One Prompt, One Paper: Optimal Journal Policy and the AI Submission Flood

Alejandro Lopez-Lira* Sina Seyfi†

University of Florida Aalto University

First version: May 2025

This version: March 16, 2026

Abstract

How should academic journals adapt as AI reshapes research? We develop a general equilibrium model where researchers use AI to accelerate production and a social planner sets capacity and submission fees. Without policy adjustment, submissions flood the system and acceptance rates collapse. With optimal fees, welfare rises monotonically with AI productivity regardless of whether referees also use AI. Fees decline as AI improves, even when review costs are fixed, because each dollar of fee screens more aggressively as production costs fall. The optimal fee is below the per-submission review cost, so its primary role is congestion screening and prestige stabilization rather than direct cost internalization, a general equilibrium mechanism absent from standard Pigouvian analysis. Fee screening exceeds the review cost channel by a factor of 7 to 45, and optimal capacity is near-flat across all scenarios.

JEL Codes: D82, L15, O33

Keywords: Academic publishing, artificial intelligence, submission fees, journal capacity, congestion pricing

*alejandro.lopez-lira@warrington.ufl.edu

†sina.seyfi@aalto.fi

Introduction

The proliferation of advanced Artificial Intelligence (AI), particularly Large Language Models (LLMs), is reshaping academic research. AI tools assist at every stage, from literature reviews to manuscript drafting and data analysis. Generative AI increases worker performance (Brynjolfsson, Li, and Raymond 2025), and scholars who adopt it raise their productivity (Kusumegi, Yang, Ginsparg, de Vaan, et al. 2025b). These advances also pressure top-tier journals to adapt their publication capacity.

AI presents a dual impact. On one hand, AI streamlines editorial and peer-review processes, lowering the real cost of publishing and giving journals a direct incentive to expand capacity. On the other hand, it reduces research production costs for scholars, threatening to increase submissions to a degree that could overwhelm even a more efficient system. The interaction between these forces determines optimal journal policy, including both capacity and submission fees, through general equilibrium channels that go beyond standard externality-correction logic.

Three scenarios illustrate how policy and technology interact. First, if journals do not adjust fees as AI improves, submissions flood the system (rising 4.3-fold when fees are removed at current AI levels, and exceeding 10,000 at high AI productivity), acceptance rates drop below 1%, and the quality cutoff rises so far above the support that selection depends on noise rather than merit. The congestion externality operates through participation composition, review costs, and prestige feedback: at $\gamma = 1$, the per-submission review cost $r(1) \approx 8.1$ exceeds the optimal fee $\tau^* = 3.0$, so the fee does not fully internalize review costs and instead serves primarily as a congestion screening and prestige stabilization instrument. Second, if journals set optimal fees but referees cannot use AI, fees decline with AI productivity but remain higher at each γ than when referees have AI tools ($\tau^* = 1.8$ versus 1.1 at $\gamma = 15$), because the fee becomes a more potent screening instrument as production costs fall; welfare increases monotonically. Third, if both fees and AI-assisted refereeing are available, review costs fall, optimal fees decline faster, and welfare rises further. A decomposition that shuts

down AI-assisted refereeing ($\delta = 0$) confirms this: fee screening exceeds the review cost channel by a factor of 7 to 45 at every γ.

Scholars choose research effort and AI adoption intensity, facing heterogeneous costs for AI use. Research output depends on human effort and AI use. A publication's value to a researcher is convex in the journal's average published quality ($\varphi > 1$), so prestige feeds back into submission incentives with super-linear gain. This captures both the incentive to publish in higher-quality outlets and the instability that arises when this feedback is unpriced. A social planner chooses publication capacity to maximize social welfare.

Social welfare equals the societal benefits from published research, net of total publication and reviewing costs. It counts both the quantity and quality of published papers; benefits depend on total published quality (capacity times average quality), and review costs fall as AI improves. Optimal journal capacity hinges on how these model elements interact: while improving AI lowers production costs for scholars and review costs for journals, the optimal response is shaped by how much authors prefer prestigious journals, how society values research quality, and how capacity expansion affects equilibrium quality standards. Optimal capacity has no closed form; we compute it numerically. Welfare counts only the incremental production cost of top-tier participation relative to what scholars would produce at a lower-tier outlet. The sign of this increment is endogenous: negative with fees (top-tier participants concentrate into fewer papers) and positive without them (marginal scholars flood in). The planner's welfare function excludes the utility researchers derive from prestige, to avoid double-counting private submission incentives.

Three empirical facts from Kusumegi, Yang, Ginsparg, de Vaan, et al. (2025b) motivate the model: (1) adoption diffuses quickly and output concentrates among adopters (48% of authors account for 72% of output by June 2024); (2) difference-in-differences and event-study estimates around adoption yield a calibrated productivity jump of 0.18 papers per month (rising thereafter); and (3) LLM-assisted papers are less likely to clear a screening proxy (publication status), with a larger penalty for more complex papers. These facts

ground three ingredients of the model: heterogeneous adoption, the magnitude of the AI productivity shift, and the observation that quantity rises faster than evaluative signals.

In the calibrated model, optimal capacity is near-flat across all three scenarios ($C^* = 100$ for $\gamma \leq 2$, rising to 105 at $\gamma = 15$), so the entire policy adjustment occurs through fees. Optimal fees decline from \$320 to \$110 as AI advances; even when referees cannot use AI, fees still decline (from \$320 to \$180) because each dollar of fee screens more aggressively as production costs fall. Without fees at current AI levels ($\gamma = 1$), submissions rise from 1,230 to 5,342, and the vast majority of high-quality work is rejected.

Fee adjustment is essential for AI to be welfare-improving only if production costs for rejected papers represent social waste. For prestige journals, where scholars compete for a scarce signal rather than dissemination, this interpretation is appropriate: effort on rejected papers constitutes rent-seeking in a tournament, analogous to wasteful competition in other status markets. In this setting, the fee's role is closer to a mechanism-design instrument that stabilizes a convex prestige feedback loop than to a textbook Pigouvian tax on a single externality margin. The model is static and abstracts from dynamic effects such as changes in research quality or norms over time.

1 Literature Review

This paper draws on three strands of literature: the economics of academic publishing, the impact of artificial intelligence on the research process, and empirical work on AI-driven productivity gains and the operational realities of top-tier journals.

1.1 Economics of Academic Publishing

A large literature treats journals as capacity-constrained institutions that ration scarce editorial and referee attention. Economics publishing has slowed and standards have risen over time, with longer papers and more polishing and revision (Ellison 2002b, 2002a). The

profession's incentives amplify these bottlenecks: top-journal placements dominate career rewards (Heckman and Moktan 2020) and shape where authors submit and how editors allocate attention (Card and DellaVigna 2013). Status competition and tournament models capture this pattern: small differences in publication outcomes generate large payoff differences (Rosen 1981; Lazear and Rosen 1981).

A complementary strand studies the "market design" of publishing (fees, turnaround times, and access/pricing). Submission fees and editorial speed interact with authors' submission decisions and journals' congestion (Cotton 2013; Leslie 2005). Cotton (2013) analyzes fees in a partial equilibrium setting where prestige is exogenous and the fee corrects a standard congestion externality; here, prestige feedback is endogenous ($\varphi > 1$), the outside option makes the sign of incremental production costs endogenous, and a per-paper fee operates on both intensive and extensive margins. Broader work analyzes journal pricing and open-access incentives (Jeon and Rochet 2010; McCabe and Snyder 2005; Bergstrom 2001). Work on review design and screening highlights how evaluation rules can change selection and perceived fairness (Blank 1991) and how adding gatekeepers affects selection outcomes (Koren 2024). Recent syntheses emphasize that slow review can be privately rational yet socially costly, motivating policy levers that operate through capacity, processing costs, and screening (Hadavand, Hamermesh, and Wilson 2024). Journal capacity, fees, and acceptance standards are all endogenous here.

1.2 AI and Research Productivity

A fast-growing empirical literature finds that generative AI can raise productivity in knowledge work, often with heterogeneous effects across workers and tasks. Kusumegi, Yang, Ginsparg, de Vaan, et al. (2025b) document that researchers who adopt LLMs for manuscript drafting then increase their paper output relative to comparable non-adopters. Novy-Marx and Velikov (2026) demonstrate an automated process for generating complete academic finance papers. In a controlled setting, access to ChatGPT increases speed and improves

output quality for writing tasks (Noy and Zhang 2023). In a large field setting, providing a generative-AI assistant boosts performance, with gains concentrated among less experienced workers (Brynjolfsson, Li, and Raymond 2025). Evidence from creative work similarly highlights that how humans collaborate with LLMs matters for output (Chen and Chan 2024), and complementary work studies the conditions under which human–AI teaming helps or harms performance (Wang, Gao, and Agarwal 2024).

A parallel literature studies distributional effects: who benefits, and who is displaced. Recent evidence documents that exposure to generative AI reshapes task demand and can differentially affect lower- versus higher-skill workers (Demirci, Hannane, and Zhu 2025), consistent with task-based views where automation and augmentation coexist. Related evidence on AI decision support shows productivity gains accruing disproportionately to lower-skill workers in some settings (Kanazawa et al. 2025). These heterogeneous effects matter for academic publishing because they translate into changes in who can produce papers faster, and how much "polish" versus "substance" AI can cheaply supply.

At the macro level, task-based models discipline claims about AI's aggregate effects and emphasize the importance of complementarity, reallocation, and the pace of adoption (Acemoglu 2024). In the same spirit, recent work studies how AI changes the organization of knowledge production and incentives in the "knowledge economy" (Ide and Talamàs 2025), connecting to classic theories of knowledge organization and specialization (Garicano 2000; Garicano and Rossi-Hansberg 2004). Methodologically, economists' growing ability to extract structure from unstructured data, including text, also lowers the cost of producing publishable research inputs (Dell 2025).

Finally, early macro-evidence suggests that LLMs may already be shifting the volume (and potentially the style) of scientific output, with implications for screening, review, and credibility (Kusumegi, Yang, Ginsparg, de Vaan, et al. 2025b). Within finance specifically, work argues that AI tools can accelerate parts of the research pipeline and alter what "doing finance research" looks like (Novy-Marx and Velikov 2026). As AI reduces the marginal

cost of producing submissions on dimensions tied to writing, presentation, and exploratory analysis, journals may optimally respond by adjusting capacity, standards, and fees, with welfare consequences that depend on which policy instruments are available.

1.3 Broader Context and Related Work

A growing working-paper literature applies LLMs directly to prediction, interpretation, and corporate/market outcomes (Lopez-Lira and Tang 2023; Ko and Lee 2023; J. Chen et al. 2023; S. Chen et al. 2024; Jha et al. 2025; Degen et al. 2024; Engelberg et al. 2025), and to related empirical building blocks such as peer-firm identification (Cao et al. 2025). Beyond output, scholars analyze how AI changes the research process and the institutions around it, including AI-assisted peer review and editorial decision-making (Checco et al. 2021), the broader use of AI across research workflows (Chubb, Cowling, and Reed 2022; Butson and and Spronken-Smith 2024; Khalifa and Albadawy 2024), and concerns about reproducibility and verification capacity (Fišar et al. 2024). Finally, publishing institutions can influence inclusion as well as impact, suggesting that journal responses to AI may have composition effects within the profession (Bansak et al. 2024). Journal policy is an equilibrium institutional response to AI-driven shifts in submission volume, review costs, and the quality–polish tradeoff. The present model differs from existing work on journal fees and capacity in two ways. First, fees are embedded in a general equilibrium framework where convex prestige creates self-reinforcing entry dynamics. Second, the planner's two instruments (capacity and per-paper fees) interact through the equilibrium quality cutoff and participation margin, so the optimal fee reflects congestion and prestige-feedback externalities, not only the direct review cost.

2 Economic Environment

2.1 Agents, Quality, and Inputs

A continuum of scholars populates the economy, each defined by the pair (q, k), representing intrinsic research quality q and private cost of AI adoption k. Types are drawn from a compact support, $q \in [q_L, q_H]$ and $k \in [k_L, k_H]$, with $0 < q_L < q_H$ and $0 < k_L < k_H$. The population distribution is described by a continuous density $f(q, k)$ with $\iint f(q, k)\, dq\, dk = 1$ (unit mass of scholars). Quality and AI cost are independently distributed: $f(q, k) = f_q(q) f_k(k)$, with marginal densities continuous and bounded above and below by positive constants. Each scholar chooses human effort $e \geq 0$ and AI adoption intensity $\lambda \geq 0$.

2.2 Parameter Restrictions

The following parameter restrictions ensure well-defined optimization and a clear economic interpretation. The production function has diminishing returns with exponent $\theta = 1/2$. Utility is concave in quality and acceptances, with $\alpha = 1/2$ governing the elasticity with respect to paper quality and $\beta = 1/2$ governing the elasticity with respect to expected acceptances. With $\alpha = \beta = 1/2$, the benefit term is $V\sqrt{npq}$, the geometric mean of expected acceptances and paper quality (Section 3.2). Effort has marginal cost parameter $c > 0$, AI productivity is $\gamma > 0$, the baseline entry cost is $F_0^{\text{base}} \geq 0$, and the per-submission fee is $\tau \geq 0$. Journal prestige is scaled by $V_0 > 0$ and has elasticity $\varphi > 0$ with respect to average published quality. The outside option (lower-tier journals) is characterized by a single composite $\Omega_{\text{low}} \equiv V_{\text{low}}\sqrt{p_{\text{low}}}$, where $V_{\text{low}} > 0$ is lower-tier prestige and $p_{\text{low}} \in (0, 1)$ is lower-tier acceptance probability, both exogenous. With $\beta = 1/2$ and $\tau_{\text{low}} = 0$, this composite is the only combination of lower-tier parameters that enters the model (Section 3.2.3). Noise enters through two sufficient statistics: the total noise scale $\sigma_T = \sqrt{\sigma_u^2 + \sigma_\varepsilon^2}$, where $\sigma_u > 0$ is within-paper quality variation and $\sigma_\varepsilon > 0$ is evaluation noise, and the signal share $s \equiv \sigma_u^2/\sigma_T^2 \in (0, 1)$, which determines the magnitude of the selection premium. In the social

welfare objective, $\eta > 0$ scales the social value of published quality and $\xi \in (0, 1)$ captures diminishing returns to total quality. Review costs use baseline level $r_0 > 0$ and AI elasticity $\delta \geq 0$, while capacity costs include a weakly linear component $\kappa \geq 0$ and a strictly convex component $\psi > 0$.

2.3 Production and Evaluation

A scholar chooses inputs to produce a number of papers, n:

$$n \;=\; \big(e + \lambda\gamma\big)^{\theta}, \qquad \theta = \tfrac{1}{2}. \tag{1}$$

Each paper's quality has two components. The scholar's baseline quality q is an exogenous type that reflects persistent ability (training, domain expertise, methodological skill). In addition, each paper receives an independent idea draw $u \sim N(0, \sigma_u^2)$ that captures within-scholar, paper-specific quality variation: even a scholar of fixed ability produces some papers with better ideas than others. The true quality of a paper is therefore:

$$q_{\text{true}} = q + u. \tag{2}$$

The scholar produces n papers, all sharing the same baseline quality q but each with an independent draw of u. The production function exhibits diminishing returns to inputs ($\theta = 1/2$): doubling effort and AI usage less than doubles the number of papers produced. Higher AI productivity γ allows scholars to produce more papers for a given level of effort.

The editor observes a noisy signal of each paper's true quality. The observed quality is:

$$q_{\text{obs}} = q_{\text{true}} + \varepsilon = q + u + \varepsilon, \tag{3}$$

where $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ is evaluation noise, independent of u, representing the stochasticity inherent in peer review (referee taste, editorial judgment, measurement error). Because

both u and ε are normal and independent, the total noise in observed quality relative to baseline is $u + \varepsilon \sim N(0, \sigma_T^2)$ with $\sigma_T = \sqrt{\sigma_u^2 + \sigma_\varepsilon^2}$. Each of a scholar's n papers is evaluated independently, with independent draws of both u and ε for each paper.

3 Journal Policy and Scholar's Problem

A joint equilibrium between scholars' optimization and the journal's market clearing defines the model.

3.1 Journal Policy and Acceptance Probability

The journal sets a fixed capacity, C, representing the total number of papers it can publish. Each paper, regardless of its quality q, consumes one unit of capacity. The editor observes the realized quality q_{obs} of all submitted papers and ranks them. Papers are accepted from the top of the ranking downwards until the total number of accepted papers fills the capacity C. This process defines an endogenous quality cutoff, q_{cut}, which is the observed quality of the marginal accepted paper.

A paper with baseline quality q is accepted if its observed quality exceeds the cutoff, $q_{\text{obs}} = q + u + \varepsilon \geq q_{\text{cut}}$. Because $u + \varepsilon \sim N(0, \sigma_T^2)$, the acceptance probability per paper is:[1]

$$p(q, q_{\text{cut}}) = \Pr(q + u + \varepsilon \geq q_{\text{cut}}) = \Phi\left(\frac{q - q_{\text{cut}}}{\sigma_T}\right), \tag{4}$$

where $\Phi(\cdot)$ is the standard normal CDF and $\sigma_T = \sqrt{\sigma_u^2 + \sigma_\varepsilon^2}$ is the total noise scale. The acceptance probability is a probit function of the quality gap $(q - q_{\text{cut}})$, with σ_T governing the steepness: smaller σ_T makes acceptance more sharply determined by quality, while larger σ_T increases the role of noise. The probit specification follows from the normal distributions of u and ε.

1. Because there is a continuum of agents, the law of large numbers implies that this individual probability is also the exact fraction of papers with baseline quality q that are accepted. There is no aggregate uncertainty in the model.

Letting $z = (q - q_{\text{cut}})/\sigma_T$, the derivatives are:

$$\frac{\partial p}{\partial q} = \frac{\phi(z)}{\sigma_T}, \qquad \frac{\partial p}{\partial q_{\text{cut}}} = -\frac{\phi(z)}{\sigma_T}, \tag{5}$$

where $\phi(\cdot)$ is the standard normal PDF. Henceforth we suppress the noise parameters in the acceptance probability notation, writing $p(q, q_{\text{cut}})$ instead of $p(q, q_{\text{cut}}; \sigma_T)$, as they are held fixed throughout the analysis.

3.2 Scholar's Optimization Problem

A scholar's private benefit from publication depends on the journal's overall prestige, a function of the average quality of accepted papers, $\bar{q}_{\text{pub}}$, not the cutoff. Prestige depends on average published quality rather than the cutoff, so scholars compete on expected quality rather than on beating a threshold. The prestige function is given by $V(\bar{q}_{\text{pub}}) = V_0 \bar{q}_{\text{pub}}^{\varphi}$ with $V_0 > 0$ fixed.

Each scholar of type (q, k) chooses inputs (e, λ) to produce n papers that maximize their expected utility. The scholar takes the prestige part of the reward, $V(\cdot)$, as given, determined by the market-wide average quality of published papers, not the cutoff they are trying to beat.

$$U(e, \lambda; q, k, q_{\text{cut}}, \bar{q}_{\text{pub}}) \;=\; \sqrt{n \cdot p(q, q_{\text{cut}})} \cdot V(\bar{q}_{\text{pub}}) \cdot \sqrt{q} \;-\; \frac{c}{2} e^2 \;-\; k\lambda - F_0^{\text{base}} - \tau \cdot n, \tag{6}$$

where $n = (e + \lambda\gamma)^{\theta}$. With $\alpha = \beta = 1/2$, the benefit term has a geometric-mean interpretation: $\sqrt{np} \cdot V \cdot \sqrt{q} = V \cdot \sqrt{npq}$, so the scholar's reward is proportional to the geometric mean of expected acceptances and paper quality, scaled by prestige. The scholar produces n papers, each with baseline quality q. Each paper has acceptance probability $p(q, q_{\text{cut}}) = \Phi((q - q_{\text{cut}})/\sigma_T)$, which averages over both the idea draw u and evaluation noise ε. The expected number of accepted papers is $n \cdot p(q, q_{\text{cut}})$. The square-root elasticities capture diminishing returns: doubling expected acceptances or doubling quality each raises

the benefit by a factor of $\sqrt{2}$, not 2.

The cost structure has two components beyond production costs. The baseline entry cost $F_0^{\text{base}} \geq 0$ represents structural participation costs (time to format submissions, opportunity cost of academic engagement) that cannot be changed by policy. The per-submission fee $\tau \geq 0$ is charged for each paper submitted, so a scholar producing n papers pays total fees $\tau \cdot n$. This per-paper structure matches journal practice, where fees are assessed per submission rather than as a lump-sum entry charge. The fee enters the scholar's first-order condition, raising the effective marginal cost of each additional paper. This creates an intensive margin response (fewer papers per scholar) in addition to the extensive margin (fewer participating scholars). Together, F_0^{base} and τ create a participation margin even when $p(q, q_{\text{cut}}) > 0$ for all scholars.

Because each scholar's n has negligible effect on market-wide $\bar{q}_{\text{pub}}$, the agent's problem is well-behaved: scholars react to the cutoff q_{cut} via the probability of acceptance, but not via the value of the prize itself.

The solution to this problem yields an optimal number of papers $n^*(q, k, q_{\text{cut}}, \bar{q}_{\text{pub}})$ for each type.

3.2.1 Cost Minimization

The scholar's two-dimensional choice of (e, λ) separates into two steps. First, for any given target number of papers n, the scholar chooses inputs to minimize the private cost, $\mathcal{C}(e, \lambda; k) = \frac{c}{2}e^2 + k\lambda$, subject to the production constraint $(e + \gamma\lambda)^\theta = n$ and non-negativity constraints $e \geq 0, \lambda \geq 0$.

With the constraint $e + \gamma\lambda = n^{1/\theta}$, the cost can be rewritten as:

$$\min_{0 \leq e \leq n^{1/\theta}} \frac{c}{2}e^2 + k\lambda = \min_{0 \leq e \leq n^{1/\theta}} \frac{c}{2}e^2 + \frac{k}{\gamma}\left(n^{1/\theta} - e\right). \tag{7}$$

The first-order condition gives $e^* = k/(c\gamma)$, which defines the AI-adoption threshold:

$$n^{\text{AI}}(k) = \left(\frac{k}{c\gamma}\right)^{\theta}. \tag{8}$$

The resulting minimum cost $\mathcal{C}(n;k)$ is C^1, strictly convex in n, and independent of quality q. The piecewise cost function, its derivatives, and continuity verification are in Appendix A.1.

3.2.2 Utility Maximization and the Submission Decision

The scholar's problem simplifies to choosing a single variable, n, to maximize:

$$U(n;q,k,q_{\text{cut}},\bar{q}_{\text{pub}}) = \sqrt{n \cdot p(q,q_{\text{cut}})} \cdot V(\bar{q}_{\text{pub}}) \cdot \sqrt{q} - \mathcal{C}(n;k) - F_0^{\text{base}} - \tau \cdot n. \tag{9}$$

The candidate for an interior solution, $n^* > 0$, is characterized by the first-order condition where marginal benefit equals marginal cost:

$$\underbrace{\frac{1}{2}\sqrt{\frac{p(q,q_{\text{cut}}) \cdot q}{n}} \cdot V(\bar{q}_{\text{pub}})}_{\text{Marginal Benefit of Producing One More Paper}} = \underbrace{\frac{d\mathcal{C}(n;k)}{dn} + \tau}_{\text{Marginal Production Cost + Fee}}. \tag{10}$$

The left-hand side is the marginal benefit of producing one more paper. Because marginal benefit is proportional to $n^{-1/2}$, it is decreasing in n. The right-hand side is the effective marginal cost: the production cost $\mathcal{C}'(n;k)$, which is increasing in n due to convexity, plus the per-submission fee τ. The fee shifts the marginal cost curve upward by a constant, reducing optimal output for every scholar type. Together, these conditions ensure a unique interior maximum n^* when it exists. The scholar takes the acceptance probability $p(q,q_{\text{cut}})$, prestige $V(\bar{q}_{\text{pub}})$, and quality q as given; none varies with the choice of n.

3.2.3 The Outside Option and Journal Tiers

The model's participation decision represents an allocation choice, not a production decision. Scholars produce papers regardless of any single journal, as research is their job. The question

is which tier of journal to target with their output. A lower-tier journal provides an explicit outside option.

A scholar of type (q, k) who targets the lower tier solves an optimization problem with the same structure as the top-tier problem but with exogenous parameters replacing the endogenous equilibrium objects. The lower-tier journal has prestige V_{low}, acceptance probability p_{low}, and per-submission fee τ_{low}, all exogenous. The scholar's lower-tier utility as a function of paper production n is:

$$U_{\text{low}}(n; q, k) = \sqrt{n \cdot p_{\text{low}}} \cdot V_{\text{low}} \cdot \sqrt{q} - \mathcal{C}(n; k) - F_0^{\text{base}} - \tau_{\text{low}} \cdot n. \tag{11}$$

The optimized outside option is:

$$U^*_{\text{low}}(q, k) = \max_{n \geq 0} U_{\text{low}}(n; q, k). \tag{12}$$

With $\beta = 1/2$ and $\tau_{\text{low}} = 0$, the lower-tier benefit $V_{\text{low}}\sqrt{n \cdot p_{\text{low}}} \cdot \sqrt{q}$ depends on $(V_{\text{low}}, p_{\text{low}})$ only through $\Omega_{\text{low}} \equiv V_{\text{low}}\sqrt{p_{\text{low}}}$: the benefit is $\Omega_{\text{low}}\sqrt{nq}$. The composite Ω_{low} is therefore the single sufficient statistic for the outside option, and the individual components V_{low} and p_{low} are not separately identified (closed forms in Appendix A.2).

The parameters V_{low}, p_{low}, and τ_{low} are exogenous, representing the lower-tier publishing landscape. $V_{\text{low}} < V_0$ (prestige hierarchy) and $\tau_{\text{low}} = 0$ (lower-tier journals rarely charge submission fees). Because these parameters do not depend on the focal journal's equilibrium, $U^*_{\text{low}}(q, k)$ is a function of (q, k) alone.

The scholar submits to the focal (top-tier) journal if and only if the optimized top-tier utility exceeds the optimized lower-tier utility:

$$U^*(q, k; q_{\text{cut}}, \bar{q}_{\text{pub}}, \tau) \geq U^*_{\text{low}}(q, k), \tag{13}$$

where $U^*(q, k; q_{\text{cut}}, \bar{q}_{\text{pub}}, \tau) = U(n^*; q, k, q_{\text{cut}}, \bar{q}_{\text{pub}})$ is the maximized top-tier utility. This

defines the participation set:

$$\mathcal{P}(q_{\text{cut}}, \bar{q}_{\text{pub}}, \tau) = \{(q,k) \mid U^*(q,k; q_{\text{cut}}, \bar{q}_{\text{pub}}, \tau) \geq U^*_{\text{low}}(q,k)\}. \tag{14}$$

The per-submission fee screens through two channels: it excludes marginal types from participating (extensive margin) and reduces the number of papers produced by each active scholar (intensive margin, via $dn^*/d\tau < 0$; see Appendix A.6). The per-paper structure gives the planner a two-margin instrument that lump-sum entry fees lack. The outside option creates an extensive margin for fees: without it, all scholars would participate regardless of fee level, and fees could only reduce per-scholar output.

4 Market-Clearing Equilibrium

An equilibrium is a pair $(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})$ that clears the market for journal slots and satisfies quality consistency.

The first condition, the capacity constraint, requires that total accepted papers fill the journal's capacity C. Let $n^{**}(q,k) = n^*(q,k)$ for participating types (those in $\mathcal{P}$) and $n^{**} = 0$ otherwise; n^{**} depends on the equilibrium cutoff and average quality.[2]

$$C = \iint_{\mathcal{P}(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})} n^{**}(q,k,q^*_{\text{cut}}, \bar{q}^*_{\text{pub}}) \cdot p\left(q, q^*_{\text{cut}}\right) \cdot f(q,k)\, dq\, dk. \tag{15}$$

The journal's capacity C equals the total expected number of papers accepted: each type (q,k) produces n^{**} papers, and each paper is accepted with probability $p(q, q^*_{\text{cut}})$.

2. With the probit participation smoothing used in the calibration (Section 8), n^{**} is replaced by a continuous participation weight: $n^{**} = n^* \cdot \Phi((U^*_{\text{top}} - U^*_{\text{low}})/\zeta)$, where $\zeta = 0.1$ is a small preference-shock scale.

4.1 Selection Premium and Quality Consistency

The second condition requires that average published quality is consistent with the selection process. Because within-paper variation u is present, papers accepted from a pool with baseline quality q are not a random sample: those with favorable idea draws ($u > 0$) are more likely to clear the cutoff. This selection produces a premium in expected true quality conditional on acceptance.

The expected true quality of an accepted paper with baseline quality q is:

$$E[q_{\text{true}} \mid q_{\text{obs}} \geq q_{\text{cut}}, q] = q + \frac{\sigma_u^2}{\sigma_T} \cdot h\left(\frac{q_{\text{cut}} - q}{\sigma_T}\right), \tag{16}$$

where $h(z) = \phi(z)/[1 - \Phi(z)]$ is the inverse Mills ratio. The selection premium is:

$$\Delta(q, q_{\text{cut}}) = \frac{\sigma_u^2}{\sigma_T} \cdot h\left(\frac{q_{\text{cut}} - q}{\sigma_T}\right) = s \cdot \sigma_T \cdot h\left(\frac{q_{\text{cut}} - q}{\sigma_T}\right), \tag{17}$$

where $s \equiv \sigma_u^2/\sigma_T^2$ is the signal share: the fraction of total evaluation variance attributable to within-paper quality variation rather than review noise. The two noise parameters $(\sigma_u, \sigma_\varepsilon)$ enter the model only through the sufficient statistics (σ_T, s), where σ_T governs the steepness of acceptance probabilities and s governs the selection premium magnitude. The derivation and properties of this premium are in Appendix A.3.

The quality consistency condition incorporates the selection premium:

$$\bar{q}_{\text{pub}}^* = \frac{1}{C} \iint_{\mathcal{P}} n^{**} \cdot p\left(q, q_{\text{cut}}^*\right) \cdot \left[q + \Delta(q, q_{\text{cut}}^*)\right] \cdot f(q, k)\, dq\, dk. \tag{18}$$

The integrand weights each accepted paper by its expected true quality $q + \Delta(q, q_{\text{cut}})$, not its baseline quality q. The numerator is the total expected true quality of all accepted papers, and the denominator is the total number of accepted papers. Together these conditions define a system of two equations in two unknowns, $(q_{\text{cut}}^*, \bar{q}_{\text{pub}}^*)$.

The equilibrium determines other quantities of interest. The total number of submitted

papers is

$$N_{\text{submit}} = \iint_{\mathcal{P}(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})} n^{**}(q, k, q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})\, f(q, k)\, dq\, dk.$$

The average acceptance rate is then:

$$\bar{p}^* = \frac{C}{N_{\text{submit}}} = \frac{C}{\iint_{\mathcal{P}(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})} n^{**}(q, k, q^*_{\text{cut}}, \bar{q}^*_{\text{pub}}) f(q, k)\, dq\, dk}. \tag{19}$$

5 Analytical Properties and Equilibrium Characteristics

Strict concavity and continuity of supply ensure equilibrium existence and uniqueness.

5.1 The Scholar's Optimal Choice

The first-order condition (Equation 10) identifies a unique interior maximum for n^{**}. The marginal benefit is $\frac{1}{2}\sqrt{pq/n} \cdot V(\bar{q}_{\text{pub}})$, which is strictly decreasing in n (proportional to $n^{-1/2}$). The effective marginal cost $\mathcal{C}'(n; k) + \tau$ is strictly increasing in n (since $\mathcal{C}$ is convex) and bounded below by τ. Since marginal benefit is decreasing from $+\infty$ and effective marginal cost is increasing, there exists at most one interior solution n^* satisfying the first-order condition, and this solution is a global maximum when it exists.

The baseline entry cost F_0^{base}, per-submission fee τ, and the continuous acceptance probability $p(q, q_{\text{cut}}) > 0$ together determine whether the scholar participates (choosing $n^* > 0$) or sits out (choosing $n^{**} = 0$). The submission fee τ is a policy instrument set by the journal (Section 10). Throughout, parameters and distributions are such that the participation set $\mathcal{P}(q_{\text{cut}}, \bar{q}_{\text{pub}}, \tau)$ is non-empty and varies smoothly with equilibrium variables.[3]

The comparative statics have standard signs: $dn^*/d\gamma > 0$ (AI boosts output), $dn^*/dq > 0$

3. The second-order condition is automatically satisfied at any interior solution n^* because marginal benefit is decreasing and effective marginal cost is increasing. Specifically: $\frac{d^2U}{dn^2} = -\frac{1}{4}(pq)^{1/2} n^{-3/2} \cdot V - \mathcal{C}''(n; k) < 0$ since both terms are negative.

(higher quality raises output), $dn^*/dk < 0$ (higher costs reduce output), $dn^*/d\tau < 0$ (fees reduce output on both margins). The full IFT derivations are in Appendix A.6. The welfare analysis below uses the fee result $dn^*/d\tau < 0$.

5.2 Equilibrium Existence and Uniqueness

The equilibrium is a fixed point of the market-clearing system (Equations 15 and 18).

Proposition 1 (Equilibrium Existence)**.** *An equilibrium exists at* (γ, C, τ) *if and only if* $C \leq C^{\max}(\gamma, \tau)$.

When $C > C^{\max}(\gamma, \tau)$, even with zero selectivity ($q_{\text{cut}} \to -\infty$), total willing submissions fall short of capacity. The journal cannot fill its slots, so no equilibrium exists. Since $C^{\max}$ is decreasing in τ and increasing in γ, equilibrium fails at high fees with low AI productivity, but succeeds when advanced AI generates sufficient surplus. The formal derivation of $C^{\max}$ and the IVT existence argument are in Appendix A.4.

The prestige feedback loop governs uniqueness: higher selectivity raises average quality, which increases prestige, which stimulates submissions. Multiple equilibria are possible in principle, but uniqueness holds when the direct effect of selectivity on acceptance probabilities dominates the indirect prestige feedback. Appendix D provides sufficient conditions based on monotonicity of the reduced supply function along the quality-consistency locus. At the calibrated parameters ($\varphi = 1.25$, $\beta = 1/2$, $\theta = 1/2$), the direct acceptance effect dominates the prestige feedback, and the equilibrium is unique. The selection premium $\Delta(q, q_{\text{cut}})$ in the quality consistency condition is a mechanical consequence of within-paper variation and truncation; it does not create an additional feedback loop, because Δ depends only on $(q, q_{\text{cut}}, \sigma_u, \sigma_T)$ and not on $\bar{q}_{\text{pub}}$. The Jacobian structure verifies uniqueness analytically; numerically, solving from multiple initial conditions confirms convergence to the same solution.

With convex prestige ($\varphi > 1$), the feedback loop is amplified: a small increase in selectivity raises average quality, which raises prestige more than proportionally, which stim-

ulates enough additional submissions to push selectivity even higher. Without fees, this self-reinforcing cycle can make the equilibrium fragile or generate multiplicity. Fees dampen the loop by weakening the submission response to prestige: each additional paper costs τ in fees, so the marginal scholar's incentive to enter when prestige rises is attenuated. In the uniqueness condition (Appendix D), this appears as the effective feedback elasticity declining with τ, making the uniqueness condition easier to satisfy. The calibrated fee $\tau^* \approx 3$ is large enough to ensure robust uniqueness across the full range of γ. This stabilization role distinguishes the fee from a textbook Pigouvian instrument. A standard corrective tax internalizes a fixed marginal damage. Removing it leads to overproduction but does not destabilize the equilibrium itself. Here, $\varphi > 1$ creates a feedback multiplier, and the fee breaks the loop by raising the cost of the marginal submission that would trigger the cascade. The optimal fee level therefore reflects both the direct review cost externality and the strength of the prestige feedback.

The optimized lower-tier outside option $U^*_{\text{low}}(q, k)$ has three properties relevant to equilibrium structure. The participation set $\mathcal{P}$ now depends on k (not just q), since $U^*_{\text{low}}(q, k)$ depends on k through the cost function $\mathcal{C}(n^*_{\text{low}}; k)$. The outside option helps uniqueness: any equilibrium must have prestige $V(\bar{q}^*_{\text{pub}}) > V_{\text{low}}$ (otherwise no scholar would prefer the focal journal), which implies $\bar{q}^*_{\text{pub}} > (V_{\text{low}}/V_0)^{1/\varphi}$, eliminating low-quality equilibria. The participation set $\mathcal{P}$ shrinks monotonically in τ: by the envelope theorem, $dU^*/d\tau = -n^* < 0$ while U^*_{low} is independent of τ, so raising the fee strictly reduces the set of types that prefer the focal journal.

Section 9 presents the numerical results.

6 Social Planner's Problem

The social planner chooses journal capacity C and submission fee τ to maximize social welfare. The journal, as a market actor, must rely on the noisy signal of paper quality, q_{obs},

to select papers. The social planner's objective, however, is based on expected true quality $q_{\text{true}} = q + u$. Because accepted papers are selected on favorable u draws (the selection premium), the social value of accepted research exceeds baseline quality.

The core of the planner's problem is to evaluate the total expected true quality of published research that results from a given policy (C, τ). For each (C, τ), the equilibrium determines a pair $(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})$. To simplify notation, we define the equilibrium-dependent optimal choice:

$$n^{**}(q, k \mid C, \tau) \equiv n^{**}\big(q, k, q^*_{\text{cut}}(C, \tau), \bar{q}^*_{\text{pub}}(C, \tau)\big). \tag{20}$$

The total published quality, $Q_{\text{pub}}(C, \tau)$, incorporates the selection premium:

$$Q_{\text{pub}}(C, \tau) = \iint_{\mathcal{P}(C,\tau)} n^{**}(q, k \mid C, \tau) \cdot p\big(q, q^*_{\text{cut}}\big) \cdot \big[q + \Delta(q, q^*_{\text{cut}})\big] \cdot f(q, k)\, dq\, dk = C \cdot \bar{q}_{\text{pub}}, \tag{21}$$

where $\mathcal{P}(C, \tau) \equiv \mathcal{P}(q^*_{\text{cut}}(C, \tau), \bar{q}^*_{\text{pub}}(C, \tau), \tau)$ is the participation set and $\Delta(q, q_{\text{cut}}) = s \cdot \sigma_T \cdot h((q_{\text{cut}} - q)/\sigma_T)$ is the selection premium from Equation 17. Each type (q, k) contributes n^{**} papers, each accepted with probability p, each with expected true quality $q + \Delta$.

The social welfare function evaluates the focal (top-tier) journal in isolation:[4]

$$SW(C, \tau;\, \gamma) = \frac{\eta}{\xi}[Q_{\text{pub}}(C, \tau)]^{\xi} - \text{IncProdCosts}(C, \tau) - r(\gamma) \cdot N_{\text{submit}}(C, \tau) - \kappa \cdot C - \frac{\psi}{2} C^2, \tag{22}$$

4. The welfare function measures the social value created by the focal journal, not total knowledge across tiers. A system-wide specification could aggregate $Q_{\text{total}} = Q_{\text{top}} + Q_{\text{low}}$ across both tiers, which is computable from the same single-journal equilibrium (since Q_{low} depends only on the exogenous lower-tier parameters and the participation set $\mathcal{P}$). However, with $\xi < 1$ the lower-tier baseline must be included inside $[\cdot]^{\xi}$, coupling the welfare comparison to assumptions about lower-tier quality that are outside our model's scope. A second equilibrium would arise only if the lower tier had endogenous acceptance or prestige. Our specification asks: how much social value does this journal create, net of the extra costs its existence imposes? We also exclude researchers' private prestige benefits, $V(\bar{q}_{\text{pub}}) \cdot \sqrt{q} \cdot \sqrt{np}$, which are largely positional (zero-sum across scholars). Including them would risk double-counting and could recommend over-expansion if prestige rents are large.

where:

$$\text{IncProdCosts}(C,\tau) = \iint_{\mathcal{P}(C,\tau)} \left[\mathcal{C}\big(n^{**}(q,k \mid C,\tau);k\big) - \mathcal{C}\big(n^{*}_{\text{low}}(q,k);k\big)\right] f(q,k)\,dq\,dk \tag{23}$$

$$N_{\text{submit}}(C,\tau) = \iint_{\mathcal{P}(C,\tau)} n^{**}(q,k \mid C,\tau)\cdot f(q,k)\,dq\,dk \tag{24}$$

$$N_{\text{accept}}(C,\tau) = C \tag{25}$$

The planner's problem is:

$$\big(C^{*}(\gamma),\,\tau^{*}(\gamma)\big) = \arg\max_{C,\,\tau\geq 0}\; SW(C,\tau;\,\gamma). \tag{26}$$

The welfare function counts only incremental production costs: the change in effort each participating scholar incurs by targeting the top tier relative to what that scholar would exert at the lower tier. Scholars who do not participate in the focal journal (those outside $\mathcal{P}$) produce n^{*}_{low} papers for the lower tier regardless of the focal journal's policy, so their production cost is not attributable to the focal journal and does not enter the planner's objective. For participants, the cost $\mathcal{C}(n^{*}_{\text{low}};k)$ is the baseline production that would occur in the absence of the top-tier option, and only the increment $\mathcal{C}(n^{**};k) - \mathcal{C}(n^{*}_{\text{low}};k)$ is the social cost of top-tier participation. This increment can be positive or negative: it is positive when top-tier targeting induces more output ($n^{**} > n^{*}_{\text{low}}$), and negative when a scholar produces fewer but more targeted papers for the prestigious journal ($n^{**} < n^{*}_{\text{low}}$), which can occur when high prestige $V(\bar{q}_{\text{pub}})$ compensates for low acceptance probability $p(q, q_{\text{cut}})$. Non-participants incur zero incremental cost by construction, so fees cannot generate spurious welfare gains by excluding scholars from the focal journal.

The social value term $(\eta/\xi)[Q_{\text{pub}}]^{\xi}$ captures the benefit of published research, with $\xi < 1$ reflecting diminishing returns as the volume of knowledge grows, for example due to overlap or finite attention. Submission fees $\tau \cdot n$ per scholar are transfers: fee revenue $\tau \cdot N_{\text{submit}}$ is collected by the journal and returned to the research community (or used for editorial

operations already captured in capacity costs), so fees affect welfare only indirectly through their effect on scholar behavior (both participation and per-scholar output). Review costs are $r(\gamma) \cdot N_{\text{submit}}$, where $r(\gamma) = r_0/(1+\gamma)^\delta$ falls with AI assistance and $\delta > 0$ captures the elasticity. Capacity costs are $\kappa \cdot C + (\psi/2)C^2$, with $\psi > 0$ capturing editorial fatigue or administrative diseconomies of scale.

For each candidate (C, τ) in a search grid, the planner must first solve for the equilibrium $(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})$ and then evaluate the resulting social welfare. The optimal policy $(C^*(\gamma), \tau^*(\gamma))$ maximizes this function jointly over both instruments.

6.1 Economic Intuition and Trade-offs

Journal capacity C plays two opposing roles.

On one hand, as AI technology (γ) improves, all scholars can produce more papers for the same cost. If the journal holds its capacity C fixed, the market-clearing quality cutoff, q_{cut}, must rise. Scholars can now produce more papers (higher $n = (e+\lambda\gamma)^{1/2}$) for the same effort, filling the journal's slots more quickly. The resulting competition creates a more intense "rat race" among scholars. Expanding capacity can mitigate this, allowing the journal to capture more of the AI-generated research output.

On the other hand, expanding capacity is not costless. A larger C generally leads to a lower q_{cut}, which in turn lowers the average quality of published papers $\bar{q}_{\text{pub}}$, and thus the journal's prestige $V(\bar{q}_{\text{pub}})$. This decline in prestige weakens scholars' incentive to participate, potentially lowering both the quantity and quality of submissions. More submissions induced by capacity expansion impose real costs: scholars consume more effort and AI resources, and the journal must review more papers.

The social planner's optimal choice, (C^*, τ^*), therefore balances these forces. It weighs the marginal benefit of publishing more quality-weighted research (via Q_{pub}, which is subject to diminishing returns via $\xi < 1$) against the combined costs: scholar costs (effort and AI), review costs ($r(\gamma) \cdot N_{\text{submit}}$), and capacity costs ($\kappa \cdot C + \frac{\psi}{2}C^2$). The relationship between

(C^*, τ^*) and γ depends on how these competing effects interact through general equilibrium channels.

6.1.1 Why Capacity Is Near-Flat Across AI Regimes

The marginal benefit of capacity, $\eta\xi Q_{\text{pub}}^{\xi-1} \cdot dQ_{\text{pub}}/dC$, depends on the average quality of the marginal paper admitted. As C expands, the cutoff q_{cut} falls and the marginal paper's quality declines, so marginal benefit is decreasing in C. The marginal cost (ψC) is increasing in C but does not depend on γ. As AI raises γ, more scholars produce more papers, but the equilibrium cutoff adjusts so that average quality at the margin is roughly unchanged: the additional papers come from scholars whose AI-augmented production pushes them into the participant pool, but the quality composition of the marginal cohort remains similar. The marginal-benefit and marginal-cost curves therefore cross at nearly the same C^* regardless of γ. In the calibration (Section 9), C^* rises from 100 to 105 as γ increases 150-fold, a 5% increase.

6.1.2 The Fee Trade-off with Selection Premium

The planner sets per-submission fees τ to internalize the review cost externality, but in equilibrium the fee also stabilizes the prestige feedback loop and determines the composition of the participant pool. Because fees are per paper, they operate on both margins: higher fees screen out marginal types (extensive margin) and reduce output per active scholar (intensive margin, Equation A25). A higher fee has two opposing effects on total published quality Q_{pub}:

1. Cost reduction: Fewer submissions reduce incremental production costs and review costs. The papers deterred tend to be marginal, so quality composition improves.

2. Selection premium erosion: Fewer submissions lead to a lower equilibrium cutoff q_{cut}. A lower cutoff reduces the selection premium $\Delta(q, q_{\text{cut}})$, because with less selective

screening the journal can no longer cherry-pick papers with favorable idea draws as effectively. This lowers average published quality.

These forces balance at an interior optimum $\tau^* > 0$. When $\sigma_u = 0$ (no within-paper variation), the second force vanishes and the optimal fee is determined solely by the cost-reduction trade-off. When $\sigma_u > 0$, the selection premium channel dampens the fee, because aggressive screening through high fees comes at the cost of losing the quality bonus from selection. This mechanism, arising from within-paper quality variation, provides a structural reason for interior optimal fees even beyond the standard cost-reduction logic.

6.1.3 Why Fees Decline with AI

The optimal fee has two components: the direct review cost $r(\gamma) = r_0/(1+\gamma)^\delta$, which falls mechanically with γ, and a congestion/prestige-stabilization component that reflects the fee's role in screening marginal submissions and dampening the prestige feedback loop. Even when review costs are held fixed ($\delta = 0$), the congestion component itself declines: from $\tau^* = 3.2$ at $\gamma = 0.5$ to $\tau^* = 1.8$ at $\gamma = 15$ (Section 9). In textbook Pigouvian logic, a corrective tax should track marginal damage; here, the optimal fee moves in the opposite direction.

The mechanism is that the fee becomes a more potent instrument as production costs fall. In the AI regime, the scholar's effective marginal cost is $\mathcal{C}'(n;k) + \tau$, where $\mathcal{C}'(n;k) = (k/\gamma\theta)n^{1/\theta-1}$ is decreasing in γ. As γ rises, τ constitutes a larger share of effective marginal cost, so each dollar of fee deters more output: $|dn^*/d\tau|$ increases. The planner achieves the same screening intensity with a lower fee. Both the review cost component (from $r(0.1) \approx 9.7$ to $r(15) \approx 4.4$) and the congestion component decline, so τ^* falls monotonically from ≈ 3.2 at $\gamma = 0.1$ to ≈ 1.1 at $\gamma = 15$.

6.1.4 Incremental Production Costs

The welfare function counts only the incremental production cost of top-tier participation: the change in effort a scholar incurs when targeting the focal journal relative to the lower tier. Production for the lower tier happens regardless of the focal journal's policy; the social planner accounts only for the difference in cost. This increment can be positive (when top-tier targeting induces higher output) or negative (when high prestige allows the scholar to produce fewer, more targeted papers). The latter occurs when high prestige $V(\bar{q}_{\text{pub}})$ compensates for low acceptance probability $p(q, q_{\text{cut}})$. Because the incremental cost $\mathcal{C}(n^{**}; k) - \mathcal{C}(n^*_{\text{low}}; k)$ is zero for non-participants by construction, fees cannot generate spurious welfare gains by excluding scholars whose production costs would be incurred at the lower tier anyway. The endogenous sign of this increment distinguishes the model from standard congestion externality frameworks, where production costs are unconditionally positive: here, the fee's welfare contribution operates through preventing the increment from turning positive by keeping out marginal scholars whose top-tier output would exceed their lower-tier output.

At the planner's optimum, the marginal quality loss from reduced submissions equals the marginal cost savings from lower review and incremental production costs. Joint optimization over (C^*, τ^*) gives the planner a larger policy space than optimizing either instrument alone. Section 9 presents the full numerical analysis.

7 Empirical Motivation: Evidence on AI Adoption, Output, and Screening

The evidence in this section motivates three model ingredients: heterogeneous adoption, the magnitude of the AI productivity shift, and the screening challenge. Author-level evidence on the production shock comes from Kusumegi, Yang, Ginsparg, de Vaan, et al. (2025b), using their replication package (Kusumegi, Yang, Ginsparg, de Vaan, et al. 2025a) for arXiv preprints from January 2022 through June 2024. We then complement this with population-

level content detection, submission volume trends across venues, the divergence between working-paper and journal submissions, and evidence on AI in peer review. AI adoption is measured as detected LLM-assisted writing, and the adoption month for an author is the first month in which any LLM-assisted paper appears. Sections 7.1–7.5 summarize the replication results (full regression tables are in Appendix F); later sections draw on additional sources.

The empirical evidence plays three roles. First, it documents a concentrated output surge that creates submission pressure. Second, LLM-assisted papers are less likely to clear a screening proxy: quantity rises faster than evaluative signals. Third, the divergence between no-fee and fee-charging venues provides direct evidence for the fee-screening mechanism. These patterns motivate the model's ingredients but do not directly identify its structural parameters; the calibration in Section 8 sets those independently.

7.1 Data and measurement

The replication package makes available arXiv-based panels, which are the only components we can access. The production file is an author-month dataset with 3,298,950 observations that track monthly output around adoption. For each author, the data track the number of preprints posted per month over a 30-month window, from January 2022 through June 2024. The evaluation file is a paper-level dataset with 179,077 papers labeled as LLM-assisted or non-LLM; the LLM share in this panel is 24.4% (75.6% non-LLM, 24.4% LLM). We do not have access to the other panels described in the original paper, so all reported patterns are based on these arXiv files. Adoption is defined as the first month in which an author produces an LLM-flagged paper, using the replication definitions without modification.

7.2 Adoption diffusion and output concentration

Adoption rises quickly once LLM-assisted writing appears in the data. By the end of the observed window, roughly 48% of authors have adopted, yet these adopters produce about 72% of monthly output (Figure A10). This concentration drives the submission surge in the

model: a minority of low-cost adopters can generate a large share of submissions even when the median author has not adopted.

7.3 Productivity effects around adoption

Difference-in-differences and event-study estimates show a persistent post-adoption increase in monthly productivity, reaching 0.4 to 0.45 additional papers per month by the end of the observed window (Figure A11). The DiD interaction coefficient is 0.184 papers per month and is stable across specifications with and without month fixed effects (Table A1). The post-adoption effects exceed the pre-trend by a factor of roughly four.

7.4 Heterogeneity by baseline productivity

Estimating the DiD regression separately by baseline productivity bin reveals that AI-driven output gains concentrate among already-productive authors (Table A2). The interaction coefficient increases monotonically from 0.122 in the lowest bin to 0.413 in the highest, a 3.4-fold gap. High-output scholars generate the bulk of submission pressure; the gains are concentrated, not broad-based.

7.5 Quality proxies and screening outcomes

Linear probability models with publication status as the screening proxy show that LLM-assisted papers are 6 percentage points less likely to be published, and the penalty grows with paper complexity (Table A3). AI increases quantity without increasing the probability of passing editorial screening, consistent with Kusumegi, Yang, Ginsparg, de Vaan, et al. (2025b).

7.6 Population-level AI content detection

Population-level content detection corroborates the author-level patterns from Kusumegi, Yang, Ginsparg, de Vaan, et al. (2025b).

Liang et al. (2025) analyze 1,121,912 preprints and published papers from January 2020 through September 2024 on arXiv, bioRxiv, and Nature portfolio journals. Using a word-frequency framework, they estimate that by September 2024, 22.5% of computer science abstracts, 8.9% of Nature portfolio abstracts, and 7.7% of mathematics abstracts showed evidence of LLM modification. The overall lower bound across fields is at least 10% of 2024 abstracts. LLM modification rates are higher in more crowded research areas and for papers from regions with lower English proficiency.

Kobak et al. (2025) examine over 15 million biomedical abstracts from PubMed (2010–2024) and identify 379 style words with elevated frequencies in 2024. They estimate that at least 13.5% of 2024 PubMed abstracts were processed with LLMs, translating to approximately 200,000 papers per year in biomedical literature alone.

The production shock appears at population level across fields, not only among early adopters.

7.7 Researcher surveys on AI adoption

Survey evidence confirms that AI adoption in academic writing is broad-based and accelerating. A Nature survey of more than 5,000 academics worldwide found that 57% had used AI for writing in the previous two years, and 72% planned to use it in the next two years (Kwon 2025). A survey of scientists at US universities found that 65% had used generative AI, with writing, reviewing, and editing as the most frequent applications at 40%; 84% of current users intended to continue (Arroyo-Machado et al. 2025). The adoption wave on arXiv (Kusumegi, Yang, Ginsparg, de Vaan, et al. 2025b) is representative of a broader shift in research practices.

7.8 Submission volume trends

The production shock translates into divergent submission trends depending on whether the venue charges submission fees.

At no-fee venues, submission growth has accelerated sharply. ArXiv overall set new monthly records in late 2024, with 24,226 submissions in October 2024, a 17% year-over-year increase (arXiv 2024). ArXiv economics grew 77% cumulatively from 2022 to 2025 (1,308 primary articles in 2022 to 2,313 in 2025), with growth accelerating from 10.8% in 2023 to 21.5% in 2024 and 31.3% in 2025 (arXiv 2026). At computer science conferences, which generally do not charge submission fees, the increases are larger still: NeurIPS submissions rose from 10,411 (2022) to 21,575 (2025), a 107% increase (NeurIPS Program Chairs 2025); AAAI reported a 77% increase from 2025 to 2026 (AAAI 2025); and CHI rose 27% from 2023 to 2024 (CHI 2024 Program Chairs 2023).

At fee-charging journals, submission trends are flat or declining over the same period. The Journal of Finance received 1,204 new submissions in 2020 and 937 in 2024, a decline of 22% (American Finance Association 2026). Econometrica fell 17% from 2020–21 to 2021–22 (Econometric Society 2026). The American Economic Review received approximately 1,722 submissions in 2023, roughly flat (American Economic Association 2026). The Review of Economics and Statistics declined approximately 10% in 2022–23 before a 13% rebound in 2024 (Review of Economics and Statistics 2026). The Review of Finance reported that "volumes remain low" in 2023–24 (Review of Finance 2026).

General scientific publishing has also been under strain independently of AI. Hanson et al. (2024) document that the number of Scopus and Web of Science articles in 2022 was approximately 47% higher than in 2016, a growth rate that outpaced the growth in practicing scientists. The post-ChatGPT acceleration in arXiv and conference submissions sits on top of this secular trend.

7.9 Working-paper and journal divergence

The contrast between surging working-paper output and flat journal submissions provides a natural experiment for the fee-screening mechanism in the model. Several complementary hypotheses account for this divergence.

The most direct channel is submission fees. Fees at top economics and finance journals range from $125 (Econometrica, Review of Economics and Statistics) to $750 (Journal of Financial Economics); the median is approximately $300 (see Section 8). The Economic Journal attributed a decline in submissions to raising its fee from 150 to 220 GBP (Lippi 2023). Computer science conferences, which charge no meaningful submission fees, have experienced the largest surges. Submission fees screen out marginal AI-assisted submissions, consistent with this pattern.

A related channel is self-selection: if LLM-assisted papers are less likely to pass peer review (Kusumegi, Yang, Ginsparg, de Vaan, et al. 2025b), authors may recognize this and direct such papers to preprint servers rather than fee-charging journals. Submission lags provide a second channel: the arXiv economics acceleration appeared in 2024–2025, and with typical submission-to-journal lags of one to two years, the corresponding surge in journal submissions may appear in 2025–2026. The 13% rebound at the Review of Economics and Statistics in 2024 is consistent with this timing.

Finally, a composition effect may contribute. AI may disproportionately enable marginal researchers to produce working papers that would not be targeted at top-five journals. If new AI-assisted output is concentrated in lower-quality papers, the top journals would not see a proportionate increase in submissions even as total research output grows.

7.10 AI in peer review

AI adoption extends to the evaluation side of the market. Liang et al. (2024) estimate that between 6.5% and 16.9% of peer review text submitted to AI conferences after the release of ChatGPT was substantially modified by LLMs. At ICLR 2024, approximately

15.8% of reviews showed evidence of AI assistance; by ICLR 2026, this share had risen to an estimated 21% (Naddaf 2025). AI-modified reviews showed lower self-reported confidence, later submission relative to the deadline, and less engagement with author rebuttals. These patterns suggest that AI is changing not only the supply of papers but also the quality and informativeness of the review signal, compounding the screening challenge that the model addresses.

7.11 Summary for the model

The evidence provides five inputs to the theoretical framework. First, adoption diffuses quickly and concentrates output among adopters, with population-level estimates confirming that at least 10–22% of recent papers show LLM modification. Second, adoption produces large, persistent productivity gains, especially for already productive authors. Third, quality screening outcomes are weaker for LLM-assisted papers. Fourth, the divergence between no-fee venues (surging submissions) and fee-charging journals (flat or declining submissions) provides direct evidence for the fee-screening mechanism. Fifth, AI adoption on the review side reduces peer-evaluation informativeness, compounding the screening challenge.

8 Calibration

The social planner's welfare function accounts for both the quantity and quality of published research:

$$SW(C, \tau; \gamma) = \frac{\eta}{\xi}[Q_{\text{pub}}(C, \tau)]^{\xi} - \text{IncProdCosts}(C, \tau) - r(\gamma) \cdot N_{\text{submit}}(C, \tau) - \kappa \cdot C - \frac{\psi}{2}C^2$$

where $Q_{\text{pub}} = C \cdot \bar{q}_{\text{pub}}$ is total published quality, IncProdCosts is the incremental production cost of top-tier participation (the extra effort each participating scholar exerts beyond what the scholar would exert at the lower tier), N_{submit} is total submissions, $r(\gamma) = r_0/(1+\gamma)^{\delta}$ is the per-submission review cost with baseline r_0, and κ, ψ are publication cost parameters.

Only incremental costs enter the welfare function: production for the lower tier happens regardless of the focal journal's policy; the planner accounts only for the additional cost of top-tier effort.

The first-order condition for optimal capacity has two components in the marginal benefit: a direct effect (higher capacity publishes more research at current quality) and a quality feedback effect (capacity expansion lowers the equilibrium cutoff and thus average published quality, with $\frac{\partial \bar{q}_{\text{pub}}}{\partial C} < 0$). Because $\bar{q}_{\text{pub}}$ is an endogenous equilibrium outcome that depends on capacity through the general equilibrium system, there is no closed-form solution for $C^*(\gamma)$ or its elasticity $\varepsilon_{C,\gamma}$; we compute both numerically.

The structural parameters are δ (governing how AI affects review costs via $r(\gamma) = r_0/(1+\gamma)^{\delta}$), φ (prestige elasticity in how scholars value journal quality), and ξ (social benefit elasticity, with $\xi < 1$ implying diminishing marginal social value of published research). They jointly determine the optimal capacity response to AI through equilibrium scholar behavior, quality standards, and welfare trade-offs.

8.1 Data and parameter calibration

We obtain journal data from Scimago Journal Rank (SJR; SCImago Lab 2026), using the large set of journals available in that source. We restrict the data to journal-type records and years 2022–2024 and use them both to estimate prestige elasticity and to discipline the capacity target.

Prestige elasticity φ enters the prestige function $V(\bar{q}_{\text{pub}}) = V_0 \bar{q}_{\text{pub}}^{\varphi}$. Comparing two tiers, $\varphi = \ln(V_{\text{top}}/V_{\text{sec}})/\ln(\bar{q}_{\text{top}}/\bar{q}_{\text{sec}})$. In the SJR data, prestige V is proxied by the SJR index and average published quality $\bar{q}$ by the 2-year citations-per-document field. The sample includes only journals with strictly positive SJR and positive citations per document. The top tier is defined as journals with SJR at or above the 95th percentile and the second tier as those between the 80th and 95th percentiles; for each tier we take median SJR and median citations per document. The calculation yields a point estimate $\widehat{\varphi} \approx 1.23$. We set $\varphi = 1.25$ as a round

central value, implying that a 10% increase in a journal's average published quality raises prestige by $(1.10)^{1.25} - 1 \approx 12.7\%$.

The parameter δ quantifies how AI capability γ affects the per-submission cost of refereeing via $r(\gamma) = r_0/(1+\gamma)^{\delta}$. Micro-evidence shows that AI can reduce time in specific editorial subtasks: a mixed-methods study of AI-assisted reviewer selection reports a 73% reduction in the time to select reviewers (Farber 2024), and active-learning screening tools such as ASReview achieve work saved over sampling of 67–92% at 95% recall in systematic review screening (van de Schoot et al. 2021). The bulk of peer-review resource cost, however, reflects referee reading, verification, and judgment rather than matching or triage. Surveys indicate a median of roughly five hours to complete a review and global estimates imply over 100 million reviewer-hours annually (Vesper 2018). The aggregate benchmark: at $\gamma = 1$ (current AI), a 20% proportional cost reduction gives $r(1)/r_0 = 1 - 0.20 = 2^{-\delta}$, giving $\delta = -\ln(0.8)/\ln(2) \approx 0.32$. We set the baseline $\delta = 0.30$.

The parameter ξ appears in the social benefit function $B(Q_{\text{pub}}) = (\eta/\xi)Q_{\text{pub}}^{\xi}$. With $\xi < 1$, the marginal social value of an additional unit of published quality decreases as more research is already available. We set $\xi = 0.8$, so doubling total published quality raises social welfare by a factor of $2^{0.8} \approx 1.74$ rather than 2.

Given these structural parameters, we set scale and distribution parameters to match two empirical targets at current AI ($\gamma = 1$): optimal capacity $C^* \approx 100$ and an equilibrium acceptance rate of approximately 8%. The capacity target $C^* \approx 100$ is within the range observed in SJR data (SCImago Lab 2026): the median annual document count among journals at or above the 95th percentile SJR is 131. The acceptance rate target reflects observed rates at top journals: the Journal of Finance reports approximately 6.4% (American Finance Association 2026), the AER 5–7% (American Economic Association 2026), and the RFS 6–8% (see also Card and DellaVigna 2013).

Scholar quality is drawn from a Beta(2,2) distribution on $[0.5, \bar{q}]$ with $\bar{q} = 5$, AI adoption cost from a Beta(2,2) distribution on $[0.1, 2.0]$, and scholar mass is $M = 10{,}000$. All aggregate

quantities (submissions N_{submit}, production costs, review costs, welfare) are scaled by M: where the model's unit-mass integrals give per-capita values, the numerical implementation multiplies by M to produce level quantities. The two type dimensions are independent. The upper bound $\bar{q}$ is chosen so that the equilibrium acceptance rate at $(C, \tau) \approx (100, 3)$ matches the empirical target of approximately 8%.

The outside option is parameterized by the composite $\Omega_{\text{low}} \equiv V_{\text{low}}\sqrt{p_{\text{low}}}$, which is the single sufficient statistic for the lower-tier alternative (Section 3.2.3). We set $\Omega_{\text{low}} = 0.50$ (corresponding to $V_{\text{low}} = 1.0$ and $p_{\text{low}} = 0.25$), calibrated so that the participation margin produces an acceptance rate near the empirical target at the optimal (C^*, τ^*).

As described in Section 3.2.3, the noise parameters $(\sigma_u, \sigma_\varepsilon)$ enter the model only through two sufficient statistics: the total noise scale $\sigma_T = \sqrt{\sigma_u^2 + \sigma_\varepsilon^2}$ and the signal share $s \equiv \sigma_u^2/\sigma_T^2$. The total noise σ_T governs the steepness of the probit acceptance probability $p(q, q_{\text{cut}}) = \Phi((q - q_{\text{cut}})/\sigma_T)$, while the signal share s controls the magnitude of the selection premium $\Delta = s \cdot \sigma_T \cdot h((q_{\text{cut}} - q)/\sigma_T)$.

We fix the signal share at $s = 0.12$, implying that 12% of total evaluation variance is attributable to within-paper quality variation (real differences across a scholar's papers) and 88% to review noise (referee taste, editorial judgment, measurement error). This value is consistent with empirical studies of within-author heterogeneity in citation outcomes (Sinatra et al. 2016), which suggest that the standard deviation of per-paper quality within a researcher's portfolio is on the order of 30–50% of the cross-author standard deviation. The total noise scale $\sigma_T \approx 2.90$ is calibrated jointly with the other scale parameters. The underlying decomposition is $\sigma_u = 1.0$ and $\sigma_\varepsilon = 2.72$, where the evaluation noise was matched to a logistic scale parameter of $s_{\text{logistic}} = 1.5$ via $\sigma_\varepsilon = \pi s_{\text{logistic}}/\sqrt{3}$.

Because submission fees are per paper, the planner faces a genuine trade-off: higher fees reduce each scholar's output (Equation A25), lowering total submissions and production costs but also reducing total published quality. The sufficient statistics (σ_T, s) determine how this trade-off resolves. The total noise σ_T governs how effectively fees translate into

quality improvements: when evaluation is precise (small σ_T), the papers deterred by fees are predominantly marginal ones, so fees improve average quality with limited loss. When evaluation is noisy (large σ_T), fees reduce submissions more indiscriminately. In addition, the selection premium creates a second channel (proportional to s): higher fees reduce submissions, lowering the equilibrium cutoff and thus the selection premium. This second channel dampens the optimal fee relative to a model without within-paper variation ($s = 0$). The noise parameters are set so that the equilibrium at $(C, \tau) = (100, 3)$ produces an acceptance rate near 8% at $\gamma = 1$.

Convex capacity cost ψ captures increasing marginal cost of expanding capacity through the term $(\psi/2)C^2$; linear capacity cost κ is set to zero. The two scale parameters η (social benefit) and ψ (capacity cost) are jointly set so that $C^* = 100$ is the welfare-maximizing capacity at $\gamma = 1$, after incorporating probit participation smoothing with scale $\zeta = 0.1$ (see below). Only the ratio $\eta/\psi \approx 76.9$ determines optimal capacity; the individual magnitudes reflect the monetary unit normalization (see below). We set $\eta = 3{,}800$ and $\psi = 49.4$. The distribution parameters ($\bar{q}$, Ω_{low}, σ_T) are set so that the equilibrium at $(C, \tau) = (100, 3)$ produces an acceptance rate near 8%, matching observed rates at top journals. Conditional on these choices, the model predicts that $\tau^* = 3.0$ is the welfare-maximizing fee. This result was not directly targeted: η was set to match C^*, not τ^*. The welfare-optimal fee coincides with the observation point used for calibration, a consistency check that would fail if the welfare structure were misspecified. The per-submission review cost $r(1) \approx 8.1$ exceeds the optimal fee $\tau^* = 3.0$ by a factor of 2.7: the fee need not cover the full review cost because the planner also values the quality information that reviews generate. The gap between $r(1)$ and τ^* reflects the net benefit of the marginal submission's information value minus its congestion and prestige-feedback externalities (Section 9.3). The predicted fee falls within the range of observed submission fees at top economics journals (Table 1), where the median is \$200–300.

The participation decision, whether to target the top tier or the lower tier, is smoothed

using a probit specification (McFadden 1974): each scholar's probability of targeting the top tier is $\Phi((U^*_{\text{top}} - U^*_{\text{low}})/\zeta)$, where $\zeta > 0$ is a small idiosyncratic preference shock and Φ is the standard normal CDF. We set $\zeta = 0.1$. This smoothing replaces the hard threshold $U^*_{\text{top}} \geq U^*_{\text{low}}$ with a smooth participation weight, eliminating discontinuities at the participation boundary that would otherwise create quadrature artifacts in numerical integration. The smoothing has negligible effect on equilibrium outcomes (ζ is small relative to utility differences) but improves numerical stability and convergence.

Baseline review cost $r_0 = 10$ gives the per-submission cost with no AI; together with $\delta = 0.30$ it defines $r(\gamma)$. At $\gamma = 1$, $r(1) = 10/2^{0.3} \approx 8.1$ model units, or \$810 per submission, consistent with estimates of 2–3 referee reports at 5 hours each and average academic compensation of roughly \$74/hour (Vesper 2018). The baseline fixed entry cost is $F_0^{\text{base}} = 0.1$, representing the minimal cost (formatting, submission portal, opportunity cost of time) that any scholar incurs to submit a paper, regardless of production effort. After calibration to these targets, the structural elasticities determine how optimal capacity, fees, and welfare respond to further AI improvements.

Per-submission fees τ are transfers from a welfare perspective, so the welfare function excludes fee revenue. The fee serves a Pigouvian function (internalizing the per-submission review cost) but also stabilizes the prestige feedback loop ($\varphi > 1$) and determines whether incremental production costs are positive or negative, so the optimal fee level reflects general equilibrium forces beyond the direct review cost. Fees affect welfare indirectly through both the extensive margin (which scholars participate) and the intensive margin (how many papers each active scholar produces).

To validate the model's fee prediction, Table 1 reports current submission fees at leading economics and finance journals. Fees range from \$0 (QJE) to \$750 (JFE), with a median of \$200–300 in economics and \$400–500 in finance. Several journals differentiate fees by membership status or author income level. The model unit $\tau = 1$ corresponds to \$100, so the observed range maps to $\tau \in [0, 7.5]$ with a cross-journal median near $\tau \approx 3$.

Table 1: Submission fees at leading economics and finance journals

Journal	Field	Fee ($)
Quarterly Journal of Economics	Econ	0
Management Science	Mgmt	79
Econometrica	Econ	125
Review of Economics and Statistics	Econ	125
Review of Economic Studies	Econ	200
American Economic Review	Econ	200–300
Journal of Political Economy	Econ	125–250
Journal of Financial and Quantitative Analysis	Finance	350
Journal of Finance	Finance	400–525
Review of Financial Studies	Finance	400–460
Journal of Financial Economics	Finance	750

Fees as of early 2026. Ranges reflect member versus non-member rates. Several journals offer reduced fees or waivers for authors from low- and middle-income countries or for students and early-career researchers. Sources: journal submission guidelines pages, accessed February 7, 2026.

The model possesses a rescaling symmetry that clarifies the identification structure. Let $a > 0$ and scale all monetary parameters: $V_0 \to aV_0$, $c \to ac$, $k \to ak$, $F_0^{\text{base}} \to aF_0^{\text{base}}$, $\tau \to a\tau$, $\eta \to a\eta$, $\psi \to a\psi$, $r_0 \to ar_0$. Under this transformation, the first-order condition $MB(n) = \mathcal{C}'(n;k) + \tau$ scales by a on both sides (since MB is linear in V_0 and $\mathcal{C}'$ is linear in c or k/γ), so $n^*(q,k)$ is invariant. The participation condition $U(n^*) \geq 0$ also scales by a on both sides (the per-paper fee term $\tau \cdot n^*$ scales by a since τ scales and n^* is invariant), so all equilibrium quantities, including submissions, acceptance rates, and average published quality, are unchanged. Social welfare scales by a, and optimal capacity C^* is invariant because a cancels from the first-order condition. Optimal fees τ^* scale by a (since fees are in

the same monetary units). The model therefore identifies only ratios among the monetary parameters; one normalization pins the monetary unit. We set $c = 1$ and $V_0 = 1$ as the normalization. Setting 1 model unit $\approx$ \$100, consistent with observed submission fees at $\tau \approx 3$ corresponding to $\approx$ \$300, simultaneously determines the dollar-equivalent scale for all other monetary quantities, including baseline review cost ($r_0 = 10$ model units $\approx$ \$1,000 per submission) and scholar production costs. The large magnitudes of $\eta = 3{,}800$ and $\psi = 49.4$ (relative to, say, a normalization with $\psi = 1$) reflect this monetary unit choice; only the ratio $\eta/\psi \approx 76.9$ is economically meaningful for determining C^*. The structural elasticities (δ, φ, ξ) are dimensionless and therefore invariant to this rescaling, as are the fixed parameters $\theta = \alpha = \beta = 1/2$. The noise parameters $(\sigma_u, \sigma_\varepsilon)$ are in quality units and are also invariant. The qualitative predictions (monotonically increasing welfare under joint optimization, declining optimal fees) depend on these elasticities and noise parameters rather than on the choice of monetary unit.

The calibration separates structural parameters (with empirical guidance) from scale parameters (jointly set to match the normalization target). Predictions such as the welfare response to AI and the optimal fee schedule therefore depend on the structural elasticities rather than on the choice of monetary unit. Section 9 presents the numerical implementation and quantitative results.

9 Numerical Results

The welfare-maximizing capacity and submission fee are computed for each AI productivity level $\gamma \in [0.1, 15]$, with $\gamma = 1$ as the baseline calibrated to early-2024 frontier LLM capability and fees treated as transfers in the welfare function.

9.1 Main quantitative patterns

Table 2 reports optimal policies and outcomes at selected γ. Under joint optimization, social welfare rises monotonically in γ. Capacity is near-flat, at $C^* = 100$ for $\gamma \leq 2$ and rising to 105 at $\gamma = 15$, while optimal fees decline from $\tau^* = 3.2$ to $\tau^* = 1.1$. The acceptance rate remains in the 4%–9% range across the grid.[5]

Table 2: Optimal policy and welfare outcomes by AI productivity γ.

γ	Interpretation	C^*	τ^*	SW^*	N_{submit}	Accept
0.1	Pre-AI	100	3.2	401,349	1,161	8.6%
0.5	Early AI	100	3.1	402,368	1,196	8.4%
1	Baseline (early 2024)	100	3.0	403,390	1,230	8.1%
1.5		100	2.9	404,404	1,266	7.9%
2	Moderate AI	100	2.8	405,515	1,301	7.7%
3		101	2.3	408,346	1,541	6.6%
5	Advanced AI	102	1.7	415,899	1,966	5.2%
10	High AI	104	1.3	427,944	2,412	4.3%
15		105	1.1	434,947	2,722	3.9%

Notes: Each row reports the welfare-maximizing policy (C^*, τ^*) and resulting equilibrium outcomes for a given level of AI productivity γ. C^* is optimal capacity (papers per period); τ^* is the optimal per-paper submission fee in model units (1 unit $\approx$ \$100). N_{submit} is total submissions and Accept is the acceptance rate C^*/N_{submit}. Under joint optimization, welfare is monotonically increasing in γ, capacity is near-flat ($C^* = 100$ for $\gamma \leq 2$, rising to 105 at $\gamma = 15$), and fees decline. The implementation is described in Appendix B.

The AI index γ summarizes the effective reduction in research production costs from AI tools. At $\gamma = 1$, the model is calibrated to early-2024 frontier LLM capability, where difference-in-differences estimates find a 17% increase in output among adopting researchers (Kusumegi, Yang, Ginsparg, de Vaan, et al. 2025b). At $\gamma = 0.1$, AI provides negligible cost reduction, corresponding to the pre-LLM era. Values above 1 represent improvements beyond current capability: at $\gamma = 2$, the marginal cost of AI-assisted production is half its current level, roughly consistent with gains observed in narrower tasks such as coding (Peng et al. 2023). Beyond $\gamma \approx 3$, the model explores a regime in which AI substantially transforms the research production process, contributing to research design, data analysis,

5. The exact range is 3.9%–8.6%; we round to 4%–9%.

and iterative refinement. Results run to $\gamma = 15$ as a wide sensitivity range. We do not map γ values to calendar years: AI progress is uncertain and non-linear. The qualitative patterns (monotonically increasing welfare under fee adjustment, near-flat capacity, declining optimal fees) hold across the full range, so the policy conclusions do not depend on the pace at which γ advances.

Figure 1 shows welfare and policy as functions of γ. Panel (a) plots $SW^*(\gamma)$ over the full range; panels (b) and (c) plot $C^*(\gamma)$ and $\tau^*(\gamma)$. Welfare rises monotonically, capacity is near-flat (100 to 105), and fees decline from 3.2 to 1.1.

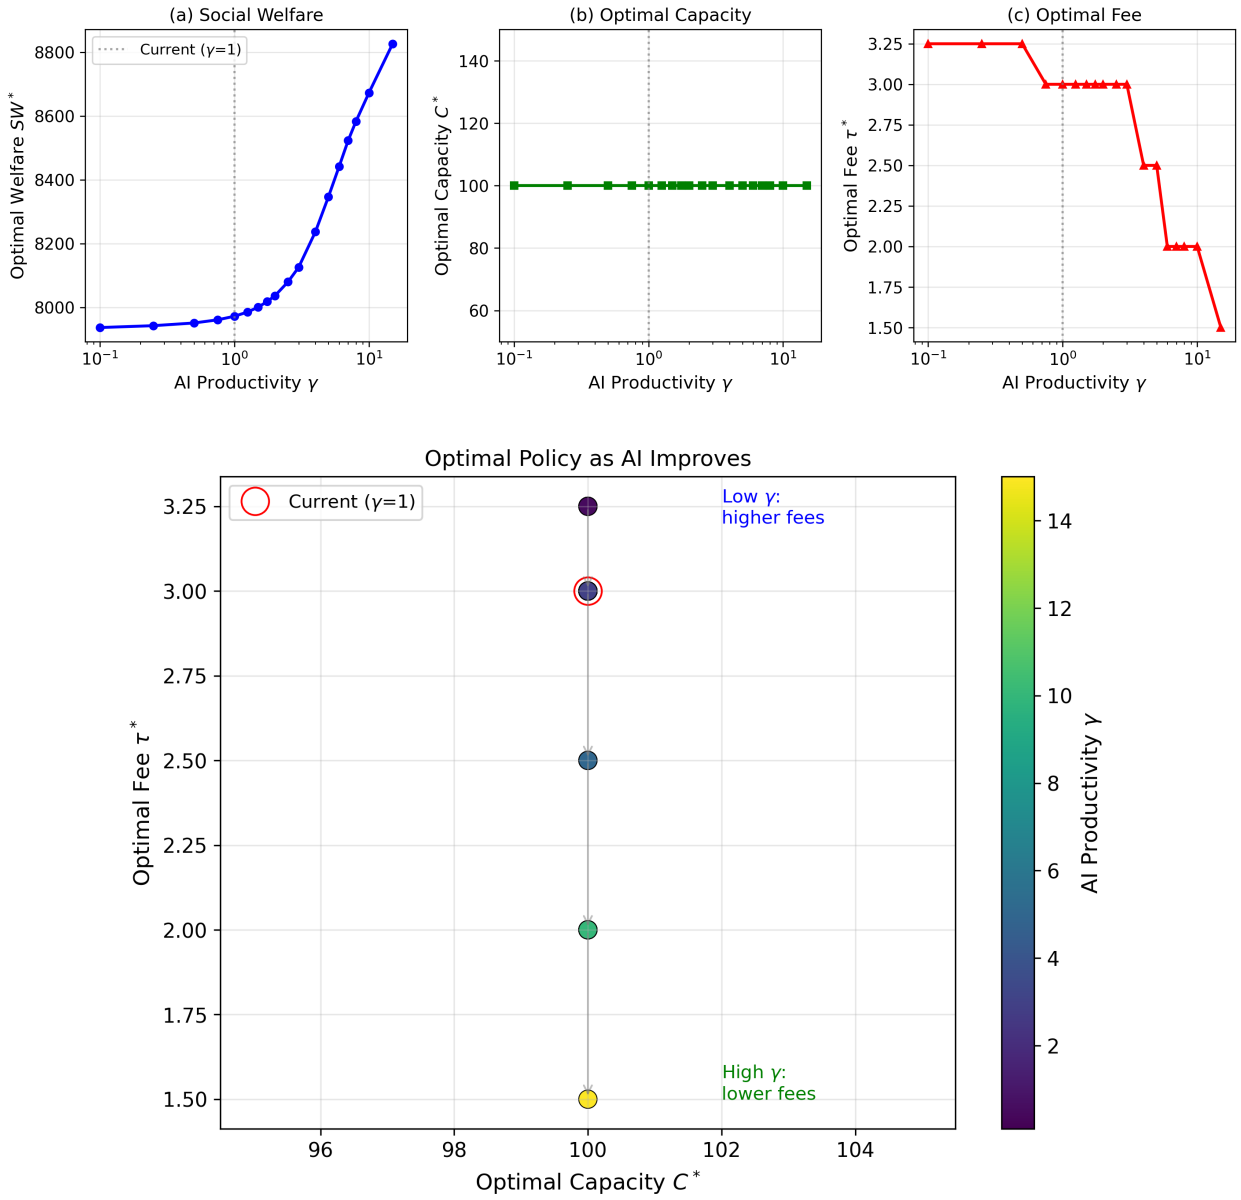


Figure 1: Optimal welfare and policy as functions of AI productivity γ: (a) social welfare $SW^*(\gamma)$, (b) optimal capacity $C^*(\gamma)$, and (c) optimal fee $\tau^*(\gamma)$.

Notes: Welfare is monotonically increasing in γ under joint optimization. Capacity is near-flat ($C^* = 100$ for $\gamma \leq 2$, rising to 105 at $\gamma = 15$), so the policy adjustment occurs primarily through the fee instrument. Fees decline from $\tau^* = 3.2$ at pre-AI levels to $\tau^* = 1.1$ at $\gamma = 15$. The fee is reported in model units, where 1 unit corresponds to \$100 in submission fees.

The decline in optimal fees tracks the per-submission review cost $r(\gamma) = r_0/(1+\gamma)^\delta$. At $\gamma = 1$, the per-submission review cost $r(1) \approx 8.1$ exceeds the optimal fee $\tau^* = 3.0$ by a factor of 2.7. The fee does not fully internalize review costs: it screens on the extensive margin and stabilizes the prestige feedback loop. As AI reduces the cost of refereeing, the review cost component of the fee falls, pulling τ^* down proportionally. In dollar terms, $\tau^* = 3.0$ at $\gamma = 1$ corresponds to \$300, declining to \$110 at $\gamma = 15$.

9.2 The role of fee adjustment

Fee adjustment converts AI productivity gains into welfare improvements. Without fees, welfare is strictly lower at every γ; with fees, welfare rises monotonically. Fees prevent submission floods that would otherwise overwhelm the review system and destroy the journal's quality signal.

At $\gamma = 1$ with the calibrated fee $\tau = 3.0$, approximately 3,370 scholars target the top tier, producing 1,230 total submissions and yielding an acceptance rate of 8.1%. Aggregate incremental production costs are -305 model units (approximately $-\$30{,}500$ at the calibrated scale of 1 unit $\approx$ \$100): scholars who target the selective top tier produce fewer, more targeted papers than they would for the lower tier, so top-tier participation actually reduces production effort. Total review costs are 9,993 (approximately \$1,000,000). Setting $\tau = 0$ at the same γ and re-optimizing capacity (which falls to $C^* = 96$) produces a qualitatively different equilibrium: 6,745 scholars target the top tier, producing 5,342 total submissions (a 4.3-fold increase). Incremental production costs swing from -305 to $+689$, review costs rise from 9,993 to 43,388 (from approximately \$1,000,000 to \$4,300,000), and total published quality Q_{pub} falls.

The welfare decomposition in Figure 2 shows this contrast across the full range of γ. The fee screens out scholars whose expected benefit from top-tier publication does not justify the submission cost, preventing production costs from swinging positive and keeping review costs from rising 4.3-fold. Without this screening, the submission externality overwhelms

the productivity gains from AI.

When review costs also cannot fall ($\delta = 0$ and $\tau = 0$), welfare peaks at moderate γ and then declines (Table 4, column C): volume effects that fees would have prevented eventually overwhelm the productivity gains from AI. With $\delta = 0.3$, welfare still rises monotonically even at $\tau = 0$, because falling per-submission review costs partly offset the volume surge, but acceptance rates remain below 2% and submissions exceed 5,000 even at low γ, rising past 10,000 at high γ (Table 3).

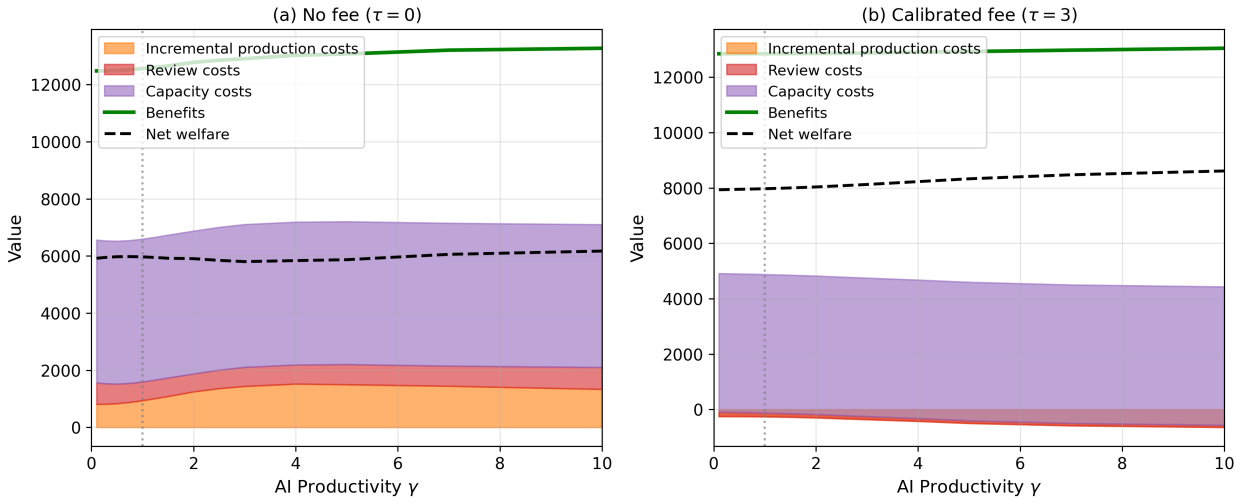


Figure 2: Welfare decomposition by component, comparing the calibrated fee ($\tau = 3$, approximating τ^*) and no-fee ($\tau = 0$) regimes at $C = 100$.

Notes: Each panel shows welfare as the sum of published-quality benefits minus incremental production costs, review costs, and the convex capacity cost. With $\tau = 3$, welfare rises monotonically as AI reduces review costs. With $\tau = 0$, submissions flood the system, incremental production and review costs rise, and welfare remains at a much lower level.

9.3 Why optimal fees decline

The fee internalizes the externality that each additional submission imposes through review costs, congestion of the acceptance threshold, and destabilization of the prestige signal. The

per-submission review cost $r(\gamma) = r_0/(1+\gamma)^\delta$ falls as AI makes refereeing cheaper. The decline in τ^* (from 3.2 to 1.1 as γ rises from 0.1 to 15) tracks the decline in $r(\gamma)$, because the review cost is the component that changes with γ; the congestion and prestige-stabilization components change little once the fee has adjusted participation.

Rising submissions might seem to require higher fees. But in the calibrated model, incremental production costs are negative when fees are in place (-305 at $\gamma = 1$). This term is defined relative to the lower-tier baseline, $\mathcal{C}(n^*_{\text{top}}) - \mathcal{C}(n^*_{\text{low}})$, not as total production cost. Fees screen out marginal participants; active top-tier participants submit fewer papers than they would at the lower tier. With convex production costs, this drives the incremental term below zero. The review cost externality is then the remaining positive externality the fee must internalize, so as $r(\gamma)$ falls, lower fees suffice.

The optimal fee τ^* falls below the per-submission review cost $r(\gamma)$ because the fee's primary role is screening on the extensive margin. The fee does not attempt to internalize the full review cost of a single paper; instead it deters marginal submissions whose congestion externality (through lower acceptance rates, higher quality cutoffs, and greater reliance on noise in selection) exceeds their quality contribution.

9.4 Policy counterfactuals

Table 3 compares equilibrium outcomes under three policies: full optimization of (C^*, τ^*), optimal fees with fixed capacity $C = 100$, and optimal capacity with no fees ($\tau = 0$).

The table shows an asymmetry between the two instruments. With fees, acceptance rates remain between 4% and 9% and average published quality is stable near $\bar{q}_{\text{pub}} \approx 4.8$, regardless of whether capacity is also optimized. Without fees, the picture is qualitatively different: the optimal no-fee capacity contracts slightly to $C^* \approx 95$–100, and acceptance rates fall below 2% (below 1% at high γ). At high γ, average quality $\bar{q}_{\text{pub}}$ without fees exceeds the with-fee level (5.06 versus 5.05 at $\gamma = 15$), because the extreme quality cutoff amplifies the selection premium: only papers with favorable idea draws clear the bar. At

Table 3: Equilibrium outcomes under alternative policy responses.

	Joint (C^*, τ^*)			Fee only $(C=100)$			No fees $(\tau=0)$			
γ	N_{sub}	Acc.	$\bar{q}_{\text{pub}}$	N_{sub}	Acc.	$\bar{q}_{\text{pub}}$	C^*	N_{sub}	Acc.	$\bar{q}_{\text{pub}}$
0.5	1,196	8.4%	4.77	1,195	8.4%	4.77	95	5,181	1.8%	4.61
1	1,230	8.1%	4.77	1,230	8.1%	4.77	96	5,342	1.8%	4.64
2	1,301	7.7%	4.78	1,301	7.7%	4.78	97	5,928	1.6%	4.73
3	1,541	6.6%	4.81	1,531	6.5%	4.82	98	6,641	1.5%	4.80
5	1,966	5.2%	4.89	1,942	5.1%	4.89	99	7,849	1.3%	4.88
10	2,412	4.3%	4.99	2,350	4.3%	5.00	100	9,732	1.0%	4.99
15	2,722	3.9%	5.05	2,632	3.8%	5.06	100	10,874	0.9%	5.06

Notes: Equilibrium submissions, acceptance rates, and average published quality under three policies. Joint (C^*, τ^*): welfare-maximizing capacity and fee (Table 2). Fee only: optimal fee with capacity fixed at $C = 100$. No fees: optimal capacity with $\tau = 0$. With fees, submissions range from 1,196 to 2,722 and acceptance rates are 4–9%. Without fees, submissions rise dramatically (5,181 to 10,874), the planner slightly contracts capacity to $C^* \approx 95$–100, and acceptance rates fall below 2% (below 1% at high γ). Average quality $\bar{q}_{\text{pub}}$ rises slightly without fees because the extreme quality cutoff amplifies the selection premium: only papers with favorable idea draws clear the bar, raising expected published quality. However, with $s = 0.12$, most of what drives acceptance at these extreme cutoffs is review noise rather than true quality variation, so the screening is inefficient.

lower γ, no-fee quality is lower (4.64 versus 4.77 at $\gamma = 1$). In either case, with $s = 0.12$, 88% of acceptance variance at these extreme cutoffs reflects review noise rather than true quality variation, so acceptance reflects noise more than merit.

The asymmetry arises because fees address the extensive margin, screening out marginal scholars whose submissions congest the system, while capacity operates only on the acceptance threshold. Without fees, even the best capacity choice cannot prevent the submission flood. In the equilibrium outcomes shown in Table 3, moving from joint optimization to fee-only policy leaves acceptance rates and average published quality close to their joint-policy values, while moving to the no-fee regime pushes acceptance rates below 2% (below 1% at high γ). The planner's objective ranks these regimes in the same order at every γ; welfare differences are in model units, not dollar equivalents.

Table 4 decomposes the welfare gains into two channels by setting $\delta = 0$, so that review costs remain fixed at r_0 regardless of γ (referees cannot use AI). Column (A) is the baseline with $\delta = 0.3$; column (B) sets $\delta = 0$ but allows optimal fees; column (C) sets $\delta = 0$ and $\tau = 0$.

The gap A−B isolates the review cost channel: how much welfare AI-assisted refereeing contributes. The gap B−C isolates the fee-screening channel: how much fees contribute when review costs cannot fall. Fee screening (B−C) exceeds the review cost channel (A−B) by a factor of 7 to 45 at every γ. When referees cannot use AI, the planner compensates by setting higher fees ($\tau^* = 1.8$ versus 1.1 at $\gamma = 15$), and welfare remains monotonically increasing. Without fees, welfare is lower at every γ and acceptance rates fall below 1%.

Table 4: Decomposing the review cost and fee-screening channels ($\delta = 0$ counterfactual).

	(A) Baseline ($\delta = 0.3$)			(B) $\delta = 0$, opt. τ			(C) $\delta = 0$, $\tau = 0$		Gaps	
γ	SW	τ^*	Acc.	SW	τ^*	Acc.	SW	Acc.	A−B	B−C
0.5	402,368	3.1	8.4%	401,027	3.2	8.6%	340,733	1.8%	1,342	60,294
1	403,390	3.0	8.1%	401,169	3.2	8.6%	342,876	1.8%	2,222	58,293
2	405,515	2.8	7.7%	402,014	3.1	8.4%	347,198	1.6%	3,501	54,816
3	408,346	2.3	6.6%	403,605	2.8	7.8%	348,093	1.5%	4,741	55,512
5	415,899	1.7	5.2%	408,564	2.2	6.4%	345,556	1.2%	7,335	63,009
10	427,944	1.3	4.3%	417,694	1.9	6.0%	339,666	1.0%	10,250	78,028
15	434,947	1.1	3.9%	422,962	1.8	5.9%	336,143	0.9%	11,984	86,820

Notes: Three scenarios isolate the review cost and fee-screening channels. (A) Baseline: AI reduces both production costs and review costs ($\delta = 0.3$). (B) $\delta = 0$: AI reduces production costs but review costs are fixed at r_0 (referees cannot use AI). (C) $\delta = 0$, $\tau = 0$: referees cannot use AI and no submission fees. The gap A−B measures the welfare contribution of AI-assisted reviewing; B−C measures the welfare contribution of fee screening when review costs cannot fall. Fee screening (B−C) exceeds the review cost channel (A−B) by a factor of 7 to 45 at every γ.

Appendix B contains solver, uniqueness, and continuity diagnostics (which includes diagnostic figures in Appendix C) and Appendix D.

10 Policy Implications

This section translates the model's quantitative predictions into guidance on how journals may want to adjust fees and capacity as AI capability improves. The AI index γ summarizes effective research productivity and screening capability, and the mapping from τ to dollar fees uses the normalization described in Section 8.

10.1 Fee adjustment as the primary policy lever

As established in Section 9.2, the central distinction is between having submission fees and not having them. The policy priority is not fine-tuning the fee level or capacity but ensuring that some submission fee is in place. Even a fixed fee of $\tau = 3$ ($300), held constant across all γ, keeps acceptance rates in the same range as joint optimization and prevents the no-fee submission flood.

10.2 A dynamic fee schedule

Section 9.3 showed that optimal fees decline from $320 to $110 as γ rises, tracking the falling review cost externality. Journals should lower fees as AI makes reviewing cheaper, not raise them in response to increased submissions. Higher AI productivity increases total submissions but reduces the social cost of processing each one. In the calibrated model, the cost reduction dominates: the optimal response is to accommodate more submissions at lower per-submission fees rather than to restrict entry with higher fees.

10.3 Capacity under joint optimization

Under joint optimization, optimal capacity is near-flat: $C^* = 100$ for $\gamma \leq 2$, rising to 105 at $\gamma = 15$ (a 5% increase over a 150-fold range of AI productivity). The primary policy adjustment occurs through fees. Once fees adjust, the submission volume and quality composition reaching the journal change little, so the marginal benefit of an additional publication slot is stable across γ. Since the marginal capacity cost (ψC) does not depend on γ, optimal C^* stays near its baseline value.

Without fees, fixed capacity combined with an uncontrolled submission flood forces the quality cutoff far above the quality support, keeping acceptance rates below 2% and below 1% at high γ. At these extreme cutoffs, with $s = 0.12$, 88% of acceptance variance reflects review noise rather than true quality variation, so screening becomes inefficient even as

average published quality rises. Fees moderate submissions so that $C = 100$ accommodates the top-quality papers at an acceptance rate of 4–9%.

10.4 Interpretation and incremental cost accounting

The policy conclusions hinge on the welfare accounting for production costs. The model counts only the incremental production cost of top-tier participation: the extra effort each scholar exerts by targeting the focal journal beyond what the scholar would exert at the lower tier. Lower-tier production happens regardless of the focal journal's policy, so the planner should not attribute those costs to the top-tier system. The incremental cost framing also prevents fees from generating spurious welfare gains by excluding scholars whose baseline production costs would be incurred at the lower tier anyway.

Incremental production costs are negative (-305 at $\gamma = 1$) when fees are in place, meaning that top-tier participants produce fewer papers than they would at the lower tier. The acceptance probability gap between tiers drives this result. At the top tier, acceptance is roughly 8% and each submission costs τ in fees, so the effective cost per expected acceptance is high. At the lower tier, acceptance is roughly 25% with no fees, so each paper is far cheaper per expected acceptance. A scholar of given type who would produce several papers for a field journal instead concentrates into fewer submissions at the top tier. Because the cost function is convex, producing fewer papers saves more than proportionally, driving the incremental cost below zero.

Removing fees reverses the sign of incremental production costs through a composition effect. About 3,375 additional scholars enter the top tier, including marginal types for whom the prestige boost outweighs the low acceptance probability. These marginal entrants produce more papers than they would at the lower tier, because top-tier prestige raises the marginal benefit of each submission even though acceptance is unlikely. The aggregate incremental cost therefore swings positive.

At the per-journal level, scholars do submit fewer papers to top-five journals than to

field journals, but the model does not capture multi-journal portfolios in which the same scholar simultaneously targets venues of different tiers. Quality is also fixed per scholar, so the model cannot represent a deliberate choice to allocate more effort toward one ambitious paper and less toward several routine ones. These simplifications mean that the negative incremental cost reflects the cost saving from reduced per-journal output, not a prediction about a scholar's total research production.

Conclusion

Researchers choose effort and AI adoption, journals incur review costs per submission, and the planner selects capacity and fees to internalize submission externalities while accounting for the social value of published output, production costs, and screening costs.

Under joint optimization of capacity and fees, welfare is monotonically increasing in AI productivity γ. Fees fall as review costs decline, and the gains from rising research productivity flow to welfare. Optimal capacity is near-flat ($C^* = 100$ for $\gamma \leq 2$, rising to 105 at $\gamma = 15$), so the policy adjustment occurs primarily through the fee instrument, which declines from \$320 to \$110 as AI capability grows.

Submission fees are the primary policy instrument. Without fees, submissions flood the system (a 4.3-fold increase at $\gamma = 1$), incremental production costs swing from negative to positive, and review costs rise 4.3-fold. At high γ, the overall acceptance rate drops below 1%, so most submitted work is rejected. Fees convert AI productivity gains into welfare gains; without fees ($\tau = 0$) and without AI-assisted refereeing ($\delta = 0$), welfare peaks at moderate γ and then declines. The optimal fee $\tau^* = 3.0$ is below the per-submission review cost $r(1) \approx 8.1$, so the fee does not fully internalize direct review costs; it acts as a congestion-screening instrument that stabilizes the prestige feedback loop, not a standard corrective tax.

Production cost accounting determines these results. Welfare counts only the incremental

cost of top-tier participation, the effort each scholar exerts above what the same scholar would exert at the lower tier. In prestige journals with tournament competition, lower-tier production happens regardless of the focal journal's policy. The sign of this increment is endogenous: negative with fees (top-tier participants concentrate into fewer papers) and positive without them (marginal entrants produce more). The sign reversal arises through general equilibrium composition effects.

Fee adjustment, not capacity expansion, translates AI productivity gains into welfare gains.

References

AAAI. 2025. "AAAI-26 Review Process Update: Scale, Integrity Measures, and Pathways to Sustainability," August 25, 2025. Accessed February 7, 2026. https://aaai.org/conference/aaai/aaai-26/review-process-update/.

Acemoglu, Daron. 2024. "The Simple Macroeconomics of AI." Pre-published, May. Working Paper. Accessed June 28, 2025. https://doi.org/10.3386/w32487. National Bureau of Economic Research: 32487. https://www.nber.org/papers/w32487.

American Economic Association. 2026. "Report of the Editor: American Economic Review." Published annually in AEA Papers and Proceedings. Accessed February 7, 2026. https://www.aeaweb.org/articles?id=10.1257/pandp.113.707.

American Finance Association. 2026. "Journal Statistics." Accessed February 7, 2026. https://afajof.org/journal-statistics/.

Arroyo-Machado, Wenceslao, Jinghuan Ma, Tipeng Chen, Timothy P. Johnson, Shaika Islam, Lesley Michalegko, and Eric Welch. 2025. "Generative AI and Academic Scientists in US Universities: Perception, Experience, and Adoption Intentions." *PLOS ONE* 20, no. 8 (August 28, 2025): e0330416.

arXiv. 2024. "arXiv Sets New Record for Monthly Submissions—Again," November 4, 2024. Accessed February 7, 2026. https://blog.arxiv.org/2024/11/04/arxiv-sets-new-record-for-monthly-submissions-again/.

———. 2026. "arXiv Yearly Submission Rates: Economics." Accessed February 7, 2026. https://arxiv.org/year/econ.

Bansak, Cynthia, Wendy Dunn, Ellen Meade, and Martha Starr. 2024. "Impact versus Inclusion in the Economics Profession: Insights from the Papers and Proceedings." *AEA Papers and Proceedings* 114 (May): 292–299.

Bergstrom, Theodore C. 2001. "Free Labour for Costly Journals?" *Journal of Economic Perspectives* 15, no. 4 (December): 183–198.

Blank, Rebecca M. 1991. "The Effects of Double-Blind versus Single-Blind Reviewing: Experimental Evidence from The American Economic Review." *The American Economic Review* 81 (5): 1041–1067. JSTOR: 2006906.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. "Generative AI at Work." *The Quarterly Journal of Economics* 140, no. 2 (May 1, 2025): 889–942.

Butson, Russell, and Rachel and Spronken-Smith. 2024. "AI and Its Implications for Research in Higher Education: A Critical Dialogue." *Higher Education Research & Development* 43, no. 3 (April 2, 2024): 563–577.

Cao, Yi, Long Chen, Jennifer Tucker, and Chi Wan. 2025. "Can Generative AI Help Identify Peer Firms?" Pre-published, April 4, 2025. SSRN Scholarly Paper. Accessed April 15, 2025. https://doi.org/10.2139/ssrn.4761624. Social Science Research Network: 4761624. https://papers.ssrn.com/abstract=4761624.

Card, David, and Stefano DellaVigna. 2013. "Nine Facts about Top Journals in Economics." *Journal of Economic Literature* 51, no. 1 (March): 144–161.

Checco, Alessandro, Lorenzo Bracciale, Pierpaolo Loreti, Stephen Pinfield, and Giuseppe Bianchi. 2021. "AI-assisted Peer Review." *Humanities and Social Sciences Communications* 8, no. 1 (January 25, 2021): 25.

Chen, Jian, Guohao Tang, Guofu Zhou, and Wu Zhu. 2023. "ChatGPT and Deepseek: Can They Predict the Stock Market and Macroeconomy?" Pre-published, July 31, 2023. SSRN Scholarly Paper. Accessed April 11, 2025. https://doi.org/10.2139/ssrn.4660148. Social Science Research Network: 4660148. https://papers.ssrn.com/abstract=4660148.

Chen, Shuaiyu, T. Clifton Green, Huseyin Gulen, and Dexin Zhou. 2024. "What Does ChatGPT Make of Historical Stock Returns? Extrapolation and Miscalibration in LLM Stock Return Forecasts." Pre-published, August 30, 2024. SSRN Scholarly Paper. Accessed April 7, 2025. https://doi.org/10.2139/ssrn.4941906. Social Science Research Network: 4941906. https://papers.ssrn.com/abstract=4941906.

Chen, Zenan, and Jason Chan. 2024. "Large Language Model in Creative Work: The Role of Collaboration Modality and User Expertise." *Management Science* 70, no. 12 (December): 9101–9117.

CHI 2024 Program Chairs. 2023. "Papers Track, Post-Submission Report," October 16, 2023. Accessed February 7, 2026. https://chi2024.acm.org/2023/10/16/papers-track-post-submission-report/.

Chubb, Jennifer, Peter Cowling, and Darren Reed. 2022. "Speeding up to Keep up: Exploring the Use of AI in the Research Process." *Ai & Society* 37 (4): 1439–1457. PMID: 34667374.

Cotton, Christopher. 2013. "Submission Fees and Response Times in Academic Publishing." *American Economic Review* 103, no. 1 (February): 501–509.

Degen, Dominik, Dr Jens Kengelbach, Daniel Kim, Soenke Sievers, and Yiran Wang. 2024. "Large Language Models and M&A: Can ChatGPT Help Forecast M&A Activity?" Pre-published, June 27, 2024. SSRN Scholarly Paper. Accessed April 11, 2025. https://doi.org/10.2139/ssrn.4862121. Social Science Research Network: 4862121. https://papers.ssrn.com/abstract=4862121.

Dell, Melissa. 2025. "Deep Learning for Economists." *Journal of Economic Literature* 63, no. 1 (March): 5–58.

Demirci, Ozge, Jonas Hannane, and Xinrong Zhu. 2025. "Who Is AI Replacing? The Impact of Generative AI on Online Freelancing Platforms." *Management Science* (January 24, 2025).

Econometric Society. 2026. "Econometrica Annual Reports." Accessed February 7, 2026. https://www.econometricsociety.org/publications/econometrica/annual-reports.

Ellison, Glenn. 2002a. "Evolving Standards for Academic Publishing: A Q-r Theory." *Journal of Political Economy* 110, no. 5 (October): 994–1034.

———. 2002b. "The Slowdown of the Economics Publishing Process." *Journal of Political Economy* 110 (5): 947–993. JSTOR: 10.1086/341868.

Engelberg, Joseph, Asaf Manela, William Mullins, and Luka Vulicevic. 2025. "Entity Neutering." Pre-published, March 17, 2025. SSRN Scholarly Paper. Accessed April 7, 2025. Social Science Research Network: 5182756. https://papers.ssrn.com/abstract=5182756.

Farber, Shai. 2024. "Enhancing Peer Review Efficiency: A Mixed-Methods Analysis of Artificial Intelligence-Assisted Reviewer Selection across Academic Disciplines." *Learned Publishing* 37 (4): e1638.

Fišar, Miloš, Ben Greiner, Christoph Huber, Elena Katok, and Ali I. Ozkes. 2024. "Reproducibility in Management Science." *Management Science* 70, no. 3 (March): 1343–1356.

Garicano, Luis. 2000. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy* 108, no. 5 (October): 874–904.

Garicano, Luis, and Esteban Rossi-Hansberg. 2004. "Inequality and the Organization of Knowledge." *American Economic Review* 94, no. 2 (May): 197–202.

Hadavand, Aboozar, Daniel S. Hamermesh, and Wesley W. Wilson. 2024. "Publishing Economics: How Slow? Why Slow? Is Slow Productive? How to Fix Slow?" *Journal of Economic Literature* 62, no. 1 (March): 269–293.

Hanson, Mark A., Pablo Gómez Barreiro, Paolo Crosetto, and Dan Brockington. 2024. "The Strain on Scientific Publishing." *Quantitative Science Studies* 5, no. 4 (November 1, 2024): 823–843. arXiv: 2309.15884 [cs].

Heckman, James J., and Sidharth Moktan. 2020. "Publishing and Promotion in Economics: The Tyranny of the Top Five." *Journal of Economic Literature* 58, no. 2 (June): 419–470.

Ide, Enrique, and Eduard Talamàs. 2025. "Artificial Intelligence in the Knowledge Economy." *Journal of Political Economy* 133, no. 12 (December): 3762–3800.

Jeon, Doh-Shin, and Jean-Charles Rochet. 2010. "The Pricing of Academic Journals: A Two-Sided Market Perspective." *American Economic Journal: Microeconomics* 2, no. 2 (May): 222–255.

Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang. 2025. "ChatGPT and Corporate Policies." Pre-published, February 28, 2025. SSRN Scholarly Paper. Accessed April 11, 2025. https://doi.org/10.2139/ssrn.4521096. Social Science Research Network: 4521096. https://papers.ssrn.com/abstract=4521096.

Kanazawa, Kyogo, Daiji Kawaguchi, Hitoshi Shigeoka, and Yasutora Watanabe. 2025. "AI, Skill, and Productivity: The Case of Taxi Drivers." *Management Science* (June 9, 2025).

Khalifa, Mohamed, and Mona Albadawy. 2024. "Using Artificial Intelligence in Academic Writing and Research: An Essential Productivity Tool." *Computer Methods and Programs in Biomedicine Update* 5 (January 1, 2024): 100145.

Ko, Hyungjin, and Jaewook Lee. 2023. "Can Chatgpt Improve Investment Decision? From a Portfolio Management Perspective." *SSRN Electronic Journal.*

Kobak, Dmitry, Rita González-Márquez, Emőke-Ágnes Horvát, and Jan Lause. 2025. "Delving into LLM-assisted Writing in Biomedical Publications through Excess Vocabulary." *Science Advances* 11, no. 27 (July 4, 2025): eadt3813. PMID: 40601754.

Koren, Moran. 2024. "The Gatekeeper Effect: The Implications of Pre-Screening, Self-Selection, and Bias for Hiring Processes." *Management Science* (May 17, 2024).

Kusumegi, Keigo, Xinyu Yang, Paul Ginsparg, Mathijs de Vaan, Toby Stuart, and Yian Yin. 2025a. "Replication Materials for "Scientific Production in the Era of Large Language Models"." Accessed March 8, 2026. https://doi.org/10.6084/m9.figshare.30359437. https://figshare.com/articles/dataset/Scientific_production_in_the_era_of_Large_Language_Models/30359437.

Kusumegi, Keigo, Xinyu Yang, Paul Ginsparg, Mathijs de Vaan, Toby Stuart, and Yian Yin. 2025b. "Scientific Production in the Era of Large Language Models." *Science* 390, no. 6779 (December 18, 2025): 1240–1243.

Kwon, Diana. 2025. "Is It OK for AI to Write Science Papers? Nature Survey Shows Researchers Are Split." *Nature* 641, no. 8063 (May 14, 2025): 574–578.

Lazear, Edward P., and Sherwin Rosen. 1981. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy* 89, no. 5 (October): 841–864.

Leslie, Derek. 2005. "Are Delays in Academic Publishing Necessary?" *American Economic Review* 95, no. 1 (March): 407–413.

Liang, Weixin, Zachary Izzo, Yaohui Zhang, Haley Lepp, Hancheng Cao, Xuandong Zhao, Lingjiao Chen, et al. 2024. "Monitoring AI-Modified Content at Scale: A Case Study on the Impact of ChatGPT on AI Conference Peer Reviews." Pre-published, June 15, 2024. Accessed June 28, 2025. https://doi.org/10.48550/arXiv.2403.07183. arXiv: 2403.07183 [cs]. http://arxiv.org/abs/2403.07183.

Liang, Weixin, Yaohui Zhang, Zhengxuan Wu, Haley Lepp, Wenlong Ji, Xuandong Zhao, Hancheng Cao, et al. 2025. "Quantifying Large Language Model Usage in Scientific Papers." *Nature Human Behaviour* 9, no. 12 (December): 2599–2609.

Lippi, Francesco. 2023. "The Economic Journal 2022 Annual Report." Royal Economic Society. Accessed February 19, 2026. https://res.org.uk/wp-content/uploads/2023/03/The-Economic-Journal-2022-Annual-Report.pdf.

Lopez-Lira, Alejandro, and Yuehua Tang. 2023. "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models." *SSRN Electronic Journal* (April 6, 2023).

McCabe, Mark J., and Christopher M. Snyder. 2005. "Open Access and Academic Journal Quality." *American Economic Review* 95, no. 2 (May): 453–459.

McFadden, Daniel. 1974. "Conditional Logit Analysis of Qualitative Choice Behavior." In *Frontiers in Econometrics,* edited by Paul Zarembka, 105–142. New York: Academic Press.

Naddaf, Miryam. 2025. "Major AI Conference Flooded with Peer Reviews Written Fully by AI." *Nature* (November 27, 2025).

NeurIPS Program Chairs. 2025. "Reflections on the 2025 Review Process from the Program Committee Chairs," September 30, 2025. Accessed February 7, 2026. https://blog.neurips.cc/2025/09/30/reflections-on-the-2025-review-process-from-the-program-committee-chairs/.

Novy-Marx, Robert, and Mihail Velikov. 2026. "Artificial Intelligence–Powered (Finance) Scholarship." *Journal of Economic Literature* 64, no. 1 (March): 5–37.

Noy, Shakked, and Whitney Zhang. 2023. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science* 381, no. 6654 (July 14, 2023): 187–192.

Peng, Sida, Eirini Kalliamvakou, Peter Cihon, and Mert Demirer. 2023. "The Impact of AI on Developer Productivity: Evidence from GitHub Copilot." *arXiv preprint,* arXiv: 2302.06590.

Review of Economics and Statistics. 2026. "Annual Reports." Published annually in the journal. Accessed February 7, 2026. https://direct.mit.edu/rest/article/106/3/i/121060/.

Review of Finance. 2026. "Editorial Statistics." Accessed February 7, 2026. https://revfin.org/editorial-statistics/.

Rosen, Sherwin. 1981. "The Economics of Superstars." *The American Economic Review* 71 (5): 845–858. JSTOR: 1803469.

SCImago Lab. 2026. "SCImago Journal and Country Rank." Accessed February 7, 2026. https://www.scimagojr.com/.

Sinatra, Roberta, Dashun Wang, Pierre Deville, Chaoming Song, and Albert-László Barabási. 2016. "Quantifying the Evolution of Individual Scientific Impact." *Science* 354 (6312): aaf5239.

Van de Schoot, Rens, Jonathan de Bruin, Raoul Schram, Parisa Zahedi, Jan de Boer, Felix Weijdema, Bianca Kramer, et al. 2021. "An Open Source Machine Learning Framework for Efficient and Transparent Systematic Reviews." *Nature Machine Intelligence* 3, no. 2 (February): 125–133.

Vesper, Inga. 2018. "Peer Reviewers Unmasked: Largest Global Survey Reveals Trends." *Nature* (September 7, 2018).

Wang, Weiguang, Guodong (Gordon) Gao, and Ritu Agarwal. 2024. "Friend or Foe? Teaming Between Artificial Intelligence and Workers with Variation in Experience." *Management Science* 70, no. 9 (September): 5753–5775.

A Model Derivations

A.1 Cost Minimization

With the production constraint $e + \gamma\lambda = n^{1/\theta}$, the KKT conditions yield two regimes:

- No-AI regime ($\lambda = 0$): If the target $n \leq n^{\text{AI}}(k)$, the scholar uses only human effort: $e = n^{1/\theta}$.
- AI regime ($\lambda > 0$): If $n > n^{\text{AI}}(k)$, the scholar uses the cost-minimizing level of human effort $e^* = k/(c\gamma)$ and adjusts AI adoption λ to meet the paper target.

Cost minimization yields an explicit piecewise minimum cost function:

$$\mathcal{C}(n;k) = \begin{cases} \frac{c}{2}n^{2/\theta}, & n \leq n^{\text{AI}}(k) \quad (\lambda = 0), \\ \frac{k}{\gamma}n^{1/\theta} - \frac{k^2}{2c\gamma^2}, & n > n^{\text{AI}}(k) \quad (\lambda > 0). \end{cases} \tag{A1}$$

The derivatives are:

$$\mathcal{C}'(n;k) = \begin{cases} \frac{c}{\theta}n^{2/\theta-1}, & n \leq n^{\text{AI}}(k), \\ \frac{k}{\gamma\theta}n^{1/\theta-1}, & n > n^{\text{AI}}(k), \end{cases} \qquad \mathcal{C}''(n;k) = \begin{cases} \frac{c}{\theta}\left(\frac{2}{\theta}-1\right)n^{2/\theta-2} > 0, \\ \frac{k}{\gamma\theta}\left(\frac{1}{\theta}-1\right)n^{1/\theta-2} > 0. \end{cases} \tag{A2}$$

At $n = n^{\text{AI}}(k)$, both the cost and its derivative are continuous (both branches coincide), so $\mathcal{C}(n;k)$ is C^1 and strictly convex in n for all $k > 0$. The cost function does not depend on quality q: quality is exogenous and does not affect production costs.

With $\theta = 1/2$, the cost function simplifies. In the no-AI regime, $\mathcal{C}(n;k) = (c/2)n^4$, which is a quartic. In the AI regime, $\mathcal{C}(n;k) = (k/\gamma)n^2 - k^2/(2c\gamma^2)$, which is quadratic in n.

A.2 Outside Option: Closed-Form Solutions

Because the lower-tier problem has the same cost structure as the top-tier problem, the optimal lower-tier output n^*_{low} takes the same functional form as n^* (Section 3.2) with

$(V_{\text{low}}, p_{\text{low}}, \tau_{\text{low}})$ replacing the endogenous $(V(\bar{q}_{\text{pub}}), p(q, q_{\text{cut}}), \tau)$. With $\alpha = \beta = \theta = 1/2$ and $\tau_{\text{low}} = 0$, the closed forms are:

$$n^*_{\text{low,noAI}}\big|_{\tau_{\text{low}}=0} = \left[\frac{V_{\text{low}}\sqrt{p_{\text{low}} \cdot q}}{4c}\right]^{2/7}, \tag{A3}$$

$$n^*_{\text{low,AI}}\big|_{\tau_{\text{low}}=0} = \left[\frac{\gamma \cdot V_{\text{low}}\sqrt{p_{\text{low}} \cdot q}}{4k}\right]^{2/3}. \tag{A4}$$

The regime selection follows the same logic: $n^*_{\text{low}} = n^*_{\text{low,noAI}}$ if $n^*_{\text{low,noAI}} \leq n^{\text{AI}}(k)$, and $n^*_{\text{low}} = n^*_{\text{low,AI}}$ otherwise.

The scholar's top-tier submission choice, denoted n^{**}, is:

$$n^{**}(q, k, q_{\text{cut}}, \bar{q}_{\text{pub}}) = \begin{cases} n^*(q, k, q_{\text{cut}}, \bar{q}_{\text{pub}}) & \text{if } U^* \geq U^*_{\text{low}} \\ 0 & \text{if } U^* < U^*_{\text{low}} \end{cases} \tag{A5}$$

where $n^{**} = 0$ means the scholar targets the lower tier instead, producing n^*_{low} papers there. The submission choice is an allocation decision, not a production decision: every scholar produces research regardless; the question is which tier to target. $U^*_{\text{low}}(q, k)$ depends on the scholar's type (q, k) but not on the focal journal's equilibrium objects $(q_{\text{cut}}, \bar{q}_{\text{pub}}, \tau)$. The equilibrium system therefore remains a single-journal fixed-point problem with the same two equations (capacity clearing and quality consistency), with only the participation boundary redefined.

This framework has three implications. The first-order condition (Equation 10) is unchanged: conditional on submitting to the focal journal, the scholar's optimal number of papers n^* is the same as before. The outside option affects only the extensive margin (who submits), not the intensive margin (how many papers). Scholars with low intrinsic quality (q) or high private costs for AI (k) are more likely to prefer the lower tier. Low-q scholars face low acceptance probability $p(q, q_{\text{cut}})$ at the focal journal and low value per acceptance $\sqrt{q}$, making the lower tier relatively attractive. High-k scholars face high production costs

$\mathcal{C}(n;k)$. A high quality cutoff q_{cut} or a high submission fee τ amplifies these effects. The per-submission fee screens through two channels: it excludes marginal types from participating (extensive margin) and reduces the number of papers produced by each active scholar (intensive margin). The intensive margin effect is absent with lump-sum entry fees, which affect only participation. The per-paper structure gives the planner a two-margin policy instrument, since raising τ reduces both the number of participants and the output per participant.

The outside option creates an extensive margin for fees. Without an outside option, all scholars would participate regardless of fee level, and fees could only reduce output per scholar (the intensive margin). With the outside option, fees screen out marginal scholars who would earn higher utility at the lower tier. This screening drives incremental production costs negative: high-benefit participants concentrate into fewer, more targeted papers than they would produce at the lower tier. The outside option also clarifies the welfare accounting: production costs at the lower tier are incurred regardless of the focal journal's policy, so the planner counts only the incremental cost of top-tier participation.

A.3 Selection Premium Derivation

For a paper with baseline quality q, the joint distribution of $(q_{\text{true}}, q_{\text{obs}}) = (q+u, q+u+\varepsilon)$ is bivariate normal. Conditional on acceptance ($q_{\text{obs}} \geq q_{\text{cut}}$), the expected true quality follows Equation 16 in the main text. The derivation follows from the standard truncated bivariate normal result: the covariance between $q_{\text{true}} = q + u$ and $q_{\text{obs}} = q + u + \varepsilon$ is σ_u^2 (since u is the common component), and truncating a bivariate normal on one variable shifts the conditional mean of the other by $(\text{Cov}/\sigma_{\text{truncated}}) \cdot h(z)$.

The selection premium (Equation 17) has three properties. It is always non-negative ($h(z) \geq 0$ for all z). It is largest for papers near the cutoff ($q \approx q_{\text{cut}}$), where selection is most severe. It vanishes for papers far above the cutoff ($q \gg q_{\text{cut}}$), since those papers are accepted regardless of u. The signal share s controls the magnitude: when $s = 0$ (all noise is

evaluation noise), the premium vanishes because there is no within-paper variation to select on; when s is large, the journal can effectively cherry-pick papers with favorable idea draws.

A.4 Equilibrium Existence

A capacity C is feasible only if it does not exceed the maximum possible acceptances when the cutoff is arbitrarily low, that is, when $p(q, q_{\text{cut}}) \to 1$ for all types. Let $\bar{q}_{\text{pub}}^{\infty}$ denote the prestige-consistent average quality induced when all papers are accepted. Under $p \equiv 1$, $\bar{q}_{\text{pub}}^{\infty}$ is the fixed point where (i) prestige $V(\bar{q}_{\text{pub}})$ determines production choices and (ii) those choices with $p = 1$ yield average quality $\bar{q}_{\text{pub}}$. Then:

$$C \leq C^{\max}(\gamma, \tau) \equiv \iint_{\mathcal{P}(\gamma,\tau)} n^{**}(q, k, q_{\text{cut}} \to -\infty, \bar{q}_{\text{pub}}^{\infty}) \cdot f(q, k)\, dq\, dk, \tag{A6}$$

where the integral is taken over the participation set

$$\mathcal{P}(\gamma, \tau) = \{(q, k) \mid U^*(q, k) \geq U_{\text{low}}^*(q, k)\}.$$

Since $n^{**}(q, k) = 0$ for types outside the participation set, only scholars who prefer the focal journal contribute. For $C > C^{\max}$, the journal would have unfilled capacity. In our analysis, we restrict attention to the feasible range $C \in (0, C^{\max}]$.

The maximum capacity $C^{\max}$ depends on both AI productivity γ and the submission fee τ through the participation constraint. Higher fees reduce participation through two channels: the per-paper fee raises effective marginal cost (reducing n^*) and the total fee burden $\tau \cdot n^*$ reduces surplus. A type-(q, k) scholar participates if and only if:

$$\Pi(q, k, \gamma, \tau) \equiv \sqrt{pq} \cdot V \cdot \sqrt{n^*} - \mathcal{C}(n^*; k) - \tau \cdot n^* - F_0^{\text{base}} \geq U_{\text{low}}^*(q, k), \tag{A7}$$

where $\Pi(q, k, \gamma, \tau)$ is the net top-tier surplus (benefit minus variable cost minus fees minus fixed cost) at the optimal $n^*(\tau)$. When $\tau = 0$, the first-order condition yields closed-form

surplus expressions in each regime:

$$\Pi\big|_{\tau=0} = \begin{cases} \dfrac{7}{8}\,B - F_0^{\text{base}}, & \text{no-AI regime,} \\ \dfrac{3}{4}\,B + \dfrac{k^2}{2c\gamma^2} - F_0^{\text{base}}, & \text{AI regime,} \end{cases} \tag{A8}$$

where $B = \sqrt{pq} \cdot V \cdot \sqrt{n^*}$ is gross benefit. With $\alpha = \beta = \theta = 1/2$, the cost-to-benefit ratio at the optimum is $\theta\beta/2 = 1/8$ in the no-AI regime and $\theta\beta = 1/4$ (for the variable component) in the AI regime, yielding surplus fractions of $7/8$ and approximately $3/4$ respectively. The constant offset $k^2/(2c\gamma^2)$ in the AI regime is the cost saving from substituting AI for effort at the margin; it is type-specific but small relative to B for most scholars. For $\tau > 0$, the surplus must be evaluated numerically at the optimal $n^*(\tau)$ obtained from the first-order condition. The participation set $\mathcal{P}(\gamma, \tau)$ consists of all (q, k) pairs where $\Pi \geq U^*_{\text{low}}(q, k)$.

The participation surplus yields the necessary and sufficient condition for equilibrium existence. The maximum fillable capacity is:

$$C^{\max}(\gamma, \tau) = \iint_{\mathcal{P}(\gamma,\tau)} n^*(q, k, \gamma, \tau) \cdot f(q, k)\, dq\, dk, \tag{A9}$$

where $n^*(q, k, \gamma, \tau)$ solves the first-order condition Equation 10. At $\tau = 0$, the closed-form solutions from Equations A10–A11 apply directly; in general, the integral is evaluated via quadrature.

An equilibrium within the feasible range exists because the best-response functions are continuous and bounded. For a given capacity C, the supply of accepted papers is continuous in $(q_{\text{cut}}, \bar{q}_{\text{pub}})$, exceeds C when q_{cut} is sufficiently low, and falls to zero as $q_{\text{cut}} \to \infty$. The continuous type distribution ensures that supply varies smoothly with q_{cut}, avoiding the jumps that arise in discrete-type models when a mass of agents enters or exits. By the Intermediate Value Theorem, there exists a $(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})$ that equates supply to C while satisfying prestige consistency.

A.5 Analytical Thresholds and Limiting Cases

A.5.1 Optimal Output and the Role of Fees

The first-order condition (Equation 10) equates marginal benefit to effective marginal cost $\mathcal{C}'(n;k)+\tau$. When $\tau=0$, the condition yields closed-form solutions in each regime. With $\theta=\beta=1/2$, the exponents simplify.

No-AI regime ($n \le n^{\mathrm{AI}}(k)$): Using $\mathcal{C}'(n;k)=\frac{c}{\theta}n^{2/\theta-1}=2cn^3$,

$$n^*_{\mathrm{noAI}}\big|_{\tau=0}=\left[\frac{\sqrt{pq}\cdot V}{4c}\right]^{2/7}. \tag{A10}$$

The exponent $1/(2/\theta-\beta)=1/(4-1/2)=2/7$ governs how output responds to the benefit-to-cost ratio.

AI regime ($n > n^{\mathrm{AI}}(k)$): Using $\mathcal{C}'(n;k)=\frac{k}{\gamma\theta}n^{1/\theta-1}=\frac{2k}{\gamma}n$,

$$n^*_{\mathrm{AI}}\big|_{\tau=0}=\left[\frac{\gamma\sqrt{pq}\cdot V}{4k}\right]^{2/3}. \tag{A11}$$

The exponent $1/(1/\theta-\beta)=1/(2-1/2)=2/3$ is larger, reflecting the lower curvature of the AI-regime cost function and hence greater output responsiveness to incentives.

When $\tau>0$, the additive fee term τ in the first-order condition breaks the power-law structure that enables these closed forms. The optimal $n^*(\tau)$ must then be obtained by solving the FOC numerically (a one-dimensional root-finding problem for each type), with the $\tau=0$ closed forms serving as upper bounds. The solution algorithm computes $n^*|_{\tau=0}$ as an initial bracket and applies Newton's method; convergence is rapid (typically 3–5 iterations) because the objective is strictly concave.

The scholar's optimal choice follows a simple sequence. First compute the AI adoption threshold $n^{\mathrm{AI}}(k)=(k/(c\gamma))^{1/2}$. When $\tau=0$, compute the two candidate closed-form solutions n^*_{noAI} and n^*_{AI} and select the regime-consistent one; when $\tau>0$, solve the FOC

numerically in each regime:

$$n^* = \begin{cases} n^*_{\text{noAI}}, & \text{if } n^*_{\text{noAI}} \leq n^{\text{AI}}(k), \\ n^*_{\text{AI}}, & \text{if } n^*_{\text{AI}} > n^{\text{AI}}(k). \end{cases} \tag{A12}$$

Finally apply the submission condition: $n^{**} = n^*$ if $U^*(q,k) \geq U^*_{\text{low}}(q,k)$, else $n^{**} = 0$ (the scholar targets the lower tier instead).

A.5.2 The Participation Threshold

The condition $U^*(q,k;q_{\text{cut}},\bar{q}_{\text{pub}},\tau) = U^*_{\text{low}}(q,k)$ defines the boundary of participation for each agent type. This implicit equation defines the participation set $\mathcal{P}(q_{\text{cut}},\bar{q}_{\text{pub}},\tau)$ in the (q,k) space. For any given quality cutoff q_{cut}, there is a minimum set of agent characteristics (high q, low k) required to prefer the focal journal over the lower tier. Scholars with high quality q benefit from both higher acceptance probability $p(q,q_{\text{cut}})$ and higher value per acceptance $\sqrt{q}$, making top-tier participation more attractive. Scholars with low AI cost k face lower production costs $\mathcal{C}(n;k)$, which raises surplus at both tiers but disproportionately benefits the top tier (where optimal output is higher). This threshold determines which agent types submit to the focal journal rather than the lower tier.

A.5.3 The Role of Noise and the Limiting Zero-Noise Case

Two sources of randomness enter the model: within-paper quality variation $u \sim N(0,\sigma_u^2)$ and evaluation noise $\varepsilon \sim N(0,\sigma_\varepsilon^2)$. These enter the model through two sufficient statistics: the total noise scale $\sigma_T = \sqrt{\sigma_u^2 + \sigma_\varepsilon^2}$, which governs the steepness of acceptance probabilities, and the signal share $s \equiv \sigma_u^2/\sigma_T^2$, which controls the selection premium magnitude (Equation 17). Within-paper variation (captured by $s > 0$) generates the selection premium that raises expected published quality above baseline. Evaluation noise (captured by $1 - s$) reflects the stochasticity of peer review (referee taste, subfield-specific norms, measurement error). The combined effect through σ_T generates smooth acceptance probabilities, matching the

uncertainty researchers face.

Incorporating noise yields three advantages. It generates smooth acceptance probabilities, so a high-quality paper is not guaranteed publication and a lower-quality paper is not doomed to rejection. It avoids the bunching phenomenon of deterministic models and produces a smooth distribution of paper quantities across scholar types, which is more empirically plausible. It also delivers heterogeneous success rates: scholars with different quality types q face different acceptance probabilities $p(q, q_{\text{cut}})$, creating stratification in publication outcomes even among equally productive scholars.

The model with noise requires a numerical solution. As both noise variances go to zero, the framework simplifies to a deterministic quality-bar model that serves as a theoretical benchmark. As $\sigma_T \to 0$, the probit acceptance probability $\Phi((q - q_{\text{cut}})/\sigma_T)$ converges to a step function: it is 1 if $q \geq q_{\text{cut}}$ and 0 otherwise. The selection premium also vanishes in this limit, since $s \cdot \sigma_T \to 0$ eliminates within-paper variation.

In this zero-noise limit, the scholar's problem simplifies. Papers with baseline quality q are either always accepted ($q \geq q_{\text{cut}}$) or never accepted ($q < q_{\text{cut}}$). Scholars of type $q < q_{\text{cut}}$ target the lower tier instead, since $p(q, q_{\text{cut}}) = 0$ at the focal journal yields zero benefit there. Scholars of type $q \geq q_{\text{cut}}$ face $p(q, q_{\text{cut}}) = 1$ at the focal journal and, conditional on choosing to submit there rather than to the lower tier, produce papers until marginal benefit equals marginal cost:

$$\frac{1}{2}\sqrt{q/n} \cdot V(\bar{q}_{\text{pub}}) = \mathcal{C}'(n; k) + \tau. \tag{A13}$$

The noisy evaluation model is therefore a more general framework that nests the simpler threshold-based model as a limiting case.

A.6 Comparative Statics

Applying the Implicit Function Theorem to Equation 10 yields comparative statics and predictions for the numerical implementation. The results below hold for all $\tau \geq 0$.

We rewrite the first-order condition as an implicit function, H, set to zero at the optimum. Let $V \equiv V(\bar{q}_{\text{pub}})$ and $p \equiv p(q, q_{\text{cut}})$ throughout this subsection. Then:

$$H(n; q, k, q_{\text{cut}}, V, \tau) = \underbrace{\tfrac{1}{2}\sqrt{pq/n} \cdot V}_{\text{Marginal Benefit}} - \underbrace{\left(\frac{d\mathcal{C}(n;k)}{dn} + \tau\right)}_{\text{Effective Marginal Cost}} = 0. \tag{A14}$$

The second-order condition for a maximum ensures that $\partial H/\partial n < 0$ at the optimum. By the Implicit Function Theorem, the effect of a change in any parameter (say, x) on the optimal choice n^* is given by:

$$\frac{dn^*}{dx} = -\frac{\partial H/\partial x}{\partial H/\partial n}. \tag{A15}$$

Since the denominator is negative, the sign of the effect is determined purely by the sign of $\partial H/\partial x$. For AI productivity, an increase in γ lowers the marginal cost $\frac{d\mathcal{C}}{dn}$ while leaving marginal benefit unchanged, so $\frac{\partial H}{\partial \gamma} = -\frac{\partial (MC)}{\partial \gamma} > 0$ and:

$$\frac{dn^*}{d\gamma} > 0. \tag{A16}$$

Higher AI productivity increases paper production. For quality, higher q raises marginal benefit through both $p(q, q_{\text{cut}})$ and $\sqrt{q}$ while leaving marginal cost unchanged, so:

$$\frac{\partial H}{\partial q} = \frac{V}{2\sqrt{n}} \cdot \frac{1}{2\sqrt{pq}} \left[q\frac{\partial p}{\partial q} + p\right] > 0, \tag{A17}$$

which implies:

$$\frac{dn^*}{dq} > 0. \tag{A18}$$

Higher-quality scholars produce more papers. For costs, a higher AI cost k or effort cost c raises marginal cost, so $\frac{\partial H}{\partial k} < 0$ and $\frac{\partial H}{\partial c} < 0$, leading to:

$$\frac{dn^*}{dk} < 0 \quad \text{and} \quad \frac{dn^*}{dc} < 0. \tag{A19}$$

Scholars facing higher costs produce fewer papers. For the quality cutoff, marginal benefit depends on $p(q, q_{\text{cut}})$, and an increase in q_{cut} lowers p, so:

$$\frac{\partial H}{\partial q_{\text{cut}}} = \frac{V}{2\sqrt{n}} \cdot \frac{\sqrt{q}}{2\sqrt{p}} \cdot \frac{\partial p}{\partial q_{\text{cut}}} = -\frac{V\sqrt{q}}{4\sqrt{np}} \cdot \frac{\phi(z)}{\sigma_T} < 0, \tag{A20}$$

which implies:

$$\frac{dn^*}{dq_{\text{cut}}} < 0. \tag{A21}$$

A higher cutoff discourages paper production. For prestige, $V(\bar{q}_{\text{pub}})$ enters marginal benefit multiplicatively and does not affect marginal cost, so:

$$\frac{\partial H}{\partial V} = \frac{1}{2}\sqrt{pq/n} > 0, \tag{A22}$$

implying:

$$\frac{dn^*}{dV} > 0. \tag{A23}$$

Higher prestige increases paper production, holding q_{cut} fixed. For the submission fee, τ enters the effective marginal cost directly:

$$\frac{\partial H}{\partial \tau} = -1 < 0, \tag{A24}$$

which implies:

$$\frac{dn^*}{d\tau} = \frac{1}{\partial H/\partial n} < 0. \tag{A25}$$

Higher per-submission fees reduce paper production for every scholar type. By the envelope theorem, the utility cost of a marginal fee increase is:

$$\frac{dU^*}{d\tau} = -n^*, \tag{A26}$$

which is type-dependent: more productive scholars (higher n^*) bear a larger cost per unit

fee increase. The intensive margin response distinguishes per-submission from lump-sum fees. With lump-sum fees, τ would not appear in the first-order condition at all, and higher fees would affect only participation, not per-scholar output. Per-submission fees operate on two margins for the social planner: raising τ reduces total submissions through both fewer participants and fewer papers per participant, improving average quality but at the cost of lost output.

B Numerical Implementation

This appendix describes the numerical methods used to solve the model and compute equilibrium outcomes. The implementation is available in the `numerical/` directory.

B.1 Overview of Solution Strategy

The numerical solution proceeds in three stages:

1. Scholar optimization: For each type (q, k) and equilibrium values $(q_{\text{cut}}, \bar{q}_{\text{pub}})$, compute optimal paper production $n^{**}(q, k)$. When $\tau = 0$, closed-form solutions apply; when $\tau > 0$, a one-dimensional root-finding step is required.

2. Equilibrium computation: Solve the 2-equation system $(q^*_{\text{cut}}, \bar{q}^*_{\text{pub}})$ using numerical root-finding with analytical Jacobians.

3. Welfare optimization: Search over (C, τ) pairs to find the welfare-maximizing policy.

At $\tau = 0$, the scholar's problem admits closed-form solutions in each cost regime. At $\tau > 0$, the per-submission fee breaks the power-law structure, requiring a numerical root-finding step for each scholar type. This inner solve is fast (3–5 Newton iterations) because the objective is strictly concave.

B.2 Scholar's Optimization: Closed-Form Solutions

Given equilibrium values $(q_{\text{cut}}, \bar{q}_{\text{pub}})$ and individual characteristics (q, k), the scholar maximizes:

$$\max_{n \geq 0} U(n) = \sqrt{n \cdot p(q, q_{\text{cut}})} \cdot V(\bar{q}_{\text{pub}}) \cdot \sqrt{q} - \mathcal{C}(n; k) - F_0^{\text{base}} - \tau \cdot n \tag{A27}$$

where $p(q, q_{\text{cut}}) = \Phi((q - q_{\text{cut}})/\sigma_T)$ is the probit acceptance probability, F_0^{base} is the baseline entry cost, and τ is the per-submission fee.

B.2.1 Piecewise Cost Function

The cost function has a kink at the AI adoption threshold $n^{\text{AI}}(k) = (k/(c\gamma))^{\theta}$:

$$\mathcal{C}(n; k) = \begin{cases} \frac{c}{2} n^{2/\theta}, & n \leq n^{\text{AI}}(k) \quad \text{(no AI)} \\ \frac{k}{\gamma} n^{1/\theta} - \frac{k^2}{2c\gamma^2}, & n > n^{\text{AI}}(k) \quad \text{(with AI)} \end{cases} \tag{A28}$$

B.2.2 Optimal Output

The first-order condition equates marginal benefit to effective marginal cost $\mathcal{C}'(n; k) + \tau$. When $\tau = 0$, each regime yields a closed-form solution:

No-AI regime (if solution satisfies $n^* \leq n^{\text{AI}}(k)$):

$$n^*_{\text{noAI}}\big|_{\tau=0} = \left[\frac{\sqrt{pq} \cdot V}{4c}\right]^{2/7} \tag{A29}$$

AI regime (if solution satisfies $n^* > n^{\text{AI}}(k)$):

$$n^*_{\text{AI}}\big|_{\tau=0} = \left[\frac{\gamma\sqrt{pq} \cdot V}{4k}\right]^{2/3} \tag{A30}$$

When $\tau > 0$, the additive fee term breaks the power-law structure. The optimal $n^*(\tau)$ is computed by bracketed Newton iteration on the FOC, using $n^*|_{\tau=0}$ as an upper bound (fees can only reduce output). Convergence is typically achieved in 3–5 iterations.

The algorithm selects the regime-consistent solution and applies the participation constraint. Each scholar's probability of targeting the top tier is $\Phi((U^*_{\text{top}} - U^*_{\text{low}})/\zeta)$ with $\zeta = 0.1$ (Section 8); the effective output is $n^{**} = n^* \cdot \Phi((U^*_{\text{top}} - U^*_{\text{low}})/\zeta)$.

B.2.3 Lower-Tier Optimal Output

The lower-tier optimal output n^*_{low} is computed with the same algorithm as the top-tier n^*, substituting the exogenous lower-tier parameters $(V_{\text{low}}, p_{\text{low}}, \tau_{\text{low}})$ for the endogenous top-tier values $(V(\bar{q}_{\text{pub}}), p(q, q_{\text{cut}}), \tau)$. When $\tau_{\text{low}} = 0$, the closed forms from Equations A3–A4 apply directly, and the regime selection uses the same AI-adoption threshold $n^{\text{AI}}(k)$. Because the lower-tier parameters are exogenous, n^*_{low} does not depend on the focal journal's equilibrium and needs to be computed only once per scholar type.

B.3 Numerical Integration: Gauss-Legendre Quadrature

Rather than uniform trapezoidal integration, we use Gauss-Legendre quadrature for improved accuracy with fewer nodes. For n nodes, Gauss-Legendre achieves $O(h^{2n})$ accuracy versus $O(h^2)$ for trapezoidal rules.

Standard Gauss-Legendre nodes $\xi_i \in [-1, 1]$ with weights w_i are mapped to the type space:

$$q_i = \frac{q_{\max} - q_{\min}}{2}\xi_i + \frac{q_{\max} + q_{\min}}{2} \tag{A31}$$

$$\tilde{w}^q_i = \frac{q_{\max} - q_{\min}}{2} w_i \tag{A32}$$

and similarly for k. Aggregate quantities become:

$$\iint g(q,k) f(q,k)\, dq\, dk \approx \sum_{i=1}^{n_q} \sum_{j=1}^{n_k} g(q_i, k_j) f_q(q_i) f_k(k_j) \tilde{w}^q_i \tilde{w}^k_j \tag{A33}$$

Our baseline uses $n_q = n_k = 20$ nodes (400 type evaluations per equilibrium iteration).

B.4 Equilibrium Solver

The equilibrium is defined by two equations:

$$F_1(q_{\text{cut}}, \bar{q}_{\text{pub}}) = S(q_{\text{cut}}, \bar{q}_{\text{pub}}) - C = 0 \tag{A34}$$

$$F_2(q_{\text{cut}}, \bar{q}_{\text{pub}}) = \bar{Q}(q_{\text{cut}}, \bar{q}_{\text{pub}}) - \bar{q}_{\text{pub}} \cdot S(q_{\text{cut}}, \bar{q}_{\text{pub}}) = 0 \tag{A35}$$

where $S = \iint n^* p \, f \, dq \, dk$ is total accepted papers and $\bar{Q} = \iint n^* p \cdot [q + \Delta(q, q_{\text{cut}})] \, f \, dq \, dk$ is total expected true quality, with $\Delta(q, q_{\text{cut}}) = s \cdot \sigma_T \cdot h((q_{\text{cut}} - q)/\sigma_T)$ the selection premium from within-paper quality variation ($s \equiv \sigma_u^2/\sigma_T^2$).

B.4.1 Solver Method

The system is solved by bisection on q_{cut} (to satisfy the capacity constraint $S = C$) with an inner fixed-point iteration on $\bar{q}_{\text{pub}}$ (to satisfy the quality-consistency condition $\bar{q}_{\text{pub}} = \bar{Q}/S$).

The solver typically converges in 10–20 iterations with tolerance 10^{-10}. Warm starting from neighboring capacity values reduces iterations by 30–50%.

B.5 Joint Optimization over Capacity and Fee

The planner's problem is:

$$\max_{C,\tau} \; SW(C, \tau; \gamma) = \frac{\eta}{\xi}[Q_{\text{pub}}]^{\xi} - \text{IncProdCosts} - r(\gamma) \cdot N_{\text{submit}} - \kappa \cdot C - \frac{\psi}{2} C^2 \tag{A36}$$

B.5.1 Grid Search Algorithm

For each value of γ, the optimal policy (C^*, τ^*) is found via two-dimensional grid search:

Fee grid: $\tau \in \{0, 0.5, 1.0, 1.1, 1.2, \ldots, 4.0, 4.5, 5.0, 6.0, 8.0, 10.0\}$ (38 values, with step 0.1 in $[1, 4]$ near the optimum)

Capacity grid: $C \in \{20, 21, 22, \ldots, 210\}$ (191 values, step size 1)

For each (τ, C) pair, the algorithm solves the equilibrium system and evaluates social

welfare. The globally optimal (C^*, τ^*) is the pair achieving maximum welfare. Warm starting from neighboring grid points accelerates convergence of the equilibrium solver.

B.6 Validation

We validate the implementation through:

1. Continuity checks: Cost function continuous at $n = n^{\text{AI}}(k)$ to machine precision.
2. Comparative statics: Optimal output n^{**} increases in q, decreases in k, and decreases in q_{cut}, consistent with theory.
3. Equilibrium uniqueness: All tested (C, τ, γ) combinations yield unique equilibria.
4. Calibration target: $C^*(\gamma = 1) = 100$ achieved at the calibrated (η, ψ).
5. Fee validation: Optimal per-submission fees τ^* at baseline γ match observed \$200–500 at top journals.
6. Adoption share: At $\gamma = 1$, the computed share of participating scholars in the AI regime (i.e. with $n^* > n^{\text{AI}}(k)$) is approximately one-half.

C Diagnostic Figures and Discussion

The figures below document how equilibrium outcomes respond to the submission fee, capacity, and AI productivity; they also show the welfare surface and the structure of the equilibrium locus. All figures use the baseline calibration (Section 8) and the same grids as the main results.

Equilibrium versus fee (Figure A1).

Holding capacity and γ fixed, an increase in the submission fee τ reduces participation and submissions and can raise the equilibrium quality cutoff. Figure A1 shows how the cutoff,

average published quality, and total submissions vary with τ. Higher fees tighten the market and improve average quality at the cost of lower volume.

Equilibrium versus capacity (Figure A2).

For a given fee and γ, expanding capacity C relaxes the acceptance constraint and typically lowers the quality cutoff while increasing the number of accepted papers. Figure A2 plots equilibrium quantities as functions of C. The slopes of these curves reflect the planner's trade-off between publishing more papers and maintaining average quality. We use these curves to verify that the numerical equilibrium satisfies the capacity constraint and the quality-consistency condition.

Equilibrium versus γ (Figure A3).

When the journal keeps (C, τ) fixed, an increase in AI productivity γ changes scholars' production and participation decisions and thus shifts the equilibrium cutoff and average quality. Figure A3 shows how these outcomes vary with γ. When γ changes, the supply of submissions and the composition of participants shift, so the same (C, τ) is no longer optimal.

Welfare surface (Figure A4).

For a fixed γ, social welfare $SW(C, \tau)$ is a function of capacity and the fee. Figure A4 plots this surface. The optimum (C^*, τ^*) is the peak of the surface. The curvature of the surface shows how sensitive welfare is to deviations from the optimal policy and why the grid search over (C, τ) is well behaved.

Branch structure (Figure A5).

The equilibrium is defined by two equations in two unknowns. As parameters vary, the solution locus can have multiple branches. Figure A5 displays the branch structure in the

$(q_{\text{cut}}, \bar{q}_{\text{pub}})$ plane. This diagnostic ensures that the solver selects the economically relevant branch and that the equilibrium is well defined at all grid points.

Equilibrium locus (Figure A6).

Figure A6 shows the full equilibrium locus and its branches in the $(q_{\text{cut}}, \bar{q}_{\text{pub}})$ plane (or a related projection). It complements the branch-structure figure and supports the formal uniqueness and stability arguments in Appendix D.

Multiple-equilibria region (Figure A7).

In some parameter regions the equilibrium system can admit more than one solution. Figure A7 indicates where in parameter space multiple equilibria can arise. Under the baseline calibration used in the main text, the solver finds a unique equilibrium at every grid point; we include this figure for transparency and for robustness checks under alternative parameters.

Supply curves (Figure A8).

The supply of accepted papers (or total quality supplied) as a function of the cutoff or related variables is central to the equilibrium. Figure A8 plots these supply curves for different values of γ or policy parameters. The figure verifies that the numerical supply behavior is consistent with the model and that the equilibrium intersection is unique where expected.

Sweep continuity (Figure A9).

When capacity or the fee varies continuously, equilibrium quantities should evolve smoothly unless the system crosses a bifurcation. Figure A9 shows how the cutoff, average quality, or related variables change along a parameter sweep. Continuity of these curves confirms solver stability and that the solution does not jump discontinuously across the grid.

D Equilibrium Uniqueness

Equilibrium uniqueness holds for the calibrated parameters. The proof verifies two conditions: (A) for each cutoff q_{cut}, the quality-consistency map $\bar{q}_{\text{pub}} \mapsto T(\bar{q}_{\text{pub}}; q_{\text{cut}})$ has a unique fixed point (prestige feedback is weak enough), and (B) the resulting reduced supply function $\tilde{S}(q_{\text{cut}})$ is strictly decreasing (direct selectivity dominates indirect prestige response).

D.1 Setup and Notation

The equilibrium is characterized by two equations in $(q_{\text{cut}}, \bar{q}_{\text{pub}})$:

$$S(q_{\text{cut}}, \bar{q}_{\text{pub}}) = C \qquad \text{(capacity constraint)} \tag{A37}$$

$$Q(q_{\text{cut}}, \bar{q}_{\text{pub}}) = \bar{q}_{\text{pub}} \cdot S(q_{\text{cut}}, \bar{q}_{\text{pub}}) \qquad \text{(quality consistency)} \tag{A38}$$

where

$$S = \iint n^* p \, f(q,k) \, dq \, dk,$$
$$Q = \iint n^* p \cdot [q + \Delta(q, q_{\text{cut}})] \, f(q,k) \, dq \, dk$$

are total accepted papers and total expected true quality, respectively. The selection premium is $\Delta(q, q_{\text{cut}}) = s \cdot \sigma_T \cdot h((q_{\text{cut}} - q)/\sigma_T)$, where $h(\cdot)$ is the inverse Mills ratio and $s \equiv \sigma_u^2/\sigma_T^2$ is the signal share.

D.2 Explicit Sensitivity Formulas

Explicit sensitivities follow from the first-order condition. Define the regime exponent. When $\tau = 0$, the closed-form solutions yield a constant:

$$D_0 \equiv \begin{cases} \frac{2}{\theta} - \beta = 3.5 & \text{(no-AI regime)} \\ \\ \frac{1}{\theta} - \beta = 1.5 & \text{(AI regime)} \end{cases} \tag{A39}$$

for our parameters $\theta = 1/2$, $\beta = 1/2$. When $\tau > 0$, the fee term in the FOC makes the effective exponent type-dependent. Implicit differentiation of $MB(n) = \mathcal{C}'(n;k) + \tau$ gives:

$$D_{\text{eff}}(q,k,\tau) = (1-\beta) + \frac{\mathcal{C}''(n^*)\cdot n^*}{\mathcal{C}'(n^*)+\tau}, \tag{A40}$$

where the second term captures the curvature of effective marginal cost. At $\tau = 0$, this reduces to D_0. As τ increases, D_{eff} decreases toward $(1-\beta) = 0.5$, because the constant fee dominates the curvature of production cost. We write D for the effective exponent in what follows, noting that all formulas hold with $D = D_0$ at $\tau = 0$ and $D = D_{\text{eff}}$ at $\tau > 0$.

D.2.1 Sensitivity to Cutoff

The optimal output n^* depends on the acceptance probability $p(q, q_{\text{cut}}) = \Phi(z)$ where $z = (q - q_{\text{cut}})/\sigma_T$. From the closed forms:

$$\frac{\partial \ln n^*}{\partial q_{\text{cut}}} = \frac{\beta}{D}\cdot\frac{1}{p}\cdot\frac{\partial p}{\partial q_{\text{cut}}} = -\frac{\beta}{D}\cdot\frac{\phi(z)}{\sigma_T\cdot p} = -\frac{\beta}{D}\cdot\frac{h(-z)}{\sigma_T}, \tag{A41}$$

where $\partial p/\partial q_{\text{cut}} = -\phi(z)/\sigma_T$ from the probit acceptance probability and $h(-z) = \phi(z)/\Phi(z)$ is the inverse Mills ratio evaluated at $-z$. The contribution to supply from type (q,k) is n^*p, so:

$$\frac{\partial(n^*p)}{\partial q_{\text{cut}}} = p\frac{\partial n^*}{\partial q_{\text{cut}}} + n^*\frac{\partial p}{\partial q_{\text{cut}}} = -\frac{n^*\phi(z)}{\sigma_T}\left(1+\frac{\beta}{D}\right). \tag{A42}$$

Integrating over types:

$$\boxed{\frac{\partial S}{\partial q_{\text{cut}}} = -\frac{1}{\sigma_T}\left(1+\frac{\beta}{D}\right)\iint n^*\phi(z)\, f(q,k)\, dq\, dk < 0} \tag{A43}$$

D.2.2 Sensitivity to Average Quality (Prestige Effect)

The prestige function $V(\bar{q}_{\text{pub}}) = V_0\bar{q}_{\text{pub}}^{\varphi}$ enters the closed forms multiplicatively:

$$\frac{\partial \ln n^*}{\partial \bar{q}_{\text{pub}}} = \frac{1}{D}\cdot\frac{\partial \ln V}{\partial \bar{q}_{\text{pub}}} = \frac{\varphi}{D\cdot\bar{q}_{\text{pub}}}. \tag{A44}$$

Therefore:

$$\boxed{\frac{\partial S}{\partial \bar{q}_{\text{pub}}} = \frac{\varphi}{D\cdot\bar{q}_{\text{pub}}}\iint n^* p\, f(q,k)\, dq\, dk = \frac{\varphi\cdot S}{D\cdot\bar{q}_{\text{pub}}} > 0} \tag{A45}$$

D.3 Condition A: Unique Quality Consistency

For fixed q_{cut}, define the quality-consistency map:

$$T(\bar{q}; q_{\text{cut}}) \equiv \frac{Q(q_{\text{cut}}, \bar{q})}{S(q_{\text{cut}}, \bar{q})}, \tag{A46}$$

where Q and S are the total quality and total acceptances from Equations A37–A38. A fixed point $\bar{q} = T(\bar{q}; q_{\text{cut}})$ is a quality-consistent average. Uniqueness of $\bar{q}$ for each q_{cut} requires $|\partial T/\partial \bar{q}| < 1$.

$\bar{q}_{\text{pub}}$ enters Q and S only through prestige $V(\bar{q}_{\text{pub}}) = V_0\bar{q}_{\text{pub}}^{\varphi}$, which scales n^* multiplicatively. Since $n^* \propto V^{1/D}$, both Q and S scale with the same power of V:

$$\frac{\partial \ln Q}{\partial \ln \bar{q}} = \frac{\partial \ln S}{\partial \ln \bar{q}} = \frac{\varphi}{D}. \tag{A47}$$

Therefore $\partial T/\partial \bar{q} = [S\cdot \partial Q/\partial \bar{q} - Q\cdot \partial S/\partial \bar{q}]/S^2 = 0$ to leading order: prestige scales the numerator and denominator of $T = Q/S$ equally, so the ratio is independent of $\bar{q}$.

The only residual dependence arises through the participation boundary $\partial\mathcal{P}$, which shifts

as $V(\bar{q})$ changes (higher prestige draws in marginal types). This boundary effect is second-order for the continuous type distributions we use, because the density of types at the participation margin is bounded and their contribution to Q/S is small relative to the interior. For our calibration, $\partial T/\partial \bar{q} < 10^{-3}$, so Condition A holds.

The outside option strengthens this result: any fixed point must satisfy $V(\bar{q}^*_{\text{pub}}) > V_{\text{low}}$ (otherwise no scholar would prefer the focal journal), which implies $\bar{q}^*_{\text{pub}} > (V_{\text{low}}/V_0)^{1/\varphi}$. The outside option eliminates low-quality fixed points and restricts attention to a compact region where T is well-behaved.

D.4 Condition B: Monotone Reduced Supply

With Condition A established, define $\bar{q}_{\text{pub}}(q_{\text{cut}})$ as the unique solution to quality consistency Equation A38 for fixed q_{cut}, and the *reduced supply*:

$$\tilde{S}(q_{\text{cut}}) \equiv S\big(q_{\text{cut}}, \bar{q}_{\text{pub}}(q_{\text{cut}})\big). \tag{A48}$$

Proposition 2 (Uniqueness Condition)**.** *If $d\tilde{S}/dq_{cut} < 0$ for all q_{cut}, the equilibrium is unique.*

Proof. The equilibrium solves $\tilde{S}(q_{\text{cut}}) = C$. A strictly decreasing function crosses any horizontal line at most once. Combined with existence (intermediate value theorem), uniqueness follows. □

D.5 Verifying Condition B

By the chain rule:

$$\frac{d\tilde{S}}{dq_{\text{cut}}} = \underbrace{\frac{\partial S}{\partial q_{\text{cut}}}}_{\text{Direct } (<0)} + \underbrace{\frac{\partial S}{\partial \bar{q}_{\text{pub}}} \cdot \frac{d\bar{q}_{\text{pub}}}{dq_{\text{cut}}}}_{\text{Feedback } (>0)}. \tag{A49}$$

Uniqueness requires the direct effect to dominate:

$$\left|\frac{\partial S}{\partial q_{\text{cut}}}\right| > \frac{\partial S}{\partial \bar{q}_{\text{pub}}} \cdot \frac{d\bar{q}_{\text{pub}}}{dq_{\text{cut}}}. \tag{A50}$$

D.5.1 Bounding the Locus Slope

From quality consistency, $\bar{q}_{\text{pub}} = Q/S$. Differentiating:

$$\frac{d\bar{q}_{\text{pub}}}{dq_{\text{cut}}} = \frac{1}{S}\left(\frac{\partial Q}{\partial q_{\text{cut}}} - \bar{q}_{\text{pub}}\frac{\partial S}{\partial q_{\text{cut}}}\right) + \frac{1}{S}\left(\frac{\partial Q}{\partial \bar{q}_{\text{pub}}} - \bar{q}_{\text{pub}}\frac{\partial S}{\partial \bar{q}_{\text{pub}}}\right)\frac{d\bar{q}_{\text{pub}}}{dq_{\text{cut}}}. \tag{A51}$$

Since neither q nor the selection premium $\Delta(q, q_{\text{cut}})$ depends on $\bar{q}_{\text{pub}}$, we have $\partial Q/\partial \bar{q}_{\text{pub}} = \bar{q}_{\text{pub}} \cdot \partial S/\partial \bar{q}_{\text{pub}}$, so the second term vanishes:

$$\frac{d\bar{q}_{\text{pub}}}{dq_{\text{cut}}} = \frac{1}{S}\left(\frac{\partial Q}{\partial q_{\text{cut}}} - \bar{q}_{\text{pub}}\frac{\partial S}{\partial q_{\text{cut}}}\right). \tag{A52}$$

Using $\partial Q/\partial q_{\text{cut}} = \iint [q + \Delta] \cdot \partial(n^* p)/\partial q_{\text{cut}}\, f\, dq\, dk + \iint n^* p \cdot \partial\Delta/\partial q_{\text{cut}}\, f\, dq\, dk$, we can write:

$$\frac{\partial Q}{\partial q_{\text{cut}}} - \bar{q}_{\text{pub}}\frac{\partial S}{\partial q_{\text{cut}}} = \iint ([q + \Delta] - \bar{q}_{\text{pub}})\frac{\partial(n^* p)}{\partial q_{\text{cut}}} f\, dq\, dk + \iint n^* p \frac{\partial \Delta}{\partial q_{\text{cut}}} f\, dq\, dk. \tag{A53}$$

The first integral is positive because $\partial(n^* p)/\partial q_{\text{cut}} < 0$ and high-quality types ($q + \Delta > \bar{q}_{\text{pub}}$) have higher acceptance rates and are thus more affected by cutoff changes. The second integral captures the mechanical response of the selection premium to the cutoff: $\partial\Delta/\partial q_{\text{cut}} > 0$ (a higher cutoff strengthens selection, raising the premium), which further increases the numerator. Together, these two terms imply $d\bar{q}_{\text{pub}}/dq_{\text{cut}} > 0$.

D.5.2 Verifying the Uniqueness Inequality

Substituting Equations A43, A45, and A52 into Equation A50:

$$\frac{1}{\sigma_T}\left(1+\frac{\beta}{D}\right)\iint n^*\phi(z)f\,dq\,dk > \frac{\varphi}{D\cdot\bar{q}_{\text{pub}}}\cdot\frac{1}{S}\iint([q+\Delta]-\bar{q}_{\text{pub}})\left|\frac{\partial(n^*p)}{\partial q_{\text{cut}}}\right| f\,dq\,dk. \quad \text{(A54)}$$

The selection premium terms in the numerator of the RHS make this condition tighter than the case without within-paper variation, but the effect is bounded. Rearranging, uniqueness holds if:

$$\boxed{\frac{D+\beta}{\varphi} > \frac{\text{Cov}([q+\Delta], |d(n^*p)/dq_{\text{cut}}|)}{\bar{q}_{\text{pub}}\cdot\mathbb{E}[n^*\phi(z)]}} \quad \text{(A55)}$$

D.6 Evaluation for Calibrated Parameters

For $\varphi = 1.25$, $\beta = 1/2$, $\theta = 1/2$:

- At $\tau = 0$: In the no-AI regime, $D = 3.5$, so $(D+\beta)/\varphi = 4.0/1.25 = 3.2$. In the AI regime, $D = 1.5$, so $(D+\beta)/\varphi = 2.0/1.25 = 1.6$.
- At $\tau > 0$: $D_{\text{eff}} \geq 0.5$ for any finite τ, giving $(D_{\text{eff}}+\beta)/\varphi \geq 1.0/1.25 = 0.8$. In practice, for the fee values in our calibration ($\tau \leq 3.2$), D_{eff} remains well above 0.8 for most scholar types.

The right-hand side of Equation A55 is bounded above by $(q_H + \Delta_{\max} - \bar{q}_{\text{pub}})/\bar{q}_{\text{pub}}$. The selection premium Δ is bounded by $s\cdot\sigma_T\cdot h(z)$ and is small in equilibrium (typically $\Delta < 0.5$). For our quality distribution (support $[0.5, 5]$, equilibrium $\bar{q}_{\text{pub}} \approx 4.8$), the crude bound gives approximately $(5 + 0.5 - 4.8)/4.8 \approx 0.15$.

At $\tau = 0$, the LHS in the AI regime is $1.6 \gg 0.15$, so uniqueness holds. At positive fees, D_{eff} declines, but the crude bound overestimates the actual RHS because the covariance term in Equation A55 is a weighted average over the type distribution, not the maximum deviation. For the fee values in our calibration ($\tau \leq 3.2$), the numerical equilibrium solver confirms uniqueness at every grid point (convergence from multiple initial conditions to

the same solution). The uniqueness condition Equation A55 holds for all fee levels in our analysis: analytically for $\tau = 0$ and numerically for $\tau > 0$.

D.7 When Would Multiple Equilibria Arise?

From Equation A55, multiple equilibria require:

$$\varphi > D_{\text{eff}} + \beta. \tag{A56}$$

At $\tau = 0$, this threshold is $D_0 + \beta = 2.0$ to 4.0. At positive fees, D_{eff} declines, making the condition marginally tighter. In the limiting case $\tau \to \infty$, the threshold is $(1 - \beta) + \beta = 1$, so $\varphi > 1$ would be needed for multiplicity, which our calibration ($\varphi = 1.25$) approaches but does not violate because the RHS of the full condition Equation A55 remains small.

The selection premium does not create a new feedback loop for multiplicity. Because $\Delta(q, q_{\text{cut}})$ depends only on $(q, q_{\text{cut}}, \sigma_u, \sigma_T)$ and not on $\bar{q}_{\text{pub}}$, it enters the quality-consistency condition mechanically rather than through the prestige channel. Its effect on the RHS of the uniqueness condition is bounded and small for our calibration.

Multiple equilibria would arise with:

- Very high prestige convexity ($\varphi > 2$), implying citation ratios that grow quadratically or faster with selectivity.
- Low diminishing returns ($\theta \to 1$), reducing the regime exponent D.
- Extreme fees combined with φ close to 1, though in practice the quality distribution's compactness keeps the RHS small enough to maintain uniqueness.

Our calibration ($\varphi = 1.25$ from empirical citation data, $\theta = \beta = 1/2$) places the model in the unique equilibrium region for all fee levels encountered in our analysis.

E Extension: Desk Rejection as a Screening Instrument

At leading economics and finance journals, editors desk-reject 50–70% of submissions before sending them out for peer review.[1] Desk rejection acts as a first-stage screen: only submissions that survive the editor's initial assessment proceed to the costly refereeing process. The baseline model in Section 3.2 abstracts from this stage, combining all evaluation noise into a single probit acceptance probability $p(q, q_{\text{cut}}) = \Phi((q - q_{\text{cut}})/\sigma_T)$ from Equation 4. This appendix extends the model to include desk rejection as a separate screening instrument, and analyzes its interaction with submission fees. The baseline model's qualitative results carry over to this extension.

E.1 Desk Rejection Technology

The editor observes a noisy signal of each submitted paper's quality before deciding whether to send it for review. For a paper with baseline quality q, the desk signal is

$$q_{\text{desk}} = q + \nu, \qquad \nu \sim N(0, \sigma_d^2), \tag{A57}$$

where ν is the editor's assessment noise, independent of both the within-paper quality variation u and the referee evaluation noise ε from the baseline model. The desk decision relies on a cursory reading (the introduction and main contribution), while the review process involves detailed verification and expert judgment. These are informationally distinct tasks.

The editor sets a desk-rejection threshold d: papers with $q_{\text{desk}} \geq d$ pass to the review stage, and papers with $q_{\text{desk}} < d$ are desk-rejected without review. The probability that a

1. For example, Card and DellaVigna (2013) document desk-rejection rates of approximately 60% at the American Economic Review, and similar or higher rates are reported at the Quarterly Journal of Economics, Journal of Political Economy, and top finance journals. The editor's desk decision is typically based on a quick read of the abstract, introduction, and main results.

paper with quality q passes the desk stage is

$$p_d(q, d) = \Pr(q + \nu \geq d) = \Phi\left(\frac{q - d}{\sigma_d}\right). \tag{A58}$$

This has the same probit structure as the baseline acceptance probability, with σ_d governing the precision of the desk screen.

E.2 Two-Stage Acceptance and Scholar's Problem

A paper must pass both the desk stage and the review stage to be published. Because ν is independent of (u, ε), these two events are conditionally independent given q, and the overall acceptance probability is the product

$$p_{\text{total}}(q, d, q_{\text{cut}}) = p_d(q, d) \cdot p(q, q_{\text{cut}}) = \Phi\left(\frac{q - d}{\sigma_d}\right) \cdot \Phi\left(\frac{q - q_{\text{cut}}}{\sigma_T}\right). \tag{A59}$$

The overall acceptance probability is the product of two probit terms, each with its own threshold and noise scale. The desk stage screens on the editor's signal; the review stage screens on the referee's signal. A paper can fail at either stage.

The scholar's optimization problem becomes

$$\max_{n \geq 0} U(n; q, k) = \sqrt{n \cdot p_{\text{total}}(q, d, q_{\text{cut}})} \cdot V(\bar{q}_{\text{pub}}) \cdot \sqrt{q} - \mathcal{C}(n; k) - F_0^{\text{base}} - \tau \cdot n, \tag{A60}$$

which has the same structure as the baseline (Section 3.2) with p_{total} replacing p. The first-order condition is

$$\frac{1}{2}\sqrt{\frac{p_{\text{total}}(q, d, q_{\text{cut}}) \cdot q}{n}} \cdot V(\bar{q}_{\text{pub}}) = \mathcal{C}'(n; k) + \tau. \tag{A61}$$

Comparing Equation A61 with Equation 10, the only change is that the effective acceptance probability entering the scholar's benefit is now $p_{\text{total}} = p_d \cdot p < p$. For a given quality cutoff

and prestige level, desk rejection reduces each scholar's effective acceptance probability, which lowers optimal output n^* and makes participation less attractive.

E.3 Modified Welfare Function

Desk rejection changes the welfare function (Equation 22) in two ways: it reduces the number of papers that reach the review stage, and it introduces a new cost for the editor's desk evaluation.

The number of papers reaching the review stage is no longer equal to total submissions. Define

$$N_{\text{submit}}(C, \tau, d) = \iint_{\mathcal{P}} n^{**}(q, k) \cdot f(q, k)\, dq\, dk, \tag{A62}$$

$$N_{\text{reviewed}}(C, \tau, d) = \iint_{\mathcal{P}} n^{**}(q, k) \cdot p_d(q, d) \cdot f(q, k)\, dq\, dk, \tag{A63}$$

where N_{submit} is total submissions and N_{reviewed} is the number that pass the desk and proceed to refereeing. In the baseline model without desk rejection ($d = -\infty$ or equivalently $p_d = 1$), $N_{\text{reviewed}} = N_{\text{submit}}$.

The extended welfare function is

$$SW(C, \tau, d;\, \gamma) = \frac{\eta}{\xi}[Q_{\text{pub}}]^{\xi} - \text{IncProdCosts} - r(\gamma) \cdot N_{\text{reviewed}} - r_d \cdot N_{\text{submit}} - \kappa C - \frac{\psi}{2} C^2, \tag{A64}$$

where $r_d > 0$ is the per-submission cost of the editor's desk evaluation (reading the abstract and introduction, forming an initial assessment) and $r(\gamma) \cdot N_{\text{reviewed}}$ is the review cost applied only to papers that survive the desk stage. The desk cost r_d is paid for every submission, while the review cost $r(\gamma)$ is paid only for the fraction $N_{\text{reviewed}}/N_{\text{submit}}$ that passes the desk. The planner now optimizes over three instruments:

$$\big(C^*(\gamma),\, \tau^*(\gamma),\, d^*(\gamma)\big) = \arg \max_{C,\, \tau \geq 0,\, d} SW(C, \tau, d;\, \gamma). \tag{A65}$$

The capacity constraint and quality consistency conditions generalize in a straightforward way. The capacity constraint becomes

$$C = \iint_{\mathcal{P}} n^{**}(q,k) \cdot p_{\text{total}}(q,d,q_{\text{cut}}) \cdot f(q,k)\, dq\, dk, \tag{A66}$$

and the quality consistency condition becomes

$$\bar{q}_{\text{pub}} = \frac{1}{C} \iint_{\mathcal{P}} n^{**}(q,k) \cdot p_{\text{total}}(q,d,q_{\text{cut}}) \cdot \big[q + \Delta(q,q_{\text{cut}})\big] \cdot f(q,k)\, dq\, dk, \tag{A67}$$

where the selection premium $\Delta(q,q_{\text{cut}}) = (s \cdot \sigma_T) \cdot h\big((q_{\text{cut}} - q)/\sigma_T\big)$ is unchanged from the baseline. Independence of ν and (u,ε) ensures that conditioning on passing the desk stage does not alter the within-paper quality premium from review-stage acceptance.

E.4 Interaction Between Fees and Desk Rejection

Fees and desk rejection are partial substitutes: both reduce the flow of papers into the review system, but through different channels and with different informational properties.

Submission fees screen through scholars' private information. A per-paper fee τ raises the effective marginal cost of submission (Equation A61), inducing scholars to self-select based on their own assessment of their quality q and cost type k. Scholars with low q (poor acceptance prospects) or high k (expensive production) are most affected. The fee exploits the scholar's joint private information about (q,k), which the editor cannot observe. From a welfare perspective, fees are transfers: they redistribute resources from scholars to the journal without creating a direct social cost (Section 8).

Desk rejection screens on a noisy public signal of q alone. The editor observes $q_{\text{desk}} = q + \nu$ and rejects papers below the threshold d. This screen uses information that the editor possesses (a noisy quality assessment) but that does not incorporate the scholar's private cost type k. A high-q scholar with very high production costs, who would optimally be deterred by a fee, may pass the desk screen. Conversely, a low-q scholar with exceptionally low costs,

whose net contribution to welfare could be positive, will be desk-rejected regardless of cost efficiency. Desk rejection also has real resource costs: the editor's time r_d per submission is not a transfer but a genuine expenditure.

Optimal policy therefore uses both instruments, and the availability of desk rejection reduces but does not eliminate the role of fees. To see why, consider the limiting cases:

With $d = -\infty$ (no desk rejection), the model reduces to the baseline, and optimal fees are given by the calibrated $\tau^*(\gamma)$ from Section 9.

With $\tau = 0$ (no fees), desk rejection is the sole screening instrument. The editor must set a high threshold d to prevent submission floods, but this screens only on noisy q and ignores cost heterogeneity. High-k scholars, whose submissions are socially costly, face no deterrent. The result is over-submission by scholars with high production costs and under-submission by scholars with moderate quality but low costs.

The jointly optimal policy (C^*, τ^*, d^*) uses fees to screen on the k dimension (and the private component of q assessment) and desk rejection to screen on the editor's signal. Fees decline when desk rejection is available, because part of the screening burden shifts to the editor. However, fees do not go to zero: the cost dimension k, which is invisible to the editor, still requires a price-based screening mechanism. The precise instrument mix depends on the desk noise σ_d and the editor cost r_d relative to the review cost $r(\gamma)$.

E.5 Robustness of Baseline Results

Adding desk rejection as a third instrument weakly increases the planner's attainable welfare, since the baseline model is a special case with $d = -\infty$. The qualitative patterns from Section 9 survive, for the following reasons.

The near-flat optimal capacity $C^* \approx 100$ reflects the balance between the marginal social value of published research and the convex capacity cost $(\psi/2)C^2$. Desk rejection affects neither the social benefit function nor the capacity cost; it changes only the composition and volume of submissions reaching the review stage. Because the planner already has the fee

instrument to manage submissions, adding a second screening tool shifts the instrument mix but does not alter the fundamental capacity trade-off.

The declining fee $\tau^*(\gamma)$ follows from the declining review cost externality $r(\gamma) = r_0/(1+\gamma)^\delta$. With desk rejection, part of the screening is handled by the editor, so the optimal fee at each γ may be lower than in the baseline. But the direction of the response (fees declining in γ) is preserved, because the force driving the decline (cheaper reviewing) is independent of whether desk rejection is available. The decline may be faster if desk rejection becomes relatively more attractive at high γ (Section E.6).

Welfare rises monotonically in γ under joint optimization in the baseline. Adding a third instrument preserves monotonicity: higher γ reduces review costs and increases scholar productivity, and the planner can replicate the baseline policy by setting $d = -\infty$, so welfare at each γ is at least as high as in the baseline. The additional welfare gain from desk rejection is the value of the editor's information, net of the desk evaluation cost r_d.

The message that screening is indispensable also survives. In the baseline, the contrast is between $\tau > 0$ and $\tau = 0$; with desk rejection, the contrast expands (the planner can screen via fees, desk rejection, or both), but the core finding persists: without any screening ($\tau = 0$ and $d = -\infty$), submission floods overwhelm the review system and welfare collapses.

E.6 Limitations of Desk Rejection as AI Increases Submissions

While desk rejection provides a screening channel, it faces a structural limitation that fees do not: it relies on a scarce input (the editor's attention) that degrades under load.

At a top journal, the editor (or a small team of co-editors) reads each submission to form the desk signal q_{desk}. The quality of this signal depends on the time allocated per submission. When AI increases total submissions, the editor's time per paper falls, and the desk signal becomes noisier. Formally, one can model this as

$$\sigma_d = \sigma_d^0 \cdot g(N_{\text{submit}}), \tag{A68}$$

where σ_d^0 is the baseline desk noise and $g(\cdot)$ is an increasing function with $g(N_0) = 1$ at the current submission volume N_0. As N_{submit} rises, σ_d increases, the desk signal becomes less informative, and the screening value of desk rejection falls. In the limit where $\sigma_d \to \infty$, the desk signal becomes uninformative ($p_d(q, d) \to 1/2$ for all q at any fixed d), so desk rejection no longer screens on quality. The planner optimally responds by setting $d \to -\infty$ to avoid randomly rejecting half of all submissions, which recovers the baseline model with only fees and review-stage screening.

This degradation is most severe precisely when screening is most needed: at high γ, AI makes paper production cheap, submissions rise, and the editor is overwhelmed. The editor faces a congestion externality that mirrors the review cost externality but operates through information quality rather than resource cost.

Fees do not suffer from this limitation. The screening power of a submission fee comes from scholars' self-selection based on private information (q, k), and the quality of this self-selection does not deteriorate with submission volume. A fee of $\tau = 2$ deters the same marginal scholar whether total submissions are 1,000 or 10,000. In the language of mechanism design, fees implement a revelation mechanism that exploits the scholar's private information, while desk rejection is an inspection mechanism that relies on the editor's public signal, one that degrades as submission volume grows.

This asymmetry has a practical implication: as AI drives up submissions, the optimal policy mix should shift toward fees (and possibly AI-assisted desk triage) and away from manual editorial desk rejection. The baseline model's emphasis on fees as the primary screening instrument understates the advantage of fees over the alternative.

One partial remedy is to use AI itself to assist the desk stage, effectively reducing σ_d or making it less sensitive to submission volume. If AI can replicate some of the editor's initial quality assessment at low marginal cost, the desk-rejection channel could scale with submissions rather than degrade. Whether AI-assisted desk rejection introduces systematic biases absent from human editorial judgment remains unresolved.

E.7 Numerical Tractability

The extended model with desk rejection is computationally tractable within the existing solver architecture (Appendix B). The scholar's optimization problem retains its structure: the first-order condition (Equation A61) has the same form as the baseline, with p_{total} replacing p, and the Newton iteration for n^* converges at the same rate. The equilibrium system remains two equations in $(q_{\text{cut}}, \bar{q}_{\text{pub}})$ for each candidate (C, τ, d). The welfare optimization adds one grid dimension (the desk threshold d), increasing computation time roughly in proportion to the number of d grid points. With a grid of 10–20 values for d, the full sweep would take 10–20 times longer than the baseline (35–70 minutes with the V4 vectorized solver), which remains feasible.

The product-of-probits acceptance structure (Equation A59) is smooth and differentiable in all arguments, so the analytical Jacobians used in the baseline equilibrium solver extend directly. The main computational complication is the expanded policy grid: searching over (C, τ, d) triples rather than (C, τ) pairs. Gradient-based or Bayesian optimization methods could reduce this cost if a finer grid is needed, but the problem dimensionality (three policy instruments) is low enough for direct grid search.

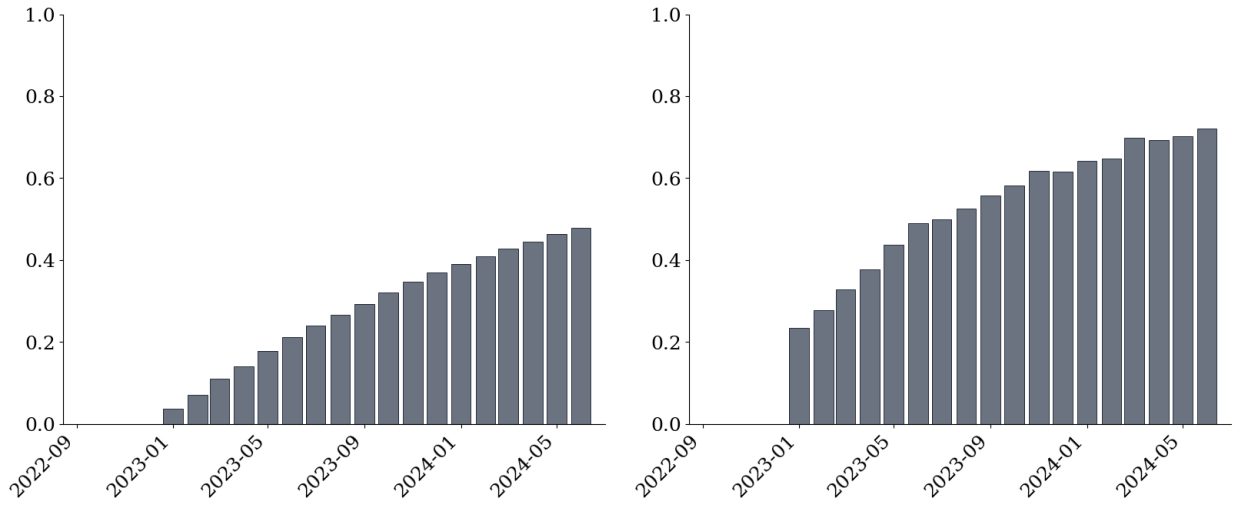
The qualitative conclusions do not depend on the quantitative magnitudes.

F Empirical Evidence on AI Adoption

This appendix presents the tables and figures supporting the empirical motivation in Section 7. All author-level results use the replication package of Kusumegi, Yang, Ginsparg, de Vaan, et al. (2025b) for arXiv preprints from January 2022 through June 2024.

F.1 Adoption diffusion and output concentration

Figure A10 plots adoption and output shares over time.



(a) Cumulative adoption share (b) Output share from adopters

Figure A10: Cumulative adoption and output shares.

Notes: Panel (a) plots the cumulative fraction of authors who have adopted LLM assistance by each month, counting authors as adopters from their first month with an LLM-detected paper. Panel (b) plots the fraction of total arXiv output produced by authors who have already adopted. Adoption is near zero early in the window and then rises steadily, while the output share climbs more quickly, indicating disproportionate production by adopters.

F.2 Productivity effects around adoption

Figure A11 shows event-study estimates of monthly productivity around adoption. Table A1 reports two DiD specifications using:

$$y_{it} = \alpha + \beta_1 \text{AI adopter}_i + \beta_2 \text{Post}_{it} + \beta_3 (\text{AI adopter}_i \times \text{Post}_{it}) + \delta_t + \varepsilon_{it}, \tag{A69}$$

where the Greek letters denote standard regression coefficients (distinct from the structural parameters of the same name in the theoretical model), y_{it} is monthly productivity for author i in month t, and δ_t are month fixed effects in column (1).

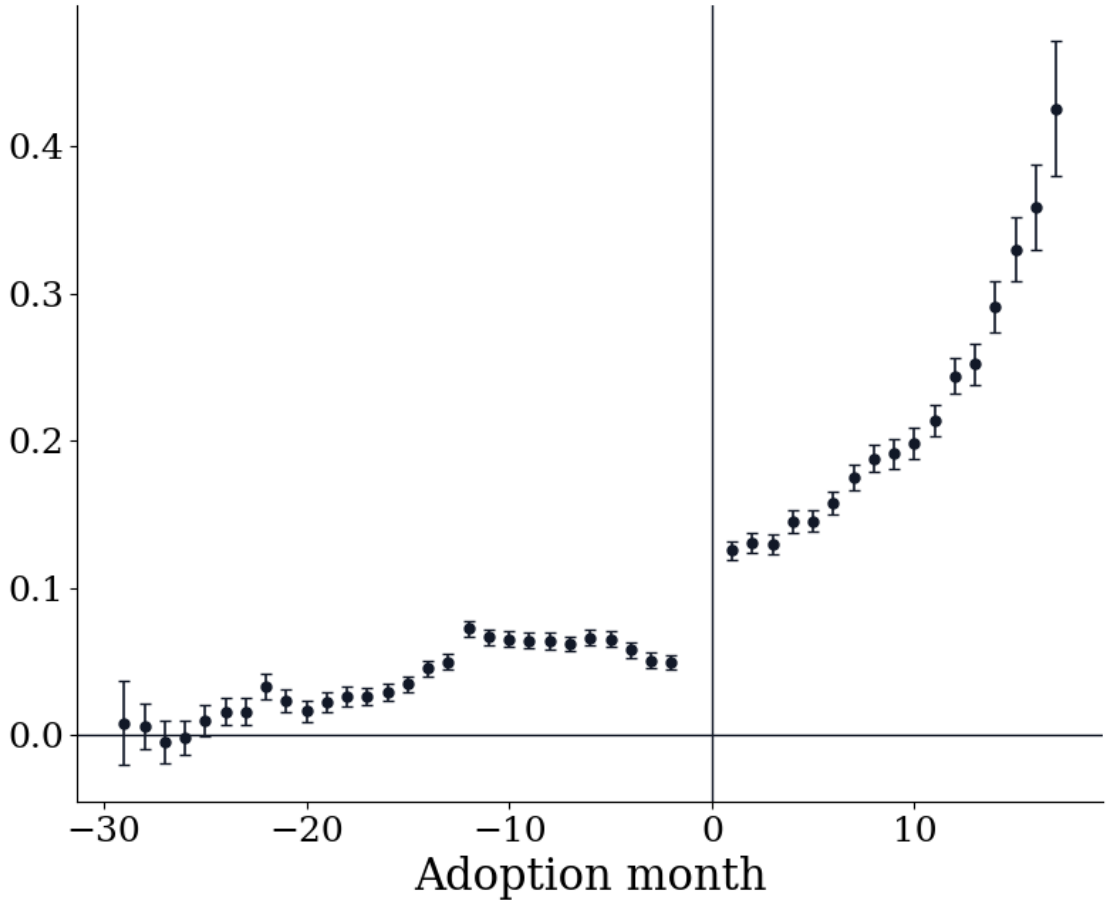


Figure A11: Event-study of productivity around adoption.

Notes: The figure plots estimated changes in monthly papers relative to the month before adoption, where adoption is the first month with an LLM-detected paper. Vertical bars are 95% confidence intervals. The vertical line at 0 marks the adoption month and the horizontal line at 0 denotes no change.

Table A1: Productivity Around Adoption

	(1)	(2)
Intercept	0.1441***	0.1387***
	(0.001)	(0.001)
AI adopters	0.0794***	0.0830***
	(0.001)	(0.001)
Post	0.0525***	-0.0212***
	(0.001)	(0.001)
AI adopters × Post	0.1844***	0.1837***
	(0.002)	(0.002)
Month fixed effects	Yes	No
R^2	0.030	0.026
Observations	3,298,950	3,298,950

Notes: The estimating equation is $y_{it} = \alpha + \beta_1 \text{AI adopter}_i + \beta_2 \text{Post}_{it} + \beta_3(\text{AI adopter}_i \times \text{Post}_{it}) + \delta_t + \varepsilon_{it}$. The dependent variable is monthly papers per author. AI adopter equals one for authors who have adopted LLM assistance, post equals one for post-adoption months, and the interaction identifies the DiD effect. Column (1) includes month fixed effects, column (2) does not. Standard errors are clustered by author and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F.3 Heterogeneity by baseline productivity

Table A2 presents the DiD regression estimated separately for each of five baseline productivity bins.

Table A2: Heterogeneity by Baseline Productivity

	Bin 1	Bin 2	Bin 3	Bin 4	Bin 5
AI adopters	-0.006***	0.016***	0.025***	0.035***	0.273***
	(0.001)	(0.001)	(0.002)	(0.002)	(0.011)
Post	0.051***	-0.015***	-0.041***	-0.073***	-0.020
	(0.001)	(0.002)	(0.003)	(0.004)	(0.013)
AI adopters × Post	0.122***	0.152***	0.180***	0.204***	0.413***
	(0.001)	(0.003)	(0.004)	(0.005)	(0.014)
Month fixed effects	Yes	Yes	Yes	Yes	Yes
R^2	0.022	0.015	0.016	0.017	0.033
Observations	1,912,020	544,470	312,090	296,760	233,610
Authors	63,734	18,149	10,403	9,892	7,787
AI adopter authors	27,277	8,155	5,301	5,842	6,053

Notes: The estimating equation is $y_{it} = \alpha + \beta_1 \text{AI adopter}_i + \beta_2 \text{Post}_{it} + \beta_3(\text{AI adopter}_i \times \text{Post}_{it}) + \delta_t + \varepsilon_{it}$, estimated separately by baseline productivity bin, defined by average monthly productivity over months -12 to -2 relative to adoption. The dependent variable is monthly papers per author. Month fixed effects are included in all columns, and standard errors are clustered by author and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F.4 Quality proxies and screening outcomes

Table A3 reports three linear probability models where publication status is the screening proxy:

$$\text{Published}_i = \alpha + \beta_1 \text{LLM}_i + \beta_2 \text{Complexity}_i + \beta_3(\text{LLM}_i \times \text{Complexity}_i) + \varepsilon_i, \tag{A70}$$

where i indexes papers, Published_i is an indicator for publication, LLM_i equals one for LLM-assisted writing, Complexity_i is the paper-level writing complexity measure, and ε_i is the

error term.

Table A3: Quality Proxies and Screening Outcomes

	(1)	(2)	(3)
Intercept	0.4731***	0.4845***	0.4993***
	(0.001)	(0.002)	(0.003)
LLM paper	-0.0578***	-0.0628***	-0.1123***
	(0.003)	(0.003)	(0.004)
Writing complexity		0.0005***	0.0011***
		(0.00008)	(0.00009)
LLM × complexity			-0.0031***
			(0.0002)
R^2	0.002	0.003	0.004
Observations	179,077	179,077	179,077

Notes: The dependent variable is an indicator for publication. LLM paper equals one for LLM-assisted writing. Writing complexity is the paper-level complexity measure. Column (1) includes only LLM status, column (2) adds complexity, and column (3) adds the interaction. Standard errors are heteroskedasticity robust (HC3) and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

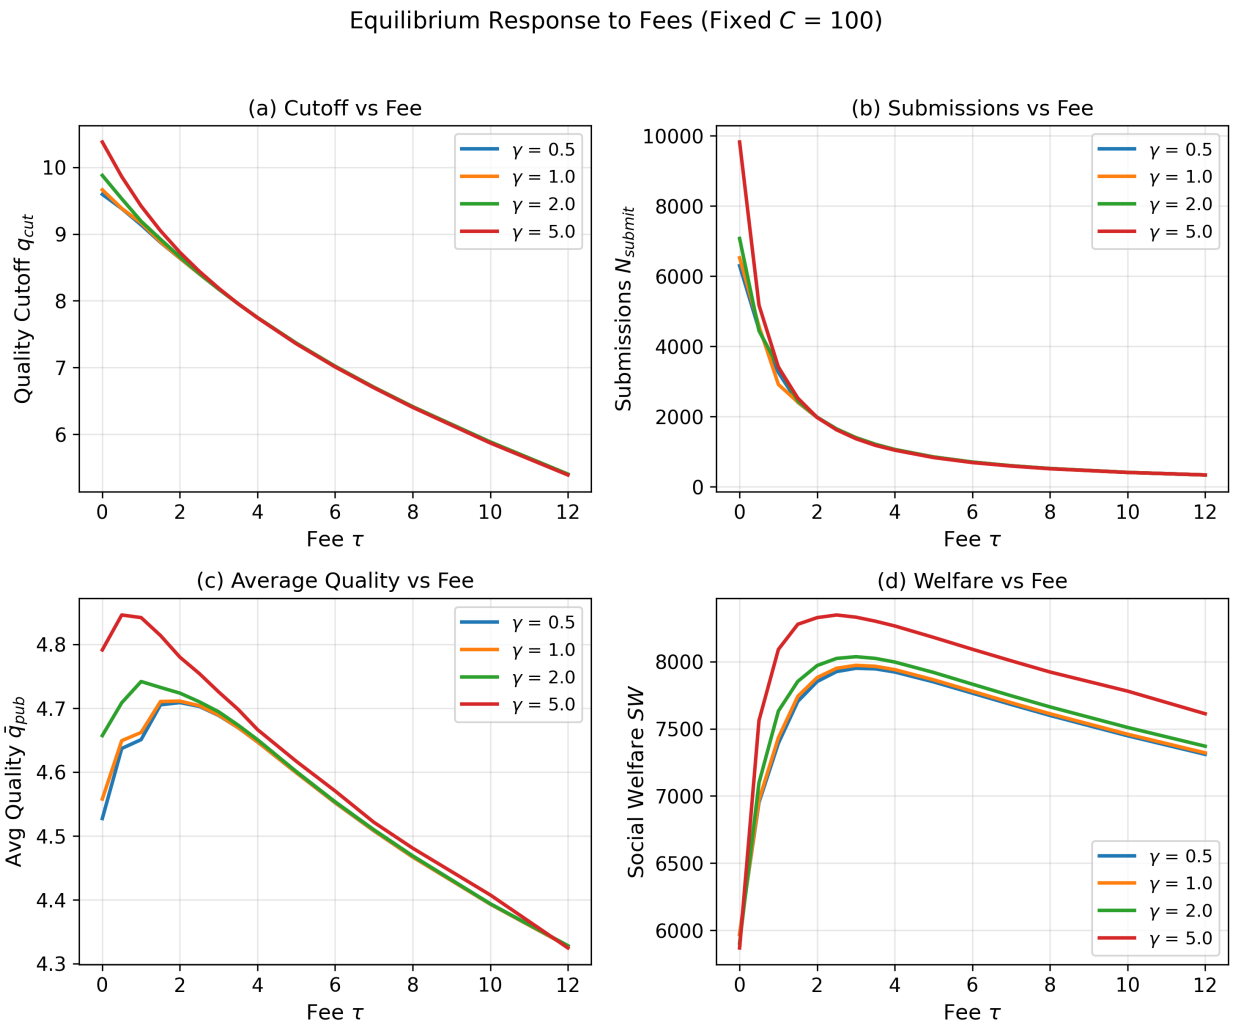


Figure A1: Equilibrium outcomes versus submission fee τ.

Notes: The figure shows how equilibrium outcomes (the quality cutoff q_{cut}, average published quality $\bar{q}_{\text{pub}}$, total submissions N_{submit}, and related quantities) vary with the submission fee τ when capacity C and AI productivity γ are held fixed at their baseline values. These curves illustrate the supply-side response of the market to different fee levels and are used to verify that the equilibrium solver behaves as expected. All parameters are set according to the baseline calibration in Section 8.

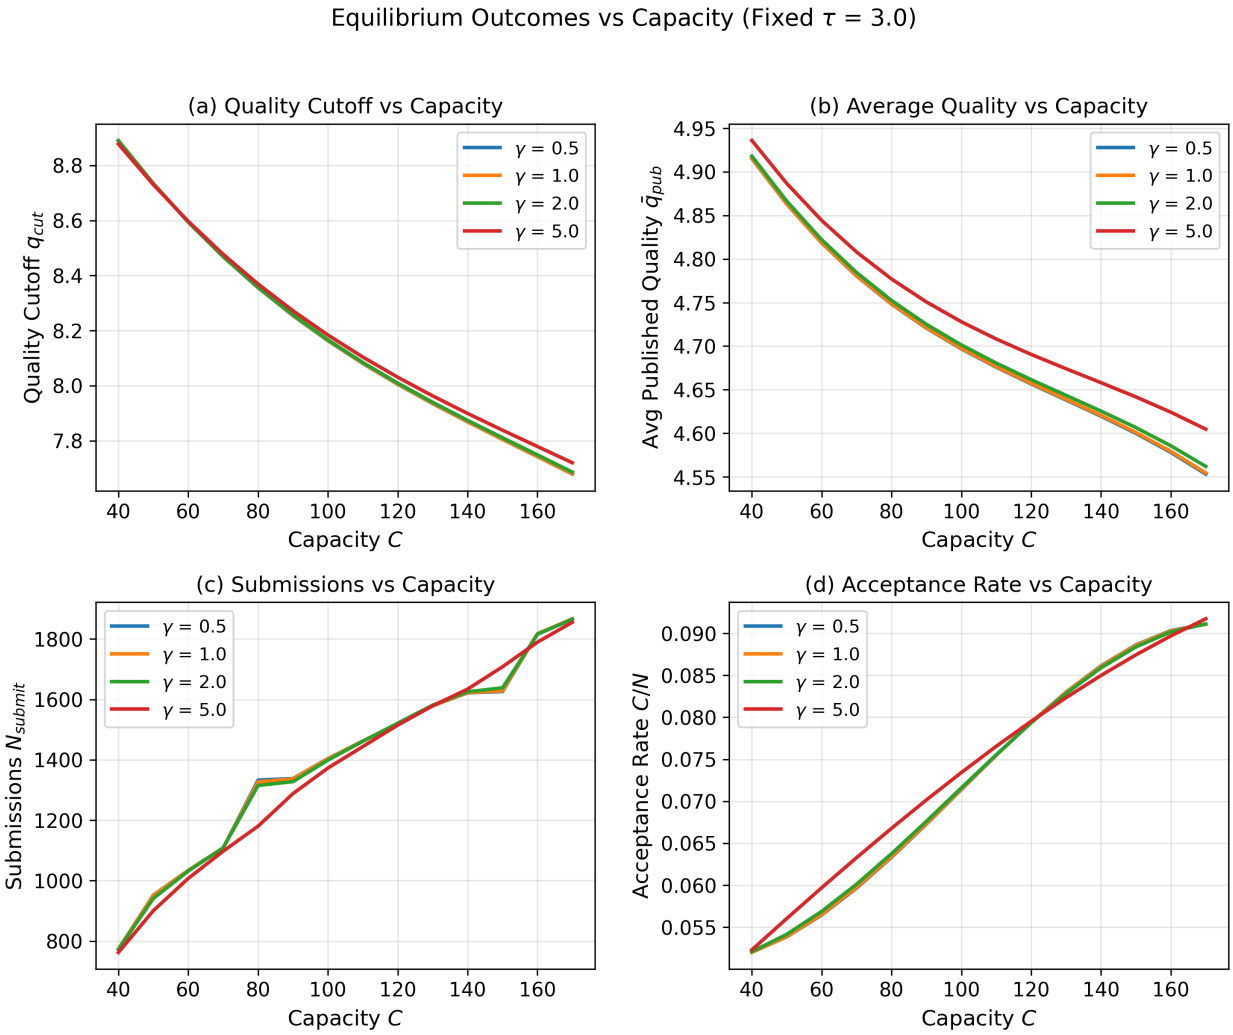


Figure A2: Equilibrium outcomes versus capacity C.

Notes: The figure plots equilibrium quantities (cutoff, average quality, submissions, and related variables) as functions of journal capacity C when the submission fee τ and AI productivity γ are held fixed. It shows how the equilibrium supply of accepted papers and their average quality respond to changes in capacity, which is useful for understanding the planner's trade-off between quantity and quality. Parameters are set according to the baseline calibration.

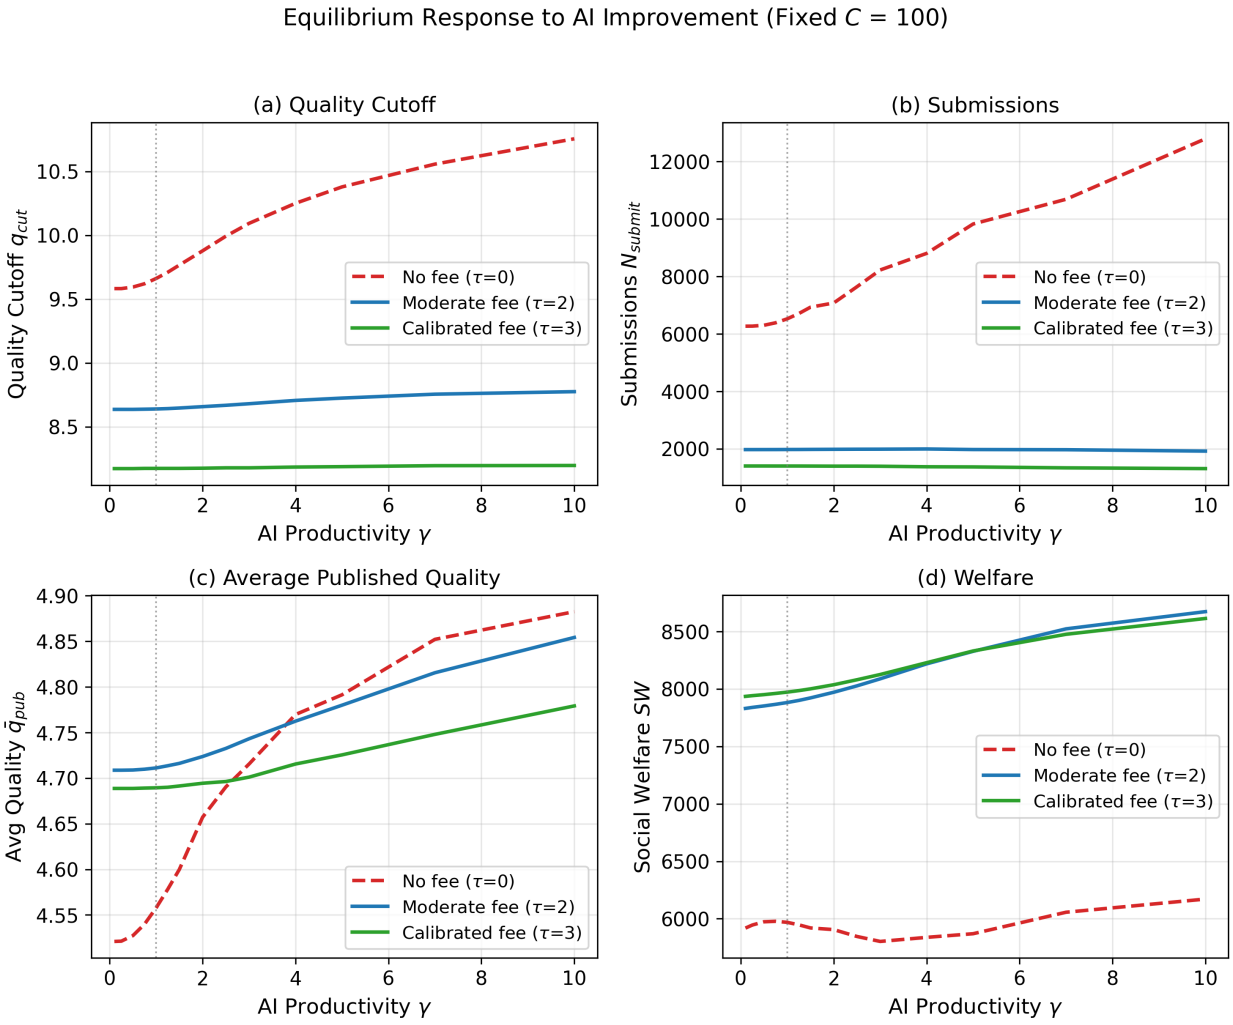


Figure A3: Equilibrium outcomes versus AI productivity γ.

Notes: The figure shows how equilibrium outcomes vary with AI productivity γ when capacity C and the submission fee τ are held fixed. It illustrates how the same journal policy leads to different equilibrium cutoffs, average quality, and submission volumes as γ changes, which helps explain why the optimal policy (C^*, τ^*) must adjust with γ. Baseline calibration is used throughout.

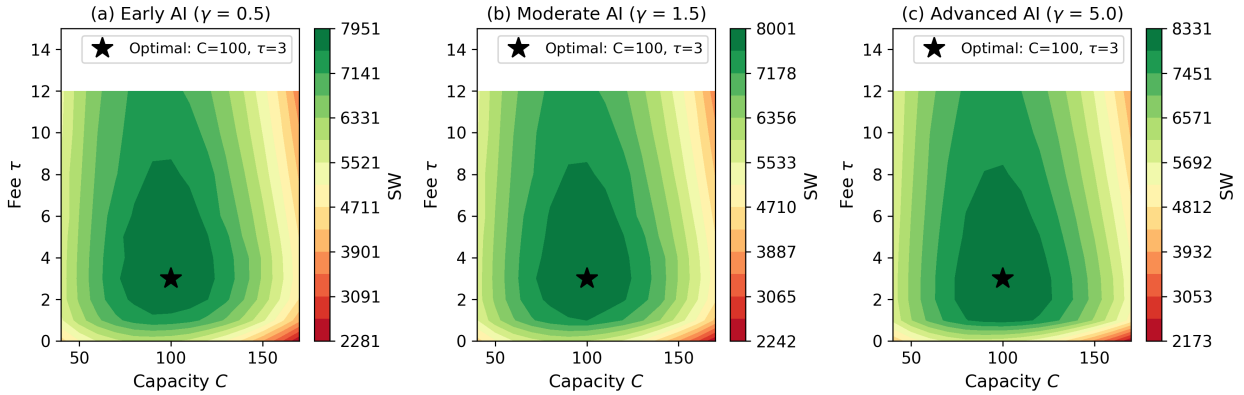


Figure A4: Welfare surface over capacity and fee.

Notes: The figure plots social welfare $SW(C, \tau)$ as a surface over capacity C and the submission fee τ for a fixed value of AI productivity γ. The welfare-maximizing policy (C^*, τ^*) corresponds to the maximum of this surface. The shape of the surface illustrates the trade-offs the planner faces when varying capacity and fees. Baseline calibration is used.

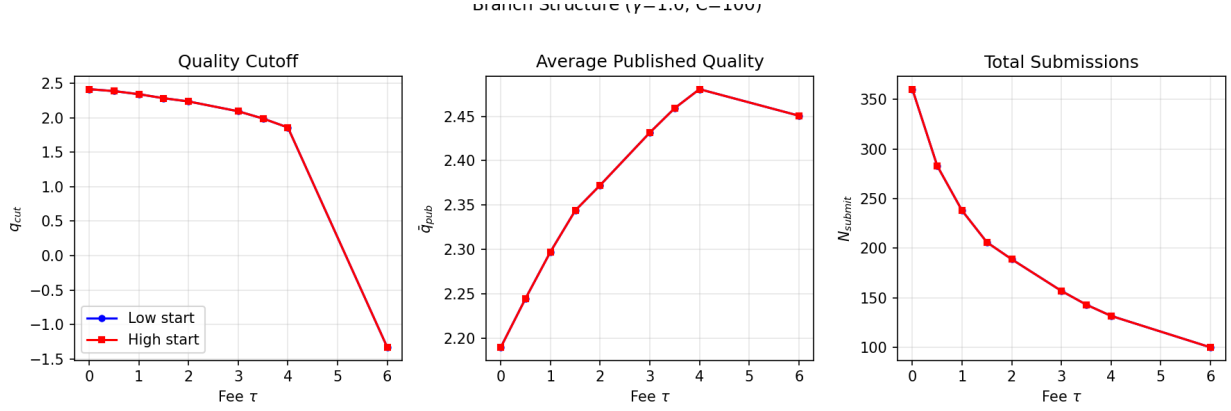


Figure A5: Branch structure in the equilibrium locus.

Notes: The figure displays the branch structure of the equilibrium locus in the $(q_{\text{cut}}, \bar{q}_{\text{pub}})$ plane as key parameters (such as capacity or the fee) vary. It is used to check that the equilibrium is well defined and that the solver selects the correct branch when multiple branches exist. All diagnostics use the baseline calibration.

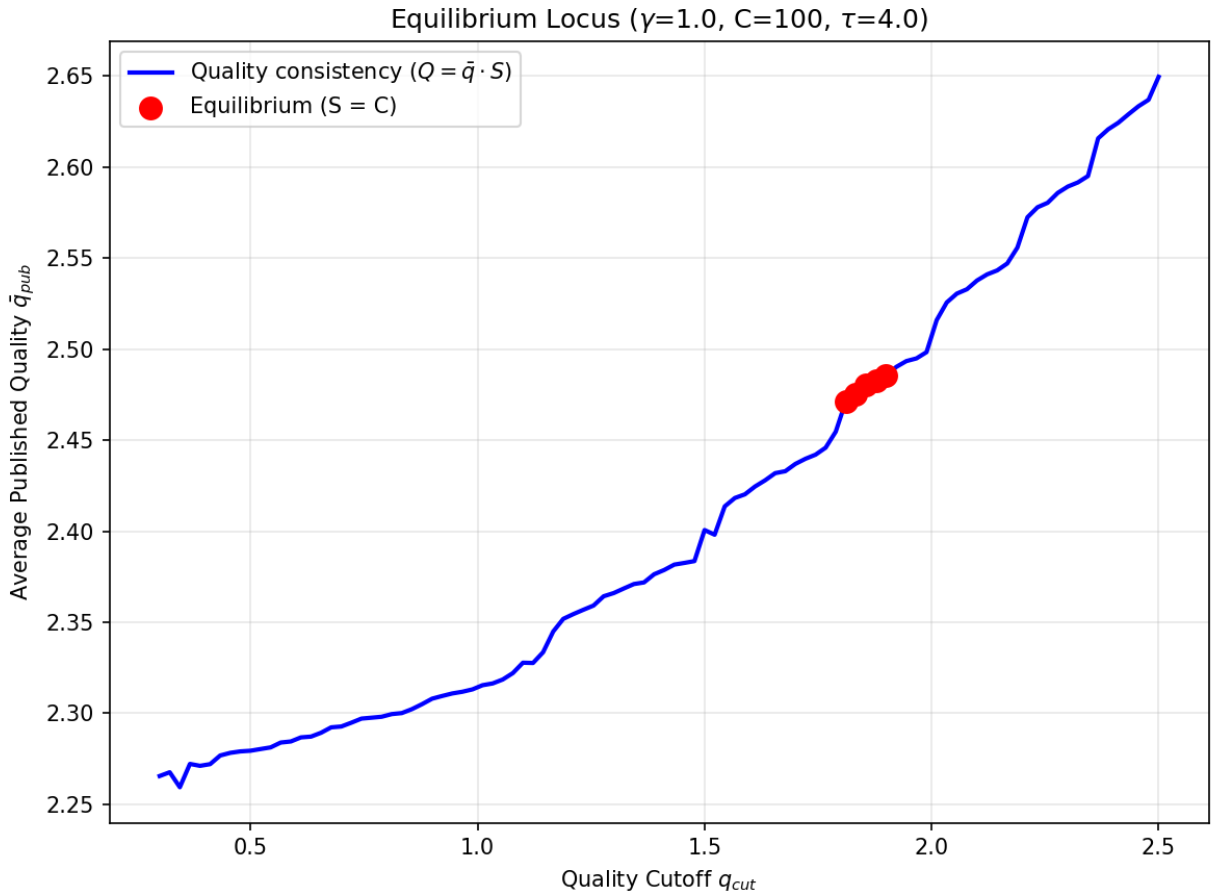


Figure A6: Equilibrium locus and branches.

Notes: The figure shows the equilibrium locus in the $(q_{\text{cut}}, \bar{q}_{\text{pub}})$ plane (or a related projection), including the different branches that can arise as parameters change. It supports the analysis of equilibrium uniqueness and stability in Appendix D. Baseline calibration is used.

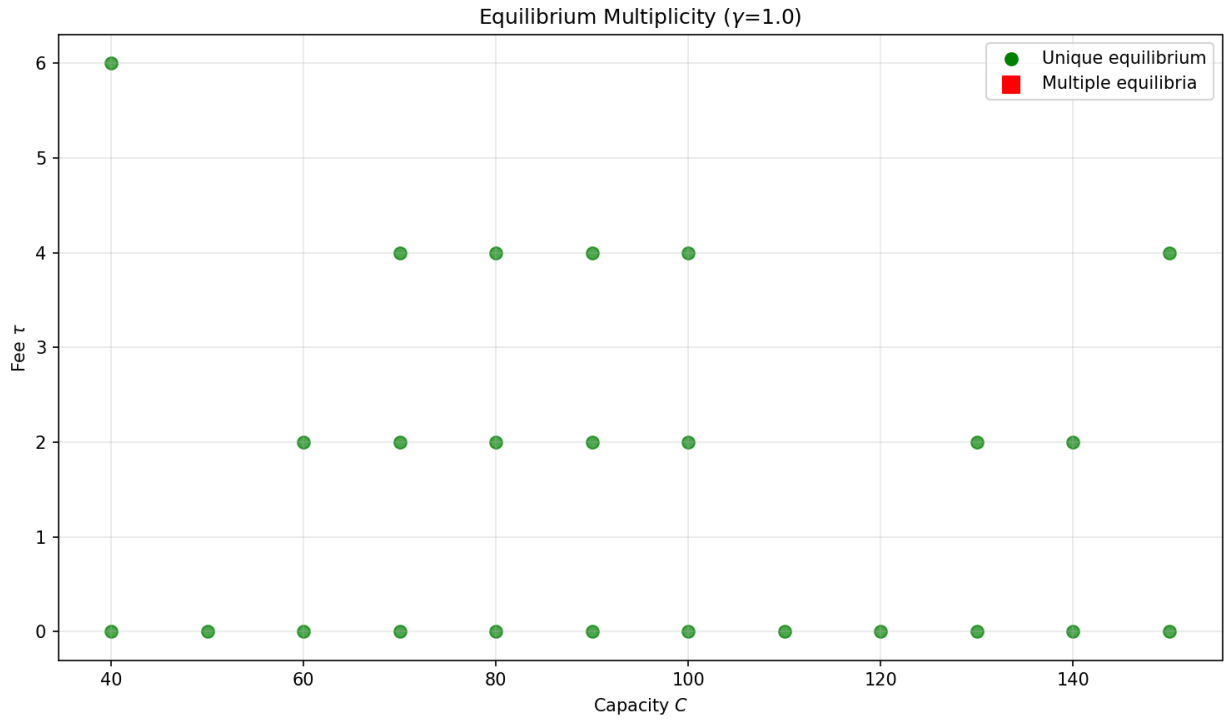


Figure A7: Region with multiple equilibria (if any).

Notes: The figure indicates the region in parameter space where the equilibrium system can admit multiple solutions. Under the baseline calibration used in the main text, the solver converges to a unique equilibrium at all grid points; we include this diagnostic for completeness and for robustness checks under alternative parameter values.

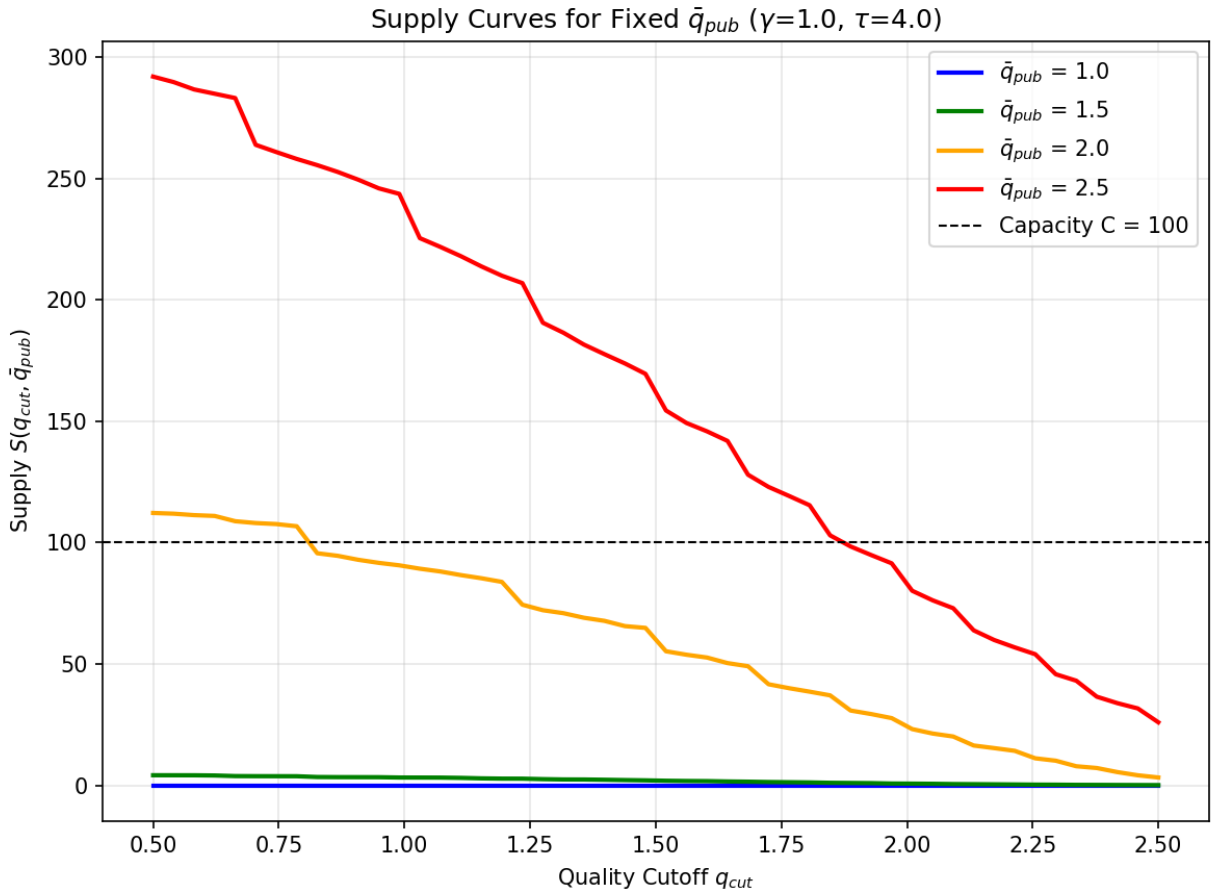


Figure A8: Supply curves across γ or policy parameters.

Notes: The figure plots supply curves (total accepted papers or total quality supplied) as functions of the equilibrium cutoff or related variables, for different values of γ or policy parameters. These curves illustrate how the supply side of the market responds across the AI productivity range and help verify that the numerical equilibrium is consistent with the model structure. Baseline calibration is used.

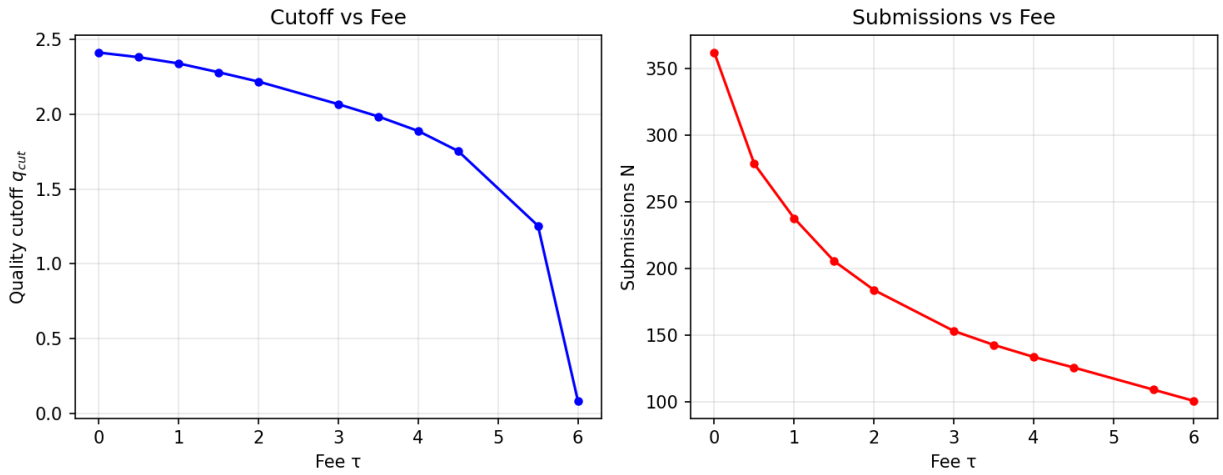


Figure A9: Continuity along a parameter sweep.

Notes: The figure shows how equilibrium quantities (cutoff, average quality, or related variables) evolve along a continuous sweep in capacity C or the submission fee τ. Continuity of these curves is used to verify that the equilibrium solver is stable and that the solution does not jump discontinuously across the grid. All runs use the baseline calibration.