathcal{W}_T^{\text{full}} = F_T - C_T(m_T,h_T) \leq -K_T < 0$. Even the planner, facing the same comprehension ceiling, attaches no value to marginal technical entry.*

2. Marginal audit subsidies do not reach the trap. *A marginal reduction in C_V raises $\mathcal{V}$ only through projects whose audit window (28) is active (all of them considered, hence meme projects in this regime) and has zero effect on any excluded project. The*

statement is about the derivative: because C_V enters the evaluation value (4), a non-marginal reduction raises excluded projects' scores and can lift them over the benchmark, in which case it operates as a processability intervention in the sense of part (3); Section 8 reports the counterfactual in which a fivefold cut produces exactly this flip.

3. Processability interventions act on the binding margin. *Any intervention that lowers effective tech complexity enough that $Z_T > \bar{Z}$ (standardization that lowers $\kappa_{T,0}$, technologies that raise β_T, certification that substitutes for private processing) makes $\Delta\mathcal{V}_T > 0$ strictly and restores strictly positive social (and possibly private) value of tech entry.*

Hence the constrained-efficient ordering of instruments in the trap regime is: restore processability first; entry and audit subsidies are ineffective until the exclusion margin is relaxed.

(Proof: Online Appendix OA.H.4.)

Part (1) locates the failure: in the trap, planner and market agree that an unprocessable project is worth zero, because neither can get investors to look at it. A subsidy cannot buy what investors cannot process. They can disagree only on entry of *considered* projects (the appropriability wedge) and on the volume of salience production (the salience wedge). Internalizing the audit free-riding margin (Grossman and Stiglitz, 1980), say through a certification utility, complements but cannot substitute for the processability interventions in part (3).

The ordering is not an artifact of the Shannon cost. Online Appendix OA.H.5 states it for an arbitrary attention-cost functional: whenever some optimum ignores a project, adding that project leaves investor value unchanged and a marginal audit subsidy has zero value on it, while any intervention that makes every optimum consider it raises value strictly. The Shannon structure delivers the explicit benchmark $\bar{Z}$ and the crowding comparative statics; the ordering itself is a property of consideration, not of entropy.

7.4 The cost of the trap at the matched point

How large is the trap? The matched parameterization of Section 6.3 lets us answer in welfare units. Evaluate (30) at the pre state (expensive salience) and the post state (the platform era, on the deterioration branch). When the investor program is re-solved at the evaluation-value index, as this section's accounting prescribes, the shock destroys 29 percent of transfer-clean welfare, or 0.34 per investor in units of the tech good-state payoff. Both readings of hype agree on this figure, because the re-solved consideration set holds no memes. The picture darkens if we instead value the behavior the matched myopic market actually delivers. Now

the loss is 37 percent of the pre-shock level when hype is experienced consumption, and 97 percent when it is pure decision weight: in that case almost none of the pre-shock surplus survives. The investors barely notice. Their own decision payoff falls by only an eighth, the collapse staying invisible in the market's own valuation metric, as Section 6 anticipates. About two thirds of the experienced loss sits in investors' fundamentals payoffs and one third in forgone tech issuer profits, of which the meme sector's contest profits recover only a fifth; funding flows net out as transfers, and the accounting residual is zero to machine precision. The numbers are robust to a small positive anticipation weight: under the bridge index of Online Appendix OA.G.2 at $\alpha = 0.01$, each moves by less than a percentage point.

8 Policy and Mechanism Design Implications

Policy acts on two margins in the model. Standardization and machine-readable formats determine whether complex projects become visible enough to matter. Audit cost and trigger design determine whether already-visible projects are worth investigating. Confusing these margins leads to the wrong instrument, and Proposition 7 makes the ordering precise: in the trap regime, processability interventions are the only instruments that act on the binding margin. Attention creates access to a funding opportunity; s_i captures the conditional conversion of that opportunity into capital after screening or verification.

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

In the attention stage, prices and allocations depend on how much of the disclosed environment investors can compress into an actionable rule, so requiring more raw disclosure need not help: longer whitepapers, more documentation, or larger code repositories add informational volume without adding usable signal unless they lower the processability wedge κ_T (Persson, 2018). In classical hidden-information problems the policy instinct is to reveal more; here the relevant question is whether the revelation technology reduces the cost of turning information into a low-dimensional decision statistic.

8.2 Lowering effective processability costs

The interventions that reach the binding margin raise the visibility of complex projects before the CSV stage. Standardized interfaces, risk taxonomies, and reporting templates reduce the dimensionality of the evaluation problem, lowering $\kappa_{T,0}$ or raising β_T. Machine-readable

audit layers (formal verification outputs, invariant checkers, standardized vulnerability taxonomies, API-readable disclosures) are compression technologies that make the same public information easier to use, not merely easier to access. Certifiers and research intermediaries compress high-dimensional technical information into low-cost verifiable claims; relative to the economics of signaling and communication (Spence, 1973; Crawford and Sobel, 1982), the intermediary matters here because it lowers the audience's processing burden, not only because it speaks.

8.3 Triggered verification infrastructure

Verification still matters for projects that reach the consideration set. Cheaper audits ($C_V \downarrow$) expand the set of visible projects worth investigating, the intensive-margin lever of classic CSV logic, and, because attention anticipates verification, sufficiently cheaper audits raise complex projects' attention scores directly; auditor liability, reputation, and competition make the triggered audit more informative; and event-based triggers (contract-permission changes, abnormal liquidity movements, governance shocks) partially bypass the attention bottleneck by forcing projects into the verification queue even when they are not salient.

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

The model contains an extensive margin, so policy comparisons should distinguish between partial and full welfare objects. The complete, transfer-clean welfare measure is $\mathcal{W}^{\text{full}}$ of (30) in Section 7, and the central distortion in the paper operates through its entry term: under exclusion, genuine tech projects cannot recover their fixed costs and exit, and Proposition 7 ranks the instruments that can and cannot repair this.

The policy table below is a complementary, *fixed-entry investor-side counterfactual*. It isolates the intensive-margin effects on already-active projects and does not claim to be a complete welfare ranking; the extensive-margin ranking is the content of Proposition 7.

Summary. What matters for the policy comparison is the margin an instrument reaches, not the name it goes by (Proposition 7): marginal audit subsidies act only inside the consideration set, and in the trap they act on nothing that matters. What works is the class of interventions that changes which projects enter the consideration set, whether by lowering the cost of processing disclosure (standardization, machine-readable formats) or by making verification cheap enough that evaluating a complex project is worth attention again, which is certification by another name.

9 Conclusion

Open code does not make a project easy to value. The model formalizes this point by placing attention before verification: investors first decide which projects enter an evaluation queue, and only then can audits or due diligence reveal quality. The consideration set is characterized in primitives (a project is noticed exactly when its salience-adjusted score beats an equilibrium benchmark set by the competing menu), and complexity destroys the component of the score that rewards quality, leaving complex projects to compete on salience alone.

This ordering gives simple narrative assets an advantage, but the model also disciplines when and how the advantage bites. Exclusion of fundamentally sound projects requires a hot market in which the outside option itself is abandoned, and it requires a salience wedge, not complexity alone. When both conditions hold, exclusion cuts financing to zero, fixed development costs go unrecovered, and tech entry collapses; the extension with ecosystem externalities shows that the loss can persist after primitives improve, because a thin technical ecosystem makes future protocols still harder to evaluate. The welfare analysis sharpens the diagnosis: in the trap, even a planner facing the same processing frictions gains nothing from subsidizing entry or audits, because neither instrument reaches projects investors cannot process, and at the calibrated 2024 state the trap costs a third of measured surplus, nearly all of it under the reading in which hype is distortion. The pump.fun evidence tracks the mechanism's ordering margin by margin: on the treated chain, audit adoption falls, most steeply among the most complex entrants, while venture composition initially stands still.

Disclosure helps when it reduces the cost of forming a valuation-relevant signal. In this setting, useful disclosure looks like standardized contract interfaces, clear governance permissions, machine-readable audit outputs, reusable audit templates, and certification layers that investors can check quickly. Public code does little if investors cannot cheaply decide what the code implies.

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.

Barber, B. M., and T. Odean 2008. All That Glitters: The Effect of Attention and News on

the Buying Behavior of Individual and Institutional Investors. *The Review of Financial Studies* 21(2), 785–818.

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

Beyer, S. 2026. The Audit Gap in Blockchain Security: A Four-Year Empirical Study of Public Audit Findings and Real-World Exploit Incidents. Working paper, arXiv:2606.15465.

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

Blackwell, D. 1953. Equivalent Comparisons of Experiments. *The Annals of Mathematical Statistics* 24(2), 265–272.

Blankespoor, E., E. deHaan, and I. Marinovic 2020. Disclosure Processing Costs, Investors' Information Choice, and Equity Market Outcomes: A Review. *Journal of Accounting and Economics* 70(2–3), 101344.

Bordalo, P., N. Gennaioli, and A. Shleifer 2012. Salience Theory of Choice under Risk. *The Quarterly Journal of Economics* 127(3), 1243–1285.

Bordalo, P., N. Gennaioli, and A. Shleifer 2013. Salience and Consumer Choice. *Journal of Political Economy* 121(5), 803–843.

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.

Caplin, A., M. Dean, and J. Leahy 2019. Rational Inattention, Optimal Consideration Sets, and Stochastic Choice. *The Review of Economic Studies* 86(3), 1061–1094.

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.

Cernera, F., M. La Morgia, A. Mei, and F. Sassi 2023. Token Spammers, Rug Pulls, and Sniper Bots: An Analysis of the Ecosystem of Tokens in Ethereum and in the Binance Smart Chain (BNB). In *Proceedings of the 32nd USENIX Security Symposium*.

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.

Eliaz, K., and R. Spiegler 2011. Consideration Sets and Competitive Marketing. *The Review of Economic Studies* 78(1), 235–262.

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.

Hirshleifer, D., and S. H. Teoh 2003. Limited Attention, Information Disclosure, and Financial Reporting. *Journal of Accounting and Economics* 36(1–3), 337–386.

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.

Kőszegi, B., and A. Szeidl 2013. A Model of Focusing in Economic Choice. *The Quarterly Journal of Economics* 128(1), 53–104.

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

Landsman, W. R., E. Lyandres, E. L. Maydew, D. Rabetti, and S. Zhang 2025. Auditing Smart Contracts. Working paper.

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., and A. Tsyvinski 2021. Risks and Returns of Cryptocurrency. *The Review of Financial Studies* 34(6), 2689–2727.

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.

Mancino, D. 2025. The Memecoin Phenomenon: An In-Depth Study of Solana's Blockchain Trends. Working paper, arXiv:2512.11850.

Manzini, P., and M. Mariotti 2014. Stochastic Choice and Consideration Sets. *Econometrica* 82(3), 1153–1176.

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.

Persson, P. 2018. Attention Manipulation and Information Overload. *Behavioural Public Policy* 2(1), 78–106.

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

Pump.fun. 2026. Fees (official documentation: creation free; first-buy, gas, and graduation fee schedule). Available at: https://pump.fun/docs/fees.

Rabetti, D. 2023. Auditing Decentralized Finance (DeFi) Protocols. Working paper, National University of Singapore. SSRN No. 4458298.

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.

Xia, P., H. Wang, B. Gao, W. Su, Z. Yu, X. Luo, C. Zhang, X. Xiao, and G. Xu 2021. Trade or Trick? Detecting and Characterizing Scam Tokens on Uniswap Decentralized Exchange. *Proceedings of the ACM on Measurement and Analysis of Computing Systems* 5(3), Article 39.

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

ONLINE APPENDIX

OA.A Assumptions, notation, and primitives

This appendix formalizes the baseline model used in the main text. The information structure has two layers. First, a comprehension channel converts each project's public disclosure into a noisy binary signal y_i whose accuracy is governed by the processability factor $\kappa_i(m_i)$. Second, the investor allocates attention *over the processed signals* under a Shannon/KL information cost with a *common* scale parameter λ. Project heterogeneity therefore enters through signal quality, not through option-specific attention costs, and no choice rule can be more state-sensitive than the comprehension channel permits. Specifications that place the RI problem directly over states (with the derived index v_i or with a multiplicative discount u_i/κ_i) are reduced forms of this structure. All main propositions maintain the common-λ, signal-level specification.

OA.A.1 Timing and objects

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

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

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

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

Legibility determines the processability factor

$$\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i), \qquad \beta_M > \beta_T \geq 0, \qquad \kappa_{T,0} > \kappa_{M,0} > 0. \tag{31}$$

We allow $\kappa_i < 1$ for unusually legible projects. Hype effort determines a market hype stock through

$$\delta_{\tau_i} H_i = \xi_{\tau_i} h_i + \rho_{\tau_i} q_i, \qquad \delta_{\tau_i} \in (0, 1], \quad \xi_{\tau_i} > 0, \quad \rho_{\tau_i} \geq 0. \tag{32}$$

Equivalently,

$$H_{i,t+1} = (1 - \delta_{\tau_i}) H_{i,t} + \xi_{\tau_i} h_{i,t} + \rho_{\tau_i} q_{i,t}.$$

The one-shot hype benchmark is recovered when $\rho_\tau = 0$ and $\xi_\tau = \delta_\tau = 1$.

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

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

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

Comprehension channel and attention-stage actions. Conditional on θ_i, the processed signal $y_i \in \{g, b\}$ has symmetric accuracy $\rho_i(\kappa_i) \in [1/2, 1]$ with $\rho_i' < 0$ and $\rho_i \to 1/2$ as $\kappa_i \to \infty$; signals are independent across projects given states. The investor chooses a single consideration action $a \in \mathcal{A} = \{0, 1, \ldots, N\}$. A stochastic rule is a conditional distribution

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

where $\boldsymbol{y} = (y_1, \ldots, y_N)$. Unconditional probabilities are

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

The payoff index at the attention stage is the evaluation value $\mathcal{W}_i(y_i)$ of (4), with $\mathcal{W}_0 \equiv 0$ for the outside option; the proofs in OA.B are written for a generic bounded own-signal-measurable index and so cover both the baseline and the myopic variant that uses the posterior valuation $\mu_i(y_i) = \mathbb{E}[U_i(\theta_i; H_i) \mid y_i]$.

Verification-stage timing. After the attention stage, verification can be triggered only for considered projects, $i \in \mathcal{C}^* = \{i : q_i > 0\}$; the operational candidate set $\mathcal{C}(\varepsilon) = \{i : q_i \geq \varepsilon\} \subseteq \mathcal{C}^*$ captures audit-capacity limits and is the object reported numerically. The belief entering the CSV stage is the processed-signal posterior $p_i(y_i)$; OA.D derives this and adds noisy verification. The type prior is the high-complexity limit of $p_i(y_i)$.

Financing primitives. Issuer financing follows

$$F_i = \alpha g(q_i) s_i,$$

where $g : [0, 1] \to \mathbb{R}_+$ is continuous and increasing with $g(0) = 0$, and $s_i \in [0, \bar{s}]$ is the conditional funding conversion rate. The linear benchmark is nested by $g(q) = q$ and $s_i = 1$.

OA.A.2 Noisy comprehension and the processable payoff index

This subsection derives the processable payoff index used in the main text. Fix project i and suppress the type subscript. Let

$$p := p_{\tau_i}, \qquad U_G := U_i(G; H_i), \qquad U_B := U_i(B; H_i), \qquad \bar{U} := pU_G + (1-p)U_B.$$

The investor processes public disclosure into a binary signal $y_i \in \{g, b\}$ with symmetric accuracy $\rho_i(\kappa_i)$:

$$\Pr(y_i = g \mid \theta_i = G) = \Pr(y_i = b \mid \theta_i = B) = \rho_i(\kappa_i).$$

Thus

$$\Pr(y_i = g \mid \theta_i = B) = 1 - \rho_i(\kappa_i), \qquad \Pr(y_i = b \mid \theta_i = G) = 1 - \rho_i(\kappa_i).$$

We assume $\rho_i(\kappa_i) \in [1/2, 1]$, $\rho_i'(\kappa_i) < 0$, and $\rho_i(\kappa_i) \to 1/2$ as $\kappa_i \to \infty$.

By Bayes' rule,

$$\Pr(G \mid g) = \frac{p\rho}{p\rho + (1-p)(1-\rho)}, \qquad \Pr(G \mid b) = \frac{p(1-\rho)}{p(1-\rho) + (1-p)\rho}.$$

Let $\mu_y = \mathbb{E}[U(\theta; h) \mid y]$. Then

$$\mu_g = \Pr(G \mid g)U_G + [1 - \Pr(G \mid g)]U_B, \qquad \mu_b = \Pr(G \mid b)U_G + [1 - \Pr(G \mid b)]U_B.$$

The processable payoff index conditional on the true state is

$$v(G) = \rho\mu_g + (1-\rho)\mu_b, \qquad v(B) = (1-\rho)\mu_g + \rho\mu_b.$$

Proposition 8 (Noisy comprehension generates state-sensitivity attenuation)**.** *Under the binary processing channel,*

$$v(\theta) = \bar{U} + A(\kappa, p)\big[U(\theta) - \bar{U}\big],$$

where

$$A(\kappa, p) = \frac{p(1-p)(2\rho(\kappa) - 1)^2}{[p\rho(\kappa) + (1-p)(1-\rho(\kappa))][p(1-\rho(\kappa)) + (1-p)\rho(\kappa)]}.$$

Equivalently, with $r(\kappa) = 2\rho(\kappa) - 1$*,*

$$A(\kappa, p) = \frac{4p(1-p)r(\kappa)^2}{1 - (2p-1)^2 r(\kappa)^2}.$$

Therefore

$$v(G) - v(B) = A(\kappa, p)(U_G - U_B).$$

Proof. First,

$$v(G) - v(B) = (2\rho - 1)(\mu_g - \mu_b).$$

Next,

$$\mu_g - \mu_b = [\Pr(G \mid g) - \Pr(G \mid b)](U_G - U_B).$$

Using Bayes' rule,

$$\Pr(G \mid g) - \Pr(G \mid b) = \frac{p(1-p)(2\rho - 1)}{[p\rho + (1-p)(1-\rho)][p(1-\rho) + (1-p)\rho]}.$$

Combining these expressions gives the formula for $A(\kappa, p)$ and the state gap. Finally, by the law of iterated expectations,

$$pv(G) + (1-p)v(B) = \bar{U}.$$

Since the mean is $\bar{U}$ and the state gap is $A(\kappa, p)(U_G - U_B)$, the affine representation follows. □

Lemma 3 (Complexity lowers state sensitivity)**.** *If $\rho'(\kappa) < 0$, then $\partial A(\kappa, p)/\partial \kappa < 0$ whenever $r(\kappa) > 0$.*

Proof. Write $A(p, r) = 4p(1-p)r^2/[1 - (2p-1)^2 r^2]$. Then

$$\frac{\partial A}{\partial r} = \frac{8p(1-p)r}{[1 - (2p-1)^2 r^2]^2} > 0$$

for $r > 0$. Since $r'(\kappa) = 2\rho'(\kappa) < 0$, the result follows. □

A convenient accuracy schedule used in the numerical illustrations is

$$r(\kappa) = \frac{1}{\sqrt{1+\kappa}}, \qquad \rho(\kappa) = \frac{1}{2} + \frac{1}{2\sqrt{1+\kappa}}.$$

Then

$$A(\kappa, p) = \frac{4p(1-p)}{1 + \kappa - (2p-1)^2} = O(1/\kappa).$$

This schedule delivers the $O(1/\kappa)$ pooling rate without assuming $v = u/\kappa$.

OA.A.3 Microfoundation of the salience term and a belief-tilt alternative

Attribute-level salience weighting. Let project i deliver the fundamental attribute $R_\tau(\theta)$ and a state-independent narrative dividend $x_\tau(H) \geq 0$ with $x'_\tau > 0$. A salience-weighted investor evaluates the bundle as

$$U^{\text{dec}}(\theta; H) = R_\tau(\theta) + w_\tau(H)\, x_\tau(H), \qquad U^{\text{exp}}(\theta; H) = R_\tau(\theta) + x_\tau(H),$$

where $w_\tau(H) \geq 1$, nondecreasing in H, is the decision weight on the narrative attribute: hype makes that attribute vivid in the attribute-salience sense of Bordalo, Gennaioli, and Shleifer (2013), or prominent in the focusing sense of Kőszegi and Szeidl (2013), while the fundamentals of a complex project remain pallid. Setting

$$\phi_\tau(H) := w_\tau(H)\, x_\tau(H)$$

delivers the additive, state-independent payoff of Section 3.1; the shape restrictions $\phi'_\tau > 0$ and $\phi''_\tau \leq 0$ are joint restrictions on (w_τ, x_τ), satisfied by the numerical specification $\phi_\tau = \chi_\tau \log(1+H)$. The welfare readings of Section 7 are the poles $w_\tau \equiv 1$ (all of ϕ_τ experienced) and $x_\tau \to 0$ with $w_\tau x_\tau$ fixed (none of it); an intermediate case experiences $x_\tau = \phi_\tau / w_\tau$ and treats the remainder $(w_\tau - 1)x_\tau$ as distortion, so the two reported cases bracket it.

A belief-tilt alternative. Salience could instead operate on beliefs: hype makes upside exemplars vivid and tilts the decision odds of the good state,

$$\frac{\tilde{p}_\tau(H)}{1-\tilde{p}_\tau(H)} = e^{\gamma(H)}\, \frac{p_\tau}{1-p_\tau}, \qquad \gamma(0) = 0, \quad \gamma' > 0.$$

Because the tilt multiplies prior odds and Bayes updating multiplies odds by the signal likelihood ratio, the two operations commute: the decision posterior after signal y is the Bayes posterior computed from the tilted prior $\tilde{p}_\tau(H)$. Three consequences follow.

1. *Scores rise in hype.* The attention score evaluated at a prior p is strictly increasing in p by Lemma 10, so the tilt raises the score and attracts attention, as the additive form does.

2. *Hype erodes discrimination.* As $\gamma \to \infty$, $\tilde{p}(y) \to 1$ at both signal realizations, so the spread of decision posterior means, $\big(\tilde{p}(g) - \tilde{p}(b)\big)\big(V^G_\tau - V^B_\tau\big)$, collapses to zero: in this reading hype itself destroys the option value of processing, reinforcing the pooling force of Corollary 1 on top of complexity.

3. *Hype de-audits.* The CSV stage evaluates the audit window of Proposition 14 at $\tilde{p}(y)$, and $\tilde{p}(y) \to 1$ exits the window from above: hyped projects are funded on faith, as in Proposition 5(1).

The belief-tilt variant therefore reproduces the baseline's qualitative behavior and adds a second pooling channel; the additive baseline is the conservative specification because it leaves a hyped project's information advantage intact.

OA.A.4 Information cost specification (discrete-choice RI over processed signals)

The baseline information cost is

$$\mathcal{K}(\pi) := \lambda \, \mathbb{E}_{\boldsymbol{y}} \left[\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{y}) \log \frac{\pi(a \mid \boldsymbol{y})}{q_a} \right], \qquad \lambda > 0. \tag{33}$$

The investor solves

$$\max_{\pi(\cdot|\cdot)} \; \mathbb{E}_{\boldsymbol{y}} \left[\sum_{a \in \mathcal{A}} \pi(a \mid \boldsymbol{y}) \, \mu_a(y_a) \right] - \mathcal{K}(\pi), \tag{34}$$

subject to feasibility of π.

Interpretation. The main text may use informal language such as "information complexity" or "cognitive cost." Formally, the primitive RI cost is the common Shannon/KL object (33) defined over processed signals; heterogeneity in project difficulty works through the accuracy of the comprehension channel, summarized by the attenuation factor in (8)–(10). The induced state-contingent rule is the mixture $\pi(a \mid \boldsymbol{\theta}) = \mathbb{E}[\pi(a \mid \boldsymbol{y}) \mid \boldsymbol{\theta}]$, so the data-processing inequality caps the state sensitivity of any feasible strategy.

OA.B Proof of Proposition 1

Throughout, fix issuer choices and hence $\{\kappa_i\}$, $\{\rho_i\}$, and the payoff indices, with the outside option normalized to zero. Because each option's payoff index depends only on its own signal and signals are independent across options, all arguments below apply verbatim to *any* bounded own-signal-measurable index; μ_a below denotes that generic index, covering both the baseline evaluation value $\mathcal{W}_a$ of (4) and the myopic posterior valuation.

Proof of Proposition 1. Step 1. Fix $q \in \Delta(\mathcal{A})$. For each $\boldsymbol{y}$, the pointwise problem

$$\max_{\pi(\cdot|\boldsymbol{y})\in\Delta(\mathcal{A})} \sum_a \pi(a \mid \boldsymbol{y})\mu_a(y_a) - \lambda \sum_a \pi(a \mid \boldsymbol{y}) \log \frac{\pi(a \mid \boldsymbol{y})}{q_a}$$

is strictly concave, with first-order conditions

$$\mu_a(y_a) - \lambda\Big(\log \frac{\pi(a \mid \boldsymbol{y})}{q_a} + 1\Big) - \eta(\boldsymbol{y}) = 0 \quad \Longrightarrow \quad \pi^*(a \mid \boldsymbol{y}) = \frac{q_a e^{\mu_a(y_a)/\lambda}}{D(\boldsymbol{y})},$$

where $D(\boldsymbol{y}) := \sum_b q_b e^{\mu_b(y_b)/\lambda}$, with pointwise value $\lambda \log D(\boldsymbol{y})$. Hence, as in Matějka and McKay (2015) and Caplin, Dean, and Leahy (2019), (34) reduces to

$$\max_{q\in\Delta(\mathcal{A})} V(q) := \lambda\, \mathbb{E}_{\boldsymbol{y}} \log\Big(\sum_b q_b e^{\mu_b(y_b)/\lambda}\Big). \tag{35}$$

Step 2. V is concave (expectation of logs of linear functions) and continuous on the compact simplex, so a maximizer exists, and

$$\frac{d}{dt} V\big(q + t(e_a - q)\big)\Big|_{t=0^+} = \lambda\, \mathbb{E}\left[\frac{e^{\mu_a(y_a)/\lambda}}{D(\boldsymbol{y})}\right] - \lambda \;\le\; 0 \quad \forall a, \qquad = 0 \text{ if } q_a > 0,$$

is necessary and sufficient by concavity; this is (21). Corners $q_a = 0$ are admissible and constitute the endogenous consideration set.

Step 3. For $q_a > 0$,

$$\mathbb{E}\big[\pi^*(a \mid \boldsymbol{y})\big] = q_a\, \mathbb{E}\left[\frac{e^{\mu_a(y_a)/\lambda}}{D(\boldsymbol{y})}\right] = q_a,$$

and $\mathbb{E}[\pi^*(a \mid \boldsymbol{y})] = 0 = q_a$ when $q_a = 0$; the weights are the unconditional choice probabilities and (19)–(20) follow. □

Remark (relation to the fixed-point formulation). Existence can alternatively be stated via Brouwer's theorem applied to the mapping $q \mapsto \mathbb{E}[\pi^*(\cdot \mid \cdot)]$. The concave-program formulation is equivalent but strictly more informative: it delivers necessity *and sufficiency* of the support conditions (21), which is what the consideration-set characterization in Proposition 2 uses, and it identifies the investor's equilibrium value with the program value, which is what the welfare analysis in Section 7 uses.

OA.C Pooling limits, the consideration set, and equilibrium crowding: proofs

This section proves the formal properties that drive the main results under the common-λ, signal-level RI specification: the pooling bound, the score-monotonicity lemma, the attention-score-benchmark characterization of the consideration set, the hot-market corollary, and the equilibrium crowding proposition. None of the arguments below invokes an assumed fixed-point mapping; everything is derived from the concave program of Proposition 1.

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

Lemma 4 (Pooling / decoupling under noisy comprehension)**.** *Suppose $V^G_{\tau_i} - V^B_{\tau_i}$ is uniformly bounded. Under the signal-level RI specification,*

$$\left|\pi^*(i \mid \theta_i = G, \boldsymbol{y}_{-i}) - \pi^*(i \mid \theta_i = B, \boldsymbol{y}_{-i})\right| \le \frac{A_i(\kappa_i, p_{\tau_i})}{4\lambda} \left|V^G_{\tau_i} - V^B_{\tau_i}\right| \qquad \text{for every } \boldsymbol{y}_{-i}.$$

Hence the difference converges to zero as $\kappa_i \to \infty$. If $r_i(\kappa_i) = O(\kappa_i^{-1/2})$, the difference is $O(1/\kappa_i)$.

Proof. Fix $\boldsymbol{y}_{-i}$ and define

$$D_{-i} = q_0 + \sum_{j \neq i} q_j \exp(\mu_j(y_j)/\lambda) > 0, \qquad f(x) = \frac{q_i \exp(x/\lambda)}{D_{-i} + q_i \exp(x/\lambda)}.$$

Then $f'(x) = \lambda^{-1} f(x)[1 - f(x)] \le 1/(4\lambda)$. Conditional on θ_i, the choice probability is the mixture over the own signal, $\pi^*(i \mid \theta_i, \boldsymbol{y}_{-i}) = \sum_y \Pr(y_i = y \mid \theta_i) f(\mu_i(y))$, and $\Pr(y_i = g \mid G) - \Pr(y_i = g \mid B) = r_i(\kappa_i)$. Hence

$$\left|\pi^*(i \mid G, \boldsymbol{y}_{-i}) - \pi^*(i \mid B, \boldsymbol{y}_{-i})\right| = r_i(\kappa_i) \left|f(\mu_i(g)) - f(\mu_i(b))\right| \le \frac{r_i(\kappa_i)\,(\mu_i(g) - \mu_i(b))}{4\lambda}.$$

By the computation in the proof of Proposition 8, $r_i(\kappa_i)(\mu_i(g) - \mu_i(b)) = A_i(\kappa_i, p_{\tau_i}) |V^G_{\tau_i} - V^B_{\tau_i}|$. Since $\rho_i(\kappa_i) \to 1/2$, $A_i \to 0$; under $r_i(\kappa_i) = O(\kappa_i^{-1/2})$, the closed form implies $A_i = O(1/\kappa_i)$. For the baseline index, replace μ_i by $\mathcal{W}_i$ and use $|\mathcal{W}_i(p(g)) - \mathcal{W}_i(p(b))| \le (V^G_{\tau_i} - V^B_{\tau_i})\,[p_i(g) - p_i(b)] = \mu_i(g) - \mu_i(b)$, which holds because $\mathcal{W}_i$ is a maximum of affine functions of the posterior with slopes in $[0, V^G_{\tau_i} - V^B_{\tau_i}]$; the bound is unchanged. Proposition 11 states the general-channel version. □

OA.C.2 Denominator effect and attention crowding

Lemma 5 (Mechanical denominator effect (holding q fixed))**.** *Fix unconditional weights $\{q_a\}_{a\in\mathcal{A}}$ and a signal vector $\boldsymbol{y}$. Introduce a new project k with $q_k > 0$ and $\exp(\mu_k(y_k)/\lambda) > 0$. 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,*

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

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

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

OA.C.3 The consideration set: proofs

Throughout this subsection, q denotes an optimal solution of the concave program (18) and $D(\boldsymbol{y}) = q_0 + \sum_j q_j e^{\mathcal{W}_j(y_j)/\lambda}$.

Lemma 6 (Association inequality)**.** *Let $X > 0$ be a random variable with $\mathbb{E}[X] < \infty$ and let h be nonincreasing with $\mathbb{E}[|Xh(X)|] < \infty$. Then $\mathbb{E}[Xh(X)] \leq \mathbb{E}[X]\,\mathbb{E}[h(X)]$.*

Proof. Let X' be an independent copy of X. Since $x \mapsto x$ is increasing and h is nonincreasing, $(X - X')(h(X) - h(X')) \leq 0$ almost surely. Expanding and taking expectations gives $2\mathbb{E}[Xh(X)] - 2\mathbb{E}[X]\mathbb{E}[h(X)] \leq 0$. □

Proof of Lemma 1 (score monotonicity). Work in posterior space: the law of iterated expectations gives $\mathbb{E}[p(y)] = p_\tau$ at every κ, and $\rho'(\kappa) < 0$ implies

$$\frac{\partial p(g)}{\partial \kappa} < 0, \qquad \frac{\partial p(b)}{\partial \kappa} > 0 :$$

the two posteriors move inward as complexity rises, holding their mean fixed. A two-point distribution on $\{b^\dagger, g^\dagger\}$ with mean p_τ has weight $P(g^\dagger) = (p_\tau - b^\dagger)/(g^\dagger - b^\dagger)$ on the upper point, so for convex φ,

$$\psi(b^\dagger, g^\dagger) := \varphi(b^\dagger) + \frac{p_\tau - b^\dagger}{g^\dagger - b^\dagger}\Big[\varphi(g^\dagger) - \varphi(b^\dagger)\Big],$$

$$\frac{\partial \psi}{\partial g^\dagger} = \frac{(p_\tau - b^\dagger)\Big[\varphi'(g^\dagger)(g^\dagger - b^\dagger) - \varphi(g^\dagger) + \varphi(b^\dagger)\Big]}{(g^\dagger - b^\dagger)^2} \;\geq\; 0,$$

and symmetrically $\partial\psi/\partial b^\dagger \le 0$ by monotonicity of chord slopes. Apply this with $\varphi(p) = e^{\mathcal{W}_\tau(p)/\lambda}$, convex as the composition of the increasing convex exponential with the convex $\mathcal{W}_\tau$: $Z_\tau(\kappa, H_\tau)$ is nonincreasing in κ, and continuous since ρ_τ is. Under the separation condition $\mathcal{W}_\tau(p(g)) > \mathcal{W}_\tau(p(b))$ the function φ is not affine on $[p(b), p(g)]$ (it is the exponential of a nonconstant piecewise-affine function), so the chord inequalities are strict at the moving endpoints and the decrease is strict. As $\kappa \to \infty$, $p(g), p(b) \to p_\tau$, so $Z_\tau \to e^{\mathcal{W}_\tau(p_\tau)/\lambda} = Z_\tau^\infty(H_\tau)$. Finally, Jensen on the martingale posterior gives

$$Z_\tau = \mathbb{E}\Big[\varphi(p(y))\Big] \;\ge\; \varphi(p_\tau) = Z_\tau^\infty(H_\tau),$$

strictly under the same conditions. The myopic score substitutes the strictly convex $\varphi(p) = e^{(pV_\tau^G + (1-p)V_\tau^B + \phi_\tau)/\lambda}$ and needs no condition beyond $r_\tau(\kappa) > 0$. □

Proof of Proposition 2 (attention-score benchmark). If $q_i = 0$, then $D(\boldsymbol{y})$ depends on $\boldsymbol{y}_{-i}$ only and $y_i \perp \boldsymbol{y}_{-i}$, so (21) factorizes:

$$1 \;\ge\; \mathbb{E}\left[\frac{e^{\mathcal{W}_i(y_i)/\lambda}}{D(\boldsymbol{y})}\right] = \mathbb{E}\Big[e^{\mathcal{W}_i(y_i)/\lambda}\Big]\,\mathbb{E}\left[\frac{1}{D(\boldsymbol{y})}\right] = \frac{Z_i}{\bar{Z}} \quad\Longrightarrow\quad Z_i \le \bar{Z}.$$

If $q_i > 0$, set $X := e^{\mathcal{W}_i(y_i)/\lambda}$; conditional on $\boldsymbol{y}_{-i}$, $D = q_i X + D_{-i}$ and $h(x) := 1/(q_i x + D_{-i})$ is decreasing, so Lemma 6 and the equality in (21) give

$$1 = \mathbb{E}\Big[\frac{X}{D}\Big] = \mathbb{E}_{\boldsymbol{y}_{-i}}\mathbb{E}\Big[X\,h(X) \mid \boldsymbol{y}_{-i}\Big] \;\le\; \mathbb{E}[X]\,\mathbb{E}\left[\frac{1}{D}\right] = \frac{Z_i}{\bar{Z}} \quad\Longrightarrow\quad Z_i \ge \bar{Z}.$$

For $a = 0$ ($\mu_0 \equiv 0$), (21) reads $\mathbb{E}[1/D] \le 1$, i.e., $\bar{Z} \ge 1$, with equality when $q_0 > 0$. □

Proof of Corollary 2 (hot market). The posterior is a martingale, $\mathbb{E}[p_i(y_i)] = p_{\tau_i}$, and $p \mapsto e^{\mathcal{W}_i(p)/\lambda}$ is convex, so Jensen gives $Z_i \ge e^{\mathcal{W}_i(p_{\tau_i})/\lambda} \ge 1$, the last step because $\mathcal{W}_i \ge 0$. If $\Pr(\mathcal{W}_i(y_i) > 0) > 0$ then $Z_i > 1$; both funding on the prior ($\bar{R}_{\tau_i} + \phi_{\tau_i}(H_i) \ge 0$, whose positive prior value propagates to some signal by convexity) and a nonempty audit window at some posterior produce this. If $q_0 > 0$, then $\bar{Z} = 1$ and Proposition 2 implies $q_i > 0$ for every project with $Z_i > 1$. Exclusion of such a project therefore requires $q_0 = 0$. □

Proof of Corollary 3 (complexity-driven exclusion). At the candidate allocation q^M ($q_0 = 0$, $q_T = 0$, $q_M = 1/N_M$), $D(\boldsymbol{y}) = \bar{X}_{N_M} = N_M^{-1}\sum_j X_j$ with $X_j = e^{\mathcal{W}_M(y_j)/\lambda}$ i.i.d. The Kuhn–Tucker conditions at q^M are:

$$\text{memes:}\quad \mathbb{E}\left[\frac{X_j}{\bar{X}_{N_M}}\right] = 1 \quad \text{(automatic: } \textstyle\sum_j N_M^{-1} X_j/\bar{X}_{N_M} \equiv 1 \text{ and the } X_j \text{ are exchangeable)},$$

$$\text{outside option:}\quad \mathbb{E}\Big[\bar{X}_{N_M}^{-1}\Big] \le 1, \qquad \text{tech:}\quad Z_T\,\mathbb{E}\Big[\bar{X}_{N_M}^{-1}\Big] \le 1 \iff Z_T(\kappa_T, H_T) \le \bar{Z}(N_M),$$

the last by the factorization step above. Sufficiency of (21) certifies global optimality, and $\bar{Z}(N_M)$ is free of κ_T because tech carries zero weight. By Lemma 1, $Z_T(\cdot, H_T)$ is continuous, strictly decreasing, with limit $Z_T^{\infty}(H_T)$: if $Z_T^{\infty}(H_T) < \bar{Z}(N_M)$ the tech condition holds exactly on $[\bar{\kappa}_T, \infty)$, where $Z_T(\bar{\kappa}_T, H_T) = \bar{Z}(N_M)$; if $Z_T^{\infty}(H_T) \ge \bar{Z}(N_M)$ it fails for every κ_T and Proposition 2 yields $q_T > 0$ at every optimum. □

Proof of Proposition 3 (equilibrium crowding). The regime characterization is the preceding proof. *(1)* $\{\bar{X}_n\}$ is a reverse martingale, $\mathbb{E}[\bar{X}_n \mid \bar{X}_{n+1}] = \bar{X}_{n+1}$, so conditional Jensen with $x \mapsto 1/x$ gives

$$\mathbb{E}\Big[\bar{X}_n^{-1}\Big] = \mathbb{E}\Big[\mathbb{E}\Big[\bar{X}_n^{-1} \mid \bar{X}_{n+1}\Big]\Big] \;\ge\; \mathbb{E}\Big[\Big(\mathbb{E}[\bar{X}_n \mid \bar{X}_{n+1}]\Big)^{-1}\Big] = \mathbb{E}\Big[\bar{X}_{n+1}^{-1}\Big],$$

so $\bar{Z}(n)$ is nondecreasing in n. Since $X_j \in [e^{\mu_M(b)/\lambda}, e^{\mu_M(g)/\lambda}] \subset (0, \infty)$, the LLN and dominated convergence give

$$\mathbb{E}\Big[\bar{X}_n^{-1}\Big] \longrightarrow Z_M^{-1}, \qquad \bar{Z}(n) \uparrow Z_M, \qquad \bar{Z}(n) = \Big(\mathbb{E}\Big[\bar{X}_n^{-1}\Big]\Big)^{-1} \le \Big(\mathbb{E}[\bar{X}_n]^{-1}\Big)^{-1} = Z_M.$$

(2) Raising ϕ_M raises W_M and lowers L_M one for one, so $\mathcal{W}_M(p)$ rises pointwise wherever the meme is funded or audited and is unchanged where it is declined; hence every X_j rises weakly pointwise, $\bar{X}_{N_M}^{-1}$ falls, $\mathbb{E}[\bar{X}_{N_M}^{-1}]$ falls, and $\bar{Z}(N_M)$ rises, strictly when funding or auditing occurs with positive probability. Under the myopic index the map is $X_j \mapsto e^{\Delta/\lambda} X_j$ exactly, so $\bar{Z}(N_M)$ scales by $e^{\Delta/\lambda}$. Both regime inequalities are preserved under (1) and (2). *(3)* follows from (1)–(2) and the strict monotonicity of $Z_T(\cdot, H_T)$ in Lemma 1. □

Proof of Corollary 4 (what the trap requires). Exclusion at high complexity requires $Z_T^{\infty}(H_T) < \bar{Z}(N_M) \le Z_M$ by Proposition 3; taking logarithms gives $\mathcal{W}_T(p_T) < \lambda \log Z_M$, and $\mathcal{W}_T(p_T) \ge \max\{\bar{R}_T + \phi_T(H_T), 0\}$ because funding on the prior and declining are both feasible evaluation actions, which is (24). Under the myopic index, $Z_M^{\mu} = e^{\phi_M(H_M)/\lambda} \widetilde{Z}_M(\kappa_M)$ and the same chain rearranges to $\phi_M(H_M) - \phi_T(H_T) > \bar{R}_T - \lambda \log \widetilde{Z}_M(\kappa_M)$. □

Remark 3 (Interior comparative statics)**.** *The results above establish crowding on the extensive margin, the margin used by the entry and verification theorems, directly from the Kuhn–Tucker conditions, without any auxiliary monotonicity assumption on a fixed-point mapping. Section OA.C.4 extends the supply and salience channels to the interior margin in the hot-market regime; the response of interior visibility to own complexity is the comparative static that remains computational.*

OA.C.4 Interior crowding in the hot-market regime

The results above characterize the extensive margin: when a project's weight is zero and when it is positive. This subsection signs the interior margin in the regime where exclusion can occur (Corollary 2), the hot market $q_0 = 0$. Two preliminaries make the interior object well defined.

Lemma 7 (Uniqueness of the attention allocation)**.** *Suppose every project's processed signal is informative, its fundamentals are state dependent, and its index separates the posteriors: $r_i(\kappa_i) > 0$, $V^G_{\tau_i} > V^B_{\tau_i}$, and $\mathcal{W}_i(p_i(g)) > \mathcal{W}_i(p_i(b))$ for all i (the last is automatic under the myopic index, and holds under the baseline whenever the project is funded or audited at the good signal). Then the maximizer of the concave program (18) is unique.*

Proof. Write $X_a := e^{\mathcal{W}_a(y_a)/\lambda}$ with $X_0 \equiv 1$, and let q, q' be two maximizers with denominators $D = \sum_a q_a X_a$ and $D' = \sum_a q'_a X_a$. The optimal set of a concave program is convex, so for $t \in (0,1)$

$$\mathbb{E}\log\Big(tD + (1-t)D'\Big) = t\,\mathbb{E}\log D + (1-t)\,\mathbb{E}\log D',$$

and strict concavity of the logarithm forces $D = D'$ almost surely:

$$\sum_{a\geq 1}(q_a - q'_a)\,X_a + (q_0 - q'_0) = 0 \qquad \text{a.s.}$$

The X_a are independent across options, and each project's X_a is nondegenerate by the separation hypothesis. Taking the variance,

$$0 = \mathrm{Var}\Big(\sum_{a\geq 1}(q_a - q'_a)X_a\Big) = \sum_{a\geq 1}(q_a - q'_a)^2\,\mathrm{Var}(X_a),$$

so $q_a = q'_a$ for every project, and $q_0 = q'_0$ follows from the simplex constraint. □

Uniqueness has two side benefits. In the symmetric-type environment it forces the optimum to be symmetric within types, because relabeling the projects of a type permutes the set of optima and the set is a singleton; the reduction to type-level weights below is therefore without loss. And it makes the continuity hypothesis of Proposition 16 automatic: the argmax correspondence of (18) is single-valued wherever signals are informative, so $q_T(\kappa_T)$ is continuous by Berge's theorem.

In the symmetric-type environment, write $\bar{X} := N_M^{-1}\sum_j X_j$ and $\bar{Y} := n_T^{-1}\sum_g Y_g$ for the equal-weight score averages of the N_M meme and n_T tech projects, and let $a_T \in [0,1]$ denote total tech attention. Under the hot-market hypothesis $q_0 = 0$, within-type symmetry reduces

the program to the scalar concave problem

$$\max_{a\in[0,1]} V(a) := \mathbb{E}\log\Big((1-a)\bar{X} + a\bar{Y}\Big), \qquad V'(a) = F(a) := \mathbb{E}\left[\frac{\bar{Y}-\bar{X}}{(1-a)\bar{X}+a\bar{Y}}\right].$$

Proposition 9 (Interior crowding in the hot market)**.** *Assume informative signals and suppose the optimum of* (18) *sets* $q_0 = 0$ *before and after each change below. Write* $D = (1-a)\bar{X} + a\bar{Y}$*. Then:*

1. First-order condition and corner. *V is strictly concave,*

$$F'(a) = -\mathbb{E}\left[\frac{(\bar{Y}-\bar{X})^2}{D^2}\right] < 0, \qquad F(0) = \mathbb{E}[\bar{Y}]\,\mathbb{E}[\bar{X}^{-1}] - 1 = \frac{Z_T}{\bar{Z}(N_M)} - 1,$$

   so the optimal a_T is unique, satisfies $a_T > 0$ if and only if $Z_T > \bar{Z}(N_M)$, reproducing the inclusion rule of Proposition 2 and Corollary 3 at the corner, and solves $F(a_T) = 0$ whenever interior.

2. Salience channel. *Raising meme salience raises the meme score variables pointwise ($\bar{X} \mapsto \bar{X}' \geq \bar{X}$ almost surely, strictly with positive probability when memes are funded or audited; under the myopic index exactly $\bar{X}' = e^{\Delta/\lambda}\bar{X}$) and lowers a_T: strictly when $a_T \in (0,1)$, while $a_T = 0$ persists. Raising tech salience acts symmetrically and raises a_T.*

3. Supply channel. *Raising N_M lowers a_T, and raising n_T raises a_T, strictly when a_T is interior.*

Proof. (1) Strict concavity: $\bar{Y}-\bar{X}$ is nondegenerate because the two averages are independent and each is nondegenerate under informative signals, so $F' < 0$. At $a = 0$ independence factorizes the expectation, and $\mathbb{E}[\bar{Y}] = Z_T$ with $\mathbb{E}[\bar{X}^{-1}] = \bar{Z}(N_M)^{-1}$ by (22) evaluated at the meme-dominant allocation. A scalar concave program has argmax 0 exactly when $F(0) \leq 0$, an interior root of F otherwise (the corner $a = 1$ requires $F(1) \geq 0$).

(2) Fix $a \in (0,1)$. By part (3)'s computation, the integrand $(\bar{Y}-\bar{X})/D$ is strictly decreasing in the realization of $\bar{X}$ (derivative $-\bar{Y}/D^2$) and strictly increasing in that of $\bar{Y}$ (derivative $\bar{X}/D^2$). Replacing $\bar{X}$ by a pointwise weakly larger $\bar{X}'$ therefore lowers F pointwise, strictly when $\bar{X}' > \bar{X}$ with positive probability, and the argmax of a scalar concave program whose derivative falls pointwise falls: strictly when interior, while the shifted $F(0)$, weakly below its old value, preserves the corner. A pointwise increase in $\bar{Y}$ raises F and a_T symmetrically.

(3) For fixed $y > 0$,

$$\frac{\partial}{\partial x}\,\frac{y-x}{(1-a)x+ay} = -\frac{y}{D^2}, \qquad \frac{\partial^2}{\partial x^2}\,\frac{y-x}{(1-a)x+ay} = \frac{2(1-a)y}{D^3} \;>\; 0,$$

so $x \mapsto \mathbb{E}_{\bar{Y}}\left[(\bar{Y}-x)/((1-a)x+a\bar{Y})\right]$ is strictly convex for $a < 1$; symmetrically, the map in y at fixed x has second derivative $-2ax/D^3 < 0$ and is strictly concave for $a > 0$. By the reverse-martingale property in the proof of Proposition 3, $\bar{X}_{N_M+1} = \mathbb{E}[\bar{X}_{N_M} \mid \bar{X}_{N_M+1}]$, so conditional Jensen applied to the strictly convex map lowers F pointwise in a and the argmax falls; contracting $\bar{Y}$ raises F by the concave case and the argmax rises. Strictness at interior optima follows from strict convexity or concavity together with the nondegeneracy of the conditional distributions of i.i.d. nondegenerate averages. □

Remark 4 (Persistence of the hot market, and the margin left to computation)**.** *The hypothesis $q_0 = 0$ is preserved by the salience channel in the deep-menu limit: with $\bar{X} \to Z_M$ almost surely, an optimum with $q_0 > 0$ has $\mathbb{E}[1/D] = 1$ by (21), and the meme condition $cZ_M\,\mathbb{E}[1/D] \le 1$ then reads $cZ_M \le 1$; hence $cZ_M > 1$ forces $q_0 = 0$ at every optimum, and raising c preserves the condition. The response of interior tech attention to tech complexity itself is the one comparative static these arguments do not sign: complexity moves the tech posterior means inward at a fixed mean (Lemma 1), but the composite of that contraction with the first-order condition has curvature of ambiguous sign, and that margin remains computational.*

OA.C.5 General comprehension channels and general payoff indices

Nothing in the paper's information structure is essentially binary, and nothing essentially depends on the attention index being the affine posterior valuation. This subsection states the two relevant results at their natural level of generality: complexity as Blackwell garbling, and the index as any convex increasing function of the posterior.

Setup. Fix a project with prior $\bar{p} := \Pr(\theta = G)$ and let complexity index a family of signal structures $\{F_\kappa\}$ that is decreasing in Blackwell informativeness: $\kappa' > \kappa$ implies $F_{\kappa'}$ is a garbling of F_κ. For binary states, Blackwell's theorem makes this equivalent to a statement about the induced distribution H_κ of the posterior probability $p(y) = \Pr(\theta = G \mid y)$: every H_κ has mean $\bar{p}$, and $\kappa' > \kappa$ implies $H_{\kappa'}$ is a mean-preserving contraction of H_κ (Blackwell, 1953). Write the uninformative limit as $H_\kappa \to \delta_{\bar{p}}$. Let the attention-stage payoff index be $\omega(p)$, any convex, nondecreasing function of the posterior with Lipschitz constant at most

L_ω; the two cases used in the paper are the posterior valuation $\mu(p) = pW - (1-p)L$, affine with $L_\omega = W + L = V^G - V^B$, and the CSV continuation value $\mathcal{W}(p) = \max\{0,\ pW - (1-p)L,\ pW - C_V\}$, a maximum of affine functions with the same Lipschitz constant. The score is $Z(\kappa) := \mathbb{E}_{H_\kappa}\left[e^{\omega(p)/\lambda}\right]$.

Proposition 10 (Scores under general channels and indices)**.** *$e^{\omega(p)/\lambda}$ is convex and nondecreasing in p, so:*

1. *$Z(\kappa)$ is nonincreasing in κ, strictly decreasing wherever the contraction is strict and H_κ places mass where ω is nonconstant;*
2. *$Z(\kappa) \geq e^{\omega(\bar{p})/\lambda}$ for every κ, with equality in the uninformative limit;*
3. *hence the exclusion threshold logic of Corollary 3 and Proposition 4 applies verbatim: $Z(\cdot)$ crosses any benchmark $\bar{Z} \in \left(e^{\omega(\bar{p})/\lambda},\ Z(\kappa_{\min})\right)$ exactly once.*

Proof. Convexity and monotonicity of the composition are immediate ($e^{x/\lambda}$ is convex increasing, ω convex nondecreasing). A mean-preserving contraction weakly lowers the expectation of every convex function (strictly under the stated conditions), which gives (1); Jensen gives (2); continuity and monotonicity give (3). □

Generalized attenuation. Define

$$A(\kappa) := \frac{\mathrm{Var}_{H_\kappa}(p)}{\bar{p}(1-\bar{p})} \ \in\ [0,1]. \tag{36}$$

A is nonincreasing in κ (variance is a convex moment), $A \to 0$ in the uninformative limit, and $A = 1$ at full information. For the binary symmetric channel, a direct computation gives $\mathrm{Var}(p) = P(g)P(b)\,[p(g) - p(b)]^2$ and

$$A(\kappa) = \frac{4\bar{p}(1-\bar{p})\, r(\kappa)^2}{1 - (2\bar{p}-1)^2 r(\kappa)^2},$$

exactly the attenuation factor (10): the paper's A_i is the normalized posterior variance of the comprehension channel, and (10) is its binary closed form.

Proposition 11 (Pooling under general channels and indices)**.** *Fix rivals' signals $\boldsymbol{y}_{-i}$ and unconditional weights, and let the conditional choice probability be induced by (19) with index ω. Then*

$$\left|\pi^*(i \mid G, \boldsymbol{y}_{-i}) - \pi^*(i \mid B, \boldsymbol{y}_{-i})\right| \ \leq\ \frac{L_\omega\, A(\kappa)}{4\lambda},$$

which reduces to the bound of Corollary 1 for the binary channel and the μ index.

Proof. Write $h(p) := f(\omega(p))$ where $f(x) = q_i e^{x/\lambda}/(D_{-i} + q_i e^{x/\lambda})$; $f' = f(1-f)/\lambda \le 1/(4\lambda)$, so h is Lipschitz with constant $L_\omega/(4\lambda)$. Bayes' rule gives the conditional posterior densities $dH^G/dH = p/\bar{p}$ and $dH^B/dH = (1-p)/(1-\bar{p})$, hence

$$\mathbb{E}[h \mid G] - \mathbb{E}[h \mid B] = \frac{\mathbb{E}[h(p)\, p] - \bar{p}\, \mathbb{E}[h(p)]}{\bar{p}(1-\bar{p})} = \frac{\mathrm{Cov}(h(p), p)}{\bar{p}(1-\bar{p})}.$$

For an independent copy p', $\mathrm{Cov}(h(p), p) = \frac{1}{2}\,\mathbb{E}\big[(h(p) - h(p'))(p - p')\big] \le \frac{L_\omega}{4\lambda} \cdot \frac{1}{2}\,\mathbb{E}[(p-p')^2] = \frac{L_\omega}{4\lambda}\,\mathrm{Var}(p)$, and the same bound holds for $-\mathrm{Cov}$. Combining the displays gives the claim. □

What depends on what. Proposition 1, the benchmark characterization of Proposition 2, the supply channel of Proposition 3, the uniqueness lemma, and the welfare results of Section 7 use only own-signal measurability, boundedness, and independence across projects, so they are free of both the channel and the index. Lemma 1, Corollary 1, and the threshold constructions need only Propositions 10 and 11 above. The binary symmetric channel with $r(\kappa) = (1+\kappa)^{-1/2}$ remains the parametric example used in all computations.

OA.C.6 Types as regions of a characteristic space

The two-type structure of the main text is expository. Index projects by the characteristic vector

$$x_i = \big(p_i,\, V_i^G,\, V_i^B,\, \kappa_{i,0},\, \chi_i,\, K_i,\, c_{im},\, c_{ih}\big) \in \mathcal{X},$$

and read "meme" as the salience-cheap, legible, low-fixed-cost region of $\mathcal{X}$ and "tech" as the fundamentals-rich, complex, high-fixed-cost region. Every project-level object in the paper is defined at the level of x_i and nowhere uses the two-point support of the type distribution: the score Z_i of (16), the inclusion rule $Z_i \gtrless \bar{Z}$ of Proposition 2, the pooling bound of Corollary 1, the hot-market condition of Corollary 2, the audit window of Proposition 5, and the entry conditions of Proposition 4 therefore hold verbatim for an arbitrary finite menu of heterogeneous projects. Corollary 4 likewise locates the salience wedge that exclusion requires in characteristic space rather than in labels.

Exchangeability is used in exactly one place: the closed form $\bar{Z}(N_M) = (\mathbb{E}[\bar{X}_{N_M}^{-1}])^{-1}$ of Proposition 3 treats the crowding menu as i.i.d. draws, which is what delivers the reverse-martingale supply monotonicity. The salience channel survives heterogeneity without any symmetry:

Lemma 8 (Salience scaling under heterogeneity, myopic index)**.** *Work under the myopic index μ, where salience scales scores exactly. Let q^* be optimal for a menu of heterogeneous projects with $q_0 = 0$, let $\mathcal{C}^*$ be its consideration set, and suppose every considered option's*

salience rises by $\Delta > 0$: $\phi_j \mapsto \phi_j + \Delta$ for all $j \in \mathcal{C}^*$. Then q^* remains optimal, and the benchmark scales to $e^{\Delta/\lambda}\bar{Z}$: every excluded project's inclusion condition tightens by the factor $e^{\Delta/\lambda}$. Under the baseline index the benchmark rises weakly by the pointwise argument in the proof of Proposition 3(2), without the exact factor.

Proof. The shift maps $e^{\mu_j(y_j)/\lambda} \mapsto e^{\Delta/\lambda}e^{\mu_j(y_j)/\lambda}$ for $j \in \mathcal{C}^*$; since excluded options and the outside option carry zero weight at q^*, $D(\boldsymbol{y}) \mapsto e^{\Delta/\lambda}D(\boldsymbol{y})$. The Kuhn–Tucker residuals (21) transform as

$$\mathbb{E}\left[\frac{e^{(\mu_j(y_j)+\Delta)/\lambda}}{e^{\Delta/\lambda}D(\boldsymbol{y})}\right] = \mathbb{E}\left[\frac{e^{\mu_j(y_j)/\lambda}}{D(\boldsymbol{y})}\right] = 1 \quad (j \in \mathcal{C}^*),$$
$$\mathbb{E}\left[\frac{e^{\mu_a(y_a)/\lambda}}{e^{\Delta/\lambda}D(\boldsymbol{y})}\right] = e^{-\Delta/\lambda}\,\mathbb{E}\left[\frac{e^{\mu_a(y_a)/\lambda}}{D(\boldsymbol{y})}\right] \leq 1 \quad (q_a = 0),$$

and the outside option's condition $\mathbb{E}[(e^{\Delta/\lambda}D)^{-1}] \leq 1$ is implied by $\mathbb{E}[D^{-1}] \leq 1$. Sufficiency of (21) for the concave program preserves optimality, and $\bar{Z} = (\mathbb{E}[D^{-1}])^{-1}$ scales by $e^{\Delta/\lambda}$. □

For the supply channel, the reverse-martingale argument uses within-class exchangeability. With heterogeneous classes the comparative static is verified computationally: the implementation of Section 6 operates entirely in a heterogeneous environment, ten tech legibility groups and ten meme virality classes (Online Appendix OA.F.1), and reproduces the crowding pattern, so no reported result depends on within-type symmetry.

OA.C.7 Hype-stock comparative statics

Meme crowding can now occur through three distinct channels: larger meme supply N_M, higher meme scores Z_M from low complexity or high salience, and higher meme hype stocks H_M generated by attention feedback. The last channel operates through

$$Z_\tau(\kappa_\tau, H_\tau) = \exp\left(\frac{\phi_\tau(H_\tau)}{\lambda}\right)\widetilde{Z}_\tau(\kappa_\tau), \qquad \frac{\partial Z_\tau}{\partial H_\tau} = \frac{\phi'_\tau(H_\tau)}{\lambda}Z_\tau > 0.$$

Proposition 12 (Dynamic hype amplification)**.** *Fix all rival projects and let the equilibrium visibility of project i be written as $q_i = Q_i(H_i)$, where $Q'_i(H_i) > 0$. Suppose the hype stock satisfies the steady-state law*

$$\delta_{\tau_i}H_i = \xi_{\tau_i}h_i + \rho_{\tau_i}Q_i(H_i).$$

At any locally stable steady state satisfying

$$D_i := \delta_{\tau_i} - \rho_{\tau_i}Q'_i(H_i) > 0,$$

the effect of hype effort on hype stock and visibility is

$$\frac{\partial H_i}{\partial h_i} = \frac{\xi_{\tau_i}}{D_i}, \qquad \frac{\partial q_i}{\partial h_i} = \frac{\xi_{\tau_i} Q_i'(H_i)}{D_i} = \frac{\xi_{\tau_i}}{\delta_{\tau_i}} Q_i'(H_i) \cdot \frac{1}{1 - \rho_{\tau_i} Q_i'(H_i)/\delta_{\tau_i}}.$$

The one-shot hype benchmark corresponds to $\rho_{\tau_i} = 0$; positive attention feedback $\rho_{\tau_i} > 0$ amplifies the marginal effect of hype effort, more strongly for meme projects when $\rho_M > \rho_T$ or $\xi_M > \xi_T$.

Proof of Proposition 12. Write $G_i(H_i, h_i) := \delta_{\tau_i} H_i - \xi_{\tau_i} h_i - \rho_{\tau_i} Q_i(H_i) = 0$. If

$$\frac{\partial G_i}{\partial H_i} = \delta_{\tau_i} - \rho_{\tau_i} Q_i'(H_i) = D_i > 0,$$

the implicit function theorem gives

$$\frac{\partial H_i}{\partial h_i} = -\frac{\partial G_i/\partial h_i}{\partial G_i/\partial H_i} = \frac{\xi_{\tau_i}}{D_i}, \qquad \frac{\partial q_i}{\partial h_i} = Q_i'(H_i)\frac{\partial H_i}{\partial h_i} = \frac{\xi_{\tau_i} Q_i'(H_i)}{D_i},$$

and rearranging yields the multiplier expression; at $\rho_{\tau_i} = 0$ the multiplier equals one. □

Proposition 13 (Self-reinforcing hype and multiple salience states)**.** *Fix rival projects and let $q_i = Q_i(H_i)$ be continuous, increasing, differentiable, and bounded. Consider the steady-state equation*

$$\delta_i H_i = \xi_i h_i + \rho_i Q_i(H_i).$$

If $\rho_i \sup_H Q_i'(H) \le \delta_i$, the steady state is unique. If instead $\rho_i Q_i'(H) > \delta_i$ on some interval and there exist $H_1 < H_2$ such that

$$\delta_i H_1 - \rho_i Q_i(H_1) > \xi_i h_i, \qquad \delta_i H_2 - \rho_i Q_i(H_2) < \xi_i h_i,$$

then the steady-state equation admits multiple solutions. Low- and high-hype steady states are locally stable whenever

$$|1 - \delta_i + \rho_i Q_i'(H_i^*)| < 1.$$

Proof. Define

$$F(H) := \delta_i H - \xi_i h_i - \rho_i Q_i(H).$$

If $\rho_i \sup_H Q_i'(H) \le \delta_i$, then

$$F'(H) = \delta_i - \rho_i Q_i'(H) \ge 0,$$

so F is monotone. Since $F(0) < 0$ when $\xi_i h_i + \rho_i Q_i(0) > 0$ and $F(H) \to \infty$ as $H \to \infty$ because Q_i is bounded, there is a unique root.

For multiplicity, suppose there are $H_1 < H_2$ satisfying the stated inequalities. Since $F(0) < 0$ and $F(H_1) > 0$, the intermediate value theorem gives one root in $(0, H_1)$. Since $F(H_1) > 0$ and $F(H_2) < 0$, there is a second root in (H_1, H_2). Since Q_i is bounded, $F(H) \to \infty$ as $H \to \infty$; together with $F(H_2) < 0$, this gives a third root above H_2. Local stability of the discrete law

$$H_{t+1} = (1 - \delta_i)H_t + \xi_i h_i + \rho_i Q_i(H_t)$$

requires the absolute value of the derivative of the transition map at H_i^* to be below one, i.e.

$$|1 - \delta_i + \rho_i Q_i'(H_i^*)| < 1.$$

□

OA.D Costly verification (CSV): posterior beliefs, the audit window, and extensions

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

OA.D.1 Baseline: posterior beliefs, perfect verification, and the audit window

Lemma 9 (Consideration adds no information beyond the own signal)**.** *Fix the attention rule $\pi(a \mid \boldsymbol{y})$. For any project i,*

$$\Pr\left(\theta_i = G \mid y_i,\ a = i\right) = \Pr\left(\theta_i = G \mid y_i\right) = p_i(y_i).$$

Proof. Conditional on y_i, the event $\{a = i\}$ is determined by ($\boldsymbol{y}_{-i}$, the randomization device), and $\theta_i \perp \boldsymbol{y}_{-i} \mid y_i$ because states and signals are independent across projects. Hence $\Pr(\theta_i = G,\ a = i \mid y_i) = \Pr(\theta_i = G \mid y_i)\Pr(a = i \mid y_i)$, and Bayes' rule gives the claim. □

The belief entering the CSV stage is therefore the processed-signal posterior $p_i(y_i)$ of (11). For a generic belief p, if the investor funds without verification the expected payoff is

$$U_i^{NV}(p) = p\Big(V_{\tau_i}^G + \phi_{\tau_i}(H_i)\Big) + (1 - p)\Big(V_{\tau_i}^B + \phi_{\tau_i}(H_i)\Big) = pW_i - (1 - p)L_i, \tag{37}$$

where $W_i := V^G_{\tau_i} + \phi_{\tau_i}(H_i) > 0$ and $L_i := -(V^B_{\tau_i} + \phi_{\tau_i}(H_i)) > 0$. If she verifies and funds only after observing the good state, the expected payoff is

$$U_i^V(p) = pW_i - C_V. \tag{38}$$

Allowing for the outside option implies continuation values

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

Proposition 14 (Audit window)**.** *Verification is optimal at belief p, i.e., $W_i^V(p) \geq W_i^{NV}(p)$ with $W_i^V(p) > 0$, if and only if*

$$C_V \leq \frac{W_i L_i}{W_i + L_i} \qquad \textit{and} \qquad \frac{C_V}{W_i} \;\leq\; p \;\leq\; 1 - \frac{C_V}{L_i}.$$

The window has the following comparative statics: the upper edge $1 - C_V/L_i$ is decreasing in the hype stock H_i (because L_i is decreasing in H_i), so hype de-audits projects near the funded-on-faith margin; the lower edge C_V/W_i is decreasing in H_i; and both edges tighten as C_V rises.

Proof. If $U_i^{NV}(p) \leq 0$, i.e. $p \leq \hat{p} := L_i/(W_i + L_i)$, then $W_i^{NV}(p) = 0$ and verification is optimal iff

$$U_i^V(p) = pW_i - C_V \geq 0 \iff p \geq \frac{C_V}{W_i}.$$

If $U_i^{NV}(p) > 0$, verification is optimal iff

$$pW_i - C_V \;\geq\; pW_i - (1-p)L_i \iff (1-p)L_i \geq C_V \iff p \leq 1 - \frac{C_V}{L_i}.$$

At $\hat{p}$ the two thresholds coincide,

$$\hat{p}W_i = (1-\hat{p})L_i = \frac{W_i L_i}{W_i + L_i},$$

so the cases paste into $[C_V/W_i,\, 1 - C_V/L_i]$, nonempty iff $C_V \leq W_i L_i/(W_i + L_i)$. The comparative statics follow from $\partial W_i/\partial H_i = \phi' \geq 0$ and $\partial L_i/\partial H_i = -\phi' \leq 0$. □

Prior-based rule as the high-complexity limit. As $\kappa_i \to \infty$, $p_i(y_i) \to p_{\tau_i}$ for both signal realizations, and the window condition evaluated at the prior reduces (in the funded

region $U^{NV} \geq 0$, $U^V > 0$) to the familiar CSV threshold

$$(1 - p_{\tau_i})\Big(- V^B_{\tau_i} - \phi_{\tau_i}(H_i)\Big) \geq C_V. \tag{39}$$

Complexity therefore imposes a second penalty even conditional on consideration: it removes the informational content of having been noticed, so audits of very complex projects are triggered on the prior alone.

Invisible innovation: an explicit parameter region. Fix meme primitives so that the meme-dominant regime of Proposition 3 obtains with benchmark $\bar{Z}(N_M) > 1$, and pick tech primitives with $\mathcal{W}_T(p_T) < \lambda \log \bar{Z}(N_M)$, so that Corollary 3 excludes tech for $\kappa_T > \bar{\kappa}_T$. The audit-window condition of Proposition 14 involves only $(p_T, V^\theta_T, \phi_T, C_V)$ and is satisfied, for example, whenever $C_V < W_T L_T/(W_T + L_T)$ and $p_T \in [C_V/W_T,\, 1 - C_V/L_T]$; the exclusion condition constrains only the level $\mathcal{W}_T(p_T)$ relative to the meme side, and the benchmark can be pushed above any given level through meme salience (Proposition 3), so the two requirements are compatible. Any such configuration exhibits a project that would be audited if considered but is never considered, which is Proposition 5(2).

OA.D.2 Extension 1: noisy verification

The baseline perfect-audit assumption is a benchmark, not a description of smart-contract auditing. Let p_i denote the belief entering the CSV stage; by Lemma 9 this is the processed-signal posterior $p_i(y_i)$, with the type prior as its high-complexity limit. Suppose an audit generates a report $\sigma_i \in \{g, b\}$ with symmetric accuracy

$$\Pr(\sigma_i = g \mid \theta_i = G) = \Pr(\sigma_i = b \mid \theta_i = B) = \xi_i, \qquad \xi_i \in (1/2, 1].$$

The report probabilities are

$$\Pr(\sigma_i = g) = p_i\xi_i + (1 - p_i)(1 - \xi_i), \qquad \Pr(\sigma_i = b) = p_i(1 - \xi_i) + (1 - p_i)\xi_i,$$

and Bayes' rule gives

$$p^g_i = \frac{p_i\xi_i}{p_i\xi_i + (1 - p_i)(1 - \xi_i)}, \qquad p^b_i = \frac{p_i(1 - \xi_i)}{p_i(1 - \xi_i) + (1 - p_i)\xi_i}. \tag{40}$$

Write

$$W_i := V^G_{\tau_i} + \phi_{\tau_i}(H_i), \qquad L_i := -(V^B_{\tau_i} + \phi_{\tau_i}(H_i)).$$

The no-verification continuation value is

$$W_i^{NV} = \max\{p_i W_i - (1 - p_i)L_i, 0\}. \tag{41}$$

Under noisy verification, the investor observes σ_i, updates to p_i^σ, and then chooses whether to fund. The verification value is therefore

$$W_i^V(\xi_i) = \sum_{\sigma \in \{g,b\}} \Pr(\sigma_i = \sigma) \max\{p_i^\sigma W_i - (1 - p_i^\sigma)L_i, 0\} - C_V. \tag{42}$$

Verification is optimal if and only if $W_i^V(\xi_i) \geq W_i^{NV}$. When $\xi_i = 1$, $p_i^g = 1$ and $p_i^b = 0$. If $W_i > 0$ and $L_i > 0$, then (42) reduces to

$$W_i^V(1) = p_i W_i - C_V,$$

which is the perfect-verification expression used in the main text.

Monotonicity in audit precision. For binary symmetric signals, a higher precision ξ_i Blackwell-dominates a lower precision. Since the investor's post-audit funding decision is chosen after observing the report, the expected post-audit continuation value is weakly increasing in ξ_i. Hence

$$W_i^V(1) \geq W_i^V(\xi_i) \qquad \text{for all } \xi_i < 1.$$

Noisy verification therefore cannot enlarge the set of projects for which verification is valuable relative to the perfect-verification benchmark.

Complexity-dependent audit precision. Smart-contract complexity can also reduce audit informativeness. Let $\xi_i = \xi(\kappa_i)$ with $\xi'(\kappa_i) \leq 0$. Then effective complexity has a double penalty:

$$\kappa_i \uparrow \quad \Rightarrow \quad A_i(\kappa_i, p_i) \downarrow \quad \Rightarrow \quad q_i \downarrow,$$

and, conditional on eligibility,

$$\kappa_i \uparrow \quad \Rightarrow \quad \xi_i(\kappa_i) \downarrow \quad \Rightarrow \quad W_i^V(\xi_i) \downarrow .$$

Complexity weakens both the probability of reaching the audit stage and the informativeness of the audit conditional on reaching it. The invisible-innovation result is unchanged because any project with $q_i < \varepsilon$ is outside the audit pipeline regardless of audit precision.

OA.D.3 Extension 2: exogenous trigger events

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

$$\mathbf{1}\{q_i > 0\} + T_i \geq 1,$$

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

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

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

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

and the value of verification becomes

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

The same threshold logic survives with a richer continuation value.

OA.D.5 Remark: relation to a prior-based specification

A simpler specification conditions the CSV stage on the consideration event and the type prior only. Lemma 9 shows that the posterior-based rule is the internally consistent baseline: the attention stage and the verification stage share the information generated by the comprehension channel, and the prior-based rule is the $\kappa_i \to \infty$ limit. The invisible-innovation mechanism is identical under either convention: if the project never enters the consideration set, neither prior-based nor posterior-based verification is triggered.

OA.E Issuer optimization and the complexity threshold

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

OA.E.1 Issuer problem and best responses

Given equilibrium visibility q_i, issuer i solves

$$\max_{m_i \geq 0,\ h_i \geq 0} \alpha g(q_i(m_i, H_i))s_i - \Big(K_{\tau_i} + c_{\tau_i m} m_i + c_{\tau_i h} h_i\Big), \qquad \delta_{\tau_i} H_i = \xi_{\tau_i} h_i + \rho_{\tau_i} q_i. \tag{43}$$

The choice variables affect visibility through two channels: legibility lowers κ_i and hype-seeding effort raises the stock H_i, which shifts salience through $\phi_{\tau_i}(H_i)$ and may be amplified by attention feedback.

OA.E.2 Symmetric scores

Under symmetry within types, define

$$U_\tau(\theta; H_\tau) := V_\tau^\theta + \phi_\tau(H_\tau), \qquad \bar{R}_\tau := p_\tau V_\tau^G + (1 - p_\tau) V_\tau^B,$$

and the signal-level type score of (17),

$$Z_\tau(\kappa_\tau, H_\tau) = \exp\left(\frac{\phi_\tau(H_\tau)}{\lambda}\right) \widetilde{Z}_\tau(\kappa_\tau), \qquad \widetilde{Z}_\tau(\kappa_\tau) = \sum_{y \in \{g,b\}} P_\tau(y) \exp\left(\frac{m_\tau(y)}{\lambda}\right),$$

with $P_\tau(y)$ the unconditional signal probability and $m_\tau(y) = \mathbb{E}[R_\tau(\theta) \mid y]$; the display is the myopic score of (17), and the baseline score is $Z_\tau(\kappa_\tau, H_\tau) = \mathbb{E}[e^{\mathcal{W}_\tau(p(y))/\lambda}]$. By Lemma 1, either score is continuous and strictly decreasing in κ_τ, so

$$Z_T(\kappa_T, H_T) \downarrow Z_T^\infty(H_T) = \exp\left(\frac{\mathcal{W}_T(p_T)}{\lambda}\right) \qquad \text{as } \kappa_T \to \infty,$$

with the myopic limit $\exp\{(\bar{R}_T + \phi_T(H_T))/\lambda\}$.

Relation to the RI-over-θ reduced form. A natural reduced form replaces $\mu_\tau(y)$ by its state-conditional mean $v_\tau(\theta)$ and computes the state-level score $\widetilde{Z}_\tau^\theta = \mathbb{E}_\theta[e^{v_\tau(\theta)/\lambda}]$; the numerical implementation does not use it (it solves the signal-level program directly, Online Appendix OA.F.2), and the comparison is recorded for interpretation. By conditional Jensen, $\widetilde{Z}_\tau^\theta \leq \widetilde{Z}_\tau$ with equality only in the degenerate cases; both scores share the same mean payoff, the same $\kappa \to \infty$ limit, the same strict monotonicity in κ_τ (the state-level version is increasing in A_τ, hence decreasing in κ_τ), and the same pooling bound. Online Appendix OA.G.2 discusses the comparison further.

OA.E.3 Tech entry thresholds

Proposition 15 (Entry collapse at the exclusion margin)**.** *Fix meme primitives such that the meme-dominant regime of Proposition 3 obtains, with benchmark $\bar{Z}(N_M)$, and suppose the exclusion condition (23) holds, $Z_T^{\infty}(H_T) < \bar{Z}(N_M)$. Then for all $\kappa_T > \bar{\kappa}_T$, where $Z_T(\bar{\kappa}_T, H_T) = \bar{Z}(N_M)$: every tech project has $q_T = 0$ (Corollary 3), hence $F_T = \alpha g(0) s_T = 0$ and $\Pi_T \leq -K_T < 0$, so no tech issuer enters at any feasible entrant count and for any (α, g, s_T). Because legibility investment enters through $\kappa_T = \kappa_{T,0} e^{-\beta_T m_T}$ at cost $c_{Tm} m_T$, entry also fails whenever the legibility outlay required to reach $\kappa_T \leq \bar{\kappa}_T$, namely $c_{Tm} \log(\kappa_{T,0}/\bar{\kappa}_T)/\beta_T$, added to K_T, exceeds the maximal attainable revenue $\alpha g(1) \bar{s}_T$; in particular, for fixed revenue bound this failure region expands as $\kappa_{T,0}$ rises or β_T falls.*

Proof. Immediate from Corollary 3, $g(0) = 0$, and the cost function (2): on the exclusion region revenue is zero while cost is at least $K_T > 0$. For the legibility clause, reaching $\kappa_T \leq \bar{\kappa}_T$ requires $m_T \geq \log(\kappa_{T,0}/\bar{\kappa}_T)/\beta_T$, and revenue is bounded above by $\alpha g(1) \bar{s}_T$ regardless of visibility. □

Proposition 16 (Entry failure at the reach margin)**.** *Suppose tech visibility is positive, $q_T > 0$. Entry fails whenever $\alpha g(q_T) \bar{s}_T < K_T$. If $g_1 \leq g_2$ pointwise, the entry region under g_1 is contained in that under g_2, and the entry region is monotone in α, $\bar{s}_T$, and K_T in the obvious directions. If, in addition, the investor optimum is unique at each κ_T in a neighborhood, then $q_T(\kappa_T)$ is continuous there (Berge's maximum theorem applied to the concave program (18)), and the reach-margin failure set $\{\kappa_T : \alpha g(q_T(\kappa_T)) \bar{s}_T < K_T\}$ is open in that neighborhood.*

Proof. The first statements are monotonicity of $\alpha g(\cdot) \bar{s}_T$. For the continuity clause: the objective of (18) is jointly continuous in (q, κ_T) (through μ_T), the feasible set is the fixed compact simplex, so the argmax correspondence is upper hemicontinuous; single-valuedness upgrades this to continuity of the selection $q_T(\kappa_T)$, and the failure set is the preimage of an open set under a continuous map. □

Remark 5 (Uniqueness)**.** *The concave program's objective is not strictly concave in q, but its maximizer is nevertheless unique whenever every project's signal is informative (Lemma 7), which the maintained accuracy schedule guarantees at every finite κ. The continuity hypothesis in Proposition 16 is therefore automatic, and issuer payoffs are single-valued continuous functions of issuer actions.*

OA.E.4 The deep-menu limit: regimes and thresholds in closed form

The finite-menu issuer equilibrium is computed by branch tracing (Online Appendix OA.F.4). This subsection solves the deep-menu limit of that protocol in closed form: the limit collapses the interior transition to a knife edge, delivers the deterioration and recovery thresholds analytically, and shows that hysteresis needs no ecosystem externality, only the meme sector's conduct feedback.

The limit economy. Work in the symmetric-type environment and replicate each type: as the number of projects per type grows with total type weights held fixed, the weighted score averages converge almost surely to their means, so the investor program becomes

$$\max_{(q_0,a_M,a_T)\in\Delta} \lambda \log\Big(q_0 + a_M Z_M(h) + a_T Z_T(m)\Big),$$

a linear function inside a logarithm. Its solution is winner take all: all attention goes to the class with the highest score among $\{1,\, Z_M(h),\, Z_T(m)\}$, ties split arbitrarily, and $\bar{Z} = \max\{1, Z_M(h), Z_T(m)\}$. Write $Z_T(m) := Z_T(\kappa_{T,0}e^{-\beta_T m})$, strictly increasing in m (Lemma 1), and $Z_M(h)$, strictly increasing in h through salience. The issuer protocol is the one the computations iterate: the meme sector's effort is the largest h with nonnegative per-meme profit at the current attention allocation, tech groups enter or exit by the sign of profit at a common legibility choice, and the two adjust until a fixed point (a *regime*) is reached.

Proposition 17 (Closed-form regimes)**.** *Let $\bar{s}_T, \bar{s}_M$ bound conversion rates, let $K_M = 0$, and define the affordability bounds*

$$\bar{m} := \frac{\alpha\bar{s}_T - K_T}{c_{Tm}}, \qquad h^* := \left(\frac{\alpha\bar{s}_M/N_M}{c_{Mh2}}\right)^{1/2},$$

the largest legibility a lone tech entrant can finance and the meme sector's zero-profit effort when memes hold all attention. Then, in the limit economy:

1. Tech-active regime. *Memes attract no attention, earn nothing, and set $h = 0$; the regime is a fixed point of the protocol if and only if*

$$Z_T(\kappa_{T,0}e^{-\beta_T m}) \;\geq\; \max\{1,\, Z_M(0)\} \quad \textit{for some } m \in [0, \bar{m}],$$

equivalently $\kappa_{T,0} \leq \kappa^d := e^{\beta_T \bar{m}}\, Z_T^{-1}\Big(\max\{1, Z_M(0)\}\Big)$, and free entry supports $N^* = \lfloor \alpha\bar{s}_T/(K_T + c_{Tm}m^*)\rfloor$ entrants at the smallest consideration-preserving legibility m^*.

2. Meme-dominant regime. *Memes hold all attention and set $h = h^*$; the regime is a fixed point if and only if no affordable tech entry clears the benchmark,*

$$Z_T(\kappa_{T,0}e^{-\beta_T\bar{m}}) \;\le\; Z_M(h^*), \quad \textit{equivalently} \quad \kappa_{T,0} \;\ge\; \kappa^r := e^{\beta_T\bar{m}}\, Z_T^{-1}\Big(Z_M(h^*)\Big).$$

3. Hysteresis without externalities. *Because $h^* > 0$, $Z_M(h^*) > Z_M(0)$ and hence $\kappa^r < \kappa^d$ whenever $Z_M(0) \ge 1$: on the band $\kappa_{T,0} \in (\kappa^r, \kappa^d]$ both regimes are fixed points, the deterioration branch collapses at κ^d and the recovery branch restarts only at κ^r. The ecosystem externality of Section OA.F.5 widens this band; it does not create it.*

Proof. Winner take all is the corner solution of a linear program read through a monotone transformation. *(1)* With $q_T = 1$ shared by N entrants, per-entrant revenue is at most $\alpha\bar{s}_T/N$ and cost is $K_T + c_{Tm}m$, so entry at $N = 1$ is affordable exactly on $m \in [0, \bar{m}]$; consideration requires the score condition; excluded memes earn zero revenue, so the largest zero-profit effort is $h = 0$, and the condition is self-consistent. Solving $Z_T(\kappa_{T,0}e^{-\beta_T m}) \ge \max\{1, Z_M(0)\}$ for the largest admissible $\kappa_{T,0}$ at $m = \bar{m}$ gives κ^d, using strict monotonicity of Z_T in its argument. *(2)* With memes holding all attention, per-meme revenue is $\alpha\bar{s}_M/N_M$, so the largest nonnegative-profit effort is h^*; a tech entrant is considered only if its score strictly exceeds $Z_M(h^*)$, and the most legible affordable configuration is $m = \bar{m}$; the displayed inequality makes every affordable entry unprofitable ($q_T = 0$, revenue zero, cost at least K_T). *(3)* Z_M strictly increasing and Z_T^{-1} strictly decreasing give $\kappa^r < \kappa^d$; on the band, (1) and (2) hold simultaneously. □

Three uses of the proposition. First, it explains the computational protocol: the finite-menu model smooths the knife edge with signal noise and group heterogeneity (which is why the figures show a staircase rather than a jump), and the deterioration and recovery branches are the finite-menu images of the two fixed points above; the thinness of the interior transition band found in the matching exercise of Section 6.3 is the finite-menu shadow of the winner-take-all limit. Second, it locates the source of path dependence in the conduct feedback itself: cheap salience raises h^*, raises $Z_M(h^*)$, and lowers κ^r one for one in closed form, which is the analytical content of the platform shock. Third, it prices the policy instruments directly: standardization ($\kappa_{T,0} \downarrow$), legibility effectiveness ($\beta_T \uparrow$), and anything that raises $\bar{m}$ move the two thresholds multiplicatively through $e^{\beta_T\bar{m}}$, while audit subsidies enter neither threshold, an analytic restatement of Proposition 7. The limit is a description of the protocol's fixed points, not an equilibrium refinement: individual deviations in the limit economy are degenerate under winner take all, and the game-theoretic existence statements for the finite menu are those of Proposition 18.

OA.E.5 Existence of equilibrium in the full game

The full game nests three layers. (i) *Attention stage:* for any menu and issuer actions, an investor optimum exists by Proposition 1 (continuous concave objective, compact simplex), and by Lemma 7 it is unique under informative signals, so visibilities and issuer revenues are single-valued and continuous in issuer actions by Berge's theorem. (ii) *Issuer actions:* given the menu, each issuer chooses (m_i, h_i) from a compact box (bounded by the zero-profit outlay $\alpha g(1)\bar{s}/\min\{c_{\tau m}, c_{\tau h}\}$); continuity gives existence of best responses. (iii) *Entry:* the timing of the computations is maintained, with salience production adjusting faster than entry: given a realized entry profile e (the set of active tech groups with their legibility choices on the grid $\mathcal{M}$), the meme sector's effort is the conduct value $h_M(e)$, the largest element of the effort grid $\mathcal{H}$ at which expected per-meme profit is nonnegative. Folding the attention stage and the conduct rule into payoffs makes entry a finite normal-form game among the $\bar{N}_T$ potential tech groups, with payoffs $\Pi_g(e)$ given by (13) at the unique attention allocation.

Proposition 18 (Existence in the entry game)**.** *Under informative signals, with actions* $\{\text{out}\} \cup \{\text{enter with } m \in \mathcal{M}\}$ *and the payoffs above:*

1. *The entry game possesses a Nash equilibrium in mixed strategies at every parameter point.*

2. *If the tech groups are symmetric and entrants share a common legibility value m, payoffs depend on the profile only through the entrant count k, through some $\pi(k)$, and a pure-strategy equilibrium exists for every shape of π: setting $k^* := \max\{k \geq 1 : \pi(k) \geq 0\}$, and $k^* := 0$ if the set is empty, the profile in which exactly k^* groups enter is an equilibrium.*

Proof. (1) The game is finite, so Nash's theorem applies. *(2)* Symmetry and Lemma 7 make the attention allocation, the conduct value, and hence the common entrant payoff functions of the count alone. At the profile with $k^* \geq 1$ entrants, an entrant earns $\pi(k^*) \geq 0$, which weakly beats the outside payoff of zero, and a deviating outsider would earn $\pi(k^* + 1) < 0$ by maximality of k^* (there is no outsider when $k^* = \bar{N}_T$). If $k^* = 0$, then $\pi(1) < 0$ and staying out is a best response for every group. □

Purity in part (2) does not extend to heterogeneous groups. Attention competition makes entry decisions strategic substitutes among tech groups, while the conduct response makes them complements: entry lowers the meme sector's zero-profit effort, which raises rivals' entry payoffs. With unequal legibility draws the two forces can leave no profile consistent, and the computations encounter exactly such a band (Section OA.F.4); mixed equilibria

exist throughout by part (1), and the deterioration and recovery branches are the two stable adjustment selections in that region. We do not claim existence of a pure-strategy equilibrium in the continuum-action game: the payoff discontinuity at the exclusion boundary $Z_T = \bar{Z}$ is real (visibility jumps from zero to a positive level as the score crosses the benchmark when $q_0 = 0$), and the numerical implementation handles it by direct search over regimes. The hysteresis analysis in Section OA.F.5 exploits exactly this multiplicity structure rather than assuming it away.

OA.E.6 Linear financing as a benchmark

The linear mapping $F_i = \alpha q_i$ is recovered by setting $g(q) = q$ and $s_i = 1$. All threshold arguments in the main text use only continuity and monotonicity of g and boundedness of s_i. Thus the location of the cutoff can change with the reach technology or conversion rule, but the invisible-innovation channel remains upstream: when q_i approaches zero, $g(q_i)$ also approaches zero.

OA.F Numerical algorithms and replication strategy

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

OA.F.1 Baseline calibration

The numerical exercises use the empirically disciplined parameterization collected in Table 7; anchored values are pinned to Table 4 in the main text, and the three parameters with the most leverage (χ_M, C_V, $\kappa_{T,0}$) are pinned by minimum distance in Section 6.3. Theoretical discussions may normalize $\alpha = 1$; the simulation output uses $\alpha = 300$ throughout.

OA.F.2 Solving the attention stage

For a given menu and issuer actions, the unconditional weights solve the concave program (18) directly:

$$\max_{q \in \Delta(\mathcal{A})} \lambda \, \mathbb{E}_{\boldsymbol{y}} \log \left(q_0 + \sum_g a_{T,g} \bar{S}_g + \sum_j a_{M,j} Z_{M,j} \right),$$

where $a_{T,g}$ is the aggregate weight of tech group g, $\bar{S}_g$ is the group's realized average of $e^{\mu/\lambda}$ over its five signals, and the meme-class aggregates enter through their exact means $Z_{M,j}$ (deep-menu limit; the coefficient of variation of a 500-project class average is below

0.3%). The expectation over the joint distribution of tech signals is evaluated by a deterministic sample-average approximation with common random numbers, and the program is maximized by SLSQP with analytic gradients. Because the program is concave, the Kuhn–Tucker conditions (21) are necessary and sufficient; the solver output is certified by the residuals $\mathbb{E}[e^{\mu_a/\lambda}/D]$, which equal one up to $O(10^{-3})$ sampling error for included options and fall strictly below one for excluded options. Corners are therefore the consideration set itself: no fixed-point iteration, no damped updating, and no ε threshold enter the computation.

OA.F.3 Solving the joint equilibrium

The issuer layer iterates to a stationary point: (i) the meme sector's seeding effort is set to the largest h_M on the grid with nonnegative expected per-meme profit (rent dissipation under free entry into salience); (ii) each active tech group best-responds in its own legibility m_g; (iii) tech groups exit when their optimized profit is negative and, on the recovery branch, re-enter when entry at their best legibility is profitable. Each candidate action profile requires one attention solve, so a full entry equilibrium at one grid point takes a few hundred concave-program solves (seconds on a laptop with warm starts).

OA.F.4 Multiplicity and branch tracing

Because meme salience effort and tech entry are strategic substitutes, the joint entry game can lack a pure-strategy equilibrium in a band of complexity levels: with tech active the meme sector's zero-profit effort is low, which invites tech entry, and with tech absent it is high, which blocks it. All figures therefore report continuation branches, deterioration (exit-only adjustment, initialized tech-active) and recovery (entry-only adjustment, initialized at the deterioration terminal), which are the numerically stable analogues of the locally stable branches discussed in OA.E. Figure 1 shows the deterioration branch of the baseline ($\eta = 0$); Figure 2 shows both branches of the extension ($\eta = 10$).

OA.F.5 Figure 2: extension with ecosystem externality and hysteresis

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

$$\kappa_g^{\text{eff}}(m_g; x_T) = \kappa_g \exp(-\beta_T m_g) \cdot \frac{1}{1 + \eta x_T}, \qquad x_T := \frac{N_T^{\text{eq}}}{\bar{N}_T}, \qquad \eta = 10.$$

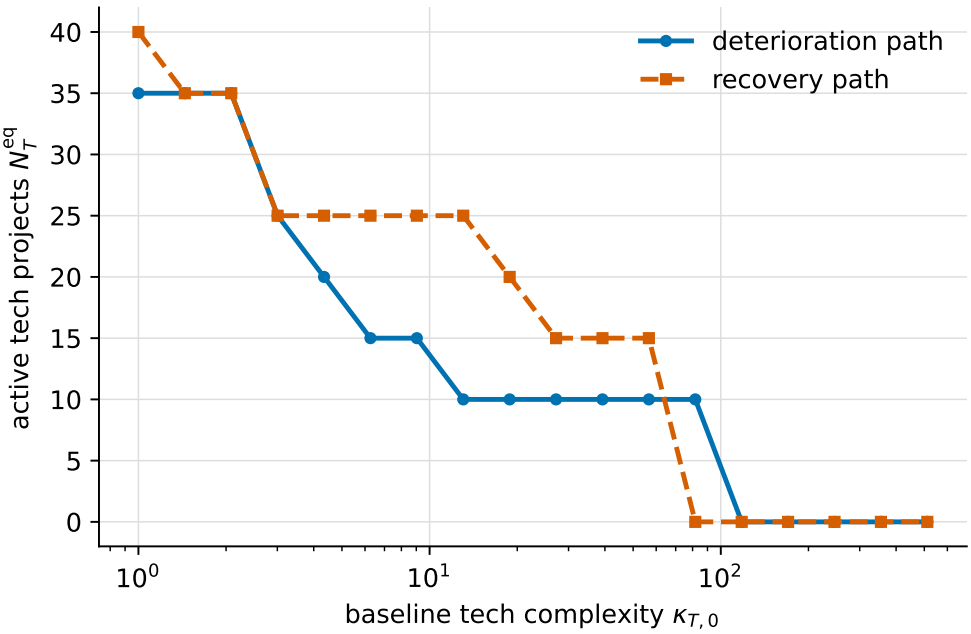


Figure 2: **Extension: hysteresis with ecosystem externality ($\eta = 10$).** Active tech projects N_T^{eq} along the deterioration path (solid) and the recovery path (dashed). The deterioration path collapses at $\kappa_{T,0} \approx 118$; the recovery path remains collapsed until $\kappa_{T,0} \approx 57$, because the meme sector's rent-dissipating salience contest keeps the attention benchmark high in the trap while a thin tech ecosystem keeps effective complexity high. The band $\kappa_{T,0} \in [57, 118]$ is bistable at the trap margin, and partial-count multiplicity extends over a wider range. The deep-menu thresholds of Online Appendix OA.E.4 are the analytic counterpart. This figure should not be read as a property of the baseline $\eta = 0$ model.

This externality captures reusable audit templates, standardized interfaces, developer education, and other forms of ecosystem maturity; x_T is the tech sector's activity rate relative to its potential. None of Propositions 1–5 requires this extension; it is introduced solely to study persistence and path dependence.

We vary $\kappa_{T,0}$ slowly upward and then downward, using the previous grid point's active set as the initial condition, with exit-only adjustment on the way up and entry-only adjustment on the way down. The resulting deterioration and recovery paths trace the locally stable branches.

Interpretation. The hysteresis loop formalizes a trust trap: when a low-tech state itself worsens the processability of technical projects, the economy remains stuck in the meme-dominant branch over a wide range of primitives: here, complexity must fall by a factor of roughly two below the collapse point before recovery begins. The deep-menu limit of Online Appendix OA.E.4 shows that the conduct feedback alone already produces a bistable

band with closed-form thresholds; the ecosystem externality widens the band rather than creating it. Two forces make the trap absorbing. The ecosystem externality raises effective complexity once the sector is thin, and the meme sector's internal salience competition keeps the attention benchmark high even after tech has exited, so a lone re-entrant faces both a degraded processing environment and a hostile benchmark. Once invisible innovation becomes systemic, rebuilding the market's processing infrastructure requires a much larger intervention than the shock that caused the collapse.

OA.F.6 Figure 3: verification alpha and market composition

Figure 3 studies how the private return to costly verification changes as the meme side strengthens. The experiment lowers the meme hype cost c_{Mh2} on a grid while holding everything else fixed: the model analogue of the pump.fun shock studied empirically in Section 5, which made virality production dramatically cheaper for every token at once.

Why the salience-cost path. A pure increase in meme supply N_M is the wrong experiment in this model: under free entry the salience contest dissipates rents, so more memes mean thinner per-meme rents and weaker individual hype, and tech survives a larger meme count. In the computations, cutting N_M to 2,000 collapses tech entry (richer memes hype harder, $h_M = 3.75$) while raising it to 20,000 leaves tech active. What weakens screening is the cheapness of salience production rather than the count of memes, and cheap salience is what launch platforms supplied.

Investor heterogeneity and verification alpha. Expert investors face the same processing environment as the market but a lower audit cost, $C_V^E = 0.02$ versus $C_V = 0.1$. For an expert, define verification alpha as the option value of the audit on the expert's consideration set,

$$\alpha^{\text{ver}} := \mathbb{E}\Big[\max\{W^V - W^{NV},\, 0\} \mid i \in \mathcal{C}^*,\ \text{expert}\Big],$$

computed at each equilibrium along the salience-cost path.

Interpretation. The figure shows the two channels through which meme dominance destroys the private value of expertise. Upstream, complex projects are filtered out of the consideration set, so audits cannot be deployed where they would be most informative. Downstream, the projects that do populate the consideration set are hype-shielded: their padded downside pushes them past the upper edge of the audit window, so even a cheap audit is worthless on them. The equilibrium also displays a substitution the theory predicts: active

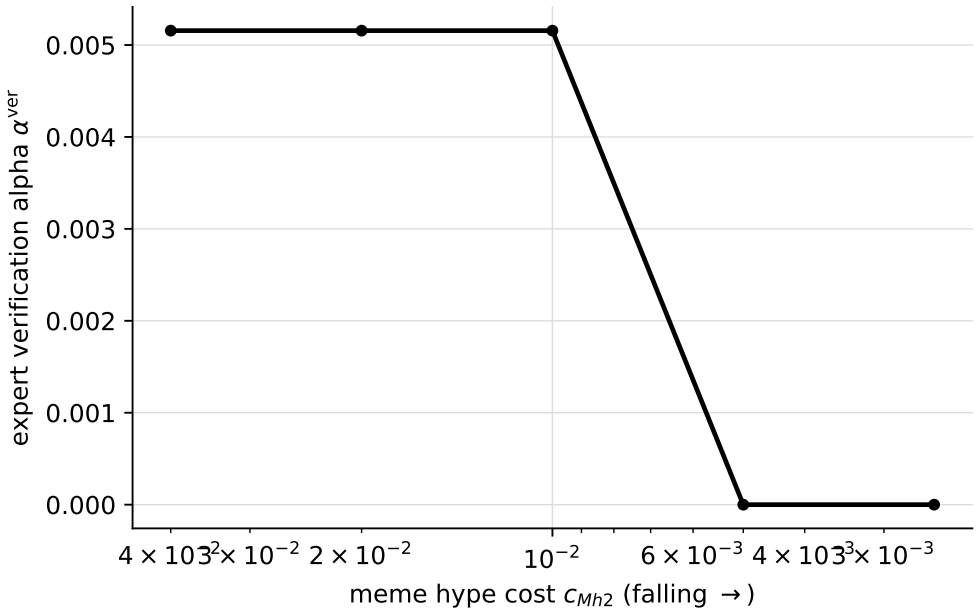


Figure 3: **Verification alpha as salience production gets cheaper.** The expert investor's option value of verification, α^{ver}, along a path that lowers the meme hype cost c_{Mh2} (cheaper to the right). While tech projects populate the consideration set, expert audits carry positive option value ($\alpha^{\text{ver}} = 0.0052$); once cheap salience lets the meme sector take over the consideration set ($c_{Mh2} \leq 0.005$, $h_M \geq 3.25$), the option value drops to zero, because hyped memes sit outside the audit window of Proposition 5(1).

tech issuers defend themselves with legibility rather than audits, so *retail* audit activity is approximately zero even in the tech regime, and verification value survives only for experts with low C_V. Diligence does not scale at the market level; expertise retains value only while complex projects remain visible. Detailed diagnostics are reported in Section OA.F.9.

OA.F.7 Figure-by-figure numerical settings

The numerical settings behind each figure and the matched table are:

1. **Figure 1 (baseline complexity trap):** $\eta = 0$; vary $\kappa_{T,0}$ on the geometric grid $[0.5, 80]$ (17 points), deterioration branch.

2. **Figure 2 (extension with hysteresis):** $\eta = 10$; vary $\kappa_{T,0}$ on the geometric grid $[1, 512]$ (18 points), forward then backward with the branch protocol above.

3. **Figure 3 (verification alpha):** vary the meme hype cost c_{Mh2} over the five-point grid from 0.04 down to 0.0025; at each equilibrium compute $\alpha^{\text{ver}} = \mathbb{E}[\max\{W^V - W^{NV}, 0\} \mid$

$i \in \mathcal{C}^*$, expert] with $C_V^E = 0.02$. The N_M-path robustness runs are produced in the same script.

4. **Table 5 (moment matching):** the grid-plus-refinement search runs at $S = 2{,}000$ and re-solves the winner and its neighborhood at $S = 20{,}000$; the first-stage search documents why $\kappa_{T,0}$, not c_{Tm}, is the third free parameter.

OA.F.8 Reported equilibrium issuer strategies and candidate-set composition

To make the numerical results independently auditable, we report the equilibrium issuer choices and consideration-set counts that underlie the main-text regimes and policy counterfactuals. These objects are mechanically upstream of the verification behavior reported in Tables 9 and 6; without them, readers cannot map the audit-window condition in Proposition 5 to the reported frequencies.

OA.F.9 Figure 3 diagnostics: exact experiment and the N_M path

The primary Figure 3 path holds all primitives fixed and lowers the meme hype cost c_{Mh2} on the grid $\{0.04, 0.02, 0.01, 0.005, 0.0025\}$, tracing the deterioration branch. At each point we record the expert's consideration-set composition and α^{ver} (the option value of the audit at $C_V^E = 0.02$).

Why Table 6 is intentionally investor-side. The main-text policy table is reported as an investor-side counterfactual because the issuer side depends on the full attention-to-conversion pipeline $F_i = \alpha g(q_i) s_i$. Different reach and conversion technologies can change the transfer-clean aggregate welfare object. We therefore report the investor-side metric in the main text and treat broader welfare accounting as model-dependent.

OA.G Robustness and alternative specifications

This section reports robustness checks and alternative specifications. The goal is to verify that the paper's qualitative mechanisms, namely (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 Strategic obfuscation and pooled complexity

The baseline treats high complexity as structural. This is the conservative case for genuine innovation because it gives complex projects the benefit of the doubt: investors do not begin by assuming that complexity is malicious. Now allow bad issuers to choose strategic obfuscation $\omega_i \geq 0$:

$$\kappa_i = \kappa_{\tau_i,0} \exp(-\beta_{\tau_i} m_i + \omega_i)\,. \tag{44}$$

The obfuscation choice raises effective complexity but does not improve fundamentals. A bad issuer's cost can be written as

$$C_B(m, h, \omega) = K_B + c_{Bm} m + c_{Bh} h + c_{B\omega} \omega,$$

with $c_{B\omega}$ potentially below the cost of genuine legibility or engineering quality. Thus strategic complexity can be cheaper than making the project genuinely easier or safer to evaluate.

Let C denote the high-complexity pool, for example projects with $\kappa_i \geq \bar{\kappa}$. Suppose investors observe membership in this pool but cannot perfectly distinguish structural complexity from strategic obfuscation. If N_G^C is the number of good structurally complex projects, N_B^C is the number of bad projects that are structurally complex even without obfuscation, and N_O is the number of strategically obfuscated bad projects, then a simple pooled prior is

$$p_C(N_O) = \frac{N_G^C}{N_G^C + N_B^C + N_O}, \qquad \frac{\partial p_C}{\partial N_O} < 0. \tag{45}$$

More generally, any inflow of obfuscated projects whose quality is below the incumbent high-complexity pool lowers the posterior quality of that pool.

For a representative high-complexity project, the pool's processed signal y has accuracy $\rho(\kappa)$, and the pool posterior means are

$$\mu_C(y; p_C) = \Pr(G \mid y; p_C)\, U(G; H) + \Big[1 - \Pr(G \mid y; p_C)\Big] U(B; H),$$

where the posterior is computed from the pooled prior p_C. The corresponding pool score is

$$Z_C(\kappa, H, p_C) = \sum_{y \in \{g,b\}} P(y; p_C) \exp\left(\frac{\mu_C(y; p_C)}{\lambda}\right), \qquad P(g; p_C) = p_C\rho + (1-p_C)(1-\rho). \quad (46)$$

Obfuscation therefore has two effects. First, by raising effective complexity it lowers state discrimination (the spread of μ_C). Second, by increasing the mass of bad projects in the high-complexity pool it lowers p_C and hence the pooled score, by the following lemma.

Lemma 10 (The pool score is increasing in pool quality)**.** *Fix κ (hence $\rho \in (1/2, 1]$) and H, and suppose $U(G; H) > U(B; H)$. Then $Z_C(\kappa, H, p_C)$ is strictly increasing in p_C on $(0, 1)$.*

Proof. Write $Z_C = P(g)e^{\mu_C(g)/\lambda} + (1 - P(g))e^{\mu_C(b)/\lambda}$. Three derivatives are signed. (i) $\partial P(g)/\partial p_C = 2\rho - 1 \geq 0$, and $\mu_C(g) \geq \mu_C(b)$, so shifting weight toward the larger exponential raises Z_C. (ii)–(iii) The posteriors $\Pr(G \mid g; p_C)$ and $\Pr(G \mid b; p_C)$ are strictly increasing in the prior p_C (monotone likelihood ratio), so both $\mu_C(g)$ and $\mu_C(b)$ are strictly increasing in p_C given $U(G; H) > U(B; H)$. All three effects are nonnegative and the posterior effects are strictly positive for $p_C \in (0, 1)$, so $\partial Z_C/\partial p_C > 0$. □

Corollary 6 (Strategic obfuscation strengthens the complexity trap)**.** *Suppose some bad issuers can choose $\omega_i \geq 0$, which raises effective complexity but does not improve fundamentals. If investors cannot perfectly distinguish structural complexity from strategic obfuscation, then an increase in obfuscated bad entry lowers the posterior quality p_C of the high-complexity pool, and by Lemma 10 strictly lowers the pool score Z_C. By Proposition 2, a lower score moves the pool toward the exclusion region $Z_C \leq \bar{Z}$: strategic obfuscation weakly shrinks the set of considered genuine complex projects and enlarges the region in which they cannot trigger verification.*

Proof sketch. Equation (45) gives $N_O \uparrow \Rightarrow p_C \downarrow$; Lemma 10 gives $Z_C \downarrow$; the attention-score-benchmark characterization gives the consideration consequence, since the benchmark $\bar{Z}$ is determined by the competing menu and is unaffected by the pool's own score when the pool is excluded. □

Remark 6. *Working at the signal level matters here: the state-level attenuation factor $A(\kappa, p)$ is non-monotone in p (it peaks at $p = 1/2$), so monotonicity of a state-level pool score in p_C is genuinely ambiguous. At the signal level the ambiguity disappears because a higher prior raises both posterior means and the weight on the good report simultaneously.*

OA.G.2 Alternative information-cost specifications and the RI-over-states reduced form

G.2.1 Signal-level baseline vs. the RI-over-$\boldsymbol{\theta}$ reduced form. The baseline places the discrete-choice RI problem, with a *common* Shannon/KL scale parameter λ, over the processed signals $\boldsymbol{y}$, with payoff index $\mu_i(y_i)$; project complexity enters through signal accuracy. A tractable reduced form instead places the RI problem directly over the states $\boldsymbol{\theta}$ with the derived index $v_i(\theta_i; H_i, m_i) = \mathbb{E}[\mu_i(y_i) \mid \theta_i]$ of (8)–(10); a multiplicative form u_i/κ_i further conflates level and state-sensitivity effects and is dominated by the v_i-form. The signal-level and v_i-form scores are ordered by conditional Jensen,

$$\mathbb{E}_\theta\left[e^{v_i(\theta)/\lambda}\right] = \mathbb{E}_\theta\left[e^{\mathbb{E}[\mu_i(y_i)|\theta]/\lambda}\right] \;\leq\; \mathbb{E}_\theta\,\mathbb{E}\left[e^{\mu_i(y_i)/\lambda} \mid \theta\right] = Z_i,$$

with equality only when the signal is degenerate, so the reduced form uniformly understates attention scores. The two specifications share the mean payoff $\bar{R}_\tau + \phi_\tau$, the $\kappa \to \infty$ limit Z_τ^∞, strict monotonicity in κ_τ, and the pooling bound with the same attenuation factor A_i. No reported number relies on the reduced form: the numerical implementation solves the signal-level program directly (Online Appendix OA.F.2), and the reduced form appears only in this comparison. Two reasons make the comparison useful: analyses built on the reduced form are conservative, because it uniformly understates scores; and under the v_i-form the attention rule conditions directly on $\boldsymbol{\theta}$ and is not feasible under the comprehension ceiling, so the signal-level formulation is the one on which the internal consistency of the RI–CSV interface (Lemma 9) rests.

As robustness checks, we consider alternative information-cost formulations, including an unweighted mutual-information benchmark $\widetilde{\mathcal{K}}(\pi) = \lambda I(\boldsymbol{y}; a)$ 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.2.2 The myopic attention variant. A second variant keeps the signal-level structure and the Shannon cost but scores options by the posterior valuation μ_i instead of the evaluation value $\mathcal{W}_i$: attention without anticipation of the audit and walk-away options. Every structural result holds (the proofs in OA.B are written for a generic index), the scores factor as (17), and the salience comparative statics become exact. Quantitatively the trap bites much sooner: the myopic investor does not credit complex projects with the option value of evaluating them, so their scores are lower and the deterioration-branch collapse of Figure 1

moves from $\kappa_{T,0} \approx 42$ under the baseline to $\kappa_{T,0} \approx 2.2$ under the myopic variant. The baseline is therefore the conservative specification for the trap, and the myopic variant the conservative one for the market's health. The moment-matching exercise of Section 6.3 runs under this variant, for the reason recorded there.

A bounded-anticipation bridge quantifies the distance between the two indices. Let the attention index be $\mathcal{W}_\alpha(y) := \alpha\,\mathcal{W}(y) + (1-\alpha)\,\mu(y)$ for $\alpha \in [0,1]$: a convex combination of a convex and an affine function of the posterior, hence inside the general class of Online Appendix OA.C.5, nesting the baseline at $\alpha = 1$ and the myopic variant at $\alpha = 0$ (the implementation reproduces both nested cases to machine precision). Re-running the matching protocol with α free adds no fit. Interior allocations exist only for anticipation weights below a complexity-dependent collapse threshold ($\alpha \approx 0.0125$ at the matched complexity, ≈ 0.105 at $\kappa_{T,0} = 5.75$), at which the meme sector's zero-profit hype effort drops discontinuously to zero and the deep-menu regime logic takes over; within the band the targeted moments are nearly invariant in α, so the data cannot distinguish no anticipation from any weight below the threshold, and the four-parameter search cannot improve on the myopic fit without degrading the untargeted pre-state check. Two conclusions carry to the text. The myopic calibration is not a knife edge in the anticipation weight: at $\alpha = 0.01$ the matched interior equilibrium and the welfare numbers of Section 7.4 move by less than a percentage point. And the baseline index's regime structure is a measured threshold rather than a binary index choice: attention that anticipates the evaluation option with more than a few percent weight cannot sustain the observed interior composition. The bridge also explains why attention-scale heterogeneity cannot substitute for bounded anticipation: on the matched menu the baseline-index score ordering is invariant to λ, with the top tech group's score exceeding the best meme class's at every $\lambda \in [0.3, 5]$, so every scale picks the same winner and mixtures of scales inherit the corner, while the myopic index reorders exactly near $\lambda = 1$, where the matched interior sits.

G.2.3 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.3 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 attenuation factor $A_i(\kappa_i, p_{\tau_i})$ is small, so 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.4 Alternative candidate-set definitions

The baseline characterizes consideration endogenously as the support $\mathcal{C}^* = \{i : q_i > 0\}$ of the optimal attention rule (Proposition 2); the threshold set $\mathcal{C}(\varepsilon) = \{i : q_i \geq \varepsilon\} \subseteq \mathcal{C}^*$ is the capacity-motivated and computational proxy. We consider two further alternatives.

G.4.1 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, support, threshold, and top-K sets coincide for a wide range of (ε, K).

G.4.2 Soft consideration weights. Instead of hard inclusion/exclusion, define verification eligibility weights $w_i = \min\{1, q_i/\varepsilon\} \in [0, 1]$ and let verification occur with probability w_i conditional on triggers. This preserves the core insight: as κ_T rises and meme salience expands, tech weights w_i shrink, reducing the scope for CSV discipline even when verification would be valuable conditional on attention.

OA.G.5 Investor heterogeneity

Section OA.F.6 introduces expert investors with lower information and/or verification costs. We extend heterogeneity along three dimensions.

G.5.1 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.5.2 Heterogeneous verification costs. Let $C_V^E < C_V^R$. Experts verify more frequently and earn positive verification option value while complex projects populate the consideration set, but once cheap salience lets the meme sector take over, the considered projects are hype-shielded and the option value drops to zero, consistent with Figure 3. The two paths for strengthening the meme side are not interchangeable: the salience-cost (platform) shock produces the takeover, while a pure N_M expansion is self-limiting under free entry because it dilutes per-meme rents and weakens individual hype (Table 12); salience and market composition affect the RI denominator through different channels.

G.5.3 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.6 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.6.1 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.6.2 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.7 Verification technology and trigger design

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

G.7.1 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.7.2 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.8 Entry dynamics, equilibrium selection, and hysteresis robustness

G.8.1 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.8.2 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 OA.F.5).

OA.G.9 Alternative reach functions and funding conversion

We verify the threshold logic under $g(q) = q^\eta$, $g(q) = 1 - e^{-\eta q}$, and $g(q) = \max\{q - \bar{q}, 0\}$. The location of the cutoff changes with the reach technology, but the invisible-innovation channel persists because $g(0) = 0$ and conversion is bounded. In all cases, projects with $q_i \to 0$ have $F_i = \alpha g(q_i) s_i \to 0$.

OA.G.10 Dynamic hype feedback

We vary ρ_M holding $\rho_T = 0$. Larger ρ_M increases meme-side score through H_M, steepens the denominator crowding effect, and can generate multiple salience states when the feedback multiplier in Proposition 12 approaches one.

OA.G.11 Numerical robustness: sample size, tolerances, and grids

The pipeline is deterministic given the common-random-number seed, and the only sampled object is the joint distribution of the tech signal vectors. We verify that Figures 1–3 are stable with respect to:

- the number of sample-average draws S: doubling S from 20,000 to 40,000 leaves every entry count unchanged and moves the reported continuous statistics (attention benchmark, investor payoffs) by less than 10^{-3};
- the solver tolerance: the Kuhn–Tucker residuals of included options are $1 \pm O(10^{-3})$, so consideration-set membership is unambiguous at the reported parameter points;
- the action grids: entry decisions away from the meme-ignition zone are unchanged when the legibility and hype grids are coarsened from step 0.25 to 0.5; near the ignition point the meme-effort grid must be fine enough to resolve the sector's zero-profit root ($h_M = 2.25$ at the baseline), and a coarser grid that misses the root delays the ignition and shifts the collapse point outward by one to two grid steps.

Across these variations, the qualitative objects, the staircase collapse in Figure 1 and the bistable band in Figure 2, are preserved; the location of the collapse point is sensitive only to the resolution of the meme-effort grid, in the direction just described.

OA.H Welfare: proofs for Section 7

OA.H.1 Generality of the attention machinery

The proofs in OA.B and OA.C use only three properties of the payoff index: (i) option a's index depends on $\boldsymbol{y}$ only through the own signal y_a; (ii) indices are bounded; (iii) signals are independent across options. Any index with these properties, in particular the CSV continuation value

$$\mathcal{W}_i(y_i) = \max\left\{W_i^V(p_i(y_i)),\, W_i^{NV}(p_i(y_i))\right\},$$

which is a bounded function of y_i, therefore inherits Proposition 1 (logit, concave program, Kuhn–Tucker conditions) and Proposition 2 (score-benchmark characterization) verbatim, with scores $Z_i = \mathbb{E}[e^{\mathcal{W}_i(y_i)/\lambda}]$ and benchmark $\bar{Z}$. In the baseline this is automatic, because $\mathcal{W}_i$ *is* the attention index; the generality matters for the myopic robustness variant, whose welfare analysis re-solves the program at $\mathcal{W}_i$. Throughout this appendix, $\mathcal{V}(\mathcal{N})$ denotes the value of the concave program for menu $\mathcal{N}$ under this index.

OA.H.2 Menu monotonicity

Lemma 11 (Menu monotonicity of the investor value)**.** *For any menu $\mathcal{N}$ and project $i \notin \mathcal{N}$:*

1. *$\mathcal{V}(\mathcal{N} \cup \{i\}) \geq \mathcal{V}(\mathcal{N})$;*
2. *if $Z_i \leq \bar{Z}(\mathcal{N})$, then $\mathcal{V}(\mathcal{N} \cup \{i\}) = \mathcal{V}(\mathcal{N})$ and the enlarged program has an optimum with $q_i = 0$;*
3. *if $Z_i > \bar{Z}(\mathcal{N})$, then $\mathcal{V}(\mathcal{N} \cup \{i\}) > \mathcal{V}(\mathcal{N})$ and every optimum of the enlarged program has $q_i > 0$.*

Proof. (1) Extending any feasible q by $q_i = 0$ leaves the objective unchanged, so $\mathcal{V}(\mathcal{N} \cup \{i\}) \geq \mathcal{V}(\mathcal{N})$. (2) Let q^* be optimal for $\mathcal{N}$, extended by $q_i = 0$; $D(\boldsymbol{y})$ is unchanged, so the enlarged program's Kuhn–Tucker conditions are the original ones plus, by factorization,

$$Z_i \, \mathbb{E}\left[\frac{1}{D(\boldsymbol{y})}\right] \leq 1 \iff Z_i \leq \bar{Z}(\mathcal{N}),$$

which holds by hypothesis; sufficiency for the concave program gives optimality of the extension and equality of values. (3) If $Z_i > \bar{Z}(\mathcal{N})$, the extension violates i's condition, so the (feasible) extension is suboptimal and the enlarged maximum is strictly larger; any candidate optimum with $q_i = 0$ restricts to an optimum for $\mathcal{N}$ and reproduces the violated condition, so every optimum has $q_i > 0$. □

OA.H.3 Proof of Proposition 6

The decomposition of $\Delta\mathcal{W}_i^{\text{full}}$ is immediate from (30). Part (1) is Lemma 11: $\Delta\mathcal{V}_i \geq 0$, strictly on $\{Z_i > \bar{Z}\}$; since the wedge is absent from $\Delta\Pi_i$, any considered project with

$$F_i \;<\; C_{\tau_i} \;<\; F_i + \alpha\Delta\mathcal{V}_i$$

is socially but not privately viable. Part (2) is the transfer accounting recorded in the text. For part (3), choose $\bar{R}_M < 0$ and $\phi_M(H_M)$ large enough that, at both signal realizations,

$$p_M(y)\left(V_M^G + \phi_M\right) + (1 - p_M(y))\left(V_M^B + \phi_M\right) > 0 \qquad \text{and} \qquad (1 - p_M(y))\left(-V_M^B - \phi_M\right) < C_V,$$

so the entrant is considered ($Z_M \to \infty$ as $\phi_M \to \infty$), attracts mass $q'_M > 0$, and is always funded without audit. Experienced welfare evaluates the same policy at payoffs stripped of ϕ: attention totaling q'_M is diverted from options with nonnegative experienced value toward

an option with experienced mean $\bar{R}_M < 0$, so

$$\Delta\mathcal{V}_M^{\text{exp}} < 0, \qquad \Delta\Pi_M = F_M - K_M > 0 \quad \text{for } K_M \simeq 0. \qquad \square$$

OA.H.4 Proof of Proposition 7

Part (1). In the trap regime $Z_T \leq \bar{Z}$, Lemma 11(2) gives $\Delta\mathcal{V}_T = 0$ and $q_T = 0$ at an optimum of the enlarged program, so $F_T = \alpha g(0) s_T = 0$ and

$$\Delta\mathcal{W}_T^{\text{full}} = \alpha \cdot 0 + \Big(0 - C_T(m_T, h_T)\Big) \leq -K_T < 0.$$

An entry subsidy is a transfer and does not change this social calculation.

Part (2). The index $\mathcal{W}_a(y_a; C_V)$ is piecewise linear and nonincreasing in C_V, with slope -1 exactly where verification is strictly optimal and 0 elsewhere. By Danskin's theorem, at points of differentiability,

$$\frac{\partial\mathcal{V}}{\partial C_V} = -\,\mathbb{E}\Big[\sum_a \pi^*(a \mid \boldsymbol{y})\,\mathbf{1}\{\text{verification optimal at } (a, y_a)\}\Big],$$

which loads only on options with $q_a > 0$; excluded projects contribute zero.

Part (3). As κ_i falls, the distribution of $p_i(y_i)$ is a mean-preserving spread, $\mathcal{W}_i$ is convex in the posterior (a maximum of affine functions), and $e^{x/\lambda}$ is convex increasing, so Z_T rises continuously and strictly by the argument of Lemma 1. Once $Z_T > \bar{Z}$, Lemma 11(3) gives $\Delta\mathcal{V}_T > 0$. $\square$

OA.H.5 The intervention ordering without Shannon

The ordering of Proposition 7 does not depend on the Shannon specification of the attention cost. Fix the comprehension channel and the verification stage of the baseline, and let the investor choose any feasible stochastic rule π mapping her processed information into an action in $\mathcal{A}$ and a verification decision, at cost $\mathcal{K}(\pi)$. Write $\Pi(\mathcal{N})$ for the feasible rules on menu $\mathcal{N}$ and

$$V(\mathcal{N}; C_V) := \max_{\pi \in \Pi(\mathcal{N})} \mathbb{E}_\pi\Big[\mathcal{W}(a, y_a; C_V)\Big] - \mathcal{K}(\pi),$$

where $\mathcal{W}(a, y_a; C_V) = \max\{W_a^V(p_a(y_a); C_V),\, W_a^{NV}(p_a(y_a))\}$ is the CSV continuation value. Only two structural properties are used:

(P1) *Free extension.* Any $\pi \in \Pi(\mathcal{N})$ extends to $\Pi(\mathcal{N} \cup \{i\})$ by never choosing i, with $\mathcal{K}$ unchanged: ignoring an option costs nothing.

(P2) *Cost-free verification margin.* $\mathcal{K}$ does not depend on C_V.

Property (P1) holds for the Shannon cost, for every uniformly posterior-separable cost, and for any cost that is a function of the induced joint distribution of (information, action); (P2) is definitional.

Proposition 19 (Cost-free intervention ordering)**.** *Assume the maximum in V is attained. Then:*

1. (Menu monotonicity and unconsidered entrants.) *$V(\mathcal{N} \cup \{i\}; C_V) \geq V(\mathcal{N}; C_V)$, with equality if and only if some optimum of the enlarged problem ignores i. Under equality, marginal entry of i changes investor value by zero and social value by $F_i - C_{\tau_i} = -C_{\tau_i} < 0$, so entry subsidies are worthless at the margin, for every cost functional $\mathcal{K}$.*

2. (Audit subsidies.) *$\mathcal{W}(a, y_a; \cdot)$ is a maximum of affine functions of C_V, hence convex and nonincreasing, so $V(\mathcal{N}; \cdot)$ is convex in C_V and, at every point of differentiability,*

$$\frac{\partial V}{\partial C_V} = -\,\mathbb{E}_{\pi^*}\Big[\mathbf{1}\{\textit{verification optimal at } (a, y_a)\}\Big],$$

   which loads only on options chosen with positive probability. A marginal audit subsidy has zero value on any project ignored at the optimum, for every $\mathcal{K}$.

3. (Processability.) *If an intervention changes project i's comprehension channel so that* every *optimum of the enlarged problem uses i with positive probability, then V rises strictly.*

Proof. (1) Weak inequality is (P1). If an optimum ignores i, its restriction is feasible and optimal for $\mathcal{N}$, giving equality; conversely, if V is unchanged, the extension by (P1) of an optimum for $\mathcal{N}$ attains the enlarged maximum and ignores i. With $q_i = 0$, $F_i = \alpha g(0) s_i = 0$. (2) For fixed π, $C_V \mapsto \mathbb{E}_\pi[\mathcal{W}] - \mathcal{K}(\pi)$ is convex nonincreasing by (P2) and the affine-maximum structure of $\mathcal{W}$; V is a pointwise maximum of such functions, and Danskin's theorem gives the derivative as that of any active selection, which is the expected mass of realized audits. Options with zero choice probability contribute zero mass. (3) If $V(\mathcal{N} \cup \{i\}) = V(\mathcal{N})$, then by the converse in (1) some optimum ignores i, contradicting the hypothesis; so V rises strictly. □

The Shannon specification therefore contributes the *explicit* objects (the score Z_i, the benchmark $\bar{Z}$, the crowding comparative statics of Proposition 3) but not the ordering itself: that entry and audit subsidies cannot reach an unconsidered project, while processability interventions can, is a property of consideration, not of entropy.

OA.H.6 Remark on the free-riding margin

With a continuum of investors, a single audit is non-rival information. The baseline assigns each considered project at most one audit and charges its cost once, which is the planner's cost-minimizing arrangement; decentralized markets may instead duplicate audits or underprovide them (Grossman and Stiglitz, 1980). Either deviation weakly lowers $\mathcal{V}$ relative to the baseline accounting, so internalizing verification (for example, through a certification utility) is a complementary instrument to the processability interventions of Proposition 7(3); it cannot substitute for them, because it operates only inside the consideration set.

OA.I Cross-chain extensions of the empirical design

Meme launchpads arrived on other chains after pump.fun: Ape.Store on Base (May 2024), four.meme on BNB (mid-2024, with its volume takeoff in February 2025), Wonton and GasPump on TON (2024H2), DegenHive on Sui and MovePump on Aptos (late 2024), and Arena Launch and Token Mill on Avalanche (mid-2025). The staggered design now uses the same event-time audit outcome and first-archive chain labels as Table 1. With protocol chain and cohort-half-year fixed effects, Ethereum, Arbitrum, Optimism, Polygon, and Fantom as controls, and the first listed deployment chain used for multi-chain protocols, the pooled coefficient is -0.040 (HC2 standard error 0.048; chain-clustered standard error 0.041). Permuting the seven event dates and five untreated labels across the twelve chains gives directional $p = 0.341$ and two-sided $p = 0.471$. Excluding Solana gives -0.012 (chain-clustered $p = 0.705$).

The closest sensitivity assigns a protocol to every deployment chain, restricts the sample to structurally intricate categories, and dates BNB by its February 2025 volume takeoff. It gives -0.154 (HC2 standard error 0.100; chain-clustered standard error 0.099), with directional event-assignment $p = 0.050$ and two-sided $p = 0.080$. This cell was selected from the grid: its directional p becomes 0.54 under Holm correction over the eleven event-assignment specifications that were taken to the high-draw stage. A stacked not-yet-treated estimator produces small protocol-level standard errors only because the same control protocols recur across event stacks; chain-level inference does not support those values, so they are logged but rejected. Across the 324-cell grid, the main staggered conclusion remains a null.

This pattern is consistent with attention being allocated over a global menu: the benchmark $\bar{Z}$ of Proposition 2 is a market-level object, so the original salience flood can reprice attention for every chain at once, while later clone launches relocate production of an already-abundant input. The same null also admits a Solana-idiosyncratic decline. We therefore

measure the salience flood directly as the share of each chain's DEX volume generated by launchpad-origin tokens.

The exposure numerator uses token-classified DEX volume where those data exist. Dune query 3919233 tags Solana pairs that ever traded on pump.fun and retains the tag after graduation; query 4319076 supplies Clanker-origin volume on Base, supplemented by Ape.Store; query 5212405 supplies four.meme bonding-curve volume on BNB. DefiLlama launchpad series cover Avalanche, TON, and Aptos. The denominator is total chain DEX volume, using the source-matched Dune total for Solana and DefiLlama elsewhere. The resulting half-year Solana path is 0.0, 7.8, 28.6, 23.7, and 7.5 percent from 2023H2 through 2025H2; the largest post-2024 half-year share on another measured chain is 2.1 percent. BNB omits post-graduation PancakeSwap trading, Base omits some launchpads, and the unavailable small Sui series is coded zero, so the non-Solana measures are lower bounds.

For the baseline continuous specification, the estimating equation is

$$\text{audited}_i = \theta\,\text{Dose}_i + \mu_{\text{chain}(i)} + \delta_{c(i)} + \lambda_{\text{cat}(i)} + \varepsilon_i, \qquad \text{Dose}_i = \frac{\sum_{m=1}^{24} 2^{-(m-1)/12}\, s_{\text{chain}(i),\,t(i)-m}}{\sum_{m=1}^{24} 2^{-(m-1)/12}}, \tag{47}$$

where $s_{k,t}$ is chain k's meme-flood share of DEX volume in month t, $t(i)$ is protocol i's listing month, and the fixed effects are chain, listing half-year, and category. The Solana-priority assignment retains one row per protocol by assigning every Solana deployment to Solana and otherwise using a fixed chain ordering. Table 14 gives $\hat{\theta}$ scaled to percentage points of audit adoption per ten-percentage-point increase in dose.

The broad-sample estimate falls from -9.16 to -4.38 points when multi-chain protocols are split across deployments, while its next-six-month look-ahead is nearly zero. The intricate prior-month estimate is -16.82 points, but its next-six-month estimate is more negative at -25.28 points. The two intricate rows also rely on one chain with positive measured dose, and neither passes the chain-path test. The broad result therefore provides the more credible continuous corroboration; the intricate result is a sensitivity pattern that cannot separate pre-listing exposure from a persistent Solana time shock. The staggered null and the continuous gradient together remain consistent with a market-level attention benchmark, but they do not rule out a Solana-specific shock.

OA.J Additional evidence exhibits

This appendix first varies the timing of deployment labels and then imposes the Solana-native restriction. It also aligns the complexity decomposition with the promoted Solana-exposure definition, reports the conservative native comparison, plots the specification grid described

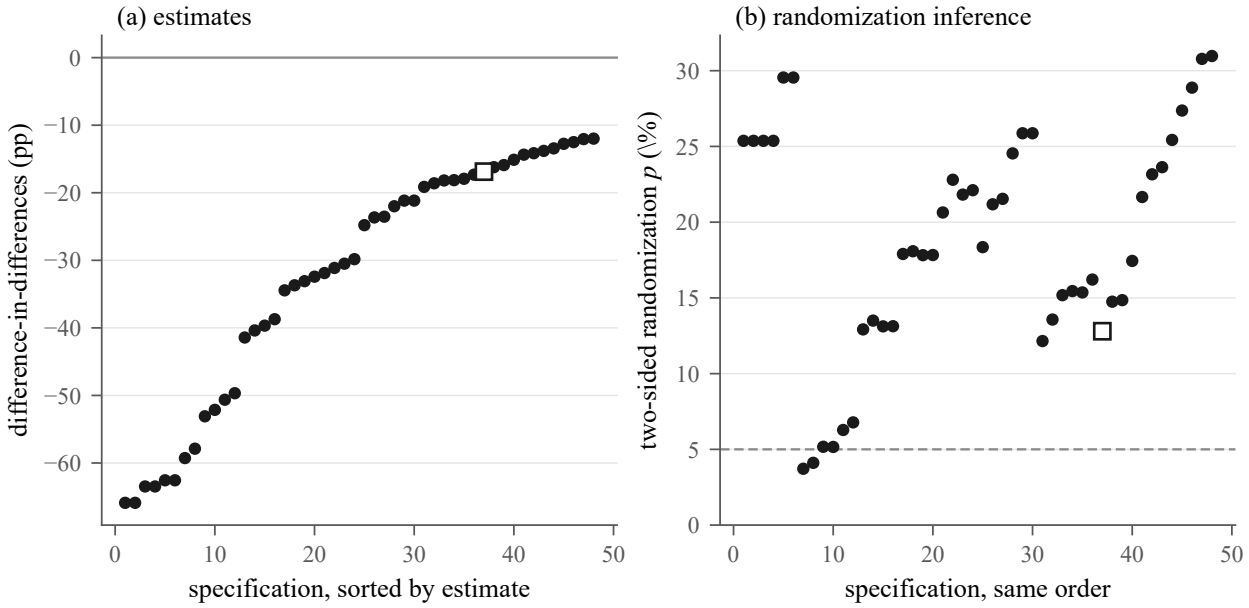


Figure 4: **Solana-native specification grid for the adoption difference-in-differences.** The 48 cells vary the post window (four definitions), the control-chain set (three), the audit coding (two), and category breadth (two), always excluding protocols deployed on both Solana and Ethereum. Panel (a): estimates sorted by size; every cell is negative, the median is −24.8 points, and the hollow square marks the native benchmark, which ranks 37th of 48, on the conservative side of the grid. Panel (b): two-sided protocol-label permutation p of each cell in the same order, with the five-percent rule dashed; the two cells below it restrict jointly to the 2024 cohorts, Ethereum-only controls, and the intricate categories, restrictions chosen after inspecting where the effect concentrates and uncorrected for multiplicity.

in Section 5, and records the chain-level synthetic-control accounting.

Synthetic-control accounting. The donor-weighted synthetic uses Ethereum, Arbitrum, Optimism, Polygon, and Avalanche. Solana's post-period adoption gap averages −13.7 points against placebo gaps of at most 1.9 in absolute value, but the donor pool caps the in-space permutation p at 1/6 and Solana's noisy pre-period (cohort cells in the twenties) rules out the standard fit-ratio test. Of 432 logged variants over donor sets, pre-period lengths, and outcome codings, none reaches $p < 0.05$; the best-fitting six-quarter pre-period variant reduces pre-RMSPE to 5.7 points with a 2024 gap of −14.0 points and in-space fit-ratio $p = 0.111$, and the rankings below 0.10 use only two or three half-year pre-periods and mechanically overfit them.

Table 2: Firm-side audit-report intensity per ten listing opportunities

	Solana side			Ethereum side			
	opps.	reports	intensity	opps.	reports	intensity	ratio
Panel A: Zellic public reports							
2022H1	24	6	2.5	121	22	1.8	1.38
2022H2	16	6	3.8	144	32	2.2	1.69
2023H1	14	0	0.0	282	37	1.3	0.00
2023H2	13	0	0.0	193	49	2.5	0.00
2024H1	21	1	0.5	176	43	2.4	0.19
2024H2	31	3	1.0	141	32	2.3	0.43
2025H1	22	2	0.9	125	29	2.3	0.39
2025H2	36	5	1.4	114	24	2.1	0.66
2026H1	34	1	0.3	121	19	1.6	0.19
Panel B: six balanced public archives							
2022H1–2023H2	402	20	0.50	4,440	612	1.38	0.36
2024H1–2025H2	660	19	0.29	3,336	634	1.90	0.15
Log ratio-in-ratios							−0.868
Poisson cell standard error							(0.325)
z statistic							−2.67
Exact two-sided time-label p							[0.029]
Firm-side FE PPML							−0.725
Standard error clustered by firm							(0.438)
z statistic							−1.65
Two-sided cluster p							[0.098]

Notes: Panel A uses Zellic's month-dated, chain-tagged public reports ($n = 416$), matched to the chain side of the audited codebase. Panel B uses Code4rena, Consensys Diligence, OpenZeppelin, Spearbit, Trail of Bits, and Zellic, the six archives whose dated public coverage spans the full 2022H1–2025H2 window. A listing opportunity is one new complex listing exposed to one covered firm; hence the denominators equal side-specific listings times six in Panel B. Intensity is reports per ten opportunities, and ratio is Solana intensity divided by Ethereum-side intensity. The log ratio-in-ratios is $\log[(19/660)/(20/402)] - \log[(634/3336)/(612/4440)]$. Its Poisson cell standard error conditions on the opportunity counts. The exact test assigns four of the eight common half-years to the pre-period in all $\binom{8}{4} = 70$ ways. The PPML includes firm-side and cohort-half-year fixed effects; its displayed standard error clusters by six firms. No stars are used. The exact p is unadjusted for the logged sample/window search and does not survive a global Holm correction. Halborn, Sherlock, Pashov, and Offside fail the balanced-coverage rule; the all-archive market calculation including Offside remains flat, as discussed in the text. Publication lags engagement by roughly one quarter, and public reports are client-approved.

Table 3: Venture funding composition around the shock, memecoin presales excluded

	Solana side			Ethereum side		
	tech	consumer	tech share	tech	consumer	tech share
2022H1	9	4	69%	34	5	87%
2022H2	5	3	62%	37	10	79%
2023H1	5	3	62%	50	10	83%
2023H2	2	2	50%	30	9	77%
2024H1	9	5	64%	48	14	77%
2024H2†	9	4	69%	40	1	98%

Notes: Deal counts from the DefiLlama raises registry, using an Internet Archive snapshot of December 23, 2024, so the full event window is covered by data recorded contemporaneously. Tech = DeFi/Infrastructure/ZK/Data/Oracle categories; consumer = NFT/Gaming/Web3/Social/Meme. Memecoin presales, the salience flood itself, are excluded here and discussed in the text. †2024H2 is right-censored in the snapshot.

Table 4: Observable anchors for the parameterization

Model object	**Observable anchor (source)**	**Baseline vs. anchor**
Meme fixed cost K_M	Token creation on pump.fun is free; on-chain deployment requires only a first buy plus gas, ≈ 0.04 SOL (Pump.fun Help Center, 2025; Pump.fun, 2026).	$K_M = 0$: pinned.
Audit cost C_V	Standard pre-launch audit \$15k–\$40k, full span \$5k–\$250k+, duration 2–6 weeks, booking lead 6–8 weeks (Sherlock, 2026; Hacken service page, 2026); person-day totals in the ConsenSys Diligence (2025) and OpenZeppelin (2025) audit reports.	Anchors pin ratios: C_V/K_M is three to four orders of magnitude, and C_V is lumpy, not marginal. The level $C_V/V_T^G = 0.1$ is a normalization.
Menu composition N_M/N_T	$\sim$5,300 new tokens per day economy-wide in early 2024 (CoinGecko Research, 2024); 514,000 pump.fun launches in November 2025 alone (Cryptonews, 2025); versus 7,766 DeFi protocols ever tracked by DefiLlama (DefiLlama protocols API).	$N_M/\bar{N}_T = 100$: anchored, and still conservative by at least an order of magnitude relative to the measured ratio.
Meme quality prior p_M	Bonding-curve graduation 1.41% in 2024, below 1% through much of 2025 (Mancino, 2025) (Coinpaper, 2024, from Dune Analytics); $\sim$60% of new tokens inactive within a day (Cernera, La Morgia, Mei, and Sassi, 2023); $\sim$50% of Uniswap listings classified as scams (Xia et al., 2021); pump-and-dump patterns in 98.6% of pump.fun tokens (Solidus Labs, 2025).	$p_M = 0.05$: anchored to the survival metrics.
Verification capacity	Largest auditor's lifetime total $\approx$6,100 projects (CertiK service page, 2026), roughly one *day* of new-token issuance; audited share of newly launched DeFi protocols fell from $\sim$50% (2019) to $\sim$30% (2023) (Rabetti, 2023).	Motivates consideration-gated CSV; ε / top-K as the capacity proxy.
Audit precision ξ	48% of 2022–26 exploit incidents had at least one prior public audit (Beyer, 2026); only 6% of rug-pulled projects in 2023 had any audit (Hacken, 2024).	Baseline $\xi = 1$ is generous to audits; noise strengthens the trap (Section 4.9).
Salience level ϕ_M	Revealed preference given the anchored p_M: buying tokens with $\bar{R}_M = -0.45$ requires $\phi_M \gtrsim 0.45$; near-zero audits among rug pulls (Hacken, 2024) require, via (28), $\phi_M \geq -V_M^B - C_V = 0.4$.	Both bounds hold in equilibrium ($\phi_M \approx 0.97$ for the considered class; all classes 0.49–0.97). Only ϕ_M and ϕ_M/λ enter behavior; $\chi_M = 0.6$ delivers the level.
λ, g, ℓ	Concentration and conversion moments reported for discipline: meme share of Solana DEX volume 67.5% at the 2024 peak versus $\sim$10% in the bear market (SolanaFloor, 2025, from Blockworks); audit-adoption premium of +37% locked value and +42% users in week one (Rabetti, 2023).	Not pinned: swept as robustness dimensions.

Notes: Anchors are order-of-magnitude moments, not estimates; Section 6.3 adds a formal minimum-distance fit of the three parameters the anchors leave free. The 98.6% figure rests on a pump-and-dump and liquidity-death definition; it does not count literal contract rug pulls, which pump.fun's fair-launch design largely precludes (Mancino, 2025) (Solidus Labs, 2025).

Table 5: Minimum-distance fit of the free parameters (myopic attention variant)

	data	model
Targeted moments (2024 platform era)		
Meme share of attention	0.675	0.743
Audit adoption, new complex protocols	0.333	0.414
Funded share of meme projects	0.014	0.010
Matched parameters		
Salience technology χ_M	0.38	
Audit cost C_V	0.06	
Baseline complexity $\kappa_{T,0}$	5.0	
Untargeted (pre state, expensive salience)		
Audit adoption, new complex protocols	0.586	0.610
Change in adoption across the shock	-25.3 pp	-19.6 pp
Meme share of attention	low	0.000
Active tech groups (of 10)	—	$3 \to 1$

Notes: Data columns: meme share of Solana DEX volume at the late-2024 peak (SolanaFloor, 2025, from Blockworks); audit adoption of newly listed complex Solana-exposed protocols at the twelve-month reading, 2024–25 and 2023 cohorts (Table 1); pump.fun bonding-curve graduation rate in 2024 (Coinpaper, 2024, from Dune Analytics). Model: myopic attention variant (under the baseline index every parameter point tried delivers a regime allocation with post-shock meme share zero or one, the deep-menu knife edge of Online Appendix OA.E.4); post state at $c_{Mh2} = 0.0025$ continued on the deterioration branch from the pre state at $c_{Mh2} = 0.04$; $G_M = 100$ virality classes (so the funded share has one-percent resolution, which is why 0.014 is matched at its 0.010 floor); all other parameters at their Table 7 anchors; $S = 20{,}000$. The pre-state adoption level and the adoption decline are not targeted. Two identification limits: the audit-adoption moment is quantized by the surviving groups' signal probabilities, so C_V is interval-identified (every $C_V \in [0.06, 0.10]$ fits identically), and the matched state sits near the collapse boundary ($\kappa_{T,0} = 7$ at these values collapses tech entirely), which is the model's own reading of the 2024 market. The pre-state meme attention share is zero because the model's cold state has no trading motives; the data's bear-market floor of roughly ten percent (SolanaFloor, 2025, from Blockworks) includes them.

Table 6: Fixed-entry investor-side counterfactuals: lowering processability costs versus lowering verification costs

	Baseline ($\kappa_{T,0} = 42.4$, $C_V = 0.1$)	**Lower κ_T / Higher β_T** ($\boldsymbol{\kappa_{T,0} = 0.5}$, $C_V = 0.1$)	**Lower C_V** ($\kappa_{T,0} = 42.4$, $\boldsymbol{C_V = 0.02}$)
Tech projects present (fixed)	50	50	50
Average tech visibility $\bar{q}_T$	0.0000	**0.0200**	**0.0200**
Consideration set: tech projects	0	**25**	**10**
Candidate-set tech fraction	0.0%	**100.0%**	**100.0%**
Meme seeding effort h_M	2.25	**0.00**	**0.00**
Average investor payoff (decision)	0.516	**0.589**	**0.580**
Average investor payoff (fundamentals)	−0.450	**+0.589**	**+0.580**
Total verification cost paid	0.0000	0.0000	0.0200

Notes: This table forces all 50 tech projects to be present (no exit) and therefore isolates intensive-margin effects; it should not be interpreted as a full social-welfare ranking, which is the content of Proposition 7.

Both interventions flip the regime, and the difference between them is the content of Proposition 7. Lowering κ_T restores processability directly. Lowering C_V fivefold restores it through the evaluation value: because attention anticipates verification, cheap audits raise complex projects' scores over the benchmark, the consideration set flips to tech (audited with probability one), and the meme contest shuts down endogenously. What the cut buys is not cheaper audits for already-considered projects, whose marginal value in the trap is exactly zero by Proposition 7(2), but consideration itself: verification acting as the certification technology of part (3). Under the myopic attention variant the same cut leaves the consideration set unchanged to four decimal places, which isolates where the anticipation channel does its work.

Verification spending is zero in the first two columns (excluded tech generates no audits, and legible tech in the second column screens without them) and equals C_V per unit of tech attention in the third, where consideration is sustained by the audit option being exercised.

Table 7: Baseline Parameters for Numerical Experiments

Category	Parameter (definition)	Baseline value
Project fundamentals		
Priors	$p_T = \Pr(G \mid T)$, $p_M = \Pr(G \mid M)$	$p_T = 0.6$ (swept); $p_M = 0.05$ (anchored)
Payoffs	(V_T^G, V_T^B), (V_M^G, V_M^B)	$(1.0, -1.0)$, $(0.5, -0.5)$ (normalization)
Salience / hype		
Salience function	$\phi_\tau(H) = \chi_\tau \log(1 + H)$	$\chi_T = 0.1$, $\chi_M = 0.6$
Virality lottery	$H_j = h_M e^{\sigma_M \zeta_j}$, G_M classes	$\sigma_M = 0.35$, $G_M = 10$
Meme hype cost	$c_{Mh2} h_M^2$ (convex)	$c_{Mh2} = 0.01$ (anchored ≈ 0, convex)
Meme conduct	sector effort h_M	largest h_M with $\mathbb{E}[\Pi_M] \geq 0$ (rent dissipation)
Tech hype; feedback	h_T; ρ_τ	$h_T = 0$; $\rho_\tau = 0$ benchmark ($\rho_M > 0$ in extension)
Legibility / processability		
Baseline complexity	$\kappa_{T,0}$ (sweep axis), $\kappa_{M,0}$	$\kappa_{T,0} \in [0.5, 80]$; $\kappa_{M,0} = 0.2$
Tech dispersion	$\kappa_g = \kappa_{T,0} e^{\sigma_T z_g}$, G groups	$\sigma_T = 0.4$, $G = 10$
Legibility technology	β_T; cost $c_{Tm} m_g$ per group	$\beta_T = 0.5$, $c_{Tm} = 0.3$; $m_M = 0$
Processing accuracy	$r_\tau(\kappa) = 2\rho_\tau(\kappa) - 1$	$(1 + \kappa)^{-1/2}$
RI scale	λ	$\lambda = 1.0$ (swept)
Verification (CSV)		
Verification cost	C_V (retail); C_V^E (expert, Fig. 3)	$C_V = 0.1$; $C_V^E = 0.02$
Consideration set	support of the attention rule	corners of (18) (exact)
Market scale and entry		
Reach and conversion	$g(q)$; ℓ	$g(q) = q$; unit-investment cutoff
Market size	α in $F_i = \alpha g(q_i) s_i$	$\alpha = 300$
Fixed costs	K_T, K_M	$K_T = 0.4$; $K_M = 0$ (anchored)
Menu size	$(\bar{N}_T, N_M)$	$\bar{N}_T = 50$, $N_M = 5{,}000$ (anchored ratio)
Numerical method		
Attention index	evaluation value (4)	baseline (μ variant in robustness)
Sample-average approximation	draws S; seed	$S = 20{,}000$; 20260703

Notes: These are simulation parameters for numerical illustrations. Values marked "anchored" are pinned to the observable moments of Table 4; values marked "swept" or "normalization" are free and varied in robustness, and the three with the most leverage (χ_M, C_V, $\kappa_{T,0}$) are pinned by minimum distance in Section 6.3. The theoretical exposition sometimes normalizes $\alpha = 1$ for convenience; every figure and table in Section 6 uses the simulation value $\alpha = 300$. Doubling S leaves entry counts unchanged and moves reported statistics by less than 10^{-3}.

Table 8: Numerical settings

Item	**Value**
Common-random-number seed	20260703
Sample-average draws S (tech signal vectors)	20,000
Attention solver	SLSQP, analytic gradients
Solver tolerance (`ftol`)	10^{-12}
KKT residual of included options (measured)	$1 \pm O(10^{-3})$
Legibility grid m_g	$0 : 0.25 : 4$
Meme effort grid h_M	$0 : 0.25 : 4$
Entry protocol	exit-only / entry-only continuation branches
Effect of doubling S	entry counts unchanged; statistics move $< 10^{-3}$

Notes: These settings apply to all numerical exercises unless otherwise stated. The pipeline is deterministic given the seed; meme-class aggregates are integrated exactly in the deep-menu limit (500 projects per class), so the only sampling object is the joint distribution of the 50 tech signals.

Table 9: Selected equilibrium outcomes in the extension: high-trust vs. low-trust branches

	High-trust / tech-active	**Low-trust / meme-dominant**
Active tech projects N_T^{eq}	10	0
Meme seeding effort h_M	0.00	2.25
Aggregate tech attention $\sum_i q_i$	1.000	0.000
Consideration set (tech / meme projects)	10 / 0	0 / 500
Attention benchmark $\bar{Z}$	1.696	1.698
Average investor payoff (decision)	0.500	0.516
Average investor payoff (fundamentals)	0.500	-0.450
Marginal tech profit (active) / counterfactual entry profit	$+7.32$	-0.400
Per-meme expected profit	0.000	$+0.009$

Notes: Locally stable branches of the *extension* with ecosystem externality ($\eta = 10$) at $\kappa_{T,0} = 81.7$, a value inside the bistable band of Figure 2. Figure 1 is instead generated in the baseline environment with $\eta = 0$.
The consideration set is the support of the optimal attention rule; in the low-trust branch it consists of the top virality class of memes (500 projects), and the investor's decision payoff barely moves (it even rises slightly, from 0.500 to 0.516) while the fundamentals payoff (salience stripped) collapses from $+0.50$ to -0.45: the decision-versus-experienced wedge of Section 7.
"Counterfactual entry profit" in the low-trust branch is the best achievable profit of the most legible tech group entering against the equilibrium menu; it equals $-K_T$ because entry attracts zero attention at any feasible legibility (exclusion margin).

Table 10: Equilibrium issuer strategies and consideration-set composition: regime comparison

	High-trust / tech-active	**Low-trust / meme-dominant**
Active tech groups (of 10)	2	0
Tech legibility m_g (active groups)	4.00 (all)	—
Meme seeding effort h_M	0.00	2.25
Consideration set: tech projects	10	0
Consideration set: meme projects	0	500 (top virality class)
Audit rate, considered tech	1.000	—
Audit rate, considered memes	—	0.000 (hype-shielded)

Notes: Locally stable branches of the extension at $\kappa_{T,0} = 81.7$, $\eta = 10$. The consideration set is the support of the optimal attention rule (corners of the concave program). The audit pattern is the model's verification logic in miniature: at this complexity level even full legibility leaves tech posteriors inside the audit window, so the active groups are audited with probability one and verification is part of how they stay evaluable, while hyped memes have padded downside ($L_i < 0$ for the considered class) and are funded on faith, never audited.

Table 11: Equilibrium issuer strategies and consideration-set composition: policy counterfactuals

	Baseline	**Lower κ_T / Higher β_T**	**Lower C_V**
Tech groups paying $m_g = 4$	0	5	2
Meme seeding effort h_M	2.25	0.00	0.00
Consideration set: tech	0	25	10
Consideration set: memes	500	0	0
Audit rate, considered projects	0.000	0.000	1.000

Notes: These entries are taken directly from the fixed-entry equilibria used to produce Table 6 (all 50 tech projects present). The purpose is diagnostic transparency rather than adding a new headline claim. Under both interventions the meme sector's salience contest shuts down endogenously: with legible or audit-supported tech in the consideration set, no level of seeding effort earns the memes nonnegative expected profit. The two interventions restore consideration through different technologies, legibility ($m_g = 4$, five groups) versus exercised audit options (two groups, audited with probability one).

Table 12: Figure 3 diagnostics

Diagnostic	**Value**
α^{ver} while tech populates the consideration set	0.0052
α^{ver} after the meme takeover ($c_{Mh2} \leq 0.005$)	0.0000
Tech entry at the takeover (N_T^{eq}: $c_{Mh2} = 0.01 \to 0.005$)	$20 \to 0$
Meme seeding effort at the takeover	$h_M = 3.25$
Retail audit activity along the whole path	≈ 0
N_M path (robustness): $N_M = 2{,}000$ / $5{,}000$ / $20{,}000$	$N_T^{\text{eq}} = 0$ / 20 / 20

Notes: The N_M path is *anti-monotone* in this equilibrium, for the rent-dissipation reason given in the main text: fewer memes mean richer per-meme rents, harder individual hype ($h_M = 3.75$ at $N_M = 2{,}000$), and a collapsed tech sector, while a larger meme count spreads the same attention pie more thinly and disciplines hype. The empirically relevant meme-wave experiment is therefore the salience-cost (platform) shock, not a pure count expansion; this also implies that raw issuance counts are a poor proxy for the salience pressure the model cares about.

Table 13: Robustness Summary: Persistence of Main Theoretical Forces

Specification	**Complexity trap persists?**	**Hysteresis persists?**
Unweighted mutual information cost (standard RI)	Yes	Yes
General convex divergence cost $\mathcal{K}_\psi$	Yes	Yes
Correlated states via common shock z	Yes	Yes (Stronger)
Strategic obfuscation ($\omega > 0$ for bad issuers)	Yes (Stronger)	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 (Stronger)	Yes (Stronger)

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

Table 14: Continuous meme-volume exposure and event-time audit adoption

Specification	Effect (HC2 SE)	t	HC2 p_1/p_2	Wild p_1/p_2	Path RI p_1/p_2	Protocol RI p_1/p_2	Protocols
Panel A: Broad complex-protocol sample, 2023–25							
Pre-listing EWMA, Solana-priority assignment	−9.16 (5.80)	−1.58	0.057/0.114	0.075/0.151	0.388/0.493	0.137/0.149	1,700
Pre-listing EWMA, deployment-split assignment	−4.38 (5.96)	−0.74	0.231/0.462	0.148/0.296	0.520/0.689	–	1,700
Next-six-month look-ahead	−0.49 (3.95)	−0.12	0.451/0.901	0.381/0.761	0.367/0.859	0.193/0.919	1,700
Panel B: Intricate protocols, 2023–24, Solana and five EVM controls							
Prior-month dose	−16.82 (9.90)	−1.70	0.045/0.089	0.078/0.156	0.250/0.500	0.009/0.013	172
Next-six-month look-ahead	−25.28 (9.93)	−2.54	0.005/0.011	0.078/0.156	0.401/0.601	0.003/0.019	172

Notes: The outcome and first-archive protocol labels follow Table 1. Effects and parenthesized HC2 standard errors are percentage points of audit adoption for a ten-percentage-point increase in chain meme-volume share. Panel A uses all twelve chains, chain, listing-half-year, and category fixed effects. The baseline dose averages the 24 pre-listing months with weights $0.5^{(\ell-1)/12}$ at lag ℓ. The deployment-split row uses a protocol's primary chain when recorded and otherwise weights all deployment chains equally; it contains 2,663 weighted protocol-chain rows. Panel B restricts the sample to the derivatives, options, bridge, synthetics, basis-trading, leveraged-farming, and restaking families, ends with the 2024 cohorts, uses Solana plus Ethereum, Arbitrum, Optimism, Polygon, and Fantom, and includes chain, listing-quarter, and category fixed effects. Only Solana has positive measured dose in Panel B. The next-six-month measure uses the listing month through month +5 and is a look-ahead diagnostic. p_1 is directional for $H_1 : \beta < 0$; p_2 is two-sided. Wild values enumerate all Rademacher sign assignments by chain. Path RI values use 100,000 permutations of the observed exposure paths across chains. Protocol RI values use 100,000 residual permutations within listing-time and category strata; repeated protocol-chain rows make that statistic undefined for the deployment-split row. Across the nine pre-listing dose definitions in each fixed family, Holm-adjusted directional HC1 p is 0.506 for the Panel A baseline and 0.340 for the Panel B prior-month row.

Table 15: Treatment-definition and label-timing sensitivity of audit adoption

Classification source	DiD	HC2 SE	Directional p	Two-sided p	N
Panel A: Solana-exposed protocols					
Current cross section	−23.9 pp	(9.6)	0.011	0.016	1,211
Outcome-date archive	−21.5 pp	(10.6)	0.028	0.045	1,204
First available archive (main)	−21.0 pp	(10.7)	0.034	0.052	1,146
Panel B: Solana-native restriction					
Current cross section	−16.9 pp	(11.2)	0.083	0.134	1,167
Outcome-date archive	−18.3 pp	(11.3)	0.065	0.108	1,175
First available archive	−19.8 pp	(11.3)	0.049	0.081	1,130

Notes: Each row uses the event-time audit outcome and the 2023 versus 2024–25 comparison of Table 1. Panel A treats every protocol with a Solana deployment as exposed. Panel B additionally requires that Ethereum not appear among deployment chains. "Outcome-date archive" reads deployment labels from the same designated snapshot as audit status; "first available archive" reads them from the first full archived protocols capture containing the protocol. HC2 standard errors, in percentage points, come from the saturated protocol-level linear probability model. Directional and two-sided p-values use 100,000 protocol-label permutations within cohort half-years and are unadjusted across rows. No row supplies chain-clustered inference.

Table 16: Audit adoption across the shock by pre-launch complexity class

Complexity class	Solana pre	Solana post	EVM pre	EVM post	DiD
Panel A: oracle dependence, promoted E1 exposure ($n = 907$)					
Oracle-dependent	78.6% (14)	60.0% (10)	75.4% (126)	68.9% (61)	−12.0 pp
Oracle-free	40.0% (15)	33.9% (56)	46.3% (352)	51.6% (273)	−11.4 pp
Triple difference					−0.6 pp
HC2 standard error					(26.0) pp
Directional / two-sided RI p					0.498/0.981
Panel B: structurally intricate categories, promoted E1 exposure ($n = 1{,}146$)					
Intricate categories	85.7% (7)	36.0% (25)	60.6% (109)	51.7% (143)	−40.9 pp
Other complex	50.0% (22)	32.6% (95)	52.0% (369)	48.9% (376)	−14.3 pp
Triple difference					−26.6 pp
HC2 standard error					(22.3) pp
t statistic					−1.20
Directional / two-sided RI p					0.136/0.289
Panel C: direct 2024 window, Ethereum control (exploratory)					
Intricate-category DiD					−61.1 pp
Other-complex DiD					−7.0 pp
Triple difference					−54.1 pp
HC2 standard error					(26.0) pp
t statistic					−2.08
Directional / two-sided RI p					0.036/0.065
Directional Holm p, two proxies					0.072

Notes: Panels A and B use the event-time audit outcome, first-available archive labels, the Solana-exposed definition, the 2023 versus 2024H1–2025H2 window, and Ethereum/Arbitrum/Optimism controls of Table 1. Cell N is in parentheses. Oracle dependence is a nonempty oracle field; intricate categories are derivatives, options, bridges, synthetics, basis trading, leveraged farming, and restaking. The oracle field is observable only in captures through 2025-03, with truncation similar across chains, and its empty/nonempty class changes for 1.2 percent of repeated observations. Panel C changes only the post window and control set; it was selected from the 180-cell complexity log and is not a familywise result. Randomization inference permutes protocol labels within cohort-half-year × complexity-class blocks; the displayed high-draw values use 100,000 draws. Under the Solana-native restriction, the oracle and intricate triple differences are −15.4 and −17.6 points, with two-sided RI $p = 0.556$ and 0.513. The fork proxy is unidentified because no Solana-native protocol is coded as a fork.