Market Design for AI: Beyond the Copyright Binary

Yan Dai* Maryam Farboodi† Negin Golrezaei‡ Sepehr Shahshahani§

March 26, 2026

Abstract

How can we design a market of human-generated content for use in training AI models that both enables technological progress and preserves individual incentives for high-quality content creation? Existing approaches take polar positions: a "free-for-all" model based on fair use and a "strong intellectual property rights" model. We show that both fail. Strong IP rights underpower creative incentives, especially for more innovative creators, a phenomenon we term the "originality penalty." Dynamically, this market failure undermines AI model performance, even for a good model. A good model induces greater reliance by humans on AI-generated content, resulting in homogenized content feeling back into training data and degrading model performance—the "curse of precision." We propose a market design with a data intermediary that internalizes cross-creator externalities and subsidizes innovative contributions, thereby restoring efficiency.

*MIT Operations Research Center. yandai20@mit.edu.
†MIT Sloan School of Management, NBER and CEPR. farboodi@mit.edu.
‡MIT Sloan School of Management. golrezae@mit.edu.
§Washington University School of Law. sepehrs@wustl.edu.

1 Introduction

Recent breakthroughs in generative artificial intelligence (AI) are laden with both promise and peril. By enabling the fast, cheap delivery of useful content, AI models have improved work and life for millions of people. But AI models' ability to generate this useful content depends on a rich reservoir of human-created content on which to train. To date, AI firms have acquired this content by scraping it from online sources—largely without human creators' consent, compensation, or credit. While this practice has enabled rapid AI progress, it undermines human creators' incentives to create. We are already seeing warning signs: Stack Overflow's traffic declined significantly after the rise of Large Language Models (LLMs) like ChatGPT, as users consumed AI answers rather than contributing to producing them. And various artists and other content creators have sued AI firms contending that their practices have deprived them of the rewards of their creative work. These trends point to a misalignment between the incentives of AI developers and human creators, and they give rise to the central research question of this project: How can we design a market of human-generated content for AI that both enables technological progress and preserves individuals' incentive to create high-quality content?

This is among the foremost policy challenges of our times, implicating in its narrow berth important questions of innovation and technology policy and, more broadly, fundamental concerns about how to preserve human creative flourishing. But the tools we have for answering the question come from a decades-old Copyright Act that was not designed with anything like AI technology in mind. As such, the current debate is trapped in a legal binary: AI harvesting is either fair use (what we call a "free for all" system) or copyright infringement (what we call a "strict property rights" regime). This paper argues that both extremes are destined to fail and points instead to a sustainable "third way" via data intermediaries.

The reason a free-for-all system fails is easy to grasp: If AI can freely take human-created content and use it to generate content that largely replaces consumers' need for the human-created content, then humans will be left with much-diminished incentives to invest in high-quality content creation. But we show that the opposite solution—strict property rights—is also doomed to fail. The reasons for this failure are subtle and unexpected. Specifically, we identify *two short-run market failures*

and a long-run market failure in a strict property rights system. First, contrary to conventional wisdom about the effects of intellectual property (IP) rights, we show that *strong individualistic IP rights result in under-powered creative incentives.* This is due to the market power of AI firms and correlation among different human creators' content. Market power allows an AI firm to profitably price content below what would be required to induce socially optimal effort on the part of human creators; content correlation implies that, by selling her own content to the AI firm, a human creator is also effectively selling "part of" other creators' content, a negative externality in the form of information leakage that further depresses the price of content. Second, we show that AI's statistical preference for "representative" data creates an *originality penalty*, meaning the market degrades the incentives of innovative content creators to a greater extent (relative to social optimum) than it degrades the incentives of run-of-the-mill creators who are highly influenced by common trends and biases. The originality penalty thus leads to the production of homogenized content. Although it is predictable that AI firms' market power may lead to under-investment by human creators, our finding of an originality penalty contradicts the economic intuition that scarcity demands a premium, revealing a market failure that is distinctive to the AI context.

Beyond the tradeoff between human creative incentives and technological progress in the short term, we extend our analysis to a continuous-time dynamic setting in which an AI firm trains and refines its models on human-generated content over time, and humans in turn are able to use AI to help them generate content. In this dynamic setting, we identify a *curse of precision* whereby a good model induces greater reliance by human content creators on AI outputs, leading to a homogenization of human content, which feed back into the AI training pipeline, degrading model performance. The upshot is that a good model sows the seeds of its own decay. This "curse of precision" is similar to the phenomenon of "model collapse" identified in the machine learning literature (Shumailov et al., 2024): when trained on self-generated data, an AI model's capability collapses. Our analysis, however, reveals a different pathway to failure and a more pessimistic prediction: Even if the AI firm acquires *fresh data* from human creators, poorly designed markets can lead to AI model failure.

Having identified the static and dynamic properties of market equilibrium under existing institutional designs, we then move beyond diagnosis to treatment by proposing a better market design.

Our proposed solution takes the form of a *data intermediary* who negotiates on behalf of individual content owners with an AI firm and apportions royalties among them. We show that, given full information, the data intermediary can overcome both market failures that characterize the individual-IP-rights regime. By managing creators as a portfolio and negotiating collectively on their behalf, the intermediary both overcomes the AI firm's market power and internalizes the negative externality arising from data correlation and information leakage. And by apportioning the payment received from the AI firm among individual creators according to their respective contributions to social surplus (using weights that we define inspired by the Aumann-Shapley value), the intermediary reverses the originality penalty that characterizes unintermediated interactions between AI firms and human content creators. We show, interestingly, that efficiency can be restored using a simple *two-part tariff* mechanism where both the lump-sum transfer that the intermediary receives from the AI firm and the amounts that the intermediary pays to individual content creators are affine functions of the effort exerted by content creators.

Our intermediary solution abstracts away from informational asymmetries and, more generally, from principal-agent problems between the intermediary and human creators. As such, it should not be viewed as a definitive policy solution. The idea is nevertheless quite useful as a first step toward an optimal market design because any problems that characterize the full-information context will also plague an incomplete-information environment. The market failures we identify will exist even in the absence of information gaps or principal-agent costs, and they are the very failures addressed by our proposed market design, so the solution is well-calibrated to the problem as defined and serves as a useful benchmark for a definitive solution.

The paper is organized as follows. The next section explains the legal-policy background and situates our contribution in the literature. Section 3 develops and solves a static model that uncovers the shortcomings of the existing legal regime and commonly proposed solutions. Section 4 extends the model to a continuous-time dynamic case, identifying a long-run market failure. Building on these results, Section 5 proposes an institutional design that overcomes these shortcomings and outperforms commonly proposed solutions. The Appendix furnishes formal statements and proofs.

2 Policy Debate and Literature

The question of how best to resolve AI's disruption of the market for creating and consuming informational content is hotly debated not just in academic literature but also in courts. The most consequential legal questions affecting the future of AI have been raised in several lawsuits filed by content owners against AI firms alleging copyright infringement.[1] To a layperson, the juxtaposition of copyright with the cutting edge of technology may appear unusual—after all, copyright is thought to be about artistic creativity, whereas the high-technology context might seem better suited to patent law or other regulatory regimes—but in fact copyright law has long been at the forefront of regulating new technologies in the United States. From the piano roll to cable to VCRs to peer-to-peer file sharing to online video streaming and beyond, the first steps in deciding the fate of emerging technologies have often been taken by courts in copyright cases (Shahshahani, 2018; Samuelson, 2024). To appreciate the reasons for this, it is essential to review the fundamental structure of American IP laws, of which copyright is a central part.

The goal of American IP laws, in the words of the U.S. Constitution, is "to promote the progress of science and useful arts, by securing for limited times to authors and inventors the exclusive right to their respective writings and discoveries" (U.S. Const. art. I, § 8, cl. 8). The reason such an "exclusive right"—in other words, a temporary monopoly—has been considered necessary to achieve the objective of progress in arts and science traces to the "public good" nature of intangible products of the mind: Unlike tangible goods (say, a chair), intangible products are "nonrivalrous"—one person's use does not diminish another person's use—and "nonexclusive"—the use of an idea once disclosed is hard to limit to the first recipient. As such, the theory of American IP laws is that exclusive rights are necessary to prevent freeriding on others' creativity and thus maintain incentives to create. But the theory also acknowledges that these exclusive rights come at a cost: Like any monopoly, IP rights raise the cost of a work protected by IP and make it more difficult for users to access. These include not only costs to end users, such as a person who wants to read a copyrighted novel or take a patented medicine, but also—of special importance in the AI context—

[1] A useful list appears at https://chatgptiseatingtheworld.com/2024/08/27/master-list-of-lawsuits-v-ai-chatgpt-openai-microsoft-meta-midjourney-other-ai-cos/. For a non-lawyer-friendly introduction, see Samuelson (2023) and Samuelson (2025).

costs to follow-on creators who want to use existing content as building blocks for creating new content (Freilich and Shahshahani, 2023; Shahshahani, 2025). Balancing creative incentives against access costs—the *incentives-access tradeoff*—is thus at the core of American intellectual property law. This has been the dominant conceptual framework for understanding IP in economics for a long time (e.g., Plant, 1934; Machlup, 1958; Arrow, 1962; Nordhaus, 1969) and in US law for an even longer time (e.g., *Wheaton v. Peters*, 33 U.S. 591, 657–58, 661 (1834); *Kendall v. Winsor*, 62 U.S. 322, 327–29 (1858)); *Sears, Roebuck & Co. v. Stiffel Co.*, 376 U.S. 225, 229–31 (1964); *Twentieth Century Music Corp. v. Aiken*, 422 U.S. 151, 156 (1975)).

The incentives-versus-access paradigm helps explain the current array of scholarly opinion on the question of whether the unauthorized use of copyrighted content in training AI models should be permitted.[2] On one side, scholars have argued that AI firms' unrestricted use of copyrighted works in training AI models undermines individuals' incentives to create high-quality content, so robust IP rights are needed to preserve creative incentives and promote the constitutional goal of artistic and scientific progress (Opderbeck, 2024; Barnett, 2026). We call this position the "full property rights" model. On the other side, scholars have argued that unrestricted access to data is essential to innovation, concluding that the use of copyrighted works in AI training constitutes "fair use"[3] and is not a copyright infringement (Sag, 2019; Lemley and Casey, 2021; Lee, 2025). We call this the "free for all" model.

Our first insight is that these polar ways of thinking about the problem fall short. It is easy to see why the free-for-all model will not work: If content can be taken without rewarding its creator,

[2]That is not the only copyright-law question implicated by AI. Grossly simplified, there are least three major legal questions: (1) Does the use of copyrighted material in an AI system's training set infringe the copyright in the material? (2) Does the output of an AI model infringe the copyright in works in its training set or other works? In this connection, does the substantial-similarity standard differ from the standard that is applied when considering non-AI cases of copyright infringement? (3) Is output produced by someone using AI copyrightable? If so, who is the author (initial owner)? It is the first of these questions that is most directly applicable to this paper, though the second question also potentially comes into play. For analysis of these and other questions related to copyright and AI, see Henderson et al. (2023); Lemley (2024); Lemley and Ouellette (2025).

[3]Fair use is a longstanding judge-made doctrine, now codified in the Copyright Act, 17 U.S.C. § 107, which holds that certain uses of copyrighted works, though otherwise impinging on the exclusive rights of the copyright owner, qualify as noninfringing in light of the purpose and character of the use, the nature of the copyrighted work, the amount and substantiality of the material used, and the effect of the use on the market for the original work. The two main theoretical justifications for the fair use doctrine are that it cures market failure in situations where the copyright owner is unlikely to license a valuable use (Gordon, 1982) and that it encourages "transformative" uses whereby a second-generation creator uses a copyrighted work as raw material in the creation of new information, insights, or aesthetics (Leval, 1990). While these academic justifications are useful and influential, the doctrine is highly fact-dependent and case results cannot be confidently predicted based only on academic theory.

people will have little incentive to invest resources into creating content, so the amount and quality of human-created content will decrease. That is the very freerider problem that American copyright and patent laws are designed to ameliorate. The problem will be especially pronounced as consumers increasingly turn to AI platforms, rather than the original source, to get the content they want.

Admittedly, this analysis does not do full justice to the concept of fair use, which is flexible and capable of adaptation to a variety of factual circumstances. The doctrine does not require that *all* uses of copyrighted work for training AI models be deemed noninfringing (though some commentators have advocated just that); depending on the circumstances, some uses may be shielded by fair use while others are held to be copyright infringement. For example, in a case involving Anthropic, the firm that developed the large language model Claude, a court recently held that the unauthorized copying and use of copyrighted books to train the model qualified as fair use with respect to books that Anthropic purchased but not with respect to books that it downloaded for free from "pirate" websites (*Bartz v. Anthropic PBC*, 787 F. Supp. 3d 1007, 1033 (N.D. Cal. 2025)). Whether particular AI uses qualify as fair use could thus be determined case by case in light of particular factual circumstances, as the recent Copyright Office Report on artificial intelligence predicts (Register of Copyrights, 2025, 74). While such flexibility can be a virtue, it also breeds uncertainty and unpredictability, a longstanding criticism of fair use doctrine (see Shahshahani (2015) for a critical survey). For example, on facts very similar to the *Anthropic* case, a different court held that Meta's use of copyrighted books downloaded from pirate websites to train the Llama models *was* fair use, but took pains to point out that the decision was based on the record before it and could well have gone the other way if the plaintiffs had developed a better factual record (*Kadrey v. Meta Platforms, Inc.*, 788 F. Supp. 3d 1026, 1036–37 (N.D. Cal. 2025)). But if the legality of use will depend on each case's particular facts and circumstances, or on the judge presiding over the case, then both AI developers and individual creators are left without much *ex ante* guidance about whether AI firms can freely take human creators' content. So, whereas a blanket fair use privilege results in under-powered creative incentives, a highly fact-specific doctrine leads to uncertainty and potential chilling effects, particularly for risk-averse tech firms and content creators.

What about the full property rights model? Proponents of this model rely on the Coase Theorem (Coase, 1960) to argue that if property rights are clearly assigned to creators, then creators and AI

firms will negotiate a price and efficient data transfers or licenses will occur. The standard counter-argument in the literature is that *transaction costs* of different flavors—the logistical difficulty of an AI company negotiating with millions of individual creators, overlapping or unclear ownership of IP rights, cultural and institutional barriers to mutual understanding, divergent estimates of value, and suchlike—may stand in the way of an efficient bargain (Ostrom, 1990; Heller, 1998; Heller and Eisenberg, 1998).

That is *not* why the full property rights model fails in our theoretical framework. We show, more pessimistically, that the system will fail even in the absence of garden-variety transaction costs—that is, even if the parties reach an agreement to sell or license human-created content to AI firms. The reason behind this failure is twofold: the *market power* of AI firms and *correlation* between different human creators' content. A monopsonist (or oligopsonist) AI firm reduces prices paid to human content creators, sacrificing the quality of content it buys in favor of higher (per-unit-of-quality) profit when it sells its AI-generated content to consumers. Though this lower price is profitable for the AI firm, it reduces human incentives to create high-quality content below what is socially optimal. Moreover, because each individual's content is similar to other individuals' content in that domain, an individual selling her own content is also effectively selling "part of" other creators' content to the AI firm, so only part of the content is valuable to the AI firm, though the creator bears the cost of producing the full content (not just of its marginal value). The price the AI firm is willing to pay is thus below what many individuals will find worthwhile to invest in creating content. So the property rights model breaks down not because it fails the technological-progress objective but, counterintuitively, because it fails the creative-incentives objective. This finding inverts common wisdom about the effect of strong intellectual property rights.

In deriving this result, we draw on recent work in information economics, particularly Acemoglu et al. (2022) and Bergemann et al. (2022), which demonstrate in the privacy context that when individual data are correlated, selling one's data generates a negative externality by revealing information about others, leading to *over-supply* ("excessive data sharing") relative to social optimum. We extend this literature in several ways. Whereas Acemoglu et al. (2022) treats data as an *endowment*, we shift attention to the *production* of creative works, where correlation implies an asymmetric substitutability. This richer analysis allows us to investigate the differential effects of

AI on human creative works with different levels of novelty or originality, demonstrating the *uneven* under-investment ("originality penalty") discussed above.

Further, we move beyond the static context of existing models to explore the dynamic effects of AI interaction with potential content creators. The central result in our dynamic analysis is the "curse of precision" discussed above whereby a good model sows the seeds of its own decay. This finding bridges two contrasting phenomena identified in machine learning: while scaling up AI models unlocks "emergent abilities" (Wei et al., 2022; Schaeffer et al., 2023; Berti et al., 2025), recursive training on self-generated data triggers "model collapse" (Shumailov et al., 2024; Alemohammad et al., 2023). We interpret and reconcile these phenomena through an economic lens: stronger AI capabilities act as a substitute for expensive human creativity, which creates a dynamic "market for lemons" (Akerlof, 1970; Tullis, 2025) that *disincentivizes* the production of original content. Consequently, our analysis suggests that "model collapse" can happen—even with consistent fresh human inputs—as an economic consequence of the model's own success. Hence the curse of precision.

Our analysis of market failure in the presence of IP also draws on and contributes to an extensive economics literature on property rights. Building on the seminal contributions of Coase (1960) and Williamson (1979), this literature has largely focused on transaction costs, and options-to-own have been proposed as potential solutions (Che, 2006; Segal and Whinston, 2016). We are interested in the efficient allocation of property rights in data, so we tailor the model to analyze data as productive assets that can be traded in a market. Sharing this view, a parallel line of macro-finance research treats data as a *nonrival byproduct* of economic activities, which leads to expansion and superstar firms (Farboodi and Veldkamp, 2020, 2021, 2023). By contrast, we revisit the classical insight of Arrow (1962) that valuable information comes from *costly creative production*. However, unlike Arrow's classic setting where under-investment stems from *inappropriability* (the difficulty of excluding freeriders who copy innovators' data), we identify a unique challenge of *statistical substitutability*: Even if strict property rights are enforced, the massive amount of correlated imitative data allows the buyer to approximate innovators' contents, thus depressing their returns and discouraging their investment in content creation.

Focusing on the production of creative content on online platforms—where creators strategically compete for user attention exploiting the recommendation system—several recent works study the

optimal contest design within an online platform (Jagadeesan et al., 2023a,b; Yao et al., 2023a,b; Immorlica et al., 2024; Golrezaei et al., 2025), a concept dating back to Glazer and Hassin (1988). While creators are strategic in both models, our model focuses on *pricing* for AI training, rather than *ranking* for user consumption. So the main friction in our model is the *imperfect substitution* of human creators, instead of the *congestion* of user attention.

For the proposed mechanism, we draw inspiration from the extensive economics literature analyzing the role of intermediaries in different markets (Rubinstein and Wolinsky, 1987; Farboodi et al., 2023) as well as various intermediate legal regimes utilizing collective management organizations (CMOs) like ASCAP and BMI. But while CMOs primarily act as centralized clearinghouses to reduce transaction costs (Cotter, 2005; Gilbert, 2017), we argue that in the market for AI training data, reducing transaction costs alone is not enough. Due to *statistical substitutability* among creators, the solution must additionally prevent information leakage and subsidize originality. We achieve this by adopting a *two-part tariff* structure (Oi, 1971): by combining price-per-effort payments with a fixed subsidy, this mechanism redistributes the surplus attained by socially optimal production back to all participants, thus advancing technological progress and sustaining human creative incentives at the same time.

Finally, many methods have been proposed in the machine learning literature to quantify the value of training data. For example, Ghorbani and Zou (2019) and Jia et al. (2019) proposed the *data Shapley value* to measure the marginal contribution of a data point to model performance, while Koh and Liang (2017) developed *influence functions* to approximate the counterfactual impact of the absence of a data point. Various data marketplaces have also been established (Agarwal et al., 2019; Schomm et al., 2013; Golrezaei and Nazerzadeh, 2014; Muschalle et al., 2013). But these works study fair attribution for a *fixed* dataset whereas we focus on the *production* of data.

3 Static Model and Short-Term Market Failures

We begin by analyzing a static market with multiple content creators and one AI firm. Under the classic copyright binary, either the unauthorized use of copyrighted works to train AI models is copyright infringement—in which case the AI firm is required to obtain authorization from each

copyright owner—or the use is fully shielded from copyright liability under fair use. We first show that, consistent with common intuition, a free-for-all system fails to provide incentives for human content creation. We further show that the strict property rights regime leads to two market failures: first, due to the buyer's significant market power, creators invest less effort into creating high-quality content than is socially optimal (an *under-investment* failure); second, this under-investment applies more forcefully to highly original or innovative content creators than to those whose content is highly correlated with others', leading to a proliferation of homogenized content (an *adverse selection* which we call the "originality penalty").

3.1 Static Model Setup

We formalize the interaction between human creators and an AI developer as a market for information where production is costly and products are correlated. Specifically, consider a market with two types of participants: There are $N \geq 2$ information producers (human content producers), indexed by $i = 1, 2, \ldots, N$, and a single buyer (e.g., an AI firm acquiring training data from creators). The economy centers on a fixed but unobserved state variable $X \in \mathbb{R}$, representing the best (or most useful) response to a given user query.[4]

Production of Content. We model the created by each creator i as a noisy signal s_i of X, namely

$$s_i = X + \xi_i,\ \xi_i = \varepsilon_i + \beta_i \eta, \quad \forall i = 1, 2, \ldots, N. \tag{1}$$

As such, ξ_i is a random variable denoting the deviation of an individual creator's content from X. The individual deviation (or noise) ξ_i has two components: an "idiosyncratic component" ε_i whose precision is an increasing function of the individual agent's effort to produce high-quality content, and a "common factor" η that is shared among all creators but weighted differently by creators (as captured by the parameter β_i). For simplicity, we assume all β_i's are public information.

We now explain the idiosyncratic and common components of the noise in greater detail. To produce content, agent i chooses an *effort level* $h_i \geq 0$ for production, incurring a convex cost

[4] In some contexts, the "best" response has a clear, objective meaning, as when a user asks for a solution to a math problem; in other contexts, the best response is more subjective and geared to a user's particular needs, as when a user asks a Chatbot to write a bedtime story for her child, in which case the variable X measures how well the model can understand and respond to the user's preferences given her prompts.

$C(h_i) := \frac{c}{2}h_i^2$ (where $c > 0$ is common knowledge). Effort reduces the variance of the idiosyncratic component such that $\varepsilon_i \sim N(0, h_i^{-1})$. This captures the property of increasing marginal production cost when creators strive for higher precision (via, e.g., research, fact-checking, creative polishing).

For the common factor (the $\beta_i\eta$ term), we suggest two (mutually compatible) interpretations. In the first interpretation, η represents a general *trend*, such as a stylistic convention, so a large β_i denotes agent i's conformity to prevailing conventions whereas a low (near-zero) β_i signifies agent i's willingness to buck the trend and strike out on her own, to march to the beat of a different drummer. That is the sense in which we characterize high-β creators as "rank and file" or "run of the mill" and low-β creators as highly "original" or "creative." And the "originality penalty" is the greater quality (effort) decline suffered by highly original creators, as compared to run-of-the-mill creators, when we move from the social optimum to the AI market equilibrium.

In the second interpretation, η represents a common *bias*, such as socially prevalent prejudice against a group or practice or thought, and β_i measures the extent to which agent i shares that bias. Then, the originality penalty can be interpreted as *bias amplification* by AI—not only in the general sense that AI-generated content is more biased (less accurate) than is socially optimal, but in the specific sense that the equilibrium quality reduction (compared to the social optimum) is *greater* for content with low common-bias than for content with high common-bias.

Depending on the context, one or the other interpretation may be more natural. As noted, the two interpretations are entirely compatible because, in the current model specification, the trend is biased (the addition of η can never decrease the deviation of the creators' signals from the true state, $s_i - X$), so higher creativity implies higher quality—all else (i.e., effort) being equal.[5]

A defining feature of our framework is the correlation structure of signals. We model the common factor η using a hierarchical Bayesian structure (one can also equivalently view our model from a minimax robust optimization perspective, as we prove in Lemma 12 of Appendix A):

$$\mu_\eta \sim \mathcal{N}(0, \gamma), \quad \eta \mid \mu_\eta \sim \mathcal{N}(\mu_\eta, \sigma_\eta^2), \tag{2}$$

[5]In principle, it's possible to conceptualize a notion of originality or creativity (or a residual of it) that is independent of accuracy or quality. But, without getting into the details of how such a notion could be mathematically operationalized, our conjecture is that if there is an equilibrium penalty for the kind of originality that contributes to quality (as we have shown there is), then *a fortiori* there must also be an equilibrium penalty for the kind of originality that does not contribute to quality.

where both $\gamma \geq 0$ and $\sigma_\eta^2 \geq 0$ are publicly known parameters. The term μ_η represents the *systematic bias* of the current information environment, while γ captures the *uncertainty* regarding this bias. This Bayesian setup is crucial for our dynamic analysis later on: as the buyer trains on data produced by its own previous models—a common practice called recursive training (Shumailov et al., 2024)—the systematic bias μ_η becomes harder to filter out, effectively increasing the γ.

Acquisition of Data. To incentivize the production of high-quality data while balancing the total purchase cost, the buyer offers a price $p_i \geq 0$ per unit of effort to each creator i.[6] The buyer—as a leader in our Stackelberg game—has full power to decide the p_i's. The creators react to the price vector $\boldsymbol{p}$ by deciding an individual effort level h_i that maximizes their own quasi-linear payoff:

$$U_i(h_i, p_i) = h_i p_i - C(h_i), \quad \forall i = 1, 2, \ldots, N. \tag{3}$$

The first-order condition gives creators' *best response* to any price vector $\boldsymbol{p}$, namely $\boldsymbol{h}(\boldsymbol{p})$, a

$$h_i(\boldsymbol{p}) = \frac{p_i}{c}, \quad \forall i = 1, 2, \ldots, N.$$

Given this one-to-one correspondence between $\boldsymbol{p}$ and $\boldsymbol{h}(\boldsymbol{p})$, it is convenient to likewise define $\boldsymbol{p}(\boldsymbol{h})$ as the price vector $\boldsymbol{p}$ inducing a best response of $\boldsymbol{h}$. Specifically, we define $p_i(\boldsymbol{h}) = h_i c$ for all i.

Aggregation of Data. After the signals $\boldsymbol{s} = (s_1, s_2, \ldots, s_N)'$ are produced according to Equations (1) and (2), the AI firm buys all these signals (by paying the promised $p_i h_i$ to each creator i) to construct an estimate of X, which we call $\hat{X}$. The firm derives revenue from the precision (quality) of this estimator, reflecting the fact that it can earn more when its model's responses to user queries are more useful. Specifically, we assume the buyer seeks to minimize the Mean Squared Error (MSE) of $\hat{X}$ subject to unbiasedness:

$$\min_{\hat{X}} \quad \mathbb{E}[(\hat{X} - X)^2] \qquad \text{s.t.} \quad \mathbb{E}[\hat{X}] = X. \tag{4}$$

Under our model, the optimal solution to Equation (4) is the best linear unbiased estimator (BLUE)

[6] We assume the buyer and creators contract directly on the effort vector $\boldsymbol{e} = (e_1, e_2, \ldots, e_N)'$. In practice, while "effort" is not directly observable, platforms routinely use observable proxies to price data quality (e.g., unit tests in coding, verified status in social media). We treat e_i as a reduced-form summary of these observable quality metrics.

of X. While modern LLM optimizes complex objectives like cross-entropy, the minimization in Equation (4) is a canonical approximation for the value of information: under Gaussian assumptions, minimizing MSE is equivalent to maximizing log-likelihood, aligning with the probabilistic foundations of generative models.

The firm's revenue depends on the *effective precision* (also known as Fisher information in statistics) of its estimator $\hat{X}$, denoted K and defined as $K(\boldsymbol{h}) = \mathrm{Var}(\hat{X})^{-1}$, with the interpretation that (a monotonic transformation of) the effective precision functions as the "total factor productivity" of the AI model. Practically speaking, a higher precision (lower variance) means that the content generated by the AI model in response to a user query is more useful to the user, which helps the firm earn greater revenue. The AI developer's payoff (profit) function is the revenue earned from precision minus the price paid to creators for their content, that is,[7]

$$\Pi(\boldsymbol{p}) := K(\boldsymbol{h}(\boldsymbol{p})) - \sum_{i=1}^{N} p_i h_i(\boldsymbol{p}),$$

where $\boldsymbol{h}(\boldsymbol{p}) = \boldsymbol{p}/c$ is the creators' best response to $\boldsymbol{p}$. Given the one-to-one correspondence between $\boldsymbol{h}$ and $\boldsymbol{p}$, it would be more convenient to write Π in terms of creators' effort $\boldsymbol{h}$, namely

$$\Pi(\boldsymbol{h}) = K(\boldsymbol{h}) - \sum_{i=1}^{N} p_i(\boldsymbol{h}) h_i = K(\boldsymbol{h}) - \sum_{i=1}^{N} c h_i^2. \tag{5}$$

We summarize the sequence of play in the following Definition.

Definition 1 (Sequence of Play)**.** *The sequence of play is as follows:*

1. *The buyer offers a unit price $\boldsymbol{p} = (p_1, p_2, \ldots, p_N)'$ to all creators simultaneously.*
2. *Creators simultaneously choose their best response $\boldsymbol{h}(\boldsymbol{p})$, which maximizes payoff $U_i(h_i; p_i)$.*
3. *From the signals $\boldsymbol{s} = (s_1, s_2, \ldots, s_N)'$ generated by Equations* (1) *and* (2)*, the buyer finds the estimator $\hat{X}$ minimizing the objective function in Equation* (4) *and receives profit $\Pi(\boldsymbol{h})$.*

[7] In our model, the buyer's profit is *linear in precision increase* rather than *linear in variance reduction* (another common option in information economics used by, e.g., Acemoglu et al. 2022). The reason behind our modeling choice is two-fold: First, standard information economics establishes that when information guides *downstream decisions*, utility is linear in precision (and thus *strictly convex*, instead of linear, in variance reduction) (Veldkamp, 2011, Chapter 8). Second, this convexity captures the "emergent abilities" of Generative AI (Wei et al., 2022): reducing error rates below critical thresholds triggers *phase transitions* that unlock new capabilities—beyond incremental improvements of existing ones—thereby yielding utility gains strictly exceeding the linear function of variance reduction.

In this framework, we contrast the *equilibrium* outcome of the game defined above with the *first best* solution. The equilibrium is the outcome of the game specified by, Definition 1 where creators and the buyer pursue their own interests. We denote the resulting effort vector and prices as $\boldsymbol{h}^*$ and $\boldsymbol{p}^*$, respectively. On the other hand, the first best solution $\boldsymbol{h}^{\text{sp}}$ is the outcome dictated by a benevolent social planner, who controls the production efforts to maximize social welfare:

$$W(\boldsymbol{h}) := \Pi(\boldsymbol{h}) + \sum_{i=1}^{N} U_i(h_i; p_i(\boldsymbol{h})) = K(\boldsymbol{h}) - \sum_{i=1}^{N} \frac{c}{2} h_i^2, \tag{6}$$

where we recall that $\boldsymbol{p}(\boldsymbol{h})$ is the price vector inducing $\boldsymbol{h}$ as the best response.

3.2 Failure of Free-for-All Regime

We now analyze the market equilibrium under existing policy frameworks, beginning with the free-for-all ("all model training is fair use") regime. In our model, it is easy to see that this regime fails to provide sufficient creative incentives, as noted in the Remark below.

Remark 1. *A regime that permits AI to take human-created content for free cannot sustain creative incentives.*

The reason is straightforward: If the buyer offers no price to content providers, i.e., $\boldsymbol{p} = (0, 0, \ldots, 0)'$, then the only best response of creators is $\boldsymbol{h}(\boldsymbol{0}) = (0, 0, \ldots, 0)'$. That is, creators never invest any effort in production. This is in line with common intuition; indeed, it is the very freerider problem that justifies the existence of intellectual property (IP) laws. Henceforth, we only focus on the "strict property rights" regime.[8]

3.3 Short-Term Market Failures: Under-Investment and Originality Penalty

The first market failure we identify is the *under-investment* in producing high-quality content. Because the buyer has full buyer power to decide the market price $\boldsymbol{p} = (p_1, p_2, \ldots, p_N)'$ (also

[8]Even if the AI firm can lawfully take human-created content for free *ex post*, the firm might enter into a contract to purchase content for a positive price, thereby incentivizing content production *ex ante*. Setting aside the practical difficulties (transaction costs) entailed in identifying and negotiating with potential content creators *ex ante*, such an arrangement, though legally different from an IP-based regime (in that it's based on contract rather than property rights), would practically operate much like the "property rights" regime analyzed below.

known as *monopsony power*), it may offer a lower price, which incentivizes lower effort levels from the creators but elicits higher profit per effort. This leads to the following result:

Proposition 2 (**Underpowered Creative Incentives).** *Under the market defined in Section 3.1, market equilibrium* $\boldsymbol{h}^*$ *and social planner solution* $\boldsymbol{h}^{sp}$ *are unique. Furthermore,* $K(\boldsymbol{h}^*) < K(\boldsymbol{h}^{sp})$.

The formal proof of Proposition 2, as well as Proposition 3 presented later, can be found in Appendix A. This market failure is consistent with common economic intuition that monopsony leads to under-production. However, the second market failure that we identify, an *originality penalty* which can be thought of as a form of *adverse selection*, is novel and distinctive to our context.

Because of the common factor η in Equation (1) (namely $s_i = X + \beta_i\eta + \varepsilon_i$, i.e., every creator i's signal s_i contains this potentially biased η to some extent), the buyer—who is only interested in recovering X but not η—faces a dual objective: *(i)* incentivizing effort to reduce idiosyncratic noise ε_i, while *(ii)* mitigating the biasedness arising from the common η. Intuitively, one may consequently expect the AI firm to offer a better price to highly original content creators (those with low β_i), since their effort helps better recover X. Such a conjecture seems especially reasonable given that our model is not characterized by heterogeneous production costs (with respect to the "common factor" term), so producing "original" content does not cost more than "conventional" content. Thus, even allowing for the fact that, due to its monopsony power, the AI firm "lowballs" human creators by offering a price that does not incentivize sufficient investment in quality (Proposition 2), one would expect the AI firm to lowball original content creators less than it lowballs run-of-the-mill creators in order to get the benefit of the former's individuated (non-redundant) signals.

But it turns out that the equilibrium outcome is precisely the opposite: in the AI market equilibrium, relative to the social optimum, the reduction in quality of content is *greater* for original creators than for run-of-the-mill creators. We call this the "originality penalty."

Proposition 3 (**Originality Penalty).** *For each creator* $i = 1, 2, \ldots, N$*, let* $R_i := h_i^*/h_i^{sp}$ *be the ratio between equilibrium and socially optimum. Then* R_i *is* strictly increasing *in* β_i *(recall that* β_i *is defined in Equation* (1) *such that* $s_i = X + \beta_i\eta + \varepsilon_i$ *for all* $i = 1, 2, \ldots, N$*), and furthermore:*

- *For any "innovator" i with* $\beta_i \approx 0$ *(who does not rely on the common factor, hence the name), regardless of the remaining* β*'s, their ratio* $R_i \approx \frac{1}{2} = 0.5$.

- *For any "imitator" i with $\beta_i \gg 0$, their ratio $R_i > \frac{1}{2}$. Furthermore, in the limiting case of a homogeneous market saturated by imitators ($N \to \infty$ and $\beta_i \equiv \beta > 0$), the ratio converges to*

$$\lim_{N\to\infty} R_i = \frac{1}{\sqrt[3]{2}} \approx 0.79.$$

The counterintuitive result in Proposition 3 arises because of the interplay of *monopsony power* (the buyer having full market power) and *diminishing returns* (the increase in precision due to imitators' effort gradually diminishes because of content correlation). "Innovative" creators, because they have a near-zero weight on the common factor ($\beta_i \approx 0$), are essentially facing the monopsonist AI firm alone. Their creative outputs are almost independent, such that marginal benefits do not decline steeply in volume. The social optimum occurs when the marginal production cost equals this marginal value. However, because the production cost scales quadratically with effort (recall $C(h) = \frac{c}{2}h^2$), buying less can reduce the price for *all* content. A monopsonist therefore strategically restricts the purchase volume in exchange for a lower unit price (and higher unit profit). Solving the tradeoff between volume and unit price gives $h_i^* \approx \frac{1}{2}h_i^{\text{sp}}$.

By contrast, "imitative" creators, given their high reliance on the common factor ($\beta_i \gg 0$) and thus high correlation with other creators, face strong diminishing returns. This is as if all creators were reporting the same news, so the first report is highly valuable but subsequent reports are nearly devoid of value. The socially optimal production is then "one news report in total," so the level of effort is low even in the first-best world (h_i^{sp}). In the market equilibrium, even if the buyer has full monopsony power, the price must still incentivize (roughly) this total production level: missing a part of the first news report results in a value loss much higher than the cost saved. Formalizing these arguments, we prove that $h_i^* \geq \frac{1}{\sqrt[3]{2}}h_i^{\text{sp}} \approx 0.79h_i^{\text{sp}}$ if β_i is sufficiently large.

Note that Proposition 3 is *not* saying that the AI firm pays a lower unit price for original creators' content than for run-of-the-mill creators' content. Rather, the meaning of the "originality penalty" is that when we move from the social optimum to the AI market equilibrium, original creators face a greater drop in incentives to invest in effort compared to run-of-the-mill creators. In other words, AI penalizes originality more than it penalizes run-of-the-mill content by disincentivizing the original creators more. In fact, given the redundancy (high correlation) of rank-and-file creators' output,

eliciting a great deal of effort from them would not be worthwhile even in the social optimum. Consequently, relative to the low optimal level of effort, the AI firm's strategic underpricing does not result in a great deal of reduced effort and reduced precision. Original creators, by contrast, should optimally exert great effort, while the AI firm's underpricing greatly reduces their incentive to do so and, thus, the precision of their content.

The originality penalrty identified in Proposition 3 poses a risk to the *sustainability* of human creative incentives: since original creators in reality (e.g., domain experts or premium publishers) typically have *outside options*, revenue generated at a sub-optimal scale may fail to cover their opportunity costs. The exclusion of original creators' potential contributions from AI training sets sows the seeds of the model's own decay. We characterize this long-term crisis in the next section.

4 Long-Term Market Failure: The Curse of Precision

The foregoing analysis shows that standard IP rights cannot solve the AI data crisis: market power and data correlation lead to under-investment and adverse selection, which essentially *refuses to debias* (by incentivizing the production of biased contents). The models in Section 3 are static—focusing on a *fixed snapshot* of the data market—but the AI training pipeline is recursive: AI models are iteratively trained on web-scraped datasets that accumulate over time, but such online data is increasingly *polluted* by content creators who utilize previous AI models. This creates a closed-loop supply chain where the system's output becomes its own future input, transforming the static "bias amplification" problem into a dynamic "model collapse."

In this section, we incorporate dynamics into our analysis to study how this recursive structure affects the long-run market equilibrium. Specifically, we introduce two dynamic state variables that evolve according to the continuous time $t \in [0, +\infty)$:

- **Data Stocks $S_O(t)$ and $S_C(t)$:** We define $S_O(t)$ and $S_C(t)$ as the *cumulative stock* of "original" and "correlated" data available at time t, which evolve according to standard capital accumulation dynamics: they increase via the inflow of new data production, and decrease due to depreciation (see Equation (8) for the formal definition). These stocks of training data together determine the *current model precision* $K(t)$ at any time t.

- **Recursive Bias $\gamma(t)$:** Recall from Equation (2) that the common bias η is centered around a *systematic bias* $\mu_\eta \sim \mathcal{N}(0, \gamma)$, which was fixed in the static case. In the dynamic model, such bias is no longer exogenous: as the AI model becomes *more powerful*, the "imitators"—interpreted as generalist creators who treat AI outputs as the ground-truth—rely on it *more heavily*. Consequently, the systematic bias μ_η may be more severe (e.g., due to hallucination of the AI model). We capture this *bias amplification* by modeling $\gamma(t)$ as a linear function of model precision $K(t)$. That is, stronger models induce more significant systematic biases.

To analyze the dynamics of this system, we consider a continuum of creators facing a uniform market price $p(t)$ that evolves over time t (whose detailed formulation is deferred to Section 4.1). The central question is whether the market *self-corrects over time*: One might assume that as high-quality human data becomes scarce (i.e., the ratio between original and correlated data $S_O(t)/S_C(t) \to 0$), the market price $p(t)$ will rise, inducing innovative creators to return. We prove, to the contrary, that market failure persists. We identify a central mechanism driving this failure, which we term the *curse of precision*. This mechanism contains two components:

1. **Upper Bound on Market Price:** Instead of rising to reflect the scarcity of original data, the market price $p(t)$ is capped by the *diminishing marginal value* of correlated data. When either the stock of correlated data accumulates or model precision increases, this marginal value vanishes, and consequently, the price is "trapped" at a low level. Therefore, as formalized in Proposition 4, the market price is effectively *decoupled* from the scarcity of innovators.

2. **Barrier of Model Precision:** Should this low market price drives innovators out of the market, the model—instead of sustaining on a massive amount of correlated data from the imitators—hits a "ceiling." In Proposition 5, we prove that even with an infinite supply of correlated data, the model precision $K(t)$ still cannot bypass a constant upper bound.

To summarize, our core insight—the *curse of precision*—is that as the AI model improves (more precisely estimating the state X, giving better answers to user queries), it sows the seeds of its own stagnation: Superior model outputs incentivize rank-and-file human creators to rely more heavily on AI outputs in producing the content that the model will train on, which lowers the price the AI firm is willing to pay content creators, which may in turn induce the exit of innovative creators,

thereby depriving the buyer of the fresh human content it needs to keep up its good performance. Thus, achieving a high level of precision at a given time is no guarantee of continued success and in fact may be an omen of impending failure.

4.1 Dynamic Model Setup

To capture the ecosystem's long-run evolution, we consider a population-level fluid model where all content creators form a *continuum*.

Creators and Market Structure. In light of the bifurcation result in Proposition 3, we divide the large population of content creators into the two representative classes:

- **Innovators Producing Original Data (Type-O):** These creators share a low β, denoted by $\beta_O \approx 0$.[9] They can be interpreted as domain experts whose contents are based on their own opinion. Therefore, their data is almost uncorrelated with the common factor η.
- **Imitators Producing Correlated Data (Type-C):** This type of creators share a high β, namely $\beta_C \gg 0$. They represent generalist content creators who heavily rely on (and trust) AI tools. They are more exposed to the common factor η due to AI tools.

The buyer sets a uniform price $p(t)$ for each time $t \geq 0$.[10] Creators enter the market if the price $p(t)$ provides sufficient utility compared to their outside options. Instead of explicitly modeling the continuum of outside options, we denote the *mass of active creators* of type $j \in \{O, C\}$ attracted by price p as $N_j(p)$. We assume both supplies are increasing in price ($N_j' > 0$ for $j \in \{O, C\}$), and the supply of imitators is *more elastic* to prices than that of innovators, namely

$$\epsilon_C(p) > \epsilon_O(p),\ \forall p > 0, \quad \text{where } \epsilon_j(p) := \frac{\mathrm{d} \ln N_j(p)}{d \ln p},\ \forall j \in \{O, C\}. \tag{7}$$

In words, this assumption captures the asymmetric impact of technology on different types of content creators. For imitators—who are mostly generalist creators—AI tools greatly lower their

[9] One may generalize the constant β_O to a distribution of β's among innovators. For simplicity, we treat β_O (and β_C below) as constants, which can be interpreted as the means of such distributions.

[10] While the micro-level model in Section 3 allowed for creator-specific prices p_i, in this macroscopic dynamic model, we assume a commoditized market with a single clearing price $p(t)$ for analytical tractability.

barrier to entry, allowing the supply of correlated content to scale rapidly in response to market prices. Conversely, the production of original insights relies on domain expertise—a scarce human intelligence that cannot be algorithmically amplified—making the supply of innovators less responsive to price changes. As we prove in Lemma 13 of Appendix B, Equation (7) is equivalent to the *imitator-innovator ratio* $\rho(p) := \frac{N_C(p)}{N_O(p)}$ being strictly increasing ($\rho' > 0$). Finally—consistent with the static model—the effort of active creators is their best response to price p, namely $e(p) := p/c$.

Accumulation and Depreciation of Data. We define $S_O(t)$ and $S_C(t)$ as the *cumulative stocks* of original and correlated data available to the model at time t (dynamic counterpart to the total effort term $\sum_i h_i$ in the static model). They evolve according to capital accumulation dynamics:

$$\frac{\mathrm{d}S_j}{\mathrm{d}t}(t) = \underbrace{N_j\big(p(t)\big)e\big(p(t)\big)}_{\text{Production Inflow}} - \underbrace{\delta S_j(t)}_{\text{Depreciation}}, \quad \forall j \in \{O, C\}, t \geq 0. \tag{8}$$

The first term captures the total inflow: mass of creators N_j multiplied by their production effort e. The second term captures the *concept drift* phenomenon in machine learning (Widmer and Kubat, 1996): due to the evolving nature of language and consensus, historical training data depreciates at a rate of $\delta \in (0, 1)$ (which we assume is fixed and known). Therefore, the buyer cannot rest on historical stocks; a continuous inflow of new data is necessary to maintain model performance.

Recursive Bias and Precision. To capture the recursive nature of AI training and the reality that generalist content creators are increasingly utilizing powerful AI tools during production, we turn Equation (2) into a dynamic model: Instead of treating the systematic bias μ_η and its uncertainty γ as exogenous parameters, the common bias η at time t, namely $\eta(t)$, is now generated as

$$\mu_\eta(t) \sim \mathcal{N}(0, \gamma(t)), \quad \eta(t) \mid \mu_\eta(t) \sim \mathcal{N}(\mu_\eta(t), \sigma_\eta^2), \quad \forall t \geq 0, \tag{9}$$

where the prior variance $\gamma(t)$ is linear in the current model precision $K(t)$ (defined later) as[11]

$$\gamma(t) = \lambda K(t), \quad \text{where } \lambda > 0 \text{ is publicly known.} \tag{10}$$

[11] In practice, imitators can only access the *previous model* (whose precision is $K(t^-)$) but not the *newly generated model* (whose precision is $K(t)$). However, as time t is continuous, it is unnecessary to distinguish between them.

Finally, the effective precision $K(t)$ is determined by the stocks $S_O(t)$, $S_C(t)$ and the bias $\gamma(t)$. The micro-level effective precision formula (derived in Lemma 11 in the Appendix) is:

$$K(\boldsymbol{h}) = \sum_{j=1}^{N} h_j - \frac{(\sum_{j=1}^{N} h_j \beta_j)^2}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{j=1}^{N} h_j \beta_j^2}.$$

Therefore, in the dynamic case where the original data has mass $S_O(t)$ and loading β_O, while the correlated data has mass $S_C(t)$ and loading β_C, the precision is defined via a fixed-point formula:

$$\begin{aligned} K(t) &= S_O(t) + S_C(t) - \frac{\big(S_O(t)\beta_O + S_C(t)\beta_C\big)^2}{\big(\sigma_\eta^2 + \gamma(t)\big)^{-1} + \big(S_O(t)\beta_O^2 + S_C(t)\beta_C^2\big)} \\ &\approx S_O(t) + \frac{S_C(t)}{1 + \big(\sigma_\eta^2 + \lambda K(t)\big) S_C(t)\beta_C^2}, \quad \forall t \geq 0, \end{aligned} \tag{11}$$

where the last step uses $\beta_O \approx 0 \ll \beta_C$ and that $\gamma(t) = \lambda K(t)$ from Equation (10).

Buyer's Objective. We finally define the objective of the buyer. In a competitive market where an AI firm cannot commit to long-term subsidies but has to focus on instantaneous profit, we assume that the buyer sets price $p(t)$ to myopically maximize *instantaneous profit*:

$$\Pi_{\text{inst}}(t) := \underbrace{\frac{\mathrm{d}K}{\mathrm{d}t}(t)}_{\text{Precision Increase}} - \underbrace{\Big(N_O\big(p(t)\big) + N_C\big(p(t)\big)\Big) p(t) e\big(p(t)\big)}_{\text{Purchase Expenditure}}, \quad \forall t \geq 0. \tag{12}$$

4.2 Upper Bound on Market Price

We analyze the long-run stability of this market. In a healthy market, if the stock of original data $S_O(t)$ drops, the scarcity should drive the price $p(t)$ up, thereby attracting innovators back. However, due to the common factor η, the marginal value of correlated data *decays over time*, thus depressing the price: unlike the marginal contribution of innovators—which is always $\frac{\partial K}{\partial S_O}(t) = 1$—that of imitators exhibits diminishing returns analogous to Proposition 3. Specifically, we derive

$$\frac{\partial K}{\partial S_C}(t) \leq i_C(t) := \big(1 + \big(\sigma_\eta^2 + \lambda K(t)\big) S_C(t)\beta_C^2\big)^{-2} \leq 1. \tag{13}$$

The formal proof is in Lemma 14. From Equation (13), there are two forces driving down $\frac{\partial K}{\partial S_C}$:

- **Data Accumulation ($S_C(t)$):** As the stock of correlated data $S_C(t)$ grows large, the marginal value of additional correlated data diminishes. This aligns well with standard market dynamics: scarcity commands a premium, whereas redundancy leads to price depreciation.
- **Recursive Bias ($K(t)$):** The $K(t)$ term additionally introduces a recursive effect. That is, a stronger model (larger $K(t)$) effectively "pollutes" the imitator data with amplified bias, which further drives down the value of imitators' data beyond the redundancy effect.

While a standard market stabilizes under the *data accumulation* effect via price adjustment, the *recursive bias* effect makes the market unrecoverable. The buyer—facing a overwhelming supply of imitators whose data has vanishing marginal value—fails to raise the price $p(t)$ to attract innovators back. Specifically, in Proposition 4, we prove that *(i)* a smaller $i_C(t)$ induces a lower market price $p(t)$, and *(ii)* as $i_C(t) \to 0^+$—driven by either the accumulation of correlated data $S_C(t)$ or the increase in model precision $K(t)$—the price is capped at a fixed level specified in Equation (14).

Proposition 4 (Upper Bound on Market Price)**.** *At any time $t \geq 0$, the market price $p(t)$ cannot exceed a threshold determined by $i_C(t)$, namely $p(t) \leq \bar{p}(i_C(t))$ for some strictly increasing function $\bar{p}()$. Specifically, when $i_C \to 0^+$, this upper bound $\bar{p}(i_C)$ converges to a "price trap" p^*. Formally,*

$$\lim_{i_C \to 0^+} \bar{p}(i_C) = p^*, \quad \textit{where } p^* = \left(1 + \frac{N_C(p^*)}{N_O(p^*)}\right)^{-1}. \tag{14}$$

The proof of Proposition 4, as well as Proposition 5 presented later, can be found in Appendix B. Proposition 4 characterizes the limiting behavior of the market price. As the stock of correlated data accumulates (i.e., $S_C(t) \to \infty$) or the recursive bias amplifies due to powerful AI models (i.e., $K(t) \to \infty$), the marginal value of imitators' data in Equation (13) vanishes. Consequently, the market price $p(t)$ is asymptotically trapped below p^*, which is induced by the *imitator-innovator ratio* $\frac{N_C(p^*)}{N_O(p^*)}$. Therefore, if AI models become widespread and lower barriers to entry, the supply of imitators N_C increases far faster than that of innovators N_O. This further pushes the trap level p^* down towards zero. This creates a risk: if such price p^* is insufficient to attract *any* innovator (interpreted as domain experts with outside options), what will happen to the AI model?

4.3 Barrier of Model Precision and "Curse of Precision"

To answer this question, we now suppose all innovators exit the market ($S_O(t) \to 0$) and see whether the buyer can survive solely on massive amounts of imitator data ($S_C(t) \to \infty$). The widely believed machine learning "scaling laws" (Kaplan et al., 2020; Hoffmann et al., 2022) suggest *quantity* of training data can substitute for *quality*: by training on a sufficiently large amount of synthetic data, the AI model might suddenly be able to discover the ground-truth. However, we prove that the quantity of synthetic data cannot compensate for the scarcity of original data:

Proposition 5 (Barrier of Model Precision)**.** *Suppose innovators exit the market ($S_O = 0$). Even if the stock of imitator data grows indefinitely ($S_C \to \infty$), the equilibrium model precision $K(t)$ only converges to a finite upper bound $\bar{K}$ determined by the recursive bias coefficient λ in Equation* (10)*:*

$$\lim_{S_C \to \infty} K(t) = \bar{K} := \Big(-\sigma_\eta^2 + \sqrt{\sigma_\eta^4 + 4\lambda/\beta_C^2} \Big) \Big/ 2\lambda. \tag{15}$$

Proposition 5 captures a failure mode similar to the "model collapse" phenomenon in empirical machine learning (Shumailov et al., 2024; Alemohammad et al., 2023). This term typically refers to the performance degeneration when an AI model is trained on *self-generated data*—specifically, the model's output distribution loses variance and *collapses* onto a few modes (hence the name). Crucially however, our result reveals that this failure is not limited to such closed-loop training. We show that even if the buyer is continuously acquiring *fresh data* from human creators—via a market containing both innovators and imitators—the model still hits a performance ceiling.

This terms a paradox we call the *"curse of precision"*: a more powerful AI induces stronger correlation among imitators (e.g., generalist creators), which drives down the market price (Proposition 4) and potentially hurts the model's long-run performance (Proposition 5). Therefore, we now aim to design a mechanism for the AI data market that can recover market efficiency and sustainability.

5 Mechanism Design: Restoring Efficiency and Sustainability

Sections 3 and 4 have revealed the inadequacies of existing approaches to the design of AI data market. If AI firms can simply take human content for free, then humans are deprived of incentives to invest in creating high-quality content. But it turns out that IP rights are also insufficient. IP rights are meant to incentivize creativity, but they underpower creative incentives, especially for innovative content. In the long run, these market failures create a trajectory toward "model collapse" where the AI model eventually starves itself of the novel human input it needs to remain robust.

In this section, we move from diagnosis to treatment. We propose and analyze the architecture of a *data intermediary* that can restore human creative incentives while simultaneously providing AI models with the inputs they need to improve.

The intermediary in our solution is analogous to collective management organizations (CMOs) that represent a group of content owners, such as performing rights organizations like ASCAP and BMI or reproduction rights organizations like the Copyright Clearance Center (CCC).[12] However, unlike traditional copyright CMOs, whose primarily function is to reduce transaction costs (Cotter, 2005; Gilbert, 2017), the first role of our intermediary is to correct the *collective action problem* among the creators: Because individual creators are atomic and their content is often substitutable, they suffer from coordination failure and tend to impose a negative externality on each other when facing the AI firm, so removing transaction costs would not be sufficient to restore socially optimal content creation. As a collective licensing body that aggregates the rights of a critical mass of creators, the intermediary internalizes the negative externalities of data sales on behalf of the content creators.[13] In addition to aligning individual incentives with social welfare, our intermediary ensures the active participation of all parties. Specifically, the intermediary's economic function is twofold:

[12]More broadly, our proposal for going beyond the copyright binary is not without precedent. The US Copyright Act contains a number of such "intermediate" regimes, including: 17 U.S.C. § 111 (governing secondary transmissions of broadcast programming by cable); 17 U.S.C. § 112 (governing ephemeral recordings); 17 U.S.C. § 114 (governing digital audio transmission of sound recordings); 17 U.S.C. § 115 (governing compulsory licenses to make and distribute phonorecords of nondramatic musical works); 17 U.S.C. § 118 (governing noncommercial broadcasts); 17 U.S.C. § 119 (governing secondary transmissions of distant TV station signals by satellite carriers); 17 U.S.C. § 122 (governing secondary transmissions of local TV station signals by satellite carriers).

[13]If the intermediary is a private party, there can be delegation problems between the content creators and the intermediary. Currently, we focus on the legal design of a regulated intermediary and abstract from that second layer of frictional delegation.

1. **Internalizing Correlation Externalities:** Individual creators, acting in isolation, are compensated based on their *marginal contribution* conditional on others' data. Consequently, the "common" component shared among creators is undervalued, leading to the static market failures identified in Propositions 2 and 3. By managing the creators as a portfolio, the intermediary not only increases bargaining power but also internalizes the information-leakage externality. The intermediary prices creators' data based on its *collective value* rather than the fragmented *incremental contribution*.

2. **Ensuring Sustainability via Transfers:** To resolve the tension between efficient production (which requires unit prices equal to marginal production costs) and incentive preservation (which requires subsidies to meet outside options), we propose a nonlinear pricing schedule. Specifically, we implement a *two-part tariff*: a *unit price* set to equate the buyer's demand with social marginal costs, thus inducing socially optimal effort, and a fixed *lump-sum transfer* (licensing fees or subsidies) to redistribute the surplus back to all agents, ensuring that their participation constraints are met without distorting production incentives.

In what follows, we formally introduce our intermediated market design. We will assume, for purposes of this analysis, that the intermediary is fully informed about the creators. Though this assumption is unlikely to be satisfied in the real world, it is a useful benchmark in that any inefficiencies identified here will persist in a more complex model with information asymmetries. Note, moreover, that the market failures we identified above persist *even in the full-information benchmark*, so the proposed solution is calibrated to the problems.

5.1 Market with an Intermediary

Consider $N \geq 2$ human creators, an buyer, and an intermediary with a fixed operational cost $F \geq 0$. The intermediary *negotiates with both sides*: (1) To the creators, the buyer promises that if they jointly invest efforts $\boldsymbol{h} = (h_1, h_2, \ldots, h_N)'$, they will each receive a *payment* of $P_i(\boldsymbol{h})$. (2) To the buyer, the intermediary offers, on behalf of all creators, that when creators jointly invest efforts $\boldsymbol{h} = (h_1, h_2, \ldots, h_N)'$, the buyer must pay a *transfer* $T(\boldsymbol{h})$ to receive all signals $\boldsymbol{s}$.

That is, the intermediary determines two functions $\boldsymbol{P}\colon \mathbb{R}^N_{\geq 0} \to \mathbb{R}^N$ and $T\colon \mathbb{R}^N_{\geq 0} \to \mathbb{R}$. We call tuple

$\langle \boldsymbol{P}, T \rangle$ a *mechanism*. As in Section 3, we assume the effort levels $\boldsymbol{h}$ are observable and contractible. Analogous to Equation (3), the payoff of a creator i under the mechanism $\langle \boldsymbol{P}, T \rangle$ is specified as

$$U_i(\boldsymbol{h}; \boldsymbol{P}) := P_i(\boldsymbol{h}) - C(h_i), \quad \forall i = 1, 2, \ldots, N,$$

where $C(h_i) = \frac{c}{2} h_i^2$ is the production cost. Meanwhile, the profit of the buyer is given as

$$\Pi(\boldsymbol{h}; T) := K(\boldsymbol{h}) - T(\boldsymbol{h}),$$

where recall that $K(\boldsymbol{h})$ is the effective precision induced by effort vector $\boldsymbol{h}$.

To capture the participation constraints of each agent, we denote the buyer's outside option profit by $\underline{\Pi} \geq 0$ and each creator i's outside option payoff by $\underline{u}_i \geq 0$. We assume all outside options are publicly known. The intermediary seeks a mechanism $\langle \boldsymbol{P}, T \rangle$ satisfying:

1. **Incentive Compatibility (IC):** All creators invest the socially optimal efforts $\boldsymbol{h}^{\text{sp}}$:

$$h_i^{\text{sp}} \in \operatorname*{argmax}_{h_i \geq 0} U_i\big((h_i, \boldsymbol{h}_{-i}^{\text{sp}}); \boldsymbol{P}\big), \quad \forall i = 1, 2, \ldots, N. \tag{IC}$$

2. **Budget Balance (BB):** The intermediary breaks even (neither deficit nor surplus):[14]

$$T(\boldsymbol{h}) = F + \sum_{i=1}^{N} P_i(\boldsymbol{h}), \quad \forall \boldsymbol{h} \in \mathbb{R}_{\geq 0}^N. \tag{BB}$$

3. **Individual Rationality (IR):** At the equilibrium $\boldsymbol{h}^{\text{sp}}$, no agent (each creator or the buyer) is worse off than their outside option:

$$U_i(\boldsymbol{h}^{\text{sp}}; \boldsymbol{P}) \geq \underline{u}_i, \quad \forall i = 1, 2, \ldots, N; \qquad \text{and} \qquad \Pi(\boldsymbol{h}^{\text{sp}}; T) \geq \underline{\Pi}. \tag{IR}$$

We model the negotiation between the intermediary—on behalf of all creators—and the buyer using Nash bargaining. Let $\alpha \in [0, 1]$ denote the intermediary's bargaining power (and $(1 - \alpha)$ that of the

[14] Note that our Budget Balance (BB) differs from the global budget balance requirement analyzed by Holmstrom (1982), which is known to be impossible when also requiring Incentive Compatibility (IC). Here, creators' total income $T(\boldsymbol{h})$ needs not equal the social value $K(\boldsymbol{h})$; the difference $\Pi = K(\boldsymbol{h}) - T(\boldsymbol{h})$ is absorbed by the buyer, who acts as the "budget breaker" to extract the surplus or bear the deficit.

buyer), which is an exogenous parameter measuring the collective bargaining power of all creators. The *Nash bargaining solution* $\langle \boldsymbol{P}, T \rangle$ is selected to maximize the Nash product:

$$\max \underbrace{\left(\sum_{i=1}^{N} \left(U_i(\boldsymbol{h}^{\text{sp}}; \boldsymbol{P}) - \underline{u}_i \right) \right)^{\alpha}}_{\text{Creators' Net Surplus}} \underbrace{\left(\Pi(\boldsymbol{h}^{\text{sp}}; T) - \underline{\Pi} \right)^{1-\alpha}}_{\text{Buyer's Surplus}} \quad \text{subject to (IC), (IR), and (BB).} \tag{16}$$

5.2 Linear Pricing Fails on Sustainability

We first focus on the simplest class of mechanisms: linear pricing. Specifically, consider

$$P_i(\boldsymbol{h}) = p_i h_i, \quad \forall i = 1, 2, \ldots, N, \tag{17}$$

where $p_1, p_2, \ldots, p_N \geq 0$ are the unit prices for each creator. From Equation (BB), the corresponding transfer function $T(\boldsymbol{h})$ must be $T(\boldsymbol{h}) = \sum_{i=1}^{N} p_i h_i + F$. We begin by analyzing Equation (IC):

Lemma 6 (Marginal Prices and IC)**.** *Under linear pricing, the creators' best response coincides with the social optimum $\boldsymbol{h}^{sp}$ if and only if the unit prices are set as $p_i = \frac{\partial K}{\partial h_i}(\boldsymbol{h}^{sp}) = c h_i^{sp}$ for all i.*

Since the unit prices $\boldsymbol{p}$ are uniquely pinned down by Lemma 6, the linear pricing mechanism loses the flexibility to redistribute surplus. Therefore, it fails to satisfy the participation constraints in Equation (IR). Specifically, as shown in Proposition 7, the social surplus is almost fully used to incentivize creators' efficient production, leaving the buyer in deficit. So the intermediary, while protecting human creators' incentives, drives the AI developer out of the market. But there is nothing the intermediary can do if we insist on a linear pricing mechanism, as such $\langle \boldsymbol{P}, T \rangle$ is the unique one ensuring Equations (IC) and (BB). This *inflexibility* explains why markets failed in Sections 3 and 4: Under linear regimes, it is impossible to balance all parties' incentives.

Proposition 7 (Linear Pricing Fails on Sustainability)**.** *There exists a unique linear pricing mechanism $\langle \boldsymbol{P}, T \rangle$ satisfying Equations* (BB) *and* (IC)*. Under this mechanism, each creator i's payoff is $U_i(e^{sp}; \boldsymbol{P}) = \frac{c}{2}(h_i^{sp})^2$, while the buyer's profit is*

$$\Pi(\boldsymbol{h}^{sp}; T) = K(\boldsymbol{h}^{sp}) - \left(\nabla K(\boldsymbol{h}^{sp}) \right)' \boldsymbol{h}^{sp} - F.$$

Mechanism 1 Two-Part Tariff Mechanism

Input: Number of creators $N \geq 2$, production cost $c > 0$, parameters $\boldsymbol{\beta}, \gamma, \sigma_\eta^2$ in Equations (1) and (2), buyer's outside option $\underline{\Pi}$, creators' outside options $(\underline{u}_i)_{i=1}^N$, and Nash bargaining power $\alpha \in [0,1]$.

1: The intermediary calculates the social optimum $\boldsymbol{h}^{\text{sp}}$ maximizing social welfare W in Equation (6).

2: The intermediary calculates the social surplus at $\boldsymbol{h}^{\text{sp}}$ defined in Equation (19), namely $S(\boldsymbol{h}^{\text{sp}})$.

3: The intermediary decides the *lump-sum transfer* L from the buyer as:

$$L = K(\boldsymbol{h}^{\text{sp}}) - \sum_{i=1}^{N} ch_i^{\text{sp}} \times h_i^{\text{sp}} - (1-\alpha) S(\boldsymbol{h}^{\text{sp}}) - \underline{\Pi}. \tag{20}$$

4: The intermediary decides the *base subsidy* B_i for each creator $i = 1, 2, \ldots, N$ as:

$$B_i = \phi_i \alpha S(\boldsymbol{h}^{\text{sp}}) + \underline{u}_i - \frac{c}{2}(h_i^{\text{sp}})^2, \tag{21}$$

where $\boldsymbol{\phi} = (\phi_1, \phi_2, \ldots, \phi_N)'$ is a probability distribution (i.e., $\phi_i \geq 0$ and $\sum_{i=1}^N \phi_i = 1$) defined as

$$\phi_i \propto \int_0^1 \frac{\partial(\alpha S)}{\partial h_i}(t\boldsymbol{h}^{\text{sp}})\, \mathrm{d}t h_i^{\text{sp}} \qquad \text{s.t.} \quad \sum_{i=1}^{N} \phi_i = 1. \tag{22}$$

5: The intermediary commits to the *two-part tariff* mechanism $\langle \boldsymbol{P}, T \rangle$ specified by Equation (18).

In the special case where $\beta_i \equiv 0$, $\Pi(\boldsymbol{h}^{sp}; T) = -F < 0$, thus violating Equation (IR).

All formal proofs omitted in Section 5 can be found in Appendix C.

5.3 Optimal Design: A Two-Part Tariff Mechanism

Since linear pricing is pinned down by incentive constraints (Lemma 6) and fails to cover fixed costs, we must introduce an additional component to ensure Individual Rationality (IR).

It turns out to be sufficient to limit attention to *two-part tariff* mechanisms (Oi, 1971), where payment and transfer functions take the following affine forms:

$$P_i(\boldsymbol{h}) = ch_i^{\text{sp}} \times h_i + B_i, \ \forall i = 1, 2, \ldots, N; \quad T(\boldsymbol{h}) = \sum_{i=1}^{N} ch_i^{\text{sp}} \times h_i + L. \tag{18}$$

In Equation (18), B_i is a *base subsidy* (or fee) for creator i, and L is a *lump-sum licensing fee* paid by the buyer. We define the *net social surplus* $S(\boldsymbol{h})$ relative to outside options as:

$$S(\boldsymbol{h}) := \underbrace{K(\boldsymbol{h})}_{\text{Social Value}} - \underbrace{\left[F + \sum_{i=1}^{N} C(h_i)\right]}_{\text{Total Cost}} - \underbrace{\left[\underline{\Pi} + \sum_{i=1}^{N} \underline{u}_i\right]}_{\text{Outside Options}}, \quad \forall \boldsymbol{h} \in \mathbb{R}^N_{\geq 0}. \tag{19}$$

At the social optimum $\boldsymbol{h}^{\text{sp}}$, the surplus $S(\boldsymbol{h}^{\text{sp}})$ is divided via Nash bargaining. The intermediary

(and all creators) secures a share $\alpha \in [0,1]$, while the remaining $(1-\alpha)S(\boldsymbol{h}^{\mathrm{sp}})$ goes to the buyer. Thus, the transfer from the buyer to the intermediary, namely $T(\boldsymbol{h}^{\mathrm{sp}})$, must satisfy

$$K(\boldsymbol{h}^{\mathrm{sp}}) - T(\boldsymbol{h}^{\mathrm{sp}}) = \Pi(\boldsymbol{h}^{\mathrm{sp}}; T) = (1-\alpha)S(\boldsymbol{h}^{\mathrm{sp}}) + \underline{\Pi}.$$

This uniquely determines the lump-sum licensing fee L from the buyer as Equation (20).

To distribute the creators' share $\alpha S(\boldsymbol{h}^{\mathrm{sp}})$ fairly, we define weights $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_N)'$ inspired by the Aumann–Shapley value (Aumann and Shapley, 1974). Specifically, in Equation (22), each ϕ_i is proportional to creator i's cumulative contribution to the surplus along the diagonal path from $\mathbf{0}$ to $\boldsymbol{h}^{\mathrm{sp}}$. We set each creator i's utility to their fair share $\phi_i \alpha S(\boldsymbol{h}^{\mathrm{sp}})$ plus their outside option $\underline{u}_i$:

$$ch_i^{\mathrm{sp}} \times h_i^{\mathrm{sp}} + B_i - C(h_i^{\mathrm{sp}}) = U_i(h_i^{\mathrm{sp}}; P_i) = \phi_i \alpha S(\boldsymbol{h}^{\mathrm{sp}}) + \underline{u}_i.$$

This yields the required base subsidy B_i for each creator $i = 1, 2, \ldots, N$ as Equation (21).

We therefore obtain a two-part tariff mechanism $\langle \boldsymbol{P}, T \rangle$ satisfying Equations (IC), (IR), and (BB) at the same time. Moreover, due to the Nash bargaining procedure, it also maximizes the Nash product defined in Equation (16) among *all feasible mechanisms*, including those non-affine ones. We summarize our findings regarding Mechanism 1 in Proposition 8:

Proposition 8 (Two-Part Tariff Attains Efficiency and Sustainability)**.** *If $S(\boldsymbol{h}^{sp}) > 0$, i.e., the market is efficient, the two-part tariff mechanism $\langle \boldsymbol{P}, T \rangle$ defined in Mechanism 1 satisfies:*

1. ***Incentive Compatibility*** (IC)*: creators' best response is $\boldsymbol{h}^{sp}$.*
2. ***Global Budget Balance*** (BB)***:*** *the intermediary strictly breaks even.*
3. ***Individual Rationality*** (IR)***:*** *all participants earn at least their outside options.*

Furthermore, this $\langle \boldsymbol{P}, T \rangle$ maximizes the Nash product in Equation (16) *among* all *feasible mechanisms satisfying Equations* (IC)*,* (IR)*, and* (BB)*.*

Remark (Aumann–Shapley Fails on Incentives). One might ask: "Why not simply use the fair Aumann–Shapley value as the payment rule (i.e., $P_i^{\mathrm{AS}}(\boldsymbol{h}) = \int_0^1 \frac{\partial K}{\partial h_i}(t\boldsymbol{h})\mathrm{d}th_i$)?" While this rule

ensures budget balance ($\sum_{i=1}^{N} P_i^{\mathrm{AS}}(\boldsymbol{h}) = K(\boldsymbol{h})$), it fails to incentivize the most *efficient* effort level $\boldsymbol{h}^{\mathrm{sp}}$. In our model, the presence of the common factor η implies imitator's effort exhibits diminishing returns (or equivalently, $K(\boldsymbol{h})$ is *concave in* $\boldsymbol{h}$). Consequently, their *marginal* contribution $\frac{\partial K}{\partial h_i}$ is strictly smaller than their *average* contribution captured by P_i^{AS}. As summarized in Proposition 9, Aumann–Shapley pricing $\boldsymbol{P}^{\mathrm{AS}}$ thus incentivizes creators to *over-produce*:

Proposition 9 (Aumann–Shapley Fails on Incentives)**.** *While Aumann–Shapley pricing $\boldsymbol{P}^{AS}$ satisfies Incentive Compatibility* (IC) *for the special case of $\beta_i \equiv 0$ (i.e., the market only contains "innovators"), it fails in the general case. Specifically, for any correlated creator i (with $\beta_i > 0$), their best response to $\boldsymbol{P}^{AS}$ is* strictly higher *than the socially optimal level h_i^{sp}.*

As highlighted in Propositions 7 and 9, due to the concavity of information value, mechanisms that rely solely on a single payment rule inevitably face a trade-off between market efficiency and sustainability. Our proposed two-part tariff mechanism, on the other hand, *decouples* incentive design from surplus distribution: the unit price p_i solely targets *efficiency* (i.e., restoring the social optimum), while the fixed transfers manage *sustainability* by ensuring Equations (BB) and (IR).

6 Conclusion

We have analyzed the economics of data markets for AI. The fundamental market design challenge is enabling AI technology to progress while maintaining human creators' incentives to innovate. Unlike standard commodity markets where transaction costs serve as the main friction, we identify *statistical substitutability* arising from correlation between human creators' content as the primary barrier to an efficient bargain between AI firms and human creators. We identify a series of market failures driven by the interplay between correlation and AI firms' market power.

First, in a static market, we identify an *originality penalty* where the market equilibrium disproportionately penalizes "innovators" who produce uncorrelated data relative to "imitators" who are highly influenced by common trends and biases—despite the innovators' higher marginal utility for estimation. This results in misaligned creative incentives where the most creative content producers are most likely to drop out of the market. So a regime based on individualistic intellectual property rights not only underpowers creative incentives; it underpowers creative incentives most

for creators whose content is the most valuable.

Extending this insight to a dynamic model incorporating AI-assisted creation, we identify a *curse of precision*: As the AI model improves, it incentivizes the production of homogenized data, which creates an upper bound on market price. If all innovative creators consequently quit the market, the AI model's precision cannot increase above a fixed ceiling. This is a distinct insight from the "model collapse" phenomenon, identified in machine learning literature, which shows that a model recursively trained on self-generated data becomes irreversibly defective. We show, more pessimistically, that even a continuous inflow of fresh human data cannot halt model collapse in the absence of a well-designed market.

This finding underscores the imperative of designing a market that both preserves human creative incentives and promotes technological progress, showing that these objectives might be in tension with each other in the short term but are codependent in the long term. To meet this challenge, we propose a *two-part tariff* mechanism that implements the social optimum while ensuring budget balance and satisfying all agents' participation constraints. This relatively simple mechanism cures the market failures we have identified and achieves long-term sustainability for both model performance and human creativity.

Our proposed market design takes an important step toward solving a great policy challenge of our times, but taking it all the way to implementation requires addressing some remaining challenges. One outstanding question is the *compulsory or voluntary* character of participation in an intermediary organization: Once the intermediary is constructed, can its benefits be realized by content creators' uncoordinated, self-interested equilibrium behavior or do collective-action problems require centralized enforcement and a legal obligation to join (à la Olson (1965) and progeny)?[15] Another outstanding question relates to information gaps and asymmetries. Implementing our solution requires the intermediary to know human creators' marginal contribution, but estimating this value would be difficult if the market for creators' content outside the AI firm is nonexistent or thin (Hayek, 1945). Integrating our economic mechanism with data valuation techniques is hence a promising direction.

[15] From a legal perspective, it is important to note that the facilitation of collective action by means of intermediaries probably requires a statutory exemption from antitrust liability, and the aforementioned intermediated regimes under the Copyright Act (see footnote 12) generally contain such exemptions.

Appendix

A Proofs for Section 3

In this section, we first solve the market equilibrium resulting from Definition 1 to derive the first-order conditions specifying $\boldsymbol{h}^*$. Likewise, we also solve the first-order conditions for the first best solution $\boldsymbol{h}^{sp}$. We then prove Propositions 2 and 3 based on these conditions. Finally, in Lemma 12, we present an alternative robust minimax optimization view of our Bayesian model Equation (2).

A.1 First-Order Conditions

Lemma 10 (First-Order Conditions for $\boldsymbol{h}^*$ and $\boldsymbol{h}^{sp}$)**.** *Define*

$$\Lambda(\boldsymbol{h}) := \frac{\sum_{j=1}^N h_j\beta_j}{(\sigma_\eta^2+\gamma)^{-1}+\sum_{j=1}^N h_j\beta_j^2}. \tag{23}$$

Then the first-order conditions of the buyer and social planner give

$$h_i^* = \frac{1}{2c}\big(1-\beta_i\Lambda(\boldsymbol{h}^*)\big)^2, \quad h_i^{sp} = \frac{1}{c}\big(1-\beta_i\Lambda(\boldsymbol{h}^{sp})\big)^2, \quad \forall i=1,2,\ldots,N. \tag{24}$$

Furthermore, in any given market, both h_i^* and h_i^{sp} are strictly increasing in β_i.

Proof. For the buyer who maximizes profit Π (recall Equation (5)), namely

$$\Pi(\boldsymbol{h}) = K(\boldsymbol{h}) - \sum_{i=1}^N h_i p_i = K\left(\boldsymbol{h}\right) - \sum_{i=1}^N c p_i^2,$$

the first-order condition gives

$$\frac{\partial}{\partial h_i}K(\boldsymbol{h}^*) = 2ch_i^*, \quad \forall i=1,2,\ldots,N.$$

Differentiating the effective precision $K(\boldsymbol{h})$, whose calculation is deferred to Lemma 11, yields

$$\frac{\partial}{\partial h_i} K(\boldsymbol{h}) = \left(1 - \beta_i \frac{\sum_{j=1}^N h_j \beta_j}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{j=1}^N h_j \beta_j^2}\right)^2, \quad \forall i = 1, 2, \ldots, N. \tag{25}$$

Therefore, the equilibrium outcome $\boldsymbol{h}^* = (h_1^*, h_2^*, \ldots, h_N^*)$ must satisfy

$$1 - \beta_i \frac{\sum_{j=1}^N h_j^* \beta_j}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{j=1}^N h_j^* \beta_j^2} = \sqrt{2ch_i^*}, \quad \forall i = 1, 2, \ldots, N.$$

For the social planner who maximizes social welfare W (recall Equation (6)), namely

$$W(\boldsymbol{h}) = \Pi(\boldsymbol{h}) + \sum_{i=1}^N U_i(h_i; p_i(\boldsymbol{h})) = K(\boldsymbol{h}) - \sum_{i=1}^N \frac{c}{2} h_i^2,$$

the first-order condition gives

$$\frac{\partial}{\partial h_i} K(\boldsymbol{h}^{\mathrm{sp}}) = ch_i^{\mathrm{sp}}, \quad \forall i = 1, 2, \ldots, N. \tag{26}$$

Equivalently, this suggests

$$1 - \beta_i \frac{\sum_{j=1}^N h_j^{\mathrm{sp}} \beta_j}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{j=1}^N h_j^{\mathrm{sp}} \beta_j^2} = \sqrt{ch_i^{\mathrm{sp}}}, \quad \forall i = 1, 2, \ldots, N.$$

By definition of $\Lambda(\boldsymbol{h})$ in Equation (23), we derive

$$h_i^* = \frac{1}{2c}\big(1 - \beta_i \Lambda(\boldsymbol{h}^*)\big)^2, \quad h_i^{\mathrm{sp}} = \frac{1}{c}\big(1 - \beta_i \Lambda(\boldsymbol{h}^{\mathrm{sp}})\big)^2, \quad \forall i = 1, 2, \ldots, N,$$

therefore proving Equation (24). In any market with at least one positive β_i (the other case is trivial since every creator has the same β_i), $\Lambda(\boldsymbol{h}^*)$ and $\Lambda(\boldsymbol{h}^{\mathrm{sp}})$ are both strictly positive. Therefore, Equation (24) reveals h_i^* and h_i^{sp} are both strictly decreasing in β_i. □

Lemma 11 (Effective Precision)**.** *Under Equations* (1) *and* (2)*, the precision* $K(\boldsymbol{h})$ *induced by*

creators' effort vector $\boldsymbol{h} = (h_1, h_2, \ldots, h_N)'$ *is:*

$$K(\boldsymbol{h}) = \sum_{j=1}^{N} h_j - \frac{(\sum_{j=1}^{N} h_j \beta_j)^2}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{j=1}^{N} h_j \beta_j^2}.$$

Proof. The game proceeds as described in Definition 1. In the final stage of data aggregation, the buyer observes $\boldsymbol{h}$ and signals $\boldsymbol{s}$. From Equations (1) and (2), the signals are generated as

$$s_i = X + \beta_i \eta + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, h_i^{-1}), \quad \mu_\eta \sim \mathcal{N}(0, \gamma), \quad \eta \mid \mu_\eta \sim \mathcal{N}(\mu_\eta, \sigma_\eta^2).$$

The noise vector $\boldsymbol{\xi}$ has mean $\mathbf{0}$ and covariance matrix:

$$\boldsymbol{\Sigma} = \mathbb{E}[\boldsymbol{\xi}\boldsymbol{\xi}'] = \boldsymbol{D}^{-1} + (\sigma_\eta^2 + \gamma)\boldsymbol{\beta}\boldsymbol{\beta}',$$

where $\boldsymbol{D} = \mathrm{diag}(h_1, \ldots, h_N)$ and $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_N)'$. Given this Gaussian model, standard statistics (Aitken, 1936) gives the Best Linear Unbiased Estimator (BLUE) optimizing Equation (4) as

$$\hat{X} = (\mathbf{1}'\boldsymbol{\Sigma}^{-1}\mathbf{1})^{-1}\mathbf{1}'\boldsymbol{\Sigma}^{-1}\boldsymbol{s}, \quad \text{where } \boldsymbol{\Sigma} = \boldsymbol{D}^{-1} + (\sigma_\eta^2 + \gamma)\boldsymbol{\beta}\boldsymbol{\beta}'. \tag{27}$$

Using the Sherman-Morrison-Woodbury identity (Sherman and Morrison, 1950; Woodbury, 1950),

$$\boldsymbol{\Sigma}^{-1} = \left(\boldsymbol{D}^{-1} + \left((\sigma_\eta^2 + \gamma)\boldsymbol{\beta}\right)\boldsymbol{\beta}'\right)^{-1} = \boldsymbol{D} - \frac{(\sigma_\eta^2 + \gamma)\boldsymbol{D}\boldsymbol{\beta}\boldsymbol{\beta}'\boldsymbol{D}}{1 + (\sigma_\eta^2 + \gamma)\boldsymbol{\beta}'\boldsymbol{D}\boldsymbol{\beta}}.$$

Substituting this into the precision formula $K(\boldsymbol{h}) = \mathbf{1}'\boldsymbol{\Sigma}^{-1}\mathbf{1}$, we derive

$$K(\boldsymbol{h}) = \mathbf{1}'\left(\boldsymbol{D} - \frac{(\sigma_\eta^2 + \gamma)\boldsymbol{D}\boldsymbol{\beta}\boldsymbol{\beta}'\boldsymbol{D}}{1 + (\sigma_\eta^2 + \gamma)\boldsymbol{\beta}'\boldsymbol{D}\boldsymbol{\beta}}\right)\mathbf{1} = \sum_{i=1}^{N} h_i - \frac{(\sum_{i=1}^{N} h_i \beta_i)^2}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{i=1}^{N} h_i \beta_i^2}. \tag{28}$$

This completes the proof. □

A.2 Underpowered Creative Incentives (Proof of Proposition 2)

Proof of Proposition 2. To prove the uniqueness of $\boldsymbol{h}^*$ and $\boldsymbol{h}^{\mathrm{sp}}$, we claim $K(\boldsymbol{h})$ is concave. This is because the first term $\sum_{i=1}^{N} h_i$ is linear, and the second term $\left(\sum_{i=1}^{N} h_i \beta_i\right)^2 / \left((\sigma_\eta^2+\gamma)^{-1} + \sum_{i=1}^{N} h_i \beta_i^2\right)$,

a quadratic-over-linear function in $\boldsymbol{h}$, is also convex (Boyd and Vandenberghe, 2004, p. 73). Given the production cost $C(h_i) = \frac{c}{2}h_i^2$ is strictly convex, both profit Π and social welfare W, defined as

$$\Pi(\boldsymbol{h}) = K(\boldsymbol{h}) - \sum_{i=1}^{N} ch_i^2, \quad W(\boldsymbol{h}) = K(\boldsymbol{h}) - \sum_{i=1}^{N} \frac{c}{2}h_i^2$$

are strictly concave in $\boldsymbol{h}$. Therefore, the market equilibrium $\boldsymbol{h}^*$ maximizing Π and the social optimum $\boldsymbol{h}^{\text{sp}}$ maximizing W are both unique, as claimed.

We now compare $K(\boldsymbol{h}^*)$ to $K(\boldsymbol{h}^{\text{sp}})$. For any $y \geq 0$, we define the minimal total production cost required to attain an overall precision y as $L(y)$, namely

$$L(y) := \min_{\boldsymbol{h} \in \mathbb{R}^N_{\geq 0}} \frac{c}{2} \sum_{i=1}^{N} h_i^2 \text{ s.t. } K(\boldsymbol{h}) \geq y.$$

Since $K(\boldsymbol{h})$ is strictly concave and $\sum_{i=1}^{N} h_i^2$ is strictly convex, $L(y)$ is convex. We thus rewrite the buyer's and the social planner's optimization problems in w.r.t. Π and W as:

$$\begin{aligned}
\max_{\boldsymbol{h}} \Pi(\boldsymbol{h}) &= \max_{\boldsymbol{h}} \left(K(\boldsymbol{h}) - c \sum_{i=1}^{N} h_i^2 \right) = \max_{y} \left(y - 2L(y) \right), \\
\max_{\boldsymbol{h}} W(\boldsymbol{h}) &= \max_{\boldsymbol{h}} \left(K(\boldsymbol{h}) - \frac{c}{2} \sum_{i=1}^{N} h_i^2 \right) = \max_{y} \left(y - L(y) \right).
\end{aligned}$$

The first-order conditions reveal $L'(y^*) = \frac{1}{2}$ and $L'(y^{\text{sp}}) = 1$. Since L is strictly convex, L' is strictly increasing and thus $y^* < y^{\text{sp}}$. This implies $K(\boldsymbol{h}^*) = y^* < y^{\text{sp}} = K(\boldsymbol{h}^{\text{sp}})$. □

A.3 Originality Penalty (Proof of Proposition 3)

Proof of Proposition 3. To derive the monotonicity of R_i in β_i, we first prove that $\Lambda(\boldsymbol{h}^*) < \Lambda(\boldsymbol{h}^{\text{sp}})$ (where $\Lambda(\boldsymbol{h})$ is defined in Equation (23)). Plugging Equation (24) into the definition of $\Lambda(\boldsymbol{h}^*)$,

$$\begin{aligned}
&(\sigma_\eta^2 + \gamma)^{-1} \Lambda(\boldsymbol{h}^*) + \sum_{j=1}^{N} h_j^* \beta_j^2 \Lambda(\boldsymbol{h}^*) = \sum_{j=1}^{N} h_j^* \beta_j, \\
\implies &2c(\sigma_\eta^2 + \gamma)^{-1} \Lambda(\boldsymbol{h}^*) = \sum_{j=1}^{N} \left(1 - \beta_j \Lambda(\boldsymbol{h}^*)\right)^2 \beta_j \left(1 - \beta_j \Lambda(\boldsymbol{h}^*)\right) = \sum_{j=1}^{N} \beta_j \left(1 - \beta_j \Lambda(\boldsymbol{h}^*)\right)^3,
\end{aligned}$$

and similarly plugging Equation (24) into $\Lambda(\boldsymbol{h}^{\text{sp}})$ gives

$$c(\sigma_\eta^2+\gamma)^{-1}\Lambda(\boldsymbol{h}^{\text{sp}}) = \sum_{j=1}^{N}\beta_j\big(1-\beta_j\Lambda(\boldsymbol{h}^{\text{sp}})\big)^3.$$

Viewing both equations as a function of Λ, the LHS's are both linear in Λ. Meanwhile, the RHS is monotonically decreasing in $\Lambda > 0$ (recall that $\boldsymbol{\beta}$ is non-negative and at least one $\beta_i > 0$) since:

$$\frac{\mathrm{d}}{\mathrm{d}\Lambda}\sum_{j=1}^{N}\beta_j\big(1-\beta_j\Lambda\big)^3 = -3\sum_{j=1}^{N}\beta_j^2\big(1-\beta_j\Lambda\big)^2 < 0.$$

Therefore, we obtain $\Lambda(\boldsymbol{h}^*) < \Lambda(\boldsymbol{h}^{\text{sp}})$. Hence for each creator $i = 1, 2, \ldots, N$, the ratio

$$R_i = \frac{h_i^*}{h_i^{\text{sp}}} = \frac{\frac{1}{2c}\big(1-\beta_i\Lambda(\boldsymbol{h}^*)\big)^2}{\frac{1}{c}\big(1-\beta_i\Lambda(\boldsymbol{h}^{\text{sp}})\big)^2} = \frac{1}{2}\left(\frac{1-\beta_i\Lambda(\boldsymbol{h}^*)}{1-\beta_i\Lambda(\boldsymbol{h}^{\text{sp}})}\right)^2$$

is monotonically increasing in β_i because

$$\frac{\mathrm{d}}{\mathrm{d}\beta}\frac{1-\beta\Lambda(\boldsymbol{h}^*)}{1-\beta\Lambda(\boldsymbol{h}^{\text{sp}})} = \frac{\Lambda(\boldsymbol{h}^{\text{sp}})-\Lambda(\boldsymbol{h}^*)}{\big(1-\beta\Lambda(\boldsymbol{h}^{\text{sp}})\big)^2} > 0.$$

It only remains to bound R_i. Since $\beta_i \geq 0$ and R_i is monotone, it has a lower bound

$$R_i \geq \frac{1}{2}\left(\frac{1-0\Lambda(\boldsymbol{h}^*)}{1-0\Lambda(\boldsymbol{h}^{\text{sp}})}\right)^2 = \frac{1}{2}, \quad \text{with equality attained when } \beta_i = 0.$$

Furthermore, in the limiting case where $N \to \infty$ and all creators have the same $\beta_i \equiv \beta > 0$, each creator must have a symmetric e^* or e^{sp}. In this case, the $\Lambda(\boldsymbol{h})$ defined in Equation (23) becomes

$$\Lambda(e\mathbf{1}) = \frac{Ne\beta}{(\sigma_\eta^2+\gamma)^{-1}+Ne\beta^2}, \quad \forall e \geq 0.$$

Therefore Equation (24) gives $\boldsymbol{h}^* = e^*\mathbf{1}$ and $\boldsymbol{h}^{\text{sp}} = e^{\text{sp}}\mathbf{1}$, such that

$$e^* = \frac{1}{2c}\left(1-\frac{Ne^*\beta^2}{(\sigma_\eta^2+\gamma)^{-1}+Ne^*\beta^2}\right)^2, \quad e^{\text{sp}} = \frac{1}{c}\left(1-\frac{Ne^{\text{sp}}\beta^2}{(\sigma_\eta^2+\gamma)^{-1}+Ne^{\text{sp}}\beta^2}\right)^2.$$

When $N \to \infty$, the $(\sigma_\eta^2 + \gamma)^{-1}$ terms are dominated. Therefore,

$$\lim_{N\to\infty} e^* = \sqrt[3]{\frac{(\sigma_\eta^2+\gamma)^2}{2cN^2\beta^4}}, \quad \lim_{N\to\infty} e^{\text{sp}} = \sqrt[3]{\frac{(\sigma_\eta^2+\gamma)^2}{cN^2\beta^4}}.$$

This gives $\lim_{N\to\infty} R_i = \lim_{N\to\infty} \frac{e^*}{e^{\text{sp}}} = \frac{1}{\sqrt[3]{2}}$ for all $i = 1, 2, \ldots, N$, as claimed. □

A.4 Minimax Robust Optimization View of Equation (2)

Lemma 12 (Minimax Robust Optimization View)**.** *The effective precision $K(\boldsymbol{h})$ derived under the Bayesian setting in Equation* (2) *(where $\mu_\eta \sim \mathcal{N}(0, \gamma)$ and $\eta \mid \mu_\eta \sim \mathcal{N}(\mu_\eta, \sigma_\eta^2)$) is identical to that derived from a Minimax Robust Optimization framework where the buyer only knows $\mu_\eta^2 \le \gamma$. Formally, let the buyer seek a linear estimator $\hat{X} = \boldsymbol{w}'\boldsymbol{s}$ to minimize the worst-case MSE:*

$$\min_{\boldsymbol{w}} \max_{\mu_\eta^2 \le \gamma} \mathbb{E}[(\boldsymbol{w}'\boldsymbol{s} - X)^2], \quad s.t.\ \boldsymbol{w}'\mathbf{1} = 1. \tag{29}$$

The solution $\boldsymbol{w}^*$ is identical to BLUE weights in the Bayesian setting, and the MSE equals $K(\boldsymbol{h})^{-1}$.

Proof. We decompose the MSE of any linear estimator $\hat{X} = \boldsymbol{w}'\boldsymbol{s}$ into variance and bias:

$$\text{MSE}(\hat{X}) = \underbrace{\text{Var}(\hat{X})}_{\text{Variance}} + \underbrace{(\mathbb{E}[\hat{X}] - X)^2}_{\text{Squared Bias}}.$$

Under the robust optimization view of Equation (1), $\boldsymbol{s} = \mathbf{1}X + \boldsymbol{\beta}\eta + \tilde{\boldsymbol{\varepsilon}}$ where $\mathbb{E}[\eta] = \mu_\eta$ and $\mu_\eta^2 \le \gamma$. The variance depends only on the fluctuations around the mean, not on μ_η itself. Therefore, when letting $\boldsymbol{\Sigma}_{noise} = \boldsymbol{h}^{-1} + \sigma_\eta^2\boldsymbol{\beta}\boldsymbol{\beta}'$, we know $\text{Var}(\hat{X}) = \boldsymbol{w}'\boldsymbol{\Sigma}_{noise}\boldsymbol{w}$ is independent to μ_η. On the other hand, since $\boldsymbol{w}'\mathbf{1} = 1$, the mean of $\hat{X}$ is $\mathbb{E}[\hat{X}] = \boldsymbol{w}'(\mathbf{1}X + \boldsymbol{\beta}\mu_\eta) = X + \mu_\eta(\boldsymbol{w}'\boldsymbol{\beta})$. The squared bias term therefore equals $\mu_\eta^2(\boldsymbol{w}'\boldsymbol{\beta})^2$.

This reveals that the inner maximization w.r.t. μ_η in Equation (29) is equivalent to $\max_{\mu_\eta^2 \le \gamma} \mu_\eta^2(\boldsymbol{w}'\boldsymbol{\beta})^2$. The maximum is clearly attained at the boundary $\mu_\eta^2 = \gamma$. Thus, the buyer's problem in Equation (29) simplifies to:

$$\min_{\boldsymbol{w}} \left(\boldsymbol{w}'\boldsymbol{\Sigma}_{noise}\boldsymbol{w} + \gamma(\boldsymbol{w}'\boldsymbol{\beta})^2\right) \quad \text{s.t. } \boldsymbol{w}'\mathbf{1} = 1.$$

Notice that $\gamma(\boldsymbol{w}'\boldsymbol{\beta})^2 = \boldsymbol{w}'(\gamma\boldsymbol{\beta}\boldsymbol{\beta}')\boldsymbol{w}$. We can combine the quadratic terms:

$$\min_{\boldsymbol{w}} \boldsymbol{w}' \left(\boldsymbol{\Sigma}_{noise} + \gamma\boldsymbol{\beta}\boldsymbol{\beta}'\right) \boldsymbol{w} \quad \text{s.t. } \boldsymbol{w}'\mathbf{1} = 1.$$

Let $\tilde{\boldsymbol{\Sigma}} = \boldsymbol{\Sigma}_{noise} + \gamma\boldsymbol{\beta}\boldsymbol{\beta}' = \boldsymbol{h}^{-1} + (\sigma_\eta^2 + \gamma)\boldsymbol{\beta}\boldsymbol{\beta}'$. The Lagrangian for this problem is:

$$\mathcal{L}(\boldsymbol{w}, \lambda) = \frac{1}{2}\boldsymbol{w}'\tilde{\boldsymbol{\Sigma}}\boldsymbol{w} - \lambda(\boldsymbol{w}'\mathbf{1} - 1).$$

The First-Order Condition (FOC) with respect to $\boldsymbol{w}$ reveals $\tilde{\boldsymbol{\Sigma}}\boldsymbol{w}^* = \lambda\mathbf{1}$, or equivalently, $\boldsymbol{w}^* = \lambda\tilde{\boldsymbol{\Sigma}}^{-1}\mathbf{1}$. Since $\mathbf{1}'\boldsymbol{w}^* = 1$, we know $\lambda = (\mathbf{1}'\tilde{\boldsymbol{\Sigma}}^{-1}\mathbf{1})^{-1}$ and therefore:

$$\boldsymbol{w}^* = (\mathbf{1}'\tilde{\boldsymbol{\Sigma}}^{-1}\mathbf{1})^{-1}\tilde{\boldsymbol{\Sigma}}^{-1}\mathbf{1}.$$

Since the $\tilde{\boldsymbol{\Sigma}}$ here equals the covariance matrix $\boldsymbol{\Sigma}$ in the Bayesian setting (see the proof of Lemma 11), this is exactly the formula for BLUE weights in the Bayesian setting. Furthermore, the worst-case MSE induced by $\boldsymbol{w}^*$ is

$$\text{MSE}^* = \boldsymbol{w}^{*'}\tilde{\boldsymbol{\Sigma}}\boldsymbol{w}^* = \left(\lambda\mathbf{1}'\tilde{\boldsymbol{\Sigma}}^{-1}\right)\tilde{\boldsymbol{\Sigma}}\left(\lambda\tilde{\boldsymbol{\Sigma}}^{-1}\mathbf{1}\right) = (\mathbf{1}'\tilde{\boldsymbol{\Sigma}}^{-1}\mathbf{1})^{-1},$$

where we used $\lambda = (\mathbf{1}'\tilde{\boldsymbol{\Sigma}}^{-1}\mathbf{1})^{-1}$. In the Bayesian setting, we have $K(\boldsymbol{h}) = \mathbf{1}'\boldsymbol{\Sigma}^{-1}\mathbf{1}$; therefore, $\text{MSE}^* = K(\boldsymbol{h})^{-1}$, as claimed. □

B Proofs for Section 4

In this section, we first prove that the assumption on supply elasticity (Equation (7)) is equivalent to the monotonicity of imitator-innovator ratio $\rho(p)$. We then prove Propositions 4 and 5, which together constitute for our "curse of precision" observation.

B.1 Supply Elasticity and Monotonicity of Ratio

Lemma 13 (Supply Elasticity and Monotonicity of Ratio)**.** *The condition in Equation* (7)*, namely* $\epsilon_C(p) > \epsilon_O(p)$ *for all* $p > 0$ *where* $\epsilon_j(p) := \frac{\mathrm{d}\ln N_j(p)}{\mathrm{d}\ln p}$ $(\forall j \in \{O, C\})$*, is equivalent to*

$$\rho'(p) > 0,\ \forall p > 0, \quad \textit{where } \rho(p) := \frac{N_C(p)}{N_O(p)}.$$

Proof. The condition that $\rho'(p) > 0$ is equivalent to

$$0 < \frac{\mathrm{d}}{\mathrm{d}p}\frac{N_C(p)}{N_O(p)} = \frac{N_C'(p)N_O(p) - N_C(p)N_O'(p)}{\big(N_O(p)\big)^2},$$

which is equivalent to $\frac{N_C'(p)}{N_C(p)} > \frac{N_O'(p)}{N_O(p)}$. Multiplying p onto both sides, we get $\epsilon_C(p) > \epsilon_O(p)$. □

B.2 Upper Bound on Market Price (Proof of Proposition 4)

Proof of Proposition 4. According to the buyer's optimization Equation (12), at any t, the market price $p(t)$ maximizes the instantaneous profit $\Pi_{\text{inst}}(t)$ (we omit all dependencies on t for readability):

$$\begin{aligned}
\Pi_{\text{inst}} &= \frac{\mathrm{d}K}{\mathrm{d}t} - \Big(N_O(p) + N_C(p)\Big)pe(p) \\
&\overset{(a)}{=} \frac{\partial K}{\partial S_O}\frac{\mathrm{d}S_O}{\mathrm{d}t} + \frac{\partial K}{\partial S_C}\frac{\mathrm{d}S_C}{\mathrm{d}t} - \frac{p^2}{c}\Big(N_O(p) + N_C(p)\Big) \\
&\overset{(b)}{\le} 1\left(N_O\frac{p}{c} - \delta S_O\right) + i_C\left(N_C\frac{p}{c} - \delta S_C\right) - \frac{p^2}{c}(N_O + N_C) \\
&= \frac{p}{c}\Big(N_O(1-p) + N_C(i_C - p)\Big) - \delta\big(S_O + S_C\big),
\end{aligned}$$

where (a) uses the Law of Total Derivatives and the best effort formula $e(p) = p/c$, and (b) uses $\frac{\partial K}{\partial S_O} = 1$ and $\frac{\partial K}{\partial S_C} \le i_C(t)$ (both are derived using the Implicit Function Theorem in Lemma 14 which immediately succeeds this proof) as well as the dynamics of $S_C(t)$ and $S_O(t)$ in Equation (8).

Therefore, the instantaneous profit $\Pi_{\text{inst}}(t)$ decomposes into a variable component (depending on $p(t)$) and a sunk cost component (the $-\delta(S_O + S_C)$ term, which arises due to depreciation). For the buyer to sustain a price $p(t)$, the variable component must be non-negative (otherwise the buyer would be better off by setting $p(t) = 0$ and bearing only the sunk cost). Thus, the market price

$p(t)$ should satisfy (again omitting all t's for readability):

$$N_O(p)(1-p) + N_C(p)(i_C - p) \geq 0.$$

Rearranging for p gives $p \leq i_C + \frac{1-i_C}{1+\rho(p)}$, where we recall $\rho(p) = \frac{N_C(p)}{N_O(p)}$ is the imitator-innovator ratio. From Equation (7), $\rho(p)$ is strictly increasing in p, and thus the RHS is decreasing in p. Since the LHS is increasing, the fixed-point equation has a unique solution, namely $\bar{p}(i_C)$, such that

$$\bar{p}(i_C) = i_C + \frac{1-i_C}{1+\rho\big(\bar{p}(i_C)\big)}, \quad \forall i_C \in (0,1]. \tag{30}$$

Again invoking the Implicit Function Theorem to Equation (30) (whose detailed calculation is deferred to Lemma 15), we are able to prove $\bar{p}'(i_C) > 0$, i.e., $\bar{p}$ is increasing in i_C. This further establishes the continuity of $\bar{p}$, which gives

$$\lim_{i_C \to 0^+} \bar{p}(i_C) = \bar{p}(0), \quad \text{where } \bar{p}(0) = 0 + \frac{1-0}{1+\rho(\bar{p}(0))} = \left(1 + \frac{N_C(\bar{p}(0))}{N_O(\bar{p}(0))}\right).$$

This matches the definition of p^* in Equation (14). Hence, $\lim_{i_C \to 0^+} \bar{p}(i_C) = p^*$, as claimed. □

Lemma 14 (Time-Derivative of Effective Precision)**.** *The effective precision $K(t)$ defined in Equation* (11) *satisfies*

$$\dot{K}(t) \leq \dot{S_O}(t) + i_C(t)\dot{S_C}(t), \quad \forall t \geq 0,$$

where $i_C(t)$ is defined in Equation (13)*, which we recall as*

$$i_C(t) := \big(1 + \big(\sigma_\eta^2 + \lambda K(t)\big) S_C(t) \beta_C^2\big)^{-2}.$$

Proof. From Equation (11), the effective precision $K(t)$ is given by the implicit equation:

$$F(K, S_O, S_C) := K - S_O - \frac{S_C}{1 + \big(\sigma_\eta^2 + \lambda K\big) S_C \beta_C^2} = 0,$$

From the Chain Rule, we know $\dot{K} = \frac{\partial K}{\partial S_O}\dot{S_O} + \frac{\partial K}{\partial S_C}\dot{S_C}$ (omitting t for clarity). By the Implicit Function Theorem (Krantz and Parks, 2002), the partial derivatives of K w.r.t. stocks S_O and S_C

are therefore given by

$$\frac{\partial K}{\partial S_j} = -\frac{\partial F/\partial S_j}{\partial F/\partial K}, \quad j \in \{O, C\}.$$

Calculating the partial derivatives of F with respect to S_O, S_C, and K gives:

$$\begin{aligned}
\frac{\partial F}{\partial S_O} &= -1, \\
\frac{\partial F}{\partial S_C} &= -\frac{1 + \big(\sigma_\eta^2 + \lambda K\big) S_C \beta_C^2 - S_C \cdot \big(\sigma_\eta^2 + \lambda K\big)\beta_C^2}{\big(1 + \big(\sigma_\eta^2 + \lambda K\big) S_C \beta_C^2\big)^2} = -\frac{1}{\big(1 + \big(\sigma_\eta^2 + \lambda K\big) S_C \beta_C^2\big)^2}, \\
\frac{\partial F}{\partial K} &= 1 - S_C \cdot \frac{\partial}{\partial K}\left(\frac{1}{1 + \big(\sigma_\eta^2 + \lambda K\big) S_C \beta_C^2}\right).
\end{aligned}$$

As $(1 + (\sigma_\eta^2 + \lambda K) S_C \beta_C^2)^{-1}$ is strictly decreasing in K, its derivative is negative. Hence $\frac{\partial F}{\partial K} > 1$, and

$$\frac{\partial K}{\partial S_O} = \frac{1}{\partial F/\partial K} < 1, \quad \frac{\partial K}{\partial S_C} = \frac{1}{\big(1 + \big(\sigma_\eta^2 + \lambda K\big) S_C \beta_C^2\big)^2} \frac{1}{\partial F/\partial K} < \frac{1}{\big(1 + \big(\sigma_\eta^2 + \lambda K\big) S_C \beta_C^2\big)^2}.$$

Plugging in the definition of $i_C(t)$ and recalling $\dot{K}(t) = \frac{\partial K}{\partial S_O}\dot{S_O} + \frac{\partial K}{\partial S_C}\dot{S_C}$ gives the conclusion. □

Lemma 15 (Monotonicity of Price Upper Bound)**.** *The price upper bound $\bar{p}(i_C)$ derived in Equation* (30)*, namely*

$$\bar{p}(i_C) = i_C + \frac{1 - i_C}{1 + \rho\big(\bar{p}(i_C)\big)}, \quad \textit{where } \rho(p) = \frac{N_C(p)}{N_O(p)},$$

is strictly increasing in i_C, i.e., $\bar{p}'(i_C) > 0$ for all $i_C \in (0, 1]$.

Proof. From Equation (30), the price upper bound $\bar{p}(i_C)$ is given by the implicit equation:

$$G(\bar{p}, i_C) := \bar{p} - i_C - \frac{1 - i_C}{1 + \rho(\bar{p})} = 0.$$

By the Implicit Function Theorem (Krantz and Parks, 2002), we have

$$\frac{\mathrm{d}\bar{p}}{\mathrm{d}i_C} = -\frac{\partial G/\partial i_C}{\partial G/\partial \bar{p}} = \frac{1 - \frac{1}{1+\rho(\bar{p})}}{1 + \frac{1-i_C}{(1+\rho(\bar{p}))^2}\rho'(\bar{p})}.$$

The numerator is positive because $\rho(p) > 0$ for all p, and the denominator is also positive because $i_C \in (0, 1]$ and also $\rho'(p) > 0$ for all p (from Equation (7) and Lemma 13). Therefore, $\bar{p}'(i_C) > 0$

for all $i_C \in (0, 1]$, as claimed. □

B.3 Barrier of Model Precision (Proof of Proposition 5)

Proof of Proposition 5. Recall the fixed-point definition for precision $K(t)$ in Equation (11), namely $K(t) = S_O(t) + \frac{S_C(t)}{1+(\sigma_\eta^2+\lambda K(t))S_C(t)\beta_C^2}$. When innovators have exited ($S_O(t) \to 0$), this equation becomes

$$K(t) = \frac{S_C(t)}{1 + \big(\sigma_\eta^2 + \lambda K(t)\big) S_C(t) \beta_C^2}.$$

In the limit of infinite correlated data (i.e., $S_C(t) \to \infty$), we have

$$1 = \lim_{S_C(t)\to\infty} K(t)\Big(S_C(t)^{-1} + \big(\sigma_\eta^2 + \lambda K(t)\big)\beta_C^2\Big) = \lim_{S_C(t)\to\infty} K(t)\big(\sigma_\eta^2 + \lambda K(t)\big)\beta_C^2.$$

Therefore, the limiting precision $\bar{K}$ must satisfy

$$\lambda\beta_C^2\big(\bar{K}\big)^2 + \sigma_\eta^2\beta_C^2\bar{K} - 1 = 0.$$

Solving for the unique positive root yields the expression in Equation (15). □

C Proofs for Section 5

In this section, we prove Propositions 7, 8, and 9.

C.1 Linear Pricing Fails on Sustainability (Proof of Proposition 7)

Proof of Lemma 6. For each creator $i = 1, 2, \ldots, N$, under the linear pricing rule (or the two-part tariff rule considered later in Proposition 8), their payoff $U_i(\boldsymbol{h}; \boldsymbol{P})$ only depends on h_i and p_i (and also B_i if considering the two-part tariff rule defined in Equation (18)), namely $U_i(\boldsymbol{h}; \boldsymbol{P}) = p_i h_i + B_i - C(h_i)$. Thus the first-order condition is simply

$$p_i = c h_i.$$

Therefore, the best response is h_i^{sp}—thus satisfying Incentive Compatibility (IC)—if and only if $p_i = ch_i^{\text{sp}}$. This gives the only linear $\boldsymbol{P}$ attaining IC, i.e., $P_i(\boldsymbol{h}) = ch_i^{\text{sp}} \times h_i$, $\forall i = 1, 2, \ldots, N$. $\square$

Proof of Proposition 7. According to Lemma 6, the $\boldsymbol{P}$ is uniquely determined by Equation (IC). We further pin down T using Budget Balance (BB). This gives

$$P_i(\boldsymbol{h}) = ch_i^{\text{sp}} \times h_i, \ \forall i = 1, 2, \ldots, N; \quad T(\boldsymbol{h}) = F + \sum_{i=1}^{N} p_i h_i.$$

This gives the unique $\langle \boldsymbol{P}, T \rangle$ ensuring Equations (IC) and (BB). We write the payoff of each creator:

$$U_i(h_i^{\text{sp}}; P_i) = p_i h_i^{\text{sp}} - \frac{c}{2}(h_i^{\text{sp}})^2 = \frac{c}{2}(h_i^{\text{sp}})^2,$$

and also the profit of the buyer:

$$\Pi(\boldsymbol{h}^{\text{sp}}; T) = K(\boldsymbol{h}^{\text{sp}}) - F - \sum_{i=1}^{N} p_i h_i^{\text{sp}}.$$

Since $p_i = ch_i^{\text{sp}}$ (from Lemma 6) and the social planner solution ensures (recall Equation (26)) $\nabla K(\boldsymbol{h}^{\text{sp}}) = c\boldsymbol{h}^{\text{sp}}$, we have

$$\Pi(\boldsymbol{h}^{\text{sp}}; T) = K(\boldsymbol{h}^{\text{sp}}) - F - (c\boldsymbol{h}^{\text{sp}})'\boldsymbol{h}^{\text{sp}} = K(\boldsymbol{h}^{\text{sp}}) - \big(\nabla K(\boldsymbol{h}^{\text{sp}})\big)'\boldsymbol{h}^{\text{sp}} - F.$$

In the special case where all creators are uncorrelated (i.e., $\beta_i = 0$ for all i), we have $K(\boldsymbol{h}) = \sum_{i=1}^{N} h_i$ (recall Lemma 11).[16] Euler's Homogeneous Function Theorem gives

$$K(\boldsymbol{h}) = \sum_{i=1}^{N} h_i \cdot \frac{\partial K}{\partial h_i}(\boldsymbol{h}) = (\nabla K(\boldsymbol{h}))'\boldsymbol{h}, \quad \forall \boldsymbol{h} \in \mathbb{R}^N_{\geq 0}. \tag{31}$$

Therefore, the buyer's profit $\Pi = -F < 0 \leq \underline{\Pi}$ in this case. $\square$

[16] Or more generally, in any *Constant Returns to Scale* (CRS) regime where $K(\alpha\boldsymbol{h}) = \alpha K(\boldsymbol{h})$ for any scalar $\alpha > 0$.

C.2 Two-Part Tariff for Efficiency and Sustainability (Proof of Proposition 8)

Proof of Proposition 8. From the specific choice of unit price $\boldsymbol{p} = c\boldsymbol{h}^{\text{sp}}$, Incentive Compatibility (IC) holds (recall the proof of Lemma 7). Since $\boldsymbol{\phi} = (\phi_1, \phi_2, \ldots, \phi_N)'$ sums to 1, at the equilibrium effort $\boldsymbol{h}^{\text{sp}}$ the total payment to the creators is (recall the B_i defined in Equation (21)):

$$\begin{aligned}
&\sum_{i=1}^{N} P_i(\boldsymbol{h}^{\text{sp}}) = \sum_{i=1}^{N} p_i h_i^{\text{sp}} + B_i \\
&= \sum_{i=1}^{N} \left[c h_i^{\text{sp}} \times h_i^{\text{sp}} + \phi_i \alpha S(\boldsymbol{h}^{\text{sp}}) + \underline{u}_i - \frac{c}{2}(h_i^{\text{sp}})^2 \right] \\
&= \alpha S(\boldsymbol{h}^{\text{sp}}) + \sum_{i=1}^{N} \left[\underline{u}_i + \frac{c}{2}(h_i^{\text{sp}})^2 \right].
\end{aligned}$$

Meanwhile, the transfer received from the buyer (recall the L defined in Equation (20)) ensures

$$T(\boldsymbol{h}^{\text{sp}}) = K(\boldsymbol{h}^{\text{sp}}) - (1-\alpha) S(\boldsymbol{h}^{\text{sp}}) - \underline{\Pi}.$$

The surplus of the intermediary (income minus operational cost minus payments) then satisfies

$$T(\boldsymbol{h}^{\text{sp}}) - F - \sum_{i=1}^{N} P_i(\boldsymbol{h}^{\text{sp}}) = K(\boldsymbol{h}^{\text{sp}}) - (1-\alpha) \cdot S(\boldsymbol{h}^{\text{sp}}) - \underline{\Pi} - F - \alpha S(\boldsymbol{h}^{\text{sp}}) - \sum_{i=1}^{N} [\underline{u}_i + C(h_i^{\text{sp}})] = 0,$$

by definition of $S(\boldsymbol{h})$. This proves Budget Balance (BB) condition. We finally verify the Individual Rationality (IR) condition. For each creator $i = 1, 2, \ldots, N$:

$$\begin{aligned}
U_i(h_i^{\text{sp}}; P_i) = P_i(\boldsymbol{h}^{\text{sp}}) - C(h_i^{\text{sp}}) &= \left[c h_i^{\text{sp}} \times h_i^{\text{sp}} + \phi_i \alpha S(\boldsymbol{h}^{\text{sp}}) + \underline{u}_i - \frac{c}{2}(h_i^{\text{sp}})^2 \right] - C(h_i^{\text{sp}}) \\
&= \underline{u}_i + \phi_i \alpha S(\boldsymbol{h}^{\text{sp}}) \geq \underline{u}_i,
\end{aligned}$$

where the last step is because $S(\boldsymbol{h}^{\text{sp}}) > 0$ (the condition of Proposition 8). The buyer's profit:

$$\begin{aligned}
\Pi(\boldsymbol{h}^{\text{sp}}; T) = K(\boldsymbol{h}^{\text{sp}}) - T(\boldsymbol{h}^{\text{sp}}) &= K(\boldsymbol{h}^{\text{sp}}) - \big(K(\boldsymbol{h}^{\text{sp}}) - (1-\alpha) S(\boldsymbol{h}^{\text{sp}}) - \underline{\Pi} \big) \\
&= \underline{\Pi} + (1-\alpha) S(\boldsymbol{h}^{\text{sp}}) \geq \underline{\Pi},
\end{aligned}$$

where we again used the condition that $S(\boldsymbol{h}^{\text{sp}}) > 0$. □

C.3 Aumann–Shapley Fails on Incentives (Proof of Proposition 9)

Proof of Proposition 9. In the special case where $\beta_i \equiv 0$ for all $i = 1, 2, \ldots, N$, the effective precision function $K(\boldsymbol{h})$ is simply $K(\boldsymbol{h}) = \sum_{i=1}^{N} h_i$ (recall Lemma 11). It is homogeneous of degree one, i.e., $K(\alpha \boldsymbol{h}) = \alpha K(e)$ and $\nabla K(\alpha \boldsymbol{h}) = \nabla K(\boldsymbol{h})$ for any $\alpha > 0$. Therefore,

$$P_i^{AS}(\boldsymbol{h}) = \int_0^1 \frac{\partial K}{\partial h_i}(t\boldsymbol{h})\ \mathrm{d}t \cdot h_i = h_i \cdot \frac{\partial K}{\partial h_i}(\boldsymbol{h}), \quad \forall i, 2, \ldots, N. \tag{32}$$

Substituting Equation (32) into the creator's utility function, creator i maximizes $U_i(h_i) = h_i \cdot \frac{\partial K}{\partial h_i}(\boldsymbol{h}) - \frac{c}{2} h_i^2$. The first-order condition is $c h_i = \frac{\partial K}{\partial h_i}(\boldsymbol{h})$. Since this coincides exactly with the social planner's first-order condition (recall Equation (26)), the private optimum coincides with the social optimum $\boldsymbol{h}^{\mathrm{sp}}$. Thus, Aumann–Shapley pricing ensures Equation (IC) in the absence of correlation.

However, in general, the effective precision $K(\boldsymbol{h})$ exhibits *Decreasing Returns to Scale* (DRS). Specifically, from Equation (25), for any scalar $0 < t < 1$ and efforts vector $\boldsymbol{h}$, we have

$$\begin{aligned}\frac{\partial K}{\partial h_i}(t\boldsymbol{h}) &= \left(1 - \beta_i \frac{\sum_{j=1}^{N} (t \cdot h_j)\beta_j}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{j=1}^{N} (t \cdot h_j)\beta_j^2}\right)^2 \\ &> \left(1 - \beta_i \frac{\sum_{j=1}^{N} h_j \beta_j}{(\sigma_\eta^2 + \gamma)^{-1} + \sum_{j=1}^{N} h_j \beta_j^2}\right)^2 = \frac{\partial K}{\partial h_i}(\boldsymbol{h}).\end{aligned}$$

Consequently, the Aumann–Shapley pricing $\boldsymbol{P}^{\mathrm{AS}}$ satisfies

$$P_i^{AS}(\boldsymbol{h}^{\mathrm{sp}}) = \int_0^1 \frac{\partial K}{\partial h_i}(t\boldsymbol{h}^{\mathrm{sp}})\, \mathrm{d}t \cdot h_i^{\mathrm{sp}} > \frac{\partial K}{\partial h_i}(\boldsymbol{h}^{\mathrm{sp}}) \cdot h_i^{\mathrm{sp}}.$$

This implies that the *average* marginal return faced by the creator strictly exceeds the *true* marginal social value at the optimum $\boldsymbol{h}^{\mathrm{sp}}$. Since the creator equates their marginal cost $C'(h_i)$ to this inflated price, and since the cost function $C(\cdot)$ is convex, the creator will rationally choose an effort level $h_i^* > h_i^{\mathrm{sp}}$. This leads to systematic over-production and efficiency loss. □

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar, "Too Much Data: Prices and Inefficiencies in Data Markets," *American Economic Journal: Microeconomics*, 2022, *14* (4), 218–256.

Agarwal, Anish, Munther Dahleh, and Tuhin Sarkar, "A Marketplace for Data: An Algorithmic Solution," in "Proceedings of the 2019 ACM Conference on Economics and Computation" 2019, pp. 701–726.

Aitken, Alexander C, "IV.—On Least Squares and Linear Combination of Observations," *Proceedings of the Royal Society of Edinburgh*, 1936, *55*, 42–48.

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

Alemohammad, Sina, Josue Casco-Rodriguez, Lorenzo Luzi, Ahmed Imtiaz Humayun, Hossein Babaei, Daniel LeJeune, Ali Siahkoohi, and Richard Baraniuk, "Self-Consuming Generative Models Go Mad," in "The Twelfth International Conference on Learning Representations" 2023.

Arrow, Kenneth, "Economic Welfare and the Allocation of Resources for Invention," *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 1962, pp. 609–626.

Aumann, Robert J and Lloyd S Shapley, *Values of Non-Atomic Games*, Princeton University Press, 1974.

Barnett, Jonathan M., "The Free Content Illusion," *Journal of Intellectual Property Law*, 2026, *33* (1).

Bergemann, Dirk, Alessandro Bonatti, and Tan Gan, "The Economics of Social Data," *The RAND Journal of Economics*, 2022, *53* (2), 263–296.

Berti, Leonardo, Flavio Giorgi, and Gjergji Kasneci, "Emergent abilities in large language models: A survey," *arXiv preprint arXiv:2503.05788*, 2025.

Boyd, Stephen and Lieven Vandenberghe, *Convex optimization*, Cambridge university press, 2004.

Che, Yeon-Koo, "Beyond the Coasian Irrelevance: Asymmetric Information," *Unpublished notes, Columbia University.[110, 114]*, 2006.

Coase, Ronald H, "The Problem of Social Cost," *Journal of Law and Economics*, 1960, *3*, 1–44.

Cotter, Thomas F, "Some Observations on the Law and Economics of Intermediaries," *Michigan State Law Review*, 2005, pp. 67–82.

Farboodi, Maryam and Laura Veldkamp, "Long-Run Growth of Financial Data Technology," *American Economic Review*, 2020, *110* (8), 2485–2523.

_ and _ , "A Model of the Data Economy," *NBER Working Paper*, 2021, (w28427).

_ and _ , "Data and Markets," *Annual Review of Economics*, 2023, *15* (1), 23–40.

_ , Gregor Jarosch, and Robert Shimer, "The Emergence of Market Structure," *The Review of Economic Studies*, 2023, *90* (1), 261–292.

Freilich, Janet and Sepehr Shahshahani, "Measuring Follow-On Innovation," *Research Policy*, 2023, *52*, 104854.

Ghorbani, Amirata and James Zou, "Data Shapley: Equitable Valuation of Data for Machine Learning," in "International Conference on Machine Learning" PMLR 2019, pp. 2242–2251.

Gilbert, Richard, "Collective Rights Organizations: A Guide to Benefits, Costs and Antitrust Safeguards," *The Cambridge Handbook of Technical Standardization Law*, 2017, *1.*

Glazer, Amihai and Refael Hassin, "Optimal Contests," *Economic Inquiry*, 1988, *26* (1), 133–143.

Golrezaei, Negin and Hamid Nazerzadeh, "Pricing schemes for metropolitan traffic data markets," in "International Conference on Data Management Technologies and Applications," Vol. 2 SCITEPRESS 2014, pp. 266–271.

_ , **MohammadTaghi Hajiaghayi, and Suho Shin**, "The Contest Behind the Feed: Optimal Contest for Recommender Systems," *Available at SSRN 5258940*, 2025.

Gordon, Wendy J., "Fair Use as Market Failure: A Structural and Economic Analysis of the Betamax Case and its Predecessors," *Columbia Law Review*, 1982, *82*, 1600–1657.

Hayek, Friedrich A., "The Use of Knowledge in Society," *American Economic Review*, 1945, *35* (4), 519–530.

Heller, Michael A, "The Tragedy of the Anticommons: Property in the Transition from Marx to Markets," *Harvard Law Review*, 1998, *111* (3), 621–688.

Heller, Michael A. and Rebecca S. Eisenberg, "Can Patents Deter Innovation? The Anticommons in Biomedical Research," *Science*, 1998, *280* (5364), 698–701.

Henderson, Peter, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A. Lemley, and Percy Liang, "Foundation Models and Fair Use," *Journal of Machine Learning Research*, 2023, *24*, 1–79.

Hoffmann, Jordan, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark et al., "Training Compute-Optimal Large Language Models," in "Proceedings of the 36th International Conference on Neural Information Processing Systems" 2022, pp. 30016–30030.

Holmstrom, Bengt, "Moral Hazard in Teams," *The Bell Journal of Economics*, 1982, pp. 324–340.

Immorlica, Nicole, Meena Jagadeesan, and Brendan Lucier, "Clickbait vs. Quality: How Engagement-Based Optimization Shapes the Content Landscape in Online Platforms," in "Proceedings of the ACM Web Conference 2024" 2024, pp. 36–45.

Jagadeesan, Meena, Michael I Jordan, and Nika Haghtalab, "Competition, Alignment, and Equilibria in Digital Marketplaces," in "Proceedings of the AAAI Conference on Artificial Intelligence," Vol. 37 2023, pp. 5689–5696.

_ , **Nikhil Garg, and Jacob Steinhardt**, "Supply-Side Equilibria in Recommender Systems," *Advances in Neural Information Processing Systems*, 2023, *36*, 14597–14608.

Jia, Ruoxi, David Dao, Boxin Wang, Frances Ann Hubis, Nick Hynes, Nezihe Merve Gürel, Bo Li, Ce Zhang, Dawn Song, and Costas J Spanos, "Towards Efficient Data Valuation Based on the Shapley Value," in "The 22nd International Conference on Artificial Intelligence and Statistics" PMLR 2019, pp. 1167–1176.

Kaplan, Jared, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei, "Scaling Laws for Neural Language Models," *arXiv preprint arXiv:2001.08361*, 2020.

Koh, Pang Wei and Percy Liang, "Understanding Black-Box Predictions via Influence Functions," in "International Conference on Machine Learning" PMLR 2017, pp. 1885–1894.

Krantz, Steven G and Harold R Parks, *The Implicit Function Theorem: History, Theory, and Applications*, Springer Science & Business Media, 2002.

Lee, Edward, "Fair Use and the Origin of AI Training," *Houston Law Review*, 2025, *63*, 105–229.

Lemley, Mark A., "How Generative AI Turns Copyright Law Upside Down," *Columbia Science & Technology Law Review*, 2024, *25* (2), 190–212.

_ **and Bryan Casey**, "Fair Learning," *Texas Law Review*, 2021, *99*, 743–785.

_ **and Lisa Larrimore Ouellette**, "Plagiarism, Copyright, and AI," *University of Chicago Law Review Online*, 2025, *2025.*

Leval, Pierre N., "Toward a Fair Use Standard," *Harvard Law Review*, 1990, *103* (5), 1105–1136.

Machlup, Fritz, "An Economic Review of the Patent System," Committee Print Study No. 15, U.S. Senate Subcommittee on Patents, Trademarks, and Copyrights, Washington 1958. 85th Congress, 2d session.

Muschalle, Alexander, Florian Stahl, Alexander Löser, and Gottfried Vossen, "Pricing Approaches for Data Markets," in "Enabling Real-Time Business Intelligence: 6th International

Workshop, BIRTE 2012, Held at the 38th International Conference on Very Large Databases, VLDB 2012, Istanbul, Turkey, August 27, 2012, Revised Selected Papers," Vol. 154 Springer 2013, p. 129.

Nordhaus, William, *Invention, Growth, and Welfare: A Thoeretical Treatment of Technological Change*, MIT Press, 1969.

Oi, Walter Y, "A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly," *The Quarterly Journal of Economics*, 1971, *85* (1), 77–96.

Olson, Mancur, *The Logic of Collective Action: Public Goods and the Theory of Groups*, Harvard University Press, 1965.

Opderbeck, David W., "Copyright in AI Training Data," *Oklahoma Law Review*, 2024, *76*, 951–1023.

Ostrom, Elinor, *Governing the Commons: The Evolution of Institutions for Collective Action*, Cambridge University Press, 1990.

Plant, Arnold, "The Economic Theory Concerning Patents for Inventions," *Economica*, 1934, *1* (1), 30–51.

Register of Copyrights, *Copyright and Artificial Intelligence – Part 3: Generative AI Training*, U.S. Copyright Office, 2025. Pre-Publication Version.

Rubinstein, Ariel and Asher Wolinsky, "Middlemen," *The Quarterly Journal of Economics*, 1987, *102* (3), 581–593.

Sag, Matthew, "The New Legal Landscape for Text Mining and Machine Learning," *Journal of the Copyright Society of the USA*, 2019, *66*, 291.

Samuelson, Pamela, "Generative AI Meets Copyright," *Science*, 2023, *381* (6654), 158–161.

_ , "Fair Use Defenses in Disruptive Technology Cases," *UCLA Law Review*, 2024, *71*, 1484–1572.

_ , "Does Using In-Copyright Works as Training Data Infringe?," *Communications of the ACM*, 2025, *68* (11), 28–30.

Schaeffer, Rylan, Brando Miranda, and Sanmi Koyejo, "Are Emergent Abilities of Large Language Models a Mirage?," *Advances in neural information processing systems*, 2023, *36*, 55565–55581.

Schomm, Fabian, Florian Stahl, and Gottfried Vossen, "Marketplaces for Data: An Initial Survey," *SIGMOD Record*, 2013, *42* (1), 15.

Segal, Ilya and Michael D Whinston, "Property Rights and the Efficiency of Bargaining," *Journal of the European Economic Association*, 2016, *14* (6), 1287–1328.

Shahshahani, Sepehr, "The Nirvana Fallacy in Fair Use Reform," *Minnesota Journal of Law, Science & Technology*, 2015, *16* (1), 273–341.

_ , "The Role of Courts in Technology Policy," *Journal of Law and Economics*, 2018, *61* (1), 37–64.

_ , "Against the Abstract-Ideas Exclusion," *Berkeley Technology Law Journal*, 2025, *40*, 447–518.

Sherman, Jack and Winifred J Morrison, "Adjustment of an Inverse Matrix Corresponding to a Change in One Element of a Given Matrix," *The Annals of Mathematical Statistics*, 1950, *21* (1), 124–127.

Shumailov, Ilia, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson, "The Curse of Recursion: Training on Generated Data Makes Models Forget," *Nature*, 2024, *631*, 755–759.

Tullis, Jane, "Sifting Through the Slop: How Generative AI Created a Market for Lemons for Text-Based Works," 2025.

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

Wei, Jason, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus, "Emergent Abilities of Large Language Models," *Transactions on Machine Learning Research*, 2022.

Widmer, Gerhard and Miroslav Kubat, "Learning in the Presence of Concept Drift and Hidden Contexts," *Machine learning*, 1996, *23* (1), 69–101.

Williamson, Oliver E, "Transaction-Cost Economics: The Governance of Contractual Relations," *Journal of Law and Economics*, 1979, *22* (2), 233–261.

Woodbury, Max A, *Inverting Modified Matrices*, Department of Statistics, Princeton University, 1950.

Yao, Fan, Chuanhao Li, Denis Nekipelov, Hongning Wang, and Haifeng Xu, "How Bad is Top-K Recommendation under Competing Content Creators?," in "International Conference on Machine Learning" PMLR 2023, pp. 39674–39701.

_ , _ , Karthik Abinav Sankararaman, Yiming Liao, Yan Zhu, Qifan Wang, Hongning Wang, and Haifeng Xu, "Rethinking Incentives in Recommender Systems: Are Monotone Rewards Always Beneficial?," *Advances in Neural Information Processing Systems*, 2023, *36*, 74582–74601.