

Stellar Skills: Superstar Firms and the Specialization of Human Capital

Eugenia Menaguale*

Updated regularly; please [click here for the latest version](#)

This version: November 6, 2025

Abstract

Technological change is the key driver of sustained economic growth, but new ideas can only transform production if the supply of technology-specific inputs supports their adoption. This paper studies how imperfect competition in markets for technology-specific capital shapes the direction of firms' technology adoption and workers' specific human capital accumulation. I argue that when production requires combining technology-specific capital and labor, providers of specialized capital have a direct stake in the formation of complementary skills. I develop a general equilibrium model in which dominant capital providers set rental rates and subsidize training in the skills required to operate their systems, internalizing their effect on the market for technology-specific labor. The framework features a policy-relevant trade-off: market power in markets for specialized capital raises misallocation, but it can stimulate investment in high-quality training. I quantify how the balance between these forces shapes technological specialization and welfare, showing that stronger competition in input markets is not always welfare-improving, particularly when workers' preferences make skills less substitutable. I test the theory in the digital infrastructure industry. Using large-scale job posting microdata and developer surveys, I study the *open source* release of platform-specific software by a major supplier of cloud computing power. I observe patterns consistent with my theory of subsidies to skill acquisition costs: the release of specific software is associated with increased adoption of the complementary platform, but also a lowered wage premium for the related specialized skills.

JEL Codes: E23, E24, O33, O34, J24, L16, L17

*Princeton University, Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ 08544.

I thank my advisors Ezra Oberfield, Steve Redding and Gianluca Violante for invaluable guidance and support. I also thank David Argente, Seth Blumberg, Leonardo D'Amico, Mayara Felix, Jeff Gortmaker, Allison Green, Gene Grossman, Michael Impink, Danial Lashkari, Simon Margolin, Jessica Min, Yuhei Miyauchi, Giovanni Montanari, Eduardo Morales, Jeremy Pearce, Richard Rogerson, Karthik Sastry, Martin Vaeth, Hal Varian for insightful conversations and seminar participants at Princeton University and the Federal Reserve Bank of New York for their helpful suggestions. I gratefully acknowledge financial support from the International Economics Section and the Louis A. Simpsons Center for the Study of Macroeconomics at Princeton University. All errors are my own. Email: eugenia.menaguale@princeton.edu.

1 Introduction

Modern production processes often rely on technology-specific know-how that ties workers' skills to particular systems and tools. In rapidly innovating sectors such as automotive, pharmaceuticals, and semiconductors, technicians operate proprietary diagnostic instruments, and software engineers master the programming environments of distinct computing platforms. Because the availability of specialized skills is crucial for firms to adopt new technologies, understanding how workers acquire specific human capital is fundamental to explaining technological change, the key driver of sustained economic growth.¹ Work in economics has emphasized how workers acquire skills in anticipation of labor market opportunities (Becker, 1962; Goldin and Katz, 2008) and at the same time that the diffusion of technology hinges on the availability of skilled human capital (Nelson and Phelps, 1966; Acemoglu, 1998; Caselli and Coleman, 2006). Yet, most of this literature classifies skills using coarse measures, such as education levels or broad occupational categories, and overlooks the growing importance of highly specialized, technology-specific expertise developed through training and certification programs sponsored by major capital providers.²

In this paper, I argue that when production requires combining technology-specific capital and labor, providers of specialized capital have a direct stake in the formation of complementary skills. I study the interplay between firms' technology adoption decisions and workers' human capital accumulation choices in a setting where skills are *technology-specific* and compatible with capital supplied in imperfectly competitive markets. This setting generates new insights into how economies specialize across technologies and provides a new perspective on *open source* strategies observed in the digital infrastructure industry. Despite not employing labor themselves, capital providers may invest to reduce workers' training costs because a larger pool of skilled users increases downstream adoption of their technology. I argue that incentives to invest in human capital accumulation depend on market structure in the input sector, leading to a novel economic trade-off between allocative distortions from market power and gains from specialization in productive technologies.

To make this argument, I develop a tractable general-equilibrium model where workers and firms

¹The idea that technological change is the key driver of sustained economic growth traces back to Solow (1956), was later formalized by Romer (1990) and modelled as innovation-driven by Aghion and Howitt (1990) and Grossman and Helpman (1993). For a review, see The Committee for the Prize in Economic Sciences in Memory of Alfred Nobel (2025).

²For example, General Motors offers its Automotive Service Educational Program (ASEP) in partnership with dozens of community colleges across the U.S., training technicians in the use of its diagnostic and car repair tools. Similarly, Amazon offers the AWS Certified Solutions Architect credential, and Microsoft provides the Azure Solutions Architect Expert certification, both recognizing proficiency in designing and deploying cloud-based systems.

make interdependent choices of technological specialization while dominant capital providers can influence workers' costs of skill acquisition. Methodologically, I show that this environment can be represented as a nested fixed-point problem, which allows me to characterize the general equilibrium impact of firm-level shocks. The framework is helpful to assess the macroeconomic consequences of industrial concentration and I use it to inform the design of policies such as competition rules, training subsidies and regulation that shapes platform interoperability.

I apply the theory to the digital infrastructure industry, which offers rich variation to study the mechanisms at the core of the model and sits at the center of current policy debates on competition and innovation. This setting allows me to measure adoption of technology-specific capital and acquisition of technology-specific knowledge by analyzing large-scale microdata from the United States that are informative about firms' adoption of computing platforms and the evolution of demand for specialized software skills. After showing that some specialized software skills differ across computing platforms, I study the *open source* release of platform-specific software by Google, a large provider of cloud computing power. The empirical findings are consistent with my theory of subsidies to skill acquisition costs: the release of specific software is associated with increased adoption by employers of the complementary computing platform, but also a lowered wage premium for the associated specialized skills. This case study serves as a proof of concept for the core assumptions and mechanisms of the theory, motivating future empirical work on other sectors. It also yields empirical moments that inform a quantitative implementation of the model, which I use to evaluate a series of counterfactual and policy experiments.

The quantitative exercises highlight the trade-off between the model's misallocation and specialization forces for social welfare. While the balance of these two forces for aggregate productivity and welfare is theoretically ambiguous, I discuss how it depends on the value of parameters regulating how firms and workers reallocate across technologies following changes in fundamentals. The results call for caution when designing competition policies aimed at distortions in markets for technology-specific capital. Imperfect competition in these markets can bolster incentives for capital providers to invest in skill formation, with potentially beneficial effects on technological specialization and aggregate productivity.

I build on [Dixit and Stiglitz \(1977\)](#), [Hopenhayn \(1992\)](#) and [Melitz \(2003\)](#) and generalize an equilibrium model of heterogeneous goods producers by introducing two additional features.

First, I endow each producer with a technology adoption choice. All good varieties can be produced with any of the available technologies, but firms face idiosyncratic productivity draws.

Production under each technology requires a combination of technology-specific labor and capital. In contrast to related models with strategic capital providers (see, e.g., [Romer, 1990](#)), a key feature of my framework is that each type of capital requires different specific labor. Therefore, I model the labor market as segmented by types of technology-specific human capital and study the problem of workers facing a discrete choice over which type of specialized skills to acquire. These individual decisions jointly determine endogenous supply curves for each segment of the labor market.

The second fundamental feature of the model is the assumption that the imperfectly competitive capital providers, renting to goods producers, internalize indirect wage effects. Specifically, I introduce a *specific-labor margin* in the problem of capital suppliers: dominant providers internalize the wage and labor-demand responses in the sector employing workers trained for their technology. The presence of this channel shapes the markup on capital and generates strategic incentives to subsidize training in complementary technological knowledge. By investing to expand the supply of specific labor, capital providers reduce the cost of operating the technology linked to the type of capital they offer. This encourages downstream firms to adopt it and reallocates demand towards the type of capital they provide.

I model capital suppliers that are granular enough to impact factor prices in markets where they do not directly participate, and sophisticated enough to internalize this. Specifically, these firms understand that a) sponsoring the accumulation of technology-specific knowledge will change the equilibrium price of labor, b) employers' adoption decisions will respond to this. Formally, the demand constraint faced by imperfectly competitive technology suppliers will depend on the equilibrium wage in the market for complementary labor, which is no longer taken as given. Intuitively, technology providers internalize the indirect network effects through the market for specific labor. However, unlike theories of predatory pricing, suppliers encourage adoption of the technology not by pricing their capital below marginal cost, but by becoming sponsors of technological knowledge³.

The theoretical analysis delivers several insights. First, markups on technology-specific capital depend on the strength and direction of the reallocation of labor and production across technologies that follows an increase in the rental rate (Proposition 1). This channel implies that capital providers can charge higher markups when they take wage effects into account. Intuitively, when capital providers raise prices on their technology, they directly reduce demand for it both on the intensive margin (quantity of units supplied) and the extensive margin (measure of producers adopting the

³[Katz and Shapiro \(1986\)](#) introduce the term “sponsor” to refer to suppliers that engage in predatory pricing to standardize the economy on the use of their technology.

technology). The latter effect implies weaker demand for specialized workers who use the technology, and therefore lower wages. This cost adjustment, arising only when the specific-labor margin is present, partially dampens quantity losses, allowing providers to sustain higher prices.

At the same time, I show that the presence of the specific-labor margin generates incentives to engage in sponsorship of technological knowledge. I prove that the incentives to invest in human capital are monotonically increasing in the productivity of the underlying technology (Proposition 2). This result highlights that strategic investment by capital providers ultimately encourages the accumulation of human capital complementary to high-quality technologies.

Finally, I study the relative importance of the misallocation and specialization channels by evaluating how social welfare varies as competition in capital markets weakens. In general, I find that the balance between these opposite forces depends on the parameters that regulate the productivity distribution and the substitutability across technologies for both producers and workers. I present a calibrated version of the model and use it to illustrate that weaker competition in markets for technology-specific capital implies higher prices on the capital input with negative effects on consumption. However, when suppliers' markups over marginal cost are sufficiently high, they stimulate capital providers to finance training subsidies. These strategies favor adoption of more efficient technologies and ultimately result in welfare gains from reduced labor disutility.

As an empirical application of the theory and to validate key model assumptions, I perform a case study on the industry of digital infrastructure, where I focus on providers of cloud computing power and the associated market for specialized software skills. I analyze a specific event: the open-sourcing on November 2015 of TensorFlow, a machine learning program developed by Google which is complementary to the cloud computing infrastructure of the same firm, Google Cloud Platform (GCP).⁴ This is an appropriate setting to test/apply my theory because, while the rationale and the consequences of open-sourcing software are multifaceted, the immediate effect of open-sourcing software is that it makes specialized code directly available to read and customize. In the context of software designed to run machine-learning models on a specific computing platform, open-sourcing software optimized for use on a given platform effectively lowers the cost of acquiring the skills required to operate that platform, by allowing developers to study and adapt ready-made code. First, using data from the Stack Overflow Developer Survey, I document that developers proficient in TensorFlow are more than twice as likely to report using GCP in their jobs, with the relationship

⁴The release of TensorFlow has been studied in the strategy literature as a quasi-exogenous shock to AI production costs and AI-related labor (Rock (2019), Impink (2024)).

being strongest among respondents who report having substantial influence over their employers’ technology adoption decisions.⁵ I then turn to a second data source to obtain higher-frequency evidence on the evolution of skill demand and technology adoption: job postings from the U.S. labor market, available from Lightcast. Applying a text-based search algorithm, I construct (i) a time-varying measure of demand for TensorFlow as a specialized skill and (ii) a time-varying measure of cloud adoption by provider, based on employers’ demand for related software skills. In the raw data, demand for workers proficient in TensorFlow rises sharply after the November 2015 open-sourcing event, across a wide range of sectors and employers. To assess whether this surge follows increased adoption of GCP, I estimate a difference-in-differences specification and find that inferred adoption of GCP rises substantially following the open-source release, while no comparable change occurs for the two main competitors, Amazon AWS and Microsoft Azure. In a subsample of postings containing salary information, I further document that the open-sourcing event is associated with a pronounced decline in the wage premium for jobs requiring specialized skills to operate GCP. Taken together, these patterns are consistent with the assumptions and predictions of my model of capital-provider-financed subsidies to specialized training.

Related Literature. This paper contributes to the growing literature in macroeconomics that seeks to understand the macroeconomic importance of firm market power, recently reviewed by Syverson (2019). Existing work studies implications for aggregate outcomes (i.e. welfare, business dynamism, labor and capital shares) of microeconomic market imperfections such as monopsony power in labor markets (Berger et al., 2022), market power in product markets (De Loecker et al. (2020, 2021); Autor et al. (2020) or both (Deb et al., 2022, 2024)). Particularly relevant is the research on endogenous growth that studies the importance of market structure for innovation and technology diffusion (Aghion and Howitt, 1990; Klette and Kortum, 2004; Akcigit and Ates, 2023). This strand of research emphasizes product-market competition as a key driver of technology adoption and innovation by producers of final goods. My work complements these theories by shifting the focus to upstream to markets for technology-specific capital. I argue that imperfect competition at this level of the production network induces strategies to expand market share that ultimately distort downstream technological specialization.

Because production combines technology-specific capital and labor, although I study a static economy, my work is very closely related to the theory in Chari and Hopenhayn (1991). In that

⁵Stack Overflow is an online community of peer-to-peer coding support. Prior to the advent of generative AI models, it was widely regarded as the leading source of support for programming questions. Each year, the platform conducts a voluntary survey of developers to document their use and attitudes towards software and technology.

paper each type of technology is a different vintage and the authors study the rate at which new technologies are adopted. Instead, in my model technologies are not necessarily vertically ranked and I analyze how market structure impacts *which* technologies get adopted rather than *when*. In fact, whereas prior work has focused on distortions to the *scale* of innovation and technology diffusion, I study distortions to the *composition* of technologies adopted, in the spirit of [Acemoglu \(2002\)](#). In this respect, my paper is most closely related to recent work by [Lensman \(2025\)](#), who shows that market structure impacts aggregate growth through the distortion of the direction of innovation in technologies. That work focuses on product-market distortions and innovation, while I emphasize input-market distortions and technology adoption. In doing so, I provide a microfoundation for human-capital accumulation through workers’ skill choices and highlight how interactions between different input markets can shape the direction of technology adoption.⁶

A central point of this paper is that technology adoption is linked to the direction of the accumulation (i.e. specialization) of human capital. The core mechanism of this paper is that capital providers shape the specialization pattern precisely by influencing incentives to acquire technology-specific skills. For this reason, my work connects to a large literature studying human capital accumulation (see, e.g., [Becker, 1962](#)) and skill-biased technical change (see, e.g., [Krusell et al., 2000](#)). [Hsieh et al. \(2019\)](#) study how distortions to the allocation of workers across occupations hinder growth, I contribute a theory of how skill acquisition choices may be distorted by large firms that seek market advantages. More closely related is the literature on the labor market effects of technology, reviewed in [Acemoglu and Autor \(2011\)](#). Particularly, this work relates to recent work on artificial intelligence and automation ([Acemoglu and Restrepo, 2019, 2020](#); [Autor et al., 2003, 2024](#)) which shows that employers’ decisions to adopt new technology impacts labor markets because of production complementarities/substitutabilities. The closer paper is that of [Contractor and Taska \(2024\)](#) who show that software adoption by employers increases wage inequality by pushing up wages of skilled workers complementary to software. I add to this analysis by studying what determines firms’ adoption decisions and by introducing the strategic behavior of tech innovators participating in non-competitive factor markets.

The mechanism I propose builds on classic models of indirect network effects ([Katz and Shapiro, 1986](#)) and tying strategies ([Adams and Yellen, 1976](#); [Carlton and Waldman, 2002](#)), where firms expand demand for one product by linking it to a complementary good, creating interdependence across markets. My framework extends this logic to a setting in which technology-specific capital

⁶Exploring whether the input-market interactions of my framework also matter for innovation and growth would be an interesting topic of future research.

must be combined with technology-specific labor. Capital providers effectively tie factor markets by shaping training costs and wages, creating a general-equilibrium channel through which market power distorts technology adoption.

Finally, both the empirical analysis and the theory support the view of open-sourcing as a means to expand demand for complementary products or platforms. This contributes to an active literature on open-sourcing that builds on the seminal theoretical work by [Lerner and Tirole \(2002, 2005\)](#). While most papers focus on incentives of private individuals to contribute to open source projects, my work is involved with the rationale behind strategies of open-sourcing proprietary software of large and established corporations.⁷ Relatedly, [Athey and Ellison \(2014\)](#) examine openness as a strategic platform choice balancing control and participation. More closely, [Kumar et al. \(2011\)](#) study firms' strategic decisions to develop features for open-source software in competitive environments, highlighting how openness shapes product market competition. Finally, [Gortmaker \(2024\)](#) analyzes the welfare effects of open-source investment and draws policy implications. I contribute to this literature by providing empirical evidence on the links between open sourcing, technology adoption and skill formation. To the best of my knowledge, I am the first to provide a general-equilibrium model of how open sourcing may impact labor market outcomes and aggregate productivity through technology adoption.

Outline. The rest of the paper is organized as follows. Section 2 provides two running examples from real-world settings. Section 3 presents the model setup, assumptions about market structure and the definition of the equilibrium. Section 4 characterizes the equilibrium and presents theoretical results. Section 5 summarizes empirical findings when the theory is applied to open-sourcing in the U.S. labor market and Section 6 includes counterfactual analysis and policy experiments from a quantitative example of the theoretical framework. Section 7 concludes.

2 Illustrative Examples

The model that will be outlined in Section 3 applies to sectors where workers' tech-specific knowledge (i.e. specialized human capital tied to a particular production technology) is an essential input in production and must be combined with tech-specific capital. I focus on environments where markets for tech-specific capital are imperfectly competitive and I allow capital providers to influence not only rental rates but also the composition of specialized skills of the economy. By strategically

⁷Lerner and Tirole introduce a related "Alumni effect", whereby open-source code is more likely to be used in educational settings, lowering future programming costs on those platforms. Their framework, however, does not analyze firms' strategic incentives to release proprietary code to increase demand for complementary products.

shaping the supply of compatible specialized labor, capital providers can increase the adoption of the technology for which they supply capital. The following examples illustrate settings in which capital providers, facing such incentives, directly or indirectly sponsor the training of workers in the skills required to operate their technology.

2.1 Automotive Technicians and Manufacturer-Sponsored Programs.

Vehicles are repaired in repair centers where specialized technicians use manufacturer-supplied diagnostic tools.⁸ While mechanical repair shops and collision centers are typically independently owned, they can be certified by carmakers. Therefore, car manufacturing companies provide tech-specific physical capital (tools and equipment) and intangible capital (certifications). In the United States, we observe a few large players in the U.S. automotive industry that engage with community colleges to support brand-specific knowledge of technicians, who will later move to work in independent service centers.⁹

2.2 Open Source Software and Adoption of Cloud Infrastructure.

Different types of machine learning models run better when integrated with different computing infrastructure. While computing resources are critical to analyze data, so is the infrastructure-specific expertise of developers, workers who write code and run data analysis programs on the cloud infrastructure that its employer offers. In the U.S. the sector of cloud infrastructure has very few players, and most firms in the economy rent their resources to fulfill the computational needs of their businesses. Some of these companies write computer programs for models that run with their infrastructure and facilitate the work of developers by making these codes available for reading and writing (i.e. by publishing them with an *open source*). I will focus on this context in the empirical application in Section 5.

2.3 General Insights from the Examples

Automotive and digital infrastructure are examples of industries for which rapid technological change and heavy reliance on human capital are defining characteristics. Interestingly, we observe that in both these industries capital providers actively invest in the specialized training of

⁸General Motors competes with a few other firms for the provision of diagnostic tools: Autel, Launch, Snap-on.

⁹As part of the ASEP program, General Motors collaborates with more than 50 community colleges in the U.S. to offer certification programs that train technicians to repair Chevrolet, Buick, GMC, and Cadillac vehicles. Honda sponsors a similar certification program, PACT. Toyota and Lexus invest in a pipeline of dealers and technicians to the T-TEN training program.

their future/potential customers’ workforce. While a detailed account of the institutional dynamics of each industry lies beyond the scope of this paper, these examples raise a broader question: can we extract general economic insights from these recurrent patterns? Empirical observation alone cannot fully reveal the incentives that drive capital providers to invest in training, nor the general equilibrium consequences of such strategies. Motivated by this observation, the next Section develops a general equilibrium model designed to capture the essential features common across these cases: (i) workers who acquire tech-specific skills; (ii) firms that adopt specialized technologies and demand corresponding forms of human and physical capital; and (iii) capital providers that compete for adoption and strategically invest to lower barriers to acquiring technological knowledge linked to their systems. The model will allow me to evaluate general-equilibrium effects of investment in skill supply and to undertake policy counterfactuals.

3 A Model of Strategic Investments in technology-specific Knowledge

In this section, I provide a new general equilibrium theory of how firms and workers make technology-specific decisions when choosing, respectively, how to produce and which skills to acquire. The static model departs from standard theories of human capital accumulation and technology adoption in two respects. First, technology-specific capital markets are imperfectly competitive, and capital providers internalize both the direct effect of their actions on producers’ technology adoption decisions and the indirect effect operating through technology-specific-labor markets. Second, providers of technology-specific capital can undertake investments to reduce access barriers faced by workers when acquiring technology-specific knowledge. In the following sections, I describe the main building blocks of the model and illustrate them with examples from real-world settings. I then proceed to formally define the problem of each agent while motivating my modelling assumptions.

3.1 Environment

I model a static economy with three types of agents: goods producers, workers and capital providers. Their interactions are summarized in Figure 1. I summarize notation in Appendix Table 3. The first model block is a measure $\bar{I}$ of firms producing differentiated varieties of consumption goods. Each firm $i \in \mathcal{I}$ chooses which technology $j \in \mathcal{J}$ to adopt to produce its variety, conditional on idiosyncratic productivity shocks $\mathbf{z}^i = \{z_j^i\}_{j \in \mathcal{J}}$. Production inputs are technology-specific: firms adopting technology j must hire technology-specific-labor n_j taking as given the wage w_j prevailing

in the corresponding segment of the labor market, and must rent technology-specific capital k_j at cost r_j from a specialized capital provider.

The second block of the model is a measure $\bar{N}$ of workers, indexed by $q \in \mathcal{Q}$. Each worker/-consumer chooses one technology-specific skillset to acquire and then provides one hour of labor inelastically in the segment of the labor market where their specialized skills are demanded. Workers are heterogeneous in their disutility $\xi^q = \{\xi_j^q\}_{j \in \mathcal{J}}$ of providing specialized labor and face a (finite) access cost $\tau_j \in [0, \bar{\tau}]$ when acquiring technology-specific knowledge. These training costs are homogeneous across workers and independent of the individual idiosyncratic characteristics. The level of training costs for technology j has a default level of $\bar{\tau}$ and it is a function of o_j , the level of investment in training by the provider of the associated capital. Additionally to labor income, workers earn profit income from owning shares in the economy's firms and rental income from land used by capital providers. Workers use these funds to finance consumption of a composite of the differentiated varieties $i \in \mathcal{I}$.


The third block consists in providers of technology-specific capital. These firms produce technology-specific capital k_j using land L_j and rent it to goods producers. I assume capital providers do not directly participate in labor markets, but can indirectly affect the equilibrium price and quantity of the technology-specific-labor compatible with their capital. In particular, they can invest resources o_j at cost $\phi(o_j)$ to reduce the access barrier faced by workers when acquiring technology-specific knowledge, τ_j , below its default level $\bar{\tau}$. For notational simplicity, I denote $\tau_j \equiv \tau(o_j)$ whenever no confusion arises.

3.2 Market structure.

In the baseline specification, I make the following assumptions about the competitive structure of each market. Labor markets are assumed to be segmented by skills, but otherwise perfectly competitive. Workers and employers take wages and training costs as given. Goods producers are monopolistically competitive: they internalize demand for their variety and set prices. Capital providers are also monopolistically competitive. They internalize total demand for their type of capital.

As discussed more extensively below, an important feature of the model is the assumption that capital providers internalize their impact on the equilibrium price and quantity of labor specific to their technology and how this influences the total demand for their type of capital. Therefore, capital providers in this economy can exert their market power on goods producers in two ways:

Figure 1: **Economic Environment and Market Interactions.**



Notes: The figure summarizes interactions among three agent types: workers ($q \in \mathcal{Q}$), goods producers ($i \in \mathcal{I}$), and technology providers ($j \in \mathcal{J}$). Each provider supplies technology-specific capital k_j and may invest o_j to lower training cost τ_j . Producers adopt one technology j , demanding both capital at rate r_j and labor n_j at wage w_j . Workers choose skill specialization based on expected wages and training costs. Providers influence outcomes via (i) direct market power in pricing, and (ii) the novel *specific-labor margin*, i.e., their ability to shift wages w_j through investment o_j that reduces $\tau_j(o_j)$. These two margins jointly determine the economy's equilibrium allocation of skills and production processes across technologies.

first, by charging a markup over marginal cost on the capital they supply and second by influencing the specialization of human capital in a direction that favors adoption of their technology.

3.3 Workers

There is a measure $\bar{N}$ of workers, indexed by $q \in \mathcal{Q}$ with preferences described in eq. 1.

$$U_j = \ln(\mathbf{C}^q) - \xi_j^q \quad (1)$$

Each worker derives utility from consuming a bundle of varieties $\mathbf{C}^q = \left(\int_{i \in \mathcal{I}} c_i^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}$ and provides inelastically one unit of labor at idiosyncratic disutility cost $\xi_j^q \sim T1EV(\eta)$, where $j \in \mathcal{J}$ denotes the type of specialized skills the worker chooses to acquire, i.e. the segment of the labor market where the worker chooses to participate.

Participating in a market entails training in technology-specific knowledge, captured by the idiosyncratic disutility cost ξ_j^q and by a time cost to access technology-specific knowledge, $\tau_j \in [0, \bar{\tau}]$. In particular, I assume that each worker choosing to specialize in skill $j \in \mathcal{J}$ provides one hour of labor, but only a fraction $(1 - \tau_j)$ of that time can be allocated to production and be paid at the

hourly wage w_j . This is because a fraction of hours τ_j must be devoted to accessing technology-specific knowledge¹⁰. These access costs are homogenous across workers and independent of their disutility of providing specialized labor. While the default level of these access costs is $\bar{\tau}$, this can be lowered through investment by providers of technology-specific capital, as discussed later.¹¹ The budget constraint of worker q is

$$PC^q = \int_{i \in \mathcal{I}} p^i c_i^q di = w_j (1 - \tau_j) + \Pi_j \quad (2)$$

where p^i are prices of goods varieties and P denotes the ideal price index of $\mathbf{C}^q$ (which I normalize to 1). I denote by $\Pi_j = \alpha_j \Pi$, where α_j is the share of financial income accruing to a worker in labor market j .¹²

3.4 Goods Producers

There is a continuum of measure $\bar{I}$ of producers¹³ of differentiated good varieties indexed by $i \in \mathcal{I}$. Each producer adopts one technology $j \in \mathcal{J}$ and then hires specialized labor and capital to operate it. While each variety can potentially be produced with any technology, I assume each match has a different productivity. In particular, production of variety $i \in \mathcal{I}$ with technology $j \in \mathcal{J}$ is given by

$$y^i = z_j^i \frac{(n_j^i)^\rho (k_j^i)^{1-\rho}}{\rho^\rho (1-\rho)^{1-\rho}} \quad (3)$$

where k_j^i is capital input and n_j^i is labor input. I assume z_j^i is an i.i.d. draw from a Fréchet distribution with parameter $\theta > 0$ and $\rho \in (0, 1)$ is the labor share in production costs. Notice that in the limit case where $\theta \rightarrow \infty$, productivity heterogeneity collapses and technologies are perfectly substitutable from the point of view of goods producers while as $\theta \rightarrow 0$ idiosyncratic productivity shocks are very dispersed.

Profits of good producer $i \in \mathcal{I}$ when choosing technology j are given by

$$\pi_j^i = p^i y^i - w_j n_j^i - r_j k_j^i \quad (4)$$

¹⁰Access costs are meant to capture any time investment necessary to acquire technology-specific knowledge which is independent of individuals talent or preferences. For instance, they could capture the time required to retrieve appropriate training materials or to enroll in specialized training course.

¹¹To simplify notation, I use τ_j for $\tau(o_j)$

¹²For tractability, I assume aggregate financial income in the economy (i.e. profits of goods producers, capital providers and rental income from land), Π , is distributed lump sum to workers, proportionally to their labor income: $\Pi_{j(q)} = \alpha w_{j(q)} (1 - \tau_j) \Pi$ with $\alpha \in (0, 1)$ and $\Pi = \Pi^F + \Pi^K + p^L \bar{L} - T$.

¹³Throughout, I equivalently refer to goods producers as “firms” or “employers”.

I denote by Π^F total profits in this sector.

3.5 Providers of Technology-Specific Capital

There is a large but finite number of capital providers.¹⁴ Each is a monopolist in the market for its technology, therefore I index capital providers and technologies interchangeably, without loss of generality. Providers produce and price technology-specific capital r_j and invest resources $o_j \in [0, \bar{\tau}]$ to lower barriers to technology-specific training.

A technology provider $j \in \mathcal{J}$ produces technology-specific capital k_j from land L_j , an input in fixed supply with price p^L .¹⁵ To emphasize the indirect impact of capital providers on labor markets through technology adoption, I abstract from labor use in production of technology-specific capital. Providers are heterogeneous in productivity with $\varphi_j \sim G(\varphi)$. The production function for technology-specific capital is linear:

$$k_j = \varphi_j L_j. \quad (5)$$

Capital providers set the rental rate r_j and investment o_j internalizing how these impact total demand for technology-specific capital emerging from goods producers' adoption decisions. Demand faced by capital provider $j \in \mathcal{J}$ is:

$$k_j^d = \int_{i \in \mathcal{I}} k_j^i di = k_j(r_j, w_j; Y, \mathbf{W}) \quad (6)$$

where Y is aggregate output and $\mathbf{W}$ is an aggregate wage index, both taken as given in the monopolistic competition setup. While in standard models of imperfect competition the wage of the other input w_j is taken as given, a core feature of my framework is that capital providers internalize how their choice of r_j and o_j impacts the equilibrium price and quantity of labor specific to technology j and how this feeds back into the goods producers adoption decisions determining total demand for specific capital. This assumption is what generates the central mechanism in this model, which I refer to as the *specific-labor margin* of capital provision.

Capital providers therefore face another constraint, which is given by the equilibrium price of specific-labor:

¹⁴I focus on distortions arising from market power of capital suppliers and abstract from other market imperfections in capital markets. A literature in macro-finance has studied the importance of financial frictions, such as borrowing constraints or costly external finance (Kiyotaki and Moore (1997), Bernanke et al. (1999), Buera et al. (2011), Midrigan and Xu (2014), Ottonello and Winberry (2024)).

¹⁵Each unit of land can only be occupied by one capital provider. In the example of providers of computing power, the land fixed input captures the costs associated to renting land required to build a data center.

$$k_j^d = \int_{i \in \mathcal{I}} k_j^i di = k_j(r_j, w_j^{eq}; Y, \mathbf{W}) \quad (7)$$

$$w_j^{eq} = w_j(r_j, \tau_j(o_j); Y, \mathbf{W}) \quad (8)$$

I assume that firms can subsidize access to technology-specific knowledge: firms can lower the cost of access to technology-specific training faced by workers by paying a cost. In particular, I assume that the cost of access to technology-specific training entering the worker's budget constraint in Equation (2) is given by

$$\tau_j = \bar{\tau} - o_j \quad (9)$$

where $\bar{\tau} > 0$ is a default cost, which I assume to be homogeneous across technologies and $o_j \in [0, \bar{\tau}]$ is provider j 's investment in technology-specific knowledge, incurred at finite cost $\phi_j(o_j) \in [0, \infty)$.¹⁶ The costs incurred by capital providers are expressed in units of the consumption good; total investment costs in the economy are denoted by Φ^o . The total fixed cost Φ^o is financed through spending on individual varieties according to

$$\phi^i = \left(\frac{p^i}{P} \right)^{-\varepsilon} \Phi^o, \quad (10)$$

so that the contribution of each variety $i \in \mathcal{I}$ reflects its CES expenditure share. Appendix B.3.3 develops a microfoundation in which an intermediary implements this allocation using the same aggregator that characterizes consumer preferences. Intuitively, capital providers can incur costs to build a systematic body of technological knowledge relevant for the type of specialized productive equipment they rent, making it easier for workers to gain expertise in the use of the technology. Alternatively, they could sponsor training programs that teach prospective workers how to operate their systems.

Net profits of technology provider $j \in \mathcal{J}$ are given by

$$V(r_j, k_j, o_j; \varphi_j) = \left(r_j - \frac{p^L}{\varphi_j} \right) \times k_j(r_j, w_j^{eq}(r_j, o_j); Y, \mathbf{W}) - \phi(o_j) \quad (11)$$

I denote by Π^K total net profits in this sector.

¹⁶In general, I assume the cost function $\phi(\cdot)$ to be increasing and convex. In Section 4, I'll focus on a simplified case in which investment is a binary choice subject to a fixed cost.

3.6 Equilibrium

I define the general equilibrium of the static economy, where labor, capital, financial and land markets clear. Goods producers and capital providers engage in monopolistic competition.

Definition 1 (General Equilibrium). *An equilibrium for the economy is an allocation*

$$\left\{ \{n_j, k_j, L_j\}_{j \in \mathcal{J}}, \{c_i\}_{i \in \mathcal{I}}, C, Y \right\},$$

investment strategies $\{o_j\}_{j \in \mathcal{J}}$, shares of financial income $\{\alpha_j\}_{j \in \mathcal{J}}$ and a price vector¹⁷

$$\left\{ \{w_j, r_j\}_{j \in \mathcal{J}}, \{p^i\}_{i \in \mathcal{I}}, q, \mathbf{W} \right\},$$

such that the following conditions hold:

- (i) **Worker optimality.** *Given wages $\{w_j\}$, training costs $\{\tau_j\}$, aggregate profits, and goods prices $\{p^i\}$, each worker chooses consumption and labor-market participation to maximize utility (Equation (1)) subject to the budget constraint (Equation (2)).*
- (ii) **Goods producers' optimality.** *Given input prices and idiosyncratic productivity draws, each goods producer chooses a technology, prices, and input demands to maximize profits (Equation (4)), subject to the technological constraint (Equation (3)) and the demand for its variety for consumption (Equation (20)) and investment costs (Equation (10)).*
- (iii) **Capital providers' optimality.** *Each capital provider chooses the supply of technology-specific capital and the training subsidy (or investment decision) to maximize net profits (Equation (11)), taking as given (a) the demand for its capital (Equation (7)), (b) the equilibrium wage equation for specialized labor (Equation (8)), (c) its production constraint (Equation (5)), and (d) the investment strategies of other providers. They take as given aggregate output and the aggregate wage index.*

¹⁷The aggregate price level P is normalized to one and thus omitted.

(iv) **Market clearing.** All markets clear:

$$y^i = c^i + \phi^i, \quad \forall i \in \mathcal{I}, \quad (12)$$

$$n_j^d = n_j^s, \quad k_j^d = k_j^s, \quad \forall j \in \mathcal{J}, \quad (13)$$

$$Y = C + \Phi^o, \quad (14)$$

$$\sum_{j \in \mathcal{J}} n_j = \bar{N} \quad (15)$$

$$\sum_{j \in \mathcal{J}} L_j = \bar{L} \quad (16)$$

$$\sum_{j \in \mathcal{J}} \alpha_j = 1 \quad (17)$$

In the next section, I characterize this equilibrium as a three-layered fixed point problem and derive key theoretical results.

4 Theoretical Analysis

In this section, I characterize the decentralized equilibrium of the economy after making functional form assumptions and study its welfare properties. In particular, I describe the solution to the discrete choice problems of workers and firms specializing in one type of technology. Solutions to these problems give us the technological specialization pattern of the economy, i.e. distributions of specialized skills across workers and technologies across producers. I characterize the strategic incentives of providers of technology-specific capital to shape these equilibrium distributions through subsidies to the costs of specialized training. In describing the solution of the model, I make reference to a comparison in which providers of technology-specific capital do not internalize their pecuniary externalities on the market for labor specific to their technology, but take the price of the other factor of production as given. Comparing the results to this benchmark offers insight into the core mechanism proposed in this paper, which I refer to as the *specific-labor* margin. I discuss how the equilibrium distribution of skills and technologies in the economy compares across the two models and derive implications for equilibrium markups and productivity. These results will inform the policy experiments and welfare analysis in Section 6.

For tractability and to simplify the exposition, I assume that capital providers face a binary decision: they either invest a fixed amount of resources, denoted $\bar{o}$, or abstain from investing. The

investment entails a fixed cost $\phi^o > 0$.

The assumption of a fixed cost of investment in training is based on the running examples described in Section 2. In the case of carmakers' partnerships with community colleges, establishing a program typically requires dedicated facilities and hiring of qualified instructors, costs that are largely invariant to small changes in enrollment. Likewise, the decision to open source specialized software generally entails allocating a team of developers to adapt, standardize, and document the codebase so that it can be disseminated as a coherent package. These expenditures do not scale with the number of eventual users.¹⁸

Assumption 1. (*Fixed cost of sponsorship*) *Sponsorship is a binary decision which entails a fixed cost.*

$$o_j \in \{1, \bar{o} > 1\} \quad , \quad \phi(o_j) = \bar{\phi} \times \mathbf{1}[o_j = \bar{o}].$$

Throughout, I also assume that the dispersion in producers' productivity θ and consumers' elasticity of substitution across varieties ε are such that $\theta > \varepsilon - 1$, a parameter restriction needed for aggregation of firm-level outcomes.

4.1 Technological specialization: skill acquisition and technology adoption.

In this section, I solve for the optimal technological specialization choice of workers and goods producers. Workers specialize by acquiring technology-specific human capital and firms specialize by choosing which technology to adopt. For each worker, a skill choice corresponds to an exclusive choice of a technology-specific-labor market. I solve the problem in Section 3.3 and aggregate individual worker choices to the level of the technology/skill to obtain total labor supply for skills compatible with a given technology.

Lemma 1. *The optimal choices implied by the worker's problem are characterized by the following expressions:*

- (i) **Worker choice probabilities.** *The share of workers (or the probability that a given worker) specializes in technology $j \in \mathcal{J}$ is*

$$\varsigma_j = \left(\frac{w_j(1 - \tau_j)}{\mathbf{W}} \right)^{1+\eta}, \quad \text{where } \mathbf{W}^{1+\eta} = \sum_{j' \in \mathcal{J}} (w_{j'}(1 - \tau_{j'}))^{1+\eta}. \quad (18)$$

¹⁸For instance, the TensorFlow package discussed in the next section was built by a dedicated team of Google Brain engineers. We can think of the direct cost of open sourcing as the payroll cost of these researchers and/or the opportunity cost of having these researchers work on alternative projects. Both costs are largely unrelated to the number of end users of the package.

(ii) **Labor supply in efficiency units.** Aggregate labor supply in each market segment j is

$$n_j = \bar{N} \varsigma_j (1 - \tau_j), \quad (19)$$

where $\bar{N}$ is the total mass of workers.

(iii) **Goods demand.** Demand for each differentiated variety $i \in \mathcal{I}$ follows

$$c_i = \left(\frac{p^i}{P} \right)^{-\varepsilon} C, \quad (20)$$

$$P = \left(\int_{i \in \mathcal{I}} (p^i)^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}, \quad (21)$$

where the price index P is normalized to one.

Proof: See Appendix B.1.

In the equations above, n_j denotes total efficiency units of labor of type j . Notice the role of the cost of access to training τ_j : other things equal, the type of labor complementary to a technology with low access costs (low τ) will be in higher supply and therefore lower wages will be required to attract workers.

Turning to the problem of goods producers described in Section 3.4, their technology adoption decision is based on the comparison of expected profits (or, equivalently, costs of production) across all available possibilities.¹⁹ Since the technology of each producer does not depend on the variety they produced (with the exception of the productivity draw), the probability of adoption is symmetric across producers of different varieties. It can be expressed as the ratio of a technology's unit cost to an aggregate cost index. After determining the assignment of producers to technologies, total input demands in each segment of capital and labor markets can be obtained by aggregating individual producers' demands at the level of each technology $j \in \mathcal{J}$. The following lemma summarizes the solution of the firm's problem.

Lemma 2. *The optimal choices implied by the producer's problem imply the following expressions:*

(i) **Technology adoption.** *The share of producers (or the probability that a given producer of variety $i \in \mathcal{I}$) adopts technology $j \in \mathcal{J}$ is*

$$s_j = \left(\frac{\bar{\pi}_j}{\bar{\Pi}} \right)^{\frac{\theta}{\varepsilon-1}} = \Phi^{-1} \left(w_j^\rho r_j^{1-\rho} \right)^{-\theta}, \quad (22)$$

¹⁹The structure of heterogeneity of producers is similar to Eaton and Kortum (2002), while in their model countries choose where to source goods, here producers choose which technology to adopt.

where $\bar{\Pi}^{\frac{\theta}{\varepsilon-1}} = \sum_{j' \in \mathcal{J}} \bar{\pi}_{j'}^{\frac{\theta}{\varepsilon-1}}$ and $\Phi = \sum_{j' \in \mathcal{J}} (w_{j'}^\rho r_{j'}^{1-\rho})^{-\theta}$ are profit and cost indices, respectively.

(ii) **Input demands.** Conditional on adopting technology j , total factor demands are

$$n_j = \bar{I} \rho (\varepsilon - 1) A r_j^{-\theta(1-\rho)} w_j^{-\rho\theta-1}, \quad (23)$$

$$k_j = \bar{I} (1 - \rho) (\varepsilon - 1) A (\varepsilon - 1) w_j^{-\theta\rho} r_j^{-\theta(1-\rho)-1}, \quad (24)$$

where $A = A(\bar{\Pi}, Y, P)$ is a term that only depends on aggregate outcomes.

Proof: See Appendix B.2.

Imperfectly competitive providers of technology-specific capital internalize total capital demand in Equation (24). Importantly, the demand for technology-specific capital is inversely related to the prevailing wage in the market for labor that can operate the technology. In the next section, I characterize the equilibrium wage equations for all labor markets and then discuss how internalizing these changes the optimal behavior of technology suppliers.

4.2 Capital providers: pricing and investment strategies.

In each labor market $j \in \mathcal{J}$, combining total labor supply and demand (Equations 19 and 23) we obtain a system of equations that can be solved for wages, for given values of the tech provider's actions.

Lemma 3. *Given choices of r_j, τ_j , wages in each skill market satisfy*

$$w_j = \left(\rho (\varepsilon - 1) \frac{A \bar{I} \bar{W}^{1+\eta}}{\bar{N}} r_j^{-\theta(1-\rho)} (1 - \tau_j(o_j))^{-(2+\eta)} \right)^{\frac{1}{\eta + \rho\theta + 2}} \quad \forall j \quad (25)$$

where $A = A(\bar{\Pi}, Y, P)$ is a term that only depends on aggregate outcomes and $\bar{I}, \bar{N}$ denote the measures of goods producers and workers.

Proof: combining total labor supply and demand, Equations (19) and (23).

The above provides a fixed point iteration procedure to solve for wages in each labor market, given aggregate variables and choices of providers of technology-specific capital.

I allow imperfectly competitive providers of technology-specific capital to internalize total capital demand, Equation (24), but also the equilibrium price of the type of specific-labor compatible with the type of capital they offer. Then, the problem of the technology provider is constrained both by total capital demand, Equation (24) and by the wage, Equation (25). In particular, the provider will

internalize how its choice of sponsoring the technology it provides by subsidizing technology-specific training impacts wages in the market for complementary labor and therefore, indirectly, employers' demand for the technology-specific capital it provides.

The solution of the problem of capital providers is characterized by a pricing decision and an investment strategy. Both of these are central to the *specific-labor margin* that this paper introduces.

4.2.1 Pricing Strategy.

This subsection characterizes optimal pricing for a technology-specific capital provider $j \in \mathcal{J}$ that internalizes how the equilibrium wage w_j of its complementary labor co-moves with its rental rate r_j . The provider maximizes net profits, Equation (11), by choice of the rental rate, taking aggregates as given.²⁰ The first result is that the presence of the specific-labor margin introduced in 3.5 implies markups that depend on the elasticity of labor supply to the rental rate.

Proposition 1 (Markup with specific-labor margin). *Define the Marshallian demand elasticity holding w_j fixed and the effective elasticity incorporating wage feedback as $\varepsilon_j^{\text{MC}} \equiv -\frac{\partial k_j}{\partial r_j} \frac{r_j}{k_j}$, $\varepsilon_j^{\text{eff}} \equiv -\frac{dk_j}{dr_j} \frac{r_j}{k_j}$. Let $\mu_j \equiv r_j/(q/\varphi_j)$ denote the gross markup. The provider's first-order condition implies a Lerner rule with the effective demand elasticity:*

$$\mu_j = \frac{\varepsilon_j^{\text{eff}}}{\varepsilon_j^{\text{eff}} - 1}, \quad \varepsilon_j^{\text{eff}} = \varepsilon_j^{\text{MC}} - \underbrace{\left(-\frac{\partial \ln k_j}{\partial \ln w_j}\right)}_{\text{input substitution}} \underbrace{\left(-\frac{\partial \ln w_j}{\partial \ln r_j}\right)}_{\text{wage channel}}.$$

Proof: see Appendix B.3.1.

Given total demand for capital, Equation (7) and the equilibrium wage equation Equation (25), in this model both input substitution and the wage channel are positive. Hence, Proposition 1 establishes that this environment of monopolistic competition with a *specific-labor margin* generally features higher markups compared to the standard set up in Dixit and Stiglitz (1977). Intuitively, capital providers recognize that raising the price of their technology-specific capital reduces demand not only through the intensive margin but also by discouraging technology adoption along the extensive margin. This contraction in adoption benefits the remaining firms, as they face less competition for specialized labor and therefore lower wages in the relevant skill segment. By internalizing this *specific-labor margin*, strategic capital providers can capture part of the resulting cost reduction through higher markups.

²⁰ All derivatives are evaluated at the symmetric environment the provider takes as given; dependence on aggregates is suppressed for clarity.

The next corollary derives the effective elasticity of labor under the functional form assumptions of this framework and illustrate how they vary with parameters.

Corollary 1. *Whenever capital providers internalize the specific-labor margin, the model in 3 results in the following effective elasticity of demand for technology-specific capital*

$$\varepsilon_j^{\text{MC}} = \theta(1 - \rho) + 1, \quad \varepsilon_j^{\text{eff}} = \varepsilon_j^{\text{MC}} - \frac{\rho(1 - \rho)\theta^2}{\eta + \rho\theta + 2}.$$

therefore $\mu_j = \frac{\varepsilon_j^{\text{eff}}}{\varepsilon_j^{\text{eff}} - \varepsilon_j^{\text{MC}}} > \frac{\varepsilon_j^{\text{MC}}}{\varepsilon_j^{\text{MC}} - 1}.$

In Appendix B, I show that Proposition 1 and Corollary 1 imply that markups decrease in the parameters governing substitutability across technology for firms and workers (respectively, θ and η) and increase in the share of labor in production (ρ).

4.2.2 Investment in technology-specific training.

In this theoretical framework, incentives to sponsor technology-specific training only arise once we make the assumption that capital providers internalize wage effects. I refer to this as the *specific-labor margin* of capital provision. If capital providers take the wage of specific-labor as given, the disutility cost of training does not enter the problem of the technology provider. With a positive cost associated to sponsoring training and no marginal benefit, no provider will engage subsidize technology-specific knowledge.

When the specific-labor margin is active, what determines incentives to invest? Do more productivity technologies, i.e. technologies whose associated capital is produced more efficiently with a given land input, attract more investment? The next proposition establishes that perceived incentives to invest for a capital provider are monotonically increasing in the provider's productivity. With binary investment opportunities and a fixed cost of investment, this ensures the existence of a minimum productivity cutoff above which providers decide to invest.

Proposition 2. *(Investment in technology-specific knowledge and productivity cutoff)*

Let the (perceived) benefit of investing $\bar{o}$ in technology-specific knowledge be $B(\bar{o}, \varphi_j) = \frac{\partial V(\varphi_j)}{\partial \sigma_j}|_{\sigma_j=\bar{o}}$ where $V(\varphi_j)$ are net profits of the capital provider, Equation (11). The following statements are true:

1. At a given level of investment $\bar{o}$, the perceived benefit of sponsorship is monotonically increasing in productivity, i.e.

$$\frac{dB(\bar{o}, \varphi_j)}{d\varphi_j} > 0$$

2. If the cost of sponsorship is independent of productivity and fixed at ϕ^o , there exists a minimum productivity cutoff for sponsoring knowledge, i.e.

$$\exists \varphi^* \quad \text{s.t.} \quad B_j(\bar{o}, \varphi_j) - \phi^o \geq 0 \iff \varphi_j \geq \varphi^*$$

Proof: See Appendix B.3.2.

4.3 Aggregation.

Given optimal behavior of all agents and market clearing, we can derive an expression for aggregate output which only depends on aggregate variables.

Below, I characterize the labor index dual to the aggregate wage index $\mathbf{W}$.

Lemma 4. Let $\mathbf{N}$ be the labor index such that $\sum_j w_j n_j = \mathbf{W}\mathbf{N}$, where $\mathbf{W}$ is the aggregate wage index defined in Equation (18). Then,

$$\mathbf{N} = \bar{N} \sum_j \left(\frac{n_j}{\bar{N}(1 - \tau_j)} \right)^{\frac{2+\eta}{1+\eta}}.$$

Notice that, because of the structure of the worker problem defined in Section 3.3, the labor index is not homogeneous of degree one.

Lemma 5. Aggregate output is characterized as

$$Y = \rho^{-1} (\bar{I})^{\frac{\varepsilon}{\varepsilon-1}} (\Phi)^{\frac{1+\theta}{\theta}} \left(\Gamma \left(1 - \frac{\varepsilon-1}{\theta} \right) \right)^{\frac{1}{\varepsilon-1}} \tilde{\Phi}^{-1} N^{eff}$$

where $N^{eff} = \sum_j n_j^s = \sum_j \varsigma_j \bar{N}(1 - \tau_j)$ denotes total efficiency units supplied in the economy, i.e. total hours of labor net of the time costs of training. In this, ς_j denotes the share of workers specializing in technology j , as derived in Lemma 1.

4.4 Solution algorithm.

The equilibrium of the economy is computed using a nested fixed-point algorithm that jointly determines (i) technology-specific investment and rental rates set by capital providers, (ii) equilibrium wages in segmented labor markets, and (iii) the aggregate land price ensuring general equilibrium. The algorithm exploits the fact that the benefit of investment of technology providers can be ranked by productivity (φ_j), which induces a monotone structure in the inner nest of the problem. The

numerical solution procedure is described in detail in Appendix B.6.

For a given land price p^L , each provider $j \in \mathcal{J}$ chooses a rental rate r_j and an investment level o_j to maximize profits. Given $\{r_j, o_j\}$, wages $\{w_j\}$ are obtained from the fixed point of the labor-market clearing condition using lemma 3. Finally, the land price p^L is updated until aggregate land demand equals the fixed supply $\bar{L}$.

The inner loop computes the joint equilibrium of the capital-providers' game and the labor market. The outer loop updates the land price via a damped adjustment rule until excess land demand is below tolerance. Convergence of the full system yields the equilibrium allocation $\{r_j^*, o_j^*, w_j^*, p^{L*}\}$ and all aggregate quantities.

4.5 Efficient benchmark and externalities.

In the previous sections I have shown that investment strategies by individual capital providers can shape the technological specialization of the economy. To evaluate whether the level of investment in specialized training that prevails in the decentralized equilibrium is efficient from a social welfare standpoint, I study the efficiency properties of the model by solving the problem of the social planner.

4.5.1 Social planner problem.

I consider a utilitarian social planner that maximizes ex-ante expected utility²¹ of all worker/consumers in the economy, weighting each individual equally, an assumption motivated by the fact that workers are ex-ante identical and only differ after the realization of the disutility shock.

Lemma 6. *Let $M = \max_j u_j - \xi_j^q$.*

Per-capita social welfare in the economy is given by

$$W = \mathbf{E}_\xi[M] = \ln(C) - \ln \mathbf{N} + \kappa(J, \eta) \quad (26)$$

where C denotes aggregate consumption, $\mathbf{N} = \bar{N} \sum_j \left(\frac{n_j}{\bar{N}(1-\tau_j)} \right)^{\frac{2+\eta}{1+\eta}}$ is the labor disutility index derived in Section 4.3 and $\kappa(J, \eta)$ is a constant of integration.

Proof: omitted, see Appendix B.4.

Next, I consider the welfare impact of investment in training specific to technology $j \in J$ to study the marginal social benefit and marginal social cost of investment.

²¹This aggregate welfare notion is pervasive in theoretical and applied work featuring discrete choice frameworks: see Donald et al. (2025), Redding and Rossi-Hansberg (2017).

Definition 2 (Social planner problem). *The social planner maximizes the social welfare function in Equation (26) by choice of $\{o_j, s_j, \varsigma_j\}_{j \in \mathcal{J}} \in [0, \bar{o}] \times \Delta \times \Delta$ subject to the following constraints*

$$\text{resource constraint} \quad Y(n_j, o_j) = C + \phi^o \sum_j 1(o_j = \bar{o}) \quad (27)$$

$$\text{agg. output (labor mkt cl.)} \quad Y = \rho^{-1} (\bar{I})^{\frac{\varepsilon}{\varepsilon-1}} (\Phi)^{\frac{1+\theta}{\theta}} \left(\Gamma \left(1 - \frac{\varepsilon-1}{\theta} \right) \right)^{\frac{1}{\varepsilon-1}} \tilde{\Phi}^{-1} N^{eff} \quad (28)$$

$$\text{labor index} \quad \mathbf{N} = \bar{N} \sum_j \left(\frac{n_j}{\bar{N}(1-\tau_j)} \right)^{\frac{2+\eta}{1+\eta}} = \bar{N} \sum_j (\varsigma_j)^{\frac{2+\eta}{1+\eta}} \quad (29)$$

$$\text{efficiency units} \quad N^{eff} = \sum_j (1-\tau_j)_j \bar{N} \quad (30)$$

$$\text{training costs} \quad \tau_j = \bar{\tau}_j - o_j \quad (31)$$

Proposition 3. *Consider a change do_j in investment in training associated to arbitrary technology $j \in \mathcal{J}$. The change in welfare, expressed in utils, is characterized by*

$$MSB_j = \underbrace{\frac{\bar{N} \left[\varsigma_j \frac{d(1-\tau_j)}{do_j} \right] Y}{N^{eff}} \frac{1}{C}}_{\text{efficiency-units}} \quad (32)$$

$$SMC_j = \frac{\phi^o}{C} + \underbrace{\frac{d \ln \tilde{\Phi} Y}{do_j} \frac{1}{C}}_{\text{business-stealing}} \quad (33)$$

Proof: omitted, see Appendix B.4.

In contrast, as derived above, each capital provider takes Φ and $\tilde{\Phi}$ as given and invests if the *perceived* private gain in its own value V_j , which is just proportional to the change in the share of firms adopting the technology, exceeds the cost ϕ^o . We can express the private marginal benefit and cost of investment in utils as

$$PMB_j = \frac{1}{C} \frac{d \ln s_j}{do_j}, \quad PMC_j = \frac{\phi^o}{C}. \quad (34)$$

Because the firm perceives only its market share in capital markets s_j , it neglects the impact of o_j on aggregate efficiency units, labor disutility, and the cost index.

Comparing equations in (32) and (33) to those in (34), it is evident that the equilibrium level of investment in specialized training generally differs from the social optimum.

4.5.2 Interpretation and limit cases.

Three forces determine the wedge between private and social incentives to invest in technology-specific trainings.

First, an aggregate *efficiency-unit externality*, first term in (32) makes the social value of investment rise above private benefits. A reduction in training costs for a given technology increases the supply of efficiency units of labor specialized in that technology, thereby raising aggregate output. The effect is amplified when workers are very sensitive to wages and training costs and reallocate skill accumulation toward the newly-sponsored technology. The social value of investment captures the gains entailed by resources being reallocated precisely to the sector that transforms time into efficiency units most efficiently.

Second, the wedge between private and social cost of investment is derived as a *business-stealing externality*, second term in (33), that tends to make private investment excessive: as s_j expands, aggregate costs rise through $\tilde{\Phi}$, because of the reallocation of production away from competing technologies.²² This effect is ignored by private agents.

The total effect of the forces above is theoretically ambiguous: the first externality points to insufficient private investment and the second one suggests excessive private investment. However, the strength of these margins depends systematically on key elasticities of substitution determining how easily production and labor reallocate across technologies. Understanding how these parameters shape the magnitude and direction of the externalities clarifies when private investment will be excessive or insufficient, indicating if the social planner needs to subsidize or tax investment in specialized training.

Role of θ , producers' substitutability across technologies. The parameter θ is the shape parameter of the Frechet distribution parametrizing producers' idiosyncratic productivity shocks across technologies. As θ increases, the productivity distribution is less dispersed and producers find it easier to substitute across technologies. In this case, the business stealing externality is particularly powerful because the cost index $\tilde{\Phi}$ is more sensitive to small reallocations of shares s_m when technologies are highly substitutable. In the limit case where $\theta \rightarrow \infty$ (technologies are perfectly substitutable in production), $\tilde{\Phi}$ responds one-for-one to reallocation, and the planner

²²While the problem of determining the optimal number of investing firms in this model parallels the analysis of optimal product variety in a monopolistic competition setting with free entry of Dixit and Stiglitz (1977), my framework departs from the standard CES environment in at least one key respect: technologies are not symmetric but are associated with different input prices. As a result, the familiar Dixit–Stiglitz result that business-stealing and variety effects offset exactly does not generally hold here.

strongly penalizes excessive investment. The effects are particularly pronounced for high values of ρ , the labor share of income in the Cobb-Douglas production function. When the wage bill is a larger share of production costs, changes in wages induce stronger reallocation of producers across technologies. Conversely, when θ is small producer reallocation across j is less sensitive to input prices and aggregate output depends more on the total measure of effective technologies and trained workers rather than their composition, making business stealing negligible in comparison to the other two externalities.

Role of η , workers' inverse substitutability across technologies. The parameter η governs the elasticity of substitution across technologies in workers' training choices. In particular, high values of η imply less dispersion in the idiosyncratic preference draws of individuals across technologies²³ and therefore higher substitutability. When this is the case, workers' specialization choice is more sensitive to differences in wages or training costs, reinforcing reallocation toward the newly-subsidized technology j and increasing the strength of the *efficiency-units externality*. This also amplifies the response of the disutility term and increases the strength of the *labor-disutility externality*, increasing the social return to investment. Symmetrically, when $\eta \rightarrow 0$, workers' reallocation following a reduction in training costs is low, eliminating the labor disutility motive for additional investment.

On balance, the business-stealing term dominates when θ and ρ are high, i.e. when technologies are easily substitutable for producers and technology specialization is particularly sensitive to the cost of specialized labor.

By contrast, when θ is low, η is small, and ρ is low (workers are very sensitive to wages and training costs but producers are not), investment generates sizable spillovers in effective labor supply and welfare, and the planner would prefer more investment than the decentralized equilibrium delivers.

In Section 6, I provide a calibrated example of the model, and discuss policy counterfactuals.

5 Empirical Exercises

In this section, I present empirical results that provide qualitative motivation for the economic mechanism described in the previous section.

Empirically investigating the objects described in my model is challenging because of data

²³The dispersion parameter of the extreme value distribution is $\frac{1}{1+\eta}$.

availability considerations. It is seldom possible to observe workers' acquisition of specialized skills and firms' adoption of specialized capital at a very granular level. Moreover, it is generally difficult to attribute these to a particular technology. Even when partial information exists, it remains difficult to track these objects over time and to quantify how the price or prevalence of a given skill or technology responds to changes in the barriers to acquiring specialized training.

To overcome these limitations, I resort to a particular case study: software skills and cloud infrastructure. This sector provides an unusually clear mapping between technology, specialized skills, and a set of identifiable firms that supply the technology-specific capital. The industry of cloud infrastructure is an example of an innovative industry to which the framework I developed should apply. It also offers a unique context in which it is possible to observe variation in barriers to accessing training in cloud-specific knowledge and to link these shifts to strategies undertaken by leading capital-providers at a particular point in time.

5.1 A Case Study: Open Source Software and Adoption of Cloud Infrastructure

When firms want to analyze large data to inform business decisions (e.g. pricing strategies, product development...), they need significant computing power. In the United States, most firms do not develop computing power in-house, but rent space on one of several types of cloud infrastructure offered by a few large providers. Once on the cloud, firms are indirectly linked to large data centers with state-of-the art chips capable of performing complex computations. While cloud infrastructure is broadly substitutable across providers for most applications, some hardware features may render a given platform particularly well suited to run specific software, optimized for specific tasks research question.

Although cloud providers develop software optimized for their own platforms to meet internal data-analysis needs, this code is generally not made available to clients. Firms that rent computing power through the cloud can only make effective use of this resource if they employ software engineers with the expertise to understand the underlying hardware and to write customized code that both addresses the firm's specific needs and is compatible with the platform.

In November 2015 Google publicly released a package enabling multiple machine learning applications.²⁴ The package was available under the Apache License 2.0, meaning it was open source: anyone could download the package for free, read and modify the contained code. Because the tool

²⁴See Dean, J., & Monga, R. (2015, November 9). *TensorFlow – Google's latest machine-learning system, open sourced for everyone.* Google Research Blog. <https://research.google/blog/tensorflow-googles-latest-machine-learning-system-open-sourced-for-everyone/>.

was considered to be at the heart of Google’s machine learning operations, the event was broadly perceived as a dramatic acceleration in the open source movement.²⁵

TensorFlow, an AI development framework enabling firms to train neural network algorithms, is a quasi-exogenous shock to AI production, reducing the costs associated with AI development and increasing the value of complementary AI-related labor (Rock, 2019).

Although TensorFlow could be used across cloud platforms, its release substantially lowered the barriers to writing software that ran efficiently on Google Cloud Platform (GCP), where specialized Tensor Processing Units (TPUs) were available. Earnings calls and financial statements from Q4 2015 do not indicate the company was also engaging in significant reductions in price on its GCP cloud computing platform.²⁶

I investigate the effects of the release of TensorFlow on the mix of cloud-specific skills and on firms’ adoption of cloud platforms in the United States over the period 2015-2019.

5.1.1 Data Sources

Data cleaning and sample selection are further described in Appendix C.

U.S. Job Posting Data I use data from United States job postings available from Lightcast over the period between 2012 and 2019. Using textual information from the body of job postings, this source makes it possible to observe very detailed information about demanded skills in an opening (e.g. not just “high-skill”, or “computer science”, but, “Python” and even “TensorFlow”). The data is available at high time frequency and it allows me to track the evolution of skills in US labor demand over time and at a very fine level. For each job posting, I construct indicators that flag required specialized software skills related to cloud platforms and use this information to study adoption of cloud platforms by US employers. A subsample of the data includes information about posted annual salaries. Summary statistics for this subsample are reported in Appendix Table 4.

Survey of Developers I complement job postings with publicly available data from the Stack Overflow Developer Survey, an annual survey of coders. The survey is administered by Stack Overflow, an online forum where developers of all levels of coding experience can share questions and answers about programming issues. It offers unique insight into coding practices, self-assessed

²⁵See, for instance, Metz, C. (2015, November 9). *Google just open sourced TensorFlow, its artificial intelligence engine*. Wired. <https://www.wired.com/2015/11/google-open-sources-its-artificial-intelligence-engine/>.

²⁶The Q4 2015 earnings call of Google/Alphabet, CEO Sundar Pichai mentions the open-source release of TensorFlow in these words “[I]n November, we open-source TensorFlow, our new machine learning system, to help accelerate discovery and development in the field”.

skill levels and technology preferences of individuals. In particular, this source offers information about the use of cloud tools in professional settings and about cloud-specific software skills. The data is available from 2011 to 2024, summary statistics for the 2024 wave are reported in Appendix Table 5.

5.1.2 Empirical results

In this section, I present empirical results documenting several findings about the evolution of firms' technology adoption and hiring decisions following the open sourcing event, which I interpret through the lens of the model in section 3 as a fall in the access barriers faced by workers when acquiring technology-specific skills.


In the context of my case study, my analyses produce four main findings:

- (i) Software developers who gained platform-specific skills specialize in the use of the related cloud platform.
- (ii) When a specialized skill becomes easier to acquire, employers' demand for it rises quickly and spreads across sectors, rather than being limited to increased hiring by the firm that facilitated access to training.
- (iii) When a specialized skill becomes easier to acquire, firms' adoption of the linked technology, at the expense of competing technologies.
- (iv) When a specialized skill becomes easier to acquire, the wage premium associated to that skill falls.

1. Specialization of software developers. I analyze software developers' reported preferences over types of cloud providers used on the job. I separate respondents based on whether they report expertise in the use of TensorFlow, which is specific to the Google Cloud Platform (GCP). As shown in Figure 2, developers with TensorFlow experience are roughly twice as likely to report both preferring and actually using GCP at work, while they are less likely to use Azure and only marginally more likely to use AWS. The results are more pronounced when restricting the sample to developers that report to have influence over the technology adoption decisions of their employer. I interpret this as evidence that software developers with specialized skills specialize in the use of the technology specific to their skills.

2. Labor Demand for TensorFlow. I analyze the text of job postings to track the count of openings that specifically ask for expertise in the use of TensorFlow. I aggregate job-posting level

Figure 2: Survey Evidence on Cloud Platform Perceptions Among TensorFlow Users




Notes: The figure reports the percentage change in the share of respondents who state that they *want to use*, *have used*, or *admire* each cloud provider for work, comparing TensorFlow users with non-users. *Left panel*: all survey respondents. *Right panel*: respondents who report having at least some influence over their employer’s technology adoption (“influencers”). Estimates are based on the 2024 U.S. wave of the Stack Exchange Developer Survey. Error bars denote 95% confidence intervals.

data at the monthly level by obtaining counts of postings that require proficiency in TensorFlow.

As shown in Figure 3, I find a steep increase in demand for this specialized skill following the open sourcing of the package. Demand for TF is not concentrated in the IT sector and does not appear to be driven by the firm that developed the package, Google. I find a pronounced increase in the number of distinct employers requiring the specialized skill, with demand spreading across the full set of NAICS two-digit industries within a few years.

3. Adoption of Cloud Platforms. I track adoption of cloud platform reflected in job postings requiring cloud-specific skills. I aggregate job-posting level data at the monthly level by obtaining counts of postings for each category of cloud-specific skill. I find that after the open sourcing of TF demand for GCP (tied to TF) rises at the expense of AWS and Azure. Specifically, I observe that job postings requiring GCP start growing faster after open sourcing of TF, a pattern not observed for competing cloud platforms. This translates in a gain in market share (as measured by job postings) by GCP, at the expense of competitors. I track the adoption of cloud platforms using job postings that specify cloud-specific skill requirements. As shown in Figure ??, following the open-sourcing of TensorFlow, demand for GCP rises markedly relative to both AWS and Azure. Specifically, the number of postings requiring GCP begins to grow at a faster rate after the open-sourcing event, whereas no comparable acceleration is observed for competing platforms. Job postings requiring AWS-specific and Azure-specific software skills experience a fall in the rate of growth. This divergence implies that GCP gained market share (as measured by the share of job

Figure 3: Evolution of Labor Demand for TensorFlow, 2015–2019



Notes: The figure illustrates the evolution of job postings that mention TensorFlow following its open-source release in November 2015. *Left panel:* monthly counts of TensorFlow-related postings across all firms, excluding Google and (in an alternative series) the IT sector. *Right panel:* number of distinct employers and industries posting TensorFlow jobs over time; the line shows the share of industries represented, and the bars show total employers. Both series indicate rapid and broad diffusion of TensorFlow adoption across firms and sectors. Source: Author’s calculations from Lightcast job postings data (2015–2019).

postings requiring cloud-specific skills) at the expense of competing platforms.

4. Salary effects. I consider the subsample of 366,547 job postings containing information on annual salaries. Summary statistics on this subsample are included in Appendix Table 8. To investigate whether the open sourcing of TensorFlow is associated with any shifts in the evolution of posted salaries, I estimate a diff-in-diff in trends specification. My empirical approach compares salaries for postings that list software skills specific to the Google Cloud Platform to those that do not, around the open sourcing event of November 2015.

Let y_{it} denote the posted annual salary of vacancy i at time t . For each cloud provider, I define indicators $p_i \in (0, 1)$ for whether vacancy i requires skills specific to platform $p \in \{GCP, AWS, Azure\}$. Let $Post_{TF} = \mathbf{1}\{t \geq \text{Nov 2015}\}$ be an indicator for the post-event period, and let $Date_t$ be a normalized time indicator such that $Date_{\text{Jan 2012}} = 0$. I estimate the following regression model.


$$y_{it} = \beta Date_t + \alpha^{GCP} GCP_i + \delta^{GCP} GCP_i \times PostTF_t + \pi^{GCP} Date_t \times GCP_i \times PostTF_t + X_i + Q_{it} + \varepsilon_{it}$$

Included controls are

$$X_i = \sum_{p \in \{AWS, Azure\}} \alpha^p p_i + \mu_{n(i)} + \nu_{o(i)}$$

$$Q_{it} = \sum_{p \in \{AWS, Azure\}} p_i \times PostTF_t + \sum_{p \in \{AWS, Azure\}} Date_t \times p_i \times PostTF_t$$

Figure 4: Adoption of Cloud Platforms Following the TensorFlow Release



Notes: The figure shows the evolution of demand for cloud-specific skills in U.S. job postings between 2012 and 2019. The black dashed line marks November 2015, the month TensorFlow was open-sourced. *Left panel:* log counts of postings mentioning each cloud platform, seasonally adjusted and detrended by the pre-event trend. *Right panel:* platform shares of total cloud-related postings, normalized by their pre-event mean. All series are net of the pre-2015 trend and mean. Source: Author's calculations from Lightcast job postings data. Additional related figures appear in Appendix D.

where $\mu_{n(i)}$, $\nu_{o(i)}$ are, respectively, NAICS-2 industry and O*NET 2019 occupation fixed effects.

I report in Table 1 regression results for OLS and alternative fixed effects specifications with heteroskedasticity-robust standard errors clustered at the fixed-effect level.

Looking at the evolution of the wage premium associated to postings demanding expertise in GCP, I observe that this falls after the open sourcing event. In particular, in the main specification with industry and occupation fixed effects, the static premium associated to expertise in GCP falls by 84%, from \$24,005 to \$3,869, with no statistically significant impact on premia associated to proficiency in the use of competing cloud platforms.

Table 1: **Effect of Cloud-Specific Skills on Posted Annual Salaries.** Each column reports estimates from OLS regressions of posted annual salary on cloud-provider skill indicators and interactions with the TensorFlow open-source release. Standard errors clustered at the fixed-effect level are reported in parentheses.

<i>Dependent variable:</i>	Posted Annual Salary			
	OLS	NAICS2 FE	O*NET 2019 FE	Both FE
GCP	27,311*** (5,875)	22,174** (8,061)	28,104*** (8,773)	24,005** (10,589)
AWS	24,478*** (3,332)	22,880*** (4,909)	20,196*** (4,495)	19,703*** (6,562)
Azure	7,098 (5,369)	1,267 (8,659)	-3,175 (3,353)	-6,547 (6,917)
Date	127 (85)	-30 (714)	766 (475)	484 (676)
Post-TensorFlow	3,909*** (469)	6,881* (3,396)	4,586** (2,192)	6,729** (3,073)
GCP $\times$ Date	-5,973** (2,528)	-3,366 (2,301)	-6,124*** (1,990)	-4,014 (2,508)
AWS $\times$ Date	3,185** (1,261)	2,391 (1,770)	2,061 (1,873)	1,595 (2,480)
Azure $\times$ Date	4,717** (1,851)	5,661 (5,668)	6,844*** (2,181)	7,381 (5,795)
GCP $\times$ Post-TF	-40,081*** (10,832)	-28,825** (11,948)	-29,332** (12,674)	-20,136** (8,515)
AWS $\times$ Post-TF	4,826 (4,734)	-744 (11,484)	-451 (7,403)	-4,754 (12,083)
Azure $\times$ Post-TF	646 (6,777)	5,684 (10,834)	6,096 (7,162)	9,880 (8,263)
Date $\times$ Post-TF	-774*** (109)	-1,196** (564)	-1,158* (652)	-1,337** (576)
GCP $\times$ Date $\times$ Post-TF	9,054*** (2,867)	5,746** (2,296)	7,309*** (2,570)	4,652*** (1,400)
AWS $\times$ Date $\times$ Post-TF	-3,120** (1,359)	-1,686 (2,850)	-1,917 (2,005)	-895 (3,244)
Azure $\times$ Date $\times$ Post-TF	-4,513** (1,949)	-5,421 (6,002)	-6,379** (2,679)	-6,997 (6,045)
Observations	366,547	366,547	366,547	366,547
R ²	0.015	0.075	0.129	0.173

Notes: Dependent variable is posted annual salary. OLS estimates with robust standard errors clustered at the fixed-effect level. Post-TF = post TensorFlow release (Nov 2015). Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. An alternative specification without time trends or interaction terms is reported in Appendix Table D.1.

Overall, reductions in provider-induced barriers to acquiring specialized software skills are associated with lower posted wage premia for those skills and higher adoption of the associated technology. These patterns are consistent with the mechanism proposed in my framework.

6 Quantitative Analysis

To gain further insight into the proposed specific-labor mechanism, I calibrate the model using moments derived from the empirical findings in Section 5 and simulate counterfactual scenarios.

6.1 Calibration.

I report the baseline calibration of key parameters in Table 2 and include more details on the calibration strategy in Appendix C.4.²⁷

I calibrate externally the elasticity of substitution between varieties of the consumption good, ε , using estimates from Simonovska and Waugh (2014); this value implies a constant markup over the marginal cost of production of each variety of 32%. The parameter of the production function of goods producers, ρ , is set to target an aggregate labor share of 67%, a standard value in the literature. The analysis of Section 4 highlights the importance of θ and η , parameters that govern the substitutability across technology for, respectively, firms and workers. Higher values of θ correspond to lower dispersion across technology-specific productivity shocks drawn by producers, and therefore higher substitutability of technologies in production of varieties of the final good. Conversely, higher values of η correspond to higher dispersion of workers' idiosyncratic disutility shocks, and therefore lower substitutability across technologies on the worker side. Since these parameters are important for the analysis, I calibrate them internally to match empirical moments derived from the analyses in Section 5. In particular, I jointly calibrate parameters $\phi^o, \bar{\tau}, \theta, \eta$ to match four empirical moments:

1. A unique investor in specialized training;
2. A reduction in the wage premium of the associated technology of 84% based on the regression results in Table 1;
3. An increase in employers' adoption of the sponsored technology of 51%, based on calculations in Appendix Table 7;
4. A coefficient of variation of wages of 0.17, based on calculations in Appendix Table 6.

I set $\bar{o} = \bar{\tau}$ to capture investments that fully remove access barriers to specialized training, in line with the application on open-sourcing. Finally, I parametrize the productivity distribution of providers of technology specific capital (which determines the vertical ranking of technologies) as a Pareto distribution with shape parameter $\beta = 1.06$. This is in line with empirical evidence on the

²⁷The goal of this calibration exercise is to provide a first quantitative implementation of the model. Data limitations prevent a tight estimation of some of the parameters and future drafts will aim to make progress on this front.

tail exponent of productivity distributions in advanced economies ([Axtell \(2001\)](#), [Luttmer \(2007\)](#), [Chen \(2022\)](#)).²⁸

In the next section, I evaluate how the equilibrium of the model varies as we change structural parameters and investigate the welfare effects of government intervention to address markup distortions and investment incentives.

Table 2: **Baseline Parameter Values**

Parameter	Value	Interpretation	Target / Source
ε	4.12	Elasticity of substitution across varieties	Simonovska and Waugh (2014)
η	10.85	Workers' inverse substitutability of technologies	Internally calibrated
θ	8.2	Shape of productivity (Fréchet)	Internally calibrated
ρ	0.8	Labor share in production (Cobb–Douglas)	Matches aggregate labor share ≈ 0.67
ϕ^o	0.05	Provider's fixed cost of investment in technology-specific training	Internally calibrated
$\bar{\tau}$	0.38	Worker's default cost of training	Internally calibrated
β	1.06	Pareto shape (capital provider productivity)	Axtell (2001)

Notes: The baseline calibration of the model combines values externally calibrated in the macro-trade literature with parameters estimated to match empirical moments.

I normalize the measure of workers and firms to $\bar{N} = \bar{I} = 100$ and the number of technologies to $J = 15$, but conduct robustness checks for different normalizations.

6.2 Training investment and technological specialization.


To illustrate the specific-labor mechanism, I perform a comparative statics exercise using the baseline calibration of the model and varying ϕ^o , the fixed cost of investment in specialized training faced by capital providers. The goal of this quantitative exercise is to study how the technological specialization of the economy varies as the private marginal cost of investment falls according to the model of Section 3.

Figure 5 illustrates the technological specialization of workers and firms and the equilibrium wage prevailing in each segment of the labor market. As the cost of investment falls, more capital providers choose to subsidize access to specialized training associated to their technology. In line with Proposition 2, the marginal provider that chooses to invest as fixed costs fall is the one with highest productivity. In doing so, they encourage supply of specific-labor, as the net hourly wage increases. In equilibrium, the share of workers choosing to specialize in subsidized technologies increases and the wage prevailing in the related segment of the labor market falls. This ultimately

²⁸Estimating a Zipf regression model on my sample of computing-power providers yields a shape parameter of $\beta = 0.53$, consistent with substantial concentration among capital providers. Since the available data are insufficient to estimate this parameter precisely, the baseline calibration follows the convention in the literature.

encourages producers to adopt the subsidized technology, which can be visualized as an increase in the share of producers choosing the most productive technology. At very low values of the fixed costs, all capital providers invest to remove barriers to training in their technology. This outcome is characterized by the same distribution of workers and producers across technology as in an economy with no investment, but the level of all wages is lower reflecting the fact that employers need to compensate a lower disutility of labor.

Figure 5: **Effects of Investment Costs on Technological Specialization and Wages.**



Notes: Each panel plots equilibrium wages (right axis, in green) and the probability that workers and producers specialize in a given technology (left axis, respectively in brown and yellow). The 15 available technologies are ordered by the productivity of their capital providers, φ_j . As the fixed cost of investment falls, more capital providers invest and subsidize access to training, lowering the wages of technology-specific-labor and raising the adoption of the corresponding technology. The calibration of remaining parameters is reported in Table 2.

6.3 Policy Experiment: Taxes/Subsidies to Investment.

To quantify how efficient is the investment level emerging from the decentralized equilibrium, I evaluate the welfare effects of a policy aimed at correcting private investment incentives. In particular, I allow the fiscal authority to impose lump sum taxes/subsidies on final consumers in order to

subsidize/tax the fixed cost of investment faced by capital providers at rate $\delta \in [-1, 1]$.

The cost to capital provider $j \in \mathcal{J}$ of investing enough resources to completely eliminate access costs to specialized training (i.e. of investing $o_j = \bar{\tau}$), in units of the consumption good, is

$$PMC^{sub}(o_j = \bar{\tau}) = \delta \times \phi^o$$


and the total amount of tax revenue to be generated²⁹ is given by

$$T = \sum_{j \in \mathcal{J}} \delta \phi^o \mathbf{1}[o_j = \bar{o}].$$

Figure 6 plots the results. As expected, subsidizing fixed costs of investment increases the number of capital providers that select into sponsorship of specialized training. As we can see in the top-left panel of the figure, subsidies are welfare improving. This suggests that the level of investment naturally emerging in the decentralized equilibrium is inefficient.

²⁹As described in Section 3, for tractability, I assume that lump sum taxes are split across consumers proportionally to their total labor income. These taxes do not distort the optimal allocation of workers across technologies because workers do not value leisure and inelastically supply one hour of labor each.

Figure 6: Effects of Subsidy to Fixed Costs of Investment.



Notes: The figure reports equilibrium outcomes from model simulations varying the investment subsidy rate δ . Subsidies are financed by lump-sum taxation. Panels display, clockwise from top left, (i) social welfare, (ii) number of investing firms, (iii) income components, and (iv) aggregate consumption. Higher subsidies reduce entry costs, encouraging firm investment and raising welfare and consumption.

6.4 Markups on Capital and Investment in Specialized Training.

Finally, I investigate how the specialization pattern of the economy varies as markups increase. In the model, given Proposition 1, a simple way to vary the level of markups is varying θ , the substitutability of technologies in production. In particular, higher values of θ imply higher values of the effective elasticity of capital demand and lower markups on technology-specific capital. I compare scenarios with markups varying from a low value of 10% to high value of 60%.³⁰

Figure 7 illustrates how consumption (panel a) and investment in specialised training (panel b) are affected. As expected, higher markups imply lower consumption of the final good. However, when markups are sufficiently high, they induce capital suppliers to invest in technology-specific training. This effect reduces labor disutility which may lead to welfare gains compared to the high markup environment, based on the extent of skill specificity.


When workers are very heterogeneous in their disutility of labor across different skill markets, skills are very specific to technology. In this scenario, training subsidies have a larger impact on labor disutility and generate large welfare gains. In contrast, when skills are very substitutable from workers' perspective, subsidies to training costs do not significantly alter the aggregate labor disutility, implying muted welfare effects through this channel.

The exercise showcases the main tension in the model between the misallocation and specialisation effects of imperfect competition. On the one hand, market imperfections arising from market power of capital providers negatively dampen consumption in the economy and are associated with low social welfare. However, rents incentivize capital providers to perform socially valuable investments in human capital. By encouraging reallocation of labor and production towards high productivity technologies, increases in markups may ultimately result in welfare gains.


These results suggest that policy-makers formulating regulations aimed at correcting markup distortions in markets associated with financing of human capital accumulation should be cautious about the potential indirect consequences on welfare through changes to the composition of skills and technology in the economy.

³⁰To the best of my knowledge, there are no empirical studies that directly estimate markups in cloud computing infrastructure markets. The only available evidence comes from financial disclosures: Amazon Web Services (AWS) reported an operating income of \$9.4b on revenues of \$23.1b in the first quarter of 2025, implying an operating margin of 39.5% (Amazon.com, Inc., Form 10-Q, Q1 2025) and a price-cost markup of 1.65.

Figure 7: Economic Outcomes and Markup Distortions under Different Skill Specificity.



Panel (a): High Skill Specificity (Low η)



Panel (b): Low Skill Specificity (High η)

Notes: Each panel summarizes equilibrium outcomes from model simulations varying the markup parameter $\mu \in [1.1, 1.6]$ under alternative assumptions about worker substitutability across technologies (η). The top-left quadrant shows the number of firms that invest; the top-right reports aggregate welfare; the bottom-left and bottom-right display equilibrium consumption and labor disutility (log scale), respectively. Labels on data points indicate values of θ , the firm-level substitutability across technologies. In Panel (a), higher markups reduce consumption and increase disutility but can raise welfare when they stimulate greater investment by capital providers. In Panel (b), higher markups again depress consumption and raise disutility, but this effect dominates when skills are broadly transferable.

7 Conclusion

Modern production processes rely on technological knowledge which is highly specific to the type of capital used in production, provided by only a few large firms. As recent research in macroeconomics highlights the importance of market power for innovation and aggregate productivity, I discuss how the unique influence of technology providers on skill formation implies that market power in input markets shapes the equilibrium specialization of firms and workers across technologies.

I develop a theoretical framework that embeds imperfect competition in technology-specific capital markets into an environment of endogenous skill accumulation and technology adoption, and I show it can be solved as a nested fixed-point problem. Because dominant providers of technology-specific capital internalize their own impact on the equilibrium price and quantity of compatible labor, they have incentives to invest in skill formation. While the inclusion of this margin implies that all capital providers charge high markups, it also predicts that providers of higher-quality capital have the strongest incentives to steer human capital accumulation toward their technologies. As a result, the economy features a novel specialization–misallocation trade-off: weaker competition among capital providers raises markups but also strengthens incentives to subsidize training, deepening specialization in productive technologies. The results call for caution when designing policies addressing competition in technology markets, as they have macroeconomic implications on the specialization of human capital and the direction of technology adoption.

I provide empirical evidence that the open sourcing of platform-specific software is linked to changes in the composition of demand for skills and technology adoption by US firms. With open-source ecosystems playing a central role in artificial intelligence, my framework offers insights for policies that shape platform interoperability and openness.

Future research should seek further empirical evidence from other industries where specialized human capital is an essential input in production and imperfectly competitive capital providers compete for demand among producers choosing which technology to adopt. Extending the framework to a dynamic setting would also be valuable. While this paper demonstrates that capital-provider-sponsored training subsidies shape the technological specialization of the workforce and characterizes the associated incentives, it remains an open question whether these mechanisms influence the rate of aggregate economic growth. Distortions in the allocation of specialized human capital and to technology adoption may affect incentives to innovate, magnifying the relevance of the specific labor margin emphasized in this study.

References

- Acemoglu, Daron**, “Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality,” *The Quarterly Journal of Economics*, 1998, *113* (4), 1055–1089.
- , “Directed Technical Change,” *The Review of Economic Studies*, 2002, *69* (4), 781–809.
- **and David Autor**, “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in “Handbook of Labor Economics,” Vol. 4, Elsevier, 2011, pp. 1043–1171.
- **and Pascual Restrepo**, “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.
- **and —**, “Robots and Jobs: Evidence From US Labor Markets,” *Journal of Political Economy*, 2020, *128* (6), 2188–2244.
- Adams, William James and Janet L Yellen**, “Commodity Bundling and the Burden of Monopoly,” *The Quarterly Journal of Economics*, 1976, *90* (3), 475–498.
- Aghion, Philippe and Peter Howitt**, “A Model of Growth Through Creative Destruction,” 1990.
- Akcigit, Ufuk and Sina T Ates**, “What Happened to US Business Dynamism?,” *Journal of Political Economy*, 2023, *131* (8), 2059–2124.
- Athey, Susan and Glenn Ellison**, “The Theory of Optimal Platform Openness,” in “American Economic Review: Papers and Proceedings,” Vol. 104 2014, pp. 184–188.
- Autor, David, Caroline Chin, Anna Salomons, and Bryan Seegmiller**, “New Frontiers: The Origins and Content of New Work, 1940–2018,” *The Quarterly Journal of Economics*, 2024, *139* (3), 1399–1465.
- , **David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen**, “The Fall of the Labor Share and the Rise of Superstar Firms,” *The Quarterly Journal of Economics*, 2020, *135* (2), 645–709.
- Autor, David H, Frank Levy, and Richard J Murnane**, “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.
- Axtell, Robert L**, “Zipf Distribution of US Firm Sizes,” *Science*, 2001, *293* (5536), 1818–1820.

- Becker, Gary S**, “Investment in Human Capital: A Theoretical Analysis,” *Journal of Political Economy*, 1962, 70 (5, Part 2), 9–49.
- Berger, David, Kyle Herkenhoff, and Simon Mongey**, “Labor Market Power,” *American Economic Review*, 2022, 112 (4), 1147–93.
- Bernanke, Ben S, Mark Gertler, and Simon Gilchrist**, “The Financial Accelerator in a Quantitative Business Cycle Framework,” *Handbook of Macroeconomics*, 1999, 1, 1341–1393.
- Buera, Francisco J, Joseph P Kaboski, and Yongseok Shin**, “Finance and Development: A Tale of Two Sectors,” *American Economic Review*, 2011, 101 (5), 1964–2002.
- Carlton, Dennis W and Michael Waldman**, “The Strategic Use of Tying to Preserve and Create Market Power in Evolving Industries,” *RAND Journal of Economics*, 2002.
- Caselli, Francesco and Wilbur John Coleman**, “The World Technology Frontier,” *American Economic Review*, 2006, 96 (3), 499–522.
- Chari, Varadarajan V and Hugo Hopenhayn**, “Vintage Human Capital, Growth, and the Diffusion of New Technology,” *Journal of Political Economy*, 1991, 99 (6), 1142–1165.
- Chen, Zhang**, “Economic Growth and the Rise of Large Firms,” *Cited on*, 2022, p. 107.
- Contractor, Zara and Bledi Taska**, “The Effect of Software Adoption on Skill Demand and Wage Inequality,” *Working Paper*, 2024.
- Deb, Shubhdeep, Jan Eeckhout, Aseem Patel, and Lawrence Warren**, “Market Power and Wage Inequality,” Technical Report, IFS Working Papers 2022.
- , —, —, and —, “Walras–Bowley Lecture: Market Power and Wage Inequality,” *Econometrica*, 2024, 92 (3), 603–636.
- Dixit, Avinash K and Joseph E Stiglitz**, “Monopolistic Competition and Optimum Product Diversity,” *American Economic Review*, 1977, 67 (3), 297–308.
- Donald, Eric, Masao Fukui, and Yuhei Miyauchi**, “Unpacking Aggregate Welfare in a Spatial Economy,” Technical Report, National Bureau of Economic Research 2025.
- Eaton, Jonathan and Samuel Kortum**, “Technology, Geography, and Trade,” *Econometrica*, 2002, 70 (5), 1741–1779.

- Goldin, Claudia Dale and Lawrence F Katz**, *The Race Between Education and Technology*, harvard university press, 2008.
- Gortmaker, Jeff**, “Open Source Software Policy in Industry Equilibrium,” Technical Report, Working Paper, Tech. Rep 2024.
- Grossman, Gene M and Elhanan Helpman**, *Innovation and Growth in the Global Economy*, MIT press, 1993.
- Hopenhayn, Hugo A**, “Entry, Exit, and Firm Dynamics in Long-Run Equilibrium,” *Econometrica: Journal of the Econometric Society*, 1992, pp. 1127–1150.
- Hsieh, Chang-Tai, Erik Hurst, Charles I Jones, and Peter J Klenow**, “The Allocation of Talent and US Economic Growth,” *Econometrica*, 2019, 87 (5), 1439–1474.
- Impink, Stephen Michael**, “Both Sides of Cloud: More Standardized Tools and More Diverse Data Analytics,” *Available at SSRN 4200675*, 2024.
- Katz, Michael L and Carl Shapiro**, “Technology Adoption in the Presence of Network Externalities,” *Journal of Political Economy*, 1986, 94 (4), 822–841.
- Kiyotaki, Nobuhiro and John Moore**, “Credit Cycles,” *Journal of Political Economy*, 1997, 105 (2), 211–248.
- Klette, Tor Jakob and Samuel Kortum**, “Innovating Firms and Aggregate Innovation,” *Journal of Political Economy*, 2004, 112 (5), 986–1018.
- Krusell, Per, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante**, “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 2000, 68 (5), 1029–1053.
- Kumar, Vineet, Brett R Gordon, and Kannan Srinivasan**, “Competitive Strategy for Open Source Software,” *Marketing Science*, 2011, 30 (6), 1066–1078.
- Lensman, Todd**, “Technology Choice, Spillovers, and the Concentration of R&D,” Technical Report, MIT 2025.
- Lerner, Josh and Jean Tirole**, “Some Simple Economics of Open Source,” *The Journal of Industrial Economics*, 2002, 50 (2), 197–234.

- **and** —, “The Economics of Technology Sharing: Open Source and Beyond,” *Journal of Economic Perspectives*, 2005, *19* (2), 99–120.
- Loecker, Jan De, Jan Eeckhout, and Gabriel Unger**, “The Rise of Market Power and the Macroeconomic Implications,” *The Quarterly Journal of Economics*, 2020, *135* (2), 561–644.
- , —, **and Simon Mongey**, “Quantifying Market Power and Business Dynamism in the Macroeconomy,” Technical Report, National Bureau of Economic Research 2021.
- Luttmer, Erzo GJ**, “Selection, Growth, and the Size Distribution of Firms,” *The Quarterly Journal of Economics*, 2007, *122* (3), 1103–1144.
- Melitz, Marc J**, “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 2003, *71* (6), 1695–1725.
- Midrigan, Virgiliu and Daniel Yi Xu**, “Finance and Misallocation: Evidence From Plant-Level Data,” *American Economic Review*, 2014, *104* (2), 422–458.
- Nelson, Richard R and Edmund S Phelps**, “Investment in Humans, Technological Diffusion, and Economic Growth,” *American Economic Review*, 1966, *56* (1/2), 69–75.
- Ottonello, Pablo and Thomas Winberry**, “Capital, Ideas, and the Costs of Financial Frictions,” Technical Report, National Bureau of Economic Research 2024.
- Redding, Stephen J and Esteban Rossi-Hansberg**, “Quantitative Spatial Economics,” *Annual Review of Economics*, 2017, *9* (1), 21–58.
- Rock, Daniel**, “Engineering Value: The Returns to Technological Talent and Investments in Artificial Intelligence,” *Available at SSRN 3427412*, 2019.
- Romer, Paul M**, “Endogenous Technological Change,” *Journal of Political Economy*, 1990, *98* (5, Part 2), S71–S102.
- Simonovska, Ina and Michael E Waugh**, “The Elasticity of Trade: Estimates and Evidence,” *Journal of International Economics*, 2014, *92* (1), 34–50.
- Solow, Robert M**, “A Contribution to the Theory of Economic Growth,” *The Quarterly Journal of Economics*, 1956, *70* (1), 65–94.
- Syverson, Chad**, “Macroeconomics and Market Power: Context, Implications, and Open Questions,” *Journal of Economic Perspectives*, 2019, *33* (3), 23–43.

The Committee for the Prize in Economic Sciences in Memory of Alfred Nobel, “Scientific Background to the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2025,” Technical Report, The Royal Swedish Academy of Sciences 2025.

Appendix

Contents

A Summary of Notation	49
B Omitted Proofs	50
B.1 Solution of Worker Problem	50
B.1.1 Choice Probabilities	50
B.1.2 Aggregate Budget Constraint	50
B.1.3 Proof of Lemma 1	51
B.1.4 Allocation of Consumption	52
B.1.5 Representative Agent Problem	55
B.2 Solution of Goods Producer Problem	55
B.2.1 Characterization of Realized Profits	55
B.2.2 Proof of Firms' Technology Choice Probabilities	55
B.2.3 Proof of Total Input Demands	56
B.2.4 Derivation of Aggregate Output	60
B.3 Solution of Tech Provider Problem	60
B.3.1 Proof of Proposition 1	60
B.3.2 Proof of Proposition 2	61
B.3.3 Microfoundation of fixed costs allocation.	63
B.4 Welfare Analysis	63
B.4.1 Proof of Lemma	63
B.5 Decomposition of Social Marginal Benefit	64
B.5.1 Proof of Proposition 3.	65
B.6 Numerical solution	65
B.7 Step 0. Initialization	65
B.8 Step 1 (innermost loop). Labor-Market Fixed Point	65
B.9 Step 2. Capital Providers' Problem	65

B.10 Step 3 (outermost loop). Land-Market Equilibrium	66
C Data Cleaning and Sample Construction	66
C.1 US Job Posting Data	66
C.2 Salary Subsample	67
C.3 Developer Surveys	68
C.4 Calibration Strategy	68
D Additional Figures and Tables	70
D.1 Salary Effects	72

A Summary of Notation

Table 3: Definitions of Main Variables and Parameters

Symbol / Notation	Definition / Description
$\bar{I}, \bar{N}, \bar{L}$	Total measures of producers, workers, and land
$i \in \mathcal{I}$	Index of differentiated goods producers; total measure $\bar{I}$
$j \in \mathcal{J}$	Technology or technology-specific capital provider
$q \in \mathcal{Q}$	Worker index; total measure $\bar{N}$
n_j	Labor input of type j (efficiency units)
k_j	Technology-specific capital rented to producers
r_j	Rental rate for specialized capital, chosen by provider j
w_j	Equilibrium wage for labor of type j in segmented market
$\tau_j(o_j), \tau_j$	Access cost to training for technology j , decreasing in provider investment o_j
o_j	Provider's investment in training subsidy
$\phi(o_j), \phi^o$	Convex/fixed cost of investment in training subsidy
p^i	Price of variety i
p^L	Price of land i
P	Price index, normalized to one
ρ	Labor share in production (Cobb–Douglas exponent)
η	Inverse dispersion of worker preference draws across technologies
θ	Fréchet shape, dispersion of firm draws across technologies
ε	Elasticity of substitution across varieties (product demand)
Y, C	Aggregate output and consumption
$\Phi, \tilde{\Phi}$	Cost indices in production; composite across technologies
s_j, ς_j	Shares of producers and workers specializing in technology j
Π^F, Π^K	Total profits of producers and capital providers
$\mathbf{N}$	Labor index
N^{eff}	Total efficiency units
$\mathbf{W}$	Aggregate wage index (dual of labor index)
$\tilde{W}$	Aggregate wage index (dual of labor measure)

B Omitted Proofs

B.1 Solution of Worker Problem

B.1.1 Choice Probabilities

Setup: $U_j = V_j + \xi_j$, $V_j = \ln\left(\frac{w_j}{(1+\tau_j)P}\left(1 + \frac{\Pi}{W}\right)\right)$, $\xi_j \stackrel{i.i.d.}{\sim} \text{T1EV}(0, s)$.

$$F(\varepsilon) = \exp(-e^{-\varepsilon/s}), \quad f(\varepsilon) = \frac{1}{s}e^{-\varepsilon/s} \exp(-e^{-\varepsilon/s}).$$

Choice prob.: $\Pr(j) = \Pr(U_j \geq U_k, \forall k) = \int_{-\infty}^{\infty} f(\varepsilon) \prod_{k \neq j} F(\varepsilon + V_j - V_k) d\varepsilon$.

Substitute F, f and set $t = e^{-\varepsilon/s}$ ($\Rightarrow d\varepsilon = -s dt/t$): $\Pr(j) = \int_0^{\infty} \exp(-t A_j) dt$, $A_j \equiv e^{-V_j/s} \sum_k e^{V_k/s}$.

$$\text{Evaluate } \int_0^{\infty} e^{-t A_j} dt = \frac{1}{A_j} \Rightarrow \Pr(j) = \frac{e^{V_j/s}}{\sum_k e^{V_k/s}}.$$

Apply V_j : $e^{V_j/s} = \left[\frac{w_j}{(1+\tau_j)P}\left(1 + \frac{\Pi}{W}\right)\right]^{1/s} \rightsquigarrow \Pr(j) = \frac{\left(\frac{w_j}{1+\tau_j}\right)^{1/s}}{\sum_k \left(\frac{w_k}{1+\tau_k}\right)^{1/s}}$ (common factors cancel).

Map parameters & compact form: $s = \frac{1}{1+\eta} \Rightarrow \frac{1}{s} = 1 + \eta$, $\mathbf{W}^{-(1+\eta)} \equiv \sum_k \left(\frac{w_k}{1+\tau_k}\right)^{1+\eta}$.

$$\varsigma_j = \left(\frac{w_j}{(1+\tau_j)\mathbf{W}}\right)^{1+\eta}$$

.

B.1.2 Aggregate Budget Constraint

The aggregate budget constraint is given by:

$$\begin{aligned} \int_q PC^q dq &= \int_q (w_{j(q)} + \Pi_{j(q)}) dq \\ P \sum_j \varsigma_j C_j &= \sum_j \varsigma_j w_j (1 + \alpha \Pi) \end{aligned}$$

where ς_j is the fraction of individuals choosing skillset j .

Let aggregate consumption be $C = \sum_j \varsigma_j C_j$. Then,

$$PC = \sum_j \varsigma_j w_j + \Pi \alpha \sum_j \varsigma_j w_j$$

Imposing that weights add up to one,

$$\begin{aligned} 1 &= \int w_{j(q)} \alpha dq \\ &= \alpha \sum_j \varsigma_j w_j \\ &= \alpha W \end{aligned}$$

where we define the wage bill $W := \int w_{j(q)} dq$. Thus, $\alpha = \frac{1}{W}$ and $\Pi_{j(q)} = \frac{w_{j(q)}}{W} \Pi$. The aggregate budget constraint then becomes $PC = W + \Pi$.

B.1.3 Proof of Lemma 1

The worker's problem is to maximize utility subject to their budget constraint:

$$\begin{aligned} \max_{j, c_i} \quad & \ln(\mathbf{C}^q) - \ln(\tau_j) + \xi_j^q \\ \text{s.t.} \quad & PC^q = \int_i p^i c_i^q di = w_{j(q)} \left(1 + \frac{\Pi}{W} \right) \\ & \mathbf{C}^q = \left(\int_i c_i^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}} \end{aligned}$$

Plugging the budget constraint into the objective function, the problem for each worker simplifies to:

$$\max_j \quad \ln \left(\frac{w_j}{P} \left(1 + \frac{\Pi}{W} \right) \right) - \ln(\tau_j) + \xi_j^q = u_j + \xi_j^q$$

For each individual q , the taste shock ξ_j^q is an i.i.d. draw from a TIEV distribution with shape parameter $1 + \eta$:

$$\xi_j^q \sim TIEV(\eta) = \exp \left[-\exp \left(-(1 + \eta) \xi_j^q \right) \right]$$

The probability of choosing skillset j is then:

$$\begin{aligned}
\text{Prob}(q \text{ chooses } j) &= \frac{\exp(u_j/(1+\eta)^{-1})}{\sum_{j'} \exp(u_{j'}/(1+\eta)^{-1})} \\
&= \frac{\exp((\ln w_j + \ln(\frac{1}{\bar{P}}(1 + \frac{\Pi}{\bar{W}})) - \ln(\tau_j))(1+\eta))}{\sum_{j'} \exp((\ln w_{j'} + \ln(\frac{1}{\bar{P}}(1 + \frac{\Pi}{\bar{W}})) - \ln(\tau_{j'}))(1+\eta))} \\
&= \frac{\left(\frac{w_j}{\tau_j}\right)^{1+\eta}}{\sum_{j'} \left(\frac{w_{j'}}{\tau_{j'}}\right)^{1+\eta}} = \left(\frac{w_j}{\tau_j \mathbf{W}}\right)^{1+\eta} =: \varsigma_j
\end{aligned}$$

This probability is independent of q and sums to one. Let $\mathbf{W}^{1+\eta} = \sum_{j'} \left(\frac{w_{j'}}{\tau_{j'}}\right)^{1+\eta}$ be an aggregate wage index, including subsidies to technology-specific training, τ .

Since each worker provides one unit of labor inelastically, total labor supply for skill j is:

$$n_j = \varsigma_j = \left(\frac{w_j}{\tau_j \mathbf{W}}\right)^{1+\eta}$$

The labor supply system is thus:

- $n_j = \left(\frac{w_j}{\tau_j \mathbf{W}}\right)^{1+\eta} \implies w_j = \mathbf{W} \tau_j n_j^{\frac{1}{1+\eta}}$
- $\mathbf{W}^{1+\eta} = \sum_{j'} \left(\frac{w_{j'}}{\tau_{j'}}\right)^{1+\eta}$

B.1.4 Allocation of Consumption

The first-order conditions for consumption allocation are:

$$\begin{aligned}
c_i : \quad & U_C \frac{d\mathbf{C}^q}{dc_i^q} - \lambda p^i = 0 \\
c_{i'} : \quad & U_C \frac{d\mathbf{C}^q}{dc_{i'}^q} - \lambda p^{i'} = 0
\end{aligned}$$

This implies:

$$\begin{aligned}\frac{p^{i'}}{p^i} &= \frac{\frac{d\mathbf{C}}{dc_{i'}}}{\frac{d\mathbf{C}}{dc_i}} = \frac{\frac{\varepsilon}{\varepsilon-1} \left(\int_i c_i^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{1}{\varepsilon-1}} \frac{\varepsilon-1}{\varepsilon} c_{i'}^{-\frac{1}{\varepsilon}}}{\frac{\varepsilon}{\varepsilon-1} \left(\int_i c_i^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{1}{\varepsilon-1}} \frac{\varepsilon-1}{\varepsilon} c_i^{-\frac{1}{\varepsilon}}} \\ p^{i'} &= \left(\frac{c_i^q}{c_{i'}^q} \right)^{\frac{1}{\varepsilon}} p^i \\ \frac{c_i^q}{c_{i'}^q} &= \left(\frac{p^i}{p^{i'}} \right)^{-\varepsilon}\end{aligned}$$

All individuals have the same allocation of consumption across varieties (it does not depend on skills).

Plugging this into the budget constraint:

$$\begin{aligned}\int_{i'} \left(\frac{c_i}{c_{i'}} \right)^{\frac{1}{\varepsilon}} p^i c_{i'} di' &= w_{j(q)} \left(1 + \frac{\Pi}{W} \right) \\ p^i c_i^{\frac{1}{\varepsilon}} \int_{i'} c_{i'}^{\frac{\varepsilon-1}{\varepsilon}} di' &= w_{j(q)} \left(1 + \frac{\Pi}{W} \right) \\ p^i c_i^{q\frac{1}{\varepsilon}} \mathbf{C}^{q\frac{\varepsilon-1}{\varepsilon}} &= w_{j(q)} \left(1 + \frac{\Pi}{W} \right) \\ \frac{c_i^q}{\mathbf{C}^q} &= \left(\frac{p^i}{P} \right)^{-\varepsilon} \left(w_{j(q)} \left(1 + \frac{\Pi}{W} \right) \right)^{\varepsilon} (P \mathbf{C}^q)^{-\varepsilon} \\ c_i^q &= \left(\frac{p^i}{P} \right)^{-\varepsilon} \mathbf{C}^q\end{aligned}$$

Aggregating across consumers, total demand for variety i is:

$$\begin{aligned}
c_i &= \int_q c_i^q dq = \int_q \left(\frac{p^i}{P}\right)^{-\varepsilon} \mathbf{C}^q dq \\
&= \int_q \left(\frac{p^i}{P}\right)^{-\varepsilon} \frac{w_{j(q)} \left(1 + \frac{\Pi}{W}\right)}{P} dq \\
&= \sum_j \text{Prob}(q \text{ chooses } j) \left(\frac{p^i}{P}\right)^{-\varepsilon} \frac{w_j \left(1 + \frac{\Pi}{W}\right)}{P} \\
&= \sum_j \frac{w_j^{1+\eta} \tau_j^{-(1+\eta)}}{\mathbf{W}^{1+\eta}} \left(\frac{p^i}{P}\right)^{-\varepsilon} \frac{w_j \left(1 + \frac{\Pi}{W}\right)}{P} \\
&= \frac{1}{P} \left(\frac{p^i}{P}\right)^{-\varepsilon} \left(1 + \frac{\Pi}{W}\right) \mathbf{W}^{-1-\eta} \sum_j w_j^{2+\eta} \tau_j^{-(1+\eta)} \\
&= \frac{1}{P} \left(\frac{p^i}{P}\right)^{-\varepsilon} (W + \Pi) \\
&= \left(\frac{p^i}{P}\right)^{-\varepsilon} C
\end{aligned}$$

This yields:

$$p^i = PC^{\frac{1}{\varepsilon}} (c^i)^{-\frac{1}{\varepsilon}} \quad (35)$$

We defined $\mathbf{W}^{1+\eta} = \sum_{j'} \left(\frac{w_{j'}}{\tau_{j'}}\right)^{1+\eta}$, and from the aggregate budget constraint, $PC = (W + \Pi)$. Notice $W = \int_q w_{j(q)} dq = \sum_j \varsigma_j w_j = \mathbf{W}^{-(1+\eta)} \sum_j (w_j)^{2+\eta} \tau_j^{-(1+\eta)}$.

Using total demand for the variety, we can guess and verify that P , the ideal price index, is:

$$P = \left(\int_i (p^i)^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}$$

This holds because:

$$\begin{aligned}
PC &= \int_i p^i c_i di = \int_i p^i p^{-\varepsilon} P^{\varepsilon} C di = P^{\varepsilon} C \int_i p^{1-\varepsilon} di \\
P^{1-\varepsilon} &= \int_i p^{1-\varepsilon} di
\end{aligned}$$

We normalize this to 1.

The demand system is:

- $c_i = \left(\frac{p^i}{P}\right)^{-\varepsilon} C \implies p^i = PC^{\frac{1}{\varepsilon}} (c^i)^{-\frac{1}{\varepsilon}}$
- $P = \left(\int_i (p^i)^{1-\varepsilon} di \right)^{\frac{1}{1-\varepsilon}}$

- $C = \int_q C^q dq$
- $\mathbf{C}^q = \left(\int_i c_i^{\frac{\varepsilon-1}{\varepsilon}} di \right)^{\frac{\varepsilon}{\varepsilon-1}}$

B.1.5 Representative Agent Problem

TODO

B.2 Solution of Goods Producer Problem

B.2.1 Characterization of Realized Profits

Each monopolistically competitive goods producer $i \in \mathcal{I}$ takes as given the demand function derived in Section B.1. The first-order conditions of the profit maximization problem imply:

$$\pi_j^i = P^\varepsilon Y \frac{1}{\varepsilon} \left(\frac{\varepsilon-1}{\varepsilon} \right)^{\varepsilon-1} \left(\frac{z_{ij}}{w_j^\rho r_j^{1-\rho}} \right)^{\varepsilon-1}$$

Then, given $z_j^i = \bar{z}_j \zeta_j^i$,

$$\pi_j^i = (\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j$$

where

$$\bar{\pi}_j = P^\varepsilon Y \frac{1}{\varepsilon} \left(\frac{\varepsilon-1}{\varepsilon} \right)^{\varepsilon-1} \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}} \right)^{\varepsilon-1}$$

B.2.2 Proof of Firms' Technology Choice Probabilities

Given $\Pr(\zeta < x) = e^{-x^{-\theta}}$, the density of the idiosyncratic productivity shock is $\Pr(\zeta = x) = \theta x^{-\theta-1} e^{-x^{-\theta}}$.

First, we obtain the density for the change of variable:

$$\begin{aligned} \Pr\left((\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j < x\right) &= \Pr\left(\zeta_j^i < \left(\frac{x}{\bar{\pi}_j}\right)^{\frac{1}{\varepsilon-1}}\right) \\ &= e^{-\left(\left(\frac{x}{\bar{\pi}_j}\right)^{\frac{1}{\varepsilon-1}}\right)^{-\theta}} = e^{-\left(\frac{x}{\bar{\pi}_j}\right)^{\frac{-\theta}{\varepsilon-1}}} \end{aligned}$$

Then, the density of $(\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j$ is:

$$\Pr\left((\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j = x\right) = \frac{d\Pr\left((\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j < x\right)}{dx} = \frac{\theta}{\varepsilon-1} x^{\frac{-\theta}{\varepsilon-1}-1} (\bar{\pi}_j)^{\frac{\theta}{\varepsilon-1}} e^{-\left(\frac{x}{\bar{\pi}_j}\right)^{\frac{-\theta}{\varepsilon-1}}}$$

The probability that firm i chooses skill j is:

$$\begin{aligned}
\Pr(i \text{ chooses } j) &= \int_0^{+\infty} \Pr\left((\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j = x\right) \prod_{j' \neq j} \Pr\left((\zeta_{j'}^i)^{\varepsilon-1} \bar{\pi}_{j'} < x\right) dx \\
&= \int_0^{+\infty} \frac{\theta}{\varepsilon-1} x^{\frac{-\theta}{\varepsilon-1}-1} (\bar{\pi}_j)^{\frac{\theta}{\varepsilon-1}} e^{-\left(\frac{x}{\bar{\pi}_j}\right)^{\frac{-\theta}{\varepsilon-1}}} \prod_{j' \neq j} e^{-\left(\frac{x}{\bar{\pi}_{j'}}\right)^{\frac{-\theta}{\varepsilon-1}}} dx \\
&= \int_0^{+\infty} \frac{\theta}{\varepsilon-1} x^{\frac{-\theta}{\varepsilon-1}-1} (\bar{\pi}_j)^{\frac{\theta}{\varepsilon-1}} e^{-x^{\frac{-\theta}{\varepsilon-1}} \sum_{j'} \bar{\pi}_{j'}^{\frac{\theta}{\varepsilon-1}}} dx \\
&= (\bar{\pi}_j)^{\frac{\theta}{\varepsilon-1}} \left(\sum_{j'} \bar{\pi}_{j'}^{\frac{\theta}{\varepsilon-1}} \right)^{-1} \left[e^{-x^{\frac{-\theta}{\varepsilon-1}} \sum_{j'} \bar{\pi}_{j'}^{\frac{\theta}{\varepsilon-1}}} \right]_0^{+\infty} \\
&= \frac{\bar{\pi}_j^{\frac{\theta}{\varepsilon-1}}}{\sum_{j'} \bar{\pi}_{j'}^{\frac{\theta}{\varepsilon-1}}} = s_j
\end{aligned}$$

Notice that:

$$\begin{aligned}
s_j &= \frac{\bar{\pi}_j^{\frac{\theta}{\varepsilon-1}}}{\sum_{j'} \bar{\pi}_{j'}^{\frac{\theta}{\varepsilon-1}}} \\
&= \frac{\left(P^\varepsilon Y^{\frac{1}{\varepsilon}} \left(\frac{\varepsilon-1}{\varepsilon} \right)^{\varepsilon-1} \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}} \right)^{\varepsilon-1} \right)^{\frac{\theta}{\varepsilon-1}}}{\sum_{j'} \left(P^\varepsilon Y^{\frac{1}{\varepsilon}} \left(\frac{\varepsilon-1}{\varepsilon} \right)^{\varepsilon-1} \left(\frac{\bar{z}_{j'}}{w_{j'}^\rho r_{j'}^{1-\rho}} \right)^{\varepsilon-1} \right)^{\frac{\theta}{\varepsilon-1}}} \\
&= \frac{\left(\frac{w_j^\rho r_j^{1-\rho}}{\bar{z}_j} \right)^{-\theta}}{\sum_{j'} \left(\frac{w_{j'}^\rho r_{j'}^{1-\rho}}{\bar{z}_{j'}} \right)^{-\theta}}
\end{aligned}$$

This does not directly depend on P, Y .

B.2.3 Proof of Total Input Demands

Producers with different levels of idiosyncratic productivity demand a different amount of labor and capital, but have the same choice probability. The total demand for each input is obtained by aggregating across individual producers of differentiated varieties.

$$n_j = \int n_j^i di$$

$$n_j^i = \begin{cases} 0 & \text{if choose } j' \neq j \\ \geq 0 & \text{if choose } j \end{cases}$$

Similarly,

$$k_j = \int k_j^i di$$

$$k_j^i = \begin{cases} 0 & \text{if choose } j' \neq j \\ \geq 0 & \text{if choose } j \end{cases}$$

Labor Demand. From the first-order condition:

$$\rho^{\frac{\varepsilon-1}{\varepsilon}} PY^{\frac{1}{\varepsilon}} \left(z_j^i \frac{(n_j^i)^\rho (k_j^i)^{1-\rho}}{\rho^\rho (1-\rho)^{1-\rho}} \right)^{\frac{\varepsilon-1}{\varepsilon}} = w_j n_j^i$$

$$w_j n_j^i = \rho \frac{\varepsilon-1}{\varepsilon} py$$

$$\pi_j^i = py - \frac{\varepsilon-1}{\varepsilon} py = \frac{1}{\varepsilon} py$$

$$\implies w_j n_j^i = \rho(\varepsilon-1) \pi_j^i = \rho(\varepsilon-1) (\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j$$

$$\implies r_j k_j^i = (1-\rho)(\varepsilon-1) (\zeta_j^i)^{\varepsilon-1} \bar{\pi}_j$$

Now, let's derive the aggregate labor demand:

$$\begin{aligned} w_j n_j &= w_j \int n_j^i(\zeta_j^i) \Pr(\zeta_j^i = \zeta) \prod_{j' \neq j} \Pr(\bar{\pi}_{j'}(\zeta_{j'}^i)^{\varepsilon-1} \leq \bar{\pi}_j \zeta^{\varepsilon-1}) d\zeta \\ &= \rho(\varepsilon-1) \bar{\pi}_j \int \zeta^{\varepsilon-1} \Pr(\zeta_j^i = \zeta) \prod_{j' \neq j} \Pr(\bar{\pi}_{j'}(\zeta_{j'}^i)^{\varepsilon-1} \leq \bar{\pi}_j \zeta^{\varepsilon-1}) d\zeta \\ &= \rho(\varepsilon-1) \bar{\pi}_j \int \zeta^{\varepsilon-1} \theta \zeta^{-\theta-1} e^{-\zeta^{-\theta}} \prod_{j' \neq j} e^{-\left(\left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}}\right)^{\frac{1}{\varepsilon-1}} \zeta\right)^{-\theta}} d\zeta \\ &= \rho(\varepsilon-1) \bar{\pi}_j \int \theta \zeta^{\varepsilon-2-\theta} e^{-\sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}}\right)^{-\frac{\theta}{\varepsilon-1}} \zeta^{-\theta}} d\zeta \end{aligned}$$

Derivations. Let $X = \sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}}\right)^{-\frac{\theta}{\varepsilon-1}} \zeta^{-\theta}$. Then $dX = -\theta \sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}}\right)^{-\frac{\theta}{\varepsilon-1}} \zeta^{-\theta-1} d\zeta$. Also,

$$\zeta^{-\theta} = X / \sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}} \right)^{-\frac{\theta}{\varepsilon-1}}. \text{ So, } \zeta = \left(\frac{X}{\sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}} \right)^{-\frac{\theta}{\varepsilon-1}}} \right)^{-\frac{1}{\theta}}. \text{ And } d\zeta = \frac{dX}{-\theta \sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}} \right)^{-\frac{\theta}{\varepsilon-1}} \zeta^{-\theta-1}}.$$

Substituting these into the integral:

$$\begin{aligned} w_j n_j &= \rho(\varepsilon - 1) \bar{\pi}_j \int_0^{+\infty} \theta \zeta^{\varepsilon-2-\theta} e^{-X} d\zeta \\ &= \rho(\varepsilon - 1) \bar{\pi}_j \int_0^{+\infty} \theta \left(\frac{X}{\sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}} \right)^{-\frac{\theta}{\varepsilon-1}}} \right)^{-\frac{\varepsilon-2-\theta}{\theta}} e^{-X} \\ &\quad \times \left(-\theta \left(\sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}} \right)^{-\frac{\theta}{\varepsilon-1}} \right)^{-\frac{1}{\theta}} X^{\frac{\theta+1}{\theta}} \right)^{-1} dX \\ &= \rho(\varepsilon - 1) \bar{\pi}_j \int_0^{+\infty} X^{\frac{-(\varepsilon-2-\theta)}{\theta} - \frac{\theta+1}{\theta}} \left(\sum_{j'} \left(\frac{\bar{\pi}_j}{\bar{\pi}_{j'}} \right)^{-\frac{\theta}{\varepsilon-1}} \right)^{\frac{\varepsilon-2-\theta}{\theta} + \frac{1}{\theta}} e^{-X} dX \\ &= \rho(\varepsilon - 1) \bar{\pi}_j \int_0^{+\infty} X^{\frac{-\varepsilon+1}{\theta}} \left(\sum_{j'} \left(\bar{\pi}_{j'} \right)^{\frac{\theta}{\varepsilon-1}} \right)^{\frac{\varepsilon-1}{\theta}} e^{-X} dX \\ &= \rho(\varepsilon - 1) \bar{\pi}_j \left(\sum_{j'} \left(\bar{\pi}_{j'} \right)^{\frac{\theta}{\varepsilon-1}} \right)^{\frac{\varepsilon-1}{\theta}} \int_0^{+\infty} X^{\frac{-\varepsilon+1}{\theta}} e^{-X} dX \end{aligned}$$

Therefore,

$$\begin{aligned} w_j n_j &= \rho(\varepsilon - 1) s_j \left(\sum_{j'} \left(\bar{\pi}_{j'} \right)^{\frac{\theta}{\varepsilon-1}} \right)^{\frac{\varepsilon-1}{\theta}} \Gamma \left(1 - \frac{\varepsilon-1}{\theta} \right) \\ &= \rho(\varepsilon - 1) s_j \bar{\Pi} \Gamma \left(1 - \frac{\varepsilon-1}{\theta} \right) \\ &= \rho(\varepsilon - 1) \bar{\pi}_j^{\frac{\theta}{\varepsilon-1}} \bar{\Pi}^{1-\frac{\theta}{\varepsilon-1}} \Gamma \left(1 - \frac{\varepsilon-1}{\theta} \right) \end{aligned}$$

where $s_j = \frac{\bar{\pi}_j^{\frac{\theta}{\varepsilon-1}}}{\sum_{j'} \bar{\pi}_{j'}^{\frac{\theta}{\varepsilon-1}}}$. Note that the Gamma function is not defined at zero, so we assume $1 - \frac{\varepsilon-1}{\theta} \neq 0$.

Let $A_1 = \bar{\Pi}^{1-\frac{\theta}{\varepsilon-1}} \Gamma \left(1 - \frac{\varepsilon-1}{\theta}\right)$. Given $\bar{\pi}_j = P^\varepsilon Y \frac{1}{\varepsilon} \left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1} \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}}\right)^{\varepsilon-1}$, we can write:

$$\begin{aligned} w_j n_j &= \rho(\varepsilon - 1) A_1 \bar{\pi}_j^{\frac{\theta}{\varepsilon-1}} \\ &= \rho(\varepsilon - 1) A_1 \left(P^\varepsilon Y \frac{1}{\varepsilon} \left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1} \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}}\right)^{\varepsilon-1} \right)^{\frac{\theta}{\varepsilon-1}} \end{aligned}$$

Let $A_2 = A_1 \left(P^\varepsilon Y \frac{1}{\varepsilon} \left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1} \right)^{\frac{\theta}{\varepsilon-1}} = \bar{\Pi}^{1-\frac{\theta}{\varepsilon-1}} \Gamma \left(1 - \frac{\varepsilon-1}{\theta}\right) \left(P^\varepsilon Y \frac{1}{\varepsilon} \left(\frac{\varepsilon-1}{\varepsilon}\right)^{\varepsilon-1} \right)^{\frac{\theta}{\varepsilon-1}}$.

$$w_j n_j = \rho(\varepsilon - 1) A_2 \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}} \right)^\theta$$

Therefore, aggregate demand for labor with skills of type j is:

$$n_j = \rho(\varepsilon - 1) A_2 \left(\frac{\bar{z}_j}{r_j^{1-\rho}} \right)^\theta w_j^{-\rho\theta-1}$$

Capital Demand. Similarly,

$$\begin{aligned} r_j k_j &= (1 - \rho)(\varepsilon - 1) \bar{\Pi}^{1-\frac{\theta}{\varepsilon-1}} \Gamma \left(1 - \frac{\varepsilon-1}{\theta}\right) \bar{\pi}_j^{\frac{\theta}{\varepsilon-1}} \\ &= (1 - \rho)(\varepsilon - 1) A_2 \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}} \right)^\theta \end{aligned}$$

Therefore, aggregate demand for capital of type j is:

$$k_j = (1 - \rho)(\varepsilon - 1) A_2 \left(\frac{\bar{z}_j}{w_j^\rho} \right)^\theta r_j^{-\theta(1-\rho)-1}$$

Finally, we can express the quantity A_2 only as a function of parameters and factor prices by using the labor market clearing condition.

B.2.4 Derivation of Aggregate Output

For the labor market to clear, we need:

$$\begin{aligned}
\bar{N} &= \sum_j n_j \\
&= \sum_j \rho(\varepsilon - 1) A_2 \left(\frac{\bar{z}_j}{r_j^{1-\rho}} \right)^\theta w_j^{-\rho\theta-1} \\
&= \rho(\varepsilon - 1) A_2 \sum_j \left(\frac{\bar{z}_j}{r_j^{1-\rho}} \right)^\theta w_j^{-\rho\theta-1} \\
&= \rho P^\varepsilon Y \Gamma \left(1 - \frac{\varepsilon - 1}{\theta} \right) \frac{1}{\varepsilon} \left(\frac{\varepsilon - 1}{\varepsilon} \right)^{\varepsilon-1} \left[\sum_j \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}} \right)^\theta \right]^{1-\frac{\theta}{\varepsilon-1}} \\
&\quad \times \sum_j \left(\frac{\bar{z}_j}{r_j^{1-\rho}} \right)^\theta w_j^{-\rho\theta-1}
\end{aligned}$$

since $n_j = \rho(\varepsilon - 1) A_2 \left(\frac{\bar{z}_j}{r_j^{1-\rho}} \right)^\theta w_j^{-\rho\theta-1}$ and

$$A_2 \left(Y, (\bar{z}, w, r)_j \right) = P^\varepsilon Y \Gamma \left(1 - \frac{\varepsilon - 1}{\theta} \right) \frac{1}{\varepsilon} \left(\frac{\varepsilon - 1}{\varepsilon} \right)^{\varepsilon-1} \left[\sum_j \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}} \right)^\theta \right]^{1-\frac{\theta}{\varepsilon-1}}$$

.

Therefore, we can express Y as:

$$Y = \bar{N} \rho^{-1} P^{-\varepsilon} \left(\Gamma \left(1 - \frac{\varepsilon - 1}{\theta} \right) \right)^{-1} \left(\frac{1}{\varepsilon} \left(\frac{\varepsilon - 1}{\varepsilon} \right)^{\varepsilon-1} \right)^{-1} \left[\sum_j \left(\frac{\bar{z}_j}{w_j^\rho r_j^{1-\rho}} \right)^\theta \right]^{-1+\frac{\theta}{\varepsilon-1}} \left(\sum_j \left(\frac{\bar{z}_j}{r_j^{1-\rho}} \right)^\theta w_j^{-\rho\theta-1} \right)^{-1}$$

Using this, we can express the input demand aggregate as a function of parameters and factor prices only as in the lemma: $A = A(\{\bar{z}, w, r\}_j) = \dots$

B.3 Solution of Tech Provider Problem

B.3.1 Proof of Proposition 1

Proof. Using the chain rule, the elasticity of demand when wages are internalized is:

$$\frac{dk_j}{dr_j} \frac{r_j}{k_j} = \frac{\partial k_j}{\partial r_j} \frac{r_j}{k_j} + \frac{\partial k_j}{\partial w_j} \frac{w_j}{k_j} \cdot \frac{dw_j}{dr_j} \frac{r_j}{w_j}.$$

To ease notation, we define the following partial (e) and total (ε) elasticities:

$$e_{kr} = -\frac{\partial k_j}{\partial r_j} \frac{r_j}{k_j} \quad e_{kw} = -\frac{\partial k_j}{\partial w_j} \frac{w_j}{k_j} \quad \varepsilon_{wr} = -\frac{dw_j}{dr_j} \cdot \frac{r_j}{w_j}$$

Then

$$\varepsilon_{kr} = e_{kr} - e_{kw} \cdot \varepsilon_{wr}.$$

Notice that under the model assumptions, the partial elasticity of capital demand is constant and does not depend on whether wages are fixed, i.e. $\varepsilon_{kr}^{LM} = e_{kr}^{LM} = e_{kr}^{MC}$. Then,

$$\varepsilon_{kr} < \varepsilon_{kr} \iff e_{kr} \cdot \varepsilon_{wr} > 0$$

Under the model assumptions $e_{kr} > 0$ and by the envelope theorem, as $\frac{d\tau_j}{dr_j} = 0$, $\varepsilon_{wr} = e_{wr} > 0$, therefore the result follows. □

Corollary 2. *The effective elasticity of demand faced by capital providers that internalize the specific-labor margin:*

(i) *increases as θ , the Frechet substitutability, increases;*

$$\frac{\partial \varepsilon^{eff}}{\partial \theta} = (1 - \rho) \frac{(1 + \eta)^2}{(\eta + \rho\theta + 2)^2} > 0$$

(ii) *falls in share of labor in production, ρ*

$$\frac{\partial \varepsilon^{eff}}{\partial \rho} = -\theta \frac{(2 + \eta)(\eta + \theta + 2)}{(\eta + \rho\theta + 2)^2} < 0$$

(iii) *increases in the worker heterogeneity parameter η*

$$\frac{\partial \varepsilon^{eff}}{\partial \eta} = \frac{\rho(1 - \rho)\theta^2}{(\eta + \rho\theta + 2)^2} > 0$$

Therefore capital providers' markup decreases in θ and η and increase in ρ .

B.3.2 Proof of Proposition 2

Proof. (Investment in technology-specific knowledge and productivity cutoff)

1. Let the gross profits of capital provider $j \in \mathcal{J}$ be

$$V_j(\varphi_j) = (1 - \mu^{-1}) \frac{(1 - \rho)}{\rho} \tilde{\Phi}^{-1} \left(\frac{w_j^\rho r_j^{1-\rho}}{\bar{z}_j} \right)^{-\theta} = (1 - \mu^{-1}) \frac{(1 - \rho)}{\rho} \frac{\Phi}{\tilde{\Phi}} s_j$$

when the specific labor margin is internalized and computed at the optimal choice of r_j . The perceived elasticity of profits to investment in technology-specific knowledge is decomposed as

$$\frac{\partial \ln V_j}{\partial \ln o_j} = -\theta(1 - \rho) \frac{\partial \ln r_j}{\partial o_j} - \rho\theta \frac{\partial \ln w_j}{\partial \ln o_j} - \rho\theta \frac{\partial \ln w_j}{\partial \ln r_j} \frac{\partial \ln r_j}{\partial \ln o_j} = -\rho\theta \frac{\partial \ln w_j}{\partial \ln o_j}$$

where the last equality follows since the rental rate on technology-specific capital is a constant markup on marginal cost and marginal cost does not vary with subsidy. Using the equilibrium wage equation,

$$\frac{\partial \ln w_j}{\partial \ln o_j} = -\frac{1 + \eta}{\eta + \rho\theta + 2}$$

and therefore at a fixed level of investment in technology-specific knowledge $\bar{o}$ we conclude

$$\frac{\partial \ln V_j}{\partial \ln o_j} = \frac{\rho\theta(1 + \eta)}{\eta + \rho\theta + 2}$$

It follows that

$$B_j(\bar{o}, \varphi_j) = \frac{\partial V_j}{\partial o_j} = \frac{\partial \ln V_j}{\partial \ln o_j} \frac{V_j(\varphi_j)}{\bar{o}}$$

$$\frac{dB_j(\bar{o}, \varphi_j)}{d\varphi_j} \propto \frac{dV_j(\varphi_j)}{d\varphi_j} > 0$$

i.e. the perceived benefit of investing in training is strictly monotone in productivity.

2. We have shown that the benefit of investing an amount $\bar{o}$ is monotonically increasing. For a fixed cost f of investing amount $\bar{o}$, the function of φ_j

$$B_j(\bar{o}, \varphi_j) - f$$

exhibits single crossing and will cross zero at most once from below as φ_j increases.

□

B.3.3 Microfoundation of fixed costs allocation.

B.4 Welfare Analysis

B.4.1 Proof of Lemma

To characterize the externalities created by technology-specific investment, consider a social planner that internalizes all general-equilibrium linkages between innovation, wages, and effective labor supply. The planner chooses $\{o_j, s_j, \varsigma_j, C\}$ to maximize aggregate welfare,

$$\max_{o_j \in \{0, \bar{o}\}, s_j, \varsigma_j, C} \ln(C) - \ln \mathbf{N}, \quad (36)$$

subject to the resource constraint

$$Y(n_j, o_j) = C + \phi^o \sum_j 1(o_j = \bar{o}), \quad (37)$$

and aggregate production and labor-supply conditions,

$$Y = \rho^{-1}(\bar{I})^{\frac{\varepsilon}{\varepsilon-1}} (\Phi)^{\frac{1+\theta}{\theta}} \left[\Gamma \left(1 - \frac{\varepsilon-1}{\theta} \right) \right]^{\frac{1}{\varepsilon-1}} \tilde{\Phi}^{-1} \sum_j \bar{N} \varsigma_j (1 - \tau_j), \quad (38)$$

$$\mathbf{N} = \bar{N} \sum_j \left(\frac{n_j}{\bar{N}(1 - \tau_j)} \right)^{\frac{2+\eta}{1+\eta}} = \bar{N} \sum_j (\varsigma_j)^{\frac{2+\eta}{1+\eta}}, \quad \tau_j = \bar{\tau}_j - o_j, \quad \sum_j \varsigma_j = \sum_j s_j = 1. \quad (39)$$

The planner invests in technology j whenever the marginal social benefit (MSB) exceeds the marginal social cost (MSC),

$$MSB_j(\varphi_j) \geq MSC_j. \quad (40)$$

In welfare units, this condition is equivalent to

$$\frac{d \ln Y}{d o_j} \frac{Y}{C} - \frac{d \ln \mathbf{N}}{d o_j} \geq \frac{\phi^o}{C}, \quad (41)$$

that is, the output gain net of disutility costs must exceed the resource cost of investment.

B.5 Decomposition of Social Marginal Benefit

Differentiating aggregate output (38) with respect to o_j yields

$$\frac{d \ln Y}{do_j} = \frac{1 + \theta}{\theta} \frac{d \ln \Phi}{do_j} - \frac{d \ln \tilde{\Phi}}{do_j} + \frac{d \ln \sum_m \bar{N} \varsigma_m (1 - \tau_m)}{do_j}. \quad (42)$$

When the price index is normalized ($P = 1$), Φ is constant, and thus the first term vanishes. The second term,

$$\underbrace{-\frac{d \ln \tilde{\Phi}}{do_j}}_{\text{(business-stealing externality)}},$$

captures how a change in o_j alters the cost index $\tilde{\Phi} = \Phi \sum_m s_m w_m^{-1}$ by reallocating expenditure shares across technologies with different wages. As s_j rises, $\tilde{\Phi}$ increases, raising aggregate costs for all other technologies—an externality the planner internalizes but individual firms ignore.

The third term,

$$\underbrace{\frac{\bar{N} \left[(1 - \tau_j) \frac{d \varsigma_j}{do_j} + \varsigma_j \frac{d(1 - \tau_j)}{do_j} \right]}{N^{eff}}}_{\text{(aggregate efficiency-unit externality)}},$$

captures the increase in effective labor supply that results both from the direct relaxation of training constraints $(1 - \tau_j)$ and from the induced reallocation of workers ς_j toward more productive technologies.

Changes in disutility are obtained from (39):

$$\frac{d \ln \mathbf{N}}{do_j} = \underbrace{\frac{\bar{N}}{\mathbf{N}} \frac{2 + \eta}{1 + \eta} \varsigma_j^{\frac{1}{1+\eta}} \frac{d \varsigma_j}{do_j} \frac{d \tau_j}{do_j}}_{\text{(labor-disutility externality)}}. \quad (43)$$

A reduction in τ_j reallocates workers and lowers average training disutility; this effect raises welfare but does not enter the private objective.

Combining these elements, the social marginal benefit of investment in technology j is

$$MSB_j(\varphi_j) = \left[\underbrace{\frac{\bar{N} \left[(1 - \tau_j) \frac{d \varsigma_j}{do_j} + \varsigma_j \frac{d(1 - \tau_j)}{do_j} \right]}{N^{eff}} \frac{Y}{C}}_{\text{efficiency-unit gain}} - \underbrace{\frac{1}{\mathbf{N}} \bar{N} \frac{2 + \eta}{1 + \eta} \varsigma_j^{\frac{1}{1+\eta}} \frac{d \varsigma_j}{do_j} \frac{d \tau_j}{do_j}}_{\text{disutility reduction}} \right] - \underbrace{\frac{d \ln \tilde{\Phi}}{do_j} \frac{Y}{C}}_{\text{business-stealing loss}}, \quad (44)$$

while the planner's perceived social marginal cost is

$$SMC_j = \frac{\phi^o}{C}. \quad (45)$$

Investment occurs whenever (44) exceeds (45).

B.5.1 Proof of Proposition 3.

B.6 Numerical solution

This appendix describes the numerical implementation of the equilibrium algorithm. All parameters satisfy the aggregation restriction $\theta > \varepsilon - 1$.

B.7 Step 0. Initialization

Given parameters p , draw the productivity distribution $\{\varphi_j\}_{j=1}^J$, sort technologies in ascending order, and initialize guesses:

$$q_L^{(0)}, \quad w_j^{(0)}, \quad \tau_j^{(0)} = \bar{\tau}, \quad o_j^{(0)} = 0.$$

The sorting by φ_j allows subsequent steps to proceed via monotone iteration.

B.8 Step 1 (innermost loop). Labor-Market Fixed Point

For given $\{r_j, \tau_j(o_j)\}$, equilibrium wages solve

$$w_j = g(r_j, \tau_j, \mathbf{W}; p), \quad (46)$$

$$\mathbf{W} = \left(\sum_j (1 - \tau_j)^{1+\eta} w_j^{1+\eta} \right)^{1/(1+\eta)}. \quad (47)$$

The fixed point is computed by iterative damping, $w^{(t+1)} = \frac{1}{2}w^{(t)} + \frac{1}{2}\tilde{w}(w^{(t)})$, until $\max_j |w^{(t+1)} - w^{(t)}| < \varepsilon_W$.

B.9 Step 2. Capital Providers' Problem

Given q_L each provider chooses (r_j, o_j) . The markup is constant under the assumptions, therefore the land price is sufficient to determine rental rates for all types of capital.

(a) Binary investment. Given the optimal value of the rental rate, this step determines the cutoff φ^* such that

$$\tilde{B}_j(\varphi_j) := V_j(o_j = \bar{o}) - V_j(o_j = 0) - \phi^o$$

crosses zero exactly once. Because $\tilde{B}_j$ has been shown to be strictly decreasing in φ_j , the algorithm proceeds monotonically from the most to the least productive firm, adding investors until $B_j < 0$. At each iteration, the benefit is evaluated at the most favorable combination of investment of competitors: I evaluate the benefit of investment of the most productive provider first, under the assumption that no other provider invests. If this evaluation results in investment for the most productive provider, I proceed to the next provider along the productivity rank and evaluate the perceived benefit of investment under the assumption that only the most productive firm has invested.

Given the minimum productivity cutoff, we determine investment strategies $\{o_j\}$ and all general-equilibrium aggregates $\mathbf{W}, A, \Phi, \tilde{\Phi}$ are recomputed using the updated $\tau_j(o_j)$ and r_j .

B.10 Step 3 (outermost loop). Land-Market Equilibrium

Aggregate land demand by capital providers is

$$L^d(q_L) = \sum_j \frac{k_j(r_j, w_j)}{\psi_j}.$$

The excess-demand function $E^L(q_L) = L^d(q_L) - \bar{L}$ determines the update

$$q_L^{(t+1)} = q_L^{(t)} \left(1 + \lambda \frac{E^L(q_L^{(t)})}{\bar{L}} \right).$$

Iterations continue until $|E^L(q_L)| < \varepsilon_L$.

C Data Cleaning and Sample Construction

C.1 US Job Posting Data

Following Acemoglu, Autor, Hazell, and Restrepo, I focus on establishments in “AI-using sectors” and exclude those in sectors likely to be producing AI technologies. Specifically, I drop establishments in the Information sector (NAICS 2-digit code 51), which includes software and data-related industries, and in the Professional and Business Services sector (NAICS 2-digit code 54), which

includes management consulting and other services likely to be involved in deploying AI technologies for clients. Notably, I exclude Google and its parent company Alphabet, whose open source strategy I analyze in the main case study.

To further refine the sample, I exclude staffing firms by removing observations with NAICS 4-digit codes 561311 (Employment Placement Agencies) or 561320 (Temporary Help Services) and retain only jobs that are not classified as internships. To ensure data quality, I drop observations with missing or generic information on industry, firm and occupation.

To reduce the dimensionality of the data, I analyze trends in job posting volume, the number of distinct employers, and industry composition over time for the group of job postings explicitly demanding TensorFlow as a specialized skill. I then restrict the analysis to a broad occupational category that is likely to be highly exposed to AI technologies: Computer and Mathematical Occupations (SOC 15). This category is chosen based on its high frequency in the universe of job postings that explicitly demand TensorFlow expertise from 2015 to 2024.

C.2 Salary Subsample

Summary statistics for the subsample of job postings with information about annual salaries are reported in table 4.

Table 4: Summary Statistics: Job Postings with Salary Information

Statistic	Value
Total observations with salary information	366,547
Share of sample	5.49%
Pre-TensorFlow period	155,073
Post-TensorFlow period	211,474
Mean posted salary	\$76,301
Median posted salary	\$70,210
Standard deviation	\$37,484
Mean salary, GCP postings	\$102,439
Mean salary, AWS postings	\$109,270
Mean salary, Azure postings	\$100,208
Number of GCP postings	863
Number of AWS postings	6,476
Number of Azure postings	4,586

Notes: This table reports summary statistics for job postings containing salary information in the sample period. The sample includes postings mentioning Google Cloud Platform (GCP), Amazon Web Services (AWS), and Microsoft Azure. The pre- and post-TensorFlow periods refer to observations before and after TensorFlow’s open-source release in November 2015.

C.3 Developer Surveys

Summary statistics for the adopted sample of US responders in 2024 are reported in Table 5.

Table 5: Summary Statistics: U.S. Developer Survey Respondents (2024)

Variable	All U.S. Respondents	U.S. Influencers
Sample size	7,157	3,867
TensorFlow users	821	410
TensorFlow usage rate (%)	11.47	10.60
GCP – Have (%)	20.94	23.33
GCP – Want (%)	15.20	16.01
GCP – Admire (%)	10.42	11.48
AWS – Have (%)	46.26	52.81
AWS – Want (%)	32.86	36.38
AWS – Admire (%)	27.74	31.76
Azure – Have (%)	28.64	33.10
Azure – Want (%)	20.67	22.94
Azure – Admire (%)	16.81	19.39
Some influence (%)	42.58	64.86
Great influence (%)	23.07	35.14

Notes: Data come from the 2024 *Stack Overflow Developer Survey*. The sample is restricted to U.S. respondents with non-missing technology data. “Influencers” are defined as respondents reporting some or great influence over technology purchase decisions. TF = TensorFlow. Percentages report platform adoption, interest (“Want”), and perceived reputation (“Admire”).

C.4 Calibration Strategy

Table 6: Wage Dispersion Measures Across Cloud Skills (2012–2019)

Measure	Symbol	Value
Coefficient of Variation	CV	0.1655
Standard Deviation of Log Wages	$\sigma_{\ln w}$	0.2095
Average Log Wage Variance	$\text{Var}(\ln w)$	0.0439
<i>Sample Information</i>		
Time Period		2012–2019
Cloud Providers		AWS, Azure, GCP
Sample		All job postings

Notes: This table reports wage dispersion measures computed from job posting data across three major cloud computing platforms (AWS, Azure, and GCP) for the period 2012–2019. The coefficient of variation (CV) measures relative wage dispersion (standard deviation divided by the mean). The variance and standard deviation of log wages capture scale-invariant wage dispersion, with $\sigma_{\ln w}$ approximating percentage variation around the geometric mean.


Table 7: **GCP Market Share Change After TensorFlow Open-Sourcing**

Measure	Detrended (SA)	Adjusted Version
Pre-TensorFlow average share	0.4909	
Post-TensorFlow average share	0.7420	
Absolute change (p.p.)	0.2511	
Relative change (%)	51.15	
<i>Sample Information</i>		
Event date	November 2015	
Time period	2012–2019	
Pre-event months	46	
Post-event months	50	

Notes: This table reports changes in Google Cloud Platform (GCP) market share before and after the open-sourcing of TensorFlow in November 2015. "Detrended (SA)" refers to seasonally adjusted and detrended market shares, while the "Adjusted Version" uses deviations from pre-event means. Market shares are computed from job postings mentioning cloud-computing skills.


D Additional Figures and Tables

In section 5 I restrict the sample to occupation SOC 15. While job postings for this occupation increase during the sample period, the pace of increase does not seem to respond significantly to the open sourcing event I study.




Notes: The figure plots monthly job postings for occupations classified under SOC 15 (Computer and Mathematical Occupations). The red line represents the pre-2015 trend, and the blue line represents the post-2015 trend. The slopes and R^2 statistics are reported for each subperiod. Although postings in this category increase steadily throughout the sample period, the rate of growth does not display a significant change following the open-sourcing event analyzed in Section 5. Shaded bands denote 95% confidence intervals around the fitted trends. Source: Author's calculations from Lightcast job postings data (2010–2019).

Figure 9: Seasonally Adjusted Job Postings by Cloud Provider, 2012–2019



Notes: Seasonally adjusted monthly counts of job postings mentioning the three major cloud platforms—Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)—between 2012 and 2019. The black dashed line marks November 2015, corresponding to the open-sourcing of TensorFlow. Data: US job postings.

Figure 10: Seasonally Adjusted Job Postings by Cloud Provider, 2012–2019 (Log Scale)



Notes: Log transformation of seasonally adjusted monthly counts of job postings mentioning the three major cloud platforms—Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP)—between 2012 and 2019. The black dashed line marks November 2015, corresponding to the open-sourcing of TensorFlow. The red lines represent the pre-2015 trends, and the blue lines represent the post-2015 trends. Data: US job postings.

D.1 Salary Effects

Table 8: Summary Statistics for Job Postings with Reported Salaries

Statistic	Value
Total observations with salary information	366,547
Share of total sample	5.49%
Pre-TensorFlow period	155,073
Post-TensorFlow period	211,474
Mean salary	\$76,301
Median salary	\$70,210
Standard deviation of salary	\$37,484
Mean salary (GCP postings)	\$102,439
Mean salary (AWS postings)	\$109,270
Mean salary (Azure postings)	\$100,208
Number of postings (GCP)	863
Number of postings (AWS)	6,476
Number of postings (Azure)	4,586

Notes: This table reports summary statistics for job postings with available salary information in the sample used for the empirical analysis. The data are divided into pre- and post-TensorFlow (TF) periods and cover postings mentioning major cloud platforms (Google Cloud Platform, AWS, and Microsoft Azure).

Dependent Variable:	Posted Annual Salary
Model:	(1)
<i>Variables</i>	
has_gcp	16,136.0** (6,391.8)
has_awsT	24,173.0*** (3,031.1)
has_azure	14,147.7 (9,980.4)
Post-TensorFlow	513.7 (1,921.2)
has_gcp $\times$ Post-TensorFlow	-8,052.2 (5,850.3)
has_aws $\times$ Post-TensorFlow	-4,819.6* (2,612.4)
has_azure $\times$ Post-TensorFlow	-8,529.0 (9,855.0)
<i>Fixed-effects</i>	
naics2	Yes
onet_2019_name	Yes
<i>Fit statistics</i>	
Observations	366,547
R ²	0.17285
Within R ²	0.00811

*Clustered (NAICS2 & O*NET 2019) standard-errors in parentheses*

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*