# Protection or Exploitation?
# AI Copyright Regulation and Endogenous Market Structure

Yiyang Sun*

*Guanghua School of Management, Peking University*

*Ross School of Business, University of Michigan*

yiyangss@umich.edu

March 27, 2026

**Abstract**

Generative AI introduces a fundamental tension in the creative economy: property rights remain undefined at both the input margin of training data and the output margin of generated content. To analyze this dual-margin disruption, we develop a general equilibrium framework featuring AI-adopting art production firms alongside human artists who supply training data, corporate labor, and independent art. We formalize copyright compliance through explicit data royalties and probabilistic judicial enforcement against infringement, yielding three equilibrium regimes: data starvation, interior AI adoption, and no AI adoption. We show that policies that appear to raise artist income on impact may lose much of that force after equilibrium responses operate through market concentration, the reallocation of production and work, and the conversion of nominal income into real welfare. To quantify these margins, we evaluate artist welfare over the full two-dimensional policy space under a maintained post-shock parameterization and then assess the robustness of the results to alternative technological and market conditions and find that in the current parameter environment, real artist welfare is highest under the least restrictive combination of copyright instruments, with both the royalty burden and judicial intervention kept low.

**JEL Classification:** E24, L16, O33, O34, D50

**Keywords:** Generative AI, Copyright, Labor Monopsony, General Equilibrium, Market Structure

*I am grateful for the ongoing guidance from Mi Luo, Matthew Schwartzman, Fanqi Shi, Jingyuan Wang, Wei Xiang, Jidong Zhou. This is a preliminary draft prepared for the 2026 NBER SI submission. All errors are my own.

# 1 Introduction

The 2023 Hollywood Writers Guild (WGA) strike represents a historic watershed in modern labor relations. Screenwriters halted productions for 148 days not merely over traditional wage disputes, but to demand strict contractual guardrails against studios using Artificial Intelligence to generate or rewrite scripts. This conflict highlights how generative AI's deployment in the creative and intellectual property (IP) economy differs fundamentally from its application in other sectors, such as autonomous driving or predictive logistics. While the latter rely on objective physical data to optimize processes, generative AI introduces an endogenous displacement mechanism: its automation capital (model weights) is constructed directly from ingesting the copyrighted, expressive labor of the very human creators it aims to displace. Compounding this technological shock is a profound legal vacuum. Because AI systems lack legal personhood, their pure generated outputs are inherently excluded from traditional copyright frameworks. This fundamental axiom forces society to entirely redefine property rights at both the "input" (training data) and "output" (generated content) margins of the AI production function (see, e.g., Samuelson, 2023).

In response to this disruption, policymakers and incumbent rights-holders have pursued a pincer intervention strategy aimed at protecting human creators. On the input margin, aggressive litigation by industry incumbents—such as the landmark lawsuits filed by Universal, Sony, and Warner against AI music startups like Suno and Udio—has weaponized copyright law. Macroeconomically, this dense legal risk manifests as a severe expected data royalty, imposing staggering upfront costs for legitimate model training. Concurrently, on the output margin, the judicial system has established stringent "human authorship" thresholds. For instance, in the landmark case *Thaler v. Perlmutter*, a U.S. federal court explicitly denied copyright protection to an artwork entirely generated by an AI algorithm, cementing human input as an essential legal prerequisite. Consequently, an AI firm's expected copyright appropriability is no longer a guaranteed constant. Instead, it becomes an endogenous probability determined jointly by prevailing judicial strictness and the firm's chosen degree of automation. Intuition suggests that exorbitant data costs extract rents for creators, while strict human authorship mandates preserve human labor demand, ostensibly achieving a symbiotic equilibrium.

However, this partial-equilibrium intuition severely overlooks the profound general equilibrium distortions induced by altering the industrial organization of the creative sector. We highlight a core macroeconomic tension: exorbitant data costs function as an insurmountable endogenous barrier to entry. By systematically purging undercapitalized independent creators and smaller entrants from the market, this compliance burden creates strong economic incentives for market concentration, yielding a highly consolidated oligopoly. Paradoxically, this regulatory-induced concentration endows surviving tech giants with oligopsonistic power in the labor market. Thus, interventions intended to enrich creators via data rents may instead severely depress equilibrium wages at the intensive margin through monopsonistic markdowns.

Concurrently, the output-side intervention further distorts factor allocation. While strict judicial thresholds aim to protect human jobs, from a profit-maximizing perspective, they act as a direct penalty on the firm's expected copyright appropriability. This expected revenue penalty tightens the firm's zero-profit condition, further accelerating the exit of marginal firms and exacerbating market consolidation. Moreover, to maintain monopoly markups and prevent their products from falling into the perfectly competitive public domain, surviving firms face intense pressure to artificially retain human labor in their production functions. In this context, human labor is retained not necessarily for its inherent productivity advantage, but as a mandatory "compliance input" to bypass legal hurdles.

The general equilibrium interaction of these dual market powers generates complex, non-monotonic welfare dynamics. In the product market, the proliferation of AI-synthesized media acts as a direct substitute

for independent human art, exerting downward pressure on independent prices. Simultaneously, the monopsonistic wage collapse reduces the opportunity cost of human creation. Crucially, this system is bounded by a static physical endowment constraint: aggregate AI data ingestion cannot physically exceed the total stock of human independent creation. If the dual market squeeze causes independent human creation to severely contract, the economy endogenously hits this corner solution. This reveals a fatal macroeconomic feedback loop: excessive factor distortions not only exploit labor but ultimately constrain the physical data supply required for the AI sector's own technological expansion.

To formalize this intuition and resolve the theoretical ambiguity, this paper asks: In the era of generative AI, how do strict copyright interventions endogenously reshape the market structure of the creative economy? Furthermore, under what parameter spaces can this regulatory-induced dual market power (product market oligopoly and labor market oligopsony) achieve a Pareto improvement for human creators, and under what conditions does it trigger general equilibrium factor distortions, structural collapse, and technological stagnation? To answer these questions, we develop a tractable general equilibrium framework featuring endogenous entry, Cournot oligopoly competition, labor monopsony, and physical data constraints, providing rigorous micro-foundations for optimal intellectual property policy.

The model architecture rests on three micro-foundational pillars. First, we capture the multidimensional roles of human creators and their endogenous time allocation. The representative household acts simultaneously as a consumer with nested CES preferences (establishing independent human art and AI-synthesized corporate media as gross substitutes) and as a producer subject to strict time and budget constraints. Creators endogenously allocate their labor between independent artistic creation and formal employment by AI firms. Consequently, aggregate creator welfare is endogenously determined by a tripartite income loop: sales from independent art, data royalties extracted from the corporate sector, and wages suppressed by oligopsonistic markdowns.

Second, AI firms engage in a two-stage game of technology adoption and Cournot competition. In the first stage, potential entrants decide whether to enter the market by sinking a fixed data acquisition cost dictated. Drawing on the task-based framework of Acemoglu and Restrepo (2018), this data investment endogenously determines the technological frontier of automation. In the second stage, having established their technological boundaries, surviving firms optimize their capital and labor allocations and compete in quantities within the product market. Crucially, firms are not wage-takers. Instead, they internalize their consolidated market shares to exert oligopsonistic monopsony power over the remaining labor supply.

Third, we introduce a novel mechanism of ex-ante symmetric entry with ex-post idiosyncratic legal shocks to model institutional friction. While all potential entrants face identical initial data costs, they are subjected to idiosyncratic judicial scrutiny after production. Firms draw a random judicial strictness parameter from an exogenous distribution. This legal friction embeds expected copyright appropriability directly into the firm's objective function. To maximize expected profits and prevent their outputs from falling into the public domain, firms face an endogenous incentive to retain human labor not solely for its inherent productivity, but to satisfy legal prerequisites for intellectual property protection. This transforms a standard Cournot model into a rich framework featuring ex-post firm heterogeneity without sacrificing macro-level tractability.

We define and solve for a Symmetric Free-Entry General Equilibrium. The GE closure requires that corporate monopoly profits flow back into consumer expenditure, markets clear simultaneously, and the system remains bounded by a macroeconomic data endowment constraint. This static physical limit ensures that aggregate AI data ingestion can never exceed the total output of independent human creation. Within this environment, the paper asks how two copyright policy instruments shape equilibrium incidence in creative markets. The first instrument is the data royalty rate $\lambda$, which governs the transfer paid for lawful use of human-created content in model training. The second is judicial strictness $\bar{\xi}$, which governs the extent to

which AI-assisted output remains legally protectable against imitation and thus changes the private return to automation. The question is how those margins are transformed once firm entry, market concentration, data use, labor demand, and relative prices adjust endogenously.

The theoretical analysis shows that these adjustments must be read through equilibrium transmission rather than through the direct policy margin alone. A higher royalty rate mechanically strengthens transfers associated with data use, and greater judicial strictness mechanically tightens the appropriability environment facing AI producers. Yet these immediate effects do not by themselves determine artist welfare. Policy also changes the equilibrium mass of firms, and that change in market structure propagates into product-market markup power and labor-market power. The paper therefore distinguishes within-structure policy incidence from endogenous market-structure transmission and uses this distinction to organize the incidence logic of the model. The resulting message is disciplined and economic: interventions that appear favorable at the point of transfer or legal protection can induce equilibrium responses that attenuate, and in some regions overturn, their apparent direct gains.

This logic is carried to the artist side through an exact decomposition of nominal artist income into independent sales income, royalty income, and corporate-sector wage income. That decomposition clarifies two central offsetting forces. The first is a market-structure offset, through which policy-induced changes in firm mass alter concentration and thereby reshape the prices and wages that determine artist income. The second is a reallocation offset, through which policy changes the composition of activity across independent creation, royalty-generating data use, and corporate employment. These nominal channels are informative about transmission, but they are not the final welfare object. Once nominal receipts are translated into purchasing power through the aggregate price index, a third offset emerges in the form of real incidence. The paper's welfare analysis is built around the interaction of these three forces.

The quantitative section then asks whether these equilibrium offsets are quantitatively important under a maintained parameterization disciplined by the U.S. motion picture and entertainment sector. The calibration evaluates outcomes over the full two-dimensional policy space spanned by $\lambda$ and $\bar{\xi}$, beginning from the regime structure implied by the theory and then tracing how policy reshapes firm entry, data use, market concentration, nominal artist-income components, and real artist welfare. It also examines how these results vary with the technological environment and with the deep parameters that govern labor-market transmission, cross-sector substitution, and the responsiveness of automation to data. The quantitative objective is to determine whether the immediate protective or redistributive appeal of tighter copyright policy survives once equilibrium market structure, resource allocation, and prices are allowed to adjust jointly. Under the baseline post-shock parameterization, the answer is often negative. Policy combinations that improve some nominal artist-income margins can nonetheless reduce real artist welfare once concentration rises, production is reallocated, and the aggregate price level moves against artists.

The paper's contribution is therefore to develop a general-equilibrium incidence framework for AI copyright regulation that links legal design to industrial organization, labor-market power, and creator welfare within a unified equilibrium system. The analysis does not treat copyright as a pure transfer problem, nor does it infer artist outcomes from any single nominal margin. It instead shows how copyright policy reshapes the organization of production, how that reorganization propagates through market structure and equilibrium reallocation, and how those adjustments govern the final distribution of gains and losses between the corporate AI sector and human creators. This perspective places endogenous market structure at the center of the economic interpretation of AI regulation and yields a policy reading that is structurally richer than either a purely technological or a purely legal account. This incidence perspective also helps situate the paper within adjacent work on market power, labor-market concentration, and technological change. The macroeconomic analysis of labor markets increasingly centers on monopsony power (Manning, 2003), with robust

empirical evidence documenting that high employer concentration in specialized local labor markets directly generates significant wage markdowns (Azar et al., 2022). To endogenize this market concentration, we build upon the canonical firm dynamics and selection framework of Hopenhayn (1992). We embed this mechanism of policy-driven survival thresholds into an industry environment characterized by ex-post idiosyncratic legal risks, where copyright compliance costs and judicial strictness act as an endogenous filtration mechanism that dictates the finite mass of surviving firms at the extensive margin. Given this endogenous firm mass, we adopt the general equilibrium macro-oligopsony structure established by Berger et al. (2022). We utilize their insight that a finite number of firms internalize their labor market shares to exert monopsonistic market power, translating an oligopolistic market structure into suppressed equilibrium wages. Building on these theoretical foundations, we apply this macro-oligopsony framework to a novel environment characterized by intellectual property frictions. Through this synthesis, we propose a mechanism of regulatory-induced monopsony, demonstrating that data acquisition costs and human copyright protections intended to shield creators function instead as fixed barriers that purge marginal competitors, ultimately endowing surviving tech giants with monopsonistic pricing power over the same labor pool.

Building on the foundational task-based framework of Autor et al. (2003), we adopt the continuum of tasks model established by Acemoglu and Restrepo (2018) to characterize the substitution of human labor by algorithmic capital. Recent empirical literature documents that artificial intelligence functionally manifests as task displacement at the firm level (Acemoglu and Restrepo, 2022), providing the microeconomic rationale for applying this automation boundary to generative media. Standard macroeconomic models typically assume the automation frontier is an exogenous technological shock or endogenize it through aggregate research and development and wage dynamics (Hémous and Olsen, 2022). We depart from this paradigm by endogenizing the microeconomic automation frontier as a direct function of data ingestion. Under this novel structure, the technology adoption decision is governed not solely by relative factor prices, but by idiosyncratic legal uncertainty and data compliance costs. This structural mechanism reveals that strict copyright enforcement distorts the technology adoption trajectory rather than merely shifting marginal costs. We demonstrate that once regulatory burdens cross a critical threshold, surviving oligopolists strategically abandon artificial intelligence and revert to pure manual production to evade legal risks, establishing how data policy can endogenously precipitate technological retreat.

Standard legal and economic frameworks evaluate intellectual property through a partial equilibrium tradeoff between stimulating innovation and generating deadweight loss through rent extraction (see, e.g., Landes and Posner 1989; Boldrin and Levine 2002). Elevating this perspective to macroeconomic industry dynamics, recent literature documents how intellectual property protections function as endogenous entry barriers that consolidate incumbent market power and depress business dynamism (Akcigit and Ates, 2021). We adopt this conceptualization of intellectual property as a structural entry deterrent to formalize copyright compliance as a fixed barrier within a general equilibrium environment. In the specific context of generative artificial intelligence, contemporary scholarship investigates dual legal frictions on both data input and output margins, though these frameworks typically model such policies as intertemporal innovation incentives within frictionless markets (Yang and Zhang, 2024). Departing from models that analyze intellectual property primarily as dynamic incentive shifters, we direct our focus to the distortion of static endogenous market structures. We formalize expected judicial strictness and data taxation as fixed and marginal compliance costs within a Cournot game. This structural departure establishes that copyright interventions do not merely mitigate data scarcity. Instead, these policies operate as insurmountable moats that trigger nonmonotonic structural breaks at the extensive margin, ultimately synthesizing oligopolistic market power endogenously.

The contemporary macroeconomic paradigm establishes data as a nonrival factor of production, where information generated as a byproduct of economic activity fuels a positive feedback loop that theoretically

drives increasing returns and unbounded growth (see, e.g., Jones and Tonetti 2020; Farboodi and Veldkamp 2021). Recent interdisciplinary research disrupts this assumption of infinite data accumulation by mathematically documenting the mechanism of model collapse, which demonstrates that recursive algorithmic training causes severe distribution degradation and necessitates a continuous influx of novel human-generated data (Shumailov et al., 2024; Villalobos et al., 2024). While computer science literature identifies this impending data wall primarily as a passive engineering constraint, standard macroeconomic frameworks overlook the inherently cannibalistic dependency of generative artificial intelligence on the exact human labor it seeks to displace. We bridge this gap by embedding a strict physical data endowment constraint within a general equilibrium architecture. Because algorithmic media and independent human creation function as gross substitutes in the product market, the proliferation of synthesized content crowds out human production, which systematically contracts the aggregate supply of viable training data. This synthesis reveals that product market competition triggers an endogenous macroeconomic feedback loop resulting in a data starvation bound, proving that data nonrivalry cannot guarantee perpetual expansion when technological adoption physically eradicates its own fundamental input.

The remainder of the paper is organized as follows. Section 2 sets up the model environment. Section 3 formalizes the firm optimization problem and defines the general equilibrium. Section 4 characterizes the theoretical results regarding market structure, technology adoption, and the physical data constraint. Section 5 introduces the quantitative analysis and calibrates the framework to evaluate welfare. Section 6 concludes. All theoretical derivations and proofs are relegated to the Appendix.

# 2 Model Environment

## 2.1 Demand, Expenditure, and Human Art Supply

**Utility and Price Indices.** The aggregate consumption bundle $C$ and the corporate composite $Q_{Corp}$ are defined as:

$$C = \left[\omega^{\frac{1}{\sigma}} Y_H^{\frac{\sigma-1}{\sigma}} + (1-\omega)^{\frac{1}{\sigma}} Q_{Corp}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{1}$$

$$Q_{Corp} = \left[\sum_{i=1}^{N} q_i^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{2}$$

where $\sigma$ is the elasticity of substitution between sectors, and $\eta$ is the elasticity of substitution across AI varieties. Following the standard macro-IO literature (Atkeson and Burstein, 2008), we impose the parametric restriction $\eta > \sigma > 1$. Economically, this implies that creative varieties synthesized within the AI sector are closer substitutes for one another than they are for authentic human art. This intuitive nesting structure mathematically guarantees the standard pro-competitive effect: an increase in the number of competitors $N$ strictly increases the perceived demand elasticity $\varepsilon(N)$, thereby disciplining oligopolistic markups.

The corresponding ideal price indices dual to the utility aggregators are:

$$P = \left[\omega P_H^{1-\sigma} + (1-\omega) P_{Corp}^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{3}$$

$$P_{Corp} = \left[\sum_{i=1}^{N} p_i^{1-\eta}\right]^{\frac{1}{1-\eta}} \tag{4}$$

Under a symmetric equilibrium with $N$ firms charging price $p$, the sectoral price index implies $P_{Corp} =$

$N^{\frac{1}{1-\eta}}p$, capturing the standard variety effect.

**Demand Functions.** Expenditure minimization implies the following iso-elastic demand schedule for each AI variety $i$:

$$q_i(p_i) = \underbrace{\left(\frac{p_i}{P_{Corp}}\right)^{-\eta}}_{\text{Within-Sector}} \cdot \underbrace{(1-\omega)\left(\frac{P_{Corp}}{P}\right)^{-\sigma}\frac{E}{P}}_{\text{Cross-Sector}} \tag{5}$$

where $E$ denotes aggregate nominal expenditure.

**Aggregate Expenditure.** To close the consumer's problem, we define aggregate nominal expenditure $E$ as the sum of all revenues in the economy. In a general equilibrium where all profits and royalties are redistributed to the representative consumer, total expenditure equals total gross production:

$$E = P_H Y_H + P_{Corp} Q_{Corp} \tag{6}$$

**Human Art Supply** Unlike the standard labor pool, artistic talent is scarce. We model the production of independent art $Y_H$ using a specific technology that exhibits decreasing returns to scale in labor time $L_H$:

$$Y_H = A_H L_H^{\alpha}, \quad \alpha \in (0,1) \tag{7}$$

Profit maximization by competitive artists (maximizing revenue net of the opportunity cost of their time evaluated at the market wage $w$) yields the following problem:

$$\max_{L_H} \quad P_H \cdot A_H L_H^{\alpha} - wL_H$$

Taking the first-order condition with respect to $L_H$ and substituting the artist time constraint yields the updated upward-sloping supply curve for human art, which is now explicitly linked to the aggregate macroeconomic wage:

$$P_H = \frac{w}{\alpha A_H^{1/\alpha}} Y_H^{\frac{1-\alpha}{\alpha}}$$

## 2.2 Technology, Data, and Copyright Protection

**Task-Based Production and Aggregation** Following Acemoglu and Restrepo (2018), we model the production of the AI-generated good $y_i$ as the aggregation of a continuum of tasks indexed by $j \in [0,1]$.

Each task $j$ can be performed using either labor $l(j)$ or AI capital $k(j)$ combined with data:

$$y_i(j) = \begin{cases} A_L l(j) + A_K k(j) & \text{if } j \in [0, I(D)] \\ A_L l(j) & \text{if } j \in (I(D), 1] \end{cases} \tag{8}$$

**The Automation Threshold $I(D)$.** The threshold $I(D)$ represents the frontier of tasks that are technologically feasible for AI automation. We assume this threshold is endogenous to the aggregate stock of training data $D$. Following the statistical learning literature, prediction accuracy follows a power-law learning curve, translating into an S-curve adoption function:

$$I(D) = \frac{D^{\mu}}{D^{\mu} + \delta} \tag{9}$$

where $\mu > 0$ governs the learning rate and $\delta > 0$ represents intrinsic complexity.

**Copyright Appropriability and Judicial Strictness.** We model copyright protection as a continuous appropriability share that determines the fraction of a firm's residual demand shielded from perfect imitation.[1] Let $I(D) \in [0,1]$ denote the AI intensity of the production process. The legal system grants exclusive rights only to the human-authored component of output, so appropriability declines with AI intensity. We summarize this force by the reduced-form schedule

$$\bar{\beta}(I, \bar{\xi}) = (1 - I(D))^{\bar{\xi}}, \tag{10}$$

where $\bar{\xi} \geq 0$ indexes judicial strictness. A higher $\bar{\xi}$ reduces appropriability for any given AI intensity, and a higher $I(D)$ reduces appropriability for any given judicial environment.

The unprotected fraction $1 - \bar{\beta}(I, \bar{\xi})$ of the product is deemed AI-generated and falls into the public domain, where it is perfectly substituted by a competitive fringe supplying at marginal cost $MC(D)$. Consequently, the firm acts as a monopolist only over the protected share of its residual demand curve.

**Pricing and Appropriability.** Because $\bar{\beta}(I, \bar{\xi})$ acts as a strictly multiplicative scaling factor on appropriable operating profit, it drops out of the first-order condition with respect to price. This property preserves symmetric ex-post pricing across active firms and allows the model to isolate the macroeconomic effects of endogenous entry and market structure without introducing within-sector heterogeneity in market shares.

**Implied Aggregate Technology.** Under optimal task allocation, the task-based production structure admits the following aggregate representation for firm $i$:

$$Y_i = B(D_i) \cdot \left(\frac{K_i}{I(D_i)}\right)^{I(D_i)} \left(\frac{L_i}{1 - I(D_i)}\right)^{1 - I(D_i)}, \tag{11}$$

where

$$B(D_i) \equiv A_K^{I(D_i)} A_L^{1 - I(D_i)}. \tag{12}$$

Appendix B derives this expression from a task-based framework in the spirit of Acemoglu and Restrepo (2018) under unit elasticity across tasks.

**Data Market** We model data as a **non-rivalrous** productive asset. Since digital data can be replicated at zero marginal cost, multiple firms can utilize the same dataset without diminishing its availability to others. Consequently, the cost of data does not depend on the number of firms $N$.

**Data Pricing.** Firms pay a royalty rate $\lambda > 0$ to artists for the right to use their works as training data. The unit price of data is linear in the price of human art:

$$P_D = \lambda P_H \tag{13}$$

[1] This fractional protection aligns with contemporary judicial practices regarding AI-generated content. For instance, in the 2023 *Zarya of the Dawn* ruling, the U.S. Copyright Office granted protection to the human-authored arrangements but explicitly denied protection to the AI-generated components, placing the latter into the public domain. One microfoundation for the reduced-form object below is an environment with idiosyncratic ex-post litigation shocks, in which firm-level appropriability is realized only after production and enforcement. In the symmetric free-entry equilibrium studied here, however, only the ex-ante expected appropriability schedule matters for pricing, entry, and data choice, so the model works directly with that reduced-form object.

where $\lambda$ is a policy parameter representing the strength of data royalty enforcement.

**The Physical Constraint and Data Depreciation.** While the technological frontier $I(\cdot)$ is conceptually shaped by the stock of training data, the economically relevant stock is maintained only through a contemporaneous flow of fresh data. In a stationary equilibrium, we therefore represent automation intensity as a reduced-form function of the current data flow $D_i$, with the corresponding stock-flow proportionality absorbed into the calibrated learning parameters. Because firms must legally clear rights for the human art they ingest, data use is bounded by the aggregate supply of human art produced within the period:

$$D_i \leq Y_H. \tag{14}$$

## 2.3 Labor Supply and Market Structure

**Utility and Labor Supply.** We adopt a quasi-linear aggregate utility function,

$$U(C, L_{total}) = C - \frac{L_{total}^{1+\phi}}{1+\phi},$$

where $L_{total}$ denotes aggregate labor supplied to the human-art and corporate-art sectors, and $\phi > 0$ is the inverse Frisch elasticity of labor supply. Let $P$ denote the ideal aggregate price index. Utility maximization implies the aggregate labor-supply schedule

$$L_{total}^S(w, P) = \left(\frac{w}{P}\right)^{1/\phi}. \tag{15}$$

This specification delivers a tractable upward-sloping labor supply schedule while shutting down wealth effects on labor supply.

**Labor Market and Residual Oligopsony Power.** Because individuals can allocate their time between independent artistic creation ($L_H$) and corporate orporate employment ($L_{Corp}$), the AI sector does not face the aggregate macroeconomic labor supply directly. Instead, it faces a **residual labor supply curve**.

From the household's utility maximization, the aggregate labor supply is driven by the real wage: $L_{total}(w, P) = (w/P)^{1/\phi}$. Concurrently, profit maximization in the independent art sector implies nominal labor demand $L_H(w, P_H) = \left(\frac{\alpha A_H P_H}{w}\right)^{\frac{1}{1-\alpha}}$. The labor supply available to the AI sector is the residual:

$$L_{Corp}^S(w, P_H, P) = L_{total}(w, P) - L_H(w, P_H) = \left(\frac{w}{P}\right)^{1/\phi} - \left(\frac{\alpha A_H P_H}{w}\right)^{\frac{1}{1-\alpha}} \tag{16}$$

Because firms are atomistic with respect to the macroeconomic price index $P$, they take $P$ as given when evaluating the local elasticity of the residual labor supply faced by the corporate sector.

Let

$$\epsilon_L \equiv \frac{\partial L_{Corp}^S}{\partial w} \frac{w}{L_{Corp}} > 0 \tag{17}$$

denote that residual labor-supply elasticity. Because the number of AI firms $N$ is finite, each firm internalizes the impact of its own hiring $L_i$ on this sector-specific wage. The firm's marginal factor cost of labor is

$$MFC_L = w(L_{Corp}, P_H, P)\left[1 + \frac{s_i}{\epsilon_L}\right], \tag{18}$$

where $s_i = L_i/L_{Corp}$ is the firm's employment share within the AI sector. Under symmetry, $s_i = 1/N$, so the perceived wage entering firm decisions can be written as

$$\widetilde{w}(N) \equiv w\left(1 + \frac{1}{N\epsilon_L}\right). \tag{19}$$

For the theory and quantitative analysis that follow, the labor-market wedge is summarized by this reduced-form local elasticity $\epsilon_L$.

## 2.4 Firm Problem and Timing

The model is solved by backward induction. Entrants choose data in stage 1, which determines automation intensity and appropriability, and firms then compete in stage 2 given their technology and the mass of active firms.

**Stage 2: Production and Operating Profit.** Fix an active firm $i$ with data choice $D_i$ in a market with $N$ firms. Taking aggregate expenditure $E$, the human-art block $\{P_H, Y_H\}$, and the aggregate price index $P$ as given, the firm chooses quantity $q_i$ to maximize expected operating profit:

$$\max_{q_i} \quad \mathbb{E}\big[\pi_{op,i}\big] = \bar{\beta}(I(D_i), \bar{\xi})\left[p_i\big(q_i, \mathbf{q}_{-i}; P_H, Y_H, E\big)\, q_i - AVC(D_i, N)\, q_i\right]. \tag{20}$$

Because appropriability enters as a strictly multiplicative factor on operating profit, it does not affect the intratemporal pricing condition. The firm's output and pricing decisions are therefore governed by the perceived marginal cost:

$$MC_{perceived}(D_i, N) = \frac{1}{B}\left(\frac{r}{I(D_i)}\right)^{I(D_i)}\left(\frac{\widetilde{w}(N)}{1 - I(D_i)}\right)^{1-I(D_i)}, \tag{21}$$

where the perceived wage $\widetilde{w}(N)$ internalizes the firm's wage impact through the residual labor supply. The corresponding actual average variable cost is

$$AVC(D_i, N) = \left[\frac{w}{\widetilde{w}(N)}(1 - I(D_i)) + I(D_i)\right] MC_{perceived}(D_i, N). \tag{22}$$

In a symmetric equilibrium, the optimal pricing rule is

$$p^* = \frac{\varepsilon(N)}{\varepsilon(N) - 1} MC_{perceived}(D^*, N), \tag{23}$$

where

$$\frac{1}{\varepsilon(N)} = \frac{1}{\eta} + \frac{1}{N}\left(\frac{1}{\sigma} - \frac{1}{\eta}\right). \tag{24}$$

Let

$$\mathcal{M}(N) \equiv \frac{\varepsilon(N)}{\varepsilon(N) - 1} \tag{25}$$

denote the endogenous product-market markup.

**Stage 1: Data Choice.** Anticipating the stage-2 operating profit implied by its data choice, an entrant chooses $D_i$ to maximize expected net profit:

$$\max_{D_i} \quad \mathbb{E}\big[\pi^*_{op,i}(D_i)\big] - \lambda P_H D_i - F_{base} \qquad \text{s.t.} \qquad D_i \le Y_H, \tag{26}$$

where $\mathbb{E}[\pi^*_{op,i}(D_i)]$ denotes expected equilibrium operating profit after substituting in the optimal stage-2 quantity and price choices induced by $D_i$.

This data problem captures two opposing margins. A higher $D_i$ raises automation intensity and alters the firm's cost structure through $I(D_i)$, but it also lowers the appropriable share of output through $\bar{\beta}(I(D_i), \bar{\xi})$ and increases the royalty bill $\lambda P_H D_i$. Let

$$MB_D(D_i) \equiv \frac{\partial \mathbb{E}[\pi^*_{op,i}(D_i)]}{\partial D_i} \tag{27}$$

denote the marginal benefit of data net of its effect on equilibrium operating profit. The optimal choice satisfies the Kuhn–Tucker conditions:

$$D^* = \begin{cases} \text{the interior solution to } MB_D(D) = \lambda P_H, & \text{if } MB_D(Y_H) < \lambda P_H, \\ Y_H, & \text{if } MB_D(Y_H) \ge \lambda P_H. \end{cases} \tag{28}$$

When the constraint binds and $D^* = Y_H$, all active firms train on the same aggregate stock of human art. We nevertheless maintain imperfect substitutability across corporate varieties, since differentiated architectures, proprietary fine-tuning, reinforcement procedures, and heterogeneous generative seeds allow identical training data to support distinct creative outputs.

**Entry and Market Structure.** Entry is free at the extensive margin. Potential firms compare the expected net profit generated by the optimal data choice with the fixed entry cost $F_{base}$. In a symmetric free-entry equilibrium, the mass of active firms $N^*$ adjusts until expected net profit is driven to zero:

$$\mathbb{E}\big[\pi^*_{op}(D^*)\big] - \lambda P_H D^* - F_{base} = 0. \tag{29}$$

This condition is the key link from policy to market structure. The royalty rate $\lambda$ and judicial strictness $\bar{\xi}$ first affect firm profitability through data cost and appropriability; free entry then maps those profitability changes into the equilibrium firm mass $N^*$, which in turn governs both the product-market markup and the labor-market wedge.

# 3 Equilibrium Characterization

A symmetric free-entry equilibrium is defined by the following conditions.

1. **Stage-2 firm optimization.** Given $D^*$ and $N^*$, firms price according to the stage-2 optimality conditions in (21) and (23), so that

$$p^* = \mathcal{M}(N^*)\, MC_{perceived}(D^*, N^*).$$

2. **Stage-1 data choice.** Given $P^*_H$ and $N^*$, the firm's optimal data choice satisfies the Kuhn–Tucker condition (28).

3. **Goods-market clearing.** Supply equals demand in both art segments:

$$Y_H^* = \omega \left( \frac{P_H^*}{P^*} \right)^{-\sigma} \frac{E^*}{P^*}, \tag{30}$$

$$(N^*)^{\frac{\eta}{\eta-1}} q^* = (1-\omega) \left( \frac{P_{Corp}^*}{P^*} \right)^{-\sigma} \frac{E^*}{P^*}. \tag{31}$$

4. **Labor-market clearing.** Aggregate labor supply equals labor used in independent artistic production and corporate creative production:

$$\left( \frac{w^*}{P^*} \right)^{1/\phi} = \left( \frac{Y_H^*}{A_H} \right)^{1/\alpha} + N^* \, L_i(q^*, D^*, \widetilde{w}(N^*), r).$$

5. **Data-market feasibility.** Firm-level data use respects aggregate data availability:

$$D^* \leq Y_H^*.$$

6. **Free entry.** The mass of active firms adjusts until expected net profit is zero:

$$\mathbb{E}\big[\pi_{op}^*(D^*)\big] - \lambda P_H^* D^* - F_{base} = 0.$$

In equilibrium, all active firms choose the same data intensity, set the same price, and produce the same quantity. The equilibrium firm mass $N^*$ is endogenous and serves as the central transmission object in the model, jointly governing the product-market markup and the labor-market wedge.

# 4 Theory: Policy, Regimes, and Welfare Incidence in General Equilibrium

This section studies copyright policy in a two-dimensional space spanned by the data royalty rate $\lambda$ and judicial strictness $\bar{\xi}$. The two instruments have distinct institutional content. The royalty rate $\lambda$ governs the transfer associated with lawful data use and thus operates as a redistribution instrument. Judicial strictness $\bar{\xi}$ instead governs appropriability by shaping the extent to which AI-assisted output remains protectable against imitation. The policy problem is therefore not one-dimensional: it concerns both the contractual terms of data access and the legal scope of downstream exclusion.

A partial-equilibrium reading would treat these instruments as simple deterrence devices or cost shifters. That interpretation is incomplete here. Firms choose data use endogenously, data is non-rival, entry is free, labor-market outcomes are endogenous, and relative prices adjust jointly with market structure. Policy therefore operates through general-equilibrium incidence rather than through a single mechanically signed channel. In particular, it acts through at least two conceptually distinct margins: an intensive margin governing the profitability of data use conditional on operation, and an extensive margin governing free entry and the equilibrium mass of firms.

The remainder of the theory section develops this logic in sequence. It first characterizes the relevant equilibrium regimes, distinguishing interior adoption, binding-data configurations, and fallback to the manual regime. It then shows that market structure is the central transmission mechanism of the model: lower equilibrium firm mass mechanically strengthens both product-market markup power and the labor-market

wedge. The analysis next explains why the full policy incidence on equilibrium objects is generally not signable in closed form. It then turns to artists, organizing incidence through a nominal income decomposition into independent sales income, data royalty income, and corporate-sector wage income.

This structure also clarifies the division of labor between theory and quantification. The theory identifies the admissible equilibrium regimes, the structural role of market structure, the limits of closed-form policy signability, and the nominal income channels through which policy reaches artists. It does not by itself deliver a complete signed ranking of policy outcomes over the two-dimensional policy space. The quantitative section therefore evaluates the resulting real-welfare and macroeconomic incidence.

To rule out a degenerate environment in which the corporate sector is infeasible absent AI-assisted production, Chapter 4 maintains a manual-regime viability condition. Assumption 1 ensures that a symmetric free-entry equilibrium exists under purely human-operated production.

**Assumption 1** (Viability of the Manual Economy)**.** *Let $w^*_{man}$ and $P^*_{Corp,man}$ denote the general equilibrium wage and corporate price index in an economy strictly bounded by $D_i = 0$ for all active firms. The intrinsic labor productivity $A_L$ and baseline fixed cost $F_{base}$ satisfy*

$$\max_{q_i} \left[ p(q_i, \mathbf{q}_{-i}) \cdot q_i - \frac{w^*_{man}}{A_L} q_i \right] \geq F_{base},$$

*such that there exists a symmetric free-entry equilibrium with strictly positive firm mass $\bar{N}_{man} > 0$.*

**Assumption 2** (Technological Superiority of AI)**.** *In the least restrictive copyright environment, $(\lambda, \bar{\xi}) = (0, 0)$, the cost-reducing benefit of AI automation is sufficiently large that the maximum expected operating profit from positive AI data use strictly dominates the baseline manual regime:*

$$\max_{D>0} \left\{ \bar{\beta}(I(D), 0) \left[ p^*(D) - AVC(D) \right] q^*(D) \right\} > \left[ p^*_{man} - \frac{w^*_{man}}{A_L} \right] q^*_{man}. \tag{32}$$

This assumption rules out parameterizations in which AI is never privately attractive even at the least restrictive policy corner. It therefore excludes trivial cases in which the efficiency gain from automation is too weak to sustain positive data use under any admissible equilibrium environment. Actual adoption or non-adoption nevertheless remains jointly determined by policy and general-equilibrium conditions rather than being imposed a priori.

## 4.1 Regime Characterization: Manual Production, Interior AI Adoption, and Data Starvation

This subsection characterizes the admissible equilibrium regimes of the model before turning to market-structure and welfare transmission. Because training data are non-rival, the relevant scarcity constraint is not aggregate congestion across firms but the firm-level feasibility condition $D_i \leq Y_H$. In a symmetric equilibrium, each active firm faces the same firm-level feasibility bound implied by equilibrium human-art supply, so the selected equilibrium data choice must lie in the interval $[0, Y^*_H]$. The central object of this subsection is therefore not a comparative-statics threshold result, but a regime classification of the selected symmetric free-entry equilibrium.

**Proposition 1** (Equilibrium regime classification)**.** *Consider a selected symmetric free-entry equilibrium with $N^* > 0$. Then the equilibrium data choice $D^*$ must belong to exactly one of the following three regimes:*

*(i) Manual regime: $D^* = 0$. Under Assumption (1), the selected equilibrium coincides with the symmetric manual free-entry equilibrium, so $N^*(\lambda; \bar{\xi}) = \bar{N}_{man}$.*

*(ii) Interior AI-adoption regime with slack data constraint:* $0 < D^* < Y_H^*$. *Then* $\mu_D = 0$, $I(D^*) \in (0,1)$, *and* $\bar{\beta}(I(D^*), \bar{\xi}) \in (0,1]$.

*(iii) Data-starvation regime with binding data constraint:* $D^* = Y_H^*$.

Proposition 1 provides a regime partition rather than a comparative-statics result. Conditional on a selected symmetric free-entry equilibrium with $N^* > 0$, the model admits exactly three economically distinct possibilities: no AI data use, positive but unconstrained AI data use, and full use of the firm-level feasible stock of human-created training data. This classification is the relevant backbone for the later analysis because market structure and artist-income incidence propagate through different channels across these three cases.

In the manual regime, the equilibrium collapses to the verified manual system. Since $D^* = 0$ implies $I(0) = 0$ and $\bar{\beta}(0, \bar{\xi}) = 1$, the manual reductions apply: $B = A_L$, $MC_{perceived}(0, N) = \widetilde{w}(N)/A_L$, and $AVC(0, N) = w/A_L$. Accordingly, once the selected equilibrium is manual, subsequent equilibrium objects are governed by the same free-entry logic as in the exogenously manual economy.

In the interior AI-adoption regime, firms use human-created data but remain strictly inside the feasible set. The strict inequality $0 < D^* < Y_H^*$ means that the physical data constraint is slack, so the associated multiplier is zero. This is the genuinely mixed regime of the model: AI adoption is operative, but its equilibrium intensity is not pinned down by the hard feasibility boundary. Later results that analyze equilibrium transmission under positive AI use should therefore be interpreted as applying within this interior region unless the physical constraint binds.

In the data-starvation regime, the selected equilibrium lies at the firm-level physical boundary $D^* = Y_H^*$. Because data are non-rival in the model, this boundary should be interpreted at the level of each firm rather than as congestion across firms. In symmetric equilibrium, the binding condition means that each active firm uses the full contemporaneous supply of human-created art available for lawful training. This is the regime in which the physical availability of human art directly truncates further expansion in firm-level data use.

**Remark 1** (Physical truncation of interior AI dynamics)**.** *The classification above also clarifies the role of the feasibility condition* $D_i \leq Y_H$. *Any candidate interior AI-adoption logic is valid only so long as the selected equilibrium remains in the interval* $0 < D^* < Y_H^*$. *If the equilibrium forces that would otherwise support a positive interior data choice instead point beyond the available human-art supply, the selected equilibrium cannot remain interior and is implemented at the corner* $D^* = Y_H^*$. *Thus the model allows candidate interior AI dynamics to be truncated by the physical data boundary, without requiring a stronger threshold or monotonicity claim.*

Remark 1 is conceptually important for the sections that follow. It clarifies that interior AI-adoption relationships should not be read as globally valid across the parameter space. Their relevance is regime-dependent: they describe the selected equilibrium only when the physical data boundary is slack. Once the boundary binds, the correct equilibrium object is the corner $D^* = Y_H^*$, and the subsequent market-structure and welfare analysis must be conducted under that binding feasibility condition.The next subsection abstracts from policy signability and instead studies how any equilibrium change in firm mass transmits mechanically into markup power and labor-market wedges.

## 4.2 Market Structure: the Core Transmission Mechanism

Market structure is the model's central transmission mechanism because, in symmetric equilibrium, the same endogenous firm mass $N^*$ disciplines both pricing power in the product market and wage-setting power in

the labor market. This observation is purely structural. It states that whenever the selected equilibrium features a lower firm mass, the sector is mechanically more concentrated on both margins.

**Proposition 2** (Double Squeeze of Consolidation)**.** *Consider any selected symmetric equilibrium with* $N^* > 0$*. Under the maintained parameter ordering* $\eta > \sigma > 1$*,*

$$\varepsilon(N)^{-1} = \eta^{-1} + N^{-1}\big(\sigma^{-1} - \eta^{-1}\big) \tag{33}$$

*implies that* $\varepsilon'(N) > 0$*, and hence*

$$\mathcal{M}'(N) < 0 \qquad for \qquad \mathcal{M}(N) \equiv \frac{\varepsilon(N)}{\varepsilon(N) - 1}. \tag{34}$$

*Moreover, for*

$$\Omega(N) \equiv \frac{\widetilde{w}(N)}{w^*} = 1 + \frac{1}{N\epsilon_L}, \tag{35}$$

*one has* $\Omega'(N) < 0$ *holding* $\epsilon_L$ *fixed. Therefore, if two selected symmetric equilibria satisfy* $0 < N_a < N_b$*, then*

$$\mathcal{M}(N_a) > \mathcal{M}(N_b) \qquad and \qquad \Omega(N_a) > \Omega(N_b). \tag{36}$$

This proposition establishes the paper's core structural mechanism. A lower equilibrium firm mass strengthens market power on two margins simultaneously. On the product-market side, fewer active firms reduce the perceived elasticity of demand and raise the endogenous markup. On the labor-market side, the same reduction in firm mass enlarges each firm's share of sectoral labor demand and deepens the wedge between marginal factor cost and the actual wage. The resulting concentration effect is therefore a double squeeze: lower $N^*$ simultaneously increases price-cost margins and labor-market distortions, with immediate implications for factor shares and later incidence analysis.

The same logic admits a convenient accounting decomposition for the Corporate-Sector labor share. Combining the firm's conditional labor demand, the pricing rule, and symmetry yields a representation in which technology and market structure enter multiplicatively.

**Proposition 3** (Corporate-Sector Labor-Share Decomposition)**.** *Consider any selected symmetric equilibrium with* $N^* > 0$*. Let physical labor demand satisfy*

$$l_i = \frac{(1 - I(D^*))\, MC_{perceived}(D^*, N^*)\, q^*}{\widetilde{w}(N^*)}. \tag{37}$$

*Then the Corporate-Sector labor share*

$$LS_{AI} \equiv \frac{w^* L^*_{Corp}}{P^*_{Corp} Q^*_{Corp}} \tag{38}$$

*can be written as*

$$LS_{AI} = \frac{1 - I(D^*)}{\Omega(N^*)\, \mathcal{M}(N^*)}. \tag{39}$$

Proposition 3 shows that the Corporate-Sector labor share is shaped by two distinct objects. The term $1 - I(D^*)$ captures the technological labor requirement embedded in the chosen production method, while the denominator $\Omega(N^*)\mathcal{M}(N^*)$ captures the equilibrium distortion generated by dual market power. The decomposition is therefore useful precisely because it isolates the mechanism: conditional on equilibrium adoption, a more concentrated market structure lowers labor's share mechanically by increasing both the monopsony wedge and the markup. This prepares the next subsection, where the interaction between adoption, entry, and general-equilibrium feedbacks makes closed-form policy derivatives generally non-signable.

When the selected equilibrium is manual, the verified manual-regime reductions discussed in Section 4.1 apply, and the selected equilibrium satisfies $N^*(\lambda; \bar{\xi}) = \bar{N}_{man}$. Section **??** therefore does not require a policy theorem here: it only requires the structural fact that, whenever the equilibrium mass of firms is lower, the corporate sector exhibits stronger product-market and labor-market power, with immediate implications for factor shares.

## 4.3 Within-Structure Policy Incidence and Endogenous Market-Structure Transmission

Section 4.2 showed that equilibrium firm mass is a central transmission object because it disciplines both the markup and the labor-market wedge. That observation changes the correct interpretation of policy incidence. Conditional on a given market structure, each policy instrument has an immediate margin that is favorable to human creators. A higher royalty rate $\lambda$ strengthens the transfer associated with lawful data use. A higher judicial strictness $\bar{\xi}$ tightens the appropriability constraint on AI-assisted output through $\bar{\beta}(I(D), \bar{\xi})$ and thereby strengthens the protective margin for human-created work. These local readings are economically valid. What they do not determine is the full equilibrium incidence once free entry allows policy to alter the mass of firms and hence the market structure through which all other objects are selected.

To isolate this distinction, fix an arbitrary reference firm mass $N$ and let $\theta \in \{\lambda, \bar{\xi}\}$ denote the policy instrument under consideration. Let $X$ collect the remaining endogenous equilibrium objects other than firm mass and firm-level data use, including at least $P_H$, $w$, $P$, $p$, $q$, and $\epsilon_L$. Conditional on $N$, write the equilibrium system compactly as

$$\Gamma(D, X; N, \theta) = 0, \tag{40}$$

where $\Gamma$ collects the interior data-choice condition together with the remaining equilibrium restrictions conditional on the given market structure. The associated conditional equilibrium is denoted by

$$(D^c(N, \theta), X^c(N, \theta)). \tag{41}$$

This is a within-structure benchmark rather than a narrow partial-equilibrium exercise. Prices, wages, quantities, and data use may all adjust inside $(D^c, X^c)$ as policy changes. The restriction is only that firm mass itself is held fixed.

Under the regularity conditions required for local implicit differentiation, the conditional comparative statics satisfy

$$\Gamma_D \frac{\partial D^c}{\partial \theta} + \Gamma_X \cdot \frac{\partial X^c}{\partial \theta} + \Gamma_\theta = 0. \tag{42}$$

Equation (42) is the formal within-structure incidence object. The derivative is taken holding $N$ fixed, so no term involving $\frac{dN}{d\theta}$ appears. This is the precise sense in which the benchmark shuts down market-structure adjustment while preserving the rest of the equilibrium system. The same logic applies to any endogenous equilibrium object selected by the conditional system. The data-choice variable $D$ is used here as a representative case because it is already central to the regime classification developed above and because it provides a natural bridge to the later quantitative analysis of AI adoption. The purpose of the example is therefore analytical rather than exclusive. It illustrates, in a concrete equilibrium object, the general distinction between within-structure incidence and endogenous market-structure transmission.

The data choice $D$ provides a useful worked example because it is already central to the regime classification and will also matter in the quantitative analysis of AI adoption. In the interior AI-adoption regime,

the conditional data choice is characterized by

$$MB_D\big(D^c(N,\theta);N,X^c(N,\theta),\theta\big)=\lambda P_H^c(N,\theta), \tag{43}$$

with the understanding that the corner condition $D^c=Y_H^c$ applies when the physical constraint binds. In the interior case, differentiating (43) with respect to $\theta$ while holding $N$ fixed yields

$$\frac{\partial D^c}{\partial\theta}=-\frac{\partial_\theta[MB_D-\lambda P_H^c]+\partial_X[MB_D-\lambda P_H^c]\cdot\frac{\partial X^c}{\partial\theta}}{\partial_D[MB_D-\lambda P_H^c]}. \tag{44}$$

This expression captures the direct policy effect on firm-level data incentives under a fixed market structure. For $\theta=\lambda$, the royalty schedule enters directly through the right-hand side of (43). For $\theta=\bar{\xi}$, appropriability enters through the operating-profit component embedded in $MB_D$. In both cases, the derivative is conditional on the given degree of concentration.

Return now to the selected symmetric free-entry equilibrium. Let the free-entry residual be

$$\Psi(N,D,X;\theta)\equiv\bar{\beta}(I(D),\bar{\xi})\big[pq-AVC(D,N)q\big]-\lambda P_H D-F_{base}, \tag{45}$$

so that the full equilibrium is characterized by

$$\Gamma(D^*,X^*;N^*,\theta)=0,\qquad \Psi(N^*,D^*,X^*;\theta)=0. \tag{46}$$

Along any differentiable equilibrium path, total differentiation gives

$$\Gamma_D\,\frac{dD^*}{d\theta}+\Gamma_X\cdot\frac{dX^*}{d\theta}+\Gamma_N\,\frac{dN^*}{d\theta}+\Gamma_\theta=0 \tag{47}$$

and

$$\Psi_D\,\frac{dD^*}{d\theta}+\Psi_X\cdot\frac{dX^*}{d\theta}+\Psi_N\,\frac{dN^*}{d\theta}+\Psi_\theta=0. \tag{48}$$

Relative to (42), the substantive difference is the appearance of the entry term involving $\frac{dN^*}{d\theta}$. Once free entry is restored, policy no longer changes data use and the remaining equilibrium objects only through the conditional system. It also changes the equilibrium market structure, and that change feeds back into the same objects through $\Gamma_N$ and $\Psi_N$.

Equivalently, because the full-equilibrium data choice can be written as

$$D^*(\theta)=D^c(N^*(\theta),\theta), \tag{49}$$

the chain rule yields

$$\frac{dD^*}{d\theta}=\left.\frac{\partial D^c(N,\theta)}{\partial\theta}\right|_{N=N^*(\theta)}+\left.\frac{\partial D^c(N,\theta)}{\partial N}\right|_{N=N^*(\theta)}\frac{dN^*}{d\theta}. \tag{50}$$

Equation (50) is the economically relevant decomposition. The first term is within-structure policy incidence. It measures how the policy instrument changes firm-level data use when the degree of market concentration is held fixed. The second term is endogenous market-structure transmission. It measures how the same policy changes data use indirectly by changing equilibrium entry and hence the market structure in which data is chosen.

The same decomposition extends immediately to any equilibrium outcome that can be written as a

conditional object $Z^c(N, \theta)$ under fixed firm mass. For any such variable,

$$Z^*(\theta) = Z^c(N^*(\theta), \theta) \tag{51}$$

and hence

$$\frac{dZ^*}{d\theta} = \left.\frac{\partial Z^c(N,\theta)}{\partial \theta}\right|_{N=N^*(\theta)} + \left.\frac{\partial Z^c(N,\theta)}{\partial N}\right|_{N=N^*(\theta)} \frac{dN^*}{d\theta}. \tag{52}$$

The case of $D$ therefore serves only to make explicit, in a substantive equilibrium object, a comparative-statics structure that applies more generally throughout the model.

This decomposition also makes clear why direct policy intuition is insufficient for full equilibrium analysis. A higher $\lambda$ may, under fixed $N$, depress data demand through the royalty-payment margin. A higher $\bar{\xi}$ may, under fixed $N$, reduce the private return to AI-assisted output through the appropriability margin. But the total effect on equilibrium data use is not exhausted by either local deterrence channel. If the same policy change alters $N^*$, it also changes markup pressure, the labor-market wedge, perceived marginal cost, and the equilibrium environment summarized by $X^*$. The induced response of market structure therefore becomes part of the incidence of policy itself rather than a secondary correction.

The point is disciplined rather than expansive. It is not necessary to claim that every policy derivative is globally ambiguous in every regime. The relevant implication is narrower. Once entry is endogenous, even an outcome with a clear within-structure response need not inherit that response in the full equilibrium, because the market-structure term in (50) may reinforce, attenuate, or offset the direct effect. The local policy margin remains informative, but it is no longer sufficient for policy evaluation.

This is the sense in which market structure is a first-order transmission mechanism for later incidence analysis. Artists do not care about $N^*$ or $D^*$ in isolation. They are affected through the income objects that these equilibrium variables help determine. Since policy operates both through within-structure changes in data use and through endogenous changes in firm mass, the next subsection turns to the artist side and decomposes nominal income into the channels through which those equilibrium responses are ultimately transmitted.

## 4.4 Artist Income Decomposition, Equilibrium Offsets, and the Transition to Real Welfare

The appropriate starting point for artist-side incidence is the exact nominal income identity implied by the selected symmetric free-entry equilibrium. Let total nominal artist income be

$$W_H^{nom} = W_H^{sales} + W_H^{data} + W_H^{wage}, \tag{53}$$

with component expressions

$$W_H^{sales} = P_H^* Y_H^*, \qquad W_H^{data} = \lambda P_H^* D^* N^*, \qquad W_H^{wage} = w^* L_{Corp}^*. \tag{54}$$

This decomposition is exact and model-based. It follows directly from the equilibrium sources through which artists are compensated in the present environment: independent sales of human art, royalty receipts from lawful data use, and wage earnings from labor supplied to the corporate AI sector. For incidence analysis, this identity is the relevant accounting backbone because it keeps the artist side tied to equilibrium objects already defined by the model, rather than to an ad hoc verbal partition of channels. At this stage, the purpose of the decomposition is analytical rather than normative. It identifies how policy reaches artists in

nominal terms before the transition to real welfare is made.

Let $\theta \in \{\lambda, \bar{\xi}\}$ denote a generic policy instrument. For each nominal income component $W_k$, with $k \in \{\text{sales}, \text{data}, \text{wage}\}$, the total derivative takes the standard form

$$\frac{dW_k}{d\theta} = \frac{\partial W_k}{\partial \theta} + \sum_{s \in \mathcal{S}_k} \frac{\partial W_k}{\partial s} \frac{ds}{d\theta}, \tag{55}$$

where $\mathcal{S}_k$ collects the equilibrium objects on which the component depends. In the present model, the relevant objects include $P_H^*$, $Y_H^*$, $D^*$, $N^*$, $w^*$, and $L_{Corp}^*$, together with other equilibrium quantities that enter these objects indirectly.

For independent sales income, policy typically enters through equilibrium adjustment rather than through an explicit policy term:

$$\frac{dW_H^{sales}}{d\theta} = Y_H^* \frac{dP_H^*}{d\theta} + P_H^* \frac{dY_H^*}{d\theta}. \tag{56}$$

Thus the sales component loads on policy through induced movements in the equilibrium price and quantity of human art. For royalty income, the derivative separates the direct mechanical transfer effect from the induced equilibrium responses:

$$\frac{dW_H^{data}}{d\theta} = \mathbf{1}\{\theta = \lambda\} P_H^* D^* N^* + \lambda D^* N^* \frac{dP_H^*}{d\theta} + \lambda P_H^* N^* \frac{dD^*}{d\theta} + \lambda P_H^* D^* \frac{dN^*}{d\theta}. \tag{57}$$

Only the royalty rate generates the direct transfer term. All remaining terms depend on the equilibrium response of the human-art price base, firm-level data use, and firm mass. For wage income,

$$\frac{dW_H^{wage}}{d\theta} = L_{Corp}^* \frac{dw^*}{d\theta} + w^* \frac{dL_{Corp}^*}{d\theta}. \tag{58}$$

Here policy incidence is mediated through the equilibrium wage and the amount of artist labor allocated to corporate employment. The point of these expressions is not to enumerate every derivative term exhaustively. It is to make clear that incidence on the artist side is built from total derivatives of model-consistent income components.

The component-wise derivatives are only the first layer of the analysis. Their economic interpretation does not come from reading sales, royalty, and wage income in isolation. It comes from regrouping derivative terms across the three components according to the common equilibrium forces that generate them. In particular, the nominal incidence of policy is naturally organized around two offset categories.

The first is a market-structure offset. This category collects the terms that operate through the endogenous response of firm mass $N^*$. In the royalty component, this appears directly through the term proportional to $\frac{dN^*}{d\theta}$. More broadly, the previous subsection established that changes in $N^*$ propagate mechanically into both product-market markup power and the labor-market wedge. Those induced movements then feed back into $P_H^*$, $w^*$, and the associated income components. The artist-side implication is that policy does not reach nominal income only through the immediate object it appears to target. It also reaches artists through the market structure selected in equilibrium. This is the downstream incidence counterpart of the earlier transmission result: once policy alters equilibrium entry, part of the artist-income response is governed by the concentration of the corporate sector rather than by the direct policy margin alone.

The second is a reallocation offset. This category collects the terms that operate through endogenous changes in data use, labor allocation, human-art output, demand composition, and closely related equilibrium quantities. In the derivative system above, these forces appear through $\frac{dD^*}{d\theta}$, $\frac{dY_H^*}{d\theta}$, $\frac{dL_{Corp}^*}{d\theta}$, and through the induced movement of prices that reflect these reallocations. Economically, these terms describe how policy

reorganizes the sources of artists' nominal income. A policy that raises one margin immediately need not raise total nominal artist income proportionately, because equilibrium adjustment can shift activity across independent production, royalty-generating data use, and corporate employment. The relevant incidence object is therefore the regrouped total derivative of nominal artist income,

$$\frac{dW_H^{nom}}{d\theta} = \frac{dW_H^{sales}}{d\theta} + \frac{dW_H^{data}}{d\theta} + \frac{dW_H^{wage}}{d\theta}, \tag{59}$$

read through the common equilibrium forces embedded in these terms. The market-structure and reallocation offsets do not replace the underlying accounting identity. They are obtained by reorganizing its derivative terms according to economic content. This is the level at which it becomes transparent that an apparently favorable nominal margin can be attenuated, amplified, or reversed once the equilibrium responses across components are taken together.

Nominal incidence, however, is not yet the welfare ranking. The final step is to map nominal artist receipts into real purchasing power. Let real artist welfare be represented by nominal artist income deflated by the aggregate price index,

$$W_H^{real} \equiv \frac{W_H^{nom}}{P^*}. \tag{60}$$

Then

$$\frac{dW_H^{real}}{d\theta} = \frac{1}{P^*}\frac{dW_H^{nom}}{d\theta} - \frac{W_H^{nom}}{(P^*)^2}\frac{dP^*}{d\theta}. \tag{61}$$

The second term is the real-incidence offset. It is conceptually distinct from the nominal-side regrouping above. Market-structure and reallocation offsets describe how policy reshapes nominal artist income across equilibrium channels. The real-incidence offset enters only at the step where nominal receipts are translated into real welfare through the aggregate price level. For that reason, the nominal decomposition remains essential for channel interpretation even though the final welfare criterion must be expressed in real terms. The quantitative section takes up this final step by determining whether, under the maintained parameterization, the equilibrium offsets are strong enough to dominate the apparent direct nominal gains associated with a given policy change.

# 5 Quantitative Analysis

## 5.1 Calibration setup, targets, and equilibrium objects

This section quantifies the general-equilibrium incidence of AI copyright policy in the two-dimensional policy space $(\lambda, \bar{\xi})$. Consistent with the revised theory, the quantitative exercise is not designed to recover globally signed comparative statics from functional form alone. Its purpose is instead to map the equilibrium regimes admitted by the model, quantify the market-structure and income-channel adjustments induced by policy, and evaluate final policy incidence in real terms. The organizing logic follows directly from the theory: regime classification comes first, market structure is the central transmission mechanism, nominal artist-income decomposition is an intermediate accounting device, and real welfare is the final evaluation object.

The calibration is anchored to the U.S. motion picture and entertainment industry. This setting is empirically useful because it is highly digitized, IP-intensive, and characterized by concentrated labor demand for specialized creative inputs. Within this interpretation, the corporate side of the model corresponds to commercially organized production by studios and related firms, while the human-art sector corresponds to independent creative output. The wage $w$ is therefore interpreted as the sector-specific compensation relevant for creative labor allocation within this industry rather than as an economy-wide macro wage. This

industry mapping is also the relevant domain for the model's labor-market wedge, since oligopsony power operates through concentration in the corporate sector's demand for creative labor.

The pre-AI baseline is defined as the stationary equilibrium corresponding to 2015–2019. In the model, this period is mapped to the limit in which relative computing costs are prohibitively high, $\rho \to \infty$, so that AI-assisted production is not commercially viable and the equilibrium satisfies $D^* = 0$ and $I(D^*) = 0$. The baseline economy is thus the manual-regime equilibrium of the model, evaluated before the post-2022 decline in computing costs associated with the generative AI shock. This mapping disciplines parameters using pre-AI market structure and expenditure patterns rather than contemporary outcomes already shaped by endogenous adoption, and it aligns the quantitative section with the theory, in which later policy comparisons are understood as movements away from a manual baseline into interior-adoption or data-starvation regions of the policy space.

A subset of parameters is assigned externally. The cross-sector elasticity of substitution is set to $\sigma = 2.5$, following empirical estimates for substitution across entertainment goods in de Roos and McKenzie (2014). Production-side productivity shifters are normalized to $A_H = A_K = A_L = 1$, so that the quantitative variation is carried by equilibrium responses rather than arbitrary scale terms. The decreasing-returns parameter in human-art production is set to $\alpha = 0.428$. In the model, $\alpha$ enters the independent-art technology $Y_H = A_H L_H^{\alpha}$, under which the labor share satisfies $wL_H/(P_H Y_H) = \alpha$. I therefore discipline $\alpha$ using the ratio of employee compensation to value added in the motion picture sector over 2015–2019 from the BEA Arts and Cultural Production Satellite Account.[2] The labor-supply elasticity facing the corporate sector is set to $\epsilon_L = 2.0$, in line with benchmark estimates from the monopsony literature (Manning, 2003; Sokolova and Sorensen, 2021). Finally, the AI learning-curve parameters are set to $\mu = 0.28$ and $\delta = 3.0$, which jointly govern how licensed data translate into the automation threshold $I(D)$. These parameters are assigned rather than estimated inside the calibration because the pre-AI moments used below do not separately identify these margins.

The remaining parameters, $(\eta, F_{base}, \omega)$, are calibrated internally by matching three pre-AI moments. The within-corporate elasticity of substitution $\eta$ is chosen to match the long-run operating margin of major Hollywood studios, targeted at 11.40% using Vogel (2020). The fixed entry cost $F_{base}$ is chosen to match market concentration, measured by the numbers-equivalent firm count $N = 6.80$ computed from pre-AI distributor market shares over 2015–2019.[3] The preference weight on human art, $\omega$, is chosen to match the independent-sector revenue share of 17.00%. Solving these moments jointly yields $\eta = 15.46$, $F_{base} = 0.0243$, and $\omega = 0.085$. Table 1 summarizes the assigned parameters, calibrated parameters, and target moments for the baseline economy.

The equilibrium objects carried forward to the subsequent subsections are those identified by the theory as central for quantitative incidence. For each policy pair $(\lambda, \bar{\xi})$, the model determines a selected symmetric free-entry equilibrium and therefore a regime classification among manual fallback, interior AI adoption, and data starvation. Because data are non-rival, the relevant physical constraint is firm-level rather than aggregate: $D_i \leq Y_H$, and in a symmetric equilibrium the binding corner is $D^* = Y_H^*$ rather than $N^* D^* = Y_H^*$. Conditional on regime, the key endogenous objects are the equilibrium mass of corporate firms $N^*$, the markup or product-market power object implied by residual demand, the labor-market wedge generated by concentrated labor demand, firm-level data use $D^*$ and the associated automation intensity $I(D^*)$, the aggregate price level, real output, and the nominal artist-income components given by independent sales income, royalty income, and corporate-sector wage income. These nominal objects are informative about

[2]Data and methodological details are available from the U.S. Bureau of Economic Analysis, Arts and Cultural Production Satellite Account (ACPSA), at `https://www.bea.gov/data/special-topics/arts-and-culture`.

[3]The target is constructed from the author's calculation using distributor market-share data for the pre-AI benchmark period from *The Numbers*.

Table 1: Calibration of the baseline pre-AI economy

| Parameter | Description | Value | Source / Target |
|---|---|---|---|
| *Panel A. Externally assigned parameters* | | | |
| $\sigma$ | Cross-sector elasticity of substitution | 2.50 | de Roos and McKenzie (2014) |
| $A_H$ | Human-sector productivity normalization | 1 | Normalization |
| $A_K$ | Corporate capital/input productivity normalization | 1 | Normalization |
| $A_L$ | Corporate labor productivity normalization | 1 | Normalization |
| $\alpha$ | Decreasing-returns parameter in human-art production | 0.428 | BEA ACPSA, motion picture sector, 2015–2019 |
| $\epsilon_L$ | Residual labor-supply elasticity faced by corporate firms | 2.00 | Manning (2003); Sokolova and Sorensen (2021) |
| $\mu$ | AI learning-curve parameter | 0.28 | Assigned technology parameter |
| $\delta$ | Task-complexity / curvature parameter | 3.00 | Assigned technology parameter |
| *Panel B. Internally calibrated parameters* | | | |
| $\eta$ | Within-corporate elasticity of substitution | 15.46 | Operating margin: 11.40% from Vogel (2020) |
| $F_{base}$ | Fixed entry cost | 0.0243 | Numbers-equivalent firm count: 6.80 from pre-AI distributor market shares (2015–2019) |
| $\omega$ | Preference weight on human art | 0.085 | Independent-sector revenue share: 17.00% |

*Notes:* The table reports the parameterization of the baseline pre-AI economy. The baseline corresponds to the manual-regime equilibrium associated with the 2015–2019 benchmark period, modeled as the limit in which relative computing costs are prohibitively high, so that $D^* = 0$ and $I(D^*) = 0$. Panel A lists parameters assigned outside the calibration based on external evidence, benchmark values, or normalizations. Panel B lists parameters jointly calibrated to match pre-AI moments on market structure, profitability, and expenditure composition. In the counterfactual exercises that follow, these parameters are held fixed while policy varies over the two-dimensional space $(\lambda, \bar{\xi})$. Consistent with the model, data are non-rival: the relevant physical constraint is firm-level, $D_i \leq Y_H$, so that in a symmetric binding equilibrium the corner condition is $D^* = Y_H^*$ rather than $N^* D^* = Y_H^*$.

transmission channels, but the final policy criterion remains real artist welfare and, more broadly, real macro incidence.

The role of the quantitative section is now best understood through the decomposition developed in the theory. For any equilibrium object of interest, the total policy response can be interpreted as the combination of a *within-structure* adjustment, holding market structure fixed, and an *endogenous market-structure transmission* effect operating through the equilibrium change in firm mass $N^*$. The calibrated exercises that follow do not replace theory where closed-form signs are unavailable. Instead, they quantify the relative strength of these already-identified margins in the maintained baseline economy. The subsequent subsections therefore proceed in that order: first by locating the equilibrium regime, then by quantifying the market-structure response itself, then by tracing how that transmission mechanism reshapes nominal artist-income components, and finally by recombining those channels in real terms for welfare evaluation.

As an additional quantitative discipline, the calibration is also evaluated against a small set of untargeted moments. In the current draft, once the pre-AI parameters are fixed, a post-shock reduction in computing costs to $\rho = 0.05$ is imposed together with a baseline copyright regime $(\lambda, \bar{\xi}) = (0.1, 1.5)$. Under this post-shock environment, the model generates an equilibrium automation exposure of approximately $I^* \approx 0.195$, close to the task-exposure estimates reported by Eloundou et al. (2023). It also generates a royalty expenditure share $\lambda P_H^* D^* / R_{Corp}$ of about 1.20%. Although this moment is not targeted in the calibration, it is informative about the empirical plausibility of the calibrated economy under post-AI conditions. To benchmark its magnitude, recent industry disclosures by The Walt Disney Company (Form 10-K, 2024–2025) indicate that broad licensed-content costs account for roughly 5.11% to 5.75% of total revenue. Since AI data acquisition is naturally interpreted as a narrower subset of broader IP and content licensing expenditures, the model-implied share of 1.20% lies well within a plausible empirical range. The next subsection turns to the quantitative regime map over $(\lambda, \bar{\xi})$, which provides the backbone for the subsequent analysis of market structure, nominal income channels, real incidence, and policy evaluation.

## 5.2 Quantitative regime map in the two-dimensional policy space

The quantitative analysis begins by locating the calibrated economy in the two-dimensional policy space spanned by the data royalty rate $\lambda$ and judicial strictness $\bar{\xi}$. This step is logically prior to the interpretation of any subsequent equilibrium object.Because policy affects both the profitability of data use and the equilibrium conditions governing entry and operation in general equilibrium, later quantitative objects must be read conditional on regime. The regime map therefore provides the structural partition of the calibrated economy before turning to market structure, income channels, or welfare.

Figure 1 reports this partition for the benchmark post-shock environment with $\rho = 0.05$. Under the maintained calibration, the policy space is divided into three equilibrium regions. In the no-adoption region, the corporate sector falls back to manual production and the selected symmetric free-entry equilibrium features $D^* = 0$. In the interior-adoption region, firms adopt AI with positive data use but remain strictly inside the firm-level feasibility set, so $0 < D^* < Y_H^*$. In the data-starvation region, the firm-level data constraint binds in symmetric equilibrium, so $D^* = Y_H^*$. Consistent with the theory, this last region should be interpreted as a binding firm-level feasibility condition under non-rival data.

The two boundaries in the figure have distinct economic content. The adoption boundary separates policy pairs under which AI use is privately sustained in equilibrium from those under which the corporate sector reverts to manual production. Moving across that boundary therefore changes the operative production regime itself. The data-starvation frontier instead separates unconstrained AI adoption from configurations in which the equilibrium data choice is truncated by the available flow of human-created art at the firm

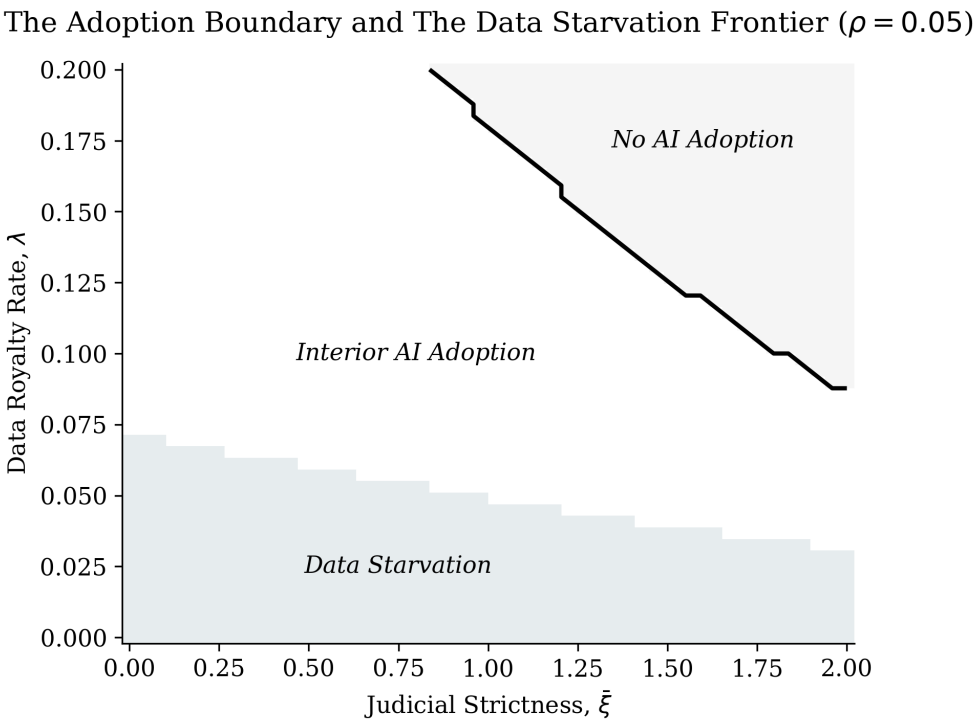



Figure 1: Quantitative regime map in the two-dimensional policy space ($\rho = 0.05$)

*Notes:* The figure reports the selected symmetric free-entry equilibrium in $(\lambda, \bar{\xi})$ space at $\rho = 0.05$, distinguishing manual fallback, interior AI adoption, and data starvation; in the binding-data region, the relevant constraint is the firm-level condition $D^* = Y_H^*$ under non-rival data.

level. The region between these two loci is the interior-adoption set, where firms use AI without hitting the binding data limit. In the calibrated environment shown here, the no-adoption region expands as the combined regulatory burden becomes sufficiently high, while the data-starvation region identifies policy combinations under which the firm-level constraint binds.

This regime map is descriptive and structural rather than evaluative. It identifies which equilibrium configuration is selected at each policy pair, and thereby fixes the domain within which later quantitative objects should be interpreted. The next subsection uses this partition to study how market structure varies within and across these regions, taking regime classification as the maintained background for the general-equilibrium transmission analysis.

## 5.3 Market structure and transmission in general equilibrium

Conditional on the regime map established in the previous subsection, the next object of interest is equilibrium firm mass $N^*$. The reason is structural. In the theoretical decomposition developed above, market structure enters as the transmission margin through which policy moves later equilibrium objects beyond their fixed-structure response. The present subsection is therefore the quantitative counterpart to that second term. In the calibrated economy, policy does not map into later income and welfare objects through a single direct deterrence margin. It first reshapes equilibrium market structure, and that change in market structure then propagates into product-market power, labor-market power, and the distribution of income across the independent and corporate sectors. The quantitative analysis therefore treats $N^*$ as a first-class

equilibrium object.

Figure 2 reports the calibrated equilibrium mass of firms over the full two-dimensional policy space. It shows directly how the endogenous degree of market fragmentation varies across $(\lambda, \bar{\xi})$. The main point is not a globally monotone comparative static. Rather, the figure shows that the incidence of policy on market structure is region-specific, so the same policy instrument can operate differently across nearby parts of the policy space.

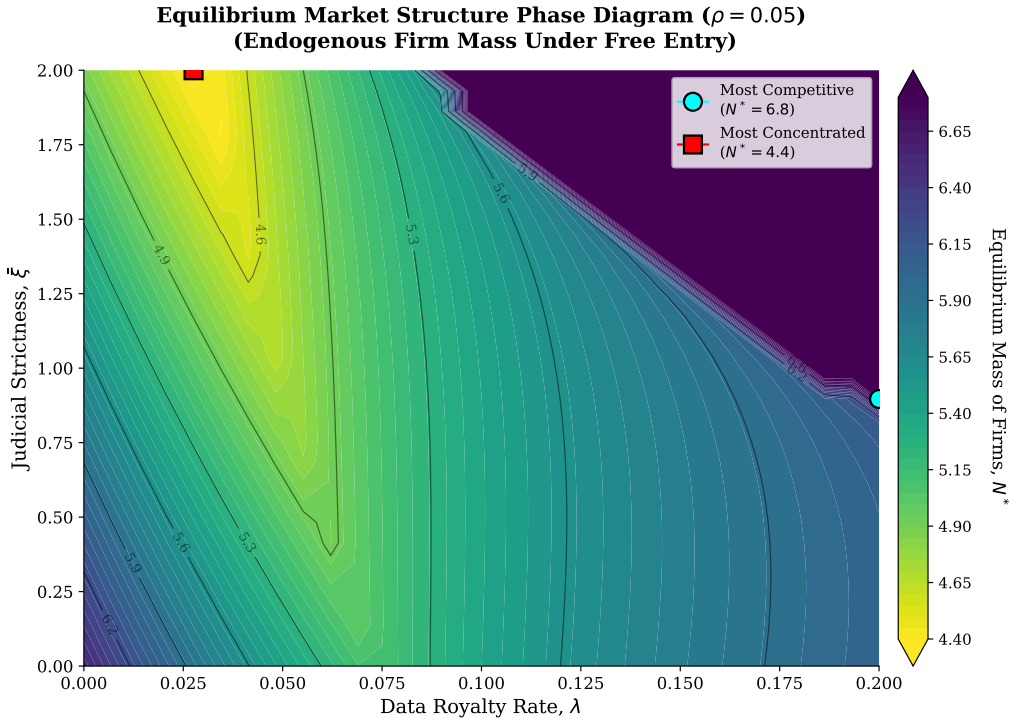



Figure 2: Equilibrium firm mass in the two-dimensional policy space

This two-dimensional heterogeneity is quantitatively important. The response of $N^*$ to $\lambda$ depends on the level of judicial strictness $\bar{\xi}$, and the calibrated surface is not well summarized by a single policy narrative. For this reason, the next step is to translate variation in $N^*$ into the corresponding changes in product-market and labor-market power.

Figure 3 quantifies this transmission mechanism directly. Panel (a) plots equilibrium firm mass against the royalty rate $\lambda$ for several values of judicial strictness $\bar{\xi}$. Panel (b) then maps the implied market structure into the product-market markup object $\mathcal{M}(N)$ and the labor-market wedge object $\Omega(N)$. Holding the maintained calibration fixed, lower firm mass is associated with stronger power on both margins: fewer firms raise the markup through reduced competitive pressure in the corporate goods market, and they also enlarge each firm's influence on residual labor supply, thereby widening the wedge between the marginal factor cost of labor and the wage. In this sense, market structure is the bridge from policy to later incidence objects.

Figure 3 also illustrates why the quantitative analysis is required. The revised theory shows that the sign of the policy effect on $N^*$ is not generally pinned down in closed form once entry, data choice, appropriability, wages, and relative prices adjust jointly in general equilibrium. The calibrated exercise does not overturn that point. Instead, it identifies the equilibrium patterns generated by the maintained parameterization. Along some policy slices, higher royalty rates reduce firm mass in a broadly consolidationary way; along others, the response is weaker or changes slope over the relevant range. The disciplined reading is therefore local

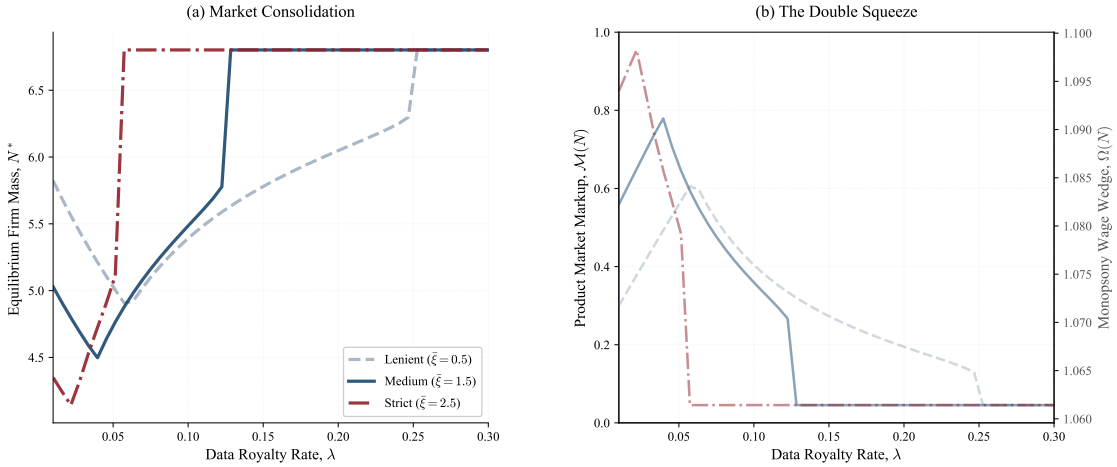



Figure 3: Market structure and the transmission of dual market power

*Notes:* Panel (a) reports the calibrated equilibrium mass of firms $N^*$ as a function of the royalty rate $\lambda$ for several values of judicial strictness $\bar{\xi}$ in the baseline post-shock technology environment $\rho = 0.05$. Panel (b) reports the corresponding product-market markup object $\mathcal{M}(N)$ and labor-market wedge object $\Omega(N)$. The figure is interpreted here purely as a calibrated mapping from policy into equilibrium market structure and from market structure into the two market-power margins. Any legacy proposition labels appearing inside the graphic are not part of the paper's substantive interpretation.

and quantitative rather than theorem-like. The calibration is used to determine which margins dominate in empirically relevant regions of the policy space, while preserving the theoretical point that incidence is regime-dependent and need not be globally monotone.

This feature is especially clear in Figure 4, which reports an informative one-dimensional policy slice in the baseline post-shock environment. The left panel shows that market consolidation can be non-monotone as $\lambda$ varies, while the right panel shows that royalty revenue can display a Laffer-type pattern. The interpretation is again restrained. A higher royalty rate raises the transfer per unit of licensed data, but it also feeds back into equilibrium entry and data usage, so the royalty base need not expand with the rate itself. Once policy becomes sufficiently burdensome along this calibrated slice, the endogenous contraction in the base offsets the higher rate and royalty revenue bends back. This is not presented as a universal property of the model, nor do the older in-figure regime labels carry independent analytical weight here. The figure is useful because it makes visible a quantitative implication of the revised theory: policy incidence may vary sharply across nearby regions of the state space, and the same instrument can operate through different combinations of entry, data intensity, and market-power adjustment.

The main implication for the remainder of the quantitative analysis is straightforward. Once policy first moves equilibrium market structure, later income and welfare objects should be decomposed along the channels opened by that endogenous transmission mechanism. The next subsection therefore turns to nominal artist-income channels—independent sales income, royalty income, and corporate-sector wage income—while the subsequent welfare analysis aggregates those channels in real terms rather than treating any single nominal object as dispositive.

Figure 5 provides the quantitative counterpart to the distinction developed in the theory between within-structure policy incidence and endogenous market-structure transmission. Over the range in which AI adoption remains active, the within-structure benchmark delivers systematically higher real artist welfare than the full equilibrium with endogenous firm mass. The gap between the two curves is the quantitative manifestation of market-structure offset. Holding firm mass fixed preserves the direct incidence of the policy

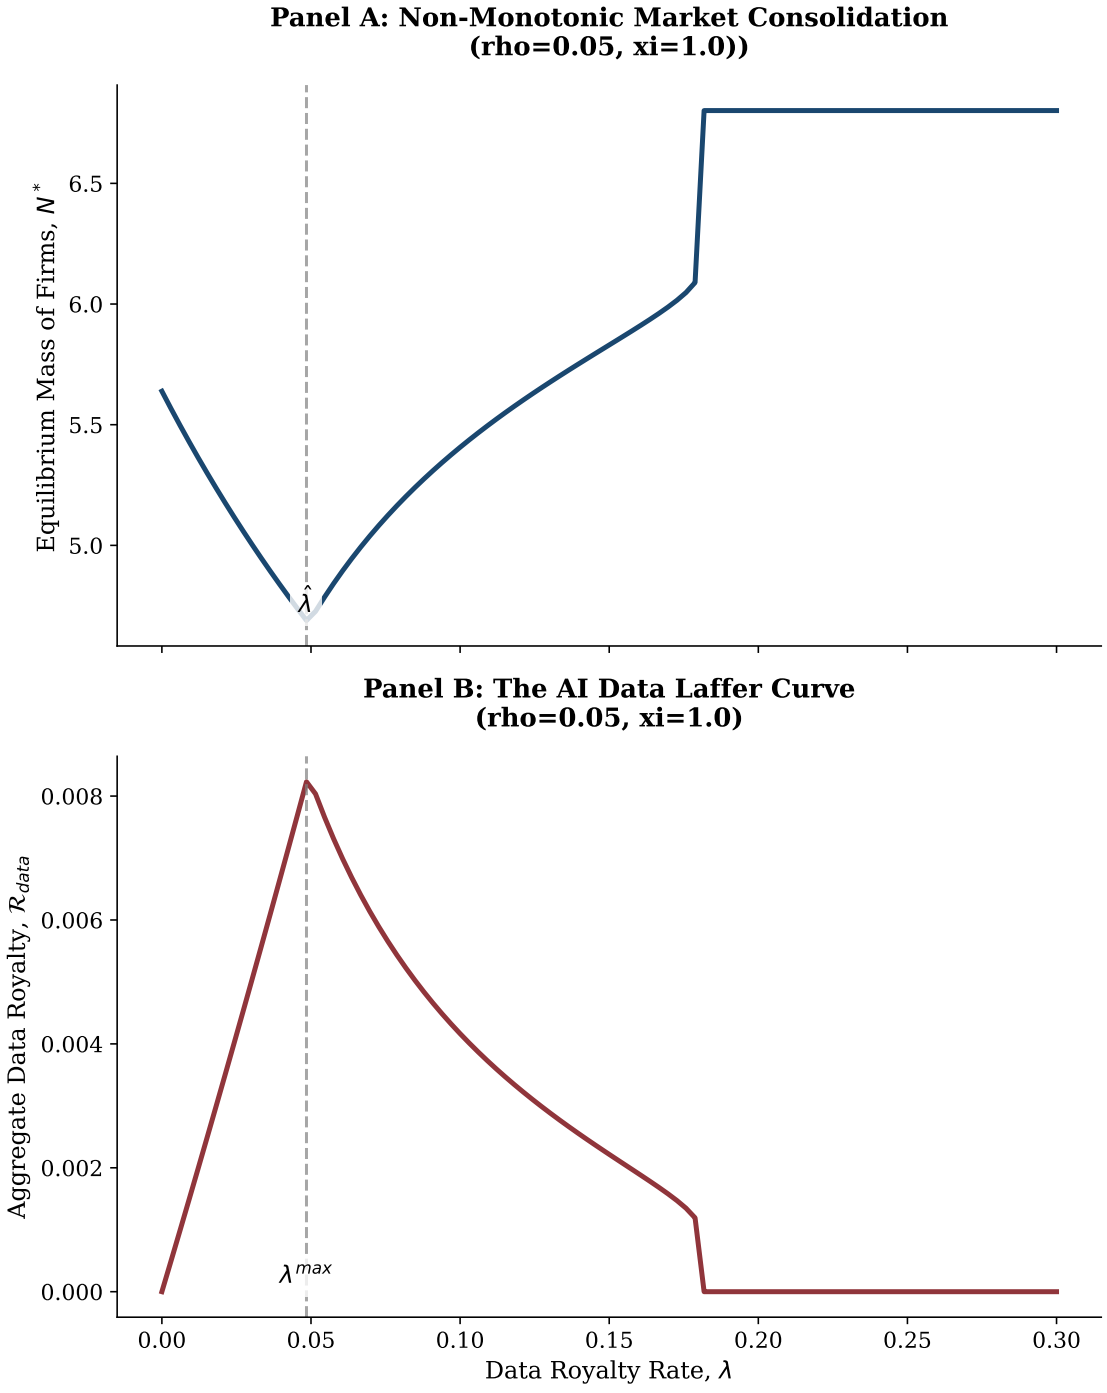



Figure 4: A calibrated policy slice: consolidation and the royalty base

*Notes:* The figure reports a calibrated one-dimensional slice of the policy space in the baseline post-shock technology environment $\rho = 0.05$. The upper panel plots equilibrium firm mass $N^*$ against the royalty rate $\lambda$ and the lower panel plots the associated royalty-revenue object.

change within a given market structure, whereas the full equilibrium also incorporates the induced response of entry and concentration. Along this policy slice, restoring that market-structure response lowers artist welfare relative to the within-structure benchmark throughout the interior region.

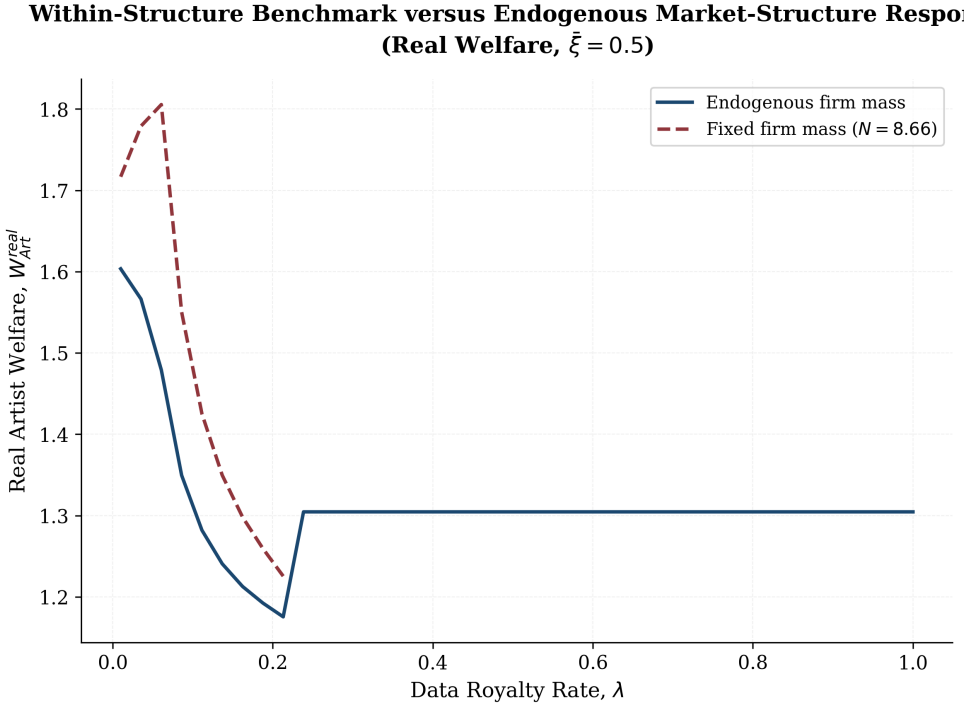



Figure 5: Within-structure benchmark versus endogenous market-structure response

*Notes:* The figure plots real artist welfare against the royalty rate $\lambda$ along the policy slice with fixed judicial strictness $\bar{\xi} = 0.5$. The solid line reports the full equilibrium with endogenous firm mass. The dashed line reports the within-structure benchmark that holds firm mass fixed at its reference level and solves the remaining equilibrium objects conditionally. The dashed curve is shown only over the range in which the economy remains in the AI-adoption region. Once the full equilibrium enters the no-AI regime, the benchmark with fixed firm mass is no longer the relevant comparison object, since in that regime $D = 0$ and firm mass is pinned down by the corner equilibrium system itself.

The figure also clarifies how the regime transition reshapes incidence. For low royalty rates, the within-structure benchmark initially displays a modest improvement in real artist welfare, while the full equilibrium declines more rapidly. This comparison shows that the local gain visible under a fixed market structure is substantially weakened once the induced adjustment in firm mass is restored. At higher royalty rates, the endogenous-equilibrium curve becomes exactly flat because the economy enters the no-AI regime. In that region, the equilibrium satisfies $D = 0$, the royalty payment term disappears, and further variation in $\lambda$ no longer changes the allocation relevant for artist welfare. The flat segment therefore reflects a regime switch rather than a smooth attenuation of marginal policy effects.

This interpretation also explains why the within-structure benchmark is reported only over the lower part of the policy slice. The benchmark is designed to isolate policy incidence conditional on an AI-adopting market structure. Once the equilibrium moves to the no-AI regime, firm mass is determined by the corner equilibrium conditions together with $F_{base}$, so an exogenously fixed-$N$ comparison no longer represents the relevant within-structure object. Quantitatively, the figure therefore separates two effects. Within the AI-adoption region, it measures the welfare gap generated by endogenous market-structure transmission. Beyond the regime boundary, the relevant change is the discrete transition to a no-AI equilibrium.

## 5.4 Quantifying nominal artist-income channels

The previous subsection established the central transmission mechanism of the calibrated economy: policy first changes equilibrium market structure, and those changes then propagate into markups, labor-market power, and the distribution of income across independent production and the corporate sector. The next step is therefore to trace those equilibrium adjustments through artists' nominal income sources. This subsection should be read as the artist-side quantitative counterpart to the theoretical distinction between within-structure incidence and endogenous market-structure transmission. For this purpose, it is useful to decompose nominal artist income into independent sales income, data royalty income, and corporate-sector wage income. Let $W_1 \equiv P_H^* Y_H^*$ denote independent sales income, $W_2 \equiv \lambda P_H^* D^* N^*$ denote aggregate data royalty income, and $W_3 \equiv w^* L_{Corp}^*$ denote corporate-sector wage income, where $L_{Corp}^*$ is equilibrium labor employed by the corporate sector. Their nominal sum is $W_{Art}^{nom} \equiv W_1 + W_2 + W_3$. The purpose of this decomposition is not to recover a closed-form sign theorem. Rather, it is to show how each nominal component responds first through its own within-structure policy margin and then through the equilibrium change in $N^*$. The decomposition is therefore informative precisely because the three components need not move together across the two-dimensional policy space, while remaining incomplete as a welfare criterion because it does not yet account for the aggregate price level.

The first component, independent sales income $W_1$, captures the nominal value of human-created output sold outside the corporate sector. In the calibration, this component tends to rise in policy configurations that reallocate expenditure toward human art. The increase reflects the joint response of the value of independent production, combining changes in the relative attractiveness of human-created content with equilibrium changes in supply. The appropriate way to interpret the independent-sales channel is that policy can shift nominal income toward independent creators, but the magnitude and even the direction of that shift depend on the equilibrium response of market structure and expenditure allocation rather than on any globally signed price comparative static.

The second component, data royalty income $W_2$, is the clearest place where the distinction between direct incidence and market-structure transmission becomes visible. Holding the royalty base fixed, a higher $\lambda$ mechanically raises transfers per unit of lawful data use. That is the direct within-structure margin. In equilibrium, however, the royalty base is endogenous. A higher $\lambda$ can reduce firm-level data demand $D^*$, reduce equilibrium firm mass $N^*$, or both. Since data are non-rival, the relevant physical constraint remains firm-level, $D_i \leq Y_H$, so the royalty base contracts through equilibrium firm behavior rather than through any aggregate scarcity condition. Quantitatively, this generates the possibility of a hump-shaped royalty schedule: at low values of $\lambda$, the transfer effect can dominate, whereas at higher values the contraction in equilibrium data use and corporate activity erodes the base on which royalties are paid. The royalty panel should therefore be read as the numerical combination of a direct transfer margin and an endogenous market-structure offset, not as a sufficient statistic for artist welfare.

The third component, corporate-sector wage income $W_3 = w^* L_{Corp}^*$, is shaped by two opposing equilibrium forces and is therefore especially useful for illustrating the decomposition logic. Within the AI-adoption region, one important margin runs through endogenous market structure: when policy reduces the equilibrium mass of active firms, the corporate sector becomes more concentrated, the labor-market wedge strengthens, and the wage bill received by artists from corporate employment is compressed. This operates through the equilibrium wage rate $w^*$, through corporate labor demand $L_{Corp}^*$, or through both jointly. At the same time, policy also changes the production mix within the corporate sector. As the royalty rate $\lambda$ rises further, AI use becomes less profitable and the corporate sector becomes more labor-intensive; near the relevant regime boundary, this reallocation toward labor can raise corporate employment and thereby increase

the nominal wage bill. The upward movement in $W_3$ at higher $\lambda$ should therefore be interpreted not as a pure wage effect or a pure employment effect in isolation, but as the product of a changing wage-employment bundle once direct policy incidence and market-structure transmission are combined. The role of judicial strictness $\bar{\xi}$ is analogous. Higher $\bar{\xi}$ weakens the profitability of AI-assisted production and therefore tends to shift the economy earlier toward a labor-intensive or no-AI configuration, which is why the transition to the high-$W_3$ region occurs at lower $\lambda$ under stricter judicial environments. Quantitatively, the wage-income panel is therefore best read as a regime-sensitive incidence object: within the AI-adoption region, greater concentration tends to depress the corporate wage bill, whereas near and beyond the adoption boundary, substitution back toward labor can dominate and raise $W_3$.

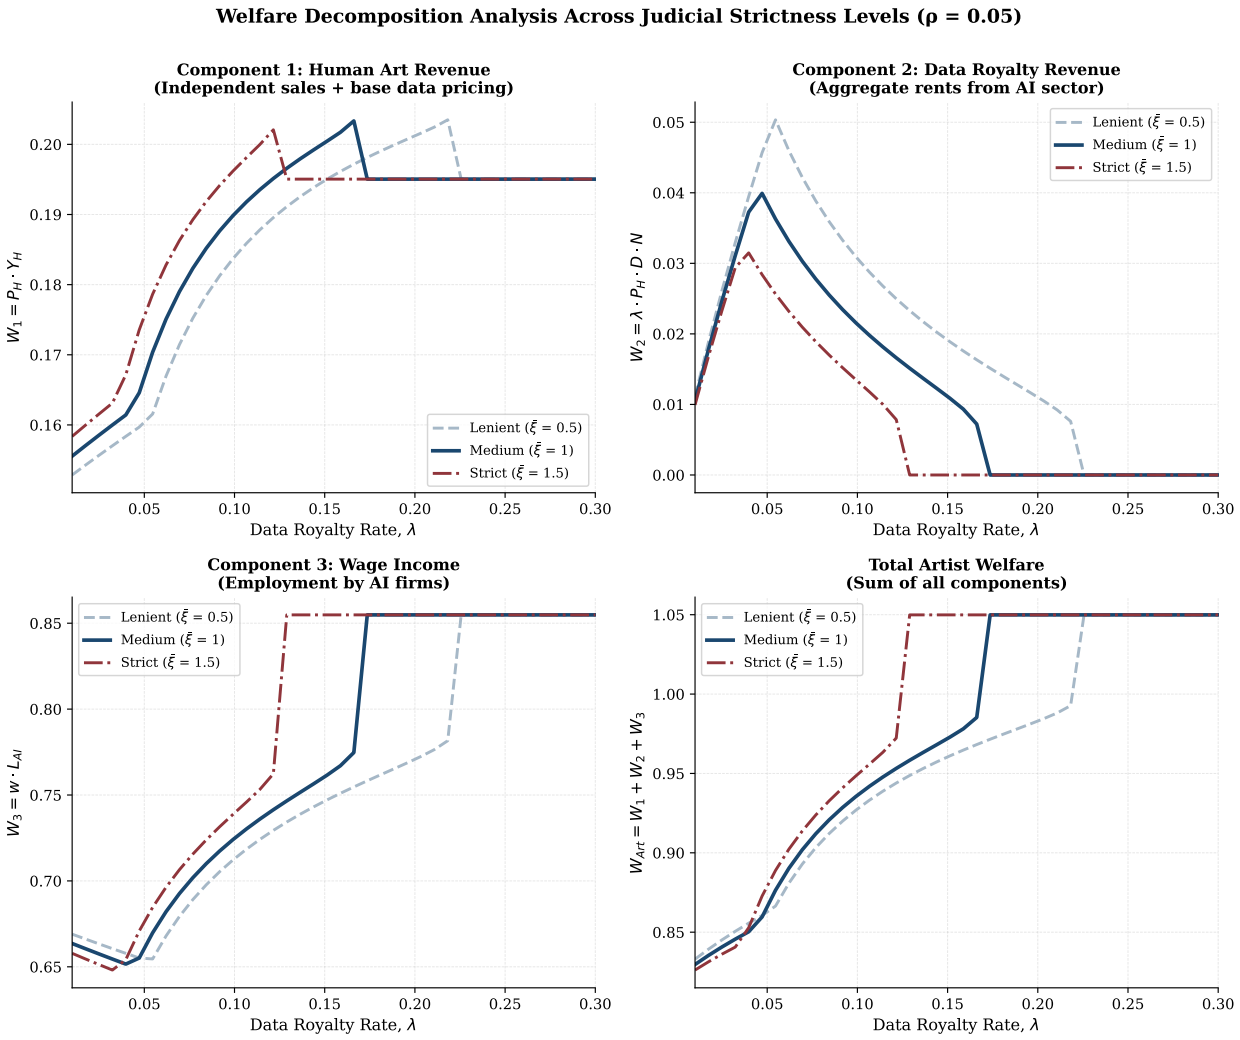



Figure 6: Nominal artist-income channels over the royalty rate

*Notes.* The figure plots the three nominal artist-income components $W_1 = P_H^* Y_H^*$, $W_2 = \lambda P_H^* D^* N^*$, and $W_3 = w^* L_{Corp}^*$ against the royalty rate $\lambda$ for several values of judicial strictness $\bar{\xi}$, holding $\rho = 0.05$ fixed. The first three panels report independent sales income, data royalty income, and corporate-sector wage income, respectively. The fourth panel is interpreted in this subsection as the nominal aggregate $W_{Art}^{nom} = W_1 + W_2 + W_3$.

Figure 6 makes these distinctions visible in the maintained calibration. The three nominal channels respond differently as policy varies, and the nominal aggregate inherits that heterogeneity rather than collapsing it. This is exactly why the decomposition is useful: it shows whether a given policy configuration

shifts artists' nominal income toward independent sales, toward royalty transfers, or toward corporate wages, and it clarifies which of those shifts are associated with the underlying market-structure response. At the same time, the nominal aggregate should not be treated as the paper's final policy criterion. A policy that raises nominal artist income through one or more channels need not improve artists' real purchasing power once the aggregate price level and broader macro incidence are taken into account. The next subsection therefore moves from nominal channel accounting to real incidence.

## 5.5 Real Macro Incidence

The previous subsection decomposed artists' *nominal* income into independent sales income, data royalty income, and corporate-sector wage income. That accounting is useful for identifying the channels through which policy operates, but it is not sufficient for welfare evaluation. The present subsection therefore recombines those nominal channels with the remaining general-equilibrium objects that matter for real incidence, most importantly the aggregate price level, technological utilization, and equilibrium reallocation. In the present environment, policy changes the aggregate price level through endogenous market structure, adoption, and sectoral adjustment, so the relevant quantitative objects are real rather than nominal. Figure 7 therefore moves from nominal channel accounting to real general-equilibrium incidence by tracing, along a policy slice at the maintained post-shock technology environment, the responses of per-firm data use, the aggregate price level, real output, and real artist welfare.

Panel A shows how technological utilization varies with the royalty rate $\lambda$ for different values of judicial strictness $\bar{\xi}$. Under the maintained calibration, data use need not move monotonically along this slice. For some values of $\bar{\xi}$, equilibrium data use rises modestly at low royalty rates and then declines as the royalty burden increases further; for higher judicial strictness, the interior-adoption region becomes narrower and the transition out of positive data use occurs at lower values of $\lambda$. This pattern is consistent with the regime structure established in the theory section. Along the numerical policy grid, smooth changes in data use correspond to movement within a given regime, while sharp transitions reflect switches between the interior-adoption, binding-data, and manual configurations selected in equilibrium.

Panels B and C translate these adoption responses into aggregate real incidence. The aggregate price level is not an ancillary object here: it summarizes how policy-induced changes in firm mass and production structure feed through to economy-wide purchasing power. When policy compresses equilibrium firm mass, the market-structure channel identified earlier implies stronger markup power and a larger labor-market wedge, both of which tend to raise the aggregate price level. At the same time, reduced technological utilization weakens productive efficiency in the corporate sector. The combined result is that real output can display non-monotonic movements along the policy slice. In particular, the real consequences of a higher royalty rate depend not only on the direct cost of data use, but also on whether the induced equilibrium adjustment primarily operates through lower data intensity within the adoption regime, through changes in firm mass, or through a transition back to the manual regime.

Panel D turns to real artist welfare, which is the relevant object for policy evaluation. Relative to the nominal decomposition in the previous subsection, the main quantitative message is not merely that prices matter in principle, but that they matter enough to change the incidence pattern materially. In particular, even where higher $\lambda$ raises some components of nominal artist income over part of the policy slice, those nominal gains do not translate one-for-one into higher purchasing power. Along the interior-adoption region, real artist welfare declines for all three values of $\bar{\xi}$, indicating that the rise in the aggregate price level and the accompanying equilibrium reallocation more than offset the nominal improvement recorded in the artist-income channels. This pattern is consistent with Panels A–C taken jointly: as the royalty burden

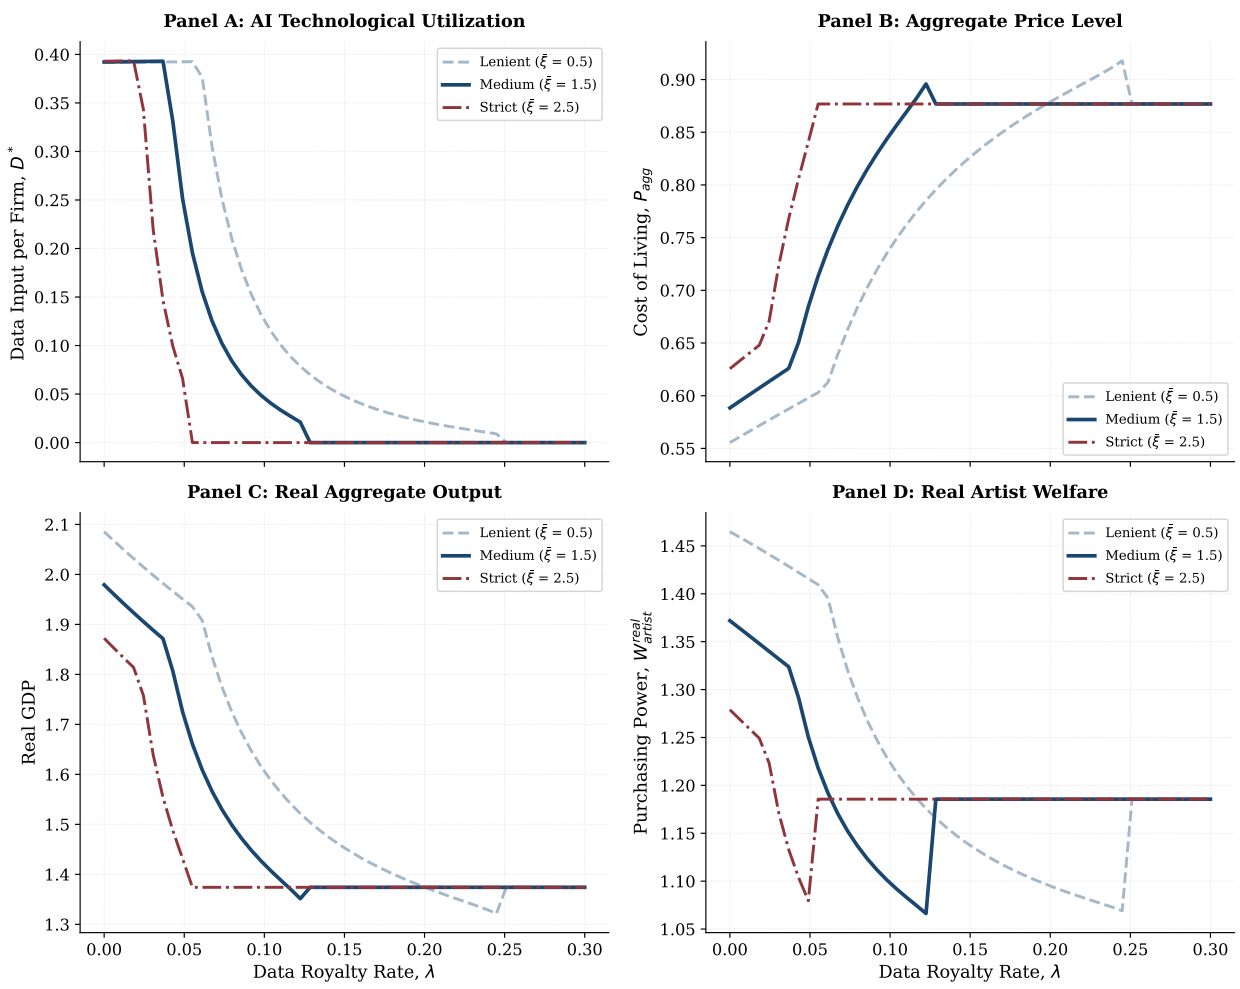



Figure 7: Real macro incidence along a policy slice

*Notes:* The figure reports four equilibrium objects plotted against the data royalty rate $\lambda$ for several values of judicial strictness $\bar{\xi}$ at $\rho = 0.05$: per-firm data use $D^*$, the aggregate price level $P^*$, real aggregate output, and real artist welfare. The figure is interpreted as a policy-slice view of real general-equilibrium incidence. Sharp visual transitions should be read as switches across equilibrium regimes on the numerical grid.

increases, data use is compressed, market structure becomes less competitive, the aggregate price level rises, and real output falls, so that artist incidence worsens in real terms even when nominal receipts need not fall contemporaneously. The flat common segment that appears once $D^* = 0$ should in turn be read as a regime outcome rather than as a separate comparative-static result: after the economy reverts to the manual regime, the data-use margin is shut down and real artist welfare converges to the corresponding manual benchmark. The broader implication is that nominal channel accounting is informative about transmission, but the relevant welfare ranking must be based on deflated outcomes in general equilibrium.

Taken together, Figure 7 shows that policy operates on real outcomes through a joint system of technological utilization, market structure, relative prices, and labor reallocation. These are calibrated equilibrium patterns rather than general theorems, but they establish the central point needed for the final quantitative step: once aggregate price incidence is taken into account, neither nominal transfers nor partial-equilibrium adoption responses are sufficient to rank policies. The next subsection therefore evaluates real artist welfare over the full two-dimensional policy space $(\lambda, \bar{\xi})$ and uses that object to compare alternative policy combinations.

## 5.6 Policy evaluation over the two-dimensional policy space

The final baseline evaluation object is real artist welfare over the full two-dimensional policy space $(\lambda, \bar{\xi})$. This is the natural quantitative counterpart to the theory. Once policy operates jointly through redistribution, appropriability, entry, prices, and labor allocation, evaluation cannot be reduced to a one-dimensional comparative static or to the movement of a single nominal transfer. The relevant object is the calibrated welfare surface itself.

Figure 8 reports that surface at the maintained post-shock technology environment $\rho = 0.05$. Under this calibration, real artist welfare is maximized at the origin, $(\lambda, \bar{\xi}) = (0, 0)$. Moving away from that point generally lowers welfare, but the decline is not mechanically one-dimensional. In the interior adoption region, the iso-welfare contours slope downward, indicating that the two policy instruments interact quantitatively in shaping overall regulatory burden. The figure also displays a lighter upper-right region corresponding to the manual regime after endogenous AI abandonment. In that region, welfare becomes locally flat with respect to further increases in policy burden, reflecting the fact that once the economy has switched to the no-adoption regime, additional tightening no longer operates through AI-use margins in the baseline equilibrium.

The economic interpretation of the surface follows directly from the earlier quantitative blocks. Welfare is relatively low in portions of the interior region where policy compresses firm mass sufficiently to worsen real incidence through the general-equilibrium transmission mechanism identified above. In those areas, the relevant force is not the nominal transfer alone, but the combination of more concentrated product markets, a stronger labor-market wedge, and the resulting effect on real purchasing power. By contrast, the manual region should not be read as uniformly desirable merely because it eliminates AI-related distortions. In the baseline calibration it appears as a plateau because, conditional on no adoption, further policy tightening no longer changes equilibrium outcomes through the AI channel; but the level of that plateau remains an equilibrium welfare outcome, not a normative benchmark. The figure is therefore best read as a synthesis of the regime map, the market-structure transmission results, and the real-incidence analysis, rather than as a stand-alone mechanism proof.

The calibrated policy implication is consequently narrow and conditional. For the maintained parameterization and $\rho = 0.05$, the baseline welfare surface assigns the highest real artist welfare to the least restrictive corner of the policy space. This does not imply a general optimal-policy theorem, nor does it imply that

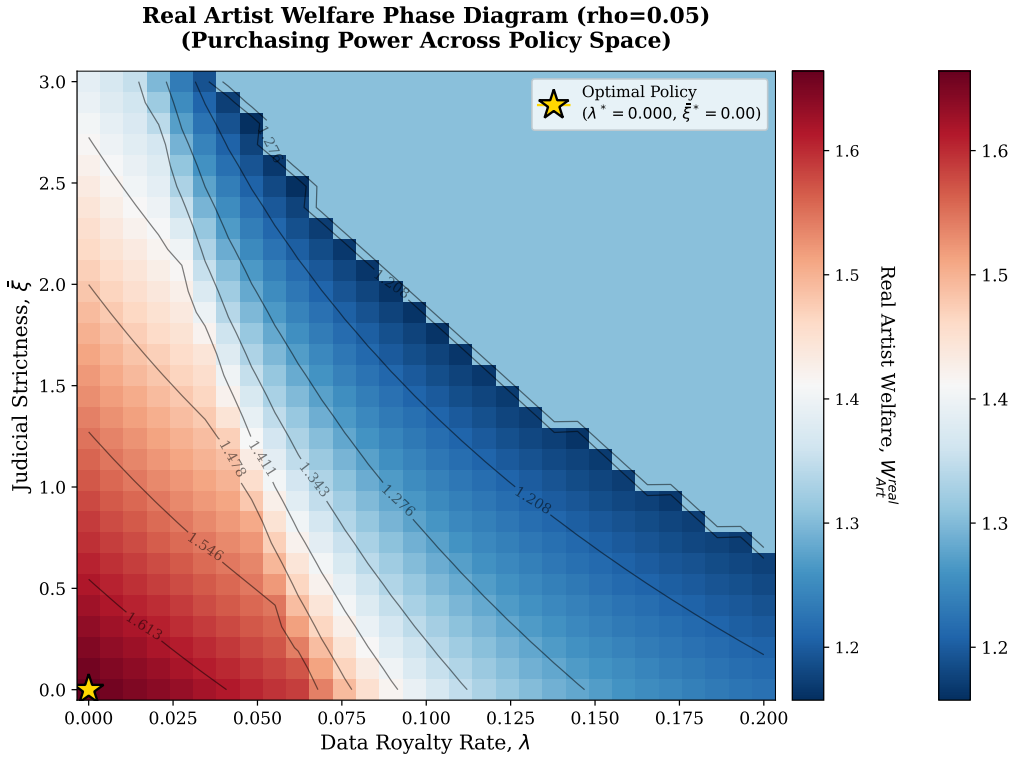



Figure 8: Real artist welfare over the two-dimensional policy space

*Notes:* The figure plots real artist welfare over the policy space spanned by the data royalty rate $\lambda$ and judicial strictness $\bar{\xi}$ at the maintained post-shock technology environment $\rho = 0.05$. The marked point indicates the welfare maximum in the baseline calibration. Contour lines summarize iso-welfare combinations of the two instruments. The lighter upper-right region corresponds to policy combinations under which the selected equilibrium reverts to the manual regime after endogenous abandonment of AI.

either instrument is welfare-reducing in all environments. What it shows is that, in this calibrated post-shock economy, the adverse general-equilibrium consequences of tighter policy dominate the nominal artist-income gains that some dimensions of regulation can generate. Put differently, nominal channel improvements are quantitatively relevant but not decisive for final evaluation once they are filtered through equilibrium prices, labor allocation, and endogenous market structure.

This calibration-contingent reading also motivates the next subsection. The location and shape of the welfare surface need not be invariant to the technological environment, since the strength of AI as a production technology affects both adoption incentives and the equilibrium margins through which policy operates. The next subsection therefore studies how the two-dimensional policy surface changes with $\rho$, and whether the baseline ranking in Figure 8 is robust to alternative technological conditions.

## 5.7 Sensitivity and Robustness

This subsection examines the robustness of the welfare results to changes in several deep parameters that load on the equilibrium offsets emphasized by the theory. I begin with the relative computing cost $\rho$, which governs the technological environment in which policy operates. I then vary the firm-level labor-supply elasticity $\epsilon_L$, the cross-sector substitution elasticity $\sigma$, and the learning parameter $\mu$. In all four exercises, the vertical axis is reported in relative form: for each alternative parameterization, the plotted object is real artist welfare relative to the baseline parameterization at the same royalty rate. This point-by-point normalization removes level differences across calibrations and isolates how changing the parameter of interest shifts welfare relative to the benchmark policy schedule.

Figure 9 varies the relative computing cost $\rho$ while holding judicial strictness fixed at its baseline value. The quantitative message is that the welfare consequences of copyright policy depend materially on the technological maturity of artificial intelligence. When computing costs are lower, the productivity gains from AI adoption are larger, and real artist welfare can lie substantially above the baseline calibration over a broad range of royalty rates. When computing costs are higher, those gains are attenuated and welfare falls below the baseline schedule at low royalty rates. The figure therefore shows that the equilibrium incidence of copyright policy is itself technology-dependent, even when evaluated relative to the same policy point in the baseline environment.

I next vary the firm-level labor-supply elasticity, which governs the strength of the labor-market wedge and therefore the intensity of market-structure transmission on the wage side.

Figure 10 shows that the welfare consequences of copyright policy depend materially on the strength of labor-market power. Relative to the baseline calibration $\epsilon_L = 2$, higher values of $\epsilon_L$ raise real artist welfare over most of the policy range, with the welfare ratio under $\epsilon_L = 7$ lying above that under $\epsilon_L = 4$. The compression of the ratio at low and intermediate royalty rates indicates that stronger labor-market competition attenuates the welfare loss generated by equilibrium offsets. The fact that the gap remains visible even on the high-$\lambda$ plateau further indicates that labor-market conditions continue to matter after the economy has entered the no-AI regime.

Figure 11 shows that the welfare profile is also sensitive to the degree of cross-sector substitution. Relative to the baseline calibration $\sigma = 2.5$, lower values of $\sigma$ reduce real artist welfare, while higher values raise it. The ordering is stable over most of the policy range, indicating that stronger substitutability attenuates the real-incidence loss associated with a given policy environment. At the same time, the non-monotonic adjustment at low royalty rates remains visible across the full range considered. The main effect of varying $\sigma$ is therefore to shift welfare relative to the benchmark schedule rather than to overturn its qualitative structure.

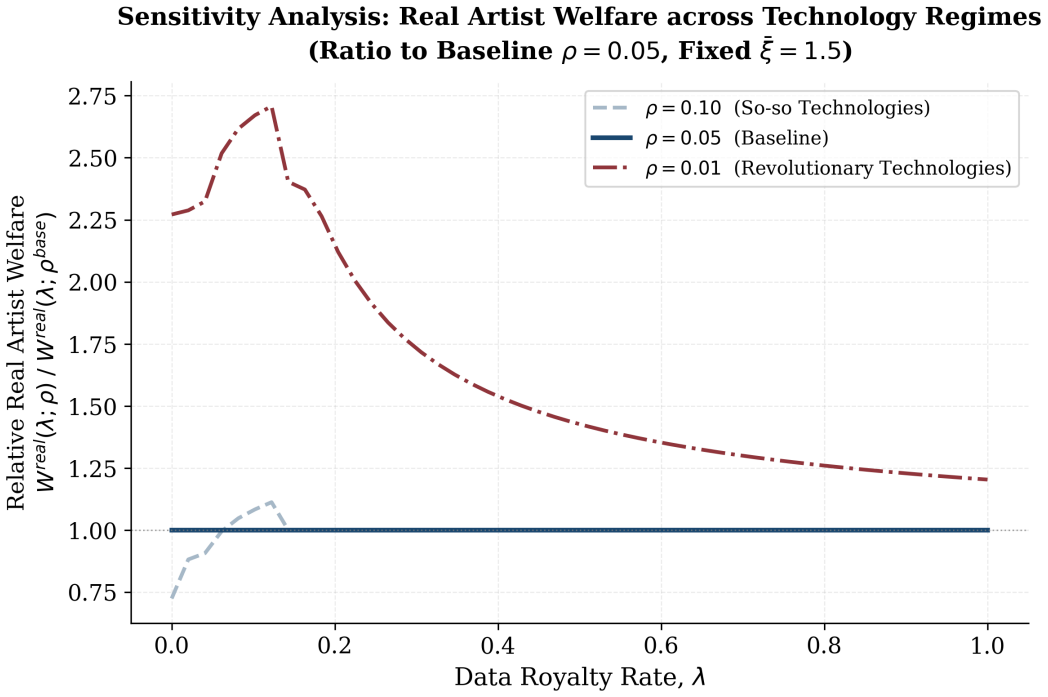



Figure 9: Sensitivity of real artist welfare to relative computing cost ($\rho$)

*Notes:* This figure plots $W_H^{real}(\lambda;\rho)/W_H^{real}(\lambda;\rho^{base})$ against the data royalty rate $\lambda$ for alternative values of the relative computing cost $\rho$, holding judicial strictness fixed at the baseline value $\bar{\xi} = 1.5$. The baseline technology parameter is $\rho^{base} = 0.05$, so the baseline curve is equal to one wherever it is defined.

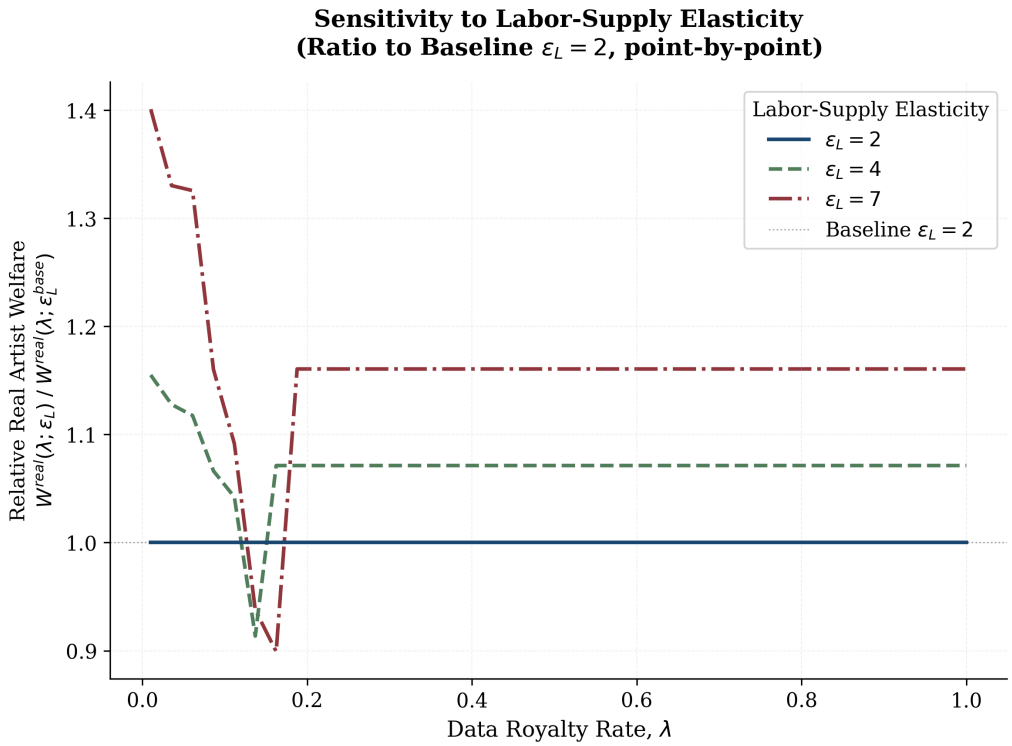



Figure 10: Sensitivity to labor-supply elasticity

*Notes:* The figure plots $W_H^{real}(\lambda;\epsilon_L)/W_H^{real}(\lambda;\epsilon_L^{base})$ against the royalty rate $\lambda$ along the policy slice with fixed judicial strictness at its baseline value, for alternative values of the firm-level labor-supply elasticity $\epsilon_L$. The baseline value is $\epsilon_L^{base} = 2$. A lower $\epsilon_L$ corresponds to stronger labor-market power on the firm side and therefore a larger wage markdown wedge. All other parameters are held at their baseline calibrated values.

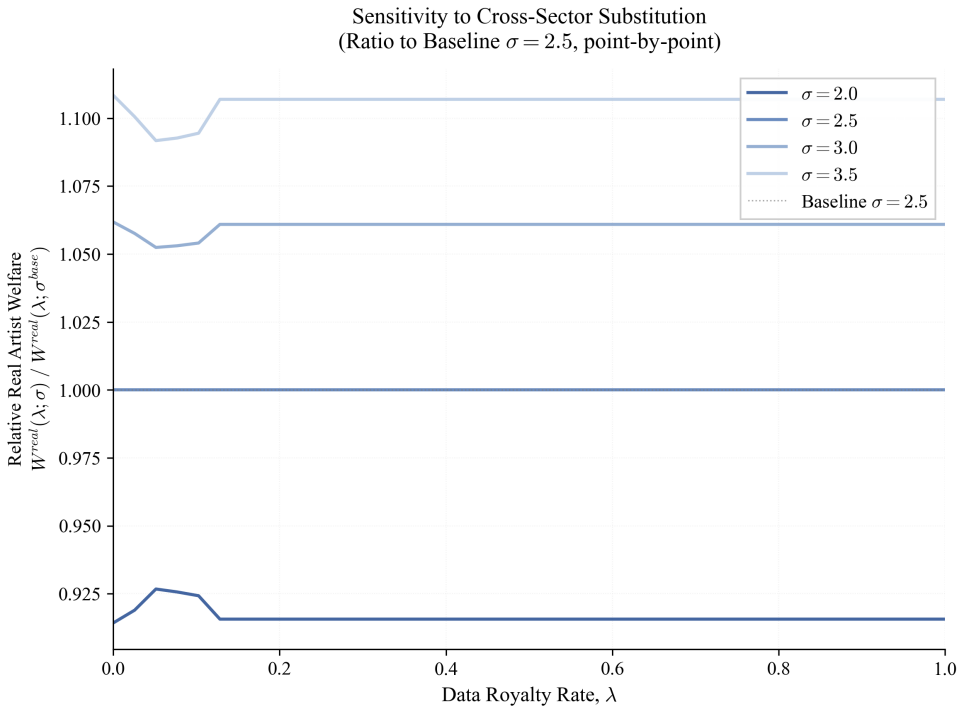



Figure 11: Sensitivity to cross-sector substitution

*Notes:* The figure plots $W_H^{real}(\lambda;\sigma)/W_H^{real}(\lambda;\sigma^{base})$ against the royalty rate $\lambda$ along the policy slice with fixed judicial strictness at its baseline value, for alternative values of the cross-sector substitution elasticity $\sigma$. The baseline value is $\sigma^{base} = 2.5$. For each value of $\sigma$, the corresponding $\eta$ is recomputed using the calibration mapping described in the text, while all other parameters are held at their baseline values.

Figure 12 shows that the learning parameter materially affects both the automation response and the welfare consequences of copyright policy. Lower values of $\mu$ raise welfare relative to the baseline calibration over most of the policy range, while higher values reduce it at low royalty rates before gradually recovering. The automation panel clarifies the mechanism: higher $\mu$ makes automation exposure more sensitive to the policy environment, so that $I(D)$ contracts more sharply as the royalty rate rises. On the welfare side, this stronger contraction of automation is associated with a weaker welfare schedule relative to the baseline along the transition path. Lower values of $\mu$ instead generate a more gradual automation response and a correspondingly stronger welfare profile. Quantitatively, $\mu$ governs the strength of the reallocation offset.

Taken together, Figures 9–12 show that the quantitative severity of the welfare loss depends on the technological environment and on the deep parameters governing labor-market transmission, relative-price adjustment, and automation responsiveness. The central message of the calibration nevertheless remains stable. Across all four exercises, policies with an immediately favorable margin can still generate substantial equilibrium offsets, and real artist welfare cannot be inferred from the direct policy margin alone.

# 6 Conclusion

This paper studies AI copyright policy in a two-dimensional space defined by the data royalty rate $\lambda$ and judicial strictness $\bar{\xi}$ in an environment with endogenous AI adoption, endogenous firm entry, endogenous market structure, and artist welfare determined in general equilibrium. In this setting, copyright policy does not operate solely through the immediate transfer associated with lawful data use or through the direct

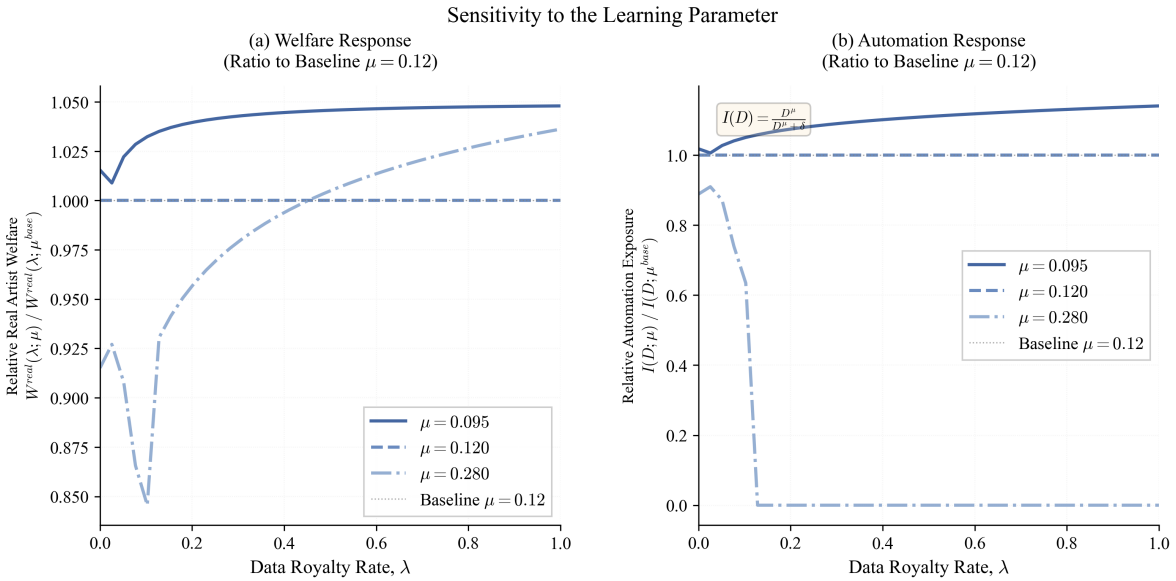



Figure 12: Sensitivity to the learning parameter

*Notes:* Panel (a) plots $W_H^{real}(\lambda;\mu)/W_H^{real}(\lambda;\mu^{base})$ against the royalty rate $\lambda$ along the policy slice with fixed judicial strictness at its baseline value, for alternative values of the learning parameter $\mu$, where the baseline value is $\mu^{base} = 0.12$. Panel (b) reports the corresponding automation-exposure ratio $I(D;\mu)/I(D;\mu^{base})$. All other parameters are held at their baseline calibrated values.

restriction imposed on AI-assisted output. Firms adjust data use and entry, labor reallocates across sectors, and prices respond jointly with concentration. Policy evaluation therefore depends on how these immediate margins propagate through the equilibrium system.

The theory clarifies the structure of that propagation. It characterizes the relevant equilibrium regimes, distinguishes within-structure policy incidence from endogenous market-structure transmission, and shows that firm mass is a central transmission object because lower equilibrium entry strengthens both product-market markup power and the labor-market wedge. On the artist side, nominal incidence is organized through independent sales income, royalty income, and corporate wage income. Its economic interpretation depends on how policy-induced adjustments are filtered through market-structure offset and reallocation offset. The final welfare implication then depends on an additional real-incidence offset, since nominal receipts must be translated into purchasing power through the aggregate price level. The central theoretical lesson is that immediate policy intuition is incomplete once entry, allocation, and prices adjust endogenously.

The quantitative analysis evaluates these forces over the two-dimensional policy space under the maintained parameterization. It maps the equilibrium regimes implied by policy, quantifies the response of firm mass and the associated transmission to prices and labor outcomes, and compares policies using real artist welfare as the final evaluation object. Under the baseline post-shock environment, the calibration implies that tighter policy can improve some nominal artist-income components while still reducing real artist welfare once concentration, equilibrium reallocation, and aggregate price adjustment are taken into account. Over the baseline policy surface, real artist welfare is maximized at the least restrictive corner, while more restrictive combinations of $\lambda$ and $\bar{\xi}$ generally lower welfare and may eventually induce endogenous reversion to the manual regime.

The sensitivity analysis shows that the quantitative severity of these welfare losses depends on the technological environment and on the deep parameters that govern labor-market transmission, cross-sector substitution, and the responsiveness of automation to data. Even so, the broader message remains stable across

the exercises. Apparent protective or redistributive gains are not sufficient for policy evaluation in this environment. The relevant object is general-equilibrium incidence, mediated by endogenous market structure, by the reallocation of activity across artistic and corporate production, and by the mapping from nominal income into real purchasing power. The paper's contribution is to make these incidence channels explicit in a unified model of AI copyright regulation, endogenous market structure, and artist welfare.

These conclusions remain conditional on the empirical discipline imposed on substitution elasticities, labor-market power, adoption technology, and the broader incidence margins through which policy affects equilibrium prices and allocation. Sharper empirical discipline on those objects, together with applications to alternative creative industries and technological environments, would further refine the scope of the policy conclusions developed here.

# References


Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*.

Acemoglu, D. and Restrepo, P. (2022). Artificial intelligence, automation, and work. *Handbook of Labor, Human Resources and Population Economics*, pages 1–35.

Akcigit, U. and Ates, S. T. (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, 13(1):257–298.

Atkeson, A. and Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5):1998–2031.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.

Azar, J., Marinescu, I., and Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, 57(S):S167–S199.

Bajari, P., Chernozhukov, V., Hortaçsu, A., and Suzuki, J. (2019). The impact of big data on firm performance: An empirical investigation. *AEA Papers and Proceedings*, 109:33–37.

Berger, D. W., Herkenhoff, K. F., and Mongey, S. (2022). Minimum wages, efficiency, and welfare. *American Economic Journal: Macroeconomics*. Based on NBER Working Paper No. 29662.

Boldrin, M. and Levine, D. K. (2002). The case against intellectual property. *American Economic Review Papers & Proceedings*, 92(2):209–212.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.

de Roos, N. and McKenzie, J. (2014). Cheap tuesdays and the demand for cinema. *International Journal of Industrial Organization*, 33:93–109.

Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models.

Farboodi, M. and Veldkamp, L. (2021). A model of the data economy. *NBER Working Paper*, (28427). Revise & Resubmit, Review of Economic Studies.

Hémous, D. and Olsen, M. (2022). The rise of the machines: Automation, horizontal innovation, and income inequality. *American Economic Journal: Macroeconomics*, 14(1):179–223.

Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127–1150.

Jones, C. I. and Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–2858.

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. (2020). Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*.

Landes, W. M. and Posner, R. A. (1989). An economic analysis of copyright law. *Journal of Legal Studies*, 18(2):325–363.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press, Princeton, NJ.

Samuelson, P. (2023). Generative ai meets copyright. *Science*, 381(6654):158–161.

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., and Gal, Y. (2024). The curse of recursion: Training on generated data makes models forget. *Nature*, 631:755–759.

Sokolova, A. and Sorensen, T. (2021). Monopsony in labor markets: A meta-analysis. *ILR Review*, 74(1):27–55.

Villalobos, P., Ho, A., Sevilla, J., Besiroglu, T., Heim, L., and Hobbhahn, M. (2024). Will we run out of data? limits of llm scaling based on human-generated data. *arXiv preprint arXiv:2211.04325*. Epoch AI working paper.

Vogel, H. L. (2020). *Entertainment Industry Economics: A Guide for Financial Analysis*. Cambridge University Press, Cambridge, 10 edition.

Yang, S. A. and Zhang, A. H. (2024). Generative ai and copyright: A dynamic perspective. *SSRN Electronic Journal*. Working paper, London Business School and USC Gould School of Law.

# A Micro-foundation of the Automation Threshold $I(D)$

In this appendix, we derive the functional form of $I(D)$ from the statistical properties of machine learning.

Let tasks $j \in [0,1]$ be heterogeneous in their intrinsic complexity $\tau_j$, drawn i.i.d. from a cumulative distribution function $G(\tau)$. Following the empirical literature on neural scaling laws (e.g., Bajari et al., 2019; Kaplan et al., 2020), we assume the prediction error of an AI model trained on dataset $D$ follows a power law:

$$\epsilon(D) = D^{-\mu} \tag{62}$$

where $\mu > 0$ is the learning rate, and we normalize the baseline error scalar to unity.

A task $j$ is successfully automated if the model's prediction error falls below the task's complexity threshold: $\epsilon(D) \leq \tau_j$. The aggregate share of automated tasks, $I(D)$, is thus the survival function of the complexity distribution evaluated at the current error rate:

$$I(D) = \Pr(\tau_j \geq \epsilon(D)) = 1 - G(D^{-\mu}) \tag{63}$$

To capture the standard S-shaped diffusion of technology (Comin and Hobijn, 2010), we assume task complexity $\tau$ follows a **Log-Logistic distribution** with shape parameter $\kappa$ and scale parameter $\delta$:

$$G(\tau) = \frac{\tau^{\kappa}}{\tau^{\kappa} + \delta} \tag{64}$$

Substituting $G(\cdot)$ into the expression for $I(D)$ yields:

$$I(D) = 1 - \frac{(D^{-\mu})^{\kappa}}{(D^{-\mu})^{\kappa} + \delta} = \frac{D^{\mu\kappa}}{D^{\mu\kappa} + \delta} \tag{65}$$

Defining the effective elasticity $\tilde{\mu} \equiv \mu\kappa$, we obtain the functional form used in the main text:

$$I(D) = \frac{D^{\tilde{\mu}}}{D^{\tilde{\mu}} + \delta} \tag{66}$$

This derivation micro-founds the concave-convex shape of automation adoption as the interaction between power-law learning and heavy-tailed task complexity.

# B Micro-foundation of the Aggregate Production Function

In this appendix, we provide the micro-foundation for the aggregate production function used in the main text. We show that the Cobb-Douglas form arises as a limiting case of the task-based framework developed by Acemoglu and Restrepo (2018), specifically when the elasticity of substitution between tasks approaches unity.

## B.1 Task Aggregation

Following Acemoglu and Restrepo (2018), the final good $Y_i$ is produced by combining a continuum of tasks $j \in [0,1]$ with a constant elasticity of substitution (CES) $\sigma$. The aggregator is given by:

$$Y_i = \left( \int_0^1 y_i(j)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}} \tag{67}$$

To ensure the model remains tractable and focuses on the substitution between aggregate capital and labor driven by the extensive margin (automation) rather than the intensive margin, we assume the elasticity of substitution between tasks $\sigma \to 1$. Using L'Hôpital's rule, Equation (67) converges to the Cobb-Douglas form:

$$\ln Y_i = \int_0^1 \ln y_i(j) dj \tag{68}$$

## B.2 Optimal Factor Allocation

The production technology for an individual task $j$ is described by:

$$y_i(j) = \begin{cases} A_L l(j) + A_K k(j) & \text{if } j \in [0, I(D)] \\ A_L l(j) & \text{if } j \in (I(D), 1] \end{cases} \tag{69}$$

We assume the standard regularity condition that for automated tasks, the effective cost of capital is lower than labor ($\frac{r}{A_K} < \frac{W}{A_L}$). Under this condition, cost minimization implies a strict corner solution for task allocation:

- **Automated Tasks ($j \leq I(D)$):** Produced exclusively with capital, so $y_i(j) = A_K k(j)$.
- **Non-Automated Tasks ($j > I(D)$):** Produced exclusively with labor, so $y_i(j) = A_L l(j)$.

Given the unit elasticity of substitution between tasks (Equation 68), the firm allocates expenditure equally across all tasks. Let $K_i$ and $L_i$ denote the firm's total demand for capital and labor, respectively.

1. **Capital Allocation:** Total capital $K_i$ is spread symmetrically across the measure $I(D)$ of automated tasks. Thus, the amount of capital used in each automated task is:

$$k(j) = \frac{K_i}{I(D)} \quad \text{for } j \in [0, I(D)] \tag{70}$$

2. **Labor Allocation:** Total labor $L_i$ is spread symmetrically across the measure $1 - I(D)$ of non-automated tasks. Thus, the amount of labor used in each manual task is:

$$l(j) = \frac{L_i}{1 - I(D)} \quad \text{for } j \in (I(D), 1] \tag{71}$$

## B.3 Deriving the Aggregate Function

Substituting the task-level outputs back into the log-aggregate production function (68):

$$\ln Y_i = \int_0^{I(D)} \ln \left( A_K \frac{K_i}{I(D)} \right) dj + \int_{I(D)}^1 \ln \left( A_L \frac{L_i}{1 - I(D)} \right) dj \tag{72}$$

$$= I(D) \ln \left( A_K \frac{K_i}{I(D)} \right) + (1 - I(D)) \ln \left( A_L \frac{L_i}{1 - I(D)} \right) \tag{73}$$

Using the property $a \ln x = \ln(x^a)$, we rewrite this as:

$$\ln Y_i = \ln \left[ \left( A_K \frac{K_i}{I(D)} \right)^{I(D)} \right] + \ln \left[ \left( A_L \frac{L_i}{1 - I(D)} \right)^{1 - I(D)} \right] \tag{74}$$

Exponentiating both sides yields:

$$Y_i = \left(A_K^{I(D)} A_L^{1-I(D)}\right) \cdot \left(\frac{K_i}{I(D)}\right)^{I(D)} \left(\frac{L_i}{1-I(D)}\right)^{1-I(D)} \tag{75}$$

Defining the productivity term $B \equiv A_K^{I(D)} A_L^{1-I(D)}$, we obtain the aggregate production function presented in the main text:

$$Y_i = B \cdot \left(\frac{K_i}{I(D)}\right)^{I(D)} \left(\frac{L_i}{1-I(D)}\right)^{1-I(D)} \tag{76}$$

This derivation confirms that our functional form is consistent with the general equilibrium implications of task-based automation models, such as Corollary 1 in Acemoglu and Restrepo (2018), under the assumption of unitary substitution elasticity between tasks.

# C Micro-foundation of Endogenous Copyright Enforcement

In the main text, we assumed a reduced-form relationship between the degree of automation $I(D)$ and the probability of copyright enforcement $\beta(D) = [1 - I(D)]^{\xi}$. In this appendix, we provide a micro-foundation for this specification based on the legal doctrine of "Human Authorship" and judicial heterogeneity.

## C.1 The Legal Doctrine

Copyright law (e.g., the standard set by the U.S. Copyright Office) requires that a work contain a sufficient amount of human creative input to be registrable. Let $H \in [0, 1]$ denote the observable *human contribution share* in a creative product. In our task-based framework, this is defined as the measure of non-automated tasks:

$$H(D) = 1 - I(D) \tag{77}$$

If a work is entirely AI-generated ($I = 1$), then $H = 0$. If it is entirely manual ($I = 0$), then $H = 1$.

## C.2 Judicial Heterogeneity

While the human share $H$ is objective, the legal threshold for protection is subjective and varies across jurisdictions, judges, or patent/copyright examiners. We model this uncertainty by introducing a random variable $\varepsilon \in [0, 1]$, which represents the **judicial strictness threshold**.

- A lenient judge (low $\varepsilon$) grants copyright even for minimal human input.
- A strict judge (high $\varepsilon$) requires significant human input to grant protection.

A firm secures copyright if and only if its human contribution exceeds the judicial threshold drawn for that specific case:

$$\text{Copyright Granted} \iff H(D) \geq \varepsilon \tag{78}$$

## C.3 Derivation of the Probability Function

The probability of enforcement, $\beta$, is the probability that the firm's human contribution satisfies the random judicial requirement:

$$\beta(D) = \Pr(H(D) \geq \varepsilon) = G(H(D)) \tag{79}$$

where $G(\cdot)$ is the Cumulative Distribution Function (CDF) of the threshold $\varepsilon$.

To capture the bias of the legal system against AI-generated content, we assume $\varepsilon$ follows a **Power Function Distribution** (a special case of the Beta distribution) over the support $[0, 1]$, with shape parameter $\xi \geq 1$. The Probability Density Function (PDF) is $g(\varepsilon) = \xi\varepsilon^{\xi-1}$, and the CDF is:

$$G(\varepsilon) = \varepsilon^{\xi}, \quad \text{for } \varepsilon \in [0, 1] \tag{80}$$

Substituting $H(D) = 1 - I(D)$ into the CDF, we obtain the functional form used in the main text:

$$\beta(D) = G(1 - I(D)) = [1 - I(D)]^{\xi} \tag{81}$$

### C.4 Interpretation of the Strictness Parameter $\xi$

The parameter $\xi$ structurally captures the rigor of the legal regime:

1. **Neutral Regime ($\xi = 1$):** The threshold $\varepsilon$ is uniformly distributed ($U[0, 1]$). The probability of protection scales linearly with human input.

2. **Strict Regime ($\xi > 1$):** The distribution of judicial thresholds is skewed towards 1 (higher standards). The function $\beta(H) = H^{\xi}$ becomes convex. This implies that a small reliance on AI ($H < 1$) leads to a disproportionately large drop in the probability of protection. For example, if $\xi = 2$ and the firm uses 20% AI ($H = 0.8$), the probability of protection falls to $0.8^2 = 0.64$.

This convexity reflects the current "Copyright Trap" observed in policy debates, where even partial AI usage can jeopardize the legal status of the entire work.

## D Proofs

### Consistency of the Manual-Regime Viability Assumption (1)

This appendix does not prove existence of the manual-regime equilibrium from primitives alone. Its purpose is instead to derive the firm's operating-profit object and the associated free-entry condition in the limiting case of a purely manual corporate economy, where all active firms are exogenously restricted to $D_i = 0$. Assumption 1 then selects the parameter region in which these manual-regime equilibrium objects are well defined and support a symmetric free-entry equilibrium with strictly positive firm mass, $N^*_{man} > 0$.

First, evaluating the automation function at $D_i = 0$ gives

$$I(0) = \frac{0^{\mu}}{0^{\mu} + \delta} = 0, \tag{A.1}$$

since $\mu > 0$ and $\delta > 0$. Hence the aggregate production technology collapses to the purely labor-based case. Using the definition

$$B \equiv A_K^{I(D)} A_L^{1-I(D)}, \tag{A.2}$$

we obtain

$$B = A_K^0 A_L^1 = A_L. \tag{A.3}$$

Second, appropriability becomes complete in the manual regime. From

$$\beta_i = (1 - I(D))^{\xi_i}, \tag{A.4}$$

substituting $I(0) = 0$ yields

$$\beta_i = (1-0)^{\xi_i} = 1 \tag{A.5}$$

for every judicial draw $\xi_i \in \mathbb{R}_+$. Therefore,

$$\bar{\beta}(0, \bar{\xi}) = \mathbb{E}_\xi[1] = 1. \tag{A.6}$$

Third, we evaluate the firm's variable-cost objects in the limit as $I(D_i) \to 0$. The perceived marginal cost is

$$MC_{perceived}(D_i, N) = \frac{1}{B} \left( \frac{r}{I(D_i)} \right)^{I(D_i)} \left( \frac{\widetilde{w}(N)}{1 - I(D_i)} \right)^{1-I(D_i)}. \tag{A.7}$$

Since $B \to A_L$ as $I(D_i) \to 0$, and

$$\lim_{x \downarrow 0} \left( \frac{r}{x} \right)^x = \exp\left( \lim_{x \downarrow 0} x \ln \frac{r}{x} \right) = \exp(0) = 1, \tag{A.8}$$

while

$$\lim_{x \downarrow 0} \left( \frac{\widetilde{w}(N)}{1-x} \right)^{1-x} = \widetilde{w}(N), \tag{A.9}$$

it follows that

$$MC_{perceived}(0, N) = \lim_{I(D_i) \to 0} MC_{perceived}(D_i, N) = \frac{\widetilde{w}(N)}{A_L}. \tag{A.10}$$

Under symmetry, any candidate symmetric equilibrium with positive firm mass $N > 0$ satisfies $s_i = 1/N$, so the effective wage perceived by each firm is

$$\widetilde{w}(N) = w \left[ 1 + \frac{1}{N \epsilon_L} \right]. \tag{A.11}$$

Hence the pricing rule in such a symmetric manual-regime equilibrium becomes

$$p^* = \mathcal{M}(N) \frac{\widetilde{w}(N)}{A_L}. \tag{A.12}$$

Fourth, the actual average variable cost is

$$AVC(D_i, N) = \left[ \frac{w}{\widetilde{w}(N)} (1 - I(D_i)) + I(D_i) \right] MC_{perceived}(D_i, N). \tag{A.13}$$

Taking the limit as $I(D_i) \to 0$ and using the expression above for $MC_{perceived}(0, N)$ gives

$$AVC(0, N) = \left[ \frac{w}{\widetilde{w}(N)} \right] \frac{\widetilde{w}(N)}{A_L} = \frac{w}{A_L}. \tag{A.14}$$

Therefore, in the purely manual regime, the firm's expected appropriable operating profit reduces to

$$\mathbb{E}[\pi_{op}^{man}(q_i, \mathbf{q}_{-i})] = \bar{\beta}(0, \bar{\xi}) \left[ p(q_i, \mathbf{q}_{-i}) q_i - AVC(0, N) q_i \right] = p(q_i, \mathbf{q}_{-i}) q_i - \frac{w}{A_L} q_i. \tag{A.15}$$

Accordingly, the firm's static operating-profit maximization problem in the manual regime is

$$\max_{q_i} \left[ p(q_i, \mathbf{q}_{-i}) q_i - \frac{w}{A_L} q_i \right]. \tag{A.16}$$

Finally, substituting $D^* = 0$, $\bar{\beta}(0, \bar{\xi}) = 1$, and $AVC(0, N^*) = \frac{w^*_{man}}{A_L}$ into the free-entry condition

$$\bar{\beta}(I(D^*), \bar{\xi}) \left[p^* q^* - AVC(D^*, N^*) q^*\right] = \lambda P^*_H D^* + F_{base}, \tag{A.17}$$

yields the manual-regime free-entry condition

$$p^* q^* - \frac{w^*_{man}}{A_L} q^* = F_{base}. \tag{A.18}$$

Hence, in a purely manual corporate economy, each active firm's operating surplus is exactly revenue net of labor cost evaluated at unit cost $\frac{w^*_{man}}{A_L}$, and free entry requires this surplus to cover the baseline fixed cost $F_{base}$. Assumption 1 is therefore a maintained parameter restriction requiring that the manual-regime objects induced by the model be such that

$$\max_{q_i} \left[p(q_i, \mathbf{q}_{-i}) q_i - \frac{w^*_{man}}{A_L} q_i\right] \geq F_{base}, \tag{A.19}$$

with equality in the symmetric free-entry equilibrium, so that the exogenously manual economy admits a strictly positive firm mass $N^*_{man} > 0$ rather than collapsing to an empty corporate sector.

## Remark on Assumption 2

Assumption 2 is imposed as a maintained parameter restriction. It is not derived here as a theorem from the model primitives alone.

The reason is that a standalone non-emptiness proof would require additional structure that is not otherwise imposed in the baseline model. In particular, such a proof would need assumptions guaranteeing how the judicial-strictness distribution $F(\xi; \bar{\xi})$ behaves as $\bar{\xi}$ varies, as well as existence, selection, and continuity of the relevant equilibrium objects under AI adoption as primitive parameters change. Since these auxiliary assumptions are not part of the maintained baseline environment, we do not claim that Assumption 2 follows from primitives alone.

Accordingly, Assumption 2 should be read in the same spirit as Assumption 1: it restricts attention to a region of the parameter space in which, at $\lambda = 0$, there exists some AI adoption choice $D > 0$ for which expected operating profit strictly exceeds that of the symmetric manual benchmark. This restriction is introduced to rule out the trivial case in which AI is never privately attractive even when data royalties are absent.

Equivalently, using the expected operating-profit object,

$$\mathbb{E}[\pi^*_{op}(D_i)] \equiv \bar{\beta}(I(D_i), \bar{\xi}) \left[p^* q^* - AVC(D_i, N^*) q^*\right], \tag{A.20}$$

Assumption 2 requires that, when $\lambda = 0$,

$$\max_{D>0} \ \mathbb{E}[\pi^*_{op}(D)] \ > \ p^*_{man} q^*_{man} - \frac{w^*_{man}}{A_L} q^*_{man}. \tag{A.21}$$

In the symmetric manual free-entry benchmark, the verified manual-regime condition implies

$$p^*_{man} q^*_{man} - \frac{w^*_{man}}{A_L} q^*_{man} = F_{base}, \tag{A.22}$$

so the maintained restriction may also be written as

$$\max_{D>0} \mathbb{E}[\pi^*_{op}(D)] > F_{base}. \tag{A.23}$$

No stronger claim is made in this appendix.

## Proofs for Section 4.1

*Proof of Proposition 1.* Fix a selected symmetric free-entry equilibrium with $N^* > 0$. By the firm-level feasibility condition in the equilibrium definition, equilibrium data use must satisfy

$$0 \leq D^* \leq Y^*_H.$$

Hence exactly one of the following three mutually exclusive cases must obtain:

$$D^* = 0, \qquad 0 < D^* < Y^*_H, \qquad \text{or} \qquad D^* = Y^*_H.$$

We consider these cases in turn.

*Case 1:* $D^* = 0$. When equilibrium data use is zero, the model collapses to the manual regime. In particular, the manual-regime reductions stated in the text imply

$$I(0) = 0, \qquad \bar{\beta}(0, \bar{\xi}) = 1, \qquad B = A_L,$$

and therefore

$$MC_{perceived}(0, N) = \frac{\widetilde{w}(N)}{A_L}, \qquad AVC(0, N) = \frac{w}{A_L}.$$

Thus the equilibrium conditions reduce to the symmetric manual free-entry system. Under Assumption 1, that system admits a strictly positive manual equilibrium firm mass $\bar{N}_{man}$, so the selected equilibrium coincides with the manual equilibrium and

$$N^*(\lambda, \bar{\xi}) = \bar{N}_{man}.$$

*Case 2:* $0 < D^* < Y^*_H$. Because the equilibrium data choice lies strictly inside the feasible interval, the complementary-slackness condition from the equilibrium system,

$$\mu_D \left(Y^*_H - D^*\right) = 0,$$

implies

$$\mu_D = 0.$$

Moreover, since $D^* > 0$, the automation threshold satisfies

$$I(D^*) \in (0, 1),$$

and hence the expected appropriability term also satisfies

$$\bar{\beta}(I(D^*), \bar{\xi}) \in (0, 1).$$

Therefore the selected equilibrium is an interior AI-adoption regime with a slack firm-level data constraint.

*Case 3:* $D^* = Y_H^*$. In this case, the firm-level physical data constraint binds by definition. From the equilibrium system, the associated multiplier satisfies

$$\mu_D \geq 0$$

together with complementary slackness

$$\mu_D \left(Y_H^* - D^*\right) = 0,$$

which is automatically satisfied when $D^* = Y_H^*$. Hence the selected equilibrium is a data-starvation regime with a binding firm-level data constraint.

Since the three cases are mutually exclusive and exhaustive on the feasible interval $[0, Y_H^*]$, the proposition follows. □

*Proof of Remark 1.* The remark is an immediate implication of the feasible-set logic embedded in the equilibrium definition. Any candidate equilibrium data choice must satisfy

$$D^* \in [0, Y_H^*].$$

Accordingly, any characterization based on interior AI adoption is valid only when the selected equilibrium satisfies

$$0 < D^* < Y_H^*.$$

If the equilibrium forces determining desired data use would otherwise require a value exceeding $Y_H^*$, feasibility rules out such an interior allocation. The equilibrium must then be implemented at the boundary,

$$D^* = Y_H^*,$$

with an associated nonnegative multiplier $\mu_D \geq 0$ satisfying complementary slackness. This is exactly the sense in which the physical human-data boundary truncates candidate interior AI dynamics. □

## Proof of Section 4.2

*Proof of Proposition 2.* By hypothesis,

$$\varepsilon(N)^{-1} = \eta^{-1} + N^{-1}\big(\sigma^{-1} - \eta^{-1}\big).$$

Under the maintained ordering $\eta > \sigma > 1$, one has $\sigma^{-1} - \eta^{-1} > 0$. Differentiating $\varepsilon(N)^{-1}$ with respect to $N$ gives

$$\frac{d}{dN}\varepsilon(N)^{-1} = -\frac{\sigma^{-1} - \eta^{-1}}{N^2} < 0.$$

Since $\varepsilon(N) > 1$, it follows that $\varepsilon'(N) > 0$.

Now consider $\mathcal{M}(N) = \varepsilon(N)/(\varepsilon(N) - 1)$. Differentiating with respect to $\varepsilon$ yields

$$\frac{d\mathcal{M}}{d\varepsilon} = -\frac{1}{(\varepsilon - 1)^2} < 0.$$

Because $\varepsilon'(N) > 0$, the chain rule implies

$$\mathcal{M}'(N) = \frac{d\mathcal{M}}{d\varepsilon}\varepsilon'(N) < 0.$$

Next, by definition,

$$\Omega(N) = 1 + \frac{1}{N\epsilon_L}.$$

Holding $\epsilon_L$ fixed,

$$\Omega'(N) = -\frac{1}{N^2\epsilon_L} < 0.$$

Therefore both $\mathcal{M}(N)$ and $\Omega(N)$ are strictly decreasing in $N$. Hence, for any two selected symmetric equilibria with $0 < N_a < N_b$,

$$\mathcal{M}(N_a) > \mathcal{M}(N_b) \qquad \text{and} \qquad \Omega(N_a) > \Omega(N_b).$$

This establishes the claim. □

*Proof of Proposition 3.* Under symmetry, aggregate Corporate-Sector labor demand is $L_{Corp} = N^* l_i$, aggregate corporate output is $Q^*_{Corp} = N^* q^*$, and the corporate price index collapses to the common price $P^*_{Corp} = p^*$. Hence

$$LS_{AI} = \frac{w^* L_{Corp}}{P^*_{Corp} Q^*_{Corp}} = \frac{w^* N^* l_i}{p^* N^* q^*} = \frac{w^* l_i}{p^* q^*}.$$

Substituting the conditional labor-demand expression,

$$l_i = \frac{(1 - I(D^*))\, MC_{perceived}(D^*, N^*)\, q^*}{\widetilde{w}(N^*)},$$

gives

$$LS_{AI} = \frac{w^*}{p^* q^*} \cdot \frac{(1 - I(D^*))\, MC_{perceived}(D^*, N^*)\, q^*}{\widetilde{w}(N^*)} = (1 - I(D^*)) \frac{w^*}{\widetilde{w}(N^*)} \frac{MC_{perceived}(D^*, N^*)}{p^*}.$$

Using the pricing rule $p^* = \mathcal{M}(N^*) MC_{perceived}(D^*, N^*)$ and the definition $\Omega(N^*) = \widetilde{w}(N^*)/w^*$, one obtains

$$\frac{w^*}{\widetilde{w}(N^*)} = \frac{1}{\Omega(N^*)} \qquad \text{and} \qquad \frac{MC_{perceived}(D^*, N^*)}{p^*} = \frac{1}{\mathcal{M}(N^*)}.$$

Therefore

$$LS_{AI} = \frac{1 - I(D^*)}{\Omega(N^*)\, \mathcal{M}(N^*)},$$

as claimed. □

# Computational Appendix: General Equilibrium System

This appendix formalizes the non-linear algebraic system governing the symmetric free-entry general equilibrium, inclusive of the oligopsonistic labor market mechanism. The model is computationally resolved by finding the root of an exact $8 \times 8$ system of equations. By Walras' Law, we normalize the nominal aggregate revenue of the corporate AI sector as the numeraire ($R_{Corp} \equiv Npq = 1$). Consequently, all real macroeconomic aggregates are obtained by deflating nominal values with the ideal aggregate price index $P$.

## 1. Exogenous Structural Parameters

The parameters are calibrated to the baseline economy and remain fixed during the solver execution:

- **Preferences & Demand:** $\sigma = 2.5$ (cross-sector elasticity), $\eta = 15.459$ (within-sector variety elasticity), $\omega = 0.0852$ (human art preference weight).
- **Production & Technology:** $\alpha = 0.428$ (artist return to scale), $A_H = 1.0$ (baseline artist productivity), $\mu = 0.28$ (learning rate), $\delta = 3.0$ (task complexity scale).
- **Costs & Market Structure:** $\rho = 0.05$ (relative cost of AI compute to human labor), $F_{base} = 0.0243$ (fixed entry cost), $\epsilon_L = 2.0$ (macroeconomic labor supply elasticity).

*Note: During general equilibrium comparative statics, policy state variables are evaluated over bounded global domains, e.g., data royalty $\lambda \in [0.1, 1.0]$ and judicial strictness $\bar{\xi} \in [1.0, 3.0]$. Localized figures in the main text may display targeted sub-intervals to highlight extensive-margin discontinuities.*

## 2. Intermediate Endogenous Functions

To evaluate the residual equations cleanly, the solver constructs intermediate structural variables sequentially from the core state vector $\mathbf{X}$.

**Prices and Labor Market Dynamics:** Given the prevailing market wage $w$, AI firm price $p$, and human art price $P_H$, the aggregate price index $P$ and aggregate labor allocations are:

$$P_{Corp}(N, p) = N^{\frac{1}{1-\eta}} p \tag{A.24}$$

$$P(P_H, P_{Corp}) = \left[\omega P_H^{1-\sigma} + (1-\omega) P_{Corp}^{1-\sigma}\right]^{\frac{1}{1-\sigma}} \tag{A.25}$$

$$L_{total}(w, P) = (w/P)^{\epsilon_L} \tag{A.26}$$

$$L_H(w, P_H) = \left(\frac{\alpha A_H P_H}{w}\right)^{\frac{1}{1-\alpha}} \tag{A.27}$$

The residual labor supply available to the corporate sector determines the perceived residual elasticity ($\epsilon_L$) and the monopsonistic effective wage ($\widetilde{w}$) used for firm optimization:

$$\epsilon_L = \frac{\epsilon_L L_{total} + \frac{1}{1-\alpha} L_H}{L_{total} - L_H} \tag{A.28}$$

$$\widetilde{w}(N, w) = w\left(1 + \frac{1}{N \cdot \epsilon_L}\right) \tag{A.29}$$

**Cost Structure and Appropriability:** AI intensity $I(D) = \frac{D^\mu}{D^\mu + \delta}$ dictates the legal appropriability share $\bar{\beta} = (1-I)^{\bar{\xi}}$. Evaluated at the effective wage $\widetilde{w}$, the perceived marginal cost (for pricing) and the actual average variable cost (inclusive of retained monopsony rents, for zero-profit conditions) are:

$$MC_{perceived} = \left(\frac{\rho w}{I}\right)^I \left(\frac{\widetilde{w}}{1-I}\right)^{1-I} \tag{A.30}$$

$$AVC_{unit} = \left[\frac{w}{\widetilde{w}}(1-I) + I\right] MC_{perceived} \tag{A.31}$$

The analytical marginal benefit of data, $MB_D$, is derived by applying the chain rule to the expected operating profit $E[\pi_{op}] = \bar{\beta}(D)[p - AVC_{unit}(D)]q$, fully internalizing the endogenous shift in the monopsony rent retention factor.

## 3. The Exact System of Equations

The robust least-squares solver identifies the strictly positive vector of endogenous variables $\mathbf{X} = [p, N, D, q, P_H, Y_H, \mu_D, w]$ by driving the following $8 \times 8$ system of residuals to zero.

1. Pricing Rule: $$p - \frac{\varepsilon(N)}{\varepsilon(N) - 1} MC_{perceived} = 0 \tag{A.32}$$
2. Numeraire (Corp Revenue): $$Npq - 1 = 0 \tag{A.33}$$
3. Human Art Inverse Supply: $$P_H - \frac{w}{\alpha A_H} \left( \frac{Y_H}{A_H} \right)^{\frac{1-\alpha}{\alpha}} = 0 \tag{A.34}$$
4. Relative Demand: $$P_H Y_H - \frac{\omega}{1-\omega} \left( \frac{P_H}{P_{Corp}} \right)^{1-\sigma} = 0 \tag{A.35}$$
5. Zero Profit (Free Entry): $$\bar{\beta}(D) \left( pq - AVC_{unit} \cdot q \right) - \lambda P_H D - F_{base} = 0 \tag{A.36}$$
6. Data Choice (KKT Condition): $$MB_D - \lambda P_H - \mu_D = 0 \tag{A.37}$$
7. Physical Constraint Slackness: $$\mu_D (Y_H - D) = 0 \tag{A.38}$$
8. Macroeconomic Labor Clearing: $$L_H + Nq(1 - I) \frac{MC_{perceived}}{\widetilde{w}} - L_{total} = 0 \tag{A.39}$$

*Note: Equation 4 exploits the numeraire identity ($P_{Corp} Q_{Corp} = Npq = 1$) to elegantly simplify the aggregate expenditure multiplier. Equation 8 enforces physical labor market clearing via Shephard's Lemma evaluated at the monopsonistic effective wage.*