

Market Design for AI: Beyond the Copyright Binary

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

First version: Feb 2026. This version: Jul 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: Free-for-all does not compensate creators, and—by modeling as a static Stackelberg game—strong intellectual property rights also underpower creative incentives. We find this especially true for more innovative creators, a phenomenon we term the “originality penalty.” Extending this insight to a dynamic model, we find another market failure undermining AI model performance, even for an initially good model: Such a model induces greater reliance by humans on AI-assisted creation, resulting in homogenized content feeding back into training, which degrades the model performance—a “curse of precision.” We further propose a market design with a data intermediary negotiating collectively with the AI firm and subsidizing innovative contributions, thus restoring efficiency.

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 (Burtch et al., 2024; del Rio-Chanona et al., 2024; Sharma and Li, 2025). A recent field experiment documents the same dynamic for Google’s AI Overviews, which divert traffic away from publishers whose content is summarized (Agarwal and Sen, 2026). Various

*MIT Operations Research Center. yandai20@mit.edu.

†Cornell, NBER and CEPR. m.farboodi@gmail.com.

‡MIT Sloan School of Management. golrezae@mit.edu.

§Washington University School of Law. sepehrs@wustl.edu.

[¶]We thank Xavier Gabaix, Mark Lemley, Ilan Morgenstern, Lisa Larrimore Oullette, Eva Tardos, and Yifan Wu, as well as the participants at Wharton Accountable AI, Stanford Market Design in the Age of AI, Marketplace Innovation Workshop, IP Researchers Europe, Society for Economic Research on Copyright Issues, Society for Institutional and Organizational Economics, ACM Economics and Computation (EC), and INFORMS M&SOM conferences for helpful conversations and comments.

artists and content creators have sued AI firms contending that their practices have deprived them of the rewards of their creative work (Section 2). 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 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 (a “strict intellectual property rights” regime). This paper argues that both extremes are destined to fail and points instead to a “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 intellectual property rights—is also doomed to fail.

The reasons for this failure are subtle and unexpected. We begin with a static Stackelberg model capturing the interaction between human creators and the AI firm. 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 driven by the market power of AI firms: a buyer with monopsony power can profitably price content below what would be required to induce socially optimal effort on the part of human creators. Correlation among different human creators’ content—the fact that, by selling her own content, a creator is effectively selling “part of” other creators’—does not by itself widen this shortfall, but, as we discuss next, it determines which creators bear it. Second—and more surprisingly—we show that AI firm’s 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. This 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 short-run tradeoff between human creative incentives and technological progress, we further 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 use AI tools to assist them generating 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 unexpected consequence is that a good model sows the seeds of its own decay. While this “curse of precision” looks similar to the phenomenon of “model collapse” identified in the empirical machine learning literature (Alemohammad et al., 2023; 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 constantly 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 overcomes the AI firm’s market power, securing creators a share of the social surplus their content generates. 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 compensates the most original creators in proportion to their contributions—rather than marking them down the most, as the unintermediated market does through the originality penalty. 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 and among legislators (Lucchi, 2026). 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 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

¹A 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, 2025).

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—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 “strict intellectual 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,

²That is not the only copyright-law question implicated by AI. Grossly simplified, there are at 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).

³Fair 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 strict intellectual 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 counterargument 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 strict intellectual 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 first reason is the market power of AI firms. 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. A second feature of this market is the correlation between different creators’ content. While it does not directly drive this under-investment—since the buyer already accounts for the correlation when setting prices—it shapes how the under-investment is distributed. Our analysis reveals that the markdown falls most heavily on the creators of the most

original 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, and it does so unevenly across creators. This finding inverts common wisdom about the effect of strong intellectual property rights.

In deriving this result, we draw on recent papers in information economics, particularly those by [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. \(2022\)](#) treat 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:⁴ Even with strict intellectual property rights, the correlation across creators’ outputs lets a buyer with market power mark down what it pays for content, and—as we show—it marks down the most original, least redundant content the most, discouraging precisely the creators whose contributions are hardest to replace.

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., 2023b,a](#); [Yao et al., 2023a,b](#); [Immorlica et al., 2024](#); [Golrezaei et al., 2025](#)), a concept dating back to [Glazer and Hassin \(1988\)](#).

⁴[Jain and Vazirani \(2010\)](#) similarly observe substitutability among digital creative goods, but in their setting such substitutability is semantic and arises on the consumer side—listeners treat songs within a genre as interchangeable. Our statistical substitutability instead arises on the production side, driven by the correlation across creators’ outputs.

Most related, a concurrent paper studies contest design when AI Overviews divert traffic from creators, proposing citation and compensation mechanisms to preserve their incentives (Wu et al., 2026). While creators are strategic in both their and our 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.

Our proposed mechanism draws 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 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. Because the failure stems from the buyer’s market power rather than transaction costs, the solution must counter that power; and because creators’ content is statistically substitutable, it must also reward originality. We achieve this by adopting a two-part tariff structure (Oi, 1971): By combining price-per-effort payments with a fixed subsidy, we redistribute the surplus attained by socially optimal production back to all participants, thus advancing technological progress and preserving human creative incentives at the same time.

Finally, methods have been proposed in machine learning 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. Koh and Liang (2017) developed influence functions to approximate the counterfactual impact of the absence of a data point in low dimensional environments, while Zou et al. (2025) calculated this impact in high dimensional environments. 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-Run 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 intellectual 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 $i = 1, 2, \dots, N$, and a single buyer (e.g., an AI firm acquiring training data from cre-

ators). 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.⁵

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

$$s_i = X + \xi_i, \quad \xi_i = \varepsilon_i + \beta_i \eta, \quad \forall i = 1, 2, \dots, N. \quad (1)$$

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).⁶ For simplicity, we assume all β_i are public information.

To produce content, agent i chooses an effort level $h_i \geq 0$. Effort reduces the variance of the idiosyncratic noise such that $\varepsilon_i \sim \mathcal{N}(0, h_i^{-1})$, but incurs a strictly convex production cost $C(h_i) := \frac{c}{2} h_i^2$ where $c > 0$ is common knowledge. We adopt the quadratic cost function—a standard formulation in contract theory (see, e.g., [Baker et al., 1994](#))—as a tractable abstraction of any convex cost function. We emphasize that our core findings—originality penalty, curse of precision, and the intermediated data market mechanism—extend to general convex costs. A strictly convex production cost function captures the increasing marginal cost when creators strive for higher precision via, e.g., careful research, thorough fact-checking, or creative effort.

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 or substantive convention (e.g., the seemingly star-crossed lovers improbably getting together at the end), 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 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 (e.g., pervasive prejudice against women in science or law), 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 effort 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.⁷

⁵In 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.

⁶Our additive signal structure is motivated by auction theory with interdependent values ([Milgrom and Weber, 1982](#)) and the value of public information in information economics ([Morris and Shin, 2002](#)), where agents observe noisy signals containing a statistically dependent public component. We adapt this logic to AI data markets.

⁷In 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

A defining feature of our framework is the correlation structure of signals. We model the common factor η using a hierarchical Bayesian structure:

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

where both $\gamma \geq 0$ and $\sigma_\eta^2 \geq 0$ are publicly known parameters, and $\sigma_\eta^2 + \gamma > 0$. The term μ_η represents the systematic bias of the current information environment. This hierarchical structure captures that the common bias is itself uncertain: while the buyer knows there may be a shared trend, bias, or AI-induced pattern affecting creators, it does not know its direction or magnitude *ex ante*, which is captured by the prior variance γ . This Bayesian setup is crucial for our dynamic model incorporating AI-assisted content creation in Section 4: As AI-generated content becomes pervasive, human-generated content becomes more homogeneous, the systematic bias μ_η becomes harder to filter out, and hence the prior variance γ effectively increases (see Section 4 for more details).

Aside from the Bayesian perspective, Equations (1) and (2) can be viewed through a minimax robust optimization lens: the platform estimates the ground-truth X knowing only that the magnitude of systematic bias is bounded ($\mu_\eta^2 \leq \gamma$). This establishes an alternative interpretation for our dynamic model: when AI tools are ubiquitous, imitators' reliance on AI increases, thereby inducing a larger bound γ for systematic bias; see Lemma 10 for formal equivalence.

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 .⁸ The buyer—as a leader in our Stackelberg game—has full power to decide all prices p_i . The creators react to the price vector $\mathbf{p}$ by deciding an individual effort level h_i that maximizes their own quasi-linear payoff:

$$U_i(h_i; p_i) = h_i p_i - C(h_i), \quad \forall i = 1, 2, \dots, N. \quad (3)$$

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

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

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

Aggregation of Data. After the signals $\mathbf{s} = (s_1, s_2, \dots, s_N)'$ 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 effective precision 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:

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

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.

⁸A critical premise of our Stackelberg game is that the buyer can contract on the effort vector. In reality, raw human “effort” (e.g., actual labor hours spent) is typically unobservable. However, modern AI platforms routinely contract on verifiable quality proxies: rigorous unit tests for coding tasks, detailed rubric scores in Reinforcement Learning from Human Feedback, or verified expert credentials. In consistency with various reduced-form approaches in contract theory, we treat h_i not as unobservable labor, but as the verifiable idiosyncratic signal precision, which is directly mapped from those observable quality metrics.

Under our model, the optimal solution to Equation (4) is the best linear unbiased estimator (BLUE) 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 $K(\mathbf{h}) = \text{Var}(\hat{X})^{-1}$, 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,⁹

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

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

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

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

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 $\mathbf{h}^*$ and $\mathbf{p}^*$, respectively. On the other hand, the first best solution $\mathbf{h}^{\text{sp}}$ is the outcome dictated by a benevolent social planner, who controls the production efforts to maximize social welfare:

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

where recall that $\mathbf{p}(\mathbf{h})$ is the price vector inducing $\mathbf{h}$ as the best response, and it is excluded from social welfare as it is a pure transfer among the market participants.

⁹In our model, the buyer’s profit is *linear in precision increase* rather than 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.

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 make the following observation: *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., $\mathbf{p} = (0, 0, \dots, 0)'$, then the only best response of creators is $\mathbf{h}(\mathbf{0}) = (0, 0, \dots, 0)'$. 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 intellectual property rights” regime.¹⁰

3.3 Short-Run 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 bargaining power to decide the market price $\mathbf{p} = (p_1, p_2, \dots, p_N)'$ (also 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 our first main result:

Proposition 1 (Underpowered Creative Incentives). *Under the market defined in Section 3.1, market equilibrium $\mathbf{h}^*$ and social planner solution $\mathbf{h}^{sp}$ are unique. Furthermore, $K(\mathbf{h}^*) < K(\mathbf{h}^{sp})$.*¹¹

The proof is in Appendix A. While Proposition 1 is consistent with the intuition that monopsony power leads to under-production, it paves the way for the second market failure that we identify—the originality penalty—a novel form of market failure distinctive to the context of data markets.

To see why the originality penalty is non-conventional, recall that the common factor η in Equation (1) implies that every creator i ’s signal s_i contains a potentially bias towards η to some extent. It thus follows that the buyer—who is only interested in recovering X but not η —faces a dual objective: (i) incentivizing production effort to reduce the idiosyncratic noise ε_i , while (ii) mitigating the statistical bias arising from the common factor η . 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 assumes homogeneous production costs, so producing “original” content does not cost more than “conventional” content. Thus, even allowing the AI firm to “lowball” human creators—using its monopsony power—by offering a price that does not incentivize sufficient investment in quality (Proposition 1), 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 market equilibrium, relative to the social optimum, the reduction in effort is greater for original creators than for rank-and-file creators. We call this the “originality penalty,” characterized in Proposition 2.

¹⁰Even 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.

¹¹The under-production result in Proposition 1 only applies to the overall effective precision, i.e., $K(\mathbf{h}^*)$ compared to $K(\mathbf{h}^{sp})$. The stronger per-creator under-production—namely $h_i^* \leq h_i^{sp}$ for all $i = 1, 2, \dots, N$ —is unfortunately incorrect. A counter-example is in Remark 3 of Appendix A, where—intuitively—a highly redundant creator i is set to invest 0 effort by the social planner (hence $h_i^{sp} = 0$); but due to the under-production of other creators in the equilibrium, the buyer finds it necessary to incentivize a strictly positive effort from creator i (hence $h_i^* > 0$).

Proposition 2 (Originality Penalty). *For each creator $i = 1, 2, \dots, N$, let $R_i := h_i^*/h_i^{sp}$ be the ratio of equilibrium effort to socially optimal effort. Then in a non-degenerate market without near-redundant creators,¹² creators with a higher β_i always retain a strictly higher R_i . Furthermore:*

- *For a creator i with $\beta_i \approx 0$ (an “innovator” who does not rely on the common factor), regardless of the remaining β ’s, their ratio $R_i \approx \frac{1}{2}$.*
- *For a creator i with $\beta_i \gg 0$ (an “imitator” producing homogeneous content), the ratio $R_i > \frac{1}{2}$.*
- *In the limiting case of a homogeneous market saturated by rank-and-file creators ($N \rightarrow \infty$ and $\beta_i \equiv \beta > 0$), the ratio converges to*

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

The counterintuitive result in Proposition 2 arises because of the interplay of monopsony power (i.e., the buyer has full market power) and diminishing returns (i.e., imitators’ contribution to precision diminishes due to content correlation). On a high level, this is because at the social optimum the planner asks little effort from imitators—their content is redundant, so extra effort costs the same but barely raises precision—and much more from innovators. The AI firm’s underpricing then falls unevenly: it buys less from the innovators to suppress the high cost—with convex costs, buying less means paying lower on every unit of effort—cutting them to about half the social optimum; on the other hand, it gains little from squeezing the already-low effort of those imitators.

Specifically, innovators—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, in the sense that their marginal contribution to the precision remains a constant no matter how much effort they have already invested. The social optimum is where the marginal cost of exerting effort equals this constant contribution. However, as the production cost is strictly convex in effort, buying less can reduce the price for all content; a monopsonist therefore strategically restricts the purchase quantity in exchange for a lower unit price (and higher unit profit). Solving the tradeoff between quantity and unit price gives $h_i^* \approx \frac{1}{2}h_i^{sp}$.

By contrast, imitators—given their high reliance on the common factor ($\beta_i \gg 0$) and thus high correlation with other creators—face strong diminishing returns. One can better understand this phenomenon as if all creators were reporting the same news: The first piece of report is highly valuable, but subsequent reports are nearly devoid of value. Consequently, the socially optimal level of data production would then be “one piece of information tells it all,” so the level of effort is low even in the social optimum (h_i^{sp}). In the market equilibrium, even when the AI firm 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 in Appendix A that $h_i^* \geq \frac{1}{\sqrt[3]{2}}h_i^{sp} \approx 0.79h_i^{sp}$ if β_i is sufficiently large.

We remark that, however, Proposition 2 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

¹²Formally, $\beta_i < 1/\Lambda(\mathbf{h}^{sp})$ for all i with $\Lambda(\mathbf{h})$ defined in Equation (23) of Appendix A. This implies the BLUE $\hat{X}$ defined in Equation (4) places a positive weight on every creator’s signal; equivalently, at the social optimum, a more correlated creator (higher β_i) invests less effort (smaller h_i^{sp}). See Remark 4 in Appendix A for details. We further remark that, even under this condition, the per-creator under-investment anticipated in Footnote 11 remains untrue.

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 penalty identified in Proposition 2 poses a risk to 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 leads to under-investment, and—interacting with data correlation—further to an adverse selection (the originality penalty) that essentially refuses to debias (by incentivizing the production of biased contents). However, we have so far kept the model static—focusing on a fixed snapshot of the data market—while 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, assume time is continuous, $t \in [0, +\infty)$. Consider two dynamic “data stock” state variables, $S_O(t)$ and $S_C(t)$, where $S_O(t)$ and $S_C(t)$ represent the cumulative stock of “original” and “correlated” data available at time t . Each data stock evolves according to standard capital accumulation dynamics: It increases via the inflow of new data production, and decreases due to depreciation (see Equation (8) for the formal definition). The two stocks of training data together determine the current model precision $K(t)$ at any time t .

Moreover, recall from Equation (2) that the common bias η is centered around a systematic bias $\mu_\eta \sim \mathcal{N}(0, \gamma)$. The bias was fixed in the static case. In the dynamic model, such bias is no longer exogenous: as the AI model becomes more powerful, the rank-and-file creators rely on it more heavily. Consequently, the systematic bias μ_η may be more severe. We capture this bias amplification through a *recursive bias* $\gamma(t)$, which we model 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 $S_O(t)/S_C(t) \rightarrow 0$), 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.”

The curse of precision consists of two components. First, there is an upper bound on market price. In particular, instead of a rising market price that reflects 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, so the price is “trapped” at a low level. Therefore, as formalized in Proposition 3, the market price is effectively decoupled from the scarcity of innovators. The second component is a barrier to model precision. If the low market price drives innovators out of the market, then the model—instead of sustaining on a massive amount of correlated data from the rank-and-file creators—hits a “ceiling.” In Proposition 4, we prove that even with an infinite supply of correlated data, model precision $K(t)$ still cannot surpass 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 2, we divide the large population of content creators into the two representative classes:

- **Innovators Producing Original Data (Type-O):** These creators share a near-zero loading on the common factor, denoted by $\beta_O \approx 0$.¹³ They can be interpreted as domain experts whose contents are based on their own opinion. Therefore, their data is uncorrelated with the common factor η .
- **Rank-and-file Creators Producing Correlated Data (Type-C):** This type of creators share a high β , namely $\beta_C \gg 0$. They represent 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 $t \geq 0$.¹⁴ 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

¹³One may generalize β_O to a distribution of β s among innovators. For simplicity, we treat β_O and β_C as constants.

¹⁴While the micro-level model in Section 3 allowed for creator-specific prices p_i , here we assume a single clearing price $p(t)$. We make this assumption because a creator’s type is a statistical property of her content relative to the rest of the corpus, and is hard to verify at the point of acquisition (especially when web-scale data is bought through bulk licenses, platform agreements, or scraping settlements that price a source or category rather than individual creator’s originality). However, this restriction is also not innocuous: Were the firm able to price by type, it could keep paying innovators their high marginal value, so the exit of innovators would not arise. Still, in the spirit of the market for lemons (Akerlof, 1970), we read the single price as a constraint forced by the non-verifiability of types.

the supply of rank-and-file creators is more elastic to prices than that of innovators, namely

$$\mathcal{E}_C(p) > \mathcal{E}_O(p), \quad \forall p > 0, \quad \text{where } \mathcal{E}_j(p) := \frac{d \ln N_j(p)}{d \ln p}, \quad \forall j \in \{O, C\}. \quad (7)$$

In words, this assumption captures the asymmetric impact of technology on different types of content creators. For rank-and-file creators, AI tools greatly lower their 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—scarce human intelligence that cannot be algorithmically amplified—making the supply of innovators less responsive to price changes. As we prove in Lemma 11 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 $h(p) := p/c$.

Data Accumulation. 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:

$$\frac{dS_j}{dt}(t) = \underbrace{N_j(p(t))h(p(t))}_{\text{Production Inflow}} - \underbrace{\delta S_j(t)}_{\text{Depreciation}}, \quad \forall j \in \{O, C\}, t \geq 0. \quad (8)$$

The first term captures the total inflow: mass of creators N_j multiplied by their production effort h . The second term is data depreciation. It 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. We follow Farboodi and Veldkamp (2021) in modeling data depreciation.

Recursive Bias and Precision. To capture the recursive nature of AI training and the reality that many content creators are increasingly utilizing powerful AI tools during production, we turn Equation (2) into a dynamic model. Let $K(t)$ denote the effective precision of the current 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

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

where the prior variance $\gamma(t)$ is linear in $K(t)$ ¹⁵

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

The effective precision of the current mode $K(t)$ is in turn determined by the stocks $S_O(t)$, $S_C(t)$ and the bias $\gamma(t)$. Lemma 9 in the Appendix shows that the effective precision is given by:

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

¹⁵In practice, rank-and-file creators 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.

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 model precision is defined via the following fixed-point problem:

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

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

AI firm Objective Function. We finally define the objective of the buyer. Given that AI technologies are susceptible to creative destruction, we believe it is reasonable to assume that the AI firm is less patient than creators.¹⁶ To simplify the exposition, we capture the difference in patience by modeling the AI firm as myopic. In particular, we assume that the AI firm cannot commit to long-term subsidies and focuses on instantaneous profit. As such, the AI firm sets the price $p(t)$ to myopically maximize instantaneous profit:

$$\Pi_{\text{inst}}(t) := \underbrace{\frac{dK}{dt}(t)}_{\text{Precision Increase}} - \underbrace{\left(N_O(p(t)) + N_C(p(t))\right)p(t)h(p(t))}_{\text{Purchase Expenditure}}, \quad \forall t \geq 0. \quad (12)$$

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 stays bounded away from zero ($0 < \frac{\partial K}{\partial S_O}(t) \leq 1$)—that of rank-and-file creators exhibits diminishing returns analogous to Proposition 2. Specifically, we derive an upper bound $\kappa_C(t)$ on the marginal increase in effective precision from rank-and-file data. As such, $\kappa_C(t)$ is the upper bound on the marginal value of rank-and-file data.

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

Lemma 12 formally derives Equation (13). From Equation (13), there are two forces driving down $\frac{\partial K}{\partial S_C}$. The first force is more intuitive and simply corresponds to accumulation of more data. As the stock of correlated data $S_C(t)$ grows, 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. The second force is more specific to the context of AI model training and corresponds to the recursive bias embedded in a more precise model. This force is an amplification effect introduced by $K(t)$: As the model becomes more precise, the rank-and-file creators rely even more on the common bias, which in turn amplifies the bias going forward. That is, a stronger model (larger $K(t)$) effectively “pollutes” the rank-and-file creator data

¹⁶ Alternatively, we can microfound this assumption by considering two AI firms who compete à la duopoly for consumers who have switching cost. In this environment, AI firms over-weight their near-future profit as it corresponds to higher market share and captive consumers.

with amplified bias, which further drives down the value of rank-and-file creators’ 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 AI firm—facing an overwhelming supply of rank-and-file creators whose data has vanishing marginal value—fails to raise the price $p(t)$ to attract innovators. The next proposition captures this phenomenon.

Proposition 3 (Upper Bound on Market Price). *At any time $t \geq 0$, the market price $p(t)$ cannot exceed a threshold determined by $\kappa_C(t)$, namely $p(t) \leq \bar{p}(\kappa_C(t))$ for some strictly increasing function $\bar{p}(\cdot)$. Specifically, when $\kappa_C \rightarrow 0^+$, this upper bound $\bar{p}(\kappa_C)$ converges to a “price trap” p^* :*

$$\lim_{\kappa_C \rightarrow 0^+} \bar{p}(\kappa_C) = p^*, \quad \text{where } p^* = \left(1 + \frac{N_C(p^*)}{N_O(p^*)}\right)^{-1}. \quad (14)$$

Proposition 3 illustrates that a smaller $\kappa_C(t)$ induces a lower market price $p(t)$. Moreover, it characterizes the limiting behavior of the market price: it shows that as $\kappa_C(t) \rightarrow 0^+$ —driven by the accumulation of correlated data or the recursive bias of a more precise model—the market price is capped at the fixed level p^* specified in Equation (14).

In particular, as the stock of correlated data accumulates (i.e., $S_C(t) \rightarrow \infty$) or the recursive bias amplifies due to powerful AI models (i.e., $K(t) \rightarrow \infty$), the marginal value of rank-and-file creators’ 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 widespread availability of AI lowers the barriers to entry for rank-and-file creators, the supply of rank-and-file creators N_C increases far faster than that of innovators N_O . This further pushes the trap level p^* down towards zero, which creates a risk: if p^* is insufficient to attract *any* innovators (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 ($S_O(t) \rightarrow 0$) and see whether the buyer can survive solely on massive amounts of rank-and-file creator data ($S_C(t) \rightarrow \infty$). The widely believed machine learning “scaling laws” (Kaplan et al., 2020; Hoffmann et al., 2022) suggest that the *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 4 (Barrier of Model Precision). *Suppose innovators exit the market ($S_O(t) = 0$). Even if the stock of rank-and-file creator data grows indefinitely ($S_C(t) \rightarrow \infty$), the equilibrium model precision $K(t)$ converges to a finite upper bound $\bar{K}$ determined by the recursive bias coefficient λ in Equation (10):*

$$\lim_{S_C(t) \rightarrow \infty} K(t) = \bar{K} := \left(-\sigma_\eta^2 + \sqrt{\sigma_\eta^4 + 4\lambda/\beta_C^2}\right) / 2\lambda. \quad (15)$$

Proposition 4 captures a failure mode similar to the “model collapse” 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 rank-and-file creators—the model still hits a performance ceiling.

We call this the counterintuitive result the “curse of precision”: a more powerful AI model induces stronger correlation among rank-and-file creators, which drives down the market price (Proposition 3) and potentially hurts the model’s long-run performance (Proposition 4). We then ask how to design a mechanism for the AI data market to recover market efficiency and also participation.

5 Mechanism Design: Restoring Efficiency and Participation

Sections 3 and 4 have revealed the inadequacies of existing approaches to the design of AI data markets. Intuitively, we showed that if AI firms can simply take human content for free, then humans are deprived of incentives to invest in creating high-quality content. More interestingly, we demonstrated that IP rights are insufficient to incentivize the creation of high-quality content. Although IP rights are meant to incentivize creativity, in fact they under-power creative incentives, especially for innovative creators. This 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 uses two-part pricing contracts to restore human creative incentives while simultaneously providing AI models with the inputs they need to improve.

The intermediary in our solution is similar 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). However, unlike traditional copyright CMOs, whose primary function is to reduce transaction costs (Cotter, 2005; Gilbert, 2017), the role of our intermediary is dictated by a different problem. The failures we identify—under-investment and the originality penalty—arise from the AI firm’s market power, and survive even when bargaining is frictionless (Proposition 1); reducing transaction costs alone would therefore not restore socially optimal content creation. By aggregating the rights of a critical mass of creators and negotiating as a single counterparty, the intermediary counters the buyer’s market power and secures creators a share of the surplus their content generates.¹⁷

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. First, it determines how the surplus is divided among creators. In an unintermediated market, the buyer’s markdown falls most heavily on the most original creators—the originality penalty of Proposition 2. Managing the creators as a portfolio, the intermediary instead apportions their collective compensation by each creator’s contribution to the joint value of the data—using weights inspired by the Aumann–Shapley value—so that the most original creators, who contribute the most, receive correspondingly larger shares rather than the deepest cut.

Second, the intermediary ensures participation via adequate 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

¹⁷If the intermediary is a private party, there can be principal-agent problems between the content creators and the intermediary. Currently, we focus on the design of a regulated intermediary and abstract from delegation frictions.

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 Institutional Design: Intermediated Data Market

Consider $N \geq 2$ human content creators, a single AI firm, and an intermediary with a fixed operational cost $F \geq 0$. The intermediary negotiates with both sides: (i) To the creators, the intermediary promises an individual payment $P_i(\mathbf{h})$ conditional on the N creators jointly invest efforts $\mathbf{h} = (h_1, h_2, \dots, h_N)'$. (ii) To the AI firm, the intermediary offers the N -dimensional signal generated by all the data creators investing efforts $\mathbf{h} = (h_1, h_2, \dots, h_N)'$, in return for a transfer $T(\mathbf{h})$ to receive all signals $\mathbf{s}$. That is, the intermediary determines two functions $\mathbf{P}: \mathbb{R}_{\geq 0}^N \rightarrow \mathbb{R}^N$ and $T: \mathbb{R}_{\geq 0}^N \rightarrow \mathbb{R}$. We call tuple $\langle \mathbf{P}, T \rangle$ a *mechanism*. As in Section 3, we assume the effort levels $\mathbf{h}$ are observable and contractible.

Analogous to Equation (3), the payoff of a creator i under the mechanism $\langle \mathbf{P}, T \rangle$ is specified as

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

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

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

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

To capture the participation constraints of each agent, we denote the AI firm'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 $\langle \mathbf{P}, T \rangle$ satisfying:

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

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

2. **Budget Balance (BB):** The intermediary breaks even (neither deficit nor surplus):¹⁸

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

¹⁸Note 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 $T(\mathbf{h})$ need not equal the social value $K(\mathbf{h})$; the difference $\Pi = K(\mathbf{h}) - T(\mathbf{h})$ is absorbed by the AI firm, who acts as the "budget breaker" to extract the surplus or bear the deficit.

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

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

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

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

We summarize the resulting sequence of play in the following definition, in parallel with Definition 1.

Definition 2 (Sequence of Play, Intermediated Market). The sequence of play is as follows:

1. The intermediary and the buyer bargain over the mechanism $\langle \mathbf{P}, T \rangle$ according to the Nash bargaining procedure in Equation (16), where the intermediary has a bargaining power of α . The intermediary then commits to the payment rules $\mathbf{P}$ and the transfer rule T .
2. Creators simultaneously choose efforts to maximize payoff $U_i(\mathbf{h}; \mathbf{P})$. From Equation (IC), their best response must equal the social optimum $\mathbf{h}^{\text{sp}}$.
3. From the signals $\mathbf{s}$ generated by $\mathbf{h}^{\text{sp}}$, the buyer receives profit $\Pi(\mathbf{h}^{\text{sp}}; T)$ and pays the transfer $T(\mathbf{h}^{\text{sp}})$, while the intermediary disburses payments $P_i(\mathbf{h}^{\text{sp}})$ and breaks even by Equation (BB).

5.2 Linear Pricing Fails on Participation

We start with a simple class of mechanisms—linear pricing—and show why it falls short. Specifically, consider the following family of mechanisms:

$$P_i(\mathbf{h}) = p_i h_i, \quad \forall i = 1, 2, \dots, N, \quad (17)$$

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

Lemma 5 (Marginal Prices and IC). *Under linear pricing, the creators’ best response coincides with the social optimum $\mathbf{h}^{\text{sp}}$ if and only if the unit prices are set as $p_i = \frac{\partial K}{\partial h_i}(\mathbf{h}^{\text{sp}}) = ch_i^{\text{sp}}$ for all i .*

Since the unit prices $\mathbf{p}$ are uniquely pinned down by Lemma 5, 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 Lemma 6, 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 $\langle \mathbf{P}, T \rangle$ 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.

Lemma 6 (Linear Pricing Fails on Participation). *There exists a unique linear pricing mechanism $\langle \mathbf{P}, T \rangle$ satisfying Equations (BB) and (IC). Under this mechanism, each creator i 's payoff is $U_i(\mathbf{h}^{sp}; \mathbf{P}) = \frac{c}{2}(h_i^{sp})^2$, while the buyer's profit is*

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

In the special case where $\beta_i \equiv 0$, we have $\Pi(\mathbf{h}^{sp}; T) = -F$, thus violating Equation (IR) whenever $F + \underline{\Pi} > 0$.

All formal proofs omitted in this section are in Appendix C.

5.3 Optimal Contract: Two-Part Pricing

Since linear pricing is pinned down by incentive constraints (Lemma 5) 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:¹⁹

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

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 $\mathcal{S}(\mathbf{h})$ relative to outside options as:

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

At the social optimum $\mathbf{h}^{sp}$, the surplus $\mathcal{S}(\mathbf{h}^{sp})$ is divided via Nash bargaining. The intermediary secures a share $\alpha \in [0, 1]$ for all creators, while the remaining $(1 - \alpha)\mathcal{S}(\mathbf{h}^{sp})$ goes to the buyer. Thus, the transfer from the buyer to the intermediary, namely $T(\mathbf{h}^{sp})$, must satisfy

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

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

To distribute the creators' share $\alpha\mathcal{S}(\mathbf{h}^{sp})$ fairly, we define weights $\phi = (\phi_1, \dots, \phi_N)'$ inspired by the Aumann–Shapley value (Aumann and Shapley, 1974). Specifically, in Equation (22), each ϕ_i

¹⁹A natural question is whether equipping the AI firm itself with the two-part tariff of Equation (18) would already restore efficiency. The answer is formally affirmative but economically vacuous: The firm would elicit $\mathbf{h}^{sp}$ via the marginal price, and consequently extract *all* social surplus through a negative $B_i = \underline{u}_i - \frac{c}{2}(h_i^{sp})^2 < 0$. This forces creators to wire substantial upfront fees *before* earning any production wages. On the creator side, this violates standard expectations that the *buyer* of the product (especially product protected by property rights)—instead of the seller—is the one who pays. Such a reverse payment would make many sellers suspicious and unwilling to enter into the arrangement. On the firm side, this requires a credible commitment—delivering the promised wages after the fee is collected—which no profit-maximizing buyer would have an incentive to honor. Such a scheme also leaves creators exactly at their outside options, with no share of the surplus they generate; securing them such a share is one role of the intermediary. Finally, the affirmative answer also relies heavily on the full information on creators' β and $\mathbf{h}$, and introducing private information and/or moral hazard to the interaction between the AI firm and creators will break this result. We leave all these issues for the next version of this work.

Mechanism 1 Two-Part Tariff Mechanism

Input: Number of creators $N \geq 2$, production cost $c > 0$, parameters $\beta, \gamma, \sigma_\eta^2$ 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 $\mathbf{h}^{\text{sp}}$ maximizing social welfare W in Equation (6).
- 2: The intermediary calculates the social surplus at $\mathbf{h}^{\text{sp}}$ defined in Equation (19), namely $\mathcal{S}(\mathbf{h}^{\text{sp}})$.
- 3: The intermediary decides the lump-sum transfer L from the buyer as:

$$L = K(\mathbf{h}^{\text{sp}}) - \sum_{i=1}^N C'(h_i^{\text{sp}}) \times h_i^{\text{sp}} - (1 - \alpha)\mathcal{S}(\mathbf{h}^{\text{sp}}) - \underline{\Pi}. \quad (20)$$

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

$$B_i = \phi_i \alpha \mathcal{S}(\mathbf{h}^{\text{sp}}) + \underline{u}_i - C(h_i^{\text{sp}}), \quad (21)$$

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

$$\phi_i \propto h_i^{\text{sp}} \int_0^1 \frac{\partial \mathcal{S}}{\partial h_i}(\tau \mathbf{h}^{\text{sp}}) d\tau \quad \text{s.t.} \quad \sum_{i=1}^N \phi_i = 1. \quad (22)$$

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

is proportional to creator i 's cumulative contribution to the surplus along the diagonal path from $\mathbf{0}$ to $\mathbf{h}^{\text{sp}}$. We set each creator i 's utility to their fair share $\phi_i \alpha \mathcal{S}(\mathbf{h}^{\text{sp}})$ plus their outside option $\underline{u}_i$:

$$ch_i^{\text{sp}} \times h_i^{\text{sp}} + B_i - C(h_i^{\text{sp}}) = U_i(\mathbf{h}^{\text{sp}}; \mathbf{P}) = \phi_i \alpha \mathcal{S}(\mathbf{h}^{\text{sp}}) + \underline{u}_i.$$

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

We therefore obtain a two-part tariff mechanism $\langle \mathbf{P}, T \rangle$ 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 Algorithm 1 in Proposition 7:

Proposition 7 (Two-Part Tariff Attains Efficiency and Participation). *If $\mathcal{S}(\mathbf{h}^{\text{sp}}) > 0$, i.e., the market is efficient, the two-part tariff mechanism $\langle \mathbf{P}, T \rangle$ defined in Algorithm 1 satisfies:*

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

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

5.4 Remaining Challenges and Legal Implementation

In this section, we lay out some remaining questions and challenges to be addressed as we continue to refine our market design proposal. As discussed above, the need for our solution emerges out of the failure of the two options usually available to courts. Neither a system of blanket immunity via fair use nor a legal regime of full control over access to content by copyright owners will be sufficient to preserve creative incentives, so a workable solution must transcend the traditional copyright binary. Although our diagnosis of market failures in this context is novel, our proposal for going beyond the

common copyright binary is by no means without precedent. The US Copyright Act (to say nothing of other countries’ laws) contains a number of such “intermediate” regimes.²⁰ We plan to draw lessons from these regimes by understanding their precise legal workings, the political and economic forces that shaped them, and their policy successes and failures. Although this research is only just beginning, it’s useful preliminarily to point out some comparisons between these precedents and the AI context we target.

Broadly speaking, existing intermediate regimes share two features with the present AI context. First, like the intermediate regime we propose, existing intermediate regimes were often responses to a disruptive technology that upended settled ways of doing things (Shahshahani, 2018). Second, consistent with our intuitions and preliminary results about the role of data intermediaries, these legislative compromises often involve organizations to coordinate collective action among copyright holders (or, occasionally, among other stakeholders). These organizations include the Mechanical Licensing Collective, SoundExchange, the Copyright Clearance Center, AP Images, Artists Rights Society, reprographic rights organizations, ASIP, and performance rights organizations such as ASCAP, BMI, SESAC, and Global Music Rights. Some of these organizations are created or sanctioned by legislation (or by regulations enacted pursuant to legislation), and some of them are outgrowths of “organic” coordination among copyright holders or other stakeholders.

Notwithstanding these similarities, existing intermediate regimes differ from our proposed regime in one crucial respect: Under existing regimes, the primary function of intermediaries is often to reduce transaction costs and smoothen frictions in collective action—exactly the sort of problems which one would intuit to be at the root of bargaining breakdown, but which were *not* the main mechanism driving the failure of the property-rights model in our theoretical framework (see Proposition 1 and surrounding discussion). Given that, in our model, market failures happen even in the absence of garden-variety transaction costs, the function of intermediaries in our market design must do more than reduce transaction costs. Hence the foregoing discussion of mechanisms to counter the buyer’s market power and subsidize originality.

Beyond highlighting these similarities and differences, researching existing intermediate regimes points up two important questions of institutional design. The first question relates to 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? This question implicates issues of collective action that have been of perennial interest in political economy at least since Olson (1965).²¹

Another important question is that of price setting. It is widely understood that, in many contexts, a centralized system is not as good at determining the value of assets as the decentralized information-aggregating mechanism of the market (Hayek, 1945). The way an intermediate regime, such as a compulsory-licensing scheme, sets the price at which an exchange must take place is therefore an important ingredient of the regime’s success or failure. Pricing content may be easy if

²⁰Examples include: 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).

²¹From 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 generally contain such exemptions.

that content has a market outside the use for which the compulsory license is sought, in which case one can piggyback on the market price in setting the license fee. But if the primary market for the content is the use for which the compulsory license is sought, then there is no ready recourse to the market price signal. In the context of a data market for AI, pricing may not be difficult in the beginning because the data in question have had a long history of consumer demand predating the advent of AI, which history can be relied upon for pricing; however, going forward, if we anticipate the primary (direct) market for content to be as fodder for AI models, then pricing becomes more difficult. The question has at least two aspects: (1) determining the price at which the intermediary sells the content to the AI firm, (2) determining how to apportion the price received by the intermediary among different content creators. Our analysis assumes that the originality and long-term value of content is *known* to the intermediary, but the elicitation of such information (assuming it is even privately available) is a central difficulty of centralized mechanisms with which our solution must contend.

6 Conclusion

This paper studies the economics of data markets for AI. The force that depresses content creation is the market power of AI firms; the statistical substitutability of human content—the correlation across creators’ outputs—does not create this shortfall but determines who bears it. In static markets, we show that the interaction between correlation and AI firms’ market power generates an originality penalty that disproportionately harms creators producing uncorrelated and socially valuable content, thereby weakening incentives precisely for the creators with the highest marginal contribution. Extending the analysis to a dynamic setting with AI-assisted creation, we identify a curse of precision: improvements in AI models induce increasingly homogenized data production, creating an upper bound on market prices and potentially causing innovative creators to exit. As a result, even a continued inflow of fresh human data may fail to prevent long-run model collapse absent an appropriate market design.

To address these failures, we propose a market design with two components, institutional design and contract design. In particular, we propose a data intermediary that uses a two-part tariff contract that implements the social optimum while satisfying participation constraints and maintaining budget balance. The mechanism restores long-run participation for both model performance and human creativity by aligning incentives for innovation with technological progress. We conclude by highlighting two remaining challenges for implementation: whether participation in intermediary organizations can emerge through decentralized equilibrium behavior or instead requires centralized enforcement and statutory coordination, and how intermediaries can estimate creators’ marginal contributions in environments where outside markets for content are thin or nonexistent. Integrating the proposed mechanism with data valuation techniques therefore remains an important direction for future research.

References

Daron Acemoglu, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4): 218–256, 2022.

- Anish Agarwal, Munther Dahleh, and Tuhin Sarkar. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 701–726, 2019.
- Saharsh Agarwal and Ananya Sen. Google AI overviews and publisher traffic: Evidence from a field experiment. *Available at SSRN 6513059*, 2026.
- Alexander C Aitken. Iv.—on least squares and linear combination of observations. *Proceedings of the Royal Society of Edinburgh*, 55:42–48, 1936.
- George A Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
- Sina Alemohammad, 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.
- Kenneth Arrow. Economic welfare and the allocation of resources for invention. *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pages 609–626, 1962.
- Robert J Aumann and Lloyd S Shapley. *Values of Non-Atomic Games*. Princeton University Press, 1974.
- George Baker, Robert Gibbons, and Kevin J Murphy. Subjective performance measures in optimal incentive contracts. *The Quarterly Journal of Economics*, 109(4):1125–1156, 1994.
- Jonathan M. Barnett. The free content illusion. *Journal of Intellectual Property Law*, 33(1), 2026.
- Dirk Bergemann, Alessandro Bonatti, and Tan Gan. The economics of social data. *The RAND Journal of Economics*, 53(2):263–296, 2022.
- Leonardo Berti, Flavio Giorgi, and Gjergji Kasneci. Emergent abilities in large language models: A survey. *arXiv preprint arXiv:2503.05788*, 2025.
- Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- Gordon Burtch, Dokyun Lee, and Zhichen Chen. The consequences of generative ai for online knowledge communities. *Scientific Reports*, 14(1):10413, 2024.
- Yeon-Koo Che. Beyond the Coasian irrelevance: Asymmetric information. *Unpublished notes, Columbia University*. [110, 114], 2006.
- Ronald H Coase. The problem of social cost. *Journal of Law and Economics*, 3:1–44, 1960.
- Thomas F Cotter. Some observations on the law and economics of intermediaries. *Michigan State Law Review*, pages 67–82, 2005.
- R Maria del Rio-Chanona, Nadzeya Laurentsyeve, and Johannes Wachs. Large language models reduce public knowledge sharing on online q&a platforms. *PNAS nexus*, 3(9):pgae400, 2024.
- Maryam Farboodi and Laura Veldkamp. Long-run growth of financial data technology. *American Economic Review*, 110(8):2485–2523, 2020.
- Maryam Farboodi and Laura Veldkamp. A model of the data economy. *NBER Working Paper*, (w28427), 2021.

- Maryam Farboodi and Laura Veldkamp. Data and markets. *Annual Review of Economics*, 15(1): 23–40, 2023.
- Maryam Farboodi, Gregor Jarosch, and Robert Shimer. The emergence of market structure. *The Review of Economic Studies*, 90(1):261–292, 2023.
- Janet Freilich and Sepehr Shahshahani. Measuring follow-on innovation. *Research Policy*, 52: 104854, 2023.
- Amirata Ghorbani and James Zou. Data Shapley: Equitable valuation of data for machine learning. In *International Conference on Machine Learning*, pages 2242–2251. PMLR, 2019.
- Richard Gilbert. Collective rights organizations: A guide to benefits, costs and antitrust safeguards. *The Cambridge Handbook of Technical Standardization Law*, 1, 2017.
- Amihai Glazer and Refael Hassin. Optimal contests. *Economic Inquiry*, 26(1):133–143, 1988.
- Negin Golrezaei and Hamid Nazerzadeh. Pricing schemes for metropolitan traffic data markets. In *International Conference on Data Management Technologies and Applications*, volume 2, pages 266–271. SCITEPRESS, 2014.
- Negin Golrezaei, MohammadTaghi Hajiaghayi, and Suho Shin. The contest behind the feed: Optimal contest for recommender systems. *Available at SSRN 5258940*, 2025.
- Wendy J. Gordon. Fair use as market failure: A structural and economic analysis of the Betamax case and its predecessors. *Columbia Law Review*, 82:1600–1657, 1982.
- Friedrich A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.
- Michael A Heller. The tragedy of the anticommons: Property in the transition from Marx to markets. *Harvard Law Review*, 111(3):621–688, 1998.
- Michael A. Heller and Rebecca S. Eisenberg. Can patents deter innovation? the anticommons in biomedical research. *Science*, 280(5364):698–701, 1998.
- Peter Henderson, Xuechen Li, Dan Jurafsky, Tatsunori Hashimoto, Mark A. Lemley, and Percy Liang. Foundation models and fair use. *Journal of Machine Learning Research*, 24:1–79, 2023.
- Jordan Hoffmann, 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*, pages 30016–30030, 2022.
- Bengt Holmstrom. Moral hazard in teams. *The Bell Journal of Economics*, pages 324–340, 1982.
- Nicole Immorlica, 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*, pages 36–45, 2024.
- Meena Jagadeesan, Nikhil Garg, and Jacob Steinhardt. Supply-side equilibria in recommender systems. *Advances in Neural Information Processing Systems*, 36:14597–14608, 2023a.

- Meena Jagadeesan, Michael I Jordan, and Nika Haghtalab. Competition, alignment, and equilibria in digital marketplaces. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5689–5696, 2023b.
- Kamal Jain and Vijay Vazirani. Equilibrium pricing of semantically substitutable digital goods. *arXiv preprint arXiv:1007.4586*, 2010.
- Ruoxi Jia, 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*, pages 1167–1176. PMLR, 2019.
- Jared Kaplan, 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.
- Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International Conference on Machine Learning*, pages 1885–1894. PMLR, 2017.
- Steven G Krantz and Harold R Parks. *The Implicit Function Theorem: History, Theory, and Applications*. Springer Science & Business Media, 2002.
- Edward Lee. Fair use and the origin of AI training. *Houston Law Review*, 63:105–229, 2025.
- Mark A. Lemley. How generative AI turns copyright law upside down. *Columbia Science & Technology Law Review*, 25(2):190–212, 2024.
- Mark A. Lemley and Bryan Casey. Fair learning. *Texas Law Review*, 99:743–785, 2021.
- Mark A. Lemley and Lisa Larrimore Ouellette. Plagiarism, copyright, and ai. *University of Chicago Law Review Online*, 2025, 2025. URL <https://lawreview.uchicago.edu/online-archive/plagiarism-copyright-and-ai>.
- Pierre N. Leval. Toward a fair use standard. *Harvard Law Review*, 103(5):1105–1136, 1990.
- Nicola Lucchi. The impact of Google AI summaries and Google AI overviews on publishers’ revenue and media freedom implications for the information ecosystem and democratic resilience in the European Union. *European Parliament Policy Department Briefing, PE*, 787, 2026.
- Fritz Machlup. An economic review of the patent system. Committee Print Study No. 15, U.S. Senate Subcommittee on Patents, Trademarks, and Copyrights, Washington, 1958. URL https://cdn.mises.org/An%20Economic%20Review%20of%20the%20Patent%20System_Vol_3_3.pdf. 85th Congress, 2d session.
- Paul R Milgrom and Robert J Weber. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society*, pages 1089–1122, 1982.
- Stephen Morris and Hyun Song Shin. Social value of public information. *American Economic Review*, 92(5):1521–1534, 2002.
- Alexander Muschalle, 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*, volume 154, page 129. Springer, 2013.

- William Nordhaus. *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change*. MIT Press, 1969.
- Walter Y Oi. A Disneyland dilemma: Two-part tariffs for a Mickey Mouse monopoly. *The Quarterly Journal of Economics*, 85(1):77–96, 1971.
- Mancur Olson. *The Logic of Collective Action: Public Goods and the Theory of Groups*. Harvard University Press, 1965.
- David W. Opderbeck. Copyright in AI training data. *Oklahoma Law Review*, 76:951–1023, 2024.
- Elinor Ostrom. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, 1990.
- Arnold Plant. The economic theory concerning patents for inventions. *Economica*, 1(1):30–51, 1934.
- Register of Copyrights. *Copyright and Artificial Intelligence – Part 3: Generative AI Training*. U.S. Copyright Office, 2025. Pre-Publication Version.
- Ariel Rubinstein and Asher Wolinsky. Middlemen. *The Quarterly Journal of Economics*, 102(3): 581–593, 1987.
- Matthew Sag. The new legal landscape for text mining and machine learning. *Journal of the Copyright Society of the USA*, 66:291, 2019.
- Pamela Samuelson. Generative ai meets copyright. *Science*, 381(6654):158–161, 2023.
- Pamela Samuelson. Fair use defenses in disruptive technology cases. *UCLA Law Review*, 71: 1484–1572, 2024.
- Pamela Samuelson. Does using in-copyright works as training data infringe? *Communications of the ACM*, 68(11):28–30, 2025.
- Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. Are emergent abilities of large language models a mirage? *Advances in neural information processing systems*, 36:55565–55581, 2023.
- Fabian Schomm, Florian Stahl, and Gottfried Vossen. Marketplaces for data: An initial survey. *SIGMOD Record*, 42(1):15, 2013.
- Ilya Segal and Michael D Whinston. Property rights and the efficiency of bargaining. *Journal of the European Economic Association*, 14(6):1287–1328, 2016.
- Sepehr Shahshahani. The Nirvana fallacy in fair use reform. *Minnesota Journal of Law, Science & Technology*, 16(1):273–341, 2015.
- Sepehr Shahshahani. The role of courts in technology policy. *Journal of Law and Economics*, 61(1):37–64, 2018.
- Sepehr Shahshahani. Against the abstract-ideas exclusion. *Berkeley Technology Law Journal*, 40: 447–518, 2025.
- Neha Sharma and Simin Li. Beyond substitution: Large language models drive novel knowledge emergence in online forums. *Available at SSRN 5388669*, 2025.

- Jack Sherman 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*, 21(1):124–127, 1950.
- Ilia Shumailov, Zakhar Shumaylov, Yiren Zhao, Yarin Gal, Nicolas Papernot, and Ross Anderson. The curse of recursion: Training on generated data makes models forget. *Nature*, 631:755–759, 2024.
- Jane Tullis. Sifting through the slop: How generative ai created a market for lemons for text-based works, 2025. URL <https://ssrn.com/abstract=5266660>.
- Laura L Veldkamp. *Information Choice in Macroeconomics and Finance*. Princeton University Press, 2011.
- Jason Wei, 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. ISSN 2835-8856.
- Gerhard Widmer and Miroslav Kubat. Learning in the presence of concept drift and hidden contexts. *Machine learning*, 23(1):69–101, 1996.
- Oliver E Williamson. Transaction-cost economics: The governance of contractual relations. *Journal of Law and Economics*, 22(2):233–261, 1979.
- Max A Woodbury. *Inverting Modified Matrices*. Department of Statistics, Princeton University, 1950.
- Yihang Wu, Jiajun Tang, Jinfei Liu, Haifeng Xu, and Fan Yao. Do AI overviews benefit search engines? an ecosystem perspective. *arXiv preprint arXiv:2601.22493*, 2026.
- Fan Yao, 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*, pages 39674–39701. PMLR, 2023a.
- Fan Yao, Chuanhao Li, 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*, 36:74582–74601, 2023b.
- Haolin Zou, Arnab Auddy, Yongchan Kwon, Kamiar Rahnama Rad, and Arian Maleki. Newfluence: Boosting model interpretability and understanding in high dimensions. In *ICML 2025 Workshop on Assessing World Models*, 2025. URL <https://openreview.net/forum?id=AALFCxEucZ>.

Technical Appendices

A Proofs for Section 3	30
A.1 First-Order Conditions	30
A.2 Underpowered Creative Incentives (Proof of Proposition 1)	32
A.3 Originality Penalty (Proof of Proposition 2)	33
A.4 Minimax Robust Optimization View of Equation (2)	35
B Proofs for Section 4	36
B.1 Supply Elasticity and Monotonicity of Ratio	36
B.2 Upper Bound on Market Price (Proof of Proposition 3)	37
B.3 Barrier of Model Precision (Proof of Proposition 4)	39
C Proofs for Section 5	39
C.1 Linear Pricing Fails on Participation (Proof of Lemma 6)	39
C.2 Two-Part Tariff for Efficiency and Participation (Proof of Proposition 7)	40

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 $\mathbf{h}^*$. Likewise, we also solve the first-order conditions for the first best solution $\mathbf{h}^{\text{sp}}$. We then prove Propositions 1 and 2 based on these conditions. Finally, in Lemma 10, we present an alternative robust minimax optimization view of our Bayesian model Equation (2).

A.1 First-Order Conditions

Lemma 8 (First-Order Conditions for $\mathbf{h}^*$ and $\mathbf{h}^{\text{sp}}$). *Define*

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

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

$$h_i^* = \frac{1}{2c} (1 - \beta_i \Lambda(\mathbf{h}^*))^2, \quad h_i^{\text{sp}} = \frac{1}{c} (1 - \beta_i \Lambda(\mathbf{h}^{\text{sp}}))^2, \quad \forall i = 1, 2, \dots, N. \quad (24)$$

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

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

the first-order condition gives

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

Differentiating the effective precision $K(\mathbf{h})$, whose calculation is deferred to Lemma 9, yields

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

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

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

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

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

the first-order condition gives

$$\frac{\partial}{\partial h_i} K(\mathbf{h}^{\text{sp}}) = ch_i^{\text{sp}}, \quad \forall i = 1, 2, \dots, N. \quad (26)$$

Equivalently, this suggests

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

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

$$h_i^* = \frac{1}{2c} (1 - \beta_i \Lambda(\mathbf{h}^*))^2, \quad h_i^{\text{sp}} = \frac{1}{c} (1 - \beta_i \Lambda(\mathbf{h}^{\text{sp}}))^2, \quad \forall i = 1, 2, \dots, N,$$

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), $\Lambda(\mathbf{h}^*)$ and $\Lambda(\mathbf{h}^{\text{sp}})$ are both strictly positive. $\square$

Lemma 9 (Effective Precision). *Under Equations (1) and (2), the precision $K(\mathbf{h})$ induced by creators' effort vector $\mathbf{h} = (h_1, h_2, \dots, h_N)'$ is:*

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

Proof. The game proceeds as described in Definition 1. We first take $h_i > 0$ for all i ; the final precision formula extends to zero effort by continuity. In the final stage of data aggregation, the buyer observes $\mathbf{h}$ and signals $\mathbf{s}$. 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:

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

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

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

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

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

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

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

This completes the proof. $\square$

A.2 Underpowered Creative Incentives (Proof of Proposition 1)

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

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

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

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

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

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

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

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

Remark 3 (Individual Creators May Over-Produce). Proposition 1 compares the overall effective precision K under $\mathbf{h}^*$ and $\mathbf{h}^{\text{sp}}$. The stronger per-creator version—namely $h_i^* \leq h_i^{\text{sp}}$ for all i —is incorrect: Consider a 2-creator example with $\beta_1 = 1$ and $\beta_2 = 2$ (hence creator 2 is highly imitative), production cost coefficient $c = 1$, and $(\sigma_\eta^2 + \gamma)^{-1} = \frac{1}{4}$. We first solve the social planner solution $\mathbf{h}^{\text{sp}}$: According to Equations (23) and (24), $\mathbf{h}^{\text{sp}}$ (which is unique according to Proposition 1) satisfies

$$\Lambda(\mathbf{h}^{\text{sp}}) = \frac{h_1^{\text{sp}} + 2h_2^{\text{sp}}}{\frac{1}{4} + h_1^{\text{sp}} + 4h_2^{\text{sp}}}, \quad h_1^{\text{sp}} = (1 - \Lambda(\mathbf{h}^{\text{sp}}))^2, \quad h_2^{\text{sp}} = (1 - 2\Lambda(\mathbf{h}^{\text{sp}}))^2.$$

Solving it gives $h_1^{\text{sp}} = \frac{1}{4}$, $h_2^{\text{sp}} = 0$, and $\Lambda(\mathbf{h}^{\text{sp}}) = \frac{1}{2}$. On the other hand, $\mathbf{h}^*$ (also unique) satisfies

$$\Lambda(\mathbf{h}^*) = \frac{h_1^* + 2h_2^*}{\frac{1}{4} + h_1^* + 4h_2^*}, \quad h_1^* = \frac{1}{2}(1 - \Lambda(\mathbf{h}^*))^2, \quad h_2^* = \frac{1}{2}(1 - 2\Lambda(\mathbf{h}^*))^2.$$

Prove by contradiction: Suppose that $h_i^* \leq h_i^{\text{sp}}$ for all i , which implies $h_2^* = 0$ (since efforts are non-negative), we must have $\Lambda(\mathbf{h}^*) = \frac{1}{2}$, which equals $\Lambda(\mathbf{h}^{\text{sp}})$. Since $\Lambda(\mathbf{h})$ is monotone in h_1 once fixing h_2 , $h_1^* = h_1^{\text{sp}} = \frac{1}{4}$. But $\frac{1}{4} \neq \frac{1}{2}(1 - \frac{1}{2})^2$. This contradicts $h_2^* = 0$ and hence implies $h_2^* > 0 = h_2^{\text{sp}}$.

We further remark that, even if $h_2^{\text{sp}} > 0$, it is still possible that $h_2^* > h_2^{\text{sp}}$: Consider the same example but now $\beta_2 = 1.7$, that is, creator 2 remains highly imitative but to a less extent. Numerically solving Equations (23) and (24), we have $\mathbf{h}^{\text{sp}} \approx (0.245, 0.020)$ and $\mathbf{h}^* \approx (0.162, 0.036)$. In either case, the intuition is the same: Creator 2 should invest no or little effort in the social optimum; but in the equilibrium, due to creator 1's under-investment, their effort becomes (slightly) valuable.

A.3 Originality Penalty (Proof of Proposition 2)

Proof of Proposition 2. To derive the monotonicity of R_i in β_i , we first prove that $\Lambda(\mathbf{h}^*) < \Lambda(\mathbf{h}^{\text{sp}})$ (where $\Lambda(\mathbf{h})$ is defined in Equation (23)). Plugging Equation (24) into the definition of $\Lambda(\mathbf{h}^*)$,

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

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

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

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 β is non-negative and at least one $\beta_i > 0$) since:

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

Therefore, we obtain $\Lambda(\mathbf{h}^*) < \Lambda(\mathbf{h}^{\text{sp}})$. Plug Equation (24) into the definition of $R_i = h_i^*/h_i^{\text{sp}}$:

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

We now view this ratio R_i as a function of β_i (in this fixed market, i.e., we fix $\Lambda(\mathbf{h}^*)$ and $\Lambda(\mathbf{h}^{\text{sp}})$, and consider the β_i 's corresponding to different creators). Formally, let $g(\beta) := \frac{1-\beta\Lambda(\mathbf{h}^*)}{1-\beta\Lambda(\mathbf{h}^{\text{sp}})}$ and $R(\beta) := \frac{1}{2}g(\beta)^2$, so $R_i = R(\beta_i)$ for all $i = 1, 2, \dots, N$. Since we know $\Lambda(\mathbf{h}^*) < \Lambda(\mathbf{h}^{\text{sp}})$,

$$0 < 1 - \beta\Lambda(\mathbf{h}^{\text{sp}}) \leq 1 - \beta\Lambda(\mathbf{h}^*), \quad \forall 0 \leq \beta < 1/\Lambda(\mathbf{h}^{\text{sp}}), \quad (29)$$

hence we know $g(\beta) > 0$ for all $\beta < 1/\Lambda(\mathbf{h}^{\text{sp}})$. Moreover, we also have

$$g'(\beta) = \frac{d}{d\beta} \frac{1 - \beta\Lambda(\mathbf{h}^*)}{1 - \beta\Lambda(\mathbf{h}^{\text{sp}})} = \frac{\Lambda(\mathbf{h}^{\text{sp}}) - \Lambda(\mathbf{h}^*)}{(1 - \beta\Lambda(\mathbf{h}^{\text{sp}}))^2} > 0, \quad \forall \beta,$$

where the last inequality again uses $\Lambda(\mathbf{h}^*) < \Lambda(\mathbf{h}^{\text{sp}})$. Henceforth $R'(\beta) = g(\beta)g'(\beta) > 0$ for all $\beta \in [0, 1/\Lambda(\mathbf{h}^{\text{sp}}))$. According to the condition that $\beta_i < 1/\Lambda(\mathbf{h}^{\text{sp}})$ for all i , R_i is increasing in β_i .

It only remains to lower and upper bound R_i . Since $\beta_i \geq 0$ and R_i is monotone in β_i , we have

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

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

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

Therefore Equation (24) implies $\mathbf{h}^* = h^*\mathbf{1}$ and $\mathbf{h}^{\text{sp}} = h^{\text{sp}}\mathbf{1}$ for scalars $h^* \geq 0$ and $h^{\text{sp}} \geq 0$, such that

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

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

$$h^* \sim \sqrt[3]{\frac{1}{2cN^2\beta^4(\sigma_\eta^2 + \gamma)^2}}, \quad h^{\text{sp}} \sim \sqrt[3]{\frac{1}{cN^2\beta^4(\sigma_\eta^2 + \gamma)^2}}.$$

This gives $\lim_{N \rightarrow \infty} R_i = \lim_{N \rightarrow \infty} \frac{h^*}{h^{\text{sp}}} = \frac{1}{\sqrt[3]{2}}$ for all $i = 1, 2, \dots, N$, as claimed. $\square$

Remark 4 (No Near-Redundant Creator). We now discuss the condition that $\beta_i < 1/\Lambda(\mathbf{h}^{\text{sp}})$ for all i (see Footnote 12 of Proposition 2). We first discuss its two implications, which is the reason why we call it “a market without near-redundant creator” in Footnote 12:

1. **BLUE $\hat{X}$ places all positive weights.** Recall from the proof of Lemma 9 that $\Sigma^{-1} = \mathbf{D} - \frac{(\sigma_\eta^2 + \gamma)\mathbf{D}\beta\beta'\mathbf{D}}{1 + (\sigma_\eta^2 + \gamma)\beta'\mathbf{D}\beta}$. Hence the weight $\hat{X} = (\mathbf{1}'\Sigma^{-1}\mathbf{1})^{-1}\mathbf{1}'\Sigma^{-1}\mathbf{s}$ assigns to creator i is

$$w_i \propto [\Sigma^{-1}\mathbf{1}]_i = h_i - \frac{(\sigma_\eta^2 + \gamma)h_i\beta_i \sum_{j=1}^N h_j\beta_j}{1 + (\sigma_\eta^2 + \gamma) \sum_{j=1}^N h_j\beta_j^2} = h_i(1 - \beta_i\Lambda(\mathbf{h})), \quad \forall i = 1, 2, \dots, N,$$

which is positive when $\beta_i < 1/\Lambda(\mathbf{h})$. Imposed at $\mathbf{h}^{\text{sp}}$, the condition thus makes the optimal estimator weight every creator positively. Should it be violated—thus the BLUE puts a negative weight—this means their highly correlated signal is only used to cancel out the common factor η , hence the name “near-redundant.”

2. **Socially optimum is monotone in correlation.** By Equation (24), $h_i^{\text{sp}} = \frac{1}{c}(1 - \beta_i \Lambda(\mathbf{h}^{\text{sp}}))^2$. On the range $\beta_i \in [0, 1/\Lambda(\mathbf{h}^{\text{sp}}))$, $1 - \beta_i \Lambda(\mathbf{h}^{\text{sp}})$ is positive and decreasing in β_i , which means h_i^{sp} is decreasing in β_i : the planner asks original creators to invest more. Past the threshold this reverses— h_i^{sp} turns back up—and the reason mirrors the previous item: the near-redundant creator becomes valuable again, but purely as a hedge against the common η .

Both readings carry over to the market equilibrium $\mathbf{h}^*$: Since $\Lambda(\mathbf{h}^*) < \Lambda(\mathbf{h}^{\text{sp}})$ (established in the proof of Proposition 2), $\beta_i < 1/\Lambda(\mathbf{h}^{\text{sp}})$ implies $\beta_i < 1/\Lambda(\mathbf{h}^*)$, so $\mathbf{h}^*$ likewise features positive weights and β -monotone efforts. We then state the technical impact of this condition:

1. **Role in the proof of Proposition 2.** This condition ensures Equation (29), which says $g(\beta) = \frac{1 - \beta \Lambda(\mathbf{h}^*)}{1 - \beta \Lambda(\mathbf{h}^{\text{sp}})}$ is positive at all $\beta_1, \beta_2, \dots, \beta_N$. Combined with the subsequent formula that $g'(\beta) > 0$ for all β , this implies the monotonicity of $R(\beta) = \frac{1}{2}g(\beta)^2$. Note that, one may think the facts that $g(0) = \frac{1}{1} = 1$ and $g'(\beta) > 0$ suffice to ensure the monotonicity of $g(\beta)$ and consequently $R(\beta)$. This is untrue, because at $\beta = 1/\Lambda(\mathbf{h}^{\text{sp}})$, the function $g(\beta)$ flips from $+\infty$ to $-\infty$ (similar to $\tan x$: non-monotone but has an always positive derivative $\sec^2 x$).
2. **Proposition 2 Fails without it.** Consider the same example in Remark 3, but now $\beta_2 = 2.2$ (i.e., even more redundant). Now $\mathbf{h}^{\text{sp}} \approx (0.252, 0.009)$, which implies $\beta_2 \Lambda(\mathbf{h}^{\text{sp}}) = 1.096 > 1$. In this case, we have $\mathbf{h}^* \approx (0.173, 0.004)$, hence $R_1 = 0.687$ and $R_2 = 0.490$ —both the monotonicity (implying $R_2 > R_1$) and the lower bound (implying $R_2 > 0.5$) fail.
3. **Per-creator Proposition 1 fails even with it.** One may naturally ask: Given this condition, would the per-creator under-investment variant of Proposition 1—discussed in Footnote 11 and Remark 3—now be true? The answer remains negative. In the second example of Remark 3 (with $\beta_2 = 1.7$), we have $\mathbf{h}^{\text{sp}} \approx (0.245, 0.020)$ and hence $\Lambda(\mathbf{h}^{\text{sp}}) \approx 0.505$. In this case $\beta_2 \Lambda(\mathbf{h}^{\text{sp}}) \approx 0.858 < 1$, but we still have $h_2^* > h_2^{\text{sp}}$. Should we really want a per-creator under-investment result, we would need $\beta_i \leq \frac{\sqrt{2}-1}{\sqrt{2}\Lambda(\mathbf{h}^{\text{sp}}) - \Lambda(\mathbf{h}^*)}$ for all i (a sufficient condition is $\beta_i \leq (1 - \frac{1}{\sqrt{2}})/\Lambda(\mathbf{h}^{\text{sp}}) \approx 0.293/\Lambda(\mathbf{h}^{\text{sp}})$, which is strictly stronger than $\beta_i < 1/\Lambda(\mathbf{h}^{\text{sp}})$).

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

Lemma 10 (Minimax Robust Optimization View). *The effective precision $K(\mathbf{h})$ 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 $\mu_\eta^2 \leq \gamma$. Formally, let the buyer seek a linear estimator $\hat{X} = \mathbf{w}'\mathbf{s}$ to minimize the worst-case MSE:*

$$\min_{\mathbf{w}} \max_{\mu_\eta^2 \leq \gamma} \mathbb{E}[(\mathbf{w}'\mathbf{s} - X)^2], \quad \text{s.t. } \mathbf{w}'\mathbf{1} = 1. \quad (30)$$

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

Proof. We decompose the MSE of any linear estimator $\hat{X} = \mathbf{w}'\mathbf{s}$ 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), $\mathbf{s} = \mathbf{1}X + \boldsymbol{\beta}\eta + \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} = (\varepsilon_1, \dots, \varepsilon_N)'$, $\mathbb{E}[\eta] = \mu_\eta$, and $\mu_\eta^2 \leq \gamma$. The variance depends only on the fluctuations around the mean, not on

μ_η itself. Therefore, letting $\mathbf{D} = \text{diag}(h_1, \dots, h_N)$ and $\Sigma_{\text{noise}} = \mathbf{D}^{-1} + \sigma_\eta^2 \beta \beta'$, we know $\text{Var}(\hat{X}) = \mathbf{w}' \Sigma_{\text{noise}} \mathbf{w}$ is independent of μ_η . On the other hand, since $\mathbf{w}' \mathbf{1} = 1$, the mean of $\hat{X}$ is $\mathbb{E}[\hat{X}] = \mathbf{w}' (\mathbf{1} X + \beta \mu_\eta) = X + \mu_\eta (\mathbf{w}' \beta)$. The squared bias term therefore equals $\mu_\eta^2 (\mathbf{w}' \beta)^2$.

This reveals that the inner maximization w.r.t. μ_η in Equation (30) is equivalent to $\max_{\mu_\eta^2 \leq \gamma} \mu_\eta^2 (\mathbf{w}' \beta)^2$. The maximum is clearly attained at the boundary $\mu_\eta^2 = \gamma$. Thus, the buyer's problem in Equation (30) simplifies to:

$$\min_{\mathbf{w}} (\mathbf{w}' \Sigma_{\text{noise}} \mathbf{w} + \gamma (\mathbf{w}' \beta)^2) \quad \text{s.t. } \mathbf{w}' \mathbf{1} = 1.$$

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

$$\min_{\mathbf{w}} \mathbf{w}' (\Sigma_{\text{noise}} + \gamma \beta \beta') \mathbf{w} \quad \text{s.t. } \mathbf{w}' \mathbf{1} = 1.$$

Let $\Sigma = \Sigma_{\text{noise}} + \gamma \beta \beta' = \mathbf{D}^{-1} + (\sigma_\eta^2 + \gamma) \beta \beta'$. The Lagrangian for this problem is:

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

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

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

Since this is the covariance matrix Σ in the Bayesian setting (see the proof of Lemma 9), this is exactly the formula for BLUE weights in the Bayesian setting. Furthermore, the worst-case MSE induced by $\mathbf{w}^*$ is

$$\text{MSE}^* = \mathbf{w}^{*'} \Sigma \mathbf{w}^* = (\nu \mathbf{1}' \Sigma^{-1}) \Sigma (\nu \Sigma^{-1} \mathbf{1}) = (\mathbf{1}' \Sigma^{-1} \mathbf{1})^{-1},$$

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

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 rank-and-file creator-innovator ratio $\rho(p)$. We then prove Propositions 3 and 4, which together constitute for our “curse of precision” observation.

B.1 Supply Elasticity and Monotonicity of Ratio

Lemma 11 (Supply Elasticity and Monotonicity of Ratio). *The condition in Equation (7), namely $\mathcal{E}_C(p) > \mathcal{E}_O(p)$ for all $p > 0$ where $\mathcal{E}_j(p) := \frac{d \ln N_j(p)}{d \ln p}$ ($\forall j \in \{O, C\}$), is equivalent to*

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

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

$$0 < \frac{d}{dp} \frac{N_C(p)}{N_O(p)} = \frac{N'_C(p) N_O(p) - N_C(p) N'_O(p)}{(N_O(p))^2},$$

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 $\mathcal{E}_C(p) > \mathcal{E}_O(p)$. $\square$

B.2 Upper Bound on Market Price (Proof of Proposition 3)

Proof of Proposition 3. 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): Let $a := \frac{\partial K}{\partial S_O}$ and $b := \frac{\partial K}{\partial S_C}$. By Lemma 12, $0 < a \leq 1$ and $0 < b \leq \kappa_C$.

$$\begin{aligned} \Pi_{\text{inst}} &= \frac{dK}{dt} - \left(N_O(p) + N_C(p)\right)ph(p) \\ &\stackrel{(a)}{=} \frac{\partial K}{\partial S_O} \frac{dS_O}{dt} + \frac{\partial K}{\partial S_C} \frac{dS_C}{dt} - \frac{p^2}{c} \left(N_O(p) + N_C(p)\right) \\ &\stackrel{(b)}{=} a \left(N_O \frac{p}{c} - \delta S_O\right) + b \left(N_C \frac{p}{c} - \delta S_C\right) - \frac{p^2}{c} (N_O + N_C) \\ &= \frac{p}{c} \left(N_O(a - p) + N_C(b - p)\right) - \delta(aS_O + bS_C), \end{aligned}$$

where (a) uses the Law of Total Derivatives, and (b) uses the best effort formula $h(p) = p/c$ and 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 $-\delta(aS_O + bS_C)$ 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, using $a \leq 1$ and $b \leq \kappa_C$, the market price $p(t)$ must satisfy (again omitting all t 's for readability):

$$N_O(p)(1 - p) + N_C(p)(\kappa_C - p) \geq N_O(p)(a - p) + N_C(p)(b - p) \geq 0.$$

Rearranging for p gives $p \leq \kappa_C + \frac{1 - \kappa_C}{1 + \rho(p)}$, where we recall $\rho(p) = \frac{N_C(p)}{N_O(p)}$ is the rank-and-file creator-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 $\bar{p}(\kappa_C)$, such that

$$\bar{p}(\kappa_C) = \kappa_C + \frac{1 - \kappa_C}{1 + \rho(\bar{p}(\kappa_C))}, \quad \forall \kappa_C \in (0, 1]. \quad (31)$$

Again invoking the Implicit Function Theorem to Equation (31) (whose detailed calculation is deferred to Lemma 13), we are able to prove $\bar{p}'(\kappa_C) > 0$, i.e., $\bar{p}$ is increasing in κ_C . This further establishes the continuity of $\bar{p}$, which gives

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

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

Lemma 12 (Marginal Effective Precision). *The effective precision $K(t)$ defined in Equation (11) satisfies*

$$0 < \frac{\partial K}{\partial S_O}(t) \leq 1, \quad 0 < \frac{\partial K}{\partial S_C}(t) \leq \kappa_C(t), \quad \forall t \geq 0,$$

where $\kappa_C(t)$ is defined in Equation (13), which we recall as

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

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

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

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

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

Calculating the partial derivatives of $\mathcal{F}$ with respect to S_O , S_C , and K gives:

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

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

$$0 < \frac{\partial K}{\partial S_O} = \frac{1}{\partial \mathcal{F} / \partial K} \leq 1, \quad 0 < \frac{\partial K}{\partial S_C} = \frac{1}{(1 + (\sigma_\eta^2 + \lambda K) S_C \beta_C^2)^2} \frac{1}{\partial \mathcal{F} / \partial K} \leq \frac{1}{(1 + (\sigma_\eta^2 + \lambda K) S_C \beta_C^2)^2}.$$

Plugging in the definition of $\kappa_C(t)$ gives the conclusion. $\square$

Lemma 13 (Monotonicity of Price Upper Bound). *The price upper bound $\bar{p}(\kappa_C)$ derived in Equation (31), namely*

$$\bar{p}(\kappa_C) = \kappa_C + \frac{1 - \kappa_C}{1 + \rho(\bar{p}(\kappa_C))}, \quad \text{where } \rho(p) = \frac{N_C(p)}{N_O(p)},$$

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

Proof. From Equation (31), the price upper bound $\bar{p}(\kappa_C)$ is given by the implicit equation:

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

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

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

The numerator is positive because $\rho(p) > 0$ for all p , and the denominator is also positive because $\kappa_C \in (0, 1]$ and also $\rho'(p) > 0$ for all p (from Equation (7) and Lemma 11). Therefore, $\bar{p}'(\kappa_C) > 0$ for all $\kappa_C \in (0, 1]$, as claimed. $\square$

B.3 Barrier of Model Precision (Proof of Proposition 4)

Proof of Proposition 4. 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 ($S_O(t) \rightarrow 0$), this equation becomes

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

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

$$1 = \lim_{S_C(t) \rightarrow \infty} K(t) \left(S_C(t)^{-1} + (\sigma_\eta^2 + \lambda K(t)) \beta_C^2 \right) = \lim_{S_C(t) \rightarrow \infty} K(t) (\sigma_\eta^2 + \lambda K(t)) \beta_C^2.$$

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

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

Solving for the unique positive root yields the expression in Equation (15). $\square$

C Proofs for Section 5

In this section, we prove Lemmas 5 and 6, and Proposition 7.

C.1 Linear Pricing Fails on Participation (Proof of Lemma 6)

Proof of Lemma 5. For each creator $i = 1, 2, \dots, N$, under the linear pricing rule (or the two-part tariff rule considered later in Proposition 7), their payoff $U_i(\mathbf{h}; \mathbf{P})$ only depends on h_i and p_i (and also B_i if considering the two-part tariff rule defined in Equation (18)), namely $U_i(\mathbf{h}; \mathbf{P}) = p_i h_i + B_i - C(h_i)$. Thus the first-order condition is simply

$$p_i = ch_i.$$

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 $\mathbf{P}$ attaining IC, i.e., $P_i(\mathbf{h}) = ch_i^{\text{sp}} \times h_i, \forall i = 1, 2, \dots, N$. $\square$

Proof of Lemma 6. According to Lemma 5, the $\mathbf{P}$ is uniquely determined by Equation (IC). We further pin down T using Budget Balance (BB). This gives

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

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

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

and also the profit of the buyer:

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

Since $p_i = ch_i^{\text{sp}}$ (from Lemma 5) and the social planner solution ensures (recall Equation (26)) $\nabla K(\mathbf{h}^{\text{sp}}) = c\mathbf{h}^{\text{sp}}$, we have

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

In the special case where all creators are uncorrelated (i.e., $\beta_i = 0$ for all i), we have $K(\mathbf{h}) = \sum_{i=1}^N h_i$ (recall Lemma 9).²² Euler's Homogeneous Function Theorem gives

$$K(\mathbf{h}) = \sum_{i=1}^N h_i \cdot \frac{\partial K}{\partial h_i}(\mathbf{h}) = (\nabla K(\mathbf{h}))' \mathbf{h}, \quad \forall \mathbf{h} \in \mathbb{R}_{\geq 0}^N. \quad (32)$$

Therefore, the buyer's profit is $\Pi = -F \leq 0$ in this case, which violates Equation (IR) whenever $F + \underline{\Pi} > 0$. $\square$

C.2 Two-Part Tariff for Efficiency and Participation (Proof of Proposition 7)

Proof of Proposition 7. From the specific choice of unit price $\mathbf{p} = c\mathbf{h}^{\text{sp}}$, Incentive Compatibility (IC) holds due to Lemma 5. Since $\phi = (\phi_1, \phi_2, \dots, \phi_N)'$ sums to 1, at the equilibrium effort $\mathbf{h}^{\text{sp}}$ the total payment to the creators is (recall the B_i defined in Equation (21)):

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

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

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

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

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

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

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

where the last step is because $\mathcal{S}(\mathbf{h}^{\text{sp}}) > 0$ (the condition of Proposition 7). The buyer's profit:

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

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

²²Or more generally, in any Constant Returns to Scale (CRS) regime where $K(a\mathbf{h}) = aK(\mathbf{h})$ for any scalar $a > 0$.