Will AI Adoption Intensify or Weaken Market Competition?*

Hamid Firooz Sylvain Leduc Zheng Liu

March 27, 2026

Abstract

We study how AI adoption affects market competition based on a general equilibrium framework. The model features heterogeneous firms choosing AI technology subject to fixed costs, where AI production requires data and energy inputs. Our model predicts non-monotonic relations between AI adoption and industry concentration. As AI adoption rises from an initially low level, large firms that are incumbent adopters gain market shares. When AI usage is sufficiently diffused, the entry of smaller, new AI adopters erodes the market share of incumbents, reducing industry concentration. This occurs irrespective of whether AI adoption is demand driven (e.g., from lower fixed costs) or supply driven (e.g., from greater AI productivity or data availability). The non-monotonic relations are robust even when firms can complement AI with their own data. Industry-level evidence suggests that the correlation between concentration and AI usage has weakened as AI adoption rises over time, broadly consistent with our model's predictions.

Keywords: Artificial intelligence; data; heterogeneous firms; industry concentration; markup; productivity.

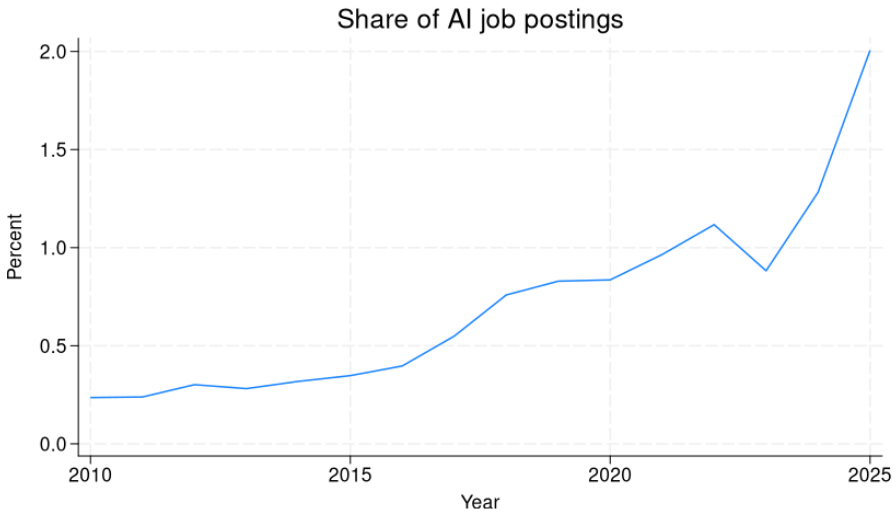
JEL Codes: E24; L11; O33.

*Firooz: North Carolina State University; email: hmdfirooz@gmail.com. Leduc: Federal Reserve Bank of San Francisco; email: Sylvain.Leduc@sf.frb.org. Liu: Federal Reserve Bank of San Francisco; email: Zheng.Liu@sf.frb.org. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

1 Introduction

Recent improvements in artificial intelligence (AI) have the potential to create a general-purpose technology that could persistently boost productivity and economic growth (Baily, Brynjolfsson and Korinek, 2023). New applications have led to promising advancements in fields as varied as information technology, medicine, manufacturing, and research and development, among others. At the aggregate level, the share of AI-related job postings in all online job postings has surged since late 2022 (see Figure 1). Rapid diffusion of AI could lead to large economic disruptions, both in the product markets and the labor markets.

Figure 1. Time series of the share of AI-related job postings



Source: Lightcast and authors' calculation.

Since AI is a relatively nascent technology, it remains unclear how it would impact on market competition. Small businesses perceive AI as a force that could level the playing field against larger firms (e.g., Salesforce SMB Trends Report (2025)). They expect to more easily transition to AI, compared to their larger competitors burdened with larger and more rigid organization structures. This optimism of small businesses stands in contrast to the view that AI is a predictive technology that works better with more data, for which larger firms have an advantage. A natural question thus arises: Will AI intensify or weaken market competition?

In this paper, we examine the impact of data-intensive technological progress, such as AI, on market competition in a macroeconomic framework with heterogeneous firms and variable markups. We introduce an AI-producing sector in the model, where a continuum of AI producers each produce a differentiated AI algorithm using data and

energy as inputs. The economy has a fixed endowment of energy, such that energy prices are endogenously determined by aggregate energy demand of AI producers. Because it is costly to enter this market, the measure of AI producers is endogenous. This is a useful feature to examine how a more competitive AI sector can affect market concentration in downstream industries.

The economy is also populated by a continuum, unit mass of intermediate goods producers, each facing a downward-sloping demand curve for its differentiated product. Final goods are a composite of intermediate goods, with variable elasticities of substitution between varieties determined by the Kimball (1995) preferences, leading to variable markups. Firms in the intermediate goods sector face idiosyncratic productivity shocks which determine their relative sizes. They each choose whether or not to adopt AI, given the presence a fixed cost of adoption. If a firm chooses to adopt AI, then its production uses unskilled labor and an "AI bundle" that combines AI algorithm, the firm's own data, and skilled labor. If a firm does not adopt AI, then its production uses unskilled labor only.

Building on Jones and Tonetti (2020), we assume that producing intermediate goods generates data as a byproduct, which the firms sell to AI producers. Because larger firms produce more data, a data-intensive technology like AI may confer them an inherent advantage, leading to more concentrated industries as larger firms adopt AI more rapidly. Therefore, we also examine environments in which AI can be combined with firms' own data to gain additional efficiencies.

We calibrate the model to U.S. data and match several stylized facts about firms' AI adoption. Based on the calibrated model, we conduct several comparative static exercises to examine the impact of AI adoption on industry concentration in economies with different AI fundamentals.[1] We consider different triggers for AI adoption driven either by AI demand, due to declining fixed cost of adoption, or by AI supply, stemming from improvements in AI productivity or increases in data availability and energy supply.

We find that the relationship between AI adoption and industry concentration is hump-shaped. Concentration initially rises with AI adoption, before falling as adoption becomes more widespread, with a turning point at around 15 percent of the AI adoption rate. This hump-shaped relation does not depend on whether AI adoption is demand or supply driven. It also arises even when firms can complement AI with their own data.

Consider for instance the case in which AI adoption is driven by demand factors, such as a decline in the fixed cost of adoption. The presence of the fixed cost generates

[1] To keep the model tractable, we abstract from dynamics, which are obviously important for studying AI investment and technology adoptions. Our static model still produces rich implications for the relations between industry concentration and AI adoption, as we show below.

an economy-of-scale effect such that larger firms are more likely to adopt AI. A decline in the fixed cost leads to two opposing effects on adoption: it could benefit large firms that are incumbent users of AI, enabling them to further expand their market share (an intensive margin effect); It could also attract some smaller non-adopting firms to become an adopter, eroding the market share of large firms (an extentive margin effect). The net effect of a lower fixed cost on industry concentration depends on the current level of AI adoption. In an economy with a high fixed cost of adoption, few firms adopt AI, and a decline in the fixed cost raises industry concentration, reflecting that the intensive-margin effect dominates. As demand for AI increases, the relative price of AI also rises, further deterring the entry of new adopters, reinforcing the increase in industry concentration. However, in an economy with a sufficiently low fixed cost, the AI adoption threshold is low, such that the extensive-margin effect through the entry of new adopters dominates, despite the rise in AI prices. This results in less industry concentration.

A similar non-monotonic pattern arises following supply-driven AI adoption. Industry concentration initially rises as the intensive margin of AI usage from incumbents outweighs the extensive margin of AI adoption. However, this non-monotonic relationship is nearly flat past peak industry concentration, as the supply of AI continues to rise. Given fixed costs of AI adoption, supply-driven declines in AI prices do not have a large impact on AI adoption, thus muting the effects of the extensive margin on industry concentration.

Our model's predictions are supported by empirical evidence. We use data on job postings to construct a measure of AI intensity across NAICS 5-digit industries, and we combine the AI intensity data with measures of industry concentration for U.S. public firms. Using a long-difference approach, we first reproduce the result in Babina et al. (2024) that increases in AI intensity are positively associated with rising industry concentration between 2010 and 2018. This positive relationship is consistent with our model's prediction that, when adoption rates are low, which was the case at the start of our sample, an increase in AI adoption leads to greater concentration. However, estimating this relationship between 2019 and 2023, we find that the relation between AI intensity and industry concentration turned insignificant, and even slightly negative, consistent with the non-monotonic relationship in our model. The more muted relationship between AI adoption and industry concentration that we find since 2019 is consistent with the pattern between these variables when AI adoption is supply driven in the model. Moreover, our calibrated model predicts that the changing relationship occurs when AI adoption reaches roughly 15 percent, which is in line with the US average adoption rate in the later part of our sample.

We use our calibrated model to study the implications of policy interventions. The model features several sources of distortions, including monopolistic competition with variable markups in the goods sector, monopolistic competition in the AI sector, and an externality stemming from the non-rivalry of data usage. Thus, decentralized equilibrium is inefficient and policy interventions might improve welfare. We examine the macroeconomic and welfare consequences of three policies: (1) taxing or subsidizing AI adoption; (2) taxing or subsidizing AI production; and (3) regulating data usage.

2 Related literature

Our work is partly motivated by the influential study of Autor et al. (2020), who document evidence of the steady rise of superstar firms in the U.S. economy since the early 2000s. Autor et al. (2020) discuss a few potential drivers of the rise of superstar firms (what they call a "winner takes most" mechanism), including greater market competition (e.g., through offshoring) or scale-biased technological changes driven by intangible capital investment and information technology. Other potential drivers of the rise in industry concentration have been studied in the literature, including uneven productivity growth across firms (Furman and Orszag, 2018; Choi et al., 2024), declines in knowledge diffusion between the frontier and laggard firms (Akcigit and Ates, 2019), a slowdown in radical innovations since the 1990s (Olmstead-Rumsey, 2019), the rise of specialized firms (Ekerdt and Wu, 2022), and digital advertising in customer accumulation (Shen, 2023). Our study focuses on how AI adoptions could affect industry concentration.

In our model, the fixed costs of AI adoption disproportionately favor large firms, contributing to sales concentration. Such economy-of-scale effects of new technology adoption have been explored in other studies, including, for example, Firooz, Liu and Wang (2025), Kwon, Ma and Zimmermann (2024), Aghion et al. (2019), Hubmer and Restrepo (2022), Ridder (2023), Lashkari, Bauer and Boussard (2024), Tambe et al. (2020), and Sui (2022). Relative to those studies, our work makes three contributions.

First, our model predicts a non-monotonic relationship between industry concentration and AI adoption. And we show that this model implication is supported by industry-level evidence. This non-monotone relationship cannot be explained by the scale-economy effect alone; it is instead the net effect of two opposing forces associated with AI adoption: an intensive-margin effect that raises concentration and an extensive-margin effect that reduces it. To our knowledge, this non-linear relation between AI adoption and industry concentration is new to the literature.

Second, our model highlights the importance of data. In our model, data is a byproduct of production in the spirit of Jones and Tonetti (2020). Data is an input factor for

producing AI algorithms, and in addition, it can also augment the efficiency for firms to operate the AI technology. This non-rivalry nature of data leads to an externality that has potential implications for AI-related policy.

Third, we study how alternative policy interventions can affect equilibrium allocations and welfare. In particular, we examine how different policies can affect market competition outcomes—measured by industry concentration and the average markup—when AI adoption increases over time.

3 The Model

To understand the economic mechanism that drives the empirical relations between industry concentration and AI usage, we consider a general equilibrium model with AI investment. The model features heterogeneous firms facing idiosyncratic productivity and monopolistic competition in the product markets. We consider a Kimball demand system such that a firm's markup varies with its relative size. Each firm produces a differentiated intermediate good using unskilled labor and a composite of AI algorithm and skilled labor. AI production requires data and energy as inputs. Energy supply is fixed. Data are generated as a byproduct of firm production. The representative household consumes final goods, which are a Kimball aggregate of the differentiated intermediate goods. The household supplies skilled and unskilled labor to firms and obtains profits from all firms.

3.1 Final goods producers

There is a continuum of monopolistically competitive intermediate good producers, indexed by their productivities ϕ drawn from a cumulative distribution function $G(\cdot)$. Final goods producers make a composite homogeneous good out of intermediate varieties and sell it to consumers in a perfectly competitive market, with the final goods price normalized to one. The final good Y is produced using a bundle of intermediate goods $y(\phi)$, according to the Kimball aggregator

$$\int_0^1 \Lambda(\frac{y(\phi)}{Y})dG(\phi) = 1. \tag{1}$$

Denote the relative output of variety ϕ by $q(\phi) := \frac{y(\phi)}{Y}$. Taking the intermediate goods price $p(\phi)$ as given, the cost-minimizing problem of final good producers delivers

the following demand schedule for intermediate variety ϕ

$$p(\phi) = \Lambda'(q(\phi))H, \tag{2}$$

where H is a demand shifter given by

$$H = \left(\int \Lambda'(q(\phi))q(\phi)dG(\phi) \right)^{-1}. \tag{3}$$

We follow Klenow and Willis (2016) and assume that[2]

$$\Lambda(q) = 1 + (\sigma - 1)\exp(\frac{1}{\varepsilon})\varepsilon^{\frac{\sigma}{\varepsilon}-1}[\Gamma(\frac{\sigma}{\varepsilon}, \frac{1}{\varepsilon}) - \Gamma(\frac{\sigma}{\varepsilon}, \frac{q^{\varepsilon/\sigma}}{\varepsilon})], \tag{4}$$

with $\sigma > 1$, $\varepsilon \geq 0$, and $\Gamma(s, x)$ denoting the upper incomplete Gamma function

$$\Gamma(s, x) = \int_x^{\infty} v^{s-1}e^{-v}dv. \tag{5}$$

Under the specification (4), we obtain

$$\Lambda'(q(\phi)) = \frac{\sigma - 1}{\sigma}\exp(\frac{1 - q(\phi)^{\frac{\varepsilon}{\sigma}}}{\varepsilon}), \tag{6}$$

which, using the demand schedule (2), implies that the demand elasticity (i.e., price elasticity of demand) faced by intermediate producer ϕ is

$$\sigma(q(\phi)) = -\frac{\Lambda'(q(\phi))}{\Lambda''(q(\phi))q(\phi)} = \sigma q(\phi)^{-\frac{\varepsilon}{\sigma}}. \tag{7}$$

Given this demand elasticity, the firm with relative production $q(\phi)$ charges a markup

$$\mu(\phi) = \frac{\sigma(q(\phi))}{\sigma(q(\phi)) - 1}. \tag{8}$$

As a result, larger firms face lower demand elasticities, have more market power, and charge higher markups.[3]

[2] Aruoba et al. (2024) employ a different Kimball aggregator and show that their calibrated model is consistent with the markup distribution in the U.S.

[3] We make the technical assumption that $q(\phi) < \sigma^{\frac{\sigma}{\varepsilon}}$ such that the effective demand elasticity is always greater than one. This assumption ensures a well-defined equilibrium under monopolistic competition. In our quantitative analysis, we find that this constraint is never binding.

3.2 Intermediate goods producers

Intermediate good producers have a fixed unit mass. They have two options for production: using AI in their production or not. To use AI, they need to pay a fixed cost $\Omega_a \phi$ in final goods unit. In this case, production takes place using AI, skilled labor, and unskilled labor as inputs, according to the production function

$$y(\phi) = \phi \left[\alpha_a \left((A(\phi)^{1-\beta} d(\phi)^{\beta})^{\gamma} s(\phi)^{1-\gamma} \right)^{\frac{\eta-1}{\eta}} + (1-\alpha_a) u(\phi)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \tag{9}$$

This production function takes into account the possibility that a firm can use its own data (d) to improve the efficiency of AI (A), since the algorithm can be better trained with more data. The parameter $\beta \in [0, 1)$ measures the importance of data for training the algorithm. This specification implies that AI and data are complementary at the firm level, consistent with the evidence of Mihet, Rishabh and Gomes (2026). The firm hires skilled labor (s) to operate AI. Thus, skilled labor in our model can be interpreted as labor with AI skills in the Lightcast data. The AI bundle, that is, the composite of AI, data, and skilled labor, is in turn a substitutable input for unskilled labor (u). The parameter $\eta > 1$ measures the elasticity of substitution between unskilled labor and the AI bundle. The parameter $0 < \gamma < 1$ denotes the cost share of AI and data in the AI bundle and the parameter α_a measures the relative importance of AI bundle in production.

If a firm does not pay the fixed cost of adopting AI, the production function would correspond to the special case with $A(\phi) = 0$. In particular, a non-AI user operates the production technology

$$y(\phi) = \phi \left[(1-\alpha_a) \right]^{\frac{\eta}{\eta-1}} u(\phi). \tag{10}$$

Data are a bi-product of production, such that

$$d(\phi) = \lambda y(\phi), \tag{11}$$

where we assume that the elasticity of data generation with respect to production is equal to one, according to Jones and Tonetti (2020). An AI-using firm can use its own data to augment (or train) the AI algorithm, which effectively boosts productivity. A non-AI user does not benefit directly from its data, although every firm can sell the data in a competitive market at the price P_d. This last assumption reflects the non-rivalry feature of data.

The intermediate producer ϕ that uses AI in production solves the following profit

maximization problem

$$\pi_a(\phi) = \max_{p,y,d,u,s,A} \left[p(\phi)y(\phi) + P_d d(\phi) - w_u u(\phi) - w_s s(\phi) - P_a A(\phi) \right], \tag{12}$$

subject to the demand schedule (2) and the production function (9), taking as given the relative price of data, the relative price of AI, and the wage rates.

In Appendix D, we derive the marginal cost and the conditional factor demand functions. We denote the marginal cost for a firm that adopts AI by $MC_a(\phi)$.

The intermediate producer ϕ that does not use AI solves the following problem

$$\pi_n(\phi) = \max_{p,y,d,u} \left[p(\phi)y(\phi) + P_d d(\phi) - w_u u(\phi) \right], \tag{13}$$

subject to the same demand schedule, the production function with $A(\phi) = 0$, and demand for data to be specified. Similarly, this producer charges a markup $\mu(\phi)$ over its marginal cost, which is given by

$$MC_n(\phi) = \frac{w_u}{\phi \left(1 - \alpha_a\right)^{\frac{\eta}{\eta-1}}}. \tag{14}$$

A firm with productivity ϕ chooses the production technology to maximize profit. Equilibrium profit is given by

$$\pi(\phi) = \max\{\pi_a(\phi) - \Omega_a \phi, \pi_n(\phi)\}, \tag{15}$$

where the fixed cost of using AI is scaled by firm productivity. The fixed cost of using AI is a reduced form that captures complementary investment (e.g., in equipment and human talent) required to implement AI. It also captures the cost of intangible assets, such as concerns of privacy and data safety.

Since operating the AI technology requires paying a fixed cost, we conjecture that there exists a threshold productivity level ϕ^*, such that a firm operates the AI technology if and only if $\phi \geq \phi^*$.

The profit-maximizing decisions for a firm that uses AI and one that does not imply, respectively, that

$$\frac{\Lambda'(q(\phi))H}{\mu(\phi)} + \lambda P_d = MC_a(\phi), \tag{16}$$

$$\frac{\Lambda'(q(\phi))H}{\mu(\phi)} + \lambda P_d = MC_n(\phi). \tag{17}$$

3.3 AI producers

There is a continuum of AI producers of measure M, each producing a differentiated AI variety indexed by $\nu \in [0, M]$. With entry costs of AI producers, the measure of AI producers M is endogenous.

3.3.1 The AI bundle

The AI algorithm used by intermediate good producers is a CES aggregate over differentiated varieties, with the aggregation function

$$A = \left[\int_0^M a(\nu)^{\frac{\kappa-1}{\kappa}} d\nu \right]^{\frac{\kappa}{\kappa-1}}. \tag{18}$$

The representative AI aggregating firm faces perfect competition and takes as given the price of the AI bundle P_a and the price of the individual AI variety $p_a(\nu)$. The optimizing decisions of the AI aggregating firms imply the downward-sloping demand schedule for each AI variety

$$a(\nu) = \left(\frac{p_a(\nu)}{P_a} \right)^{-\kappa} A, \tag{19}$$

The zero-profit condition implies that

$$P_a = \left[\int_0^M p_a(\nu)^{1-\kappa} d\nu \right]^{\frac{1}{1-\kappa}}. \tag{20}$$

3.3.2 The AI varieties

The AI algorithm of variety ν is produced using data and energy as inputs, according to the CES production function

$$a(\nu) = Z \left[\alpha E(\nu)^{\frac{\zeta-1}{\zeta}} + (1-\alpha) D(\nu)^{\frac{\zeta-1}{\zeta}} \right]^{\frac{\zeta}{\zeta-1}}, \tag{21}$$

where Z is an exogenous productivity specific to the AI sector, $E(\nu)$ is the energy input, and $D(\nu)$ is the data input. The parameter ζ denotes the elasticity of substitution between the two types of inputs and $0 < \alpha < 1$ governs the energy intensity of AI production.

The AI producers face perfect competition in the input markets and monopolistic competition in the product markets. Each AI producer takes the energy price P_e and the data price P_d as given, and chooses input quantities to solve the cost-minimizing

problem

$$\min \quad P_e E(\nu) + P_d D(\nu),$$

subject to the production function (21). The solution implies the conditional factor demand functions

$$D(\nu) = \left[\frac{P_d}{\lambda_a(1-\alpha)}\right]^{-\zeta} Z^{\zeta-1} a(\nu), \tag{22}$$

$$E(\nu) = \left[\frac{P_e}{\lambda_a \alpha}\right]^{-\zeta} Z^{\zeta-1} a(\nu), \tag{23}$$

where λ_a denotes the marginal cost of AI production, which is given by

$$\lambda_a = \frac{1}{Z}\left[\alpha^{\zeta} P_e^{1-\zeta} + (1-\alpha)^{\zeta} P_d^{1-\zeta}\right]^{\frac{1}{1-\zeta}}. \tag{24}$$

The AI producer ν chooses its price $p_a(\nu)$ to maximize profit, taking as given the demand schedule (19). The optimizing decisions imply that

$$p_a(\nu) = \frac{\kappa}{\kappa - 1}\lambda_a \equiv p_a. \tag{25}$$

Thus, in a symmetric equilibrium, all AI producers set the same price p_a. In such an equilibrium, the price index of AI algorithms from Eq. (20) is given by

$$P_a = M^{\frac{1}{1-\kappa}} p_a. \tag{26}$$

Since $\kappa > 1$, an increase in the number of AI producers (i.e., an increase in M) would reduce the price P_a of the AI bundle used by intermediate goods producers. Furthermore, the demand function for AI algorithms in Eq. (19) implies that

$$a(\nu) = M^{\frac{\kappa}{1-\kappa}} A \equiv a. \tag{27}$$

3.4 Free entry condition for AI producers

AI producers face entry cost $\mathcal{F}$. Free entry implies that the profit of active AI producers should equal to the entry cost. Thus, in a symmetric equilibrium,

$$(p_a - \lambda_a)a = \mathcal{F}. \tag{28}$$

From equations (25) and (27), this free-entry condition implies that

$$\frac{\lambda_a}{\kappa - 1} M^{\frac{\kappa}{1-\kappa}} A = \mathcal{F}. \tag{29}$$

This equation pins down the equilibrium mass of AI producers, M.

3.5 Households

The representative household has the following utility function

$$U = \ln C - \chi_u \frac{u^{1+\xi_u}}{1+\xi_u} - \chi_s \frac{s^{1+\xi_s}}{1+\xi_s}, \tag{30}$$

where C denotes consumption at time t, $\xi_u, \xi_s \geq 0$ are inverse Frisch elasticities of labor supply, and $\chi_u, \chi_s > 0$ are weights on the disutility from working.[4]

The household is a hand-to-mouth consumer, facing the budget constraint

$$C = w_u u + w_s s + \Pi, \tag{31}$$

where Π denotes total profits from intermediate goods producers and AI producers:

$$\Pi = \int_0^{\phi^*} \pi_n(\phi) dG(\phi) + \int_{\phi^*}^1 \pi_a(\phi) dG(\phi) + \int_0^M [\pi_{ai}(\nu) - \mathcal{F}] d\nu, \tag{32}$$

with the AI producer's profit given by $\pi_{ai}(\nu) \equiv [p_a(\nu) - \lambda_a] a(\nu)$.

The household takes wages w_u and w_s as given and maximizes the utility function (30) subject to the budget constraint (31). The optimizing consumption-leisure choice implies the labor supply equations

$$w_u = \chi_u u^{\xi_u} C, \tag{33}$$

$$w_s = \chi_s s^{\xi_s} C. \tag{34}$$

4 Market clearing and equilibrium

The equilibrium consists of aggregate allocations C, u, s, Y, and A, aggregate prices P_a, P_d, P_e, w_u, and w_s,, the mass of AI producers M, intermediate goods producers' choice variables $A(\phi)$, $s(\phi)$, $u_a(\phi)$, $u_n(\phi)$, $y_a(\phi)$, $y_n(\phi)$, $p_a(\phi)$, $p_n(\phi)$, and the AI producers'

[4]The utility function implies imperfect substitution between skilled and unskilled labor, which is a reduced form that captures costly skill transformation.

decision variables $a(v)$, $D(v)$, $E(v)$, and $p_a(v)$ such that (i) the representative household maximizes its utility; (ii) goods producing firms maximize their profits; (iii) AI producers maximize their profits; (iv) the markets for the final goods, the AI bundle, data, energy, and labor all clear; and (v) the free entry condition for AI producers is satisfied.

Final goods market clearing implies that

$$C + P_e \bar{E} + (1 - G(\phi^*))\Omega_a + M\mathcal{F} = Y, \tag{35}$$

where $\bar{E}$ denotes the exogenous supply of energy in the economy.

Labor market clearing implies that

$$u = \int_0^{\phi^*} u_n(\phi) dG(\phi) + \int_{\phi^*}^1 u_a(\phi) dG(\phi) \tag{36}$$

$$s = \int_{\phi^*}^1 s(\phi) dG(\phi). \tag{37}$$

AI market clearing implies that

$$\int_{\phi^*}^1 A(\phi) dG(\phi) = A. \tag{38}$$

Data market clearing implies that

$$\int_0^M D(v) dv = \int_0^1 d(\phi) dG(\phi) \equiv D, \tag{39}$$

which, using the data demand function (22), can be written as

$$\left[\frac{P_d}{\lambda_a(1-\alpha)}\right]^{-\zeta} M Z^{\zeta-1} a = D. \tag{40}$$

Energy market clearing implies that

$$\int_0^M E(v) dv = \bar{E}, \tag{41}$$

which, using the energy demand in (23), can be written as

$$\left[\frac{P_e}{\lambda_a \alpha}\right]^{-\zeta} M Z^{\zeta-1} a = \bar{E}. \tag{42}$$

5 Calibration

Table 1. Calibrated Parameters and Matched Moments

Parameter	Notation	Value	Sources/Matched Moments
Panel A: Parameters calibrated to match external sources			
Energy weight in AI	α	0.2	
Elas. of subs. b/w energy and data	ζ	0.5	
Weight of AI vs skilled labor	γ	0.3	
Elas. of subs. b/w AI varieties	κ	3	
Energy supply	$\bar{E}$	1	
Scale of data generation	λ	1	Jones and Tonetti (2020)
Productivity standard deviation	σ_ϕ	0.1	Bloom et al. (2018)
Demand elasticity parameter	σ	10.86	Edmond, Midrigan and Xu (2023)
Super elasticity	ϵ	1.73	Edmond, Midrigan and Xu (2023)
Entry cost of AI production	$\mathcal{F}$	0.5	
Inverse Frisch elasticity (unskilled)	ξ_u	0.5	
Inverse Frisch elasticity (skilled)	ξ_s	0.5	
Disutility of work (unskilled)	χ_u	1	Normalization
Panel B: Parameters calibrated to match moments in data			
AI input weight	α_a	0.4069	share of AI jobs
Disutility of work (skilled)	χ_s	12.6652	skill premium
Elas. of subs. b/w AI and labor	η	2.9243	growth of AI jobs (3.8×)/fall in AI price (26×)
AI productivity	Z	0.4908	Employment share of AI adopters
AI adoption fixed cost	Ω_a	0.0274	fraction of AI adopters

Panel C: Matched Moments

Moments	Data	Model
share of AI jobs	1.6%	1.4%
rise of AI jobs /fall in AI price	0.147	0.1650
Employment share of AI adopters	40%	41.47%
fraction of AI adopters	20%	13.1%
skill premium	1.65	1.51
rise of AI adopters /fall in AI price	0.20	0.1988

Note: The last row in Panel C shows a non-targeted moment.

We simulate the model based on calibrated parameters. For our baseline model, we assume that firms' internal data do not directly augment the power of AI for production (i.e., $\beta = 0$). This assumption partly reflects that AI is still a nascent technology and firms are reluctant to feed their own data to AI algorithms for privacy or proprietary concerns. In the near future, as those concerns about data security are gradually alleviated (e.g., with new regulations or technological breakthroughs), firms would become more willing to use their own data to train their AI algorithms. Thus, we also consider a case with a higher value of $\beta = 0.3$ to illustrate the role of data in our model mechanism. For our baseline analysis, we fix β at zero.

We calibrate a subset of the parameters based on existing studies. Those parameters include α, the relative weight of energy input for AI production; ζ, the elasticity of substitution between energy and data inputs for AI production; γ, the relative importance of AI input in the composite of AI and skilled labor for goods production; κ,

the elasticity of substitution between AI varieties; $\bar{E}$, the exogenous energy supply; λ, the elasticity of data with respect to production; σ, the demand elasticity parameter; ϵ, the super elasticity in the Kimball preferences; $\mathcal{F}$, the fixed cost of entry into the AI sector; ξ_u and ξ_s, the inverse Frisch elasticity parameters for unskilled and skilled labor, respectively; and χ_u the disutility of working for unskilled labor. We set $\alpha = 0.2$ and $\zeta = 0.5$, such that engery and data are complementary input factors for AI production, and AI production is data intensive. We set $\gamma = 0.3$, such that skilled labor accounts for 70 percent of the cost share in the composite of AI and skilled labor for producing intermediate goods. We set $\kappa = 3$, implying a markup of 50 percent for AI producers. Free entry implies that aggregate markup of all AI producers equals the fixed cost of entry, which we set to $\mathcal{F} = 0.5$. We normalize the energy supply to $\bar{E} = 1$. We calibrate the scale parameter of data generation to $\lambda = 1$, following Jones and Tonetti (2020). We assume a log-normal productivity distribution with a mean normalized to zero and set the standard deviation to $\sigma_\phi = 0.1$, according to the estimation by Bloom et al. (2018).[5] We set the demand elasticity parameter to $\sigma = 10.86$ and the super elasticity to $\epsilon = 1.73$, in line with the micro-level estimates of Edmond, Midrigan and Xu (2023). Finally, we calibrate the inverse Frisch elasticity parameters $\xi_u = 0.5$ and $\xi_s = 5$ and normalize the disutility of working for unskilled labor to $\xi_u = 1$. The calibrated values of these parameters are shown in Table 1 (Panel A).

The remaining parameters are calibrated to match several moments in the micro-level data. These parameters include α_a, the weight of the AI-skill composite in the production function; χ_s, the dis-utility of working for skilled labor; η, the elasticity of substitution between the AI-skill bundle and unskilled labor; Z, the total factor productivity (TFP) for AI production; and Ω_a, the fixed cost of AI adoption. The calibrated values and the descriptions of the targeted moments are shown in Panel B.

We calibrate the input share of AI in the production of intermediate goods (α_a) to target the observed share of AI jobs. We measure the share of AI jobs using data from Lightcast, which covers the near-universe of online job postings. We identify AI-related job postings using textual analysis, following the approach of Acemoglu et al. (2022) and Babina et al. (2024).[6] We detect AI-related jobs by identifying postings that require AI-

[5] Bloom et al. (2018) estimate a two-state Markov-switching model of firm-level volatility. They report that the low-volatility standard deviation is 0.051, while the high-volatility standard deviation is 0.209. Their estimated transition probabilities imply that the unconditional probability of the low-volatility state is 68.7%. Accordingly, the implied average standard deviation is approximately 0.10 (= $0.051 \times 68.7\% + 0.209 \times (1 - 68.7\%)$).

[6] In addition to using the set of keywords for traditional AI-related skills proposed by Acemoglu et al. (2022), we add our own set of keywords for skills related to Generative AI (Gen AI). We created the set of keywords for Gen AI-related skills using a large language model (ChatGPT 5.1). We entered the prompt "I'm interested in identifying job postings that require gen AI related skills. Some keywords such as Gen AI, LLM, or the alike would be relevant. Can you come up with some related keywords?" The specific

or GenAI- related skills. As shown in Figure 1 (Panel (b)), the share of AI jobs was about 1.6 percent in 2024. We calibrate α_a to match this share of AI jobs.

We calibrate the parameter χ_s, the dis-utility of work for skilled workers, to match the average skill wage premium in the data. According to the Household Survey of the Bureau of Labor Statistics, the median weekly earnings of workers with Bachelor's degree is on average about 1.65 times those of high school graduates with no college for the period from 2000 to 2025. Thus, we calibrate χ_s such that the model-predicted skill premium is about 1.65.

We calibrate the elasticity of substitution between AI and labor input (η) to target the ratio of the cumulative growth of AI jobs to the cumulative decline in the relative price of AI from 2010 to 2024. In our Lightcast data, the share of AI jobs grew by 3.8 times from 0.42 percent in 2010 to 1.6 percent in 2024. Over the same period, the AI price, measured by the relative price of GPU, fell by 26 times (Adams et al., 2026).

We calibrate the AI productivity (Z) and the fixed cost of AI adoption (Ω_a) to target the employment share of AI adopters and the frqction of AI adopting firms in the data. Since AI is still a nascent technology, there is a wide range of estimates of the adoption rate across existing surveys, and the adoption rate has been rapidly rising in recent years. The US Census Business Trends and Outlook Survey (BTOS) reports that, as of December 2025, about 18 percent of businesses have used AI "in any of its business functions," and they expected the adoption rate to rise to 21 percent over the next 6 months. The average AI adoption rate in the European union was similar, at about 20 percent (Bick et al., 2026).[7] According to the BTOS, AI adoption rates are higher in larger firms. For example, firms with at least 250 employees had an adoption rate of about 34 percent in December 2025, compared to the average of 18 percent. We calibrate our model parameters to target an average AI adoption rate of 20 percent and an employment share of AI adopting firms of 40 percent, which lie within the range of the survey estimates.

The calibrated model broadly matches the targeted moments in the data, as shown in Panel C. There are some discrepancies between the model predictions and the targeted moments. For example, the model understates the share of firms that adopt AI (13.1% vs. 20%) and also the share of AI jobs (1.4% vs. 1.6%). The model modestly overstates

keywords generated by the LLM and used in our search for GenAI-related skills are described in the Appendix (see Tables A.1 and A.2).

[7]Some other surveys suggest higher AI adoption rates. According to "The AI Index 2025 Annual Report," the proportion of survey respondents reporting AI use by their organizations jumped to 78% in 2024 from 55% in 2023. During the same period, the share of respondents who reported using generative AI in at least one business function more than doubled, jumping from 33% in 2023 to 71% in 2024. Similarly, a survey of senior executives of firms in US, UK, Germany, and Australia conducted by Yotzov et al. (2026) between November 2025 and January 2026 reports that about 70% of firms actively use AI, although those senior executives reported that AI has not yet generated meaningful impact on productivity and employment during the past 3 years.

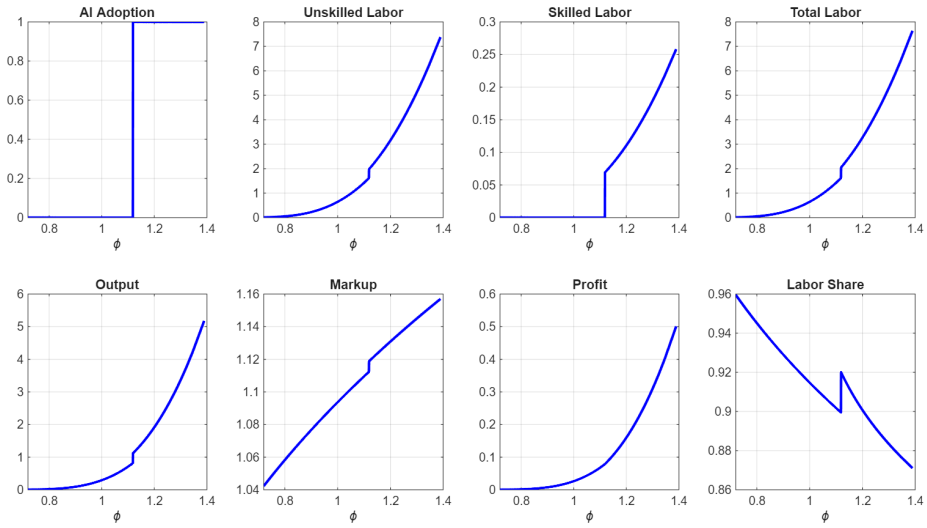
the increases in AI jobs relative to the decreases in the relative price of AI (0.165 vs. 0.147) and also the employment share of AI adopters (41.47% vs. 40%). Overall, the model does well in matching the targeted moments. In addition, the model matches closely the observed ratio of the increases in AI adopters to the declines in the relative price of AI, which is a non-targeted moment.

6 Model implications

We use our calibrated model to examine the implications of AI adoption for market power, productivity, and employment. We first discuss the decision rules of firms with different levels of productivity to illustrate the model mechanism. We then examine the impact of AI adoption—driven by either supply or demand factors—on market competition, focusing on industry concentration and average markup. Finally, we present some comparative statics analysis to illustrate the broader macroeconomic effects of AI adoptions.

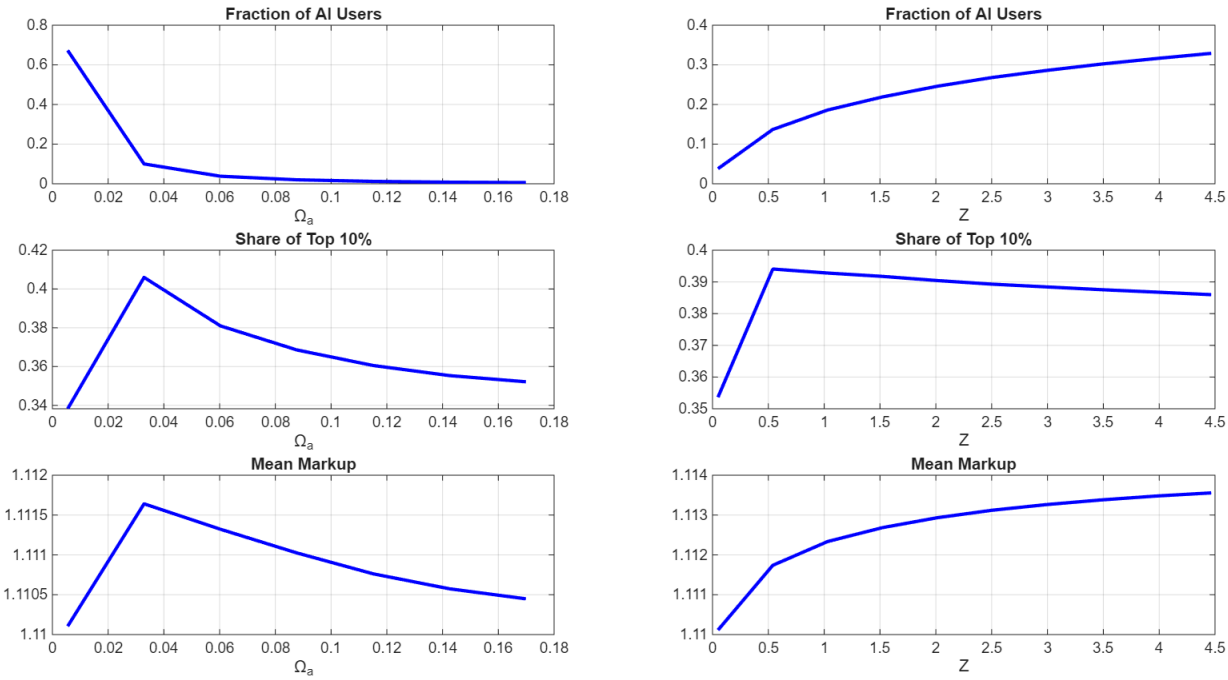
6.1 Firm decision rules

Figure 2. Firm Decision Rules



We begin with firms' decision rules. Figure 2 shows the decision rules for firms with different levels of productivity (ϕ). The figure shows that, for firms with sufficiently high productivity ($\phi > \phi^*$), it is profitable to pay the fixed cost and adopt AI. Regardless of whether a firm adopts AI, a firm with higher productivity produces more output and employs more workers. An AI adopter employs both workers with AI-related skills

Figure 3. AI adoption and market power



(skilled labor) and the other type of workers (unskilled labor). When a firm switches from a non-adopter to an adopter (i.e., with productivity ϕ^*), it experiences a jump in total employment and output, because it starts to hire skilled workers to operate the AI algorithm, and also because the use of AI enhances productivity, raising the firm's demand for both types of workers. More productive firms have higher markups, since they face less elastic demand under Kimball preferences, resulting in higher profits. As firms' productivity rises, their labor share falls.

6.2 AI adoption and market power

We now examine how increases in AI adoption affect market competition measured by industry concentration and average markup.

In our model, AI adoption can be driven by changes in AI demand, stemming from changes in the fixed cost of AI adoption (Ω_a). Adoption can also be driven by changes in AI supply, stemming from changes in the productivity in producing AI (Z). We examine the impact of variations in both types of drivers—one at a time—on the share of AI adopters, industry concentration, and the sales-weighted average markup. Figure 3 shows the simulation results.

6.2.1 Demand-driven AI adoption

First, consider the effects of demand-driven AI adoption. We vary the fixed cost of AI adoption (Ω_a)), while keeping all other parameters at their calibrated values. A decline in the fixed cost would not only increase the intensity of AI usage by incumbent users, but also attract new adopters, raising aggregate demand for AI and the share of AI adopters rises, as shown in the top left panel of the figure.

Demand-driven AI adoption has a hump-shaped relation with industry concentration, as shown in the middle left panel. At high levels of Ω_a (e.g., 0.04 or higher), a decline in Ω_a increases concentration measured by the sales share of the top 10% of firms. When Ω_a is sufficiently low (e.g., less than 0.03), a further decrease in Ω_a reduces concentration.[8]

The relationship between industry concentration and demand-driven AI adoption works through two channels. First, reducing the fixed cost benefits incumbent users of AI, which are large firms with high productivity (see Figure 2). Increased intensity of AI usage raises labor productivity for incumbent users, further increasing their market share. This reflects an intensive-margin effect. Second, with a lower fixed cost of adoption, some smaller firms would switch from non-AI users to AI users. Adopting AI helps those firms increase labor productivity and output, eroding the market share of the top firms (i.e., an extensive-margin effect). At high levels of the fixed cost, switching from a non-adopter to an adopter is costly, such that the intensive margin effect dominates and a decline in the fixed cost raises industry concentration. At a sufficiently low level of the fixed cost, further declines in the fixed cost would attract more new adopters, diluting the market share of the large incumbent AI users. These economic forces lead to a hump-shaped relation between the fixed cost of AI adoption and industry concentration. Under our calibration, the turning point for the relation between concentration and AI adoption occurs at an adoption rate of about 15 percent.

The impact of the fixed cost of AI adoption on average markup is also hump-shaped, as shown in the lower left panel of the figure.[9] When the initial level of Ω_a is high, lowering it would raise average markup; when Ω_a is sufficiently low, further reducing it would reduce markup. The hump-shaped relation between average markup and the fixed cost of AI adoption reflects the joint effects of three underlying forces. First, a decline in the fixed cost raises productivity of incumbent AI users, which are large firms with high markups. As those firms expand production, average markup rises. Second,

[8]The results are robust when we measure industry concentration using the sales share of the top 1% of firms. See the appendix for details.

[9]We compute the average markup as a sales-weighted average of firm-level markups in the intermediate goods sector.

a decline in the fixed cost attracts new AI adopters, which are smaller firms with lower markups. An increase in the share of the new adopters contributes to lowering the average markup. Third, a decline in the fixed cost raises the aggregate demand for AI, increasing the relative price of AI. Since firms are using AI as an input factor, the increase in the AI price raises the marginal cost of AI users, reducing their markup.

6.2.2 Supply-driven AI adoption

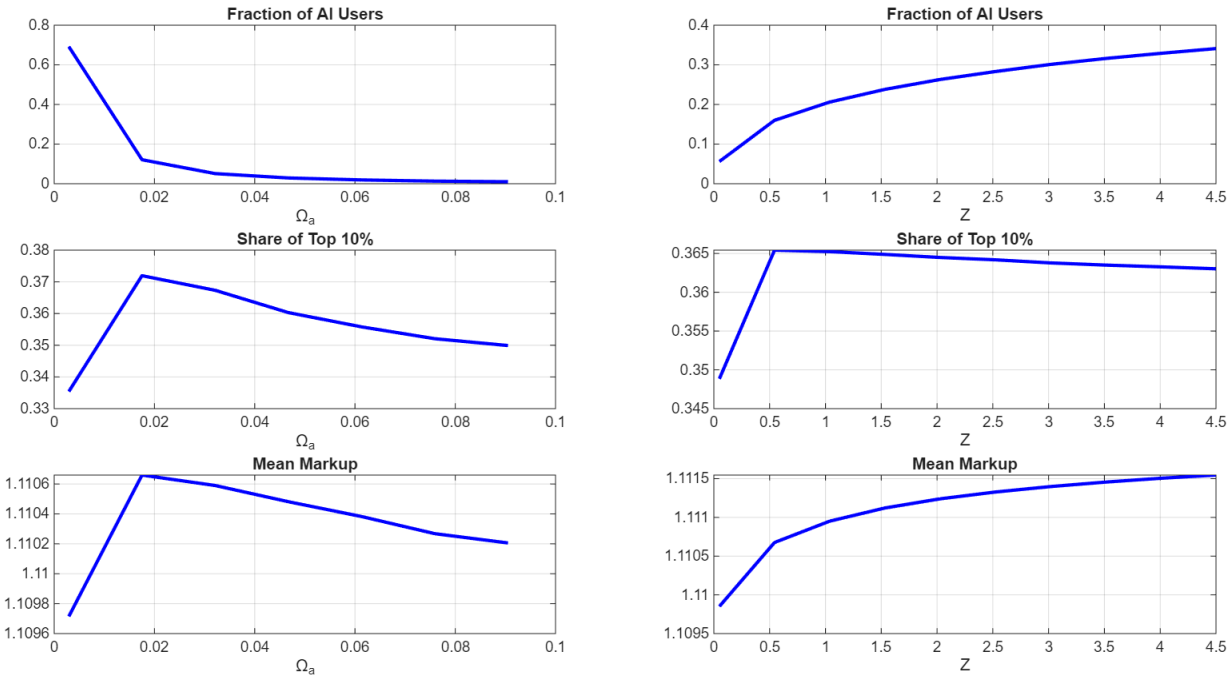
Next, consider the effects of supply-driven AI adoption. Here, we vary AI productivity (Z), which directly affects AI supply.

An increase in AI supply driven by an increase in Z lowers the relative price of AI, encouraging adoption, such that the share of AI adopters rises (upper right panel).

Supply-driven AI adoption also leads to hump-shaped changes in industry concentration (middle right panel), similar to the effects of demand-driven AI adoption. At an initially low level of Z, an increase in Z raises concentration; when Z is sufficiently high, a further increase in Z would reduce concentration. Similar to the case with demand-driven AI that we have discussed, the hump-shaped relation here reflects the offsetting effects of the extensive- vs. intensive-margin effects of AI adoption. Under our calibration, the turning point for the relation also occurs at an AI adoption rate of about 15 percent, similar to the case with demand-driven AI adoption. Unlike the demand-driven adoption, the relation turns relatively flat at high-levels of the adoption rate.

Different from the effects of demand-driven AI adoption, increases in AI supply driven by AI productivity would lead to an unambiguous increase in the average markup (lower right panel). As AI supply increases, the relative price of AI falls, lowering the marginal cost facing AI users and boosting their markup (i.e., a relative price effect). Similar to the case with demand-driven AI adoption, an increase in Z would increase the intensity of AI usage by incumbent users, who are large firms with high markups, contributing to increases in the average markup (i.e., an intensive margin effect). The increases in AI supply and the declines in the AI price would also induce some smaller firms with lower markups to switch from non-adopters to adopters, reducing the average markup (i.e., an extensive margin effect). Under our calibration, the extensive margin effect reinforces the relative price effect, which more than offset the intensive margin effect, resulting in monotonic increases in the average markup as AI productivity rises.

Figure 4. AI adoption and market power ($\beta = 0.3$)

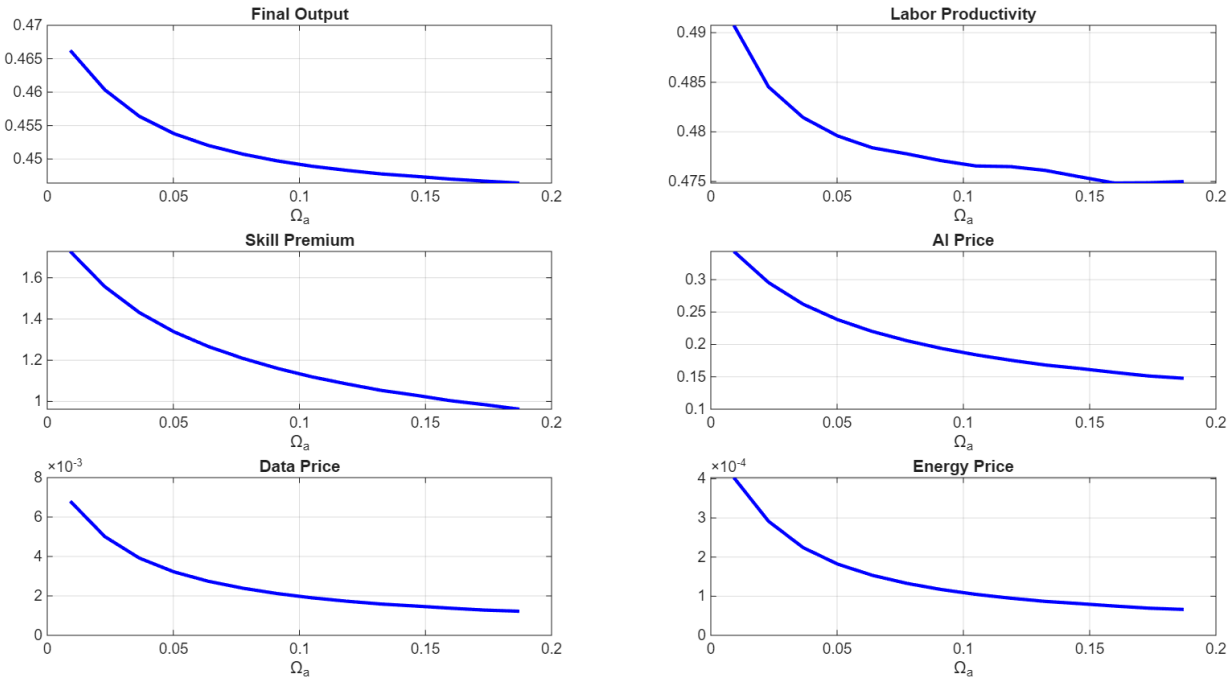


6.3 The role of data

In our baseline model, we assume that firms' internal data do not directly affect the effectiveness of their AI algorithms for production. We argue that this assumption partly reflects firms' concerns about data security as AI is still a nascent technology. If data security concerns are alleviated by new regulations or technological breakthroughs, firms would be more willing to use their own data to train the algorithms, potentially improving the labor productivity of AI users and increasing adoption. Since larger firms generate more data, allowing firms to enhance the efficiency of AI with their own data would potentially reinforce the market power of large firms, with potentially important implications for industry concentration and markup. We now examine the implications for market competition in a counterfactual economy where firms use their internal data to augment AI productivity. Specifically, we consider the case with $\beta = 0.3$ and compare the results with the baseline model with $\beta = 0$.

Figure 4 shows how changes in the AI adoption cost (left column) or AI productivity (right column) drive changes in the adoption rate, industry concentration, and the average markup in the counterfactual economy with $\beta = 0.3$. A decline in the adoption cost raises the share of AI users, resulting in hump-shaped changes in industry concentration (measured by sales share of the top 10% of firms) and the average markup. A rise in AI productivity also raises the adoption rate and leads to hump-shaped changes in concentration, although it raises the average markup monotonically. These patterns are

Figure 5. Macroeconomic effects of varying fixed costs of AI adoption (Ω_a)



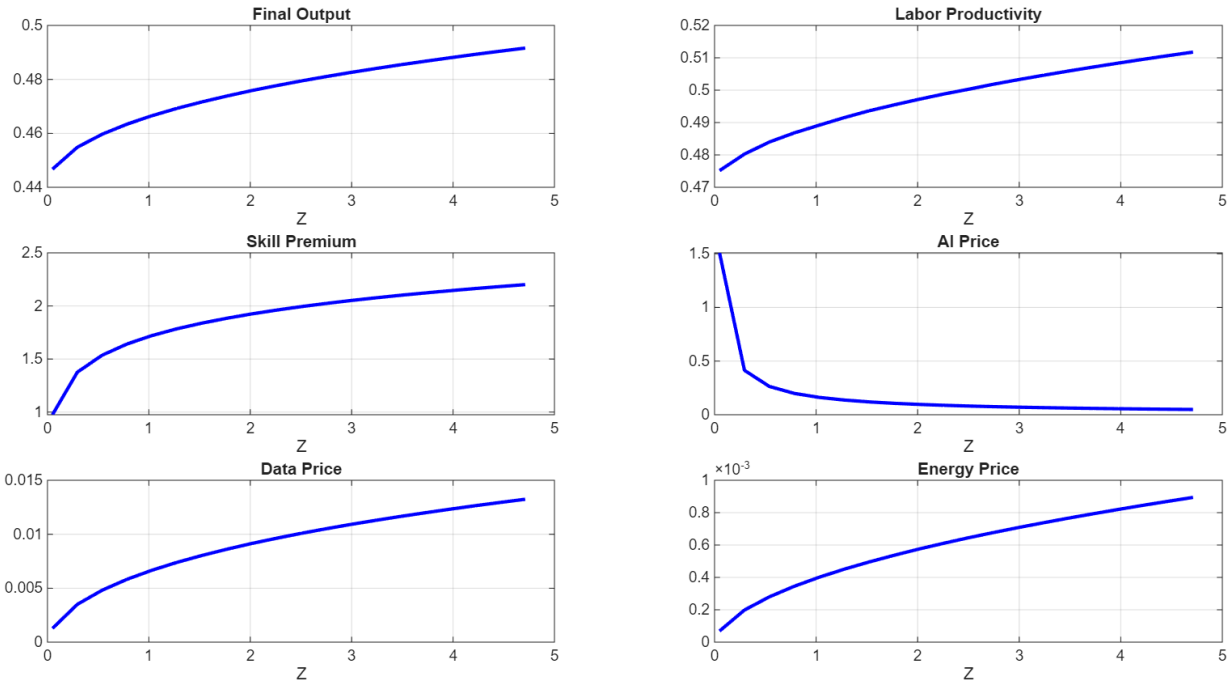
qualitatively the same as those obtained in the baseline model.

Overall, our model predicts a hump-shaped relation of AI adoption with industry concentration, regardless of whether the increases in adoption are driven by supply or demand factors. The relationship between AI adoption and the average markup, however, depends on the underlying reasons for the increased adoption. An economy experiencing supply-driven AI adoption would also experience increases in the average markup, because the decline in AI prices reduces the marginal cost for firms. In an economy with primarily demand-driven AI adoption, the average markup rises with AI adoption when the usage of AI is concentrated in a few large firms; ultimately, when AI usage is sufficiently widespread, the average markup would decline with further increased demand-driven adoption.

6.4 Broader macroeconomic implications of AI adoption

In addition to the relations between AI adoption and market competition, our model has broader macroeconomic implications of AI adoption. We now discuss these implications, again focusing on two drivers of AI adoptions: the fixed cost of adoption (i.e., a demand-side driver) and AI productivity (i.e., a supply-side driver).

Figure 6. Macroeconomic effects of varying AI productivity (Z)



6.4.1 Fixed costs of AI adoption

Figure 5 shows the macroeconomic effects of variations in the fixed cost of AI adoption, Ω_a. As discussed above, lowering Ω_a stimulates demand for AI usage, raising the share of AI adopters. The increased adoption of AI raises labor productivity and aggregate output, as shown in the figure. Since AI is complementary to skilled workers, the increased AI demand also raises the skill wage premium. Increased demand for AI drives up the relative price of AI, as well as demand for data and energy, raising the relative prices of those critical inputs as well.

6.4.2 AI productivity

While changes in the fixed cost of AI adoption drive changes in AI demand, changes in AI productivity (Z) drive changes in AI supply. Figure 6 shows the macroeconomic effects of varying AI productivity (Z). An increase in AI productivity induces more AI adoption (see Figure 3), raising aggregate labor productivity, aggregate output, and the skill premium, as shown in Figure 6. Unlike demand-driven AI adoption, this supply-driven adoption reduces the relative price of AI. The expansion of AI production raises demand for data and energy, pushing up their relative prices as well.

In Appendix D, we show that the macroeconomic effects of changes in the adoption cost or AI productivity are similar when we allow firms to use their own data to enhance

their use of AI (see Figures D.1 and D.2).

7 Supporting empirical evidence

Our model predicts a hump-shaped relation of AI adoption with industry concentration. At early stages of the adoption, AI usage is concentrated in a few large firms and increased adoption would improve labor productivity of incumbent AI users and thereby expanding their market share further. Ultimately, when the new technology becomes relatively widespread, further increases in AI adoption would attract smaller firms to switch from non-adopters to adopters, diluting the market share of the top firms. Furthermore, our model predicts that such hump-shaped relations hold regardless of whether increased AI adoption is driven by increased supply of or demand for AI. We now present some empirical evidence that supports these model predictions.

Existing empirical evidence suggests that, in the early years of AI adoption, large firms are more likely to adopt AI than small firms and that AI-powered growth concentrates among large firms. For example, Babina et al. (2024) find that the long-differences in industry concentration during the period of 2010-2018 are positively correlated with the long-differences in the share of AI workers among Compustat firms. We replicate this finding in Table 2 below (Column (1)), where we measure the industry-level intensity of AI usage by the share of AI job postings. The table shows that, in the period of 2010-2018, the long-differences in AI intensity for NAICS 5-digit industries are positively correlated with the long-differences in the top firm's sales share in the same industries after controlling for industry fixed effects and the initial conditions in NAICS 2-digit industries. The estimated correlation is statistically significant at the 5 percent level, confirming the findings of Babina et al. (2024).

However, when we re-run the same regressions using the more recent sample from 2019 to 2023, the correlation between industry concentration and AI intensity becomes negative and statistically insignificant (see Column (2)). We obtain similar results when we measure industry concentration using the Herfindahl–Hirschman index (HHI), as shown in Columns (3) and (4). The weaker or even negative correlations between industry concentration and AI intensity in the more recent sample period (after 2019) are perhaps surprising, because AI adoption has accelerated in recent years. If AI-powered growth benefits large firms disproportionately, as Babina et al. (2024) argue, then we should continue to observe a positive and even stronger correlation between industry concentration and AI intensity. This evidence, however, is in line with our model's predicted hump-shaped relation between AI adoption and industry concentration.

Table 2. Industry concentration and AI intensity

	Δ Top Firm Share		Δ ln(HHI)	
	(1) 2010-2018	(2) 2019-2025	(3) 2010-2018	(4) 2019-2025
AI Job Share (std)	0.0564***	-0.0177	0.2040***	-0.0248
	(0.0160)	(0.0177)	(0.0673)	(0.0420)
Industry FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes
Observations	170	160	170	160
R^2	0.2495	0.1549	0.3280	0.1448

Notes: This table shows the long-difference regressions of changes in industry concentration on contemporaneous changes in the industry-level AI intensity. Industry concentration is measured by the sales share of the top firm in each NAICS 5-digit industry in Compustat and also the Herfindahl–Hirschman index (HHI). AI intensity is measured as the share of AI-related job postings for a firm using online job postings data from Lightcast. The long differences of industry-level AI intensity are calculated by averaging the long differences of firm-level AI intensity, weighted by firm sales in the starting year of the samples (2010 for Columns (1) and (2) and 2019 for Columns (3) and (4)). The regressions control for log sales, log employment, and log wages for each NAICS 2-digit sector in 2010, and are weighted by the number of industry job postings in the starting year . Robust standard errors clustered at NAICS-2 sector level. *** p<0.01, ** p<0.05, and * p<0.10.

8 AI policies

The model features monopolistic competition with variable markups in the goods sector, monopolistic competition in the AI sector, as well as an externality stemming from the non-rivalry of data usage. Thus, decentralized equilibrium is inefficient and policy interventions might improve welfare. We examine the macroeconomic and welfare consequences of three policies: (1) taxing or subsidizing AI adoption; (2) taxing or subsidizing AI production; and (3) regulating data usage.

8.1 AI adoption policy

Our model implies that a hump-shaped relationship between AI adoption and industry concentration. At low levels of AI adoption, an increase in adoption induced by a decline in the fixed cost of adoption or an improvement in AI productivity would enable incumbent AI using firms to gain market share. Since incumbent users are large firms who face low demand elasticities, the increased AI adoption would lead to greater industry concentration and higher average markup. It would also raise aggregate productivity since larger firms are more productive. In an economy with widespread AI adoption, a further increase in adoption would induce smaller firms to adopt AI, eroding the market share of incumbent large firms.

In general, it is not clear a priori whether encouraging AI adoption would improve or reduce social welfare, depending on the parameter calibration. We now use our calibrated model to examine a policy that imposes a tax or a subsidy on AI adoption.

Denote by τ_a the net tax rate. If τ is negative, then it would a subsidy. If an intermediate goods producer adopts AI, then its revenue would be $(1-\tau_a)p(\phi)y(\phi)$. Non-adopters are not subject to this tax policy. The tax revenue (or subsidy expenditure) is transferred to the representative household in a lump-sum fashion, so it does not affect the aggregate resource constraint.

With this policy, the profit-maximizing problem (12) for AI adopters is replaced by

$$\pi_a(\phi) = \max_{p,y,d,u,s,A} \left[(1-\tau_a)p(\phi)y(\phi) + P_d d(\phi) - w_u u(\phi) - w_s s(\phi) - P_a A(\phi) \right], \tag{43}$$

subject to the same demand schedule (2) and the production function (9). Non-AI users are not subject to the tax policy.

8.2 AI production policy

We next consider a policy that restricts (or encourages) AI production. Denote by τ_p the tax rate on AI producers. Then the optimal decisions (25) for AI producers would be replaced by

$$(1-\tau_p)p_a(\nu) = \frac{\kappa}{\kappa-1}\lambda_a. \tag{44}$$

In an symmetric equilibrium, $p_a(\nu) = p_a$, so this condition becomes

$$(1-\tau_p)p_a = \frac{\kappa}{\kappa-1}\lambda_a. \tag{45}$$

The AI production tax does not affect the free entry condition (29) because the relative price of AI has internalized the effects of the taxes. The tax revenue (or subsidy spending) is rebated to (or collected from) the household in a lump-sum fashion.

8.3 Data regulations

Data is a by-product of firm production and firms can sell data to AI producers. This non-rivalry nature of data can distort equilibrium allocation. Now consider a policy that regulate data usage (e.g., privacy policy) such that sales of data are subject to taxes at the rate τ_d. Suppose that the taxes are imposed on firms. Then, the profit-maximizing problem of AI adopters becomes

$$\pi_a(\phi) = \max_{p,y,d,u,s,A}\left[p(\phi)y(\phi) + (1-\tau_d)P_d d(\phi) - w_u u(\phi) - w_s s(\phi) - P_a A(\phi)\right], \tag{46}$$

subject to the demand schedule (2) and the production function (9).

For non-AI adopters, the profit maximizing problem becomes

$$\pi_n(\phi) = \max_{p,y,d,u}\left[p(\phi)y(\phi) + (1-\tau_d)P_d d(\phi) - w_u u(\phi)\right], \tag{47}$$

subject to the demand schedule (2) and the production function (10).

9 Conclusion

AI has been rising rapidly. However, since AI is a relatively nascent technology, it remains unclear how it would affect market competition. In this paper, we study the implications of increasing AI adoption on industry concentration in a general equilibrium framework with heterogeneous firms. The model predicts that large firms with high productivity

and high markups are more likely to adopt AI because the fixed costs of adoption lead to a scale-economy effect. We find that the relationship between industry concentration and AI adoption is not monotone. As AI adoption increases from an initially low level, large firms that are incumbent users of AI gain market shares. When AI is sufficiently diffused, however, entry of smaller adopters erodes the market share of incumbents, reducing industry concentration. This non-monotone relation does not depend on whether AI adoption is demand- or supply-driven or whether firms can complement AI with their own data. Industry-level evidence suggests that the correlations between industry concentration and AI usage have weakened in recent years as AI rises rapidly, broadly in line with our model's predictions.

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2022. "Artificial intelligence and jobs: evidence from online vacancies." *Journal of Labor Economics*, 40(S1): S293–S340.

Adams, Jonathan J., Min Fang, Zheng Liu, and Yajie Wang. 2026. "The rise of AI pricing: Trends, driving forces, and implications for firm performance." *Journal of Monetary Economics*, 157: 103875.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. 2019. "A Theory of Falling Growth and Rising Rents." National Bureau of Economic Research Working Paper 26448.

Akcigit, Ufuk, and Sina T Ates. 2019. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." National Bureau of Economic Research Working Paper 25755.

Aruoba, S. Borağan, Eugene Oue, Felipe Saffie, and Jonathan Willis. 2024. "Non-Constant Demand Elasticities, Firm Dynamics and Monetary Non-Neutrality: Role of Demand Shocks." National Bureau of Economic Research Working Paper No. 32518.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics*, 135(2): 645–709.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. "Artificial intelligence, firm growth, and product innovation." *Journal of Financial Economics*, 151: 103745.

Baily, Martin, Erik Brynjolfsson, and Anton Korinek. 2023. "Machines of mind: The case for an AI-powered productivity boom." Brookings Institution.

Bick, Alexander, Adam Blandin, David J. Deming, Nicola Fuchs-Schundeln, and Jonas Jessen. 2026. "Mind the Gap: AI Adoption in Europe and the US." Unpublished manuscript, Federal Reserve Bank of St. Louis.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J Terry. 2018. "Really Uncertain Business Cycles." *Econometrica*, 86(3): 1031–1065.

Choi, Jaedo, Andrei A Levchenko, Dimitrije Ruzic, and Younghun Shim. 2024. "Superstars or Supervillains? Large Firms in the South Korean Growth Miracle."

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2023. "How Costly Are Markups?" *Journal of Political Economy*, 131(7): 1619–1675.

Ekerdt, Lorenz K.F., and Kai-Jie Wu. 2022. "The Rise of Specialized Firms." *Working Paper*.

Firooz, Hamid, Zheng Liu, and Yajie Wang. 2025. "Automation and the rise of superstar firms." *Journal of Monetary Economics*, 103733.

Furman, Jason, and Peter Orszag. 2018. "Slower Productivity and Higher Inequality: Are They Related?" Peterson Institute for International Economics Working Paper Series WP18-4.

Hubmer, Joachim, and Pascual Restrepo. 2022. "Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital–Labor Substitution." National Bureau of Economic Research Working Paper 28579.

Jones, Charles I, and Christopher Tonetti. 2020. "Nonrivalry and the Economics of Data." *American Economic Review*, 110(9): 2819–2858.

Kimball, Miles S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit and Banking*, 27(4): 1241–1277.

Klenow, Peter J., and Jonathan L. Willis. 2016. "Real Rigidities and Nominal Price Changes." *Economica*, 83(331): 443–472.

Kwon, Spencer Y., Yueran Ma, and Kaspar Zimmermann. 2024. "100 Years of Rising Corporate Concentration." *American Economic Review*, 114(7): 2111–40.

Lashkari, Danial, Arthur Bauer, and Jocelyn Boussard. 2024. "Information Technology and Returns to Scale." *American Economic Review*, 114(6): 1769–1815.

Mihet, Roxana, Kumar Rishabh, and Orlando Gomes. 2026. "Is it AI or data that drives firm market power?" *Journal of Monetary Economics*, 157: 103878.

Olmstead-Rumsey, Jane. 2019. "Market Concentration and the Productivity Slowdown." University Library of Munich, Germany MPRA Paper 93260.

Ridder, Maarten De. 2023. "Market Power and Innovation in the Intangible Economy." *American Economic Review*, , (forthcoming).

Shen, Shihan. 2023. "Customer Acquisition, Rising Concentration and US Productivity Dynamics." *Rising Concentration and US Productivity Dynamics (September 15, 2023)*.

Sui, Xiaomei. 2022. "Uneven Firm Growth in a Globalized World." *Working Paper*.

Tambe, Prasanna, Lorin Hitt, Daniel Rock, and Erik Brynjolfsson. 2020. "Digital Capital and Superstar Firms." National Bureau of Economic Research Working Paper 28285.

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven J Davis, Kevin M Foster, Aaron Jalca, Brent H Meyer, Paul Mizen, Michael A Navarrete, Pawel Smietanka, Gregory Thwaites, and Ben Zhe Wang. 2026. "Firm Data on AI." National Bureau of Economic Research Working Paper 34836.

Appendices

A Keywords for AI and Gen AI related skills

Table A.1. Traditional AI/ML Skills

Machine Learning	Latent Semantic Analysis
Artificial Intelligence	Sentiment Analysis/Opinion Mining
Computer Vision	Latent Dirichlet Allocation
Machine Vision	Predictive Modeling
Deep Learning	Kernel Methods
Virtual Agents	Keras
Image Recognition	Gradient Boosting
Natural Language Processing	OpenCV
Speech Recognition	XGBoost
Pattern Recognition	Libsvm
Object Recognition	Word2vec
Neural Networks	Machine Translation
AI Chatbot	Sentiment Classification
Supervised Learning	Mahout
Text Mining	Recommender Systems
Unsupervised Learning	Support Vector Machines
Image Processing	Random Forests

Table A.2. Generative AI Skills

Generative AI	GenAI
Large Language Model	LLM
Foundation Model	Foundation Models
Prompt Engineering	Prompt Design
AI Agent	AI Agents
AI Assistant	Retrieval-Augmented Generation
RAG	Multimodal Model
Multimodal Models	Transformer Model
Transformer Models	Transformers
Autoregressive Model	GPT
GPT-4	GPT-4o
GPT-3.5	Claude
LLama	LLama2
LLama3	Gemini
Mistral	Mixtral
DeepSeek	Palm
Phi	Fine-tuning
Finetuning	Instruction Tuning
Supervised Fine-tuning	SFT
NLP	NLU
NLG	Attention Mechanism
Sequence Modeling	Reinforcement Learning
Machine Learning Engineering	Data Augmentation
RLHF	PPO
DPO	LoRA
QLoRA	Parameter-Efficient Fine-tuning
Model Distillation	Model Quantization
Quantization	Model Optimization
Inference Optimization	GPU Acceleration
Model Serving	Triton Kernels
ONNX Runtime	TensorRT
vLLM	DeepSpeed
Synthetic Data	Synthetic Data Generation
Data Augmentation	Text Embeddings
Embedding Model	Embedding Models
Vector Database	Vector Databases
Semantic Search	LangChain
LlamaIndex	Chatbot
Chatbot Development	Conversational AI
AI Copilots	Copilot
Workflow Automation	Workflow Orchestration
Agentic System	Agentic Systems
AI-powered Application	AI Integration
Text Generation	Code Generation
Image Generation	Vision-Language Model
Vision-Language Models	VLM
AI/ML Product	AI Product
AI Strategy	Responsible AI
AI Governance	Model Evaluation
AI Adoption	AI Transformation
GenAI Solutions	Enterprise GenAI
AI Roadmap	Diffusion Model
Diffusion Models	

B Sketch of the solution algorithm

We assume the productivity distribution $G(\cdot)$ is log-normal with $\ln(\phi) \sim \mathcal{N}(0, \sigma_\phi^2)$ and dicretize the productivity distribution with 1000 grid points. We guess the demand shifter H, aggregate output Y, wages w_u and w_s, energy price P_e, data price P_d, and mass of AI producers M.

Given the price of energy and data, we find the price of AI varieties using equation (25). The aggregate AI price is then given by equation (26). Given wages and aggregate AI price, we compute intermediate producers' marginal cost for firms using

AI and those not using AI, given in **??** and equation (14), respectively. Given H and D, we use equations (16) and (17) to find relative quantities, given that markups in equation (8) are also in function of relative quantities. We can therefore solve for relative quantities, once assuming the firm is using AI and once assuming the firm is not using AI. Given Y, we find firms' outputs. We then determine which firms use AI by comparing profits in equation (15). We can then compute $d(\phi)$ using equation (11), and aggregate it up to compute total data supply D, according to equation (39). We also use the definition of H in equation (3) to see if it equals our guess.

We now know each intermediate producers decision variables, so we can compute aggregate objects u, s, and A using (36), (37), and (38). We then compute D using equation (27) and (40).

Mass of AI producers solves the free entry condition (29).

Energy price clears the energy market in equation (42).

Wages satisfy labor supply equations (33) and (34), where C satisfies the household budget constraint (31).

Y ensures that the definition of Kimball in equation (1) is satisfied.

satisfies the household budget constraint (31).

P_d clears the data market in equation (40).

C Additional model results

Figure C.1. AI adoption and market power

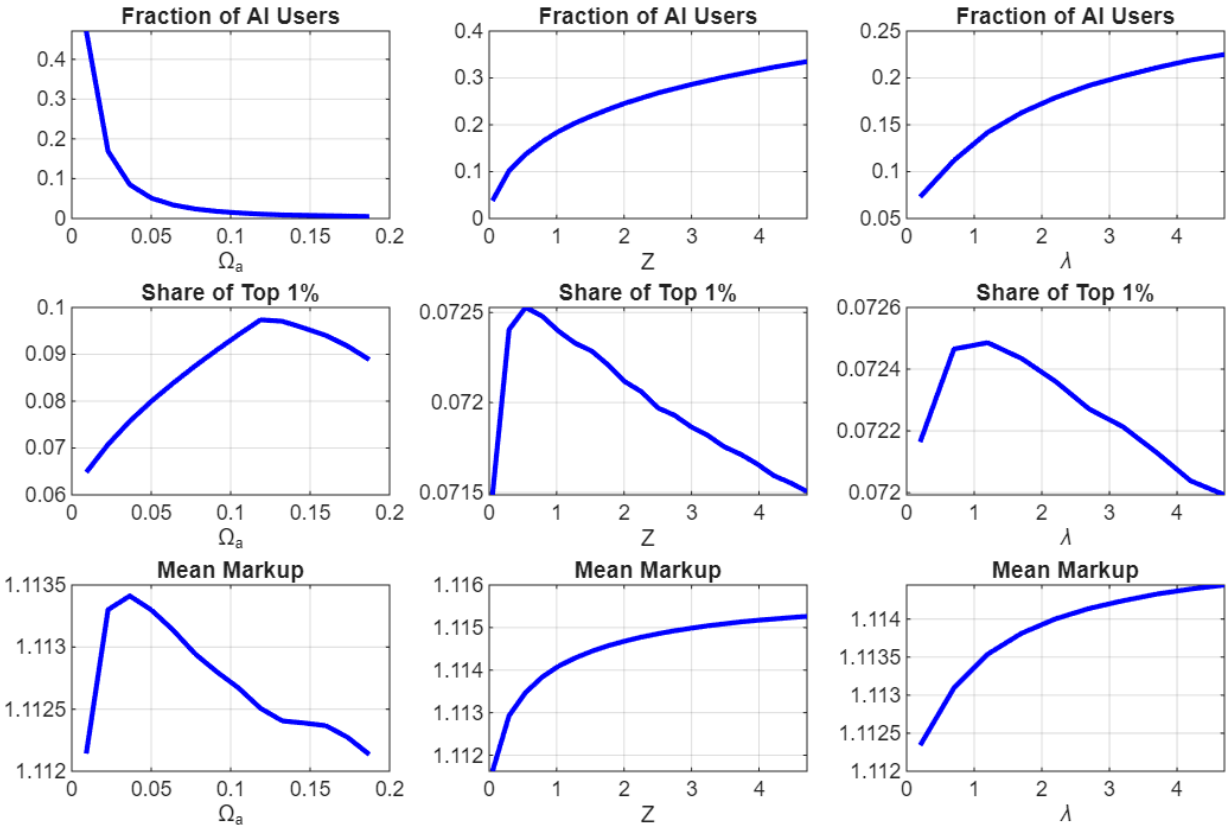
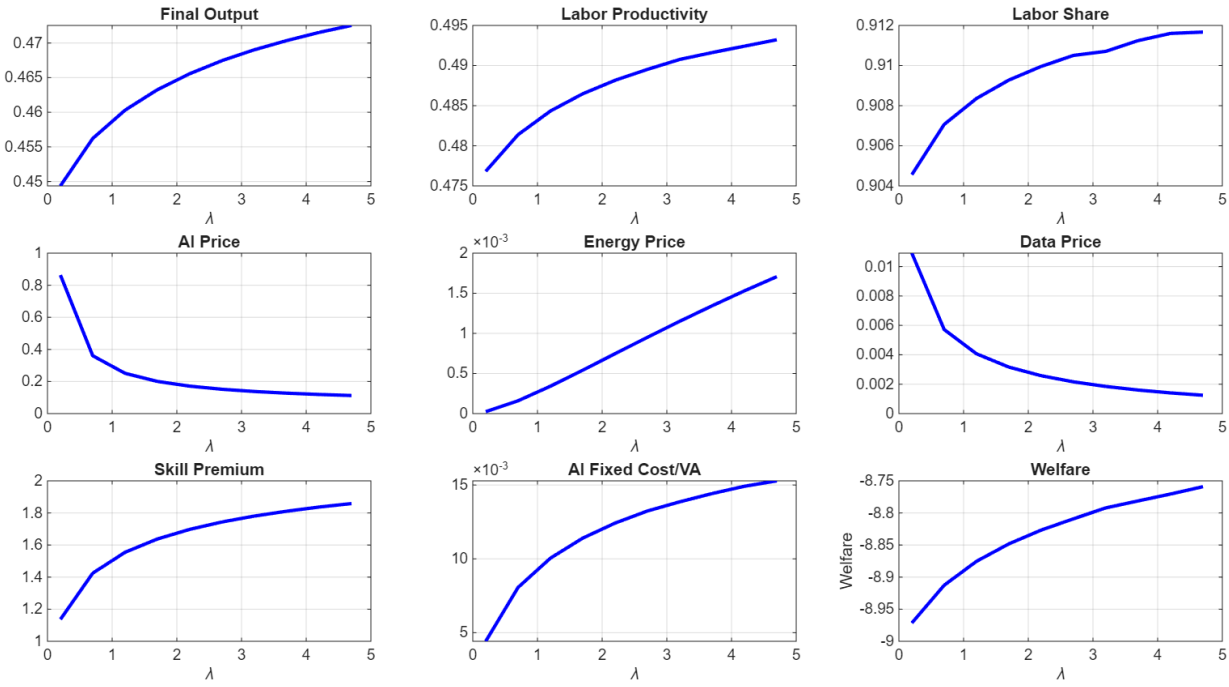


Figure C.2. Counterfactual Analysis for λ



D Augmenting AI with Data

Suppose the AI-adopter production function is given by

$$y(\phi) = \phi \left[\alpha_a \left((A(\phi)^{1-\beta} d(\phi)^{\beta})^{\gamma} s(\phi)^{1-\gamma} \right)^{\frac{\eta-1}{\eta}} + (1-\alpha_a) u(\phi)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}, \tag{48}$$

with $0 \leq \beta \leq 1$, where the firm internalizes that own data satisfies

$$d(\phi) = \lambda y(\phi). \tag{49}$$

This formulation nests our baseline model when we impose $\beta = 0$.

We derive the firm's marginal cost and conditional factor demands.

Nested representation

Define the inner composite

$$x(\phi) \equiv \left(A(\phi)^{1-\beta} d(\phi)^{\beta} \right)^{\gamma} s(\phi)^{1-\gamma}. \tag{50}$$

Then the production function can be written as

$$y(\phi) = \phi \left[\alpha_a x(\phi)^{\frac{\eta-1}{\eta}} + (1-\alpha_a) u(\phi)^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}. \tag{51}$$

Using (49),

$$x(\phi) = \left(A(\phi)^{1-\beta} (\lambda y(\phi))^{\beta} \right)^{\gamma} s(\phi)^{1-\gamma} = (\lambda y(\phi))^{\gamma\beta} A(\phi)^{\gamma(1-\beta)} s(\phi)^{1-\gamma}. \tag{52}$$

For notational convenience, define

$$a \equiv \gamma(1-\beta), \qquad b \equiv 1-\gamma, \qquad \delta \equiv a + b = 1 - \gamma\beta. \tag{53}$$

Then (52) becomes

$$x(\phi) = (\lambda y(\phi))^{\gamma\beta} A(\phi)^{a} s(\phi)^{b}. \tag{54}$$

Thus the problem is naturally nested:

1. minimize the cost of producing the inner bundle x from A and s;
2. minimize the cost of producing y from x and u.

Cost of the inner AI-skill bundle

For a given target output y, the firm takes $d = \lambda y$ as given. Conditional on y, the inner problem is

$$\min_{A,s} \; P_a A + w_s s \quad \text{subject to} \quad x = (\lambda y)^{\gamma\beta} A^a s^b. \tag{55}$$

Equivalently,

$$\frac{x}{(\lambda y)^{\gamma\beta}} = A^a s^b. \tag{56}$$

Since $a + b = \delta < 1$, this is a Cobb–Douglas technology with decreasing returns to scale in (A, s). The associated cost function is

$$C_x(x, y) = \delta \left[\left(\frac{P_a}{a} \right)^a \left(\frac{w_s}{b} \right)^b \right]^{\frac{1}{\delta}} \left(\frac{x}{(\lambda y)^{\gamma\beta}} \right)^{\frac{1}{\delta}}. \tag{57}$$

Define

$$\tilde{c} \equiv \left[\left(\frac{P_a}{a} \right)^a \left(\frac{w_s}{b} \right)^b \right]^{\frac{1}{\delta}}. \tag{58}$$

Then (57) can be written as

$$C_x(x, y) = \delta \, \tilde{c} \, (\lambda y)^{-\gamma\beta/\delta} x^{1/\delta}. \tag{59}$$

The marginal cost of one extra unit of the inner composite is

$$p_x(\phi) \equiv \frac{\partial C_x(x, y)}{\partial x} = \tilde{c} \, (\lambda y(\phi))^{-\gamma\beta/\delta} x(\phi)^{\gamma\beta/\delta}. \tag{60}$$

By Shepard's lemma, the conditional demands for A and s satisfy

$$P_a A(\phi) = a \, p_x(\phi) x(\phi), \qquad w_s s(\phi) = b \, p_x(\phi) x(\phi), \tag{61}$$

or equivalently,

$$A(\phi) = \frac{a \, p_x(\phi) x(\phi)}{P_a}, \qquad s(\phi) = \frac{b \, p_x(\phi) x(\phi)}{w_s}. \tag{62}$$

Outer CES cost minimization

Given the shadow price $p_x(\phi)$ of the inner composite, the outer problem is

$$\min_{x,u} \; C_x(x, y) + w_u u \quad \text{subject to} \quad y = \phi \left[\alpha_a x^{\frac{\eta-1}{\eta}} + (1 - \alpha_a) u^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}}. \tag{63}$$

Let

$$\Psi(\phi) \equiv \frac{1}{\phi}\left[\alpha_a^{\eta} p_x(\phi)^{1-\eta} + (1-\alpha_a)^{\eta} w_u^{1-\eta}\right]^{\frac{1}{1-\eta}} \tag{64}$$

denote the shadow expenditure index associated with the outer CES aggregator. Then the conditional demands are

$$x(\phi) = \frac{y(\phi)}{\phi}\left(\frac{\alpha_a \phi \Psi(\phi)}{p_x(\phi)}\right)^{\eta}, \tag{65}$$

and

$$u(\phi) = \frac{y(\phi)}{\phi}\left(\frac{(1-\alpha_a)\phi\Psi(\phi)}{w_u}\right)^{\eta}. \tag{66}$$

It is useful to define the shadow expenditure share on the inner composite:

$$\theta_x(\phi) \equiv \frac{\alpha_a^{\eta} p_x(\phi)^{1-\eta}}{\alpha_a^{\eta} p_x(\phi)^{1-\eta} + (1-\alpha_a)^{\eta} w_u^{1-\eta}}. \tag{67}$$

Using the standard CES demand system,

$$p_x(\phi)x(\phi) = \theta_x(\phi)\Psi(\phi)y(\phi), \qquad w_u u(\phi) = \left(1-\theta_x(\phi)\right)\Psi(\phi)y(\phi). \tag{68}$$

Because the inner cost function is nonlinear, the shadow price $p_x(\phi)$ is endogenous. Combining (60) and (65) yields the consistency condition

$$p_x(\phi)^{1+\eta\gamma\beta/\delta} = \tilde{c}\,\lambda^{-\gamma\beta/\delta}\alpha_a^{\eta\gamma\beta/\delta}\phi^{(\eta-1)\gamma\beta/\delta}\Psi(\phi)^{\eta\gamma\beta/\delta}. \tag{69}$$

Equations (64) and (69) jointly determine $p_x(\phi)$ and $\Psi(\phi)$.

Marginal cost

The firm's total variable cost is

$$VC_a(\phi, y) = C_x(x(\phi), y(\phi)) + w_u u(\phi). \tag{70}$$

Using (59) and (60),

$$C_x(x(\phi), y(\phi)) = \delta\, p_x(\phi)x(\phi). \tag{71}$$

Therefore, using (68),

$$VC_a(\phi, y) = \delta\, p_x(\phi)x(\phi) + w_u u(\phi) \tag{72}$$
$$= \delta\, \theta_x(\phi)\Psi(\phi)y(\phi) + \big(1 - \theta_x(\phi)\big)\, \Psi(\phi)y(\phi) \tag{73}$$
$$= \big[1 - \gamma\beta\, \theta_x(\phi)\big]\, \Psi(\phi)y(\phi), \tag{74}$$

where the last equality uses $\delta = 1 - \gamma\beta$.

Hence the marginal cost of the AI adopter is

$$MC_a(\phi) = \frac{\partial VC_a(\phi, y)}{\partial y} = \big[1 - \gamma\beta\, \theta_x(\phi)\big]\, \Psi(\phi). \tag{75}$$

Unlike the case with the production function $x = (Ad)^{\gamma} s^{1-\gamma}$, the variable cost here is linear in output. Therefore the AI-adopting firm has constant marginal cost.

Conditional factor demand

Using (68), the conditional demand for unskilled labor is

$$u(\phi) = \frac{\big(1 - \theta_x(\phi)\big)\, \Psi(\phi)y(\phi)}{w_u}. \tag{76}$$

Similarly, the conditional demand for the inner composite is

$$x(\phi) = \frac{\theta_x(\phi)\Psi(\phi)y(\phi)}{p_x(\phi)}. \tag{77}$$

Using (62), the conditional demands for AI and skilled labor are

$$A(\phi) = \frac{a\, p_x(\phi)x(\phi)}{P_a} = \frac{a\, \theta_x(\phi)\Psi(\phi)y(\phi)}{P_a}, \tag{78}$$

and

$$s(\phi) = \frac{b\, p_x(\phi)x(\phi)}{w_s} = \frac{b\, \theta_x(\phi)\Psi(\phi)y(\phi)}{w_s}. \tag{79}$$

Substituting back for $a = \gamma(1 - \beta)$ and $b = 1 - \gamma$,

$$A(\phi) = \frac{\gamma(1 - \beta)\, \theta_x(\phi)\Psi(\phi)y(\phi)}{P_a}, \qquad s(\phi) = \frac{(1 - \gamma)\, \theta_x(\phi)\Psi(\phi)y(\phi)}{w_s}. \tag{80}$$

Finally, data usage is

$$d(\phi) = \lambda y(\phi). \tag{81}$$

D.1 Counterfactuals

Figure D.1. Macroeconomic effects of varying fixed costs of AI adoption ($\beta = 0.3$)

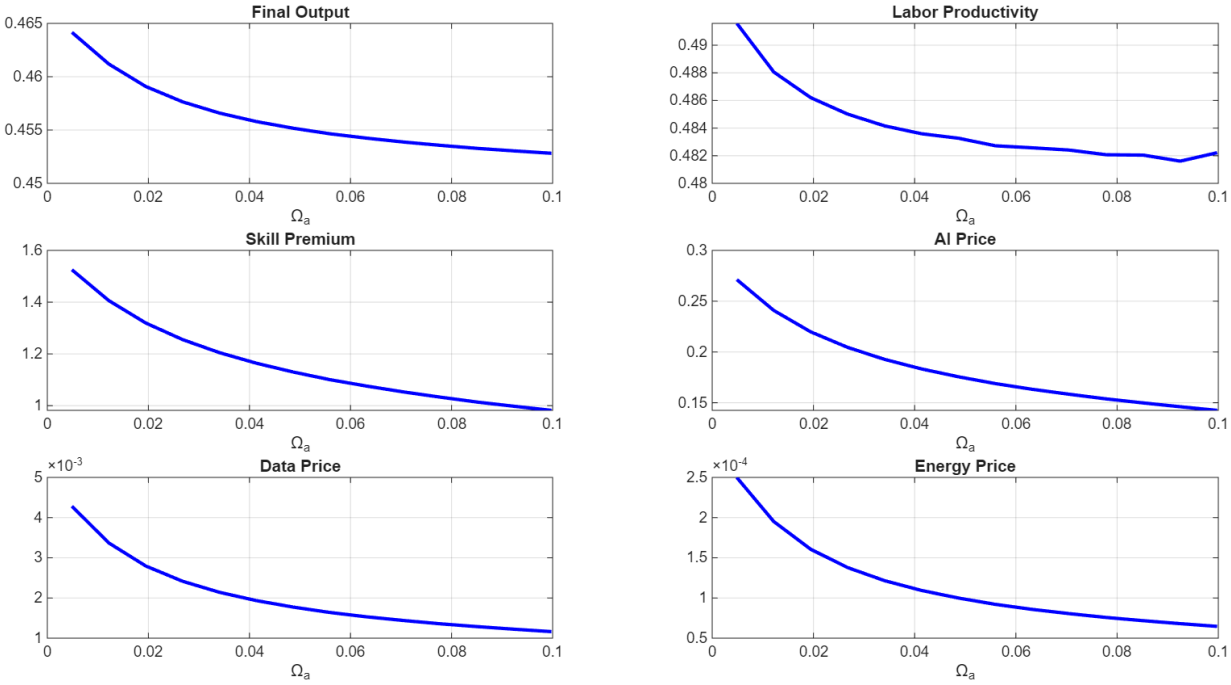


Figure D.2. Macroeconomic effects of varying AI productivity ($\beta = 0.3$)

