The Provenance Turn:

Signal Decay and the Verification Architecture of Markets Under Generative AI

Shivam Srivastava*[1]

[1]*Carnegie Mellon University, Pittsburgh, Pennsylvania, USA*

June 2, 2026

Abstract

Modern markets coordinate through costly signals. From Akerlof (1970) and Spence (1973) onward, information economics has explained how parties who cannot directly observe one another's quality, type, or intent nonetheless transact through signals whose informativeness derives precisely from their cost: signals that low types would not rationally send. Generative artificial intelligence is hollowing out this foundation, driving the marginal cost of signal production toward zero across credentials, brand expressions, application essays, product reviews, research summaries, financial disclosures, and identity performances alike, and dissolving the cost asymmetries on which separating equilibria depend.

I theorize this phenomenon as *signal decay* and argue that markets are responding through a deeper reorganization I call the *provenance turn*: a shift from coordination by signal cost to coordination by verifiable origin. Drawing on information economics, marketing, finance, information systems, the sociology of markets, and an emergent computer-science literature on attestation, I develop provenance as a market-level construct, introduce the core constructs of *verification capital* and *provenance elasticity*, propose a four-type typology of verification architectures (cryptographic, witnessed, institutional, reputational), and organize markets along a Saturation $\times$ Tractability framework that locates four distinct trajectories under signal decay. From this I derive a set of falsifiable propositions spanning provenance premia, channel migration, regulatory reconstitution, and marketing–finance integration. The framework reframes generative AI's market impact as a *coordination* problem rather than a productivity problem, and shifts attention from what firms produce to how markets establish trust in what is produced. I close with an empirical agenda specifying field experiments, difference-in-differences and event-study designs, conjoint and audit studies, and digital-trace analyses for testing the framework.

Keywords: generative AI, signaling theory, information economics, provenance, verification capital, trust infrastructure, attestation, content authenticity, credence goods, market sociology, AI governance, market microstructure, platform markets.

*Corresponding author: shivamsr@andrew.cmu.edu

Notation and Terminology

Formal symbols.

Symbol	Meaning
S	A signal type (e.g., an application essay, a product review, an analyst note).
t	Time; the index along which production technology diffuses.
$q_S(t)$	Posterior probability a receiver assigns to the sender being a high type, conditional on observing signal S at time t.
q_0	Unconditional prior probability of the high type, absent any signal.
$\delta_S(t)$	*Signal decay*: the rate at which $q_S(t)$ converges to the prior q_0 as low-cost production technology diffuses, $\delta_S(t) = -\frac{d}{dt}[q_S(t) - q_0]$.
η_P	*Provenance elasticity*: percentage change in willingness-to-pay (or to engage) per one-percent change in provenance-attestation credibility, holding signal content fixed.
ρ_T	*Trust infrastructure density*: per-transaction availability of verification mechanisms in a market.
c	Marginal cost of producing a signal; under generative AI, $c \to 0$ across sender types.

Structural taxonomy.

Term	Meaning
Verification architecture	The institutional infrastructure through which provenance claims are constructed, transmitted, and challenged. Four ideal types: *cryptographic*, *witnessed*, *institutional*, and *reputational* (Section 6).
Saturation Severity	Horizontal axis of the framework: the degree to which a signal type has been rendered uninformative by cheap generative production across the sender population (Section 8).
Verification Tractability	Vertical axis of the framework: the feasibility of establishing provenance for the relevant artifact given current technology, infrastructure, and cost (Section 8).

Named higher-order constructs are defined in Section 4 and listed in Appendix A.

1. Introduction: The Puzzle of Effective but Uninformative Markets

Consider a series of seemingly unrelated developments. Elite universities, whose admissions decisions were already among the most reputation-rich in the world, are reintroducing in-person interviews and proctored writing assessments at scale. Auction houses with centuries-old authentication regimes are accelerating investment in physical archives and chain-of-custody documentation, even as their digital records expand. Dating platforms are deploying video verification, live-photo capture, and short-form recorded interactions in place of curated written profiles. Major scientific journals have launched AI-detection initiatives even as their editorial boards concede that such tools are unreliable at the level of individual manuscripts (Sadasivan et al., 2023; Weber-Wulff et al., 2023). Hiring platforms are quietly de-weighting written application materials and re-emphasizing referrals, live skills assessments, and structured interviews. Premium media brands are stamping articles with cryptographic provenance certificates under the Coalition for Content Provenance and Authenticity (C2PA) standard (Rosenthol, 2022). Equity-research desks at major investment banks are layering human attestation onto AI-generated analyst notes (Lopez-Lira & Tang, 2023). News organizations, foundation-model labs, and content platforms have signed onto cross-industry provenance consortia at a pace unmatched by any prior trust-infrastructure cycle (Longpre et al., 2024).

These responses span industries that share little at the level of product, technology, or competitive structure. They share, however, a single underlying phenomenon: *the signals on which each market historically coordinated are losing their informational value*. The application essay, the curated profile, the writing sample, the product review, the LinkedIn summary, the analyst note, the polished image, the literature review, the legal brief, the medical note, the abstract of a working paper, the influencer endorsement, the corporate press release—each is a classical signal whose credibility flowed from its cost of production. Each is now produced, in increasing fraction, by generative AI at near-zero marginal cost (Brynjolfsson et al., 2025; Eloundou et al., 2023; Korinek & Stiglitz, 2021). The cost asymmetries that signaling theory has relied upon since Spence (1973) are collapsing across vast swathes of the economy, simultaneously and across otherwise uncorrelated industries.

Yet markets continue to function. They are reorganizing—not around signal quality, which is increasingly indistinguishable across sender types, but around something different. I theorize this reorganization as the *provenance turn*: a deep shift in the architecture of market coordination, from costly signals (which proved type by being expensive to produce) to verifiable origin (which proves type by being expensive to falsify at the source). Where signaling asked the receiver to infer type from the cost of what was sent, provenance asks the receiver to infer type from the verifiable history of how what was sent came to exist. The two logics differ in subtle but consequential ways: signaling places the informational burden on the sender's production decision; provenance places it on the verification infrastructure surrounding the sender's act of production.

1.1 Why this is a coordination problem, not a productivity problem

The dominant framing of generative AI in business and economics has emphasized productivity: tasks accelerated, jobs reshaped, output expanded (Brynjolfsson et al., 2025; Eloundou et al., 2023; Noy & Zhang, 2023). That framing is necessary but insufficient. A purely productivity-centered view treats AI as a faster substitute for labor inside otherwise stable market institutions. My claim is sharper: *generative AI is restructuring the institutions of market coordination themselves*. The signaling apparatus that organizes hiring, admissions, credit, brand competition, reputation, peer review, journalism, and capital allocation is degrading in a way no single firm initiated, no single firm can reverse, and no productivity gain compensates for. The locus of theoretical action shifts from the firm's production frontier to the market's verification frontier.

1.2 What is genuinely new

A natural objection—developed at length in Section 15—is that signal-collapse episodes are not new. Cheap printing, photography, and the internet each degraded an established signaling regime and were eventually absorbed by new verification institutions. If the provenance turn is just the latest instance of a recurring pattern, what warrants separate theoretical treatment?

The answer rests on three structural features that distinguish the present episode, and that I argue are jointly necessary to motivate the framework. *First, cross-domain simultaneity*: prior episodes degraded one signaling channel at a time (printing attacked the authority of the written word; photography attacked pictorial evidence), whereas generative AI attacks heterogeneous signal classes—text, image, audio, code, credential—concurrently, so that receivers cannot reallocate trust to an unaffected channel. *Second, symmetric diffusion*: the production technology is available to high and low types on identical terms at near-zero marginal cost, which removes the residual cost gap that earlier technologies preserved (a forged painting still required skill; a fabricated document still required effort). *Third, recursive exposure*: the reputation systems that ordinarily backstop a degraded signal—reviews, endorsements, citation counts—depend on inputs that are themselves now cheaply generable, so the usual fallback is compromised by the same shock (Mechanism 5). These three features are empirically checkable rather than rhetorical: cross-domain adoption series, type-conditioned adoption gaps, and the synthetic-content share of reputational inputs are all measurable, and Section 14 specifies how. The claim of novelty therefore stands or falls on evidence, not assertion.

1.3 Five claims

I make five claims.

1. *Signal decay* is a first-order market-level phenomenon: a market-coordination failure that propagates across heterogeneous signal types simultaneously and that classical signaling theory by itself cannot resolve.
2. The response is not better detection of AI-generated content—detection of generative output

is provably bounded (Chakraborty et al., 2023; Sadasivan et al., 2023)—but the construction of *verification architectures* that establish provenance independent of signal content.

3. Verification architectures are heterogeneous; they vary in cost structure, scalability, failure mode, and institutional location, and these differences determine which markets reconstitute around provenance and which migrate or collapse.
4. Competitive outcomes are increasingly determined not by what firms produce but by their access to and control over verification infrastructure—an inversion that strategic theory has not yet absorbed and that I formalize as the firm's *verification capital*.
5. The public-policy frontier shifts from anti-deception regulation toward *attestation governance*, a domain for which existing legal categories are poorly adapted and that will absorb significant institutional design effort over the coming decade.

1.4 Theoretical contribution

The paper sits in dialogue with six traditions: information economics on signaling and credence goods (Akerlof, 1970; Darby & Karni, 1973; Dulleck & Kerschbamer, 2006; Riley, 2001; Spence, 1973; Stiglitz, 1975); marketing scholarship on trust, brand, customer journey, and AI-mediated decisions (Davenport et al., 2020; Huang & Rust, 2021; Kannan & Li, 2017; Lamberton & Stephen, 2016; Lemon & Verhoef, 2016; Longoni & Cian, 2022; Puntoni et al., 2021; Rust, 2020); finance research on market microstructure, signaling in IPOs and disclosure, and AI risk in financial markets (Allen & Faulhaber, 1989; Cong et al., 2021; Goldstein et al., 2021; Leland & Pyle, 1977; Myers & Majluf, 1984; O'Hara, 2015); sociology of markets on legitimacy, categorization, and audience-based valuation (Carroll, 2015; Hannan et al., 2007; Hsu, 2006; Suchman, 1995; Zuckerman, 1999); information-systems research on platforms and AI governance (Benbya et al., 2021; Constantinides et al., 2018; Gawer, 2022; Parker et al., 2016; Schneider et al., 2023; Van Alstyne et al., 2016); and a nascent computer-science and policy literature on attestation, content authenticity, and data provenance (Aaronson, 2023; Christ et al., 2024; Epstein et al., 2023; Kirchenbauer et al., 2023; Longpre et al., 2024).

The contribution is conceptual. I do not test the propositions empirically. I introduce a vocabulary, a typology, a multi-level framework, and a set of propositions whose joint purpose is to render the provenance turn a tractable object of social-scientific inquiry. The level of analysis is the market, not the firm; the central *explanandum* is coordination, not advantage.

1.5 Plan of the paper

Section 2 synthesizes the signaling foundation of modern markets and identifies the assumption that generative AI violates. Section 3 develops signal decay as a market-level construct. Section 4 introduces the constructs of the verification era. Section 5 defines provenance and its boundaries. Section 6 develops the verification-architecture typology. Section 7 presents a multi-level model. Section 8 develops the Saturation $\times$ Tractability framework. Section 9 specifies six mechanisms.

Section 10 develops marketing implications; Section 11, finance implications; Section 12, governance and policy. Section 13 advances seventeen propositions. Section 14 sketches empirical research designs. Section 15 addresses boundary conditions and competing explanations. Section 16 discusses theoretical implications. Section 17 concludes.

2. Theoretical Foundations: From Signaling to Verification

This section synthesizes six traditions whose convergence—and joint insufficiency—motivates the framework. Table 1 summarizes the synthesis.

2.1 Information economics: the Spence–Akerlof tradition

Information economics in the lineage of Akerlof (1970), Spence (1973), and Stiglitz (1975) rests on a foundational insight: when receivers cannot directly observe sender types, markets can nevertheless separate types if senders bear costs that differ systematically across types. The classical separating equilibrium requires that the cost of producing a given signal be sufficiently higher for low types than for high types that low types rationally decline to send it. The signal is then informative because it is expensive. The educational credential, in Spence's formulation, separates productive from less productive workers not because schooling necessarily produces productivity, but because the cost of completing it is differentially higher for low-productivity types.

Generalizations of the argument span advertising as signal of product quality (Milgrom & Roberts, 1986; Nelson, 1974); warranties as bonded quality claims (Grossman, 1981); brand investment as posted bond (Wernerfelt, 1988); underpricing in initial public offerings as signal of issuer quality (Allen & Faulhaber, 1989; Welch, 1989); and the design of online reputation systems as a digital-era substitute for repeated bilateral play (Dellarocas, 2003; Resnick & Zeckhauser, 2002; Tadelis, 2016). A complementary line addresses *credence goods*, whose quality consumers cannot determine even after consumption (Darby & Karni, 1973; Dulleck & Kerschbamer, 2006). For credence goods, signaling alone is insufficient; markets coordinate through a layered apparatus of credentials, third-party certification, professional licensing, regulatory disclosure, and bonded reputation. In each case, the apparatus depends on cost asymmetries: producing a fraudulent certification, maintaining a deceptive reputation, or supplying credible disclosure at false quality levels must be sufficiently expensive that low types decline.

2.2 Marketing: trust, customer journey, AI-mediated decision-making

Marketing has long treated brands, advertising, and reviews as costly signals that anchor consumer choice (Erdem & Swait, 2004; Keller, 2003; Kirmani & Rao, 2000). The contemporary marketing literature has extended this view to AI-mediated environments. Huang & Rust (2021) and Rust (2020) chart how AI restructures the customer journey and the feeling economy. Longoni et al. (2019); Longoni & Cian (2022) and Castelo et al. (2019) document algorithm aversion—and its asymmetric reversal—across decision contexts. Puntoni et al. (2021) theorize consumer experiences of AI as

a multi-faceted relational phenomenon involving objectification, replacement, and surveillance. Davenport et al. (2020) sketch AI's strategic implications for marketing practice. Berger et al. (2020) and Netzer et al. (2019) treat natural language as a data infrastructure for marketing inference. A parallel strand considers the platform-pricing and demand-side implications of AI agents acting on behalf of consumers, though the empirical literature here remains nascent.

A second strand explores trust in algorithmic systems (Dietvorst et al., 2015; Glikson & Woolley, 2020; Lee & Rich, 2021; Logg et al., 2019) and the creator economy as a platform-mediated signaling market (Hofstetter & Gollnhofer, 2024; Leung et al., 2022). Across this work two themes recur: (i) the diffusion of generative tools across *both* sender and receiver sides of markets, and (ii) the re-routing of trust through unfamiliar intermediaries (recommender systems, agentic assistants, content-authentication services). Yet the theoretical core remains signal-based: brands, reviews, and content are still treated as informationally meaningful artifacts even as the marginal cost of producing each falls toward zero. The shared, unspoken assumption is the same one that organizes the rest of information economics: signal cost varies meaningfully across sender types. That assumption is the one under attack.

2.3 Finance: market microstructure, disclosure, AI risk

Finance has rich theory on signaling under information asymmetry (Leland & Pyle, 1977; Myers & Majluf, 1984), on disclosure as a price-relevant signal (Healy & Palepu, 2001; Verrecchia, 2001), and on the informational role of intermediaries (Diamond, 1984; Gurun & Butler, 2012). The market-microstructure literature (Kyle, 1985; O'Hara, 2015) models how information quality determines liquidity, depth, and price discovery. Recent work formalizes signaling and certification by underwriters, analysts, and auditors (Dechow et al., 2010; Ljungqvist, 2007).

Generative AI introduces a new class of risk that I term *synthetic information risk* (Section 4): the price of a security becomes exposed to non-trivial probabilities of fabricated press releases, deepfaked CEO statements, synthetic analyst notes, AI-generated regulatory filings, prompt-injected agentic trading errors, and coordinated narrative attacks generated and amplified at near-zero cost (Cong et al., 2021; Goldstein et al., 2021; Lopez-Lira & Tang, 2023). The May 2023 Pentagon-explosion deepfake episode, which briefly moved U.S. equity markets before being debunked, and industry analyses of deepfake threat scenarios in the financial sector (Bateman, 2020) illustrate the texture of the problem in different layers of the information stack. A small but growing body of work on tokenization and on-chain provenance (Catalini & Gans, 2020; Halaburda et al., 2022; Harvey et al., 2021; Yermack, 2017) foregrounds verifiable origin as a mechanism for restoring transactional trust in environments where conventional certifiers are absent.

2.4 Sociology of markets: legitimacy, audiences, categories

Parallel sociological work has examined how markets construct the categories through which valuation occurs. Zuckerman (1999) showed that securities markets penalize firms that cannot be

coherently categorized by analysts. Hsu (2006) extended the argument to multi-category producers. Hannan et al. (2007) developed a general theory of audiences as the constitutive evaluators of market categories. Carroll (2015) examined authenticity as a quasi-signal whose value derives from the costliness of credible biographical fit. Suchman (1995) provided the foundational treatment of legitimacy as a market resource. Rao (1994, 2009) foregrounded the role of certification contests and category-creating institutions. Across these accounts the audience is not a passive Bayesian updater but an active classifier whose cognitive and institutional infrastructure determines what counts as a credible signal.

2.5 Information systems and AI governance

The information-systems literature on platform governance (Constantinides et al., 2018; Gawer, 2022; Parker et al., 2016; Tiwana et al., 2010; Van Alstyne et al., 2016) treats platforms as institutional infrastructure designing the rules of multi-sided markets. AI governance research (Benbya et al., 2021; Berente et al., 2021; Rai et al., 2020; Schneider et al., 2023) treats attestation, audit, and provenance as emerging classes of governance instrument. A technical literature on watermarking and AI content authentication (Aaronson, 2023; Christ et al., 2024; Kirchenbauer et al., 2023) develops cryptographic and statistical methods for binding generated content to its model and producer. The Coalition for Content Provenance and Authenticity (C2PA) and the Content Authenticity Initiative have proposed cross-industry standards (Rosenthol, 2022); the IEEE, NIST, and the EU AI Act increasingly recognize provenance as a regulatory category. Longpre et al. (2024) audit dataset provenance at scale; Epstein et al. (2023) survey the art-science of generative media.

2.6 The common assumption under strain

Across information-economic, marketing, finance, sociological, and IS treatments, *one assumption is shared*: signals cost something to produce, and that cost varies meaningfully across sender types. Generative AI is the first general-purpose technology that systematically attacks this assumption across heterogeneous signal types *simultaneously*. A polished cover letter, a fluent research summary, a coherent product description, a confident analyst note, a passable code sample, a stylized image, a personalized recommendation, a moderately complex legal argument—each can now be produced at marginal cost approaching zero by senders of any type. The separating equilibrium degenerates toward a pooling equilibrium by default; receivers cannot infer type from signal; the market for the underlying good or relationship moves toward the adverse-selection outcome Akerlof (1970) described for used cars.

What is novel is not that signals can be faked—that has always been possible—but that the cost of faking has collapsed near-uniformly across signal types and across sender populations, simultaneously, in a manner that no individual market participant initiated and no individual market participant can reverse (Section 1.2). This is a coordination problem of unusual structure, and it demands theoretical apparatus that information economics has not previously been required to develop.

2.7 Why each tradition is individually insufficient

Each existing tradition supplies a piece of the puzzle but *individually* cannot account for the phenomenon. Signaling theory predicts the collapse of separating equilibria but offers no architecture for what replaces them. Credence-goods theory explains how third-party certifiers stabilize quality claims but assumes the certifier's own attestation cannot be cheaply forged—an assumption that generative AI itself attacks (synthetic certificates, AI-mediated audit). Platform theory explains the coordinating role of platforms but does not yet treat provenance as a first-class platform resource. Reputation-system theory explains how repeated play stabilizes trust but presumes the integrity of the rating inputs, themselves now generable at scale. Institutional theory explains legitimacy through category conformity but does not specify what happens when the artifacts conferring legitimacy are produced at zero marginal cost. AI-governance scholarship debates capability and safety but underspecifies the market-coordination consequences of unverifiable content.

The common omission is the same in different costume: *the simultaneous degradation of the cost foundation of multiple signaling regimes is a market-coordination shock, not a productivity shock*. It restructures the rules by which strangers transact, not the inputs they assemble. Table 1 situates the existing literatures relative to the framework.

3. Signal Decay as a Market-Level Construct

Definition 1 (Signal Decay)**.** Signal decay *is the rate at which a given signal type loses its informational value to receivers under the widespread availability of low-cost production technology accessible to all sender types. Formally, let $q_S(t)$ denote the posterior probability that the receiver assigns to the sender being a high type conditional on observing signal S at time t. Signal decay $\delta_S(t)$ is the rate at which $q_S(t)$ converges to the unconditional prior, q_0, as the production technology diffuses:*

$$\delta_S(t) \;=\; -\frac{d}{dt}\left[\, q_S(t) - q_0 \,\right].$$

Decay is a property of a signal–market pair, not of an individual signal or an individual market. The same underlying signal type—the personal essay, say—may decay rapidly in one market (college admissions, where senders are sophisticated AI users) and slowly in another (a craft fellowship application, where senders less commonly use such tools or where the evaluation process penalizes detected AI use).

3.1 A comparative-static reading

The formalism is not merely descriptive notation; it yields a sign prediction that organizes the propositions in Section 13. Let receivers update by Bayes' rule on a signal whose likelihood ratio between high and low types is $\ell(t)$, so that the informativeness of S is increasing in ℓ. Suppose adoption of the cheap production technology by low types proceeds at rate $a_L(t)$ and by high types at rate $a_H(t)$. As the cost gap closes, $\ell(t)$ falls toward unity and $q_S(t) \to q_0$. Two comparative

Table 1. Synthesis of literatures bearing on the provenance turn.

Tradition	Core insight	Representative works (2023–2026 emphasis)	What it cannot explain
Signaling & information economics	Costly signals separate types when low types face higher costs.	Akerlof (1970); Allen & Faulhaber (1989); Milgrom & Roberts (1986); Riley (2001); Spence (1973); Stiglitz (1975); Tadelis (2016)	What replaces separating equilibria when the cost foundation collapses.
Credence goods	Quality unknowable post-consumption; markets coordinate via certifiers and licensure.	Darby & Karni (1973); Dulleck & Kerschbamer (2006)	What happens when the certifier's attestation can itself be cheaply forged.
Marketing: trust, journey, AI	Brands and content as costly signals; AI restructures journey, feeling economy, identity.	Berger et al. (2020); Davenport et al. (2020); Huang & Rust (2021); Lemon & Verhoef (2016); Longoni & Cian (2022); Puntoni et al. (2021); Rust (2020)	Market-level coordination when the content layer becomes informationally indistinguishable.
Finance: microstructure, disclosure, AI	Liquidity and price discovery depend on information quality and certifying intermediaries.	Bateman (2020); Cong et al. (2021); Goldstein et al. (2021); Healy & Palepu (2001); Leland & Pyle (1977); Lopez-Lira & Tang (2023); O'Hara (2015)	Synthetic information risk and the strategic value of verifiable origin as a financial primitive.
Sociology of markets	Audiences classify; legitimacy is a constructed market resource.	Carroll (2015); Hannan et al. (2007); Hsu (2006); Rao (2009); Suchman (1995); Zuckerman (1999)	How audience evaluation reconstitutes when surface artifacts no longer reveal type.
IS, platforms, AI governance	Platforms govern multisided markets; AI governance includes audit, attestation, watermarks.	Aaronson (2023); Benbya et al. (2021); Berente et al. (2021); Christ et al. (2024); Constantinides et al. (2018); Gawer (2022); Kirchenbauer et al. (2023); Longpre et al. (2024); Parker et al. (2016); Schneider et al. (2023); Van Alstyne et al. (2016)	The strategic and institutional consequences of verification becoming a first-order coordinating layer.
Reputation systems	Repeated play stabilizes trust through accumulated ratings.	Dellarocas (2003); Resnick & Zeckhauser (2002); Tadelis (2016)	Reputational inputs that are themselves AI-generated at scale (recursive saturation).

statics follow directly. *First,* decay is increasing in the symmetry of adoption: $\partial \delta_S / \partial \left[a_L - a_H \right]$ is positive when low-type adoption catches up to high-type adoption, because the likelihood ratio collapses fastest when the two populations look alike. *Second,* decay is bounded below by zero and is identically zero whenever a residual cost asymmetry survives ($a_L < a_H$ at the relevant margin), which is exactly the *pre-decay signaling* region of the framework in Section 8. The first statement motivates the saturation axis and Proposition 1; the second delimits the boundary condition under which classical signaling persists. The model is deliberately reduced-form—a fuller verification game appears as a research agenda item in Section 16—but it is enough to show that the construct does predictive work rather than relabeling.

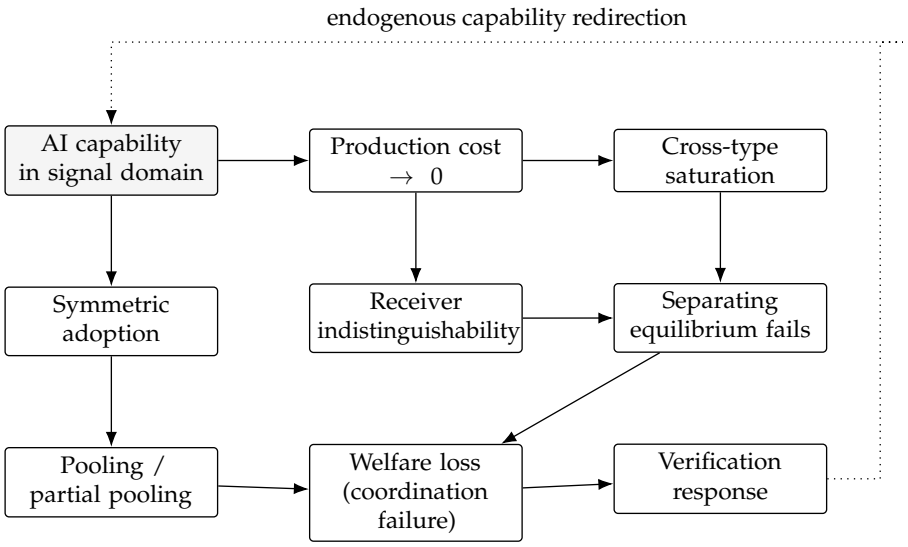


Figure 1. The signal-decay mechanism. AI capability in a signal domain drives the marginal cost of production toward zero, producing symmetric cross-type adoption and receiver indistinguishability. The separating equilibrium fails; markets shift toward pooling or partial pooling; welfare loss follows; verification responses emerge. The dotted feedback loop captures the slower process by which verification responses alter where AI capability is profitably deployed.

3.2 Three constituent components

Three components jointly constitute signal decay:

1. **Saturation:** the rate at which low and high types adopt the cheap production technology, narrowing the cost gap between them. Saturation is high when adoption is symmetric across sender types and low when adoption remains correlated with the underlying type.
2. **Indistinguishability:** the receiver's diminishing ability to separate types from signal content alone, even with sophisticated analysis. Indistinguishability rises with generative capability and with the heterogeneity of legitimate stylistic variation in the underlying signal class.
3. **Welfare loss:** the magnitude of the coordination failure that follows. Welfare loss equals the difference between the social surplus realized under the original separating equilibrium and the surplus realized under the pooling or partial-pooling equilibrium that replaces it. Welfare loss is the empirical quantity policymakers and platform designers should ultimately seek to estimate.

Figure 1 traces the causal sequence.

3.3 Distinguishing adjacent constructs

Signal decay must be distinguished from three adjacent concepts. It is not *signal fatigue* (Nelson, 1974), the boredom-like decline in receiver attention to repeated signals; decay can occur on a first encounter if receivers know the production technology is widely available. It is not *signal jamming* (Riley, 2001), the strategic act of an interested party introducing noise into a signal channel; decay arises from uncoordinated mass adoption, not strategic interference. It is not *adverse selection* in the

original Akerlof (1970) sense, which arises from quality heterogeneity invisible to buyers; decay arises from the failure of an otherwise functioning signaling apparatus when its cost foundation erodes. Adverse selection may follow signal decay, but the two are causally distinct: adverse selection is a market outcome, whereas signal decay is a precondition for that outcome.

3.4 Decay regimes

Markets exhibit qualitatively different decay regimes. *Rapid decay* markets (e.g., undergraduate writing assignments, mass-market job applications, generic dating profiles, low-stakes product reviews) lose informational content within months once AI tooling diffuses. *Slow decay* markets (e.g., high-end visual art, embodied skilled trades, specialized scientific instrumentation) retain informational content because current generative capability has not penetrated the relevant production technology. *Oscillating decay* markets (e.g., academic writing, financial analyst reports, journalism) alternate between detection improvements and generation improvements, producing punctuated rather than monotonic decay. *Recursive decay* markets are those whose reputational inputs are themselves vulnerable to saturation—e.g., the customer-review system on which seller reputation depends, or the news ecosystem on which platform-content moderation depends.

3.5 Empirical signatures

Signal decay leaves observable traces. I highlight five. *First,* measured rater agreement on signal-quality dimensions falls—human raters disagree more, and the gap between expert and lay rater shrinks. *Second,* the predictive validity of signal-based selection systems falls relative to held-out outcomes (e.g., the predictive validity of admissions essays for later academic performance declines). *Third,* channel composition of high-stakes transactions shifts toward verifiable channels (interviews replace essays, video replaces text). *Fourth,* adoption of provenance infrastructure (C2PA, signed credentials, attestation logs) accelerates within affected markets. *Fifth,* relational and reputation-mediated exchange grows relative to arms-length, signal-mediated exchange.

4. A Vocabulary for the Verification Era

The provenance turn is not fully legible through the inherited vocabulary of signaling and reputation. Four constructs do the central analytical work and are defined formally below; a further set of supporting terms is introduced in prose at the end of the section and collected in Appendix A.

Definition 2 (Verification Capital)**.** Verification capital *is the firm-level (or institution-level) stock of resources, relationships, and infrastructure that allow it to credibly bind its outputs to verifiable origin and to attest to the origin of others. It includes ownership stakes in attestation infrastructure, seats on standards consortia, accumulated audit relationships, cryptographic keys with credible custody chains, and the human capital required to administer attestation processes. Verification capital is partially substitutable across signal domains and partially specific. It is the strategic analog of brand equity in the signaling era, and it is empirically distinct: where brand equity is measured through awareness, association, and willingness-to-pay*

for the marque, verification capital is measured through consortium membership, attestation-infrastructure ownership, and audit relationships (Proposition 12).

Definition 3 (Provenance Elasticity)**.** Provenance elasticity η_P *is the percentage change in willingness-to-pay (or willingness-to-engage) of receivers in response to a one-percent change in the credibility of the provenance attestation accompanying an artifact, holding signal content constant. Provenance elasticity is the empirical signature of the provenance premium; it varies across markets, audiences, and decision stakes.*

Definition 4 (Synthetic Information Risk)**.** Synthetic information risk *is the expected loss that a market participant bears from the non-trivial probability that informationally consequential artifacts they encounter were fabricated by generative systems. It includes deepfaked statements by principals, synthetic earnings press releases, AI-generated regulatory filings, fabricated certifications, and coordinated narrative attacks. It is closest to operational risk among existing categories but has a coordination character—its realizations are correlated across participants exposed to the same fabricated artifact—that distinguishes it.*

Definition 5 (Signal Liquidity)**.** Signal liquidity *is the ease with which receivers can convert observed signals into reliable inferences about sender type at a given moment. It is the informational analogue of market liquidity: high when many signal types carry calibrated content, and falling catastrophically when many signal types lose content simultaneously. Signal-liquidity crashes mark the most severe phases of the provenance turn and are the system-level expression of correlated signal decay.*

These four constructs recur throughout as the working apparatus. A set of supporting terms names finer-grained phenomena that the framework organizes but that do not require independent formalization. *Trust infrastructure density* (ρ_T) is the per-transaction availability of verification mechanisms in a market, a market-level state variable that rises with infrastructure investment and degrades with attestation failures. *Verification arbitrage* is the profit opportunity captured by an intermediary that acquires verification capital cheaply in low-saturation markets and deploys it at a premium in high-saturation ones. *Provenance network effects* are the rising value of a provenance standard as more counterparties recognize it—a chain-of-custody variant of two-sided network effects (Parker et al., 2016; Rochet & Tirole, 2003) that favors concentration in standards-setting. *Verification frictions* are the transaction costs that attestation imposes (key management, latency, certification fees); a market with high welfare loss from decay tolerates higher frictions before exit. *Signal inflation* is the receiver-side response of demanding more or more elaborate signals as indistinguishability rises—welfare-destroying because it raises participation cost without restoring content. *Verification externalities* are spillovers of infrastructure investment across markets, positive (shared standards) or negative (high-trust norms make low-trust attestation visible by contrast). The *information authenticity premium* is the price differential between artifacts of equivalent surface content that differ only in attestation credibility—the observable counterpart of provenance elasticity. *Verification congestion* occurs when demand for attestation exceeds an architecture's capacity, raising latency and fees and creating rents for incumbent verifiers. Finally, an *AI-induced coordination shock* is a discrete episode in which a generative-capability advance suddenly degrades a previously functional signal class (ChatGPT's late-2022 release for written application materials; subsequent video generation for

Table 2. Comparative architecture: signaling markets versus provenance markets.

Dimension	Signaling regime	Provenance regime
Coordinating logic	Signal cost asymmetric across types	Verifiable origin asymmetric across types
Locus of informational burden	Sender's production decision	Verification infrastructure surrounding production
Welfare-relevant variable	Quality of separating equilibrium	Tractability of verification
Capital that confers advantage	Brand equity, reputation, certification	Verification capital, infrastructure stakes
Failure mode	Pooling on signal	Attestation failure or congestion
Key intermediaries	Brand owners, certifiers, ratings firms	Attestors, standards consortia, registries
Regulatory frame	Anti-deception, disclosure	Attestation governance, liability for false attestation
Network economics	Brand-recognition effects	Provenance network effects (standards adoption)
Premium structure	Brand premium, quality premium	Information authenticity premium
Strategic frontier	Differentiation	Verification differentiation

video-mediated identity); coordination shocks are the natural unit of analysis for the temporal dynamics of the turn.

4.1 From signal to provenance: a comparative reframing

Table 2 synthesizes the contrast between the signaling regime and the emerging verification regime along ten dimensions. The table is a heuristic, not a typology: most empirical markets carry features of both regimes, but the comparison clarifies what is moving.

5. Provenance: Definition and Boundaries

Definition 6 (Provenance)**.** Provenance *is the verifiable causal history that connects an artifact, claim, or interaction to its origin, including the identity of the producer, the conditions of production, and the integrity of the chain through which the artifact reaches the receiver.*

The concept has a long intellectual history in art-historical scholarship and archaeology, where it has authenticated objects whose value depends on their relationship to a particular source. It has been formalized in database theory as data provenance (Buneman et al., 2008), in supply-chain management as chain-of-custody documentation (Tian, 2017), and most recently in computer science and policy as content provenance for AI-generated media (Aaronson, 2023; Longpre et al., 2024; Rosenthol, 2022). I bring it into market analysis as the foundational coordination logic that emerges when signaling fails.

5.1 Three dimensions

Three dimensions jointly constitute provenance. *Origin attribution* is the binding of an artifact to a particular producer or process; it answers, "who or what produced this." *Process attestation* is the binding of an artifact to a verifiable sequence of conditions under which it was produced; it answers, "how was this produced." *Chain integrity* is the binding of an artifact to its original state through the transmission path; it answers, "has this been altered since production." A complete proof of the claim requires all three; failure on any one renders the others inferentially weak.

5.2 Boundary distinctions

Provenance is distinct from neighboring concepts in three theoretically important ways. It is *not reputation*, which inheres in the producer across many artifacts rather than in a particular artifact's relationship to its source; a producer of impeccable reputation can produce an artifact whose provenance is contested. It is *not authenticity*, which often carries normative and biographical content (Carroll, 2015) beyond verifiable causal history; an artifact may have impeccable provenance and still be judged inauthentic on stylistic or biographical grounds. It is *not authorship*, which is a legal category concerning rights to control distribution and derivative use; provenance speaks to epistemic warrant, not legal entitlement. The conceptual neighborhood is dense, but the construct does analytical work the neighbors cannot.

5.3 Provenance versus detection

Provenance is also distinct from detection. Detection asks the receiver to determine, from the artifact alone, whether it was produced by AI or by an unaided human. Provenance asks a different question: regardless of whether AI was involved, what is the verifiable causal history of the artifact, and what does that history warrant inferring about its quality, source, or context? Detection is fundamentally adversarial, depends on generative–detection technology asymmetries, and faces well-documented limits (Chakraborty et al., 2023; Sadasivan et al., 2023). Provenance is fundamentally infrastructural, depends on the availability of attestation systems independent of the artifact's surface features, and does not require the detection problem to be solvable to be useful. The shift from detection to provenance is, accordingly, one of the central conceptual moves the framework asks the literature to adopt.

5.4 Provenance as market primitive

Provenance is not a side-show institution layered on top of otherwise-functional markets. In high-saturation markets, provenance becomes the market primitive: the elementary unit of coordination from which willingness-to-pay, willingness-to-transact, and willingness-to-trust are constructed. The provenance turn names the institutional adjustment by which this primitive comes to occupy the architectural position that signal cost occupied for the preceding half-century of market design.

Table 3. Typology of verification architectures.

Architecture	Mechanism	Representative examples	Cost & scalability	Principal failure mode
Cryptographic	Provenance bound to artifact via tamper-evident signatures attached at production.	C2PA-stamped media; signed software releases; blockchain-anchored credentials; signed LLM outputs.	High fixed cost; low marginal cost per attestation; scales globally once adopted.	Key compromise; standards fragmentation; pre-signature tampering.
Witnessed	Provenance established through observable production before identifiable third parties.	Live performance; in-person interview; proctored examination; livestreamed creation.	Low fixed cost; high marginal cost; hard to scale beyond synchronous channels.	Observer compromise; staged or selectively witnessed production; deepfaked livestreams.
Institutional	Provenance certified by an accredited intermediary whose standing depends on attestation quality.	Peer-reviewed publication; FDA approval; chartered accounting; professional licensure.	Moderate–high fixed cost; moderate marginal cost; scales through intermediary authority.	Intermediary capture; institutional decline; conflict of interest; recursive contamination.
Reputational	Provenance inferred from a sustained track record across many prior artifacts.	Open-source maintainer history; long-tenured journalist; established research lab.	Slow to establish; near-zero marginal cost once established; hard to bootstrap.	Reputation laundering; AI saturation of inputs; succession failure; reputational asset sale.

6. A Typology of Verification Architectures

Provenance is not self-establishing. It requires institutional infrastructure—technological, legal, social, or some combination—through which provenance claims are constructed, transmitted, and challenged. I use the term *verification architecture* for this infrastructure and propose a typology of four ideal types. The types differ in cost structure, scalability, failure mode, and institutional location, and these differences determine which markets they can sustain. The typology is summarized in Table 3 and visualized in Figure 2.

6.1 Distinctive properties of each architecture

Cryptographic verification scales most cheaply once standards are in place but depends on adoption and on the security of the credentialing chain; failures concentrate on standards-setting and key management. Witnessed verification is highly credible but expensive per unit and difficult to scale; failures concentrate on observer compromise. Institutional verification scales through the credibility of intermediaries but transmits intermediary failures directly to all markets they serve; failures concentrate in intermediary capture or collapse. Reputational verification is robust to single failures but slow to establish and difficult to bootstrap; failures concentrate in reputational gaming and in the homogenization of reputational systems themselves under generative AI.

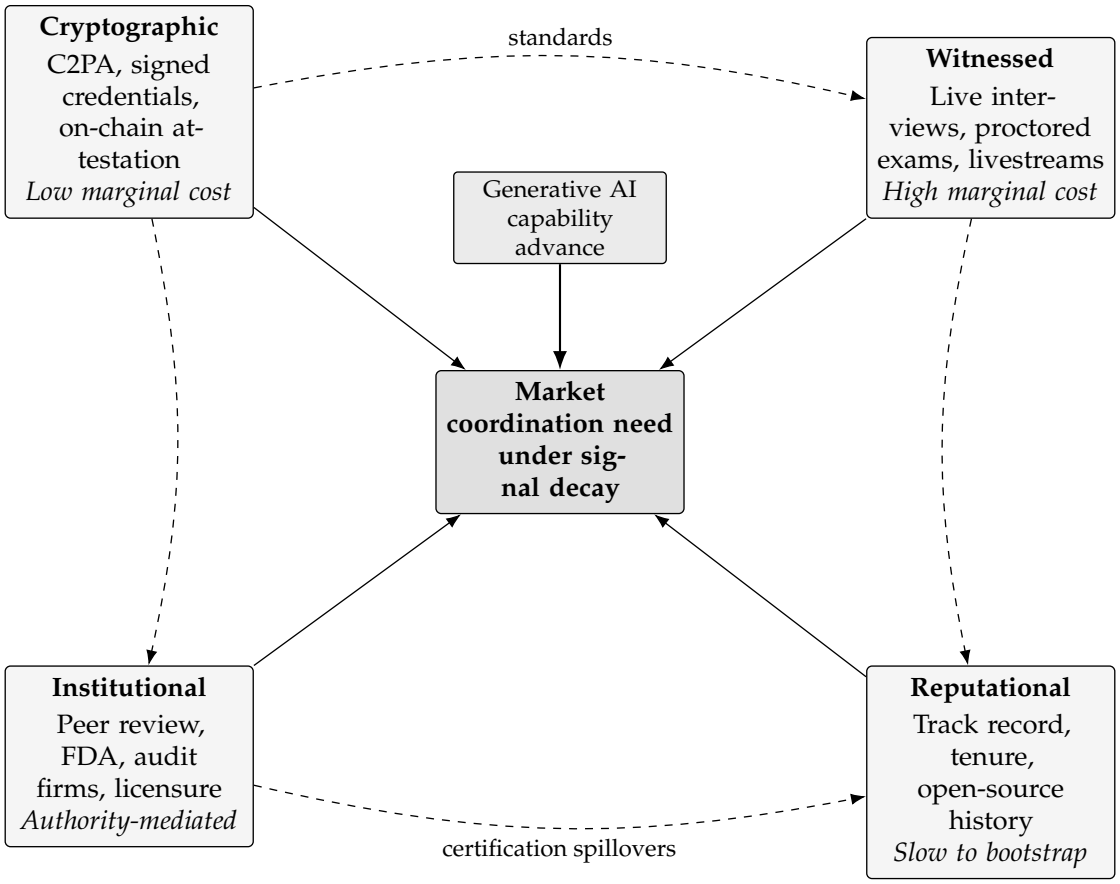


Figure 2. The verification-architecture ecosystem. Four ideal-type architectures jointly furnish provenance for a market under signal decay. Dashed arrows depict verification externalities—spillovers in standards, practices, and norms between architectures. Real markets combine multiple architectures whose relative weight is endogenous to verification tractability, cost structure, and institutional history.

6.2 Hybrid architectures in practice

Real markets routinely combine architectures. A scientific publication relies on institutional verification (the journal), reputational verification (the author's track record), and increasingly cryptographic verification (registered reports, pre-registration platforms, immutable data deposits). A premium media article may carry cryptographic provenance (C2PA), institutional provenance (the publisher's editorial process), and witnessed provenance (named bylines tied to identifiable journalists who appear publicly). A regulated drug carries institutional verification (FDA), cryptographic verification (chain-of-custody documentation), and witnessed verification (FDA inspectors). The mix is empirically determined by the cost structure of each architecture in the specific market, and it shifts over time as costs change.

6.3 The dynamics of architectural choice

Three drivers shape architectural choice. *First,* the verification *tractability* of the underlying artifact: cryptographic verification requires the artifact to be born digital or digitizable at the point of creation; witnessed verification requires the production to be synchronously observable. *Second,* the *stakes* of the underlying transaction: high-stakes transactions (pharmaceutical efficacy, securities issuance, professional licensure) sustain institutional verification's fixed costs; low-stakes transactions cannot. *Third,* the *adversarial sophistication* of the sender population: as sophistication rises, single-architecture verification becomes vulnerable and architectural redundancy becomes optimal.

The architectural choice problem is itself a research object. There is preliminary evidence that markets undergoing acute signal decay first reach for the architecture they already partially have—journalism reaches for institutional reputation, finance reaches for institutional certification, platforms reach for cryptographic standards. This path-dependence has implications both for which markets reconstitute fastest and for which sub-architectures concentrate market power.

7. A Multi-Level Framework of the Provenance Turn

The provenance turn is a multi-level phenomenon. AI capability evolves at the technology layer; signal-production costs collapse at the firm and individual layers; coordination failures emerge at the market layer; verification infrastructure develops at the institutional layer; regulation reorganizes at the policy layer. Single-level accounts capture pieces of the dynamic but miss the cross-level couplings that drive its empirical texture. Figure 3 integrates the levels into a single framework, which functions as the master conceptual figure of the paper.

7.1 Individual level: production and evaluation

At the individual level, the elementary phenomena are (i) adoption of generative tools for signal production and (ii) recalibration of evaluative trust by receivers. Adoption is approximately symmetric across sender types, by construction of generative AI as a general-purpose technology: high types do not retain a clear absolute or relative cost advantage in signal production. Receivers' trust recalibration follows. Recent behavioral work documents that receivers' confidence in their ability to distinguish AI from human content has fallen sharply over 2023–2026 (Jakesch et al., 2023; Köbis & Mossink, 2021); receiver-side trust is a falling stock.

7.2 Firm level: verification capital and the signal portfolio

At the firm level, the relevant phenomena are (i) the portfolio of signals through which the firm communicates type and quality to its audiences (consumers, regulators, investors, employees, partners) and (ii) the firm's accumulating verification capital. Firms vary widely in their verification-capital endowments. Premium media brands have substantial inherited verification capital in editorial process and bylines; consumer-platform challengers have low verification capital but high agility

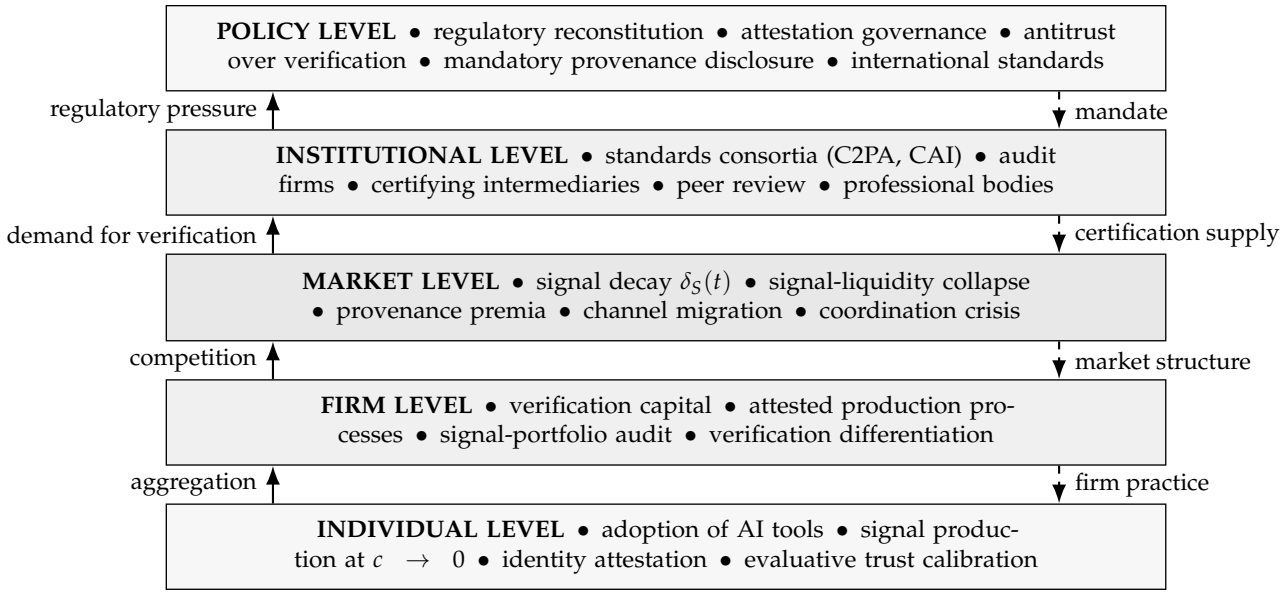


Figure 3. Master multi-level framework of the provenance turn. Generative AI capability acts on the individual level; bottom-up causal arrows (solid) aggregate through the firm, market, institutional, and policy layers. Top-down constraint arrows (dashed) flow in the reverse direction. The framework's distinctive predictions arise at level couplings.

in adopting cryptographic standards; legacy financial intermediaries have rich institutional verification capital but face provenance challenges in their own AI-generated outputs. The firm-level strategic problem is to translate verification capital into competitive advantage—or to recognize when its absence structurally disqualifies the firm from a verification-mediated market.

7.3 Market level: coordination dynamics

At the market level, the relevant phenomena are signal decay, signal-liquidity dynamics, provenance premia, channel migration, and coordination crises. Markets vary in their decay trajectories, in the speed with which verification infrastructure emerges to absorb decay, and in the channel-migration costs they face. The market level is the natural unit of analysis for the framework's central propositions.

7.4 Institutional level: verification infrastructure

At the institutional level, standards consortia, audit firms, certifying intermediaries, peer-review systems, and professional bodies supply the verification infrastructure on which provenance markets depend. The institutional layer evolves at the pace of cross-firm coordination; its emergence is uneven, path-dependent, and politically contested. Major firms and states position themselves at the institutional layer to shape the rules of the verification regime they expect to inhabit.

7.5 Policy level: regulatory reconstitution

At the policy level, regulators face the task of governing verification infrastructure—authorizing attestation providers, assigning liability for false attestation, mandating provenance disclosure, and coordinating across jurisdictions. The policy layer is the slowest-moving, the most politically contested, and—if the framework is correct—the most consequential for the long-run structure of post-decay markets.

7.6 Cross-level couplings

The framework's distinctive analytical content lies at the couplings between levels. Three are particularly consequential. *First,* the market-to-institutional coupling: market-level demand for verification creates institutional-supply incentives, but supply is lagged and uneven across markets, generating temporary disequilibria visible as coordination crises. *Second,* the institutional-to-firm coupling: institutions designing verification standards confer competitive advantage on firms positioned to influence them, creating a regulatory-capture vulnerability the literature must monitor. *Third,* the firm-to-market coupling: firms with strong verification capital both compete more effectively and re-shape market structure toward attestation-mediated coordination, accelerating the broader transition.

8. The Saturation $\times$ Tractability Framework

I propose a two-dimensional framework that locates markets undergoing the provenance turn. The horizontal axis is *Saturation Severity*: the degree to which the relevant signal type has been rendered uninformative through cheap generative production across the sender population. The vertical axis is *Verification Tractability*: the degree to which provenance for the relevant artifact, claim, or interaction can feasibly be established given current technology, institutional infrastructure, and cost constraints. The four quadrants correspond to qualitatively distinct market trajectories.

8.1 Operationalizing the two axes

Because a 2×2 invites the charge of being merely suggestive, I specify measurable proxies for each axis before deploying it. *Saturation Severity* is proxied by the product of (i) the share of the sender population with low-cost access to generative tooling in the relevant domain and (ii) one minus the type-conditioned adoption gap $|a_H - a_L|$ defined in Section 3.1; it rises as access broadens and as high- and low-type adoption converge. *Verification Tractability* is proxied by an index over (i) digital-native capture of the artifact at the point of creation (enabling cryptographic attestation), (ii) the presence of an accredited institutional certifier in the market, and (iii) the per-transaction cost of witnessed production relative to transaction value. Table 4 codes the twelve worked industries of Appendix B onto coarse (low/medium/high) values of each proxy; the coding is offered as illustrative calibration, not measurement, and is precisely the exercise Design 1 and Design 3 in

Table 4. Illustrative coding of worked industries onto the two axes (Appendix B).

Industry	Saturation	Tractability	Quadrant
Mass-market hiring funnels	High	Low	Coordination Crisis
Generic online dating	High	Low	Coordination Crisis
Low-stakes peer review at scale	High	Low–Med	Coordination Crisis
Higher-education admissions	High	Med	Reconstitution (in transit)
Consulting / professional services	High	Med	Reconstitution (in transit)
News media / journalism	High	Med	Reconstitution (in transit)
Creator economy	High	Med (bifurcating)	Reconstitution / Crisis split
Scientific publishing	High	Med–High	Provenance Reconstitution
Financial markets / equity research	Med–High	High	Provenance Reconstitution
Healthcare (non-clinical content)	Med	High	Anticipatory Verification
Pharma chain-of-custody	Low–Med	High	Anticipatory Verification
Niche creative / embodied trades	Low	Low	Pre-Decay Signaling

Section 14 would replace with estimated quantities.

8.2 Four predictions

The framework yields four predictions. *First,* lower-left markets—low saturation, low tractability—are temporary configurations that persist only as long as generative capability has not advanced into their signal domain, and most are expected to migrate rightward. *Second,* upper-left markets—low saturation, high tractability—represent verification investment that exceeds current need, whether defensive or inherited from prior coordination problems. *Third,* upper-right markets—high saturation, high tractability—are the new equilibrium destination: signal cost has collapsed, but verification is cheap enough to support coordination through provenance. *Fourth,* lower-right markets—high saturation, low tractability—are the danger zone: they face acute coordination pressure but cannot easily build the verification infrastructure that would absorb it, and they migrate to thicker channels, restructure, or collapse into adverse selection.

8.3 Movement within the matrix

A market's position is not fixed. Saturation rises mechanically with the diffusion of generative capability. Verification tractability rises with investment in attestation infrastructure, standards development, and institutional learning. Strategic actors can move markets rightward (typically not desired by incumbents but often inevitable) or upward (a productive response). The most consequential strategic and policy question of the provenance turn is which markets move upward fast enough to absorb the rightward drift, and which are left in the lower-right danger zone.

8.4 Coordination-crisis exits

For markets caught in the lower-right quadrant, three exits exist (Figure 5). The choice among them is determined by the relative cost of building verification infrastructure (which moves the market

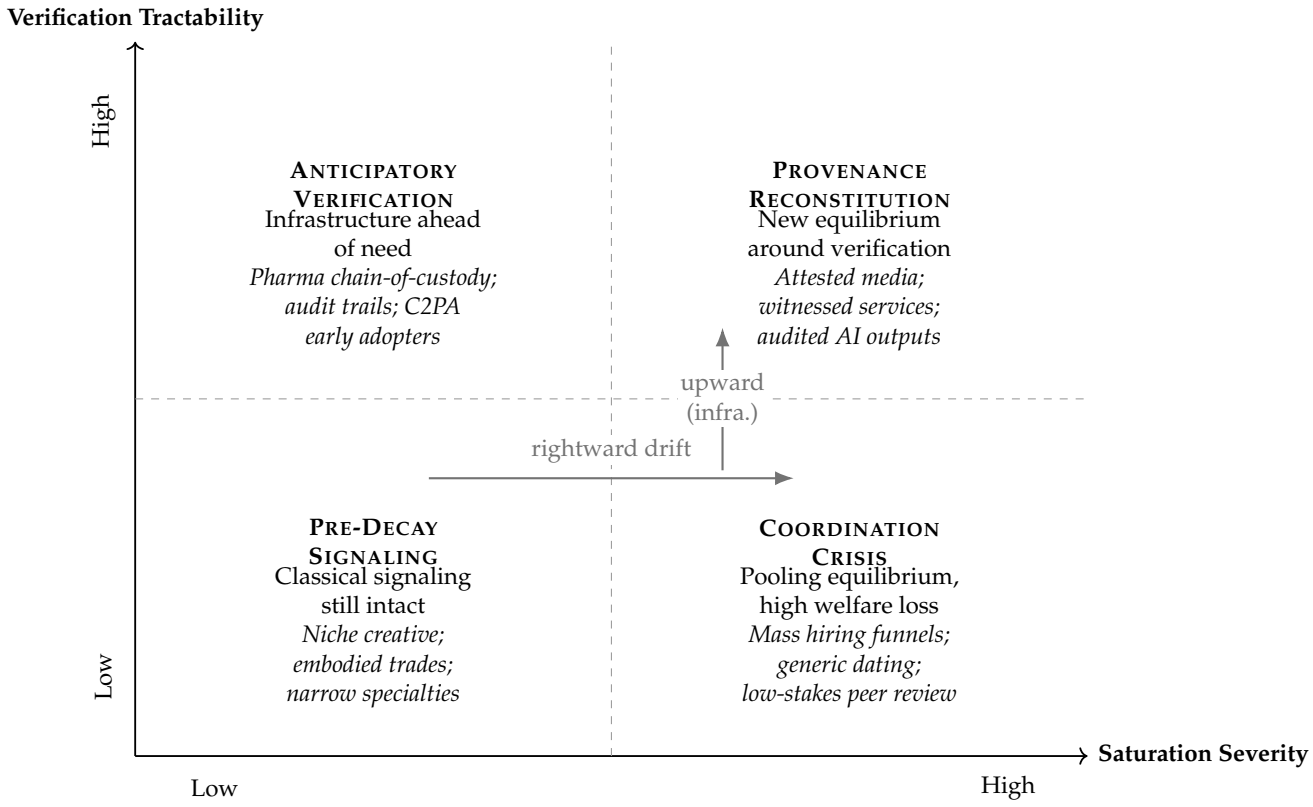


Figure 4. The Saturation × Tractability framework. Four quadrants name qualitatively distinct market trajectories under signal decay. The temporal drift of typical markets is rightward (rising saturation as generative capability diffuses); the productive response is upward (investment in verification tractability). Markets in the lower-right "coordination-crisis" quadrant either build verification infrastructure rapidly, migrate to provenance-rich channels, or face partial market collapse.

upward into Reconstitution), of migrating to thicker channels (which moves the transaction off the current platform), and of accepting the welfare loss of a partial-pooling equilibrium.

9. Six Mechanisms of the Provenance Turn

Six mechanisms connect the diffusion of generative AI to the reorganization of markets around provenance. They are analytically distinct but operate jointly in most affected markets.

Mechanism 1: Cost Asymmetry Collapse

The foundational mechanism is the collapse of cost asymmetries across sender types for a given signal. Generative AI provides uniform production capacity at near-zero marginal cost. The separating condition—that low types bear higher costs than high types for the same signal—is violated by construction, and the signal no longer separates. The mechanism applies to any signal whose production has been reduced to a generative AI task, an expanding but not universal set.

Table 5. Four trajectories under signal decay: Saturation × Tractability.

	Low Saturation Severity	**High Saturation Severity**
High Verification Tractability	ANTICIPATORY VERIFICATION. Infrastructure exists but is not yet needed. Logic: investment ahead of need, inherited or defensive. Examples: pharma chain-of-custody; financial audit trails; premium media early C2PA adoption. Trajectory: stable until saturation rises; smooth transition into Reconstitution.	PROVENANCE RECONSTITUTION. Signaling has degraded but markets coordinate through cheap, credible verification. Logic: new equilibrium organized around the provenance premium; rents accrue to verification infrastructure. Examples: standards-attested digital content; witnessed professional services; audited AI outputs. Trajectory: success destination of the turn.
Low Verification Tractability	PRE-DECAY SIGNALING. Signaling continues to function; AI has not penetrated the signal domain. Logic: classical separating equilibrium remains in force. Examples: niche creative domains; embodied performance markets; narrow technical specialties. Trajectory: temporary; expected to drift rightward as capability broadens.	COORDINATION CRISIS. Signaling has degraded but no verification infrastructure absorbs the disruption. Logic: pooling equilibrium with high welfare loss; pressure toward channel migration, relational exchange, or partial collapse. Examples: mass-market hiring funnels; generic online dating; low-stakes peer review at scale. Trajectory: danger zone; exits via migration, infrastructural development, or unraveling.

Mechanism 2: Detection Asymmetry Exhaustion

A natural first response is to invest in detection. It is bounded by a structural asymmetry: generation and detection technologies are typically developed by the same community, often the same firms, and generation has consistently outpaced detection (Chakraborty et al., 2023; Sadasivan et al., 2023). Receivers who rely on detection face an arms race they are likely to lose, and a regime in which dependable detection is impossible exhausts the response, forcing markets toward a coordination logic that does not depend on inferring AI involvement from surface features.

Mechanism 3: Verification Infrastructure Emergence

Where detection fails, verification infrastructure emerges. Cryptographic standards such as C2PA develop in concert with industry consortia; witnessed-production formats expand; institutional intermediaries adapt credentialing to attest to process rather than (or in addition to) outcome; reputation systems incorporate provenance signals. The emergence is uneven: it requires standards-setting effort, institutional coordination, and—often—regulatory scaffolding. Markets in which one or more is missing face the danger-zone trajectory.

Mechanism 4: Channel Migration

Where verification infrastructure cannot be built quickly enough, activity migrates to channels that carry provenance natively—written application materials yield to in-person interviews, polished images to live demonstrations, asynchronous communication to synchronous video. Migration is costly (it sacrifices scale and asynchrony) but converts the problem from signal verification (which has failed) to identity verification at a point in time (which remains tractable).

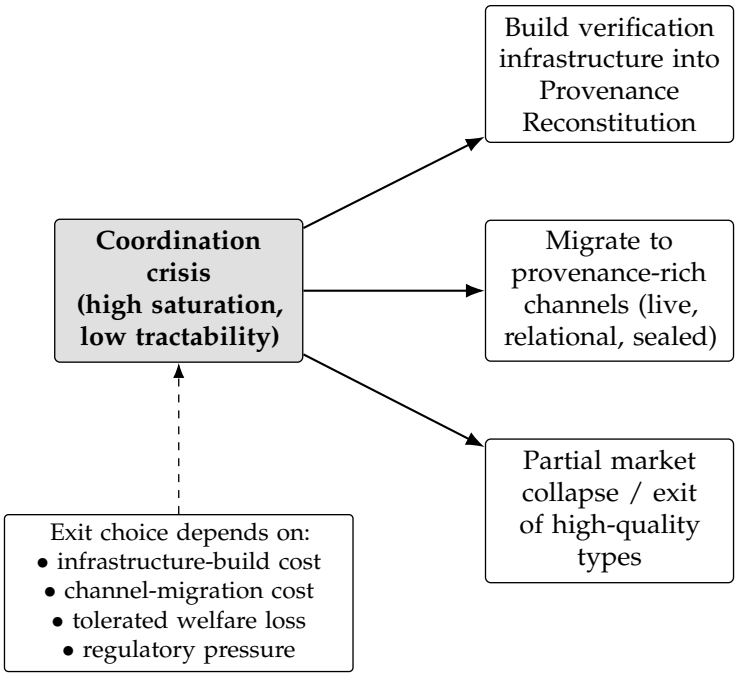


Figure 5. Coordination-crisis exits. Markets in the lower-right quadrant exit the crisis through one of three pathways. The choice is endogenous to the relative cost of infrastructure-building, channel migration, and welfare-loss tolerance, and to the regulatory environment.

Mechanism 5: Recursive Reputation Saturation

Reputation systems are not exempt from signal decay. They depend on inputs—reviews, endorsements, citation counts, contribution histories—that are themselves generable at scale by AI. As inputs saturate, the content of reputational outputs degrades, and the platforms that mediate reputation (Amazon, LinkedIn, Yelp, citation graphs) become recursive coordination targets. Recursive saturation explains why a naive "reputation will suffice" response fails: reputation systems are themselves signaling systems whose inputs are now cheap.

Mechanism 6: Verification Arbitrage

The collapse of signal cost and the emergence of verification infrastructure create arbitrage opportunities for intermediaries with cross-market verification capital. A firm that has built witnessed-production capability for one market (academic proctoring, say) can extend it to adjacent markets (proctored hiring assessments, witnessed contractor work); a cryptographic-provenance operator for digital media can extend to research data, legal documents, and financial filings. The arbitrage shapes verification-infrastructure consolidation and predicts which firms emerge as dominant cross-market attestors.

9.1 Joint operation

The six operate jointly. Cost-asymmetry collapse forces the problem; detection-asymmetry exhaustion eliminates the easy response; infrastructure emergence and channel migration are the two principal durable responses; recursive saturation explains why reputation-based hedges are themselves vulnerable; and verification arbitrage shapes the industrial organization of the emerging verification economy. The mix observed in a given market is determined by its verification tractability, by channel-migration costs, and by the institutional and regulatory environment.

10. Marketing Implications

The provenance turn restructures the marketing function. I articulate six implications and integrate them with the established literatures on customer journey, brand equity, and consumer trust.

Implication M1: Customer journeys reorganize around verification touchpoints. Lemon & Verhoef (2016) provided a now-canonical map of the customer journey through awareness, consideration, purchase, and post-purchase phases. In high-saturation markets, the journey acquires an additional architecture: verification touchpoints embedded throughout, at which the consumer (or the consumer's AI agent) checks provenance against the brand's claims—at search (synthetic-listing flags), evaluation (provenance-stamped reviews), purchase (chain-of-custody documentation), and post-purchase (origin attestation in receipt and loyalty flows). Marketing managers must architect the verification journey alongside the brand journey.

Implication M2: Brand equity bifurcates into reputational and verifiable components. Classical brand equity (Aaker, 1996; Keller, 2003) is a reputational construct accumulated through repeated exposure and consistent quality. In high-saturation markets it bifurcates. One component remains reputational and is vulnerable to recursive saturation (Mechanism 5). The other is *verifiable*: the brand's accumulated verification capital—standards-consortium positions, cryptographic identity infrastructure, witnessed-production formats. Firms that fail to develop the verifiable component will find their reputational component decay faster than competitors'.

Implication M3: The creator economy stratifies by verification access. The creator economy (Hofstetter & Gollnhofer, 2024; Leung et al., 2022) is a paradigmatic high-saturation market: synthetic creators, AI-generated content, and bot-driven engagement collapse the cost asymmetry on which creator brands depended. The market stratifies: creators with verification infrastructure (signed identity, livestreaming capacity, platform attestation) capture a premium; those without commodify. Platforms become focal points of verification governance, with disproportionate rents accruing to those that supply credible attestation across many creators.

Implication M4: AI agents reorganize the demand side. As AI agents act on behalf of consumers (Davenport et al., 2020), the demand-side actor is increasingly an agent rather than a human. Agents are systematically more reliant on provenance-rich signals than humans, both because they process attestation metadata at scale and because they cannot use the embodied trust cues humans rely on. Agentic commerce magnifies the provenance premium and accelerates the institutional emergence of cryptographic verification.

Implication M5: Synthetic media create new categories of brand risk. Deepfaked CEO statements, synthetic brand-affiliated endorsements, and AI-generated reviews create a new class of brand risk. It is partially mitigable by cryptographic provenance for official content, but the asymmetry between unattested impersonation and attested brand content is severe. Firms will invest in brand-monitoring infrastructure to detect unattested brand-adjacent content and in legal infrastructure to assign liability when impersonation causes harm.

Implication M6: Marketing measurement requires new instruments. Survey-based brand tracking, A/B-test conversion measurement, and attribution modeling all assume the integrity of the digital signals they aggregate. As signal decay proceeds, those inputs are contaminated by synthetic responses, AI-generated reviews, and bot-driven traffic. Measurement infrastructure that uses cryptographic provenance to filter inputs and estimates synthetic-information contamination will become a core marketing-research capability.

10.1 Marketing–Finance Integration

The provenance turn closes part of the gap between marketing and finance. Two integrative arguments deserve emphasis.

First, verification capital is a real asset that affects firm valuation. Investors increasingly price firms not only on cash-flow expectations but on their position in verification infrastructure. A media firm with a credible C2PA implementation and a standards-consortium voice trades at a different multiple than one without; an asset manager with auditable AI in its research workflow trades at a different multiple than one without. The marketing function's investments in verification infrastructure are no longer cost centers; they sit on the firm's quasi-balance-sheet of intangible assets.

Second, disclosure regimes are converging. Marketing claims (advertising substantiation, comparative claims, sustainability claims) and financial disclosures (regulatory filings, earnings statements) are converging on a common verification frame in which the underlying claim is tied to verifiable origin. Firms that present a unified attestation infrastructure across marketing and finance disclose more credibly than firms maintaining separate, less rigorous marketing-claim verification. The implication is operational: legal, marketing, and investor-relations functions will increasingly share verification infrastructure and the head-of-trust role that administers it.

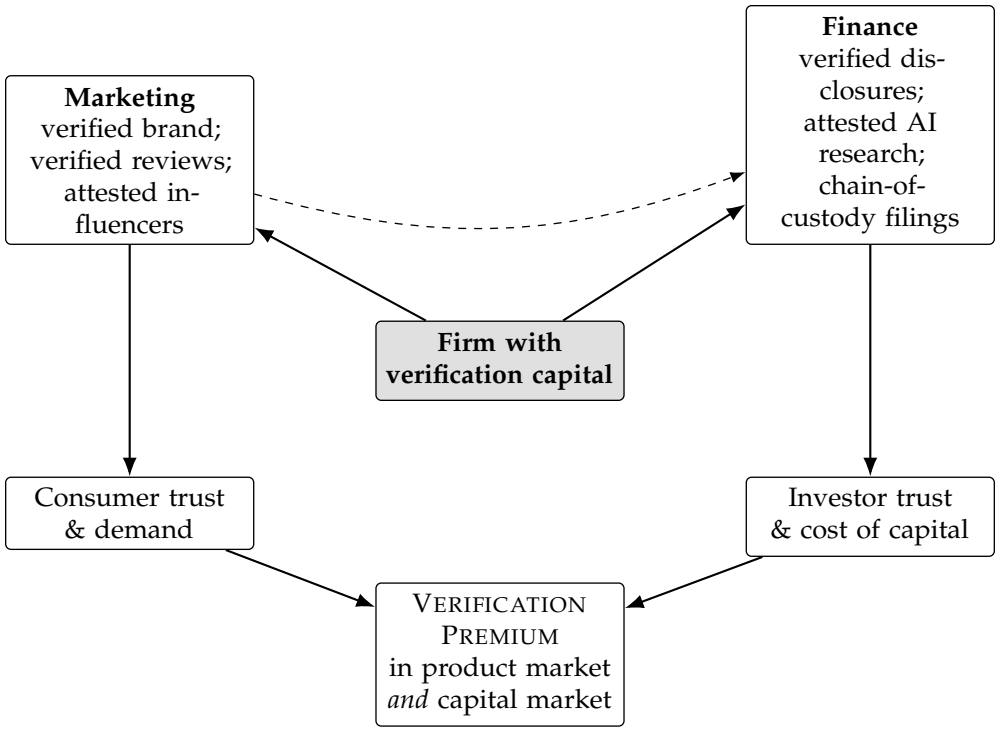


Figure 6. Marketing–finance integration through verification capital. A firm's verification infrastructure simultaneously underwrites marketing claims to consumers and disclosures to investors, generating premia in both product and capital markets. The dashed arrow captures the increasingly shared infrastructure between the two functions.

11. Finance Implications

The provenance turn restructures the institutional foundations of capital markets. I articulate five implications.

Implication F1: Synthetic information risk becomes a price-relevant risk factor. As deepfaked statements, synthetic press releases, AI-generated analyst reports, and fabricated filings proliferate, participants face a new risk factor whose moments depend on the prevalence and detection latency of synthetic information (Cong et al., 2021; Goldstein et al., 2021). The 2023 Pentagon-explosion deepfake episode, which briefly moved equity markets, is the prototype event (Bateman, 2020). Trading systems will price the local probability of synthetic information; microstructure will adjust quote depth and bid–ask spreads in response. Volatility around firm-specific news will rise where attestation is weak and fall where it is strong.

Implication F2: Verifiable disclosure becomes a competitive disclosure standard. Voluntary disclosure theory (Verrecchia, 2001) predicts that firms with better news disclose; the provenance turn extends the prediction. Firms able to credibly attest to the origin of their disclosures will disclose more, sooner, and in greater detail, because the marginal credibility of their disclosure

is higher. The implication is testable: cross-sectional adoption of cryptographic provenance for disclosure should predict disclosure frequency and credibility shocks (analyst-revision responses, abnormal returns).

Implication F3: Equity-research production shifts toward attested AI. The growing use of LLMs in equity research (Lopez-Lira & Tang, 2023) faces a credibility ceiling without attestation: a report whose AI use is unattested cannot be evaluated for known failure modes. Desks that adopt attested-AI workflows—signed prompts, logged retrieval, witnessed human review—will capture readership and trading flow disproportionately, and regulatory-disclosure standards for AI-assisted research will be promulgated by FINRA, the SEC, ESMA, and analogous bodies.

Implication F4: Tokenization and on-chain provenance reconfigure secondary markets. Tokenized assets (Catalini & Gans, 2020; Harvey et al., 2021; Yermack, 2017) carry native provenance: every transfer is recorded on a tamper-evident ledger. As synthetic information risk rises in conventional markets, the comparative advantage of tokenized markets in provenance widens. I expect partial migration of trading and ownership of certain asset classes (real estate, private credit, fund interests, certain commodities) to tokenized infrastructure where regulatory regimes accommodate it.

Implication F5: Auditing emerges as a frontline AI-governance function. As firms rely on AI for financial production—reporting, treasury, risk, compliance—the audit perimeter expands to encompass AI systems. Audit firms develop AI-attestation practices verifying model provenance, training-data integrity, and prompt-and-output chains. The Big Four become among the most important AI-governance infrastructure outside the major model labs, and audit-firm competitive dynamics are reshaped by the speed and scale of AI-attestation capability build.

12. Governance and Policy Implications

The provenance turn reorders the public-policy agenda for AI. Three policy domains warrant new attention; a fourth is implicit in their joint operation.

Domain P1: Standards governance. Cryptographic provenance standards function as market infrastructure with substantial network effects. Their governance—who sets them, who can certify, how disputes are adjudicated, how cross-border interoperability is maintained—becomes a competition-policy question of first-order significance. Regulators with little prior experience in standards governance must develop it quickly, drawing on analogies from financial standards-setting (IFRS, IOSCO) and internet governance (ICANN, IETF).

Domain P2: Attestation liability. When an attestation provider certifies a provenance that proves false, who bears the loss? Existing tort, contract, and consumer-protection frameworks are poorly adapted to attestation failures, particularly across jurisdictions. A new body of attestation-liability law will develop, and the choices made in its formative period will shape verification markets for decades. The credit-rating-agency liability literature (Partnoy, 2006) offers a useful but incomplete precedent.

Domain P3: Mandatory provenance disclosure. Several jurisdictions are considering mandates that AI-generated content carry machine-readable provenance markers. Such mandates can address coordination failures private action cannot solve, but they raise hard design choices: which content must be marked, who is responsible, how non-compliance is detected, and how the marking infrastructure is made tamper-resistant. The EU AI Act, the U.S. Executive Order on AI, and emerging Indian, Chinese, and Brazilian frameworks all gesture toward provenance mandates differently; the global landscape is fragmenting.

Domain P4 (cross-cutting): Verification-market antitrust. As verification infrastructure consolidates, the competition-policy stakes become acute. A small number of attestation providers may end up gating large parts of the digital economy. Antitrust authorities will need doctrine for verification markets analogous to what they have developed for credit-rating agencies, payment networks, and operating-system platforms. The Sherman Act, the EU Digital Markets Act, and India's Competition Act are all imperfect instruments; new institutional thinking is required.

12.1 Reframing the policy frame

The underlying frame should shift from *preventing deception* (the natural framing under existing consumer-protection law) to *enabling verification* (the framing the provenance turn suggests). Anti-deception regulation focuses on identifying bad actors and assigning liability after harm; verification-enabling regulation focuses on building infrastructure that makes harm legible *ex ante*. The two imply different regulatory tools, agency competencies, and international coordination requirements.

13. Propositions

I derive seventeen propositions across six domains: signal-decay dynamics (P1–P3); provenance premia (P4–P6); market-structural responses (P7–P9); institutional and competitive consequences (P10–P12); marketing-specific consequences (P13–P14); and finance-specific consequences (P15–P17). Each is intended as a falsifiable claim with a stated disconfirming condition. For each I note the theoretical logic, boundary conditions, candidate moderators, and an empirical strategy. The propositions are summarized in Table 7.

Table 6. Emerging governance mechanisms for the verification era.

Mechanism	Form	Function	Comparable precedent
Cryptographic provenance standards	Voluntary industry consortium with cross-recognition	Tamper-evident binding of content to producer; cross-platform interoperability	W3C, ICANN, IETF
Attestation-provider licensing	State-administered or self-regulatory registry of certified attestors	Quality assurance and liability anchor for verification claims	SEC-registered ratings agencies; auditor registration
Mandatory provenance disclosure	Statutory marking of AI-generated content; safe-harbor regimes	Coordination of receiver expectations; preserved market signal	EU AI Act provisions; U.S. EO on AI
Attestation-liability frameworks	Tort doctrine and contractual standards for false attestation	Risk allocation and incentive design for providers	Credit-rating-agency liability law
Verification-market antitrust	Competition oversight of standards bodies and dominant attestors	Prevention of verification-layer market power	Payment-network antitrust; DMA gatekeeper rules
Cross-border attestation recognition	Multilateral agreements on mutual recognition of attestations	International tradeability of attested artifacts	Apostille Convention; mutual recognition agreements

Signal-Decay Dynamics

Proposition 1 (Heterogeneous Decay Rates)**.** *Within a given period, signal types vary systematically in their decay rates according to (i) the maturity of generative AI in the relevant production domain, (ii) the rate of adoption among the relevant sender population, and (iii) the cost the receiver bears for each unsuccessful evaluation. Cross-signal-type studies controlling for these factors will explain a substantial share of variance in market-level disruption.*

Boundary conditions: requires AI available to the sender population at low cost; held off by domain-specific capability lags. *Moderators:* sender sophistication; receiver stakes; lock-in to existing channels. *Disconfirmed if:* decay rates are statistically indistinguishable across signal domains once access is held constant. *Empirical strategy:* cross-signal panel with decay curves derived from receiver classification accuracy over time.

Proposition 2 (Saturation Threshold Effects)**.** *Signal decay exhibits threshold dynamics rather than smooth degradation. Markets remain qualitatively functional under signaling until a saturation threshold (varying by market) is crossed, after which the market shifts into either a provenance regime, a channel-migrated regime, or an adverse-selection regime. Market behavior in the neighborhood of the threshold should show discontinuities in coordination outcomes.*

Disconfirmed if: coordination outcomes degrade linearly in saturation with no detectable break. *Empirical strategy:* regression-discontinuity around documented capability shocks (GPT-4 release, image-generation milestones); structural-change tests on transaction-quality series.

Proposition 3 (Receiver Burden Reallocation)**.** *Under signal decay, the coordination cost of the market is reallocated from senders (who previously bore the cost of credible signaling) to receivers (who now bear the cost of verifying provenance or migrating to thicker channels). Total coordination cost may rise or fall depending on verification tractability, but the locus shifts predictably toward receivers.*

Disconfirmed if: measured verification effort does not shift toward the receiver side following a capability shock. *Empirical strategy:* time-use studies of receivers (admissions officers, hiring managers, journal editors) before and after shocks; cost-of-verification audits.

Provenance Premia

Proposition 4 (Provenance Premium)**.** *Consider two artifacts identical in observable content that differ only in carrying a credible provenance attestation. In a market at saturation level σ and tractability level τ, the price premium $\pi(\sigma, \tau)$ commanded by the attested artifact satisfies $\partial\pi/\partial\sigma > 0$ (the premium rises with saturation) and $\partial\pi/\partial\tau < 0$ above a threshold tractability (it shrinks once low-cost verification commoditizes the basic claim), yielding an inverted-U in τ. At σ near zero the premium is statistically zero. The non-monotonicity in τ and the zero at low σ are what distinguish this claim from the definitional statement that attestation is valued.*

Disconfirmed if: the premium is flat in saturation, or monotone (rather than inverted-U) in tractability, or significantly positive in markets with negligible saturation. *Empirical strategy:* hedonic price analysis comparing C2PA-stamped versus unstamped premium-media subscriptions across markets of differing saturation; conjoint experiments isolating willingness-to-pay for attested versus unattested versions of identical content; field experiments on platform listings.

Proposition 5 (Witnessed-Channel Premium)**.** *Channels that natively carry provenance—live performance, in-person interaction, proctored assessment, sealed-process production—command rising premia under signal decay, even when the substantive content of the underlying transaction is unchanged.*

Disconfirmed if: live-versus-asynchronous price gaps do not widen as saturation rises. *Empirical strategy:* event studies around capability shocks; price-time-series for live versus asynchronous formats in matched product categories.

Proposition 6 (Provenance-Elasticity Asymmetry)**.** *Provenance elasticity η_P is higher in markets with (i) high receiver-side stakes, (ii) high prior trust in the underlying signal class, and (iii) regulatory or fiduciary requirements for verifiability, and lower in hedonic-experiential markets where consumers value the artifact independent of its origin.*

Disconfirmed if: estimated elasticities do not order across these market types. *Empirical strategy:* between-market comparative conjoint studies; field experiments varying attestation salience across product categories.

Market-Structural Responses

Proposition 7 (Verification-Architecture Concentration)**.** *The provision of verification infrastructure exhibits scale economies and network externalities, producing concentration in attestation provision. Markets dependent on cryptographic or institutional verification will see consolidation among a small number of attestation providers, with associated implications for market power and regulatory attention.*

Disconfirmed if: attestation provision remains fragmented (low Herfindahl) as markets mature. *Empirical strategy:* Herfindahl indices of verification providers across markets; case studies of standards-consortium membership dynamics.

Proposition 8 (Strategic Channel Migration)**.** *Firms whose competitive position relies on signal-rich asynchronous channels will, under sustained signal decay, reorganize value capture around migration to provenance-rich alternatives. The shift from text-heavy job applications to video-and-interview processes, from product reviews to influencer-verified live demonstrations, and from email outreach to relationship-mediated introductions are early manifestations.*

Disconfirmed if: channel mix is stable in high-saturation markets relative to low-saturation controls. *Empirical strategy:* channel-composition time series on major platforms (LinkedIn, Indeed, Amazon, dating platforms); within-firm channel-mix changes by saturation.

Proposition 9 (Coordination Crisis in Low-Tractability Markets)**.** *Markets in the lower-right quadrant of the Saturation $\times$ Tractability framework will exhibit coordination crisis: rising disputes over quality, falling transaction volumes, increasing reliance on relational rather than impersonal exchange, and migration of high-value transactions outside the original market structure.*

Disconfirmed if: lower-right markets sustain stable dispute and completion rates. *Empirical strategy:* time-series of dispute rates, contract-completion rates, and channel composition of high-value transactions in candidate lower-right markets.

Institutional and Competitive Consequences

Proposition 10 (Regulatory Reconstitution)**.** *Regulatory attention will shift from anti-deception enforcement (attributing misrepresentation to identifiable bad actors) toward attestation governance (certifying verification infrastructures and assigning liability for their failure). The shift will be visible in the outputs of competition, consumer-protection, and sectoral bodies, and will create new categories of regulated entity without clear analogues in pre-AI regulation.*

Disconfirmed if: regulatory output remains dominated by anti-deception instruments through the observation window. *Empirical strategy:* content analysis of regulatory filings, hearing transcripts, and rulemaking dockets 2023–2030; emergence of attestation-provider registries.

Proposition 11 (Verification Governance as Competitive Asset)**.** *Firms' competitive position becomes increasingly determined by their position in verification governance: whether they sit inside or outside standards-setting consortia, whether they operate attestation infrastructure or depend on others', and their*

ability to credibly bind themselves to provenance commitments.

Disconfirmed if: consortium membership and attestation-infrastructure ownership carry no measurable competitive consequence. *Empirical strategy:* match standards-consortium membership to subsequent stock-price reactions, customer adoption, and partnership formation; difference-in-differences on standards-adoption events.

Proposition 12 (Verification-Capital Premium)**.** *Construct a firm-level verification-capital index V from consortium memberships, attestation-infrastructure ownership, audit relationships, and cryptographic-identity deployments. Conditional on conventional brand equity B, productivity, and growth, V carries an* independent *positive association with firm value, and the partial association $\partial(\text{Tobin's } q)/\partial V \mid B$ is increasing in the average saturation severity of the firm's markets. The claim is non-trivial precisely because it requires V to retain explanatory power after B is controlled: if verification capital were merely relabeled brand equity, the partial coefficient would vanish.*

Boundary conditions: requires V and B to be separately measurable, which the proposed index and standard brand-equity measures permit; identification rests on cross-sectional and within-firm variation in V holding B fixed. *Disconfirmed if:* the partial coefficient on V is indistinguishable from zero once B is included, or does not increase with market saturation. *Empirical strategy:* cross-sectional and panel Tobin's q regressions with both indices entered jointly; falsification check via variance-inflation and orthogonalization of V against B.

Marketing-Specific Consequences

Proposition 13 (Synthetic-Influencer Crowd-Out)**.** *In markets with high creator-side AI adoption, the share of engagement and revenue captured by authenticated (witnessed and/or cryptographically attested) creators rises monotonically with saturation, displacing unauthenticated creators independent of content quality.*

Disconfirmed if: authenticated-creator share is flat in saturation, or explained away by content quality. *Empirical strategy:* platform-level panel on creator revenue, engagement, and authentication status; instrumental-variable strategy using rollout timing of authentication features.

Proposition 14 (Verification-Touchpoint Conversion Lift)**.** *Insertion of verification touchpoints (provenance badges, witnessed claims, cryptographically attested credentials) into the customer journey produces measurable conversion lift in high-saturation markets, with diminishing lift where saturation is low.*

Disconfirmed if: touchpoint lift does not vary with saturation. *Empirical strategy:* platform-side randomized experiments on attestation visibility; field experiments with retailers and media subscribers.

Finance-Specific Consequences

Proposition 15 (Synthetic-Information Volatility Premium)**.** *In equity markets, the volatility of returns around firm-specific news rises with the prevalence of synthetic information in the firm's industry and falls with the firm's adoption of verifiable-disclosure infrastructure.*

Disconfirmed if: news-window volatility is unrelated to synthetic-content prevalence or to disclosure-attestation adoption. *Empirical strategy:* event studies around firm news using synthetic-content prevalence as a treatment proxy; difference-in-differences on adoption of cryptographic disclosure.

Proposition 16 (Audit-Firm AI-Attestation Wedge)**.** *Audit firms that develop attested-AI auditing capabilities capture disproportionate share of audit contracts among AI-intensive clients and command price premia in those engagements.*

Disconfirmed if: AI-attestation capability carries no share or fee advantage among AI-intensive clients. *Empirical strategy:* audit-firm–client matched panel; cross-sectional regression of audit-fee growth on AI-attestation capability measures.

Proposition 17 (Tokenized-Market Provenance Wedge)**.** *Asset classes with native on-chain provenance command lower information-asymmetry risk premia than conventionally traded counterparts of equivalent fundamentals, controlling for liquidity and counterparty risk.*

Disconfirmed if: no premium differential survives controls for liquidity and counterparty risk. *Empirical strategy:* compare tokenized and non-tokenized real-estate, private-credit, and fund-interest pricing controlling for fundamentals; structural microstructure modeling of effective-spread decomposition.

14. An Empirical Research Agenda

The framework supports a substantial empirical agenda. I sketch eight research designs that, together, would test its central claims. Each is feasible with currently available methods and data sources; several are already underway in adjacent literatures.

Design 1: Decay-curve estimation across signal types. For a set of signal types (application essays, product reviews, analyst notes, dating profiles, code submissions), compile receiver-classification accuracy panels across the 2022–2026 capability era. Estimate decay curves by signal type and test the heterogeneity prediction (P1) against signal-domain capability maturity, sender adoption, and receiver stakes.

Design 2: Capability-shock event studies. Use major capability releases (GPT-4, GPT-5, multimodal milestones) as exogenous shocks. Estimate structural breaks in transaction-quality outcomes (conversion, dispute, completion) in candidate high-saturation markets versus low-saturation controls. Tests P1, P2, P9.

Table 7. Seventeen propositions and proposed empirical strategies.

#	Claim	Empirical strategy
P1	Heterogeneous decay rates across signal types	Cross-signal panel, decay-curve estimation
P2	Saturation threshold effects (non-linear decay)	Regression-discontinuity around capability shocks
P3	Receiver burden reallocation	Time-use studies of evaluators pre/post shock
P4	Provenance premium (inverted-U in tractability)	Hedonic pricing; conjoint experiments
P5	Witnessed-channel premium	Time-series of live vs. asynchronous channel prices
P6	Provenance-elasticity asymmetry across markets	Between-market conjoint comparisons
P7	Verification-architecture concentration	Herfindahl indices of verification providers
P8	Strategic channel migration by firms	Channel-composition panels on major platforms
P9	Coordination crisis in low-tractability markets	Dispute / completion / channel-composition series
P10	Regulatory reconstitution toward attestation governance	Content analysis of regulatory output
P11	Verification-governance position as competitive asset	Standards-consortium event studies
P12	Verification-capital premium, orthogonal to brand equity	Joint Tobin's q regressions; orthogonalization check
P13	Synthetic-influencer crowd-out by authentication	Platform creator-revenue panels
P14	Verification-touchpoint conversion lift	Randomized field experiments
P15	Synthetic-information volatility premium	Event studies around firm news
P16	Audit-firm AI-attestation wedge	Audit-firm–client matched panel
P17	Tokenized-market provenance wedge	Tokenized vs. conventional pricing comparison

Design 3: Hedonic provenance-premium estimation. Identify product categories with adjacent attested and unattested variants (premium media with and without C2PA labels; products with and without on-chain provenance). Estimate the price premium attributable to attestation across markets of differing saturation, controlling for fundamentals. Tests the inverted-U structure of P4.

Design 4: Conjoint experiments isolating provenance elasticity. Run large-scale conjoint experiments across consumer segments, varying attestation type, provider, and product characteristics. Estimate provenance elasticities by segment and category. Tests P4 and P6.

Design 5: Field experiments on verification touchpoints. Partner with platforms (e-commerce, hiring, dating, media) to randomly vary the visibility of verification touchpoints. Estimate conversion, retention, and willingness-to-pay effects. Tests P14.

Design 6: Digital-trace audits of synthetic-content prevalence. Construct industry-level prevalence estimates of synthetic content (reviews, news, social posts) using best-available detection plus provenance-stamp tracking. Use the prevalence series as treatment in panel identification of synthetic-information effects on prices, volatilities, and engagement. Supports P9, P13, P15.

Design 7: Standards-consortium event studies. Identify firms joining (and leaving) major provenance consortia (C2PA, CAI, regional analogues). Use event-study methodology on stock returns around membership events. Tests P11 and P12.

Design 8: Cross-jurisdictional regulatory comparison. Compile a panel of jurisdictional adoption of provenance mandates and attestation-licensing regimes. Use difference-in-differences on firm-level outcomes in regulated versus unregulated jurisdictions. Tests P10 and informs Domain P3.

14.1 Audit studies of large language models

A distinctive class of design uses the LLMs themselves as instruments. *LLM audit studies* can (i) assess how AI agents evaluate attested versus unattested artifacts, (ii) measure the synthetic-content production frontier across model versions, and (iii) test whether AI-mediated demand-side actors are more or less reliant on provenance than human consumers (Implication M4). The methodological challenges (rapidly changing models, prompt sensitivity, training contamination) are real but manageable with appropriate design (Köbis & Mossink, 2021).

14.2 Comparative-historical research

Comparative-historical research on prior signal-collapse episodes (cheap printing; mass advertising; the early internet) would clarify what the provenance turn shares with predecessors and what is distinctive (Section 1.2). Each prior episode produced new verification institutions, took decades to stabilize, and restructured the political economy of information. A serious comparative-historical literature is essential for calibrating expectations and identifying mechanisms that recur across episodes.

15. Boundary Conditions and Competing Explanations

15.1 Boundary conditions

The framework is bounded in four ways. *First,* it applies to markets coordinated through signals; markets coordinated primarily through direct observation (most physical commodities) or repeated bilateral exchange in small populations (some industrial supply relationships) are not centrally affected. *Second,* it assumes generative AI continues to advance in the relevant signal domain; a domain-specific slowdown would arrest decay there. *Third,* it assumes receivers care about the

Table 8. Future empirical designs and their primary contributions.

#	Research question	Identification	Variables & data	Tests
1	Heterogeneous decay rates?	Cross-signal panel	Receiver-classification accuracy series	P1, P2
2	Capability shocks ⇒ disruption?	Event study around model releases	Transaction outcomes (conversion, dispute)	P1, P2, P9
3	Provenance premium in product markets?	Hedonic pricing	Attested vs. unattested matched products	P4
4	Provenance-elasticity heterogeneity?	Conjoint experiment	Stated WTP for varied attestations	P4, P6
5	Verification-touchpoint lift?	Field experiment	Platform-side randomization of visibility	P14
6	Synthetic-information effects on prices/vols?	Panel DiD with prevalence treatment	Returns, volatilities, news-day effects	P15
7	Standards-consortium value?	Event study around membership	Stock-price reaction; adoption signals	P11, P12
8	Regulatory regime effects?	Cross-jurisdictional DiD	Firm outcomes in regulated vs. unregulated	P10

underlying type the signal conveyed; markets in which receivers are indifferent to type are not centrally affected, because the breakdown of separating equilibria imposes no welfare loss. *Fourth*, the turn unfolds at different rates across regulatory environments: jurisdictions with strong privacy regimes may resist comprehensive attestation on civil-liberties grounds; those with weak rule of law may resist on enforcement grounds; those with strong state capacity may pre-empt private attestation with state infrastructure. The predictions describe structural pressure on markets; the realized institutional response is jurisdictionally heterogeneous.

15.2 Competing explanations

Four competing explanations deserve direct engagement.

The detection-will-suffice view holds that AI content can ultimately be distinguished reliably and that markets will absorb the disruption through detection embedded in evaluation. The view understates a structural asymmetry: detection of generative output is provably harder than generation under broad conditions (Chakraborty et al., 2023; Sadasivan et al., 2023), and the cost of detection failures concentrates on receivers while generation failures are borne by senders. The asymmetry favors generation; the equilibrium it produces is not one in which detection prevails but one in which markets bypass detection through provenance.

The reputation-systems-suffice view holds that existing reputation mechanisms already absorb the disruption. Reputation provides partial coverage, but signal decay strikes precisely at the inputs reputation relies on: reviews are generated, professional bona fides are displayable through AI-mediated channels, brand-trust signals are replicable. Reputation systems are not exempt; they

require provenance infrastructure for their own inputs, recursively (Mechanism 5).

The eventual-equilibration view holds that markets have always absorbed information-environment disruptions and will absorb this one. The historical analogy is correct as far as it goes: cheap printing, photography, and the internet each produced signal-collapse episodes absorbed through new verification institutions (copyright registries, photojournalism standards, web-domain authentication). The provenance turn is the latest such episode and will plausibly resolve through new institutions. But "eventual" does substantial work, and the distinguishing features of Section 1.2—cross-domain simultaneity, symmetric diffusion, recursive exposure—imply that the transition is faster, broader, and less self-correcting than its predecessors. The framework is an account of that transition: what the institutions will look like, which markets reorganize first, and where coordination crises concentrate.

The productivity-only view holds that generative AI is primarily a productivity technology and that productivity gains will swamp coordination problems. The view misreads the level of analysis: productivity gains accrue largely within firms (firm-task-level estimates dominate that literature); coordination problems accrue at the market-institution level. The two are not directly comparable, and a productivity-only view systematically under-attends to the institutional restructuring the present paper concerns.

16. Discussion: Theoretical Implications

The framework has consequences for several established conversations.

16.1 For information economics and signaling theory

The separating-equilibrium tradition requires a complement: a theory of *verification equilibria*, in which coordination is achieved not through differential signal costs but through verification infrastructures with their own cost structures and failure modes. The mathematical apparatus of signaling games is well-developed; an analogous apparatus for verification games—incorporating attestation providers, verification cost, chain-integrity failures, and standards-network effects—is largely absent and represents a significant agenda. The comparative static of Section 3.1 is a first step; a fuller treatment would model receivers, senders, and an attestation provider in a three-player game in which (i) the sender chooses production and attestation effort, (ii) the attestation provider chooses verification effort, and (iii) the receiver chooses an action based on both signal content and attestation. Equilibrium concepts adapted from the screening literature appear suitable; the analytical questions include when receivers invest in verification, when senders self-bind to verifiable processes, when third-party attestation markets emerge, and how attestation-provider competition interacts with signaling equilibria upstream.

16.2 For credence-goods theory

Credence-goods theory (Darby & Karni, 1973; Dulleck & Kerschbamer, 2006) should expand to encompass artifacts whose quality was once observable through signaling content but is no longer. A well-crafted application essay was once a quasi-search good in the Nelson sense—its quality evaluable through inspection. Under signal decay it becomes a credence good in the technical sense, because its content cannot be evaluated without external verification of its origin. The boundaries of the credence-goods category are expanding rapidly, with implications for certification regimes and professional licensing.

16.3 For audience-based and category-based theory

Audience-based theories of markets (Hannan et al., 2007; Hsu, 2006; Zuckerman, 1999) imply that audiences must be augmented with infrastructure for evaluating provenance, and that the relative power of different audience members shifts as that infrastructure develops. Audiences with access to credible provenance signals exercise more determining influence than those without, generating a stratification of audience power prior theory has not addressed. Sophisticated buy-side audiences (institutional investors, professional curators, hiring committees with witnessed-assessment capacity) gain power relative to retail or mass audiences where provenance access is uneven.

16.4 For platform and IS theory

Platform theory (Constantinides et al., 2018; Gawer, 2022; Parker et al., 2016; Van Alstyne et al., 2016) now includes provenance as a first-class platform resource. Platforms that operate or curate verification infrastructure (signing services, identity attestation, witnessed-engagement formats) capture a coordinating role structurally analogous to but partially distinct from the matching and pricing roles platforms have long played. The IS research agenda should include verification governance as a core platform-design problem alongside privacy, security, and content moderation.

16.5 For institutional theory

Institutional theory of legitimacy (DiMaggio & Powell, 1983; Scott, 2014; Suchman, 1995) should incorporate provenance attestation as a constitutive component of legitimacy in AI-mediated environments. Organizations will face pressure to bind claims to verifiable origin; the institutional infrastructure of legitimacy will incorporate provenance machinery as a core feature. Future category-creation contests (Rao, 1994) will be partially adjudicated by attestation rather than narrative alone.

16.6 The locus of theoretical action

The unifying point is that the locus of theoretical action shifts from the sender's signaling decision to the verification infrastructure surrounding the sender's act of production. This is a different

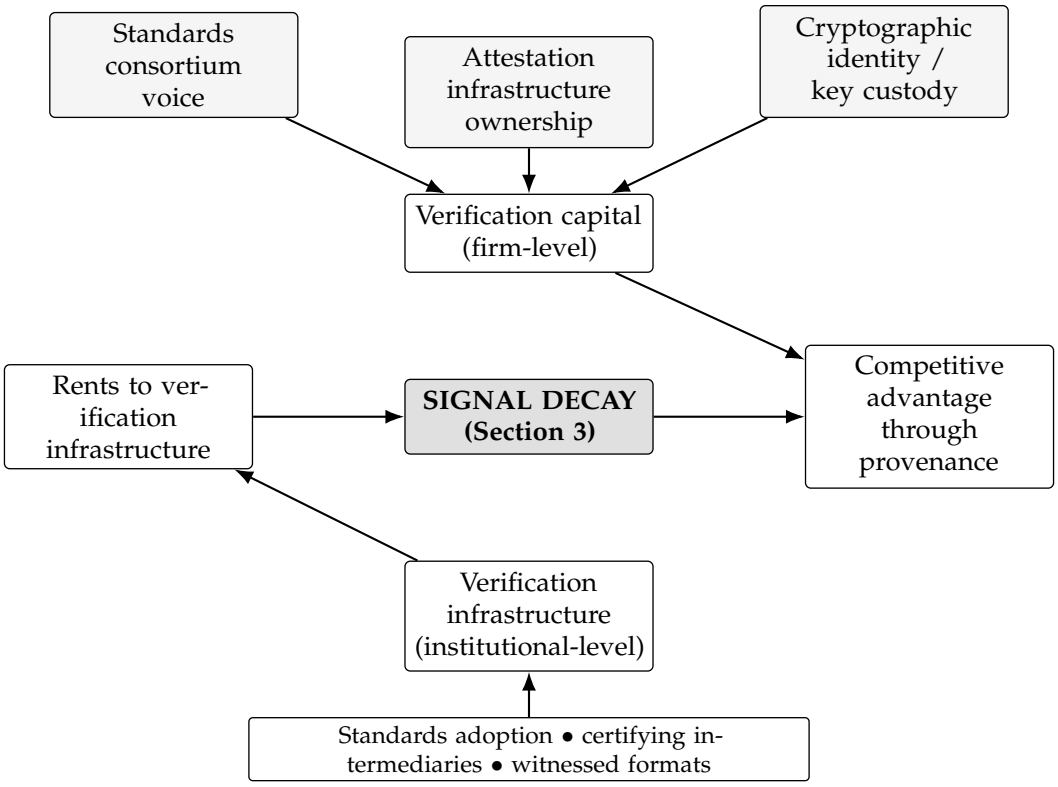


Figure 7. Verification capital and competitive advantage. Signal decay creates demand for verification; institutional-level infrastructure generates rents; firm-level verification capital captures those rents and confers advantage in product, capital, and talent markets.

research object. It calls for different empirical methods (panel data on attestation markets; LLM audit studies; cross-jurisdictional comparison), different policy attention (standards governance; attestation liability; verification antitrust), and different strategic frameworks (verification capital; channel-mix optimization; standards-consortium positioning). Figure 7 summarizes how these elements combine into competitive advantage. The present paper is intended as a charter for that reorientation.

17. Conclusion

Information economics has long assumed that markets coordinate through costly signals. The assumption has been productive: it has organized our understanding of education, advertising, brands, professional credentials, online reputation, capital-market disclosures, and much else. Generative AI is the first general-purpose technology that systematically attacks the cost foundations of signaling across heterogeneous domains at once, and the result is a coordination problem of a kind information economics has not previously had to address as a first-order phenomenon. Markets are not failing—but the architecture by which they succeed is being rebuilt, in front of us.

I argue that the rebuilding is organized around *provenance*: the verifiable causal history of artifacts, claims, and interactions, supported by infrastructures of cryptographic, witnessed, in-

stitutional, and reputational verification. The shift from signal to provenance is not a marginal adjustment; it is a deep reorganization of how strangers establish enough trust to transact. It implies new strategic priorities for firms (verification capital, channel-mix realignment, brand-equity bifurcation), new regulatory frontiers for governments (standards governance, attestation liability, mandatory provenance disclosure, verification-market antitrust), new theoretical objects for scholars (verification equilibria, provenance elasticity, synthetic information risk, recursive saturation), and new sources of advantage for those who position themselves early at the verification layer.

The framework is preliminary, and the empirical record is still forming. The phenomena are unmistakable in some markets and incipient in others, and the institutional response is uneven across jurisdictions and industries. What is not in doubt is that the cost foundation of signaling is degrading, that markets must coordinate through some other logic if they are to keep functioning, and that the logic emerging in the most visible cases is provenance. The research opportunity is to specify, predict, and discipline what comes next—before institutional arrangements solidify around configurations that are neither efficient nor fair. The challenge is interdisciplinary by necessity. I hope the paper invites the marketing, finance, information-systems, economics, sociology, computer-science, and policy communities to engage it as a shared object.

Acknowledgments

I thank colleagues across the marketing, finance, information-systems, sociology, computer-science, and policy communities for conversations that shaped the framework. All remaining errors are my own.

Appendix A. Glossary of Constructs Introduced

Signal decay. Rate at which a signal type loses informational value to receivers under widespread availability of low-cost production technology accessible to all sender types (Section 3).

Provenance. Verifiable causal history connecting an artifact, claim, or interaction to its origin, including producer identity, conditions of production, and chain integrity (Section 5).

Verification capital. Firm-level stock of resources, relationships, and infrastructure enabling credible binding of outputs to verifiable origin (Section 4).

Provenance elasticity. Percentage change in willingness-to-pay per one-percent change in the credibility of provenance attestation (Section 4).

Synthetic information risk. Expected loss from the non-trivial probability that informationally consequential artifacts were AI-fabricated (Section 4).

Signal liquidity. Ease of converting observed signals into reliable inferences about sender type at a given moment (Section 4).

Trust infrastructure density. Per-transaction availability of verification mechanisms in a market (Section 4).

Verification arbitrage. Profit opportunity from cross-market deployment of verification capital (Section 4).

Provenance network effects. Rising value of a provenance standard as more counterparties recognize it (Section 4).

Verification frictions. Transaction costs imposed by verification architectures on production, transmission, and consumption (Section 4).

Signal inflation. Receiver-side response of demanding more or more elaborate signals as indistinguishability rises (Section 4).

Verification externalities. Spillovers of verification-infrastructure investment across markets (Section 4).

Information authenticity premium. Price differential between content of equivalent surface form differing only in origin-attestation credibility (Section 4).

Verification congestion. Excess demand for attestation relative to architectural capacity, raising latency and fees (Section 4).

AI-induced coordination shock. Discrete episode in which a generative-capability advance degrades a previously functional signal class (Section 4).

Appendix B. Worked Industry Examples

The framework is illustrated with brief analytic notes on industries spanning the saturation–tractability matrix; Table 4 codes them onto the two axes.

Higher education. Admissions essays are a paradigmatic high-saturation, moderate-tractability market. Elite universities have responded with witnessed assessment (proctored writing, video interviews) and institutional attestation (school-counselor narratives). The turn is well underway; the open question is whether mass-market institutions, lacking the verification capital of elite peers, can sustain the response.

Consulting and professional services. High saturation, moderate tractability. Major firms are layering witnessed delivery (in-person workshops, client-side observation) and institutional attestation (signed reports under partner accountability) over AI-augmented production. Verification capital is increasingly central to competitive positioning.

Recruitment and hiring. A canonical high-saturation, low-tractability market. Mass-market funnels have seen rapid decay in resumes and cover letters. Responses include skills-assessment platforms (witnessed), referrals (reputational), and proctored video interviews. Channel migration from text to live formats is accelerating.

Dating platforms. High saturation, low tractability par excellence. Platforms have responded with live-photo verification (cryptographic + witnessed), short-form video, and identity attestation through government-ID linkage. Verification capital sits with platforms that bind profiles to verifiable identities.

Creator economy. High saturation with bifurcating tractability. Authenticated creators (KYC, signed content, livestreaming) capture rising shares; synthetic and unauthenticated creators commodify. Platforms are the focal verification governors.

Financial markets and equity research. High saturation pressure but high inherited tractability (audit, regulatory disclosure, professional licensing). Incumbents are layering attested-AI workflows over legacy infrastructure; finance is among the markets best positioned to reconstitute around provenance.

News media and journalism. High saturation, moderate tractability (institutional reputation, cryptographic standards). C2PA adoption among premium outlets illustrates the early response. Local journalism faces the more acute challenge because its inherited institutional verification is thinner.

Healthcare. High tractability (regulatory regime, professional licensure) with growing saturation pressure in non-clinical content (patient communications, marketing, claims documentation). Verification capital is concentrated in the regulatory–institutional layer; new tensions appear at the AI-clinical-documentation interface.

Scientific publishing. High saturation, moderate-to-high tractability. Pre-registration, registered reports, open data, and contributor-credit standards have emerged as cryptographic and institutional responses. Peer review itself faces recursive saturation.

Foundation-model providers. A central node in the verification ecosystem. Providers' positioning on watermarking, content authentication, and attestation interoperability shapes the verification possibilities of downstream markets and is a first-order determinant of the trajectory.

Tooling and creative production. Software vendors are integrating provenance standards (e.g., Content Credentials) at the production layer. Tools that bake attestation into authoring become focal points of verification capital.

Social and content platforms. A heterogeneous group whose verification investments will substantially determine which large segments of the digital economy reconstitute and which remain in coordination crisis.

References

Aaker, D. A. (1996). *Building strong brands*. Free Press.

Aaronson, S. (2023). *Watermarking of large language models* [Lecture]. Simons Institute for the Theory of Computing, University of California, Berkeley.

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics, 84*(3), 488–500. https://doi.org/10.2307/1879431

Allen, F., & Faulhaber, G. R. (1989). Signalling by underpricing in the IPO market. *Journal of Financial Economics, 23*(2), 303–323. https://doi.org/10.1016/0304-405X(89)90060-3

Bateman, J. (2020). *Deepfakes and synthetic media in the financial system: Assessing threat scenarios*. Carnegie Endowment for International Peace.

Benbya, H., Pachidi, S., & Jarvenpaa, S. L. (2021). Special issue editorial: Artificial intelligence in organizations: Implications for information systems research. *Journal of the Association for Information Systems, 22*(2), 281–303. https://doi.org/10.17705/1jais.00662

Berente, N., Gu, B., Recker, J., & Santhanam, R. (2021). Managing artificial intelligence. *MIS Quarterly, 45*(3), 1433–1450. https://doi.org/10.25300/MISQ/2021/16274

Berger, J., Humphreys, A., Ludwig, S., Moe, W. W., Netzer, O., & Schweidel, D. A. (2020). Uniting the tribes: Using text for marketing insight. *Journal of Marketing, 84*(1), 1–25. https://doi.org/10.1177/0022242919873106

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics, 140*(2), 889–942. https://doi.org/10.1093/qje/qjae044

Buneman, P., Khanna, S., & Tan, W.-C. (2008). Why and where: A characterization of data provenance. In J. Van den Bussche & V. Vianu (Eds.), *Database theory—ICDT 2001* (pp. 316–330). Springer. https://doi.org/10.1007/3-540-44503-X_20

Carroll, G. R. (2015). Authenticity: Attribution, value, and meaning. In R. A. Scott & S. M. Kosslyn (Eds.), *Emerging trends in the social and behavioral sciences* (pp. 1–13). Wiley. https://doi.org/10.1002/9781118900772.etrds0026

Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research, 56*(5), 809–825. https://doi.org/10.1177/0022243719851788

Catalini, C., & Gans, J. S. (2020). Some simple economics of the blockchain. *Communications of the ACM, 63*(7), 80–90. https://doi.org/10.1145/3359552

Chakraborty, S., Bedi, A. S., Zhu, S., An, B., Manocha, D., & Huang, F. (2023). *On the possibilities of AI-generated text detection* (arXiv:2304.04736). arXiv. https://doi.org/10.48550/arXiv.2304.04736

Christ, M., Gunn, S., & Zamir, O. (2024). Undetectable watermarks for language models. In *Proceedings of the 37th Annual Conference on Learning Theory (COLT)* (pp. 1125–1139). PMLR.

Cong, L. W., Tang, K., Wang, J., & Zhang, Y. (2021). *AlphaPortfolio: Direct construction through deep reinforcement learning and interpretable AI* [Working paper]. SSRN. https://doi.org/10.2139/ssrn.3554486

Constantinides, P., Henfridsson, O., & Parker, G. G. (2018). Introduction—Platforms and infrastructures in the digital age. *Information Systems Research, 29*(2), 381–400. https://doi.org/10.1287/isre.2018.0794

Darby, M. R., & Karni, E. (1973). Free competition and the optimal amount of fraud. *The Journal of Law and Economics, 16*(1), 67–88. https://doi.org/10.1086/466756

Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science, 48*(1), 24–42. https://doi.org/10.1007/s11747-019-00696-0

Dechow, P., Ge, W., & Schrand, C. (2010). Understanding earnings quality: A review of the proxies, their determinants and their consequences. *Journal of Accounting and Economics, 50*(2–3), 344–401. https://doi.org/10.1016/j.jacceco.2010.09.001

Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science, 49*(10), 1407–1424. https://doi.org/10.1287/mnsc.49.10.1407.17308

Diamond, D. W. (1984). Financial intermediation and delegated monitoring. *The Review of Economic Studies, 51*(3), 393–414. https://doi.org/10.2307/2297430

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General, 144*(1), 114–126. https://doi.org/10.1037/xge0000033

DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review, 48*(2), 147–160. https://doi.org/10.2307/2095101

Dulleck, U., & Kerschbamer, R. (2006). On doctors, mechanics, and computer specialists: The economics of credence goods. *Journal of Economic Literature, 44*(1), 5–42. https://doi.org/10.1257/002205106776162717

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). *GPTs are GPTs: An early look at the labor market impact potential of large language models* (arXiv:2303.10130). arXiv. https://doi.org/10.48550/arXiv.2303.10130

Epstein, Z., Hertzmann, A., & the Investigators of Human Creativity. (2023). Art and the science of generative AI. *Science, 380*(6650), 1110–1111. https://doi.org/10.1126/science.adh4451

Erdem, T., & Swait, J. (2004). Brand credibility, brand consideration, and choice. *Journal of Consumer Research, 31*(1), 191–198. https://doi.org/10.1086/383434

Gawer, A. (2022). Digital platforms and ecosystems: Remarks on the dominant organizational forms of the digital age. *Innovation: Organization & Management, 24*(1), 110–124. https://doi.org/10.1080/14479338.2021.1965888

Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals, 14*(2), 627–660. https://doi.org/10.5465/annals.2018.0057

Goldstein, I., Spatt, C. S., & Ye, M. (2021). Big data in finance. *The Review of Financial Studies, 34*(7), 3213–3225. https://doi.org/10.1093/rfs/hhab038

Grossman, S. J. (1981). The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics, 24*(3), 461–483. https://doi.org/10.1086/466995

Gurun, U. G., & Butler, A. W. (2012). Don't believe the hype: Local media slant, local advertising, and firm value. *The Journal of Finance, 67*(2), 561–598. https://doi.org/10.1111/j.1540-6261.2012.01725.x

Halaburda, H., Haeringer, G., Gans, J., & Gandal, N. (2022). The microeconomics of cryptocurrencies. *Journal of Economic Literature, 60*(3), 971–1013. https://doi.org/10.1257/jel.20201593

Hannan, M. T., Pólos, L., & Carroll, G. R. (2007). *Logics of organization theory: Audiences, codes, and ecologies*. Princeton University Press.

Harvey, C. R., Ramachandran, A., & Santoro, J. (2021). *DeFi and the future of finance*. Wiley.

Healy, P. M., & Palepu, K. G. (2001). Information asymmetry, corporate disclosure, and the capital markets:

A review of the empirical disclosure literature. *Journal of Accounting and Economics, 31*(1–3), 405–440. https://doi.org/10.1016/S0165-4101(01)00018-0

Hofstetter, R., & Gollnhofer, J. F. (2024). The creator's dilemma: Resolving tensions between authenticity and monetization in social media. *International Journal of Research in Marketing, 41*(3), 427–435. https://doi.org/10.1016/j.ijresmar.2024.06.005

Hsu, G. (2006). Jacks of all trades and masters of none: Audiences' reactions to spanning genres in feature film production. *Administrative Science Quarterly, 51*(3), 420–450. https://doi.org/10.2189/asqu.51.3.420

Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science, 49*(1), 30–50. https://doi.org/10.1007/s11747-020-00749-9

Jakesch, M., Hancock, J. T., & Naaman, M. (2023). Human heuristics for AI-generated language are flawed. *Proceedings of the National Academy of Sciences, 120*(11), e2208839120. https://doi.org/10.1073/pnas.2208839120

Kannan, P. K., & Li, H. A. (2017). Digital marketing: A framework, review and research agenda. *International Journal of Research in Marketing, 34*(1), 22–45. https://doi.org/10.1016/j.ijresmar.2016.11.006

Keller, K. L. (2003). *Strategic brand management: Building, measuring, and managing brand equity* (2nd ed.). Prentice Hall.

Kirchenbauer, J., Geiping, J., Wen, Y., Katz, J., Miers, I., & Goldstein, T. (2023). A watermark for large language models. In *Proceedings of the 40th International Conference on Machine Learning (ICML)* (pp. 17061–17084). PMLR.

Kirmani, A., & Rao, A. R. (2000). No pain, no gain: A critical review of the literature on signaling unobservable product quality. *Journal of Marketing, 64*(2), 66–79. https://doi.org/10.1509/jmkg.64.2.66.18000

Köbis, N., & Mossink, L. D. (2021). Artificial intelligence versus Maya Angelou: Experimental evidence that people cannot differentiate AI-generated from human-written poetry. *Computers in Human Behavior, 114*, 106553. https://doi.org/10.1016/j.chb.2020.106553

Korinek, A., & Stiglitz, J. E. (2021). *Artificial intelligence, globalization, and strategies for economic development* (NBER Working Paper No. 28453). National Bureau of Economic Research. https://doi.org/10.3386/w28453

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica, 53*(6), 1315–1335. https://doi.org/10.2307/1913210

Lamberton, C., & Stephen, A. T. (2016). A thematic exploration of digital, social media, and mobile marketing: Research evolution from 2000 to 2015 and an agenda for future inquiry. *Journal of Marketing, 80*(6), 146–172. https://doi.org/10.1509/jm.15.0415

Lee, M. K., & Rich, K. (2021). Who is included in human perceptions of AI? Trust and perceived fairness around healthcare AI and cultural mistrust. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems* (Article 138). ACM. https://doi.org/10.1145/3411764.3445570

Leland, H. E., & Pyle, D. H. (1977). Informational asymmetries, financial structure, and financial intermediation. *The Journal of Finance, 32*(2), 371–387. https://doi.org/10.2307/2326770

Lemon, K. N., & Verhoef, P. C. (2016). Understanding customer experience throughout the customer journey. *Journal of Marketing, 80*(6), 69–96. https://doi.org/10.1509/jm.15.0420

Leung, F. F., Gu, F. F., & Palmatier, R. W. (2022). Online influencer marketing. *Journal of the Academy of Marketing Science, 50*(2), 226–251. https://doi.org/10.1007/s11747-021-00829-4

Ljungqvist, A. (2007). IPO underpricing. In B. E. Eckbo (Ed.), *Handbook of corporate finance: Empirical corporate finance* (Vol. 1, pp. 375–422). North-Holland. https://doi.org/10.1016/B978-0-444-53265-7.50021-4

Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes, 151*, 90–103. https://doi.org/10.1016/j.obhdp.2018.12.005

Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research, 46*(4), 629–650. https://doi.org/10.1093/jcr/ucz013

Longoni, C., & Cian, L. (2022). Artificial intelligence in utilitarian vs. hedonic contexts: The "word-of-machine" effect. *Journal of Marketing, 86*(1), 91–108. https://doi.org/10.1177/0022242920957347

Longpre, S., Mahari, R., Chen, A., Obeng-Marnu, N., Sileo, D., Brannon, W., Muennighoff, N., Khazam, N., Kabbara, J., Perisetla, K., Wu, X., Shippole, E., Bollacker, K., Wu, T., Villa, L., Pentland, S., & Hooker, S. (2024). A large-scale audit of dataset licensing and attribution in AI. *Nature Machine Intelligence, 6*, 975–987. https://doi.org/10.1038/s42256-024-00878-8

Lopez-Lira, A., & Tang, Y. (2023). *Can ChatGPT forecast stock price movements? Return predictability and large language models* (arXiv:2304.07619). arXiv. https://doi.org/10.48550/arXiv.2304.07619

Milgrom, P., & Roberts, J. (1986). Price and advertising signals of product quality. *Journal of Political Economy, 94*(4), 796–821. https://doi.org/10.1086/261408

Myers, S. C., & Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics, 13*(2), 187–221. https://doi.org/10.1016/0304-405X(84)90023-0

Nelson, P. (1974). Advertising as information. *Journal of Political Economy, 82*(4), 729–754. https://doi.org/10.1086/260231

Netzer, O., Lemaire, A., & Herzenstein, M. (2019). When words sweat: Identifying signals for loan default in the text of loan applications. *Journal of Marketing Research, 56*(6), 960–980. https://doi.org/10.1177/0022243719852959

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science, 381*(6654), 187–192. https://doi.org/10.1126/science.adh2586

O'Hara, M. (2015). High frequency market microstructure. *Journal of Financial Economics, 116*(2), 257–270. https://doi.org/10.1016/j.jfineco.2015.01.003

Parker, G. G., Van Alstyne, M. W., & Choudary, S. P. (2016). *Platform revolution: How networked markets are transforming the economy and how to make them work for you*. W. W. Norton.

Partnoy, F. (2006). How and why credit rating agencies are not like other gatekeepers. In Y. Fuchita & R. E. Litan (Eds.), *Financial gatekeepers: Can they protect investors?* (pp. 59–99). Brookings Institution Press.

Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing, 85*(1), 131–151. https://doi.org/10.1177/0022242920953847

Rai, A. (2020). Explainable AI: From black box to glass box. *Journal of the Academy of Marketing Science, 48*(1), 137–141. https://doi.org/10.1007/s11747-019-00710-5

Rao, H. (1994). The social construction of reputation: Certification contests, legitimation, and the survival of organizations in the American automobile industry, 1895–1912. *Strategic Management Journal, 15*(S1), 29–44. https://doi.org/10.1002/smj.4250150904

Rao, H. (2009). *Market rebels: How activists make or break radical innovations*. Princeton University Press.

Resnick, P., & Zeckhauser, R. (2002). Trust among strangers in internet transactions: Empirical analysis of eBay's reputation system. In M. R. Baye (Ed.), *The economics of the internet and e-commerce* (pp. 127–157). Emerald. https://doi.org/10.1016/S0278-0984(02)11030-3

Riley, J. G. (2001). Silver signals: Twenty-five years of screening and signaling. *Journal of Economic Literature, 39*(2), 432–478. https://doi.org/10.1257/jel.39.2.432

Rochet, J.-C., & Tirole, J. (2003). Platform competition in two-sided markets. *Journal of the European Economic Association, 1*(4), 990–1029. https://doi.org/10.1162/154247603322493212

Rosenthol, L. (2022). C2PA: The world's first industry standard for content provenance [Conference paper]. *Applications of Digital Image Processing XLV, 12226*. SPIE. https://doi.org/10.1117/12.2632021

Rust, R. T. (2020). The future of marketing. *International Journal of Research in Marketing, 37*(1), 15–26. https://doi.org/10.1016/j.ijresmar.2019.08.002

Sadasivan, V. S., Kumar, A., Balasubramanian, S., Wang, W., & Feizi, S. (2023). *Can AI-generated text be reliably detected?* (arXiv:2303.11156). arXiv. https://doi.org/10.48550/arXiv.2303.11156

Schneider, J., Abraham, R., Meske, C., & vom Brocke, J. (2023). Artificial intelligence governance for businesses. *Information Systems Management, 40*(3), 229–249. https://doi.org/10.1080/10580530.2022.2085825

Scott, W. R. (2014). *Institutions and organizations: Ideas, interests, and identities* (4th ed.). Sage.

Spence, M. (1973). Job market signaling. *The Quarterly Journal of Economics, 87*(3), 355–374. https://doi.org/10.2307/1882010

Stiglitz, J. E. (1975). The theory of "screening," education, and the distribution of income. *The American Economic Review, 65*(3), 283–300.

Suchman, M. C. (1995). Managing legitimacy: Strategic and institutional approaches. *Academy of Management Review, 20*(3), 571–610. https://doi.org/10.5465/amr.1995.9508080331

Tadelis, S. (2016). Reputation and feedback systems in online platform markets. *Annual Review of Economics, 8*, 321–340. https://doi.org/10.1146/annurev-economics-080315-015325

Tian, F. (2017). A supply chain traceability system for food safety based on HACCP, blockchain & Internet of Things. In *2017 International Conference on Service Systems and Service Management*. IEEE. https://doi.org/10.1109/ICSSSM.2017.7996119

Tiwana, A., Konsynski, B., & Bush, A. A. (2010). Platform evolution: Coevolution of platform architecture, governance, and environmental dynamics. *Information Systems Research, 21*(4), 675–687. https://doi.org/10.1287/isre.1100.0323

Van Alstyne, M. W., Parker, G. G., & Choudary, S. P. (2016). Pipelines, platforms, and the new rules of strategy. *Harvard Business Review, 94*(4), 54–62.

Verrecchia, R. E. (2001). Essays on disclosure. *Journal of Accounting and Economics, 32*(1–3), 97–180. https://doi.org/10.1016/S0165-4101(01)00025-8

Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., Foltýnek, T., Guerrero-Dib, J., Popoola, O., Šigut, P., & Waddington, L. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity, 19*, 26. https://doi.org/10.1007/s40979-023-00146-z

Welch, I. (1989). Seasoned offerings, imitation costs, and the underpricing of initial public offerings. *The Journal of Finance, 44*(2), 421–449. https://doi.org/10.1111/j.1540-6261.1989.tb05064.x

Wernerfelt, B. (1988). Umbrella branding as a signal of new product quality: An example of signalling by posting a bond. *The RAND Journal of Economics, 19*(3), 458–466. https://doi.org/10.2307/2555667

Yermack, D. (2017). Corporate governance and blockchains. *Review of Finance, 21*(1), 7–31. https://doi.org/10.1093/rof/rfw074

Zuckerman, E. W. (1999). The categorical imperative: Securities analysts and the illegitimacy discount. *American Journal of Sociology, 104*(5), 1398–1438. https://doi.org/10.1086/210178