The AI Assurance Cost Paradox: Fixed Governance Costs, Heterogeneous-Firm Adoption, and Market Concentration

Kwan Hong TAN

Singapore University of Social Sciences, Singapore

ABSTRACT

Artificial intelligence can lower marginal production costs while requiring firms to incur fixed costs for integration, validation, governance, cybersecurity, documentation, and human oversight. This paper develops a heterogeneous-firm model of an AI assurance cost paradox: when safe deployment is produced mainly through firm-specific fixed obligations, assurance can raise the productivity threshold for adoption and reallocate market share toward firms already above that threshold. Under constant-elasticity demand, the model separates technology-integration cost, assurance-governance cost, gross AI productivity gain, and proportional verification burden. It derives an endogenous adoption cutoff, a positive adopter-share concentration wedge under selective diffusion, a closed-form turning point for that wedge, and a condition under which risk-equivalent shared assurance broadens adoption. A deterministic experiment with 200,000 synthetic firms evaluates the mechanism. The analysis explicitly distinguishes descriptive firm-size adoption gradients from causal evidence, provides literature-informed justification for the benchmark parameterization, and expands robustness checks across fixed cost, elasticity of substitution, Pareto shape, gross AI gains, and verification burdens. The core non-linearity also survives replacement of the Pareto productivity distribution with a matched-dispersion lognormal distribution. Allowing AI gains to vary across firms shows that the magnitude of concentration depends on how technological gains covary with productivity, which clarifies an important boundary of the baseline result. Current U.S. Census and Eurostat evidence is therefore treated as descriptive consistency rather than validation of the fixed-cost mechanism. The findings imply that AI policy should evaluate not only the level of assurance but also its cost architecture. Shared testing, certification, audit utilities, and reusable compliance infrastructure can preserve risk objectives while reducing duplicated fixed-cost incidence and widening safe diffusion.

Keywords: artificial intelligence; firm heterogeneity; market concentration; technology diffusion; AI assurance; regulation

INTRODUCTION

Artificial intelligence is increasingly treated as a general-purpose technology capable of reducing the cost of prediction, search, content production, coding, analysis, and decision support. Micro-level evidence already documents meaningful productivity gains in selected knowledge-work settings. Noy and Zhang (2023) report faster completion and higher quality in professional writing tasks, while Brynjolfsson, Li, and Raymond (2025) find a 15 percent average increase in issues resolved per hour after the introduction of a generative AI assistant in customer support. At the aggregate level, estimates remain more uncertain. Acemoglu (2025) emphasizes the importance of task exposure and economically feasible cost savings, whereas Filippucci, Gal, and Schief (2026) show that plausible sectoral adoption can still generate sizable total factor productivity gains. The common implication is that technical capability does not translate mechanically into economy-wide productivity. Diffusion matters.

Diffusion, however, is uneven. In 2025, 17 percent of small enterprises in the European Union used at least one AI technology, compared with 30.36 percent of medium enterprises and 55.03 percent of large enterprises (Eurostat, 2026). Recent U.S. evidence shows the same broad gradient. The U.S. Census Bureau reported in May 2026 that 37 percent of firms with at least 250 employees were using AI, compared with 32 percent among firms with 100 to 249 employees and less than 20 percent among firms with four or fewer employees (U.S. Census Bureau, 2026). Using nationally representative Business Trends and Outlook Survey data, Bonney et al. (2026) similarly find that AI use is substantially higher in large firms and knowledge-intensive

sectors. These patterns do not prove that fixed governance costs cause the adoption gap, but they make firm heterogeneity central to any economics of AI diffusion.

The conventional competition narrative contains a tension. On one side, AI can reduce operating costs, lower minimum efficient scale, and give smaller firms access to capabilities previously requiring specialized labor or capital. OECD (2025a) therefore notes multiple mechanisms through which AI could strengthen downstream contestability. On the other side, firms differ sharply in data maturity, integration capacity, specialized skills, managerial attention, cybersecurity readiness, and ability to absorb legal or compliance costs. OECD/BCG/INSEAD (2025) identifies skills, return-on-investment uncertainty, data maturity, and demand for clearer regulatory information as persistent adoption barriers. If substantial portions of the cost of safe AI use are fixed at the firm level, then the same governance requirement can represent a small burden for a high-revenue incumbent and a decisive barrier for a smaller or less productive competitor.

This paper formalizes that mechanism. It introduces the concept of the AI assurance cost paradox, defined as the possibility that a risk-reducing governance requirement, if delivered through predominantly firm-specific fixed costs, can unintentionally narrow AI diffusion and strengthen the market-share position of firms already above the adoption threshold. The argument is not that assurance should be weakened. Contemporary AI governance frameworks appropriately emphasize risk identification, measurement, controls, documentation, and oversight (Tabassi, 2023; European Parliament and Council, 2024). The economic question is different: for a given target level of assurance, how should the costs of achieving that target be structured across firms?

The model combines heterogeneous-firm economics with AI governance. Firms differ in productivity and face constant-elasticity demand. AI adoption multiplies effective productivity, but the net benefit is reduced by a proportional verification burden and a fixed adoption-and-assurance cost. The model produces an endogenous productivity cutoff for adoption. Embedding this cutoff in a Pareto productivity distribution yields closed-form expressions for the share of firms adopting, the pre-adoption revenue mass of those firms, and the post-adoption reallocation of revenue. The central distributional result is non-linear: selective adoption generates a positive concentration wedge, but the wedge is small when almost no firm adopts, rises when adoption is concentrated among the upper part of the productivity distribution, and falls again as adoption approaches universality.

The paper makes six contributions. First, it shifts the competition question from whether AI is intrinsically concentrating or democratizing to how the cost architecture of AI adoption shapes diffusion. Second, it decomposes total fixed adoption cost into technology-integration and assurance-governance components, and separates both from proportional verification burdens, allowing policy instruments to have different incidence across the firm distribution. Third, it derives a closed-form turning point for the concentration wedge and a simple condition under which shared assurance infrastructure broadens adoption. Fourth, it makes the benchmark parameterization more interpretable by expressing fixed cost relative to the incremental surplus of the least productive firm and by anchoring the elasticity, Pareto shape, and AI-gain ranges to established heterogeneous-firm and AI-productivity evidence. Fifth, it tests robustness beyond the baseline Pareto and uniform-gain assumptions using a matched-dispersion lognormal productivity distribution and heterogeneous AI-gain schedules. Sixth, it connects the economics of market structure to operational assurance research, including the problem of verification debt in AI-enabled organizations (Tan, 2026a) and the need for runtime controls when AI systems exercise delegated authority (Tan, 2026b).

The rest of the paper is organized as follows. Section 2 reviews the relevant literature and defines the research gap. Section 3 develops the model and propositions. Section 4 presents empirical motivation and numerical results. Section 5 discusses interpretation and mechanisms. Section 6 develops policy implications. Section 7 states limitations and an empirical research agenda. Section 8 concludes.

LITERATURE REVIEW AND RESEARCH GAP

AI productivity, diffusion, and complementary assets

The economics of AI increasingly distinguishes technical potential from realized productivity. Goldfarb and

Tucker (2019) frame digital technologies through changes in economic costs, while AI research emphasizes prediction, automation, complementarity, and organizational redesign. Experimental studies show that generative AI can improve performance in bounded tasks (Noy & Zhang, 2023; Brynjolfsson et al., 2025), but the magnitude is heterogeneous across workers and tasks. At the firm level, Babina, Fedyk, He, and Hodson (2024) associate investment in AI with firm growth and product innovation. At the industry level, McElheran, Yang, Kroff, and Brynjolfsson (2026) document the uneven adoption of industrial AI in the United States. These findings support a view in which AI is productive but complementary to organizational and intangible capital rather than a frictionless software shock.

The diffusion literature supplies a broader backdrop. Andrews, Criscuolo, and Gal (2016) document widening productivity divergence between frontier and laggard firms. In that setting, a technology whose adoption requires complementary capabilities can widen gaps even when it raises aggregate efficiency. The current AI evidence is consistent with this concern. Bonney et al. (2026) find not only that AI adoption is more common among large firms, but also that broader functional integration correlates positively with commercial performance. Eurostat (2026) reports that, among EU enterprises that had considered AI but did not use it in 2025, lack of relevant expertise was the most common reason, followed by unclear legal consequences and data-protection concerns. These obstacles map naturally into both fixed and variable adoption costs.

Firm heterogeneity and market reallocation

Heterogeneous-firm models provide a natural framework for studying selective technology adoption. Melitz (2003) shows how productivity differences generate endogenous selection and reallocation in monopolistically competitive industries. Although the present application is not a trade model, the same logic is useful: when a fixed cost must be paid before a firm can access a productivity-enhancing activity, only firms above a cutoff find adoption profitable. This sorting mechanism converts a technological opportunity into a distributional event.

The broader market-power literature demonstrates why reallocation matters. Autor, Dorn, Katz, Patterson, and Van Reenen (2020) show that shifts in market share toward high-productivity superstar firms can raise concentration and alter factor shares. De Loecker, Eeckhout, and Unger (2020) document rising markups concentrated in the upper tail of the firm distribution. These studies do not imply that concentration is always inefficient, because productive firms can gain share for efficiency-enhancing reasons. They do, however, establish that firm-level heterogeneity and reallocation are macroeconomically meaningful. The present paper adds a specific mechanism: fixed AI assurance costs can change which firms gain access to a common productivity multiplier.

AI market structure, governance, and assurance costs

Existing AI competition research has concentrated heavily on the upstream AI stack. Korinek and Vipra (2025) analyze economies of scale, foundation-model competition, and market structure in AI production. OECD (2025a) extends the discussion to downstream markets and explicitly notes that AI can lower operating costs while firms may still face fixed costs for installation, data integration, and complementary assets. The OECD also emphasizes that competitive effects depend on diffusion, because larger incumbents may be better able to absorb upfront investment even when AI lowers marginal costs.

AI governance adds a further layer of cost. The NIST AI Risk Management Framework organizes risk management around governance, mapping, measurement, and management activities (Tabassi, 2023). The European Union Artificial Intelligence Act adopts a risk-based regulatory structure and imposes obligations that can include risk management, documentation, monitoring, and human oversight depending on the system and actor (European Parliament and Council, 2024). Gans (2025) shows that learning about AI harms can itself alter the socially desirable speed of adoption, highlighting the economic relevance of the risk-adoption tradeoff. These contributions establish why assurance has benefits rather than being a pure deadweight cost. For the present paper, the additional question is whether the resources required to achieve a given assurance target are primarily fixed per adopting firm, proportional to usage or risk, or supplied through shared infrastructure.

Operational research on agentic AI makes the fixed-cost mechanism especially plausible. When AI systems move from recommendation toward planning, tool use, cross-system execution, and delegated action, organizations may require policy gates, audit trails, sandboxing, exception handling, and runtime monitoring. Tan (2026b) formalizes a runtime assurance architecture for such systems. Tan (2026a) separately shows how weak verification can create organizational verification debt. Those contributions focus on control quality and resilience. The present paper asks a different question: what happens to industry structure when the organizational capabilities required for reliable AI use have a substantial fixed-cost component?

Research gap and hypotheses

The literature therefore leaves a gap between two established observations. First, AI can reduce marginal costs and raise productivity. Second, AI adoption and complementary capabilities are distributed unevenly across firms. What is missing is a parsimonious model that treats the architecture of governance and assurance costs as an endogenous determinant of diffusion and concentration. The paper addresses this gap through four propositions rather than through a causal empirical estimate. The numerical analysis is designed as a structural stress test: it evaluates whether the propositions survive finite-firm simulation, systematic perturbation of the key parameters, an alternative productivity distribution, and moderate heterogeneity in AI productivity gains.

METHODOLOGY AND MODEL

Research design

The study uses analytical economic modelling supplemented by deterministic synthetic-firm experiments. The analytical model derives adoption thresholds and market-share reallocations in closed form. The baseline numerical experiment discretizes the assumed productivity distribution into 200,000 equally spaced quantiles, eliminating Monte Carlo sampling noise and making every reported result reproducible from the stated parameters. Robustness exercises then vary the principal structural inputs, replace the Pareto distribution with a matched-dispersion lognormal distribution, and relax the assumption of a uniform AI productivity gain. Official U.S. Census, Eurostat, OECD, and regulatory sources are used only to motivate assumptions and interpret the mechanism. They are not used to estimate causal parameters or to fit the normalized fixed-cost terms.

The model is deliberately partial equilibrium. It holds the mass of incumbent firms and the demand environment fixed while examining the first-order redistribution created by differential access to an AI productivity multiplier. This choice isolates the proposed mechanism. General-equilibrium entry, exit, wage adjustment, and endogenous model-provider pricing are discussed in Section 7 as extensions.

Firms and baseline operating profit

Consider a monopolistically competitive industry with constant elasticity of substitution $\sigma > 1$. A firm is indexed by productivity z. Standard CES demand and constant markups imply that operating profit before fixed overhead can be summarized by a positive market-size constant B multiplied by productivity raised to $\sigma - 1$. Let F_0 denote the fixed cost of ordinary operation. Baseline profit is therefore:

$$\pi_0(z) = Bz^{\sigma-1} - F_0 \qquad (1)$$

This reduced-form representation preserves the heterogeneous-firm logic while avoiding unnecessary demand-system notation. B collects market size, the aggregate price index, wage normalization, and markup constants. Since the paper focuses on comparative statics, B can be normalized in the numerical exercise without loss of the mechanism.

AI productivity, proportional verification burden, and fixed adoption cost

Suppose AI creates a gross proportional efficiency gain $g \geq 0$. Safe deployment also requires a proportional verification or monitoring burden $\tau \geq 0$. The net effective productivity multiplier from AI is defined as:

$$\lambda = \frac{1+g}{1+\tau} \quad (2)$$

The economically interesting case is $\lambda > 1$. A firm that adopts AI must also pay a total fixed cost $F_A = F_T + F_G$. F_T denotes technology and process-integration expenditures, including data preparation, systems integration, and staff adaptation. F_G denotes fixed assurance and governance expenditures, including validation, legal review, security hardening, documentation, audit readiness, and governance design. The decomposition prevents ordinary implementation frictions from being misattributed to regulation. The comparative statics use total F_A because both components act as entry-like costs, while policy experiments that share assurance infrastructure should be interpreted primarily as reducing F_G. The post-adoption operating-profit multiplier is $q = \lambda$ raised to $\sigma - 1$, so adopted-firm profit is:

$$\pi_A(z) = Bqz^{\sigma-1} - F_0 - F_A \quad (3)$$

Adoption is privately profitable when the incremental operating surplus generated by AI is at least as large as the total fixed adoption cost:

$$Bz^{\sigma-1}(q-1) \geq F_A \quad (4)$$

This condition generates an adoption threshold:

$$z_A = \left[\frac{F_A}{B(q-1)}\right]^{\frac{1}{(\sigma-1)}} \quad (5)$$

Equation (5) is the core sorting mechanism. Higher fixed cost raises the cutoff. A stronger net productivity multiplier lowers it. An increase in proportional verification burden reduces λ and therefore q, raising the cutoff. Because the fixed cost is divided by a productivity-dependent operating-surplus term, its incidence is regressive with respect to productivity.

Productivity distribution and diffusion

Assume productivity follows a Pareto distribution with minimum productivity z_{min} and shape parameter $k > \sigma - 1$. The restriction ensures a finite expected revenue term. For an interior cutoff $z_A > z_{min}$, the fraction of firms that adopt is:

$$P_A = \left(\frac{z_{min}}{z_A}\right)^k \quad (6)$$

The baseline revenue mass generated by firms above the adoption threshold is more concentrated than the firm-count share because high-productivity firms produce more. Let r_A denote the share of baseline revenue attributable to firms that will adopt. Integration over the Pareto tail gives:

$$r_A = \left(\frac{z_{min}}{z_A}\right)^{k-(\sigma-1)} \quad (7)$$

When $z_A \leq z_{min}$, all firms adopt, so $P_A = r_A = 1$. Equations (6) and (7) distinguish diffusion measured by the number of firms from diffusion measured by the economic weight of those firms.

The adopter-share concentration wedge

After adoption, the revenue of adopting firms is multiplied by q relative to non-adopters. Holding the demand environment fixed, the post-adoption share of industry revenue captured by adopters is:

$$S_A = \frac{qr_A}{1+(q-1)r_A} \quad (8)$$

Define the adopter-share concentration wedge as the increase in the revenue share of the group that adopts: $\Delta S = S_A - r_A$. Algebra yields:

$$\Delta S = \frac{(q-1)r_A(1-r_A)}{1+(q-1)r_A} \quad (9)$$

The wedge is positive for $0 < r_A < 1$. It is zero when no economically meaningful firm adopts and when all firms adopt. Differentiating Equation (9) with respect to r_A shows that the wedge reaches its unique maximum at:

$$r^* = \frac{1}{1+\sqrt{q}} \quad (10)$$

This result produces the diffusion-concentration hump. Selective adoption can increase concentration most strongly at an intermediate diffusion state, not necessarily when adoption is rarest.

Risk-equivalent policy design

Assurance is not merely a cost. It is an input into trustworthy deployment because it can reduce expected error, security, legal, and safety losses. Let $R(F_G, \tau)$ denote residual AI risk, decreasing in effective fixed assurance resources F_G and in proportional verification effort τ. Technology-integration cost F_T can affect adoption without necessarily producing assurance, so it is excluded from the risk-production function. A regulator or industry can be interpreted as targeting a maximum acceptable residual-risk level $\bar{R}$. The relevant policy problem is therefore not to minimize assurance, but to minimize the adoption threshold generated by $F_A = F_T + F_G$ subject to the risk constraint:

$$\min z_A(F_A, \tau) \quad \text{subject to} \quad R(F_G, \tau) \leq \bar{R} \quad (11)$$

A shared assurance utility can change the cost mix while leaving the risk constraint satisfied. Examples include common evaluation suites, pooled audit infrastructure, reusable control libraries, sectoral certification, shared model monitoring, and public testbeds. Suppose shared provision reduces F_G enough that total fixed cost falls from F_A to F_A', but also reduces the net AI multiplier from λ to λ' because firms pay a per-use fee or perform additional proportional checks. Diffusion still broadens if the threshold falls, which occurs when:

$$\frac{F_A'}{F_A} < \frac{q'-1}{q-1} \quad (12)$$

Equation (12) is a risk-equivalent cost-architecture condition. A shared system can impose more variable checking and still increase adoption if it removes enough duplicated fixed expense.

Robustness and external-validity strategy

The baseline model intentionally uses a compact CES-Pareto structure so that the adoption and concentration mechanisms remain transparent. The numerical analysis therefore adds three robustness layers. First, one-at-a-time sensitivity exercises vary total fixed adoption cost, the elasticity of substitution, the Pareto shape parameter, the gross AI productivity gain, and the proportional verification burden. Second, the Pareto distribution is replaced by a lognormal productivity distribution calibrated to match the baseline economy's mean pre-AI sales and sales Gini, so that robustness is not obtained simply by changing initial dispersion. Third, the common-gain assumption is relaxed by assigning equal masses of firms gross AI gains of 15, 20, and 25 percent under three deterministic patterns: approximately uncorrelated with productivity, increasing with productivity, and decreasing with productivity. These exercises are stress tests of the mechanism, not empirical calibrations. Their purpose is to identify which results depend on functional form and which depend more generally on selective diffusion.

Table 1. Main Analytical Propositions

Proposition	Formal result	Economic interpretation
P1. Adoption threshold	z_A rises with F_A and τ; z_A falls with g.	Total fixed adoption cost and verification burdens reduce diffusion; larger technical gains expand it.

P2. Regressive fixed-cost incidence	F_A / [B $z^{\sigma-1}(q\text{-}1)$] declines in z.	The same fixed obligation absorbs a larger share of AI-generated surplus for lower-productivity firms.
P3. Concentration hump	$\Delta S > 0$ for $0 < r_A < 1$; maximum at $r^* = 1/(1+\sqrt{q})$.	Selective adoption reallocates sales toward adopters most strongly at intermediate diffusion.
P4. Shared-assurance condition	$F_A'/F_A < (q'\text{-}1)/(q\text{-}1)$.	Fixed-cost reduction can outweigh a modest proportional assurance charge and broaden diffusion.

Note. All results are conditional on the model assumptions stated in Section 3. The concentration wedge is a group-share reallocation measure, not a claim that every concentration index must move identically.

ANALYSIS AND RESULTS

Descriptive empirical motivation and identification limits

The model predicts that, when adoption contains an important fixed-cost component, larger or more productive firms should be more likely to adopt. Current official statistics show a strong size gradient that is directionally consistent with this implication. Table 2 summarizes recent evidence. This consistency is descriptive only. Firm size is not identical to productivity, and the observed gradient may reflect sector composition, workforce skills, digital capital, data availability, access to finance, management quality, expected return from AI, vendor selection, or other omitted factors. The official statistics therefore cannot identify the causal effect of assurance cost, and no parameter in the model is estimated from the size gradient. Their narrower role is to reject the simplifying assumption of uniform diffusion and to motivate a framework in which adoption differs across the firm distribution.

Table 2. Recent Official Evidence on Uneven Business AI Adoption

Source	Population / period	Selected finding	Interpretation for model
Eurostat (2026)	EU enterprises with 10+ employees, 2025	AI use: 17.00% small; 30.36% medium; 55.03% large.	Adoption rises steeply with enterprise size.
U.S. Census Bureau (2026)	BTOS, Dec. 2025-May 2026	37% for 250+ employees; 32% for 100-249; less than 20% for 0-4 employees.	Large firms exhibit materially higher adoption.
Bonney et al. (2026)	Nationally representative U.S. BTOS AI supplement, Nov. 2025-Jan. 2026	18% of firms used AI; 32% employment-weighted; use substantially higher in large firms and knowledge-intensive sectors.	Economic weight of adopters exceeds simple firm-count diffusion.
OECD/BCG/INSEAD (2025)	G7 and Brazil enterprise evidence	Skills, data maturity, uncertain returns, and demand for regulatory guidance remain adoption barriers.	Complementary capabilities and governance frictions can generate fixed or quasi-fixed costs.

Note. Eurostat size classes are small = 10-49, medium = 50-249, and large = 250+ employees and self-employed persons. These figures establish uneven diffusion, not the cause of that unevenness. They are not used to calibrate fixed assurance costs, productivity, or causal effects.

Baseline parameterization

The numerical exercise uses two pure normalizations, $B = 1$ and $z_{min} = 1$, so monetary magnitudes should be interpreted relative to the model's incremental operating-surplus scale rather than as currency units. The

baseline elasticity is $\sigma = 3$ and the Pareto shape is $k = 4.5$. These values sit close to common heterogeneous-firm benchmarks. For example, Melitz and Redding (2015) use an elasticity of substitution of 4 and a Pareto productivity shape of 4.25 in a quantitative heterogeneous-firm setting. The present baseline uses a slightly lower elasticity and a slightly thinner productivity tail, then varies both parameters substantially in sensitivity analysis.

The gross AI productivity gain is set to $g = 0.20$. This is an order-of-magnitude benchmark rather than a firm-level estimate: recent task-level evidence ranges from roughly 15 percent higher throughput in field deployment (Brynjolfsson et al., 2025) to much larger time savings in bounded writing tasks (Noy & Zhang, 2023). The proportional verification burden is set to $\tau = 0.0435$, which converts the 20 percent gross gain into a 15 percent net productivity multiplier, $\lambda = 1.15$. Because no direct empirical estimate of τ is imposed, the sensitivity analysis spans zero to 10 percent. The intermediate fixed-cost benchmark $F_A = 0.55$ is also illustrative. Its economic meaning is clearer through the dimensionless burden ratio $\chi = F_A / [B(q - 1)z_{min}^{\sigma - 1}]$. The baseline implies $\chi = 1.705$, meaning that fixed adoption cost is about 1.7 times the incremental AI operating surplus available to the least productive firm before heterogeneity is taken into account. The scenario analysis varies this burden across regions of near-universal, selective, and sparse diffusion.

Table 3. Baseline Numerical Parameters

Parameter	Baseline	Role
Number of synthetic firms	200,000	Deterministic equally spaced quantiles.
B	1.000	Market-size and profit normalization.
z_{min}	1.000	Productivity-scale normalization.
σ	3.000	Elasticity of substitution; sensitivity: 2.5 to 4.0.
k	4.500	Pareto productivity shape; sensitivity: 3.5 to 6.0.
g	0.200	Gross AI productivity gain.
τ	0.0435	Proportional verification burden.
$\lambda = (1+g)/(1+\tau)$	1.150	Net AI productivity multiplier.
$q = \lambda^{\sigma - 1}$	1.3225	Revenue / operating-surplus multiplier for adopters.
F_A	0.550	Intermediate normalized fixed-cost benchmark.
χ	1.705	Dimensionless fixed-cost burden at the benchmark.
Analytical r^*	0.4651	Revenue-mass turning point for the concentration wedge.
Associated F_A^*	0.5950	Fixed-cost level at the analytical wedge peak.

Note. B and z_{min} are normalizations. The remaining values are transparent benchmark choices, not fitted structural estimates. Sensitivity analysis varies every substantive parameter in the benchmark model.

Result 1: fixed adoption costs produce selective diffusion

Equation (5) implies immediate comparative statics. With the baseline q, a fixed cost of 0.30 lies below the incremental surplus generated even by the least productive firm, so adoption is universal. Raising the fixed cost to 0.35 reduces the analytical adoption share to approximately 83.2 percent. At 0.55, adoption falls to about 30.1 percent, and at 1.20 it falls to roughly 5.2 percent. This rapid decline is generated entirely by firm heterogeneity and the fixed-cost threshold.

The model therefore distinguishes two margins that policy discussions often combine. The first is the per-adopter technical productivity effect, summarized by λ. The second is the extensive margin of which firms can

profitably reach that effect. A policy or organizational practice can improve the reliability of AI use but still reduce aggregate realized productivity if it shifts the extensive margin enough. Conversely, a shared assurance architecture can deliver more aggregate productivity even with a slightly lower λ if it brings many additional firms above the adoption threshold.

Result 2: selective adoption creates a concentration hump

The closed-form concentration wedge in Equation (9) is positive whenever adoption is selective. Figure 1 plots the wedge against the baseline revenue mass of firms that adopt. For any $q > 1$, the relationship is hump-shaped. When almost no revenue-generating firm adopts, AI cannot reallocate much industry revenue. When almost all firms adopt, the common productivity multiplier largely preserves relative shares. The largest reallocation occurs between those extremes.

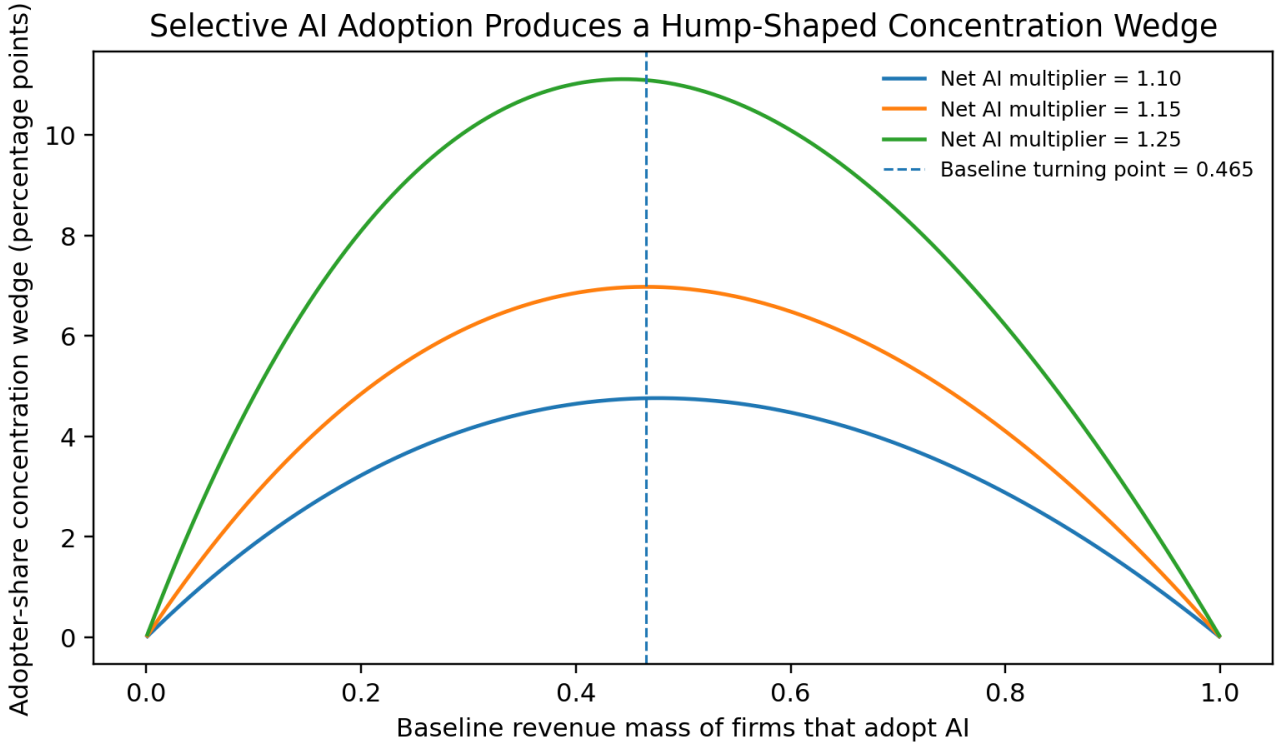


Figure 1. Selective AI adoption produces a hump-shaped adopter-share concentration wedge. The vertical line marks the analytical turning point for the baseline net multiplier $\lambda = 1.15$ and $\sigma = 3$.

For the baseline case $q = 1.3225$, the wedge reaches its maximum when r_A is approximately 0.4651. This corresponds to a normalized fixed cost of approximately 0.595. At that point, only about one-quarter of firms adopt, but they represent roughly 46.5 percent of baseline revenue. Their productivity advantage after AI raises their revenue share by almost 7.0 percentage points. The result is economically important because it rejects a monotonic intuition: the concentration effect can become smaller when fixed costs become extremely high, not because the policy becomes benign, but because so few firms obtain the technology that the aggregate reallocation channel weakens.

Result 3: a full-distribution sales measure shows similar non-linearity

Because the analytical wedge tracks the adopting group rather than the entire firm-size distribution, the numerical exercise also computes the Gini coefficient of firm sales as a full-distribution robustness measure. The baseline synthetic economy has a sales-share Gini of approximately 0.286. Universal AI adoption leaves this measure unchanged because all firms receive the same multiplier. At intermediate fixed costs, the Gini rises sharply because firms just above the adoption cutoff receive a discrete productivity advantage over firms just below it. Figure 2 shows the resulting non-monotonic pattern.

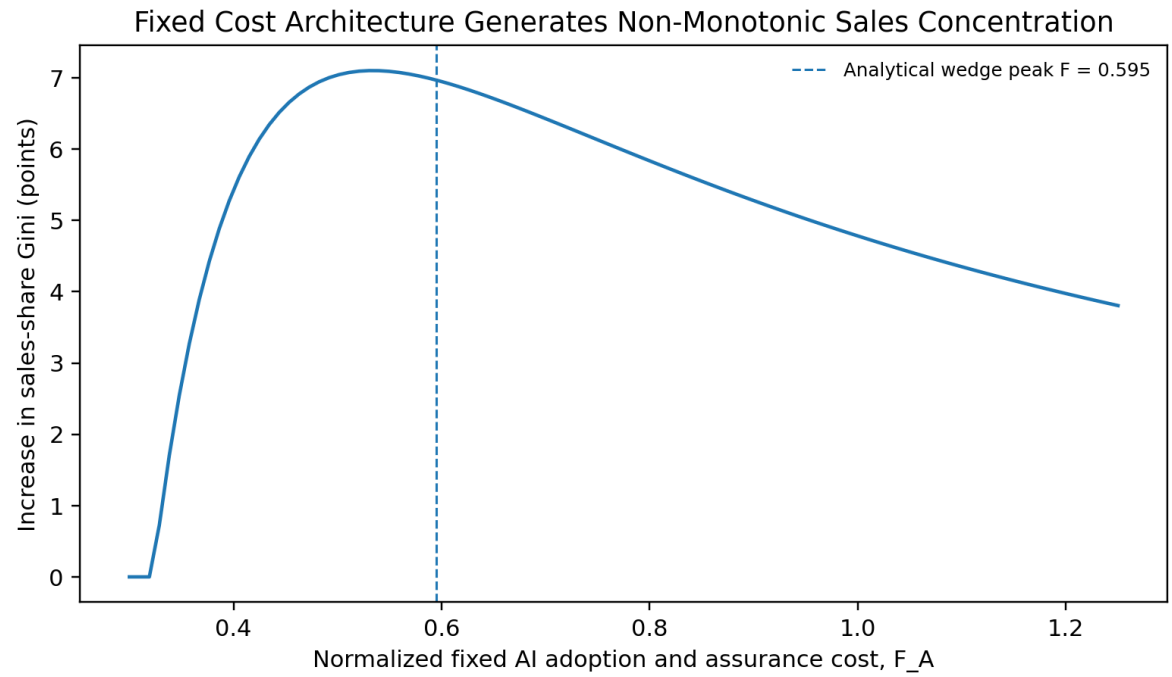


Figure 2. Increase in the sales-share Gini as the normalized fixed AI adoption-and-assurance cost changes. The deterministic synthetic-firm experiment uses 200,000 Pareto quantiles and the baseline parameters in Table 3.

The Gini increase peaks close to, though not exactly at, the analytical group-share turning point. This difference is expected because the Gini aggregates pairwise distributional distances across all firms, whereas ΔS tracks the revenue shift into the adopting group. The agreement in broad shape is therefore more important than exact coincidence. It shows that the closed-form mechanism is also visible in a full-distribution inequality measure. It does not imply that every market-concentration statistic must be hump-shaped. Measures such as the Herfindahl-Hirschman Index weight the extreme upper tail differently and can therefore respond differently as the adoption cutoff moves through a Pareto distribution.

Table 4. Baseline and Policy Scenario Results

Scenario	F_A	λ	Firm adoption	Baseline adopter revenue	ΔS (pp)	Sales Gini	Gini increase (pts)	Output index
Universal diffusion	0.300	1.15	100.0%	100.0%	0.00	0.286	0.00	1.322
Low fixed burden	0.350	1.15	83.2%	90.3%	2.19	0.313	2.71	1.291
Fragmented assurance	0.550	1.15	30.1%	51.3%	6.91	0.356	7.09	1.165
Analytical wedge peak	0.595	1.15	25.2%	46.5%	6.98	0.355	6.97	1.150
High fixed burden	0.850	1.15	11.3%	29.8%	6.15	0.341	5.55	1.096
Very high fixed burden	1.200	1.15	5.2%	19.4%	4.74	0.325	3.97	1.062
Shared assurance	0.250	1.12	100.0%	100.0%	0.00	0.286	0.00	1.254

Note. Output index is mean technology-adjusted revenue relative to the pre-AI synthetic baseline, holding the demand environment fixed. It is not a general-equilibrium welfare measure.

Result 4: shared assurance can dominate fragmented assurance

Table 4 compares a fragmented-assurance benchmark with a shared-assurance scenario. Under intermediate fragmented assurance, $F_A = 0.55$, $g = 0.20$, and $\tau = 0.0435$, giving $\lambda = 1.15$. Approximately 30 percent of firms

adopt, the sales-share Gini rises by about 7.1 points, and the technology-adjusted output index rises to roughly 1.165. In the shared-assurance scenario, the fixed cost falls to 0.25 while the same gross technical gain is paired with a higher proportional verification burden of approximately 7.14 percent, giving $\lambda = 1.12$. This representation makes the policy tradeoff explicit: shared infrastructure removes duplicated fixed assurance resources but requires more per-use checking. Because the fixed-cost reduction is large enough to satisfy Equation (12), all firms adopt. The sales distribution returns to its pre-AI Gini, while the output index rises to approximately 1.254. In this stylized comparison, broader diffusion more than offsets the smaller net per-adopter productivity multiplier.

The policy implication is not that every shared-assurance program will have this effect. It is that the fixed-versus-variable architecture can matter as much as the headline compliance burden. Risk-equivalent infrastructure that eliminates repeated setup costs across firms may improve both diffusion and competition. This is especially relevant for sectors in which many firms must perform similar evaluations, logging, documentation, vendor due diligence, or model-risk checks.

Sensitivity analysis

Table 5 expands the one-at-a-time analysis around the intermediate-cost benchmark. It varies all five core dimensions of the model: fixed adoption cost, elasticity of substitution, Pareto shape, gross AI productivity gain, and proportional verification burden. The main selection mechanism is robust, but its magnitude is not invariant. Lower gross gains or higher verification burdens push the adoption cutoff upward and eventually weaken the concentration effect because adoption becomes sparse. Higher gross gains can instead produce universal diffusion, eliminating the selective-adoption wedge. Variation in elasticity and productivity dispersion changes both the fraction of firms near the cutoff and the sales weight of adopters. These responses reinforce, rather than weaken, the central interpretation: concentration is generated by selective access to the productivity multiplier, not by a fixed numerical parameter choice.

Table 5. Expanded One-at-a-Time Sensitivity Analysis

Case	Changed parameter	Adoption cutoff	Adoption	ΔS (pp)	Gini Δ (pts)	Output
Baseline	$F_A = 0.55$	1.306	30.1%	6.91	7.09	1.165
Lower fixed cost	$F_A = 0.35$	1.042	83.2%	2.19	2.71	1.291
Higher fixed cost	$F_A = 0.85$	1.623	11.3%	6.15	5.55	1.096
Lower elasticity	$\sigma = 2.5$	1.772	7.6%	3.30	2.97	1.042
Higher elasticity	$\sigma = 4.0$	1.018	92.2%	0.91	1.33	1.507
Heavier productivity tail	$k = 3.5$	1.306	39.3%	5.87	6.48	1.216
Thinner productivity tail	$k = 6.0$	1.306	20.2%	6.55	6.37	1.111
Lower gross AI gain	$g = 0.10$	2.223	2.7%	1.28	1.03	1.015
Higher	$g = 0.30$	0.998	100.0%	0.00	0.00	1.552

gross AI gain						
No verification burden	$\tau = 0.000$	1.118	60.5%	6.08	7.04	1.333
Higher verification burden	$\tau = 0.100$	1.701	9.2%	3.52	3.11	1.050

Note. Each row changes only the stated parameter relative to the baseline. Gross gain g and verification burden τ are varied separately rather than being compressed into λ. The condition $k > \sigma - 1$ is maintained. Values are deterministic numerical results, not confidence intervals.

Robustness to an alternative productivity distribution

The Pareto assumption is analytically convenient because it yields closed-form tail integrals, but the concentration result should not depend entirely on that functional form. The robustness exercise therefore replaces Pareto productivity with a lognormal distribution while holding two baseline distributional moments approximately constant. The lognormal parameters are selected so that mean pre-AI sales equals 1.8, the same theoretical mean implied by the baseline Pareto, and the pre-AI sales Gini equals approximately 0.286. With σ = 3, this produces log productivity parameters μ = 0.2269 and s = 0.2589. The same 200,000 quantile grid, fixed-cost levels, gross AI gain, and verification burden are then applied.

Table 6 shows that the substantive non-linearity survives. At the intermediate fixed cost of 0.55, the matched lognormal economy has higher firm-count adoption than the Pareto baseline, 43.9 percent versus 30.1 percent, but still exhibits a sizeable increase in the sales Gini, 6.05 points versus 7.09 points. At higher fixed cost, both distributions show declining concentration effects as adoption becomes sparse. Figure 3 extends the comparison over a continuous fixed-cost grid. The exact peak moves because the two distributions place different mass around the adoption threshold, but both generate the same qualitative hump. This result materially reduces dependence on the Pareto tail assumption while preserving the analytical interpretation of selective diffusion.

Table 6. Distributional Robustness: Pareto and Matched-Dispersion Lognormal

Distribution	FA	Adoption	Baseline Gini	Gini Δ (pts)	Output
Pareto, k=4.5	0.35	83.2%	0.286	2.71	1.291
Matched lognormal	0.35	76.4%	0.286	2.77	1.288
Pareto, k=4.5	0.55	30.1%	0.286	7.09	1.165
Matched lognormal	0.55	43.9%	0.286	6.05	1.207
Pareto, k=4.5	0.85	11.3%	0.286	5.55	1.096
Matched lognormal	0.85	16.0%	0.286	5.33	1.102

Note. The lognormal productivity distribution is matched to the baseline Pareto economy's mean pre-AI sales and sales Gini. It is therefore a shape robustness check rather than a change in initial dispersion.

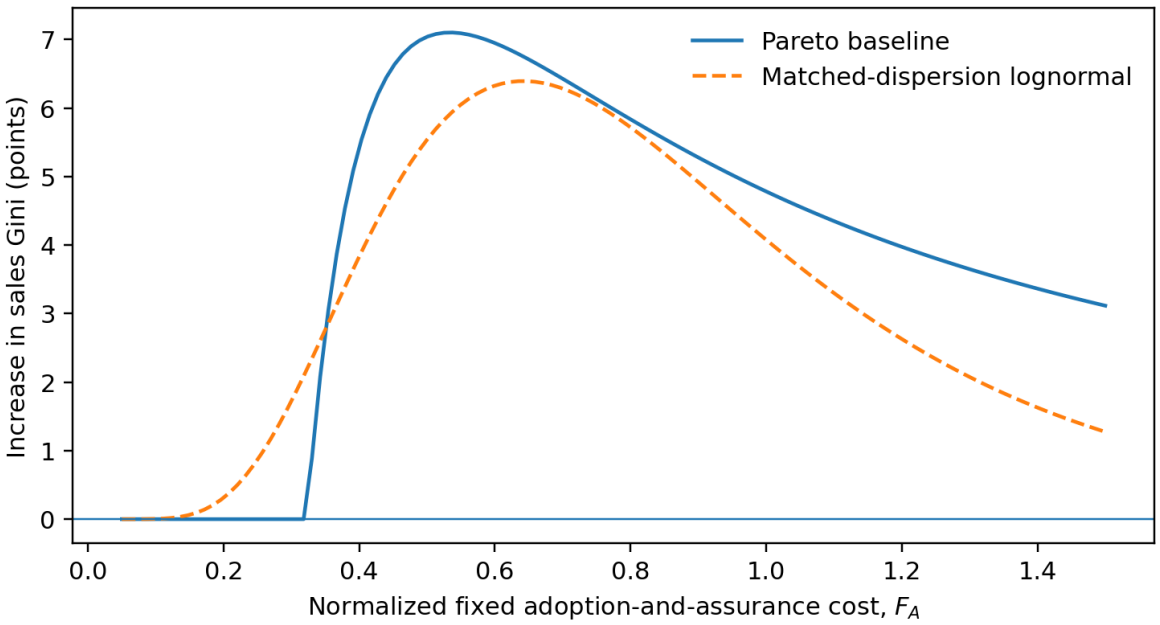


Figure 3. The increase in the sales Gini remains hump-shaped under both the baseline Pareto distribution and a lognormal productivity distribution matched on mean pre-AI sales and baseline sales Gini.

Robustness to heterogeneous AI productivity gains

The baseline assumes a common gross AI productivity gain because this is what permits a single closed-form adoption cutoff. Real firms are unlikely to receive identical gains. To test this restriction without introducing stochastic noise, the next experiment assigns one-third of firms a gross gain of 15 percent, one-third 20 percent, and one-third 25 percent, preserving the same mean gross gain of 20 percent. Three deterministic assignments are considered. In the approximately uncorrelated case, gain types cycle across adjacent productivity quantiles. In the complementarity case, gains rise with productivity. In the catch-up case, gains fall with productivity. Each firm adopts when its own incremental operating surplus exceeds the common fixed cost.

Table 7 shows that heterogeneous gains change the magnitude, but not the economic logic, of the mechanism. When gains are approximately uncorrelated with productivity, the Gini increase remains close to the uniform-gain benchmark. When gains rise with productivity, concentration is stronger because the firms already best positioned to clear the fixed-cost threshold also receive larger productivity improvements. When gains decline with productivity, the concentration effect is substantially attenuated. Thus, the uniform-gain model should not be read as a universal prediction about concentration. Its more robust implication is conditional: concentration pressure is strongest when fixed-cost selection and the distribution of AI gains jointly favor the upper part of the productivity distribution.

Table 7. Robustness to Heterogeneous Gross AI Productivity Gains

Gain structure	Assignment of g_i	Adoption	ΔS (pp)	Gini Δ (pts)	Output
Uniform g_i = 0.20 for all i	g_i = 0.20 for all firms	30.1%	6.91	7.09	1.165
Heterogeneous, uncorrelated	$g_i \in \{0.15, 0.20, 0.25\}$, one-third each	33.7%	7.47	6.82	1.185
Heterogeneous, increasing	Low: 0.15; Middle: 0.20; High: 0.25	33.3%	8.73	9.10	1.236
Heterogeneous, decreasing	Low: 0.25; Middle: 0.20; High: 0.15	12.0%	4.29	3.90	1.066

Note. All heterogeneous-gain cases have the same arithmetic mean gross gain, g = 0.20, use FA = 0.55 and τ = 0.0435, and preserve the same baseline Pareto productivity grid.

Aggregate private adoption surplus

The model also yields a closed-form expression for aggregate private adoption surplus. For an interior cutoff, integrate the incremental operating surplus of adopters over the Pareto tail and subtract the fixed cost paid by each adopter. Substituting the private adoption threshold simplifies the result to a strictly positive expression proportional to z_A raised to $\sigma - 1 - k$. Since $k > \sigma - 1$, lowering the cutoff increases aggregate private surplus. This result clarifies why broad diffusion can raise total private gains even when the per-adopter productivity multiplier is unchanged. It also explains the shared-assurance result: reducing duplicated fixed expense can increase the number of positive-surplus adoption opportunities.

$$\Delta\Pi = B(q - 1){z_{\min}}^{k}{z_A}^{\sigma - 1 - k}\left[\frac{\sigma - 1}{k - (\sigma - 1)}\right] > 0 \qquad (13)$$

DISCUSSION

Why the paradox is an incidence problem, not an anti-regulation argument

The strongest interpretation of the model is about incidence. AI assurance can create social value by reducing expected harm, improving reliability, protecting data, and making delegated decisions more accountable. Nothing in the model implies that firms should be exempt from appropriate controls. Instead, the model shows that two assurance regimes with the same expected risk outcome can have different market-structure effects if one requires every firm to replicate large fixed investments while the other supplies reusable assurance inputs.

This distinction is familiar in other regulated sectors. Shared standards, common testing protocols, accredited laboratories, central utilities, and standardized reporting can reduce duplicated compliance effort without changing the substantive risk threshold. AI is particularly suitable for this perspective because many assurance tasks are partially codifiable. Evaluation datasets, benchmark harnesses, model cards, audit logs, incident taxonomies, policy templates, access-control patterns, and monitoring components can often be reused or supplied as shared infrastructure. The economic case for reuse becomes stronger when the fixed setup burden is material relative to firm scale.

AI can be simultaneously democratizing and concentrating

The model reconciles apparently conflicting narratives about AI and competition. AI can democratize capabilities by lowering marginal costs, as emphasized in downstream competition analysis (OECD, 2025a). Yet the same technology can concentrate realized gains if complementary fixed costs screen out a large portion of firms. The direction of the net competition effect therefore depends on both the marginal-cost channel and the adoption-selection channel.

This framework also complements, rather than duplicates, upstream analyses such as Korinek and Vipra (2025). Upstream concentration arises from compute, data, model-development scale, and feedback dynamics. The mechanism here operates downstream even if every firm can purchase the same model at the same posted price. A downstream concentration effect can arise simply because firms differ in their capacity to integrate and assure the technology. Upstream competition policy alone therefore cannot guarantee broad downstream diffusion.

The importance of the extensive margin

Productivity debates often focus on the intensive margin: how much faster or cheaper an adopting worker, task, or firm becomes. The present results show why the extensive margin deserves equal attention. A 25 percent productivity improvement for 10 percent of firms can have different aggregate and distributional consequences from a 12 percent improvement for all firms. The shared-assurance scenario makes this point

starkly. Its per-adopter multiplier is lower than the fragmented benchmark, yet aggregate technology-adjusted output is higher because adoption becomes universal.

This observation is consistent with the broader productivity-diffusion literature. Frontier firms can continue improving while laggards fail to absorb new technologies, creating aggregate productivity shortfalls and widening dispersion (Andrews et al., 2016). The AI assurance cost paradox offers one microeconomic pathway through which such divergence could persist even when AI tools themselves become inexpensive.

Why intermediate fixed costs are especially important

The non-linearity has practical significance. If fixed adoption costs are extremely low, diffusion can be broad and the same productivity multiplier is shared. If costs are prohibitively high, AI contributes little to total industry transformation because adoption is rare. The most consequential region is intermediate, where a substantial but selective group adopts. In that region, the technology is important enough to reallocate market share but not accessible enough to diffuse the productivity gain widely.

This explains why policies that merely subsidize AI purchases can have ambiguous concentration effects. Starting from very low diffusion, a modest subsidy can move more upper-tail firms into adoption and increase the reallocation wedge before the subsidy is large enough to reach smaller firms. The objective should therefore be broad access to the entire bundle of complementary capabilities, not only lower software prices.

What the robustness analysis establishes

The robustness exercises sharpen the scope of the paper's claims. The selective-diffusion mechanism is not an artifact of one fixed-cost value, one Pareto tail parameter, or one elasticity. It also survives a lognormal productivity distribution matched on initial sales dispersion. However, concentration is quantitatively sensitive to the joint distribution of productivity and AI gains. This is economically informative rather than problematic: the model predicts a mechanism, not a universal sign for every possible technology shock. Empirical applications should therefore estimate both the incidence of fixed adoption and assurance costs and the covariance between firm characteristics and realized AI productivity gains. These objects are central targets for future identification.

POLICY IMPLICATIONS

Evaluate the architecture of compliance costs

Regulatory impact assessments for AI should distinguish fixed, quasi-fixed, and variable assurance costs. A requirement that looks proportionate in absolute terms may be regressive relative to firm operating surplus. Policymakers should therefore ask how much of the burden is incurred once per organization, once per model, once per use case, or per transaction. This decomposition can reveal whether a rule unintentionally shifts adoption toward incumbents even when the substantive requirement is justified.

Build shared assurance infrastructure

The most direct implication is to convert duplicative fixed costs into shared or scalable inputs where technically and legally feasible. Possible mechanisms include sectoral AI evaluation centers, accredited shared testing services, common incident-reporting infrastructure, reusable audit schemas, standardized vendor due-diligence packages, pooled red-team resources for SMEs, and reference implementations of logging and human-approval controls. Such mechanisms should be designed to preserve independence and avoid conflicts of interest. The point is not to socialize private operational risk indiscriminately, but to prevent thousands of firms from rebuilding the same assurance primitives.

Use risk-proportional obligations rather than size-blind fixed obligations

Where residual risk scales with deployment volume, autonomy, data sensitivity, or consequence severity, a

variable or risk-proportional assurance component may be economically preferable to a large size-blind fixed charge. Equation (12) formalizes the tradeoff. A moderate increase in proportional checking can be diffusion-enhancing if it allows a sufficiently large reduction in fixed cost. This design principle is especially relevant for low-risk applications where a full enterprise-specific assurance architecture would be disproportionate.

Support complementary capabilities, not just model access

Official evidence indicates that expertise, data maturity, and regulatory uncertainty are meaningful obstacles to AI use (Eurostat, 2026; OECD/BCG/INSEAD, 2025). AI diffusion policy should therefore bundle model access with capability-building. Public testbeds, advisory programs, workforce training, data-governance templates, and procurement guidance can lower effective fixed costs. The OECD's review of AI adoption institutions reaches a similar practical conclusion: dedicated public institutions can facilitate diffusion through demonstrations, advice, networking, and access to relevant information (OECD/BCG/INSEAD, 2025).

Competition authorities should monitor downstream adoption asymmetry

Competition analysis of AI should track more than upstream provider shares. Useful downstream indicators include adoption by firm-size class, the distribution of AI intensity across firms, the share of sales generated by adopters, switching and multi-homing costs, and whether compliance or assurance expenditures scale with firm size. A rapid increase in AI adoption accompanied by rising concentration may reflect efficient reallocation, exclusionary conduct, fixed-cost selection, or a combination. Distinguishing these mechanisms is necessary before policy intervention.

Preserve safety while improving cost efficiency

A final policy principle is risk equivalence. Shared assurance should be evaluated by whether it can achieve the same or better residual-risk target at lower duplicated fixed cost. This avoids a false choice between innovation and safety. In the model, the relevant comparison is not regulated versus unregulated AI. It is fragmented versus shared production of assurance at a common risk standard.

LIMITATIONS AND RESEARCH AGENDA

The paper is theoretical and simulation-based. It does not estimate the causal effect of assurance costs on adoption, and the current firm-size adoption gradients cannot identify that effect. Large and small firms differ simultaneously in management quality, sector, digital intensity, workforce composition, financing, data assets, risk exposure, and expected return from AI. The empirical section therefore treats official statistics only as descriptive evidence of heterogeneous diffusion. A causal test would require firm-level measures of integration and assurance expenditure, plausibly exogenous cost variation, or a research design around policy thresholds or shared-assurance interventions.

The model also abstracts from endogenous entry and exit. In a full Melitz-style general equilibrium, the productivity advantage of adopters would change the price index, revenue of non-adopters, survival thresholds, and potentially the mass of entrants. Those effects could strengthen the concentration mechanism by pushing marginal non-adopters out, or weaken it if AI lowers entry costs for new firms. A natural next step is to embed the assurance cutoff in an entry model with free entry and endogenous market size.

A second extension concerns endogenous AI-provider pricing. Upstream providers may use nonlinear tariffs, bundled cloud services, volume discounts, or differentiated access tiers. Such pricing can interact with fixed downstream assurance costs. A two-sided model could examine when model-provider competition offsets or amplifies the fixed-cost wedge.

Third, the model separates fixed assurance resources F_G from proportional verification burden τ, but the assurance production function remains deliberately reduced-form. In practice, assurance quality is multidimensional and may exhibit economies of scale, learning, network effects, and externalities. Future research should estimate an assurance production function linking spending and governance practices to

measurable residual risk. That would make it possible to solve the risk-constrained policy problem in Equation (11) quantitatively rather than symbolically.

Fourth, the numerical parameterization remains illustrative even after the expanded robustness analysis. The baseline is transparent and literature-informed, and the central hump survives variation in fixed cost, elasticity, Pareto shape, gross AI gains, verification burdens, and a matched lognormal productivity distribution. These checks reduce dependence on arbitrary functional-form choices, but they do not transform the exercise into an empirical calibration. The next empirical step is to combine firm-level AI adoption data with measures of integration expenditure, compliance, cybersecurity, model risk, and assurance spending. Difference-in-differences or regression-discontinuity designs around regulatory thresholds, sector-specific mandates, public AI testbeds, or shared-assurance programs could test whether lower fixed assurance costs disproportionately increase adoption among smaller or lower-productivity firms. The core causal prediction is not simply that subsidies increase adoption. It is that reductions in fixed cost should have stronger extensive-margin effects lower in the productivity or size distribution than economically equivalent reductions in proportional usage cost.

Fifth, the heterogeneous-gain exercise uses deliberately simple deterministic schedules. It shows that the covariance between productivity and AI gains can amplify or attenuate the concentration channel, but it does not estimate that covariance. Future work should measure whether complementary data, management, and workforce capabilities cause high-productivity firms to obtain systematically larger AI gains, or whether AI instead produces catch-up effects among lagging firms. That empirical relationship is likely to be as important as the fixed-cost distribution itself.

Finally, concentration is not synonymous with welfare loss. High-productivity firms gaining share can raise allocative efficiency, and AI-driven concentration may coexist with lower prices or better products. Future welfare analysis should combine consumer surplus, innovation, residual AI risk, entry, and market power. The present model identifies a reallocation mechanism that deserves measurement, not a presumption that every increase in concentration is harmful.

CONCLUSION

Artificial intelligence is unusual because it can make sophisticated capabilities cheap at the margin while still requiring substantial organizational infrastructure before firms can use those capabilities safely and reliably. This paper shows that the distinction between marginal and fixed costs is central to the economics of AI diffusion. In a heterogeneous-firm economy, fixed adoption and assurance costs create a productivity threshold. Firms above the threshold receive the AI productivity multiplier; firms below it do not. The resulting selection reallocates revenue toward adopters and can increase market concentration even when the underlying AI technology is broadly available.

The paper derives a closed-form concentration wedge and shows that its relationship with diffusion is hump-shaped under the baseline structure. A deterministic 200,000-firm experiment confirms the same broad non-linearity using the sales Gini. The expanded sensitivity analysis shows that the selection mechanism survives substantial variation in fixed cost, substitution elasticity, Pareto shape, gross AI gains, and verification burden. Replacing Pareto productivity with a matched-dispersion lognormal distribution preserves the hump, while moderate heterogeneous AI gains show that concentration becomes stronger when gains are complementary to productivity and weaker when gains favor lower-productivity firms. The paper therefore makes a conditional rather than universal claim: market-structure effects depend on who can clear the fixed-cost threshold and on where the realized AI gains occur in the firm distribution.

The central policy message is therefore architectural. AI governance should be evaluated not only by how much assurance it requires, but by how the cost of achieving assurance is produced and distributed. Requiring every firm to duplicate large fixed investments can make safe adoption regressive with respect to firm productivity. Shared testing, reusable controls, common certification infrastructure, public testbeds, and sectoral assurance utilities can potentially preserve risk standards while lowering the fixed-cost barrier. In that

sense, the most pro-competitive form of AI governance may be neither weaker governance nor indiscriminate subsidy, but assurance designed as scalable infrastructure.

This perspective also reframes the broader AI productivity debate. Aggregate gains depend not only on what AI can do for an adopter, but on how many firms can become adopters without compromising safety. The economics of AI is therefore partly the economics of assurance diffusion.

DECLARATIONS

Funding: No external funding was received for this study.

Conflict of Interest: The author declares no conflict of interest.

Ethics Approval: Not applicable. The study uses analytical modelling, public aggregate statistics, and synthetic numerical data; it involves no human participants or identifiable personal data.

Data and Reproducibility: All numerical results are generated from the equations, parameter values, and deterministic assignment rules reported in the manuscript. The baseline and alternative-distribution experiments use 200,000 equally spaced quantiles. The heterogeneous-gain cases use deterministic gain assignments rather than random draws. No random seed is required. Appendix B reports the baseline algorithm and Appendix C reports the alternative-distribution and heterogeneous-gain construction.

Author Contribution: Kwan Hong TAN is the sole author and was responsible for conceptualization, methodology, formal analysis, numerical analysis, interpretation, and writing.

REFERENCES

1. Acemoglu, D. (2025). The simple macroeconomics of AI. Economic Policy, 40(121), 13-58. https://doi.org/10.1093/epolic/eiae042
2. Andrews, D., Criscuolo, C., & Gal, P. N. (2016). The best versus the rest: The global productivity slowdown, divergence across firms and the role of public policy. OECD Productivity Working Papers, No. 5. OECD Publishing. https://doi.org/10.1787/63629cc9-en
3. Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. The Quarterly Journal of Economics, 135(2), 645-709. https://doi.org/10.1093/qje/qjaa004
4. Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. Journal of Financial Economics, 151, 103745. https://doi.org/10.1016/j.jfineco.2023.103745
5. Bonney, K., Breaux, C., Dinlersoz, E., Foster, L., Haltiwanger, J., & Pande, A. (2026). The microstructure of AI diffusion: Evidence from firms, business functions, and worker tasks (CES Working Paper 26-25). U.S. Census Bureau, Center for Economic Studies. https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-25.html
6. Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. The Quarterly Journal of Economics, 140(2), 889-942. https://doi.org/10.1093/qje/qjae044
7. De Loecker, J., Eeckhout, J., & Unger, G. (2020). The rise of market power and the macroeconomic implications. The Quarterly Journal of Economics, 135(2), 561-644. https://doi.org/10.1093/qje/qjz041
8. European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union. https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng
9. Eurostat. (2026). Use of artificial intelligence in enterprises. Statistics Explained. Data for 2025. Retrieved August 19, 2026, from https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Use_of_artificial_intelligence_in_enterprises
10. Filippucci, F., Gal, P., & Schief, M. (2026). Aggregate productivity gains from artificial intelligence: A sectoral perspective. AEA Papers and Proceedings, 116, 31-35. https://doi.org/10.1257/pandp.20261035

11. Gans, J. S. (2025). How learning about harms impacts the optimal rate of artificial intelligence adoption. Economic Policy, 40(121), 199-219. https://doi.org/10.1093/epolic/eiae053
12. Goldfarb, A., & Tucker, C. (2019). Digital economics. Journal of Economic Literature, 57(1), 3-43. https://doi.org/10.1257/jel.20171452
13. Korinek, A., & Vipra, J. (2025). Concentrating intelligence: Scaling and market structure in artificial intelligence. Economic Policy, 40(121), 225-256. https://doi.org/10.1093/epolic/eiae057
14. McElheran, K., Yang, M.-J., Kroff, Z., & Brynjolfsson, E. (2026). The adoption of industrial AI in America. AEA Papers and Proceedings, 116, 20-25. https://doi.org/10.1257/pandp.20261033
15. Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. Econometrica, 71(6), 1695-1725. https://doi.org/10.1111/1468-0262.00467
16. Melitz, M. J., & Redding, S. J. (2015). New trade models, new welfare implications. American Economic Review, 105(3), 1105-1146. https://doi.org/10.1257/aer.20130351
17. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. Science, 381(6654), 187-192. https://doi.org/10.1126/science.adh2586
18. OECD. (2025a). Artificial intelligence and competitive dynamics in downstream markets (OECD Roundtables on Competition Policy Papers, No. 331). OECD Publishing. https://doi.org/10.1787/ccf0624a-en
19. OECD/BCG/INSEAD. (2025). The adoption of artificial intelligence in firms: New evidence for policymaking. OECD Publishing. https://doi.org/10.1787/f9ef33c3-en
20. Tabassi, E. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1
21. Tan, K. H. (2026a). The coordination compression trap: A dynamic model of AI-enabled productivity, verification debt, and organisational resilience in global knowledge firms. Journal of Global Economics, Management and Business Research, 18(3), 256-278. https://doi.org/10.56557/jgembr/2026/v18i310956
22. Tan, K. H. (2026b). Runtime assurance for enterprise agentic AI systems: A policy-gated control model with quantitative autonomy-risk scoring. World Journal of Advanced Research and Reviews, 31(1), 512-522. https://doi.org/10.30574/wjarr.2026.31.1.1872
23. U.S. Census Bureau. (2026, May 26). Large firms with at least 20 employees biggest AI users. https://www.census.gov/library/stories/2026/05/ai-use-businesses.html

Appendix A. Derivation of the Concentration Turning Point

Starting from Equation (9), let the positive constant a equal $q - 1$ and suppress the adoption subscript for readability. Differentiating the concentration wedge with respect to the adopter revenue mass gives:

$$\frac{d\Delta S}{dr} = \frac{a(1 - 2r - ar^2)}{(1 + ar)^2} \tag{A1}$$

The denominator is positive on the unit interval. Setting the numerator equal to zero therefore yields:

$$ar^2 + 2r - 1 = 0 \tag{A2}$$

The economically admissible root can be written in two equivalent forms:

$$r^* = \frac{\sqrt{1 + a} - 1}{a} = \frac{1}{1 + \sqrt{q}} \tag{A3}$$

The second derivative evaluated at this root is negative. Since the wedge is zero at both endpoints and positive in the interior, the root is the unique global maximum. Combining this turning point with Equations (5) and (7) gives the fixed-cost level associated with the analytical peak:

$$F_A^* = B(q - 1)z_{\min}{}^{\sigma - 1}(r^*)^{-\frac{(\sigma - 1)}{(k - (\sigma - 1))}} \tag{A4}$$

Under the baseline parameters, $r^* = 0.4651$ and $F_A{}^* = 0.5950$.

Appendix B. Reproducibility Algorithm

The baseline deterministic synthetic-firm experiment can be reproduced without random sampling. First, set N = 200,000 and construct equally spaced quantile midpoints:

$$u_i = \frac{(i - 0.5)}{N}, \quad i = 1, \ldots, N \tag{B1}$$

Second, map each quantile into Pareto productivity:

$$z_i = \frac{z_{\min}}{(1 - u_i)^{\frac{1}{k}}} \tag{B2}$$

Third, compute baseline sales proportional to:

$$\text{sales}_i \propto {z_i}^{\sigma - 1} \tag{B3}$$

Fourth, set gross AI gain g and proportional verification burden τ, compute λ = (1 + g)/(1 + τ) and q = λ raised to σ - 1, and for each fixed-cost scenario evaluate each firm's adoption condition from Equation (4). Multiply baseline sales by q for adopters and leave non-adopters unchanged. Fifth, normalize sales to shares and compute the Gini coefficient, adopter-share wedge, top-share statistics if desired, and mean technology-adjusted sales. No random seed is required because the quantile grid is deterministic.

Appendix C. Robustness Construction

C.1 Matched lognormal productivity. Let pre-AI sales be proportional to $z^{\sigma-1}$. The baseline Pareto parameters $\sigma = 3$, $k = 4.5$, and $z_{\min} = 1$ imply theoretical mean baseline sales of 1.8 and a sales Gini of 0.2857. For lognormal productivity, ln z is normal with mean μ and standard deviation s. A lognormal variable with log standard deviation s_y has Gini $G = 2\Phi(s_y/\sqrt{2}) - 1$. Matching the baseline sales Gini gives $s_y = 0.51775$ and therefore $s = s_y/(\sigma - 1) = 0.25888$. Matching mean sales then gives $\mu = 0.22688$. Each quantile midpoint u_i is transformed through the inverse standard-normal distribution and mapped to $z_i = \exp(\mu + s\Phi^{-1}(u_i))$. The same adoption rule and fixed-cost scenarios are then applied.

C.2 Heterogeneous AI gains. The productivity grid remains the baseline Pareto grid and $\tau = 0.0435$. Firms receive gross gains g in {0.15, 0.20, 0.25}, with one-third of firms in each gain class. In the approximately uncorrelated assignment, the gain classes cycle over adjacent sorted productivity quantiles. In the complementarity assignment, the lowest, middle, and highest productivity thirds receive gains of 0.15, 0.20, and 0.25 respectively. In the catch-up assignment, the ordering is reversed. Each firm i adopts if $B\, z_i^{\sigma-1}[q_i - 1] \geq F_A$, where $q_i = [(1 + g_i)/(1 + \tau)]^{\sigma-1}$. All assignments are deterministic.