The Productivity Effects of Aging: Skill Composition and Technology Adoption

Stefano Cravero* and Ruben Piazzesi†
Job Market Paper

October 31, 2026

[Click here for the latest version]

Abstract

This paper studies how firms' technology adoption responds to the skill shortages created by population aging, and quantifies the cost for aggregate productivity. Using German administrative data, we document that younger workers complement new technologies: establishments with younger workforces adopt them at higher rates, and after adopting, hire disproportionately more young workers and raise their wages relative to older employees. We build on this evidence to quantify how aging affects aggregate productivity in an equilibrium model with endogenous technology adoption and frictional labor markets. Younger workers are more technologically proficient, while older workers have more experience accumulated on the job. Because frontier technologies are more productive but require technical skills, firms adopt them only when they can hire enough young workers to operate them. The calibrated model predicts that Germany's projected aging reduces technology adoption by 6.5% and aggregate productivity by 0.5%, enough to fully offset the gains from a more experienced workforce. Finally, we find that Germany's recent expansion of training subsidies, designed to address aging-induced skill shortages, raises long-run productivity by 0.8% and redistributes welfare from tech-proficient workers, whose comparative advantage erodes, to experienced workers who acquire technological skills through training.

Keywords: Population aging, technology adoption, human capital, labor productivity, directed search
JEL: E24, J11, J24, J64, O33

*University of Pennsylvania, Email: scravero@sas.upenn.edu
†University of Pennsylvania, Email: pruben@sas.upenn.edu.

We are grateful to Dirk Krueger, Jeremy Greenwood, and Lukas Nord for their invaluable guidance and constant support. We also thank Alessandro Dovis, Joachim Hubmer, Guillermo Ordoñez, Salomé Baslandze, Carolina Bussotti, Eliana Sena, Alexander Sawyer, and participants at the University of Pennsylvania, the Federal Reserve Bank of Atlanta Summer PhD Workshop, and the RoME Alumni Conference at the Einaudi Institute for Economics and Finance. All errors are our own.

1 Introduction

The workforce in advanced economies is aging rapidly, raising the share of older workers in the labor market.[1] This shift is changing the mix of workers' skills available to firms, with implications for technology adoption. New technologies require workers with complementary, up-to-date skills, which younger workers are more likely to possess (Deming and Noray, 2020; Lipowski, 2024). Aging therefore generates a shortage of technologically proficient workers, which can constrain the ability of firms to operate at the technological frontier. This concern has received significant attention by policy makers: the OECD, for instance, warns that older workers' slower adaptation to technological change will dampen the productivity of the economy in the future (OECD, 2025). Despite this growing attention, there is limited evidence on how large these losses could be.

This paper quantifies the aggregate productivity cost of technological skill shortages caused by an aging population. We make two contributions. First, we document empirically that younger workers complement new technologies. Using German administrative data linked to surveys of technology adoption, we show that establishments with younger workforces adopt new technologies at higher rates, and that after adoption they hire disproportionately more young workers, whose wages rise relative to those of older colleagues. Second, we develop an equilibrium model of technology adoption to quantify the aggregate implications of this evidence. In the model, workers differ along two skill dimensions: younger workers are more technologically proficient, while older workers are more experienced, a dimension that captures the standard channel through which an older workforce raises productivity in theories of learning on the job (Ben-Porath, 1967; Vandenbroucke, 2021). Firms choose between operating a standard technology and a frontier technology, which is more productive but requires more tech-proficient workers to run efficiently. Calibrating the model to Germany for 2005–2015, we find that projected aging reduces the adoption of frontier technologies by 6.5% and that the resulting technological skill shortage lowers labor productivity by 0.5%, enough to fully offset the productivity gains from a more experienced workforce. The cost of scarcer technological skills is therefore as large as the benefit that standard theories of human capital would predict from aging.

Our empirical analysis serves two purposes. First, it provides direct evidence for the proposed mechanism, that younger workers are more tech-savvy and that frontier technologies rely more heavily on them. Second, it delivers estimates of how firms respond to technology adoption, which serve as targets in the calibration of the quantitative model. Our main data source is the Linked Personnel Panel (LPP), a biennial survey that combines two components: an establishment survey recording the adoption of digital technologies, and a survey of employees within the same establishments reporting how these technologies affect their work. We link both data sets to administrative employer–employee records, and treat the digital technologies observed in the survey as the empirical counterpart of the model's frontier technology over our sample period, the 2010s.[2]

[1] In the United States, for example, the share of the labor force aged 45 and older rose from 39.3% in 2005 to 42.4% in 2025 and is projected to reach 44.1% by 2035 (BLS). In the European Union, the share of employment accounted for by workers aged 55–64 rose from 13.1% in 2010 to 19.8% in 2023 (Eurofound, 2025).

[2] Using German linked data for 2011–2016, Genz et al. (2021) similarly classify firms adopting connected digital tech-

We document three facts that support the key features of our mechanism. (i) Younger workers are more likely to perform ICT-related tasks daily and to report productivity gains after their establishment introduces a new technology, suggesting that they are more tech-savvy. (ii) Establishments with younger workforces adopt new technologies at higher rates, consistent with frontier technologies being more tech-skill intensive. (iii) We estimate how establishments' employment, age composition, and wages respond to adoption, using a stacked matched difference-in-differences design that compares adopters with non-adopting establishments with similar pre-adoption characteristics.[3] Within three years of adoption, employment at adopting establishments rises by 10% relative to matched controls. This expansion is concentrated among young workers, whose employment share rises by 2.75 percentage points, and the wages of young incumbents increase by 2% with no observed gain for older incumbents. We use these responses as calibration targets, which identify how strongly aging discourages adoption in the model.

To quantify the impact of population aging on aggregate labor productivity in general equilibrium, we develop a model of the labor market with two-sided heterogeneity. Workers belong to overlapping cohorts and face a constant retirement hazard. Their skills vary along two dimensions, technological skills and experience, which evolve systematically over the life cycle. New entrants have no experience, and a fraction of them possess high technological skills. While employed, they acquire experience at an exogenous rate and retain it for the rest of their careers. Technological skills instead depreciate: they become obsolete at an exogenous rate but can be rebuilt through training, which firms choose endogenously.[4] As a result, experience rises with age, while technological skills tend to decline. Firms combine these skills using either a standard technology or a frontier technology. The frontier technology is more productive but relies more heavily on technological skills and is costly to adopt. Each firm is characterized by its workers' composition, its technology, and an idiosyncratic productivity that evolves exogenously.

Workers and firms meet in a frictional labor market: following Schaal (2017) and Pijoan-Mas and Roldan-Blanco (2024), firms employ multiple workers and recruit unemployed workers through directed search, posting and committing to long-term contracts.[5] Search frictions make a firm's workforce a state variable: because workers cannot be hired instantly, adoption is a forward-looking decision that depends on the skills of both current employees and potential hires, consistent with the empirical evidence. Wages are set as part of the long-term contracts firms commit to, and they adjust when a firm upgrades to the frontier technology: adoption raises the value of tech-proficient workers to the firm, and with it their wages relative to those of other workers. Together, these features give the model direct counterparts to the employment, composition, and wage responses estimated in the data, so we can match them in the calibration.

We solve for a balanced growth path along which the labor force grows at a constant rate, and

nologies as being at the technological frontier.

[3] Identification of the responses relies on the assumption that, absent adoption, outcomes for treated and control firms would have evolved in parallel – supported by the absence of pre trends in our regressions.

[4] Modeling training matters for our question because aging changes firms' incentives to invest in skills (Angelini, 2023; Chari and Hopenhayn, 1991).

[5] Directed search keeps the model tractable because the equilibrium is block recursive under this assumption (Menzio and Shi, 2011).

calibrate it to the German labor market over 2005–2015. Two sets of parameters are key to quantifying the opposing forces that govern the productivity impact of aging. The first comprises the parameters of the frontier technology and the comparative advantage of tech-proficient workers at operating them, which are disciplined by the responses of establishments' employment, age composition, and relative wages to digital-technology adoption. The increase in employment informs how much more productive the frontier technology is. The shift toward young workers and their relative wage gains reveal how heavily the technology relies on technological skills and how strong young workers' comparative advantage in these skills is.[6] The second set comprises the parameters governing the experience premium and skill transitions, which are disciplined by the dynamics of wage growth and wage dispersion over the early careers of a cohort of labor market entrants. We complement these targets with cross-sectional rates of digital-technology use and employee training, which discipline adoption and training costs, and with aggregate flows into and out of unemployment, which inform the labor market parameters of the model.

Our main quantitative exercise models population aging as a decline in the inflow of young workers relative to incumbents, and uses the calibrated model to compare two balanced growth paths: a baseline with German demographics from 2005–2015, and a counterfactual with the demographic structure projected for the 2050s by the German Federal Statistical Office (Destatis). Because entrants and incumbents differ systematically in their skills, aging shifts the skill mix of the workforce toward experience: the share of experienced workers rises by 2.8%, while the share of tech-proficient workers falls by 1.7%. The resulting scarcity of technological skills makes the frontier technology less profitable, reducing the share of frontier firms by 6.5%.

We propose an exact decomposition to isolate how these changes affect productivity. The technological skill shortage alone lowers labor productivity by 0.5%, both directly, by reducing the stock of technological skills (-0.2%), and indirectly, through forgone technology adoption (-0.3%). This losses are large enough to completely offset the gains from a more experienced workforce ($+0.48\%$) predicted by standard theories of human capital, leaving aggregate labor productivity roughly unchanged ($+0.09\%$) in response to an aging population. A model without the interaction between technological skills and adoption would therefore erroneously predict that aging raises productivity significantly. Because the employment rate rises by 1.04% while productivity is essentially unchanged, total consumption increases by about 0.6%. These employment gains, however, accrue entirely to experienced workers, as the expansion of experience-intensive standard firms raises the demand for their skills, while the employment rate of inexperienced workers falls.

Motivated by these results, in the last part of the paper, we use the model to evaluate training subsidies, a widespread policy response to aging-induced shortages of technological skills.[7] Specifically, we study Germany's large expansion of training subsidies after 2019, which targeted employed workers across all ages, occupations, and regions. We calibrate the subsidy rate to match the observed take-up of the policy, and compare the transition from the baseline equilibrium to the long-run equilibrium with 2050s demographics, with and without the policy. The subsidy attenuates the decline in

[6] Relative wages are informative because they proxy for the age gradient in the productivity gains from adoption.

[7] Recent examples include Germany's *Qualifizierungschancengesetz*, Italy's *Fondo Nuove Competenze*, and Sweden's *omställningsstudiestöd* (Bundesregierung, 2018; Ministero del Lavoro e delle Politiche Sociali, 2024; CSN, 2026).

frontier adoption and raises long-run labor productivity by 0.8%.

The policy has meaningful redistributive effects among incumbent workers, generating winners and losers. Since subsidies are financed by taxes on all workers, they transfer resources from tech-proficient workers, whose comparative advantage erodes, to experienced workers, who acquire technological skills through easier access to training. Training subsidies may therefore be politically appealing: experienced workers are older on average and make up the majority of the labor force, so a policymaker attentive to their interests has a natural reason to adopt them.

Outline The rest of the paper is organized as follows: Section 1.1 discusses related literature; Section 2 discusses the data used and the empirical results; Section 3 presents the quantitative model and the main mechanisms; Section 4 outlines the calibration and quantitative properties of the model; Section 5 studies the productivity implications of aging; Section 6 evaluates German's tranining subsidies; Section 7 concludes.

1.1 Related Literature

Our paper contributes to a growing literature on demographic change, firm dynamics, and aggregate productivity. Most recent macroeconomic work on aging focuses on its effect on the size of the labor force. In particular, previous work has shown that slower labor-force growth reduces firm entry and helps explain the secular decline in U.S. business dynamism (Hopenhayn et al., 2022; Karahan et al., 2024). By shifting the firm distribution toward older incumbents, lower entry can also depress TFP growth when startups invest more in R&D or drive creative destruction (Inokuma and Sanchez, 2023; Peters and Walsh, 2021).[8] With the exception of Engbom (2020), who shows that startup creation is lower in older labor markets even conditional on labor-force growth,[9] these papers model workers as homogeneous, so that aging matters only through the number of workers. In contrast, we focus on the age composition of the workforce and show that, through this channel, aging changes the mix of skills available to firms and, in turn, their technology adoption decisions.

Our paper is the first to link firms' technology adoption to a skill composition that shifts along two opposing margins at once: aging makes experience more abundant and technological skills scarcer, generating opposing forces for productivity. Existing work has emphasized each side of this trade-off in isolation. Vandenbroucke (2021) argues that the entry of baby boomers with little accumulated human capital contributed to the U.S. productivity slowdown, a mechanism akin to our experience channel.[10] On the other hand, Angelini (2023), the closest to our paper, embeds the technology-complementary skills of young workers in a model of technology adoption. Our paper differs in two main ways. First, aging in our model also raises productivity through experience, which allows us to weigh the two forces against each other. Second, we are able to discipline the strenght of these channels with detailed firm-level data on technology adoption and workforce composition, whereas Angelini (2023) relies on cross-country aggregate investment data.

[8] See also Jones (2022) on the consequences of long-run population decline in models of idea production, and Maestas et al. (2023) and Aksoy et al. (2019) for empirical evidence on aging and economic growth.

[9] This finding is consistent with Liang et al. (2018), who document lower entrepreneurship in older societies.

[10] This channel is consistent with standard theories of human capital accumulation (e.g., Ben-Porath, 1967).

Because skills evolve over the life cycle, our paper also relates to work showing that demographic change affects wages by shifting the relative supply of experience (Katz and Murphy, 1992; Jeong et al., 2015). In these papers, the demand for experience is fixed, so demographics affect its return only through supply. We add technological skills as a second skill whose supply also changes with demographics, and we make the demand for both skills endogenous through firms' technology choices.

By modeling an endogenous technology response to aging, our paper further connects to Acemoglu and Restrepo (2022), who show that aging is associated with greater adoption of automation and other labor-saving technologies, which mitigate the resulting labor shortages.[11] We instead focus on digital technologies that complement technological skills, which become scarcer as the workforce ages. We view the two channels as complementary. Recent work also argues that young inventors are more creative, so that aging slows the production of new ideas (Kalyani, 2024). We abstract from endogenous growth, but incorporating it would likely magnify the productivity losses from technological skill shortages.

On the model side, our framework embeds insights from the vintage-human-capital literature, in which workers' skills shape the diffusion of new technologies, into a frictional labor market with multi-worker firms. In vintage models, skills are technology specific and imperfectly transferable across vintages (Chari and Hopenhayn, 1991).[12] This assumption is natural for major technological discontinuities that fundamentally reorganize production, but less so for the software and digital tools that we study, which tech-proficient workers can learn easily while retaining the value of their accumulated experience (Gathmann and Schönberg, 2010). We therefore separate two dimensions of human capital that vintage models typically combine: experience, which remains productive across technologies, and technological skills, which govern workers' ability to learn and use digital tools. We embed this structure in a directed search model with multi-worker firms, decreasing returns to scale, and dynamic long-term contracts, in the spirit of Kaas and Kircher (2015), Schaal (2017), and Pijoan-Mas and Roldan-Blanco (2024). We extend this framework in two directions: workers are heterogeneous along two skill dimensions that evolve over their careers, and we characterize a balanced growth path along which the labor force grows over time.

Finally, we contribute to a growing empirical literature documenting a complementarity between young workers and new technologies. Exploiting exogenous variation in trainee supply generated by a German education reform, Lipowski (2024) shows that a reduction in young trainees lowers firms' technology investment. Using data on robot adoption in German manufacturing, Deng et al. (2024) documents relative employment gains for young workers after adoption, while Aubert et al. (2006) finds similar patterns among French firms. Evidence from career and wage dynamics points to the same comparative advantage. Deming and Noray (2020) shows that returns to experience are flatter in occupations with more rapidly changing skill requirements, consistent with faster skill obsolescence favoring recent entrants, and Barth et al. (2023) finds that the elasticity of workers' wages to firm software investment declines with age. We contribute to this literature in two ways. First, we

[11] See also Acemoglu and Restrepo (2017).

[12] See also Atkeson and Kehoe (2007), Kredler (2014) and Adão et al. (2024).

provide new evidence from linked surveys of German establishments and their employees, which allow us to observe not only firms' adoption of digital technologies but also how workers of different ages use them and the productivity gains they report. Second, we move from documenting this complementarity to quantifying its aggregate consequences: we use the joint responses of employment, workforce age composition, and relative wages around adoption to discipline a structural model, and show that the complementarity is strong enough to fully offset the productivity gains from an older, more experienced workforce.

2 Empirical Analysis

2.1 Data and Descriptive Statistics

Our main data source is the Linked Personnel Panel linked to administrative data of the IAB. The LPP is a biennial survey of human-resource practices that follows a representative sample of German establishments with at least 50 employees, together with a random sample of their employees. Its fourth and fifth waves, fielded in June 2018 and June 2020, collected detailed establishment-level information on the use of ICT and digital technologies. For each technology, establishments report whether it was introduced during the preceding two years, had been in use for longer, or is not in use. We define technology adoption as the introduction of a technology during the two years preceding the survey. The employee survey records the tasks workers perform on a daily basis, their exposure to workplace digitalization, and how they perceive the effects of newly introduced technologies. We link the LPP survey to administrative records, which provide each establishment's employment by age and education, and each employee's employment history, including daily wages, occupation, and demographic characteristics.

Our second data source is the Longitudinal Model of the Linked Employer–Employee Data (LIAB-LM). It builds on the IAB Establishment Panel (IAB-BP), an annual survey that is representative of the cross-section of German establishments and has a substantial panel dimension, and links the survey responses to administrative records on establishments' workforce composition. We use the LIAB-LM to compute cross-sectional moments for the calibration of the quantitative model. We also use the Digitalization Survey supplement to the 2017 wave of the IAB-BP, which asked establishments about their use of several digital technologies and ICT tools. These data complement the more detailed adoption information in the LPP-ADIAB with technology-use statistics that are representative of the full establishment population.[13]

Finally, the Sample of Integrated Labour Market Biographies (SIAB) provides employment records for a two-percent sample of German workers. We use it to compute moments of early-career wage growth and wage dispersion, which serve as calibration targets.[14]

Table A1 in the Appendix compares the establishments in the two LPP waves, roughly 750 per

[13]The LPP-ADIAB, by contrast, covers only establishments with at least 50 employees and is therefore not representative of the full establishment population.

[14]Information on the three datasets can be found at: Mackeben et al. (2023) for LPP-ADIAB, Panahian Fard et al. (2024) for LIAB-LM, Graf et al. (2023), for SIAB.

wave, with the population of German establishments as represented by the IAB-BP. Relative to the average German establishment, LPP establishments are larger in terms of both employment and value added, older, and more productive, as measured by value added per worker. They also employ a higher share of highly educated workers, and manufacturing is overrepresented among them.

Table A2 in the Appendix lists the ten technologies covered by the LPP, which we group in two categories according to their primary role in production. First, digital and ICT technologies (e.g., the internet of things, big-data analytics, artificial intelligence, and virtual reality), which support workers in processing information and performing their tasks, and a large literature finds them to be complementary to workers' skills (Bresnahan et al., 2002; Akerman et al., 2015). Second, automation technologies (e.g., robots and drones), that more likely displace labor from previous tasks (Acemoglu and Restrepo, 2020). We distinguish the two groups because our analysis focuses on the complementarity between technology and workers' skills, which is most relevant for digital technologies. Pooling the two groups could therefore conflate effects that operate through different channels and may go in opposite directions. Digital technologies are also more widespread: each is used by 15 to 20 percent of establishments, compared with fewer than 10 percent for automation technologies (Table A2).

Throughout the analysis, we define an adoption event as the introduction of a technology during the two years preceding the survey. Because the survey does not report the exact date, we date adoption to $t-2$ for establishments reporting it in wave t.[15] Technologies in use for more than two years are treated as existing technologies. Unless otherwise noted, we define young workers as those under 35 and old workers as those over 45.[16]

2.2 Cross-sectional Evidence

This section provides empirical support for the main mechanism of the paper: younger workers have more technological skills, which complement new technologies, so the age composition of the workforce affects establishments' incentives to adopt them. We document two facts. First, establishments with younger workforces are more likely to adopt new technologies. Second, within establishments and conditional on occupation, younger workers are more engaged in digital and ICT-related activities and more likely to report gains from adoption, whereas older workers are more likely to report that adoption replaced their tasks and required new skills or training.

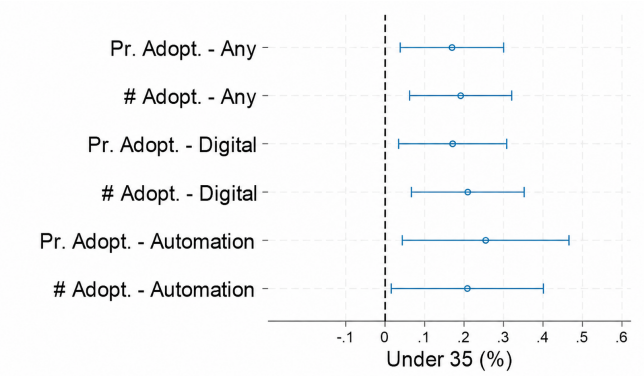
Establishments with a younger workforce adopt technologies at higher rates

We begin by documenting the relationship between establishments' workforce age composition and technology adoption. For each establishment j and survey wave t, we construct two measures of adoption. The first, $\mathbb{1}(\text{Adopt})_{jt}$, is an indicator equal to one if the establishment adopted at least one of the surveyed technologies during the preceding two years. The second, $N(\text{Adopt})_{jt}$, is the number of technologies adopted during the same period, conditional on adopting at least one.

[15] Adoption may have occurred in $t-2$, $t-1$, or t. Assigning it to $t-2$, the earliest possible date, may classify some pre-adoption observations as post-adoption, which biases the estimated short-run post-adoption coefficients toward zero.

[16] These cutoffs are common in the literature; see, e.g., Bianchi and Paradisi (2022).

Figure 1: Workforce Age Composition and Technology Adoption



Notes: The figure reports coefficients on the lagged share of employees younger than 35 in establishment-level regressions of technology adoption. Outcomes are an indicator for adopting at least one technology and the number adopted among adopters, separately for any, digital, and automation technologies. Adoption refers to technologies introduced within the two years preceding the LPP survey. Specifications control for pre-adoption establishment characteristics and include three-digit industry, state, survey-year, and previously used technology fixed effects. The bars represent 95% confidence intervals.

We regress these outcomes on the lagged age composition of the establishment's workforce. Our baseline specification is the linear probability model

$$\mathbb{1}(\text{Adopt})_{jt} = \alpha + \beta\,\text{Under 35 (\%)}_{j,t-3} + \Gamma X_{j,t-3} + \varepsilon_{jt},$$

where $\text{Under 35 (\%)}_{j,t-3}$ is the share of workers younger than 35 and $X_{j,t-3}$ is a vector of establishment-level controls. Because adoption is dated to $t-2$, controls are averaged over the three preceding years, $t-5$ to $t-3$.[17] We also include fixed effects for 3-digit industry, state, year, and the number of technologies already in use. We estimate the same specification with $N(\text{Adopt})_{jt}$ as the outcome. For both outcomes, we first pool all technologies and then distinguish digital from automation technologies.

Figure 1 reports the estimates of β. Establishments with younger workforces are more likely to adopt new technologies. A one-standard-deviation increase in the share of workers under 35 is associated with a 20 percent higher probability of adoption relative to the sample mean and, among adopters, with roughly 0.2 additional technologies. The relationship holds for both digital and automation technologies.

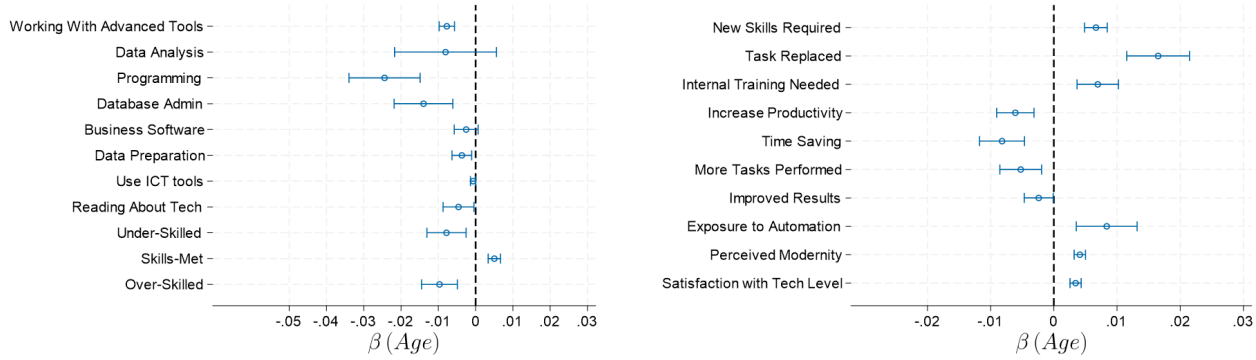
Table A4 in the Appendix shows that the results are robust to a richer set of controls and fixed effects, and Figure A2 shows that they are robust to alternative age thresholds.

Younger workers perform more digital tasks and report larger gains from adoption

The previous results establish an association between workforce age composition and technology adoption, but not why it arises. To shed light on this question, we use the employee survey of the

[17] Baseline controls are employment, establishment age, the average wage, the share of college graduates, the hiring rate, and ownership type.

Figure 2: Digital Activities and Effects of Technology Adoption by Worker Age



(a) Regular Activities

(b) Impact of tech adoption

Notes: Each point is the coefficient on worker age from a separate linear probability regression of an affirmative response to the listed LPP employee-survey question. Panel (a) concerns regular digital and ICT-related tasks; panel (b) concerns perceived consequences of newly introduced technologies and is restricted to employees of adopting establishments. The specifications include establishment and survey-year fixed effects and control for log wages, three-digit occupation, log occupation tenure, educational attainment, and task complexity. Comparisons are therefore among workers of different ages within establishments. The bars represent 95% confidence intervals.

LPP to examine whether younger workers engage more with digital technologies and whether they benefit differently when their establishment adopts new technologies.

We use two sets of questions. The first, asked to employees of all establishments, covers the digital and ICT-related activities they regularly perform, such as analyzing data, programming, using software, and keeping up with new technologies. The second, asked only to employees of adopting establishments, covers how the newly introduced technologies affected their jobs. Some of these questions capture adjustments to existing skills and tasks: whether previous tasks were replaced, new skills became necessary, or additional training was required. Others capture potential gains from adoption: higher productivity, time savings, a broader set of tasks, and higher output quality. For each question, we estimate

$$\mathbb{1}(Y_{ijt}) = \alpha + \beta\,\text{Age}_{it} + \Gamma X_{it} + \gamma_j + \delta_t + \varepsilon_{it},$$

where $\mathbb{1}(Y_{ijt})$ is the yes/no answer of worker i in establishment j and wave t, Age_{it} is the worker's age, X_{it} is a vector of worker-level controls,[18] and γ_j and δ_t are establishment and year fixed effects. Because of the establishment fixed effects, β is identified from comparisons of workers of different ages within the same establishment.

Figure 2 reports the estimates of β for both sets of questions. Panel 2a shows that younger workers are generally more likely to perform digital activities regularly. Panel 2b shows that the reported effects of adoption also differ by age. Older workers are more likely to report that the new technologies replaced tasks they previously performed and required new skills or employer-provided

[18] Baseline controls are the log wage, 3-digit occupation, log occupation tenure, educational attainment, and task complexity.

training. Younger workers are more likely to report productivity gains, time savings, and higher output quality. To interpret the magnitudes, consider the coefficient on task replacement in the second row of Panel 2b: a worker who is ten years older is approximately 15 percent more likely to report that the new technologies replaced some of their tasks.

Taken together, the cross-sectional results show that establishments with younger workforces are more likely to adopt new technologies and that, within establishments, younger workers use digital technologies more intensively and are more likely to report gains from adoption. Older workers, by contrast, are more likely to report that new technologies replaced their tasks and required new skills or training. These patterns strongly suggest that returns to adoption depend on the skills of the workforce: younger workers are better equipped to use new technologies, whereas adoption requires older workers to adjust their skills and tasks.

2.3 Event Studies around Technology Adoption

The evidence so far is cross-sectional. We now exploit variation in the timing of technology adoption across establishments to provide further evidence on the mechanism. Using an event-study design, we compare outcomes of adopting establishments before and after adoption with those of a matched sample of non-adopting establishments. The estimated responses to adoption will also serve as key targets to discipline the key parameters of the model, as detailed in Section 4.

Empirical strategy For each adopting establishment, we select one or more control establishments from the same survey wave that report no recent technology adoption and have similar pre-adoption characteristics. We restrict both adopters and potential controls to establishments that report at most one technology in use for more than two years, which limits contamination from earlier adoption events. We match establishments on two-digit industry, bins of employment, and bins of the share of workers under 35, all measured in the year before the actual or potential adoption year. We match on size and workforce age composition because these are the main determinants of the incentive to adopt in the model developed in Section 3. Each combination of these characteristics and the adoption year defines a *matching cell*. A cell may contain several adopters and controls, and controls are assigned the adoption year of the adopters in their cell.[19] Finally, we retain only establishments observed continuously for at least three years before and after the assigned adoption year, which yields a balanced panel in event time.

In the baseline analysis, we consider adoption events of digital technologies, since our mechanism relies on the complementarity between technology and workers' skills. Results for automation technologies are reported in Appendix A.4. Table A5 in the Appendix reports balance statistics for the matched sample and the number of adoption events.

Using the matched sample and the matching weights, we estimate

$$y_{jt} = \alpha_j + \gamma_t + \sum_{k \neq -1} \theta_k D^k_{jt} + \sum_{k \neq -1} \beta_k \left(D^k_{jt} \times \text{Tech Adoption}_f \right) + \varepsilon_{jt}, \tag{1}$$

[19] Within each matching cell, controls are weighted so that adopters and controls receive the same total weight.

where y_{jt} is the outcome of establishment j in year t, and α_j and γ_t are establishment and calendar-year fixed effects. Let T_j denote the assigned event year: the adoption year for adopters, and that of the adopters in the same matching cell for controls. The indicator $D_{jt}^k = \mathbb{1}\{t - T_j = k\}$ denotes event time k, and Tech Adoption$_j$ equals one for adopters and zero for controls. Because the survey reports adoption within a two-year window, T_j is the earliest possible year in which the new technology had already been installed, that is one year before the survey wave. Event time $k = 0$ is therefore the first year in which adoption may have occurred, and $k = -1$, the year before, is the omitted reference period. The coefficients θ_k capture the event-time profile of the control group, and β_k the additional change among adopters relative to their matched controls. The pre-adoption coefficients, β_k for $k < -1$, test for differential pre-trends, and the post-adoption coefficients, β_k for $k \geq 0$, trace the response to adoption.[20] For outcomes measured by age group, young workers are those under 35 and old workers those over 45, as in the cross-sectional analysis.

Identification of the parameters of intest, β_k, relies on the standard difference-in-differences assumption that, absent adoption, outcomes of adopters and matched controls would have evolved in parallel. The absence of differential pre-trends in the estimates below supports this assumption. Two concerns remain. First, adoption may respond to aggregate or industry-level shocks, such as shifts in demand. Because adopters and controls are matched within the same industry and adoption year, such shocks should affect both groups similarly. Second, adoption may respond to idiosyncratic shocks to establishment productivity or demand, which matching on observables cannot rule out. We address this concern in the quantitative analysis. In the model, adoption decisions respond to idiosyncratic productivity shocks, and we compute the model counterparts of the event-study coefficients by applying the same design to simulated data. To the extent that selection operates through productivity shocks, the model-implied coefficients reflect the same selection as the empirical ones, so the calibration does not require interpreting the estimates as causal effects.

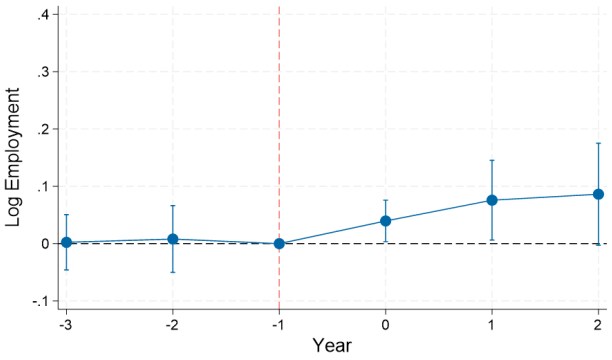
Employment and workforce composition Figures 7a–7c report the estimates for establishment employment and workforce age composition. Adopting establishments gradually expand relative to their matched controls: two years after adoption, their employment is approximately 10 percent higher (Figure 7a). This expansion is consistent with frontier technologies being more productive, which raises the optimal scale of the establishment.

The expansion is accompanied by a shift in workforce composition toward younger workers. Employment rises for both younger and older workers, but considerably more for the former (Figure 7b). As a result, two years after adoption, the share of workers under 35 is 2.65 percentage points higher and the share over 45 is 3 percentage points lower (Figure 7c). This shift corroborates the cross-sectional evidence in Section 2.2: younger workers are more proficient with new technologies, and frontier technologies require relatively more technologically proficient workers to operate.

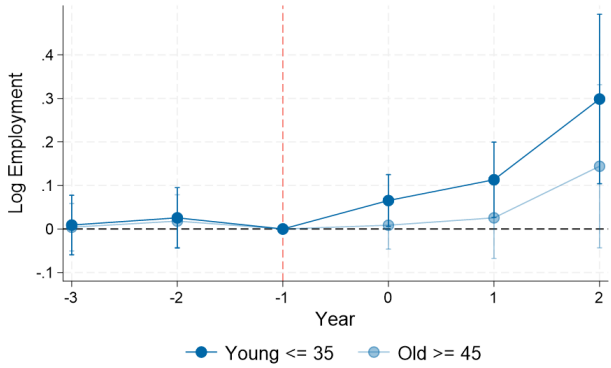
These responses will discipline two features of frontier technologies in the quantitative model: the employment response informs their productivity premium, and the shift in age composition informs their intensity in technological skills relative to established technologies.

[20] Standard errors are clustered at the establishment level.

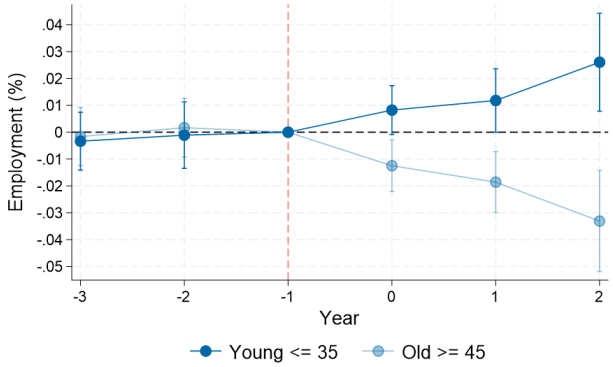
Figure 3: Employment and Age-Composition Responses to Technology Adoption



(a) Employment response to technology adoption



(b) Employment responses to technology adoption, by worker age



(c) Workforce age-composition response to technology adoption

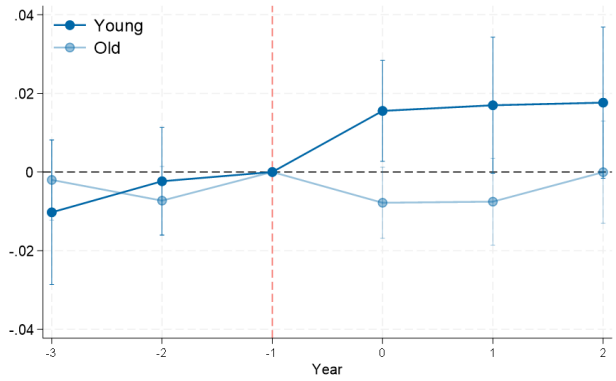
Notes: The panels report matched difference-in-differences event-study estimates for establishments adopting digital technologies, relative to non-adopting controls. Controls are selected within the same event year and matched on two-digit industry, pre-event employment, and the share of young employees. The omitted period is the year preceding the assigned adoption event ($k = -1$). Panels show log employment, log employment by age group, and changes in the shares of workers younger than 35 and older than 45. The event date is assigned using the LPP's two-year adoption window. Vertical intervals show the 95% confidence intervals, with standard errors clustered at the establishment level.

Wage responses Finally, we examine how the wage gains from adoption differ by age among incumbent workers, defined as those employed by the establishment before the event. Focusing on incumbents separates wage changes within existing employment relationships from changes in workforce composition due to subsequent hiring. We extend the establishment-level event study to a worker-level triple-difference specification:

$$\log w_{ijt} = \alpha_i + \gamma_{s(j),t} + \delta_{r(j),t} + \sum_{k \in \mathcal{K}} D_{jt}^k \left(\theta_k + \phi_k Y_i + \beta_k A_j + \tau_k A_j Y_i\right) + \varepsilon_{ijt}, \tag{2}$$

where w_{ijt} is the average daily wage of worker i in establishment j in year t, A_j equals one if worker i's establishment adopted a technology and zero if it is a matched control, and D_{jt}^k denotes event time as defined in Equation (1). The set $\mathcal{K} = \{-3, -2, 0, 1, 2\}$ excludes the year before the event, $k = -1$,

Figure 4: Wage Responses of Incumbent Workers to Adoption by Age



Notes: The figure shows the estimated response of log daily wages for incumbent workers at adopting establishments relative to matched non-adopting establishments, separately by baseline age group. Incumbents are workers employed before the adoption event; age-group membership is fixed in the pre-event year. Estimates come from the worker-level triple-difference specification with worker, three-digit-industry-by-year, and state-by-year fixed effects. The omitted event year is $k = -1$, vertical intervals show the 95% confidence intervals, and standard errors are clustered by establishment.

so all coefficients are measured relative to that year. The indicator Y_i equals one for workers under 35 in the year before the event and zero otherwise, so that age-group membership is fixed at baseline. The specification includes worker fixed effects α_i, 3-digit-industry-by-year fixed effects $\gamma_{s(j),t}$, and state-by-year fixed effects $\delta_{r(j),t}$. Standard errors are clustered at the establishment level.

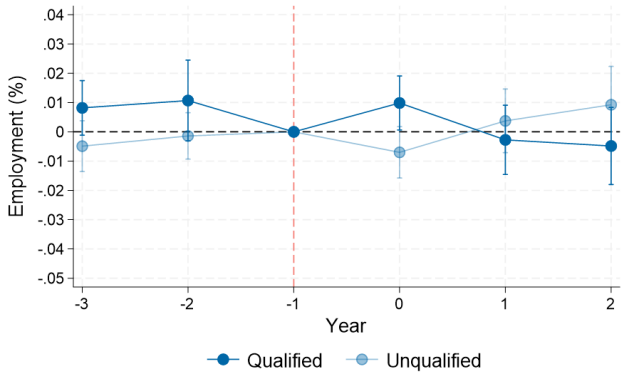
The coefficients θ_k and $\theta_k + \phi_k$ trace the wage profiles of older and younger workers in non-adopting establishments. The coefficients β_k capture the wage response to adoption of older incumbents, relative to older incumbents in non-adopting establishments, and $\beta_k + \tau_k$ the corresponding response of younger incumbents. The triple-difference coefficients τ_k measure the difference between the two responses.

Figure 4 reports the estimated responses for both age groups, β_k and $\beta_k + \tau_k$. Within two years of adoption, the wages of younger incumbents rise by approximately 2 percent, whereas the wages of older incumbents do not change significantly. The wage gains from adoption are thus concentrated among younger workers, consistent with the complementarity between their skills and frontier technologies. To the extent that wages reflect productivity, this differential response measures the age gradient in the productivity gains from adoption. In the quantitative model, the productivity gain from adoption increases with workers' technological skills, so matching the gap between the wage responses of younger and older incumbents will discipline the size of younger workers' comparative advantage in these skills.

Robustness analysis We next consider two alternative explanations for the employment patterns documented above that do not rely on a complementarity between new technologies and younger workers' skills.

First, the shift toward younger workers may reflect higher demand for college-educated workers rather than for technological skills. Because educational attainment has risen across cohorts in

Figure 5: Employment Responses to Technology Adoption by Worker Education Qualification

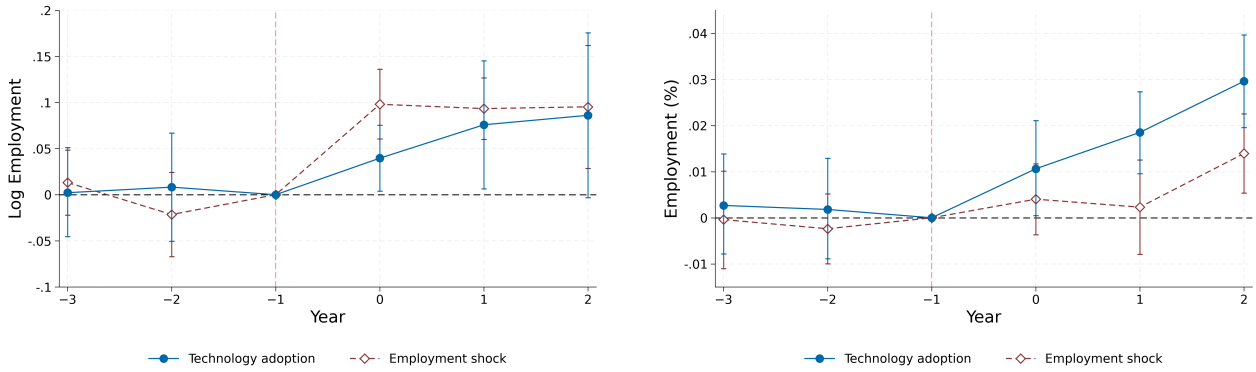


Notes: The figure reports employment responses around digital-technology adoption, separately for college-educated ("qualified") and non-college workers ("unqualified"), together with the response of the college share. The underlying sample, control establishments, event-time normalization ($k = -1$), and matched event-study design are the same as those used for the baseline establishment analysis. Vertical bars show the 95% confidence intervals, with clustering at the establishment level. This exercise assesses whether the change in workforce composition after adoption is accounted for by educational attainment rather than worker age.

Germany, younger workers are more likely to hold a college degree. If the shift instead reflected skill-biased technological change, adopting establishments should disproportionately expand their employment of college graduates. Figure 5 estimates the employment response separately for college and non-college workers. There is little evidence that employment of college graduates rises faster after adoption, and the college share of employment is essentially unchanged two years after the event. This suggests that the shift toward younger workers does not simply reflect higher demand for college-educated workers.

Second, establishments may hire disproportionately young workers whenever they expand, regardless of whether the expansion follows technology adoption. Younger workers have higher unemployment and job-mobility rates, and may therefore be easier to recruit. To assess this explanation, we estimate a placebo event study around large employment expansions not associated with technology adoption, constructed as described in Appendix A.4. Figures 6a and 6b compare the responses of employment and of the share of young workers to adoption events and to these expansion events. By construction, expansion events raise employment substantially, and two years after the event the increase is similar to that following adoption. The shift toward younger workers, however, is considerably smaller. The share of workers under 35 does not change significantly in the year of the event or the year after, and two years after the event it is about 1.4 percentage points higher, roughly half the increase following adoption (Figure 6b). Expansion alone therefore accounts for only part of the shift toward younger workers after adoption: the age-composition response is not simply a mechanical consequence of rapid growth or of the greater availability of young workers.

Figure 6: Employment and Age-Composition Responses to Large Employment Expansions



(a) **(b)**

Notes: The figure compares matched event-study responses around digital-technology adoption with responses around large employment expansions unaccompanied by adoption. Expansion events are defined using establishment employment-growth residuals after removing three-digit-industry and state trends; an event occurs when residual growth exceeds the 75th percentile. Both exercises use the matched-control design described in the text. Panel (a) reports log employment and panel (b) the share of workers under 35. Effects are relative to the preceding year, $k = -1$, and vertical bars show the 95% confidence intervals.

3 The Model

This section develops a quantitative model to measure how population aging affects aggregate labor productivity through changes in the skill composition of the workforce. Workers supply two skills, technological skills and experience, and firms combine them using either a standard technology or a frontier technology. Consistent with the empirical findings, the frontier technology is more productive but relies more heavily on technological skills. Two features make the model suitable for quantification. First, the economy is dynamic: workers accumulate experience while employed, their technological skills become obsolete over time but can be restored through training, and firms make forward-looking hiring, training, and adoption decisions. Obsolescence implies that younger workers have more up-to-date technological skills on average, in line with the empirical evidence. Second, the labor market is frictional: following Pijoan-Mas and Roldan-Blanco (2024), multi-worker firms recruit through directed search and commit to long-term contracts. Together, these features make a firm's workforce a state variable. As in the data, firms' adoption incentives therefore depend on the skill composition of their employees, and employment, workforce composition, and wages respond to adoption. As detailed in Section 4, matching the responses estimated in Section 2.3 will identify the technology parameters.

3.1 Environment

Time is continuous and infinite. The economy is populated by two sets of agents. Workers, whose total mass grows at the exogenous rate $g \geq 0$, and an endogenous mass of multi-worker firms. Both firms and workers are risk-neutral and share the common time discount rate $\rho > 0$.

Workers Workers differ only in the technological skills and experience they supply. At any point in time, a worker has either high technological skills (H), supplying $\Gamma_T \geq 1$ units, or low technological skills (L), supplying one unit, and is either experienced (e), supplying $\Gamma_E \geq 1$ units of experience, or inexperienced (i), supplying one unit. The set of worker types is therefore $\mathcal{G} \equiv \{Li, Le, Hi, He\}$. In addition, each worker is either employed or unemployed.

At each instant t, a mass $\mu_t \geq 0$ of new workers enters the labor market. All entrants start unemployed and inexperienced: a fraction $\omega \in [0, 1]$ has high technological skills (Hi) and the remainder low technological skills (Li). Employed inexperienced workers gain experience ($i \to e$) at Poisson rate $\lambda_E > 0$ and retain it for the rest of their careers. High technological skills become obsolete ($H \to L$) at rate $\lambda_T > 0$. Obsolescence is not permanent: employers can retrain low-tech workers at a rate $\lambda_{tr} \geq 0$ that they choose, subject to a convex cost $C_{tr}(\lambda_{tr})$, and successful training restores high technological skills ($L \to H$). Finally, workers retire at a constant hazard rate $\lambda_r > 0$.

Given entry and retirement, the labor force $\mathcal{L}_t$ evolves according to $\dot{\mathcal{L}}_t = \mu_t - \lambda_r \mathcal{L}_t$. We set the entry flow so that the labor force grows at a constant, possibly negative, rate g, which requires

$$\frac{\mu_t}{\mathcal{L}_t} = g + \lambda_r. \tag{3}$$

For a given retirement rate, the growth rate g thus maps one to one into the entry rate, the inflow of new workers relative to the size of the labor force.[21] This mapping is how population aging enters the model. Entrants are inexperienced, and a fraction ω of them has high technological skills, whereas incumbents have had time to accumulate experience and to lose technological skills through obsolescence. A lower g reduces the entry rate and shifts the workforce toward incumbents: the share of experienced workers rises and, unless firms offset obsolescence through training, the share of workers with high technological skills falls.

Firms Firms are multi-worker and heterogeneous along three dimensions, which constitute their physical state variables. The first is their workforce composition $\vec{n} = (n_{Li}, n_{Le}, n_{Hi}, n_{He}) \in \mathbb{N}^4$, the number of workers of each type. The second is the technology they operate, $\tau \in \{S, F\}$, which is either standard or frontier. The third is idiosyncratic productivity $z \in \mathcal{Z} \subset \mathbb{R}$, which follows an exogenous Markov process with transition function $\lambda_z(z' \mid z)$. A firm's physical state is therefore $s \equiv (\vec{n}, \tau, z)$.

A firm in state s produces output

$$y(\vec{n}, \tau, z) = z A_\tau \left(T^{\alpha_\tau} E^{1-\alpha_\tau} \right)^\eta, \tag{4}$$

where

$$T = \Gamma_T(n_{Hi} + n_{He}) + n_{Li} + n_{Le} \qquad \text{and} \qquad E = \Gamma_E(n_{Le} + n_{He}) + n_{Li} + n_{Hi}$$

are the efficiency units of technological skills and experience supplied by its workforce, and $\eta \in (0, 1)$ implies decreasing returns to scale. The two technologies differ along two margins. The first is total

[21] Non-negative entry requires $g \geq -\lambda_r$.

factor productivity A_τ: we normalize $A_S = 1$, so that A_F is the productivity premium of the frontier technology. The second is the intensity in technological skills α_τ: $\Delta\alpha \equiv \alpha_F - \alpha_S$ measures how much more intensive in technological skills the frontier technology is. The output gain from operating the frontier technology is

$$\frac{y(\vec{n}, F, z)}{y(\vec{n}, S, z)} = A_F \left(\frac{T}{E}\right)^{\eta\Delta\alpha}.$$

With $A_F > 1$ and $\Delta\alpha > 0$, this gain increases with the ratio of technological skills to experience in the workforce, and the frontier technology raises output only if this ratio is high enough. Therefore, the convenience of operating the frontier technology depends on the workforce composition of the firm, as well as on the ability to hire the right workers from unemployment.

We define technology adoption as a firm's upgrade from the standard to the frontier technology ($S \to F$). Adoption is stochastic. At each instant, a standard firm can invest a share $r \geq 0$ of its output in technology, at flow cost ry, and upgrades to the frontier at Poisson rate $q_{SF}(r)$, where q_{SF} is strictly increasing and concave. Frontier status is not an absorbing state. Because the frontier represents current best practice rather than a fixed technological vintage, frontier firms are subject to obsolescence and fall back to the standard technology ($F \to S$) at rate $q_{FS}(r)$.[22] As with adoption, frontier firms can invest a share r of their output to maintain their technology, which lowers the obsolescence rate. The rate $q_{FS}(r)$ is strictly decreasing and convex. In both cases, investment therefore has diminishing returns.

Finally, firms are destroyed at an exogenous rate $s^F > 0$, in which case they exit and all their workers become unemployed.

Given its state s, a firm chooses its recruiting, separation, training, and technology investment policies. Because the labor market is frictional, these decisions are implemented through the long-term contracts described in the next section.

Potential entrants The mass of firms is endogenous. In addition to workers and incumbent firms, there is an infinite pool of potential entrants. To open a business, potential entrants pay a one-time entry cost $\kappa > 0$. Upon entry, these firms draws idiosyncratic productivity z from an exogenous distribution $\Gamma^e(z)$ and start at the frontier with probability p_F, and at the standard technology otherwise.

Firms enter with one worker, of either of the four type, which they attract through the same search-and-matching technology as incumbent firms. After that, they immediately become incumbents in the corresponding state.

Finally, all firms in the economy are owned by a mutual fund which pools together all the profits and use them to pay entry costs for the creation of new firms.

We next describe the details of the search-and-matching technology, and how agents meet next.

[22] For example, a firm that adopted leading software or production systems in the past becomes outdated if it does not keep upgrading them.

3.2 Labor Market

Search and Matching The labor market is frictional and features directed search from unemployment. At each instant, firms (i) open one skill-specific vacancy for each type $i \in \mathcal{G}$, (ii) choose the promised utility W_i that is delivered by the posted contract, and (iii) decide the recruiting intensity rates $\{v_i\}_{i\in\mathcal{G}}$ for the vacancies subject to the convex flow cost function $C_v(v_i)$.

The labor market is segmented into submarkets indexed by worker type $i \in \mathcal{G}$ and by the long-term promised value W that the worker is expected to receive under the proposed contract. We assume that, at each instant, workers can search in at most one submarket, and each firm can post only one vacancy of each type. Denote by $\mathcal{V}_i(W)$ and $\mathcal{U}_i(W)$, respectively, the intensity-weighted measure of vacancies in submarket (i, W) and the measure of unemployed workers searching in the same market segment. At each instant, the measure of matches in submarket (i, W) is $\mathcal{M}\big(\mathcal{V}_i(W), \mathcal{U}_i(W)\big)$, where the matching function is Cobb–Douglas, with parameters common across all submarkets:

$$\mathcal{M}(\mathcal{V}, \mathcal{U}) = \mathcal{V}^{\gamma}\, \mathcal{U}^{1-\gamma}, \qquad \gamma \in (0, 1).$$

The matching function therefore has constant returns to scale. Let $\theta_i(W) \equiv \mathcal{V}_i(W)/\mathcal{U}_i(W)$ denote the (effective) tightness of submarket (i, W). By the law of large numbers, a firm with recruiting intensity v_i in submarket (i, W) fills vacancies at rate $v_i\, \eta\big(\theta_i(W)\big)$, where

$$\eta(\theta) \equiv \mathcal{M}(1, \theta^{-1}) = \theta^{\gamma-1}$$

is the job-filling rate per effective vacancy, which is common to all firms searching in the same submarket. Similarly, unemployed workers in submarket (i, W) find jobs at rate $f\big(\theta_i(W)\big) \equiv \mathcal{M}(\theta_i(W), 1) = \theta_i(W)^{\gamma}$.

Employment relationships are subject to exogenous separation shocks. An employed worker of type $i \in \mathcal{I}$ separates from the firm at Poisson rate $s_i^W > 0$, upon which the worker enters unemployment.

Dynamic Contracts. Firms and workers sign long-term, state-contingent contracts at the time of hiring. At each instant, a contract specifies the worker's flow wage, a continuation promise for each possible future state of the firm, and the firm's policies, which determine the rates at which these future states are reached. We assume that firms cannot discriminate among workers of the same type, so all workers of type $i \in \mathcal{G}$ employed by a firm receive the same contract. Contract terms therefore depend only on the worker's type and on the firm's physical state s and the vector $\vec{W} \equiv (W_i)_{i\in\mathcal{G}}$ of values promised to its incumbent workers.

The contract offered to workers of type i is

$$\mathcal{C}_i(s, \vec{W}) = \Big\{ w_i,\ \{W_i'(s')\}_{s'},\ \{v_j, \delta_j\}_{j\in\mathcal{G}},\ \{\lambda_{tr,j}\}_{j\in\mathcal{G}_L},\ r \Big\}.$$

It specifies the flow wage w_i, a continuation promise $W_i'(s')$ for each future physical state $s' = (\vec{n}', \tau', z')$, and the firm's policies: recruiting intensities v_j, firing rates δ_j, training intensities $\lambda_{tr,j}$

for low-tech workers, $j \in \mathcal{G}_L \equiv \{Li, Le\}$, and technology investment r. Recruiting, firing, and training are subject to convex flow costs $C_v(v_j)$, $C_\delta(\delta_j)$, and $C_{tr}(\lambda_{tr,j})$, while technology investment costs a share r of output. These policies determine the rates of the endogenous transitions, whereas productivity shocks, experience accumulation, obsolescence, retirement, and firm exit occur at exogenous rates. The contract is complete: it specifies a continuation promise for every transition of the firm, including those triggered by other workers, such as hiring or firing a worker of a different type. For this reason, the contract of a type-i worker includes the firm's policies toward all worker types.

Contracts are subject to two constraints. First, firms fully commit to their contracts and cannot renege on past promises. This is why the contract depends on the vector $\vec{W}$ of outstanding promises: it must deliver the promised value W_i, so the flow wage and the continuation promises, weighted by the rates of the corresponding transitions, must yield an expected discounted value equal to W_i (promise-keeping constraint). Because workers have linear utility, only the total value they receive matters to them, so the firm can shift compensation between the current wage and future promises as long as it delivers W_i. Second, workers cannot commit and may quit into unemployment at any instant. Every continuation promise must therefore be at least the value of unemployment, $W_i'(s') \geq U_i$ (participation constraint).

Because time is continuous, the probability that two events occur in the same instant is zero. We therefore index each future contingency, and the corresponding continuation promise $W_i'(s')$, by the firm's physical state after the event. Let e_i denote the unit vector corresponding to worker type $i \in \mathcal{G}$. Hiring a worker of type i moves the firm to $s_i^+ = (\vec{n} + e_i, \tau, z)$, and the separation of a type-i worker, through firing or retirement, moves it to $s_i^- = (\vec{n} - e_i, \tau, z)$.[23] A change in a worker's type from i to j, through experience accumulation, obsolescence, or training, moves the firm to $s_{i\to j} = (\vec{n} - e_i + e_j, \tau, z)$. A productivity shock moves the firm to $s_{z'} = (\vec{n}, \tau, z')$, with z' drawn from $\lambda_z(z' \mid z)$. A technology upgrade or downgrade moves the firm to $s_{\tau'} = (\vec{n}, \tau', z)$. Finally, the exogenous destruction shock ends the firm.

Timeline We conclude the description of the environment with the timing of events within an instant. A firm enters the instant with physical state $s = (\vec{n}, \tau, z)$.[24] First, the firm produces output $y(\vec{n}, \tau, z)$, chooses the contract terms for its incumbent workers, its firing and training intensities, and its technology investment, and pays wages and the associated costs. Second, it posts vacancies, each specifying the contract offered to new hires. Unemployed workers direct their search toward these vacancies, and matches form according to the matching function. Finally, at most one event occurs, either an endogenous transition at the rates implied by the contract or an exogenous shock, and the firm moves to the corresponding post-event state.

[23] A firm that separates from its last worker exits and rejoins the pool of potential entrants. The firing decision therefore also determines endogenous exit, even without a fixed operating cost.

[24] As we show in Section 3.3, the firm's problem can be written in terms of its physical state alone, without the vector $\vec{W}$ of values promised to its workers.

3.3 Equilibrium

We characterize the equilibrium in four steps. First, we solve the search problem of unemployed workers. Given a promised value W and the unemployment value U_i, workers' optimal search determines the tightness of each labor market segment, and hence the job-finding and job-filling rates. Second, taking these rates as given, we characterize the contracts offered by incumbent firms. Because utility is transferable and contracts are complete, the dynamic contracting problem reduces to the maximization of the joint surplus of the firm and its workers. Third, we derive the value of entry and impose free entry. Fourth, restricting attention to Markov contracts that do not discriminate among workers of the same type, we recover the wages and continuation values that implement the surplus-maximizing allocation. We then combine these conditions with the stationarity of the firm and worker distributions to define a balanced-growth-path equilibrium, and show that it is constrained efficient.

Unemployed Workers. Unemployed workers receive a flow payoff $b_i > 0$ and direct their search toward the submarket (i, W) that offers the highest expected value. Let $\theta_i(W)$ denote the tightness of submarket (i, W), $f(\theta)$ the job-finding rate, and U_i the value of unemployment for a worker of type $i \in \mathcal{G}$. Let $\mathcal{G}_H \equiv \{Hi, He\}$ denote the high-tech types and $\ell(i)$ the low-tech counterpart of a high-tech type, so that $\ell(Hi) = Li$ and $\ell(He) = Le$. The value of unemployment solves the Hamilton–Jacobi–Bellman (HJB) equation

$$(\rho + \lambda_r)U_i = b_i + \max_{W}\Big\{ f\big(\theta_i(W)\big) \max\{W - U_i, 0\} \Big\} + \lambda_T \big(U_{\ell(i)} - U_i\big)\mathbb{1}\{i \in \mathcal{G}_H\},$$

where the last term captures the obsolescence of technological skills during unemployment.

Workers accept a job only if it promises at least the value of unemployment. Because workers move freely across submarkets, every active submarket must deliver the same value U_i. For any promised value $W > U_i$ offered in equilibrium, tightness therefore satisfies

$$f\big(\theta_i(W)\big) = \frac{(\rho + \lambda_r)U_i - b_i - \lambda_T\big(U_{\ell(i)} - U_i\big)\mathbb{1}\{i \in \mathcal{G}_H\}}{W - U_i}. \tag{5}$$

The job-finding rate is thus decreasing in the promised value. Submarkets that offer higher values attract more workers per vacancy and are therefore less tight: workers find jobs at a lower rate, and firms fill vacancies at a higher rate.

Joint-Surplus Problem. We now turn to incumbent firms. Each firm takes as given the value of unemployment U_i of each worker type and the job-filling rate associated with each value promised to new hires. It chooses the dynamic contract $\mathcal{C}$ to maximize its value, subject to the promise-keeping constraint, because it is fully committed, and to the participation constraint, because workers can quit at any time.

As discussed in Section 3.1, the contract depends not only on the firm's physical state $s = (\vec{n}, \tau, z)$ but also on the promises $\vec{W} \equiv (W_i)_{i \in \mathcal{G}}$ outstanding to its workers. Solving the contracting problem

directly would therefore require tracking $\vec{W}$ as an additional state, which quickly makes the problem intractable. Because utility is linear and contracts are complete, however, the optimal contract maximizes the joint surplus of the firm and its workers, as in Schaal (2017) and Pijoan-Mas and Roldan-Blanco (2024). This result simplifies the problem considerably: outstanding promises are transfers between the firm and its workers, so the joint surplus does not depend on them.

Let $J(s, \vec{W})$ denote the value of a firm in physical state s with outstanding promises $\vec{W}$. The joint surplus of the firm and its workers is

$$\Sigma(s, \vec{W}) \equiv J(s, \vec{W}) + \sum_{i \in \mathcal{G}} n_i W_i.$$

Proposition 1 (Joint-surplus maximization)**.** *Under (i) linear utility and (ii) complete state-contingent contracts, the joint surplus does not depend on the outstanding promises: $\Sigma(s, \vec{W}) = \Sigma(s)$. Moreover, the firm's problem is equivalent to maximizing the joint surplus over the reduced contract space*

$$\mathcal{C}^{\Sigma} = \Big\{ \{v_i, \delta_i, W_i'(s_i^+)\}_{i \in \mathcal{G}},\ \{\lambda_{tr,i}\}_{i \in \mathcal{G}_L},\ r \Big\},$$

in which the only continuation promises are those offered to new hires.

Proof. See Appendix B.1. □

The intuition is as follows. Because the firm maximizes its value, the promise-keeping constraint binds: the firm delivers exactly the promised value W_i to each worker, through wages and continuation promises. Outstanding promises therefore reduce the firm's value one for one, $J(s, \vec{W}) = \Sigma(s) - \sum_{i \in \mathcal{G}} n_i W_i$, and raise the workers' values by the same amount. Since the firm and its workers have linear utility and the same discount rate, they value these promises identically, so the cost to the firm exactly offsets the gain to the workers. Outstanding promises therefore determine only how the surplus is divided, not its size.

Suppressing the dependence of policies on s, the joint surplus solves the HJB equation[25]

$$
\begin{aligned}
\rho\,\Sigma(s) = \max_{\mathcal{C}^\Sigma} \Bigg\{ & \underbrace{(1-r)\,y(s) - \sum_{i\in\mathcal{G}} \big[n_i C_\delta(\delta_i) + C_v(v_i)\big] - \sum_{i\in\mathcal{G}_L} n_i C_{tr}(\lambda_{tr,i})}_{\text{Net flow payoff}} \\
& + \underbrace{\sum_{i\in\mathcal{G}} v_i\,\eta\big(\theta_i(W_i'(s_i^+))\big)\big[\Sigma(s_i^+) - \Sigma(s) - W_i'(s_i^+)\big]}_{\text{Hiring}} \\
& + \underbrace{\sum_{i\in\mathcal{G}} n_i\big(\delta_i + s_i^W\big)\big[\Sigma(s_i^-) + U_i - \Sigma(s)\big]}_{\text{Firing and exogenous separations}} + \underbrace{\lambda_r \sum_{i\in\mathcal{G}} n_i\big[\Sigma(s_i^-) - \Sigma(s)\big]}_{\text{Retirement}} \\
& + \underbrace{\lambda_E \sum_{i\in\mathcal{G}_I} n_i\big[\Sigma(s_{i\to e(i)}) - \Sigma(s)\big]}_{\text{Experience accumulation}} + \underbrace{\lambda_T \sum_{i\in\mathcal{G}_H} n_i\big[\Sigma(s_{i\to \ell(i)}) - \Sigma(s)\big]}_{\text{Obsolescence}} \\
& + \underbrace{\sum_{i\in\mathcal{G}_L} n_i \lambda_{tr,i}\big[\Sigma(s_{i\to h(i)}) - \Sigma(s)\big]}_{\text{Training}} \\
& + \underbrace{\mathbb{1}\{\tau = S\}\,q_{SF}(r)\big[\Sigma(s_F) - \Sigma(s)\big]}_{\text{Upgrading}} + \underbrace{\mathbb{1}\{\tau = F\}\,q_{FS}(r)\big[\Sigma(s_S) - \Sigma(s)\big]}_{\text{Downgrading}} \\
& + \underbrace{s^F\big[\Sigma^{\text{exit}}(\vec{n}) - \Sigma(s)\big]}_{\text{Firm destruction}} + \underbrace{\sum_{z'\neq z} \Gamma_z(z' \mid z)\big[\Sigma(s_{z'}) - \Sigma(s)\big]}_{\text{Productivity shocks}} \Bigg\}
\end{aligned}
\tag{6}
$$

$$
\text{subject to} \quad W_i'(s_i^+) \geq U_i \quad \forall i \in \mathcal{G},
$$

where $\mathcal{G}_I \equiv \{Li, Hi\}$ denotes the inexperienced types, and $e(i)$ and $h(i)$ denote the experienced and high-tech counterparts of type i, defined analogously to $\ell(i)$. Upon exit, workers become unemployed and, by free entry, the firm's value is zero, so the joint value after exit is $\Sigma^{\text{exit}}(\vec{n}) \equiv \sum_{i\in\mathcal{G}} n_i U_i$.

The first line of (6) is the net flow payoff: output net of technology investment and of firing, recruiting, and training costs. The remaining terms are the expected changes in surplus from each possible event: hiring and separations, changes in workers' types, technology transitions, firm destruction, and productivity shocks. Wages do not appear because they are transfers between the firm and its workers. The only promise that remains is the value $W_i'(s_i^+)$ offered to new hires, which involves a trade-off: a higher promise raises the job-filling rate but transfers more of the surplus to the new worker. Because (6) depends only on the physical state s, the outstanding promises no longer need to be tracked.

Writing the problem in terms of the joint surplus has a further advantage: all policies can be expressed as functions of the surplus alone. Taking U_i and $\theta_i(\cdot)$ as given, each policy solves a static first-order condition that equates its marginal cost to its marginal benefit, which depends only on

[25] We set the value of retirement to zero, so the retirement term, unlike the separation term, does not include U_i. Because retirement is exogenous, a positive retirement value would not affect any optimal choice.

differences in surplus across states. The HJB equation (6) is therefore a fixed-point problem in Σ alone, which can be solved by value function iteration.

Entrants. A potential entrant pays the entry cost κ and then draws its technology and productivity: it starts with the frontier technology with probability p_F and with the standard technology otherwise, and draws productivity z from $\Gamma^e(z)$. It then posts one vacancy for each worker type and, for each vacancy, chooses the value W_i promised to the worker conditional on a match so as to maximize its expected value. If it hires a worker of type i, it becomes an incumbent in state $s_i^e \equiv (e_i, \tau, z)$, in which it employs that worker only. If it fails to hire, it does not create a firm: it loses the entry cost and returns to the pool of potential entrants, whose value is zero.

Because a firm that promises W to its only worker has value $J(s_i^e, W) = \Sigma(s_i^e) - W$, the value of an entrant with technology τ, net of the sunk entry cost, is

$$\tilde{J}_\tau^e = \sum_z \Gamma^e(z) \max_{\{W_i \geq U_i\}_{i \in \mathcal{G}}} \sum_{i \in \mathcal{G}} \eta\big(\theta_i(W_i)\big)\big[\Sigma(s_i^e) - W_i\big].$$

The ex-ante value of a potential entrant is

$$J^e = -\kappa + p_F \tilde{J}_F^e + (1 - p_F)\tilde{J}_S^e, \tag{7}$$

In equilibrium, free entry requires $J^e = 0$ at every instant.

Surplus division and wages. The joint-surplus problem determines the allocation, namely output and the transitions between physical states, but not how the surplus is divided between the firm and its workers. The only promise it pins down is the value offered to new hires. With a Cobb–Douglas matching function, it can be shown that the optimal value promised to a new hire of type i in state s is

$$W_i'(s_i^+) = \max\Big\{\gamma U_i + (1-\gamma)\big[\Sigma(s_i^+) - \Sigma(s)\big],\ U_i\Big\}, \qquad i \in \mathcal{G}. \tag{8}$$

The new hire thus receives her outside option U_i plus a share $1 - \gamma$ of the net surplus that the hire creates, $\Sigma(s_i^+) - \Sigma(s) - U_i$. The participation constraint, the reason there is a max operator, binds only when this net surplus is not positive, in which case the firm does not recruit. Because firms cannot discriminate among workers of the same type, incumbent type-i workers also receive the same value as the new hire in state s_i^+.

To recover the promises conditional on all other transitions, we restrict attention to Markov contracts, in which all contract terms depend only on the current physical state, not on the history that led to it.[26] The value promised to a type-i worker in state s' must then be the same regardless of the state s from which the firm arrived:

$$W_i'(s' \mid s) = W_i(s') \quad \text{for all } s.$$

[26]This procedure follows the method proposed by Pijoan-Mas and Roldan-Blanco (2024) for a similar environment as ours.

Otherwise, outstanding promises would reveal the firm's history, which would become an additional state variable. Since any state s with $n_i \geq 1$ can be reached by hiring a type-i worker from s_i^-, Markov consistency and (8) together determine the value of every type-i worker in every state:

$$W_i(s) = \max\Big\{\gamma U_i + (1-\gamma)\big[\Sigma(s) - \Sigma(s_i^-)\big],\ U_i\Big\}, \tag{9}$$

where $\Sigma(s_i^-) = 0$ for single-worker firms, consistent with the entrant's problem. The firm's value is the residual surplus, $J(s) = \Sigma(s) - \sum_{i\in\mathcal{G}} n_i W_i(s)$.

Finally, given the promised values, the promise-keeping constraint, which binds, pins down wages:

$$\begin{aligned}(\rho + \lambda_r) W_i(s) = w_i(s) &+ \sum_{s'} \lambda_i(s' \mid s)\big[W_i(s') - W_i(s)\big] + \sum_{j\neq i} \lambda_{i\to j}(s)\big[W_j(s_{i\to j}) - W_i(s)\big] \\ &+ \big(\delta_i(s) + s_i^W + s^F\big)\big[U_i - W_i(s)\big].\end{aligned} \tag{10}$$

The first sum in (10) covers events that leave the worker employed as type i: hiring, separations and type changes of other workers, technology transitions, and productivity shocks, where $\lambda_i(s' \mid s)$ is the rate at which such events move the firm to s'. The second sum covers changes in the worker's own type, through experience accumulation, obsolescence, or training, which occur at rate $\lambda_{i\to j}(s)$. The last term covers the worker's own separation, through firing, exogenous separation, or firm destruction. Retirement, after which the worker's value is zero, enters through the discount rate on the left-hand side. In section 4.4, we show that the quantitative model generates wage dynamics in line with the data.

Balanced Growth Path Equilibrium. So far, we have characterized individual behavior taking as given the unemployment values and labor market tightness. We now close the model by requiring these objects to be consistent with the aggregate distributions of workers and firms. Under (3), the labor force grows at the constant rate g. We therefore focus on a *balanced growth path* (BGP): an equilibrium in which value and policy functions are time-invariant and all aggregate quantities, expressed per member of the labor force, are constant. In levels, the measures of firms, employed workers, and unemployed workers all grow at rate g.

A BGP equilibrium consists of (a) the joint surplus $\Sigma(s)$; (b) the policy functions $\{v_i(s), \delta_i(s), W_i'(s_i^+)\}_{i\in\mathcal{G}}$, $\{\lambda_{tr,i}(s)\}_{i\in\mathcal{G}_L}$, and $r(s)$; (c) promised values $\{W_i(s)\}_{i\in\mathcal{G}}$ and wages $\{w_i(s)\}_{i\in\mathcal{G}}$; (d) market tightness $\theta_i(W)$ in each submarket (i, W); (e) unemployment values $\{U_i\}_{i\in\mathcal{G}}$; (f) per-capita unemployment $\{u_i\}_{i\in\mathcal{G}}$, with $u_i \equiv \mathcal{U}_{i,t}/\mathcal{L}_t$; and (g) the per-capita measure of firms $\phi \equiv \Phi_t/\mathcal{L}_t$, the per-capita entry flow $\phi^e \equiv \Phi_t^e/\mathcal{L}_t$, and a stationary distribution Λ over firm states, such that:

1. *Optimality.* Given $\{U_i\}_{i\in\mathcal{G}}$ and $\theta_i(\cdot)$, the joint surplus $\Sigma(s)$ solves (6), the policy functions attain the maximum in (6), and entrants' promises solve the entrant's problem.

2. *Promised values and wages.* Promised values $W_i(s)$ satisfy (9), and wages $w_i(s)$ satisfy the promise-keeping constraint (10).

3. *Labor market consistency.* Given $\{U_i\}_{i\in\mathcal{G}}$, tightness $\theta_i(W)$ satisfies (5) in every submarket. Moreover, for each type i, the unemployed workers searching across active submarkets add up to u_i:

$$\sum_W \frac{\mathcal{V}_i(W)}{\theta_i(W)} = u_i,$$

where $\mathcal{V}_i(W)$ is the per-capita measure of effective vacancies posted in submarket (i, W) by incumbents and entrants.

4. *Free entry.* The expected value of entry is zero: $J^e = 0$.

5. *Stationary firm distribution.* Given the optimal policies, the per-capita distribution of firms, $\phi\Lambda$, is invariant: for each state s, the per-capita flow of firms into s, from entry and from transitions out of other states, equals the flow out of s, through transitions to other states, firm destruction, and separation from the last worker, plus the dilution $g\phi\Lambda(s)$ due to labor force growth.

6. *Worker flow balance.* For each type i, per-capita unemployment u_i is constant: inflows from new entrants, separations, and obsolescence during unemployment equal outflows to employment, retirement, and obsolescence, plus the dilution gu_i. Per-capita employment of type i is $\phi \int n_i \, d\Lambda(s)$, and the labor force adds up: $\sum_{i\in\mathcal{G}} \left[u_i + \phi \int n_i \, d\Lambda(s)\right] = 1$.

We detail the numerical algorithm that solves for the equilibrium in Appendix B.2.

Efficiency Despite the presence of search frictions, the decentralized equilibrium is constraint efficient. The following proposition establishes that a social planner who faces the same matching technology, demographic and skill-transition processes, and firm-state dynamics chooses the same allocation as the one resulting in equilibrium in the decentralized economy.

Proposition 2 (Constrained efficiency)**.** *Suppose that the matching function has constant returns to scale, firms and workers write complete state-contingent long-term contracts, utility is transferable, and there are no taxes or other distortions. Then any BGP equilibrium is constrained efficient: its allocation maximizes the present discounted value of aggregate net output among all allocations that satisfy the same matching technology, demographic and skill-transition laws, and firm-state transition laws as the decentralized economy.*

Proof. See Appendix B.4. □

The result reflects two features of the environment. First, under directed search, firms internalize the congestion that their recruiting creates in each submarket: by (8), new hires receive a share $1-\gamma$ of the net surplus, which equals the elasticity of the matching function with respect to unemployment, so the hosios condition holds endogenously. Second, complete contracts and transferable utility ensure that separations, training, and technology investment maximize the joint surplus of each firm and its workers. As a result, the decentralized optimality conditions for recruiting, separations, training, technology investment, and entry coincide with those of the planner.

4 Calibration

This section describes the calibration of the quantitative model to the German economy. We draw on the IAB Establishment Panel and the Linked Personnel Panel (LPP), described in Section 2.1, and complement these datasets with the Sample of Integrated Labour Market Biographies (SIAB). The calibration combines estimates from the empirical analysis in Section 2, additional moments that we compute from these datasets, and published estimates from other studies of the German labor market. We use 2005–2015 as the baseline reference period for the calibration of labor market moments. The German labor market went through a deep process of reforms enacted between 2003 and 2005 under the Harzt reform, which substantially reshaped unemployment insurance and employment services (Krebs and Scheffel, 2013). Therefore, starting in 2005 focuses the calibration on the institutional framework introduced after the reforms. Moreover, given the paper's focus on the demographic composition of the labor force, we end the reference period in 2015 to limit the influence of the subsequent labor-market integration of the large refugee inflow to Germany in 2015–2016 (Brücker et al., 2016).[27]

4.1 Parametrization

We set the model period to one quarter. We approximate idiosyncratic productivity by a continuous-time Markov chain with $N_z = 6$ values, $\mathcal{Z} = \{z_1, \ldots, z_{N_z}\}$, and an $N_z \times N_z$ matrix of transition intensities, $\{\Gamma_z(z' \mid z)\}$. The productivity grid and transition intensities are obtained by discretizing an Ornstein–Uhlenbeck process for log productivity:

$$d \log z_t = -\rho_z \log z_t \, dt + \sigma_z \, dB_t,$$

where B_t is a standard Brownian motion, $\rho_z > 0$ governs the speed of mean reversion, and $\sigma_z > 0$ determines the volatility of productivity innovations.

For the matching technology, we adopt a Cobb–Douglas specification with matching efficiency normalized to one:

$$\mathcal{M}(\mathcal{V}, \mathcal{U}) \;=\; \mathcal{V}^{\gamma} \mathcal{U}^{1-\gamma}, \qquad \gamma \in (0,1),$$

where $\mathcal{V}$ denotes the intensity-weighted measure of vacancies and $\mathcal{U}$ the measure of unemployed workers in a given submarket. With market tightness defined as $\theta = \mathcal{V}/\mathcal{U}$, this specification implies a job-finding rate $f(\theta) = \theta^{\gamma}$ and a job-filling rate per unit of recruiting intensity $\eta(\theta) = \theta^{\gamma-1}$.

We specify quadratic flow costs for recruitment, firing, and training:

$$C_v(v) = \frac{\kappa_v}{2} v^2, \qquad C_\delta(\delta) = \frac{\kappa_\delta}{2} \delta^2, \qquad C_{tr}(\lambda_{tr}) = \frac{\kappa_{tr}}{2} \lambda_{tr}^2.$$

The parameters κ_v, κ_δ, and κ_{tr} govern the scale of the respective costs and are common across worker types. These specifications imply that the marginal costs increase linearly with the corre-

[27] Although 2015 includes the peak in refugee arrivals, integration into employment was gradual. (Brücker et al., 2016) report that 9% of refugees who arrived in 2015 were employed at the time of their 2016 survey.

sponding intensity.

For technology investment, we specify the upgrading and downgrading rates as

$$q_{SF}(r) = 1 - e^{-\psi_{SF} r}, \qquad q_{FS}(r) = e^{-\psi_{FS} r}, \qquad \psi_{SF}, \psi_{FS} > 0,$$

where $r \in [0, 1]$ is the share of current output devoted to technology investment. The upgrading rate is increasing and concave in investment, whereas the downgrading rate is decreasing and convex. Thus, investment helps standard firms reach the frontier and frontier firms retain their technological advantage, with diminishing marginal effects in both cases.[28]

4.2 Externally Set Parameters

We set eight parameters using estimates from the literature and empirical moments with a direct counterpart in the model. These are summarized in Table (1) below. We set the quarterly discount rate to $\rho = 0.125$ to imply an annual discount rate of 5% and set the matching elasticity to $\gamma = 0.5$, a conventional value in the search-and-matching literature (e.g. Petrongolo and Pissarides (2001)).

The productivity process and the degree of decreasing returns are disciplined by the estimates of Lochner and Schulz (2024), who estimate firm-level production functions and productivity dynamics using German administrative and establishment-survey data from the IAB. We use their estimated annual productivity autocorrelation of $\rho_z = 0.75$ and unconditional standard deviation of log TFP of $\sqrt{0.074} \simeq 0.2720$. To map these estimates into our production function, which abstracts from capital, we account for flexible capital adjustment. This transformation implies a stationary standard deviation of log productivity of approximately $\sigma_z = 0.3280$ and a returns-to-scale parameter of $\eta = 0.8595$. The numerical appendix (C.1) describes the capital residualization, the conversion to quarterly transition intensities, and the discretization of the productivity process.

We next set the demographic parameters to reproduce the average length of working life and the age composition of the labor force. Because the model abstracts from initial schooling decisions, we interpret entry as the beginning of working life after initial education and set the age of entry to 25. We choose the retirement rate to match a mean retirement age of 64, using evidence from the German Pension Insurance (2025).[29] With retirement arriving at a constant Poisson rate, an expected working life of 39 years implies $\lambda_r = \frac{1}{4(64-25)} \simeq 0.0064$.

Conditional on this retirement rate, the labor-force growth rate g determines the inflow of new workers relative to the existing labor force and, hence, its age composition along the balanced growth path. We therefore choose g to match the average share of young workers in the German labor force over 2005–2015. To be consistent with our empirical analysis, we define young workers as those younger than 35, and restrict the empirical sample to individuals aged 25 and older to align it with

[28] Finally, to keep the firm's state space computationally tractable, we impose type-specific employment bounds of

$$(N_{He}, N_{Hi}, N_{Le}, N_{Li}) = (10, 15, 25, 10).$$

These bounds allow firms to employ at most 60 workers in total, while retaining heterogeneity in workforce composition along both the technological-skill and experience dimensions.

[29] Official information can be found at Deutsche Rentenversicherung Bund (2025).

Table 1: Externally Set Parameters

Production Function			
Parameter		**Value**	**Target/Source**
ρ	Discount rate	0.0125	5% annual discount rate
γ	Matching elasticity	0.5000	Petrongolo and Pissarides (2001)
ρ_z	Mean-Reversion in Productivity	0.7500	Lochner and Schulz (2024)
σ_z	Productivity Dispersion	0.3280	Lochner and Schulz (2024)
η	DRS	0.8595	Lochner and Schulz (2024)
p_F	Share of frontier start-ups	0.332	Authors Calculations (IAB-BP 2017)
λ_r	Retirement Rate	0.0064	Average retirement age (2005-2015)
g	Growth rate LF	0.005	Average Young Workers (%), 2005-2015

Notes: The table lists parameters set externally before the simulated-method-of-moments calibration. The baseline corresponds to the German labor market in 2005–2015. The discount rate, retirement hazard, and labor-force growth rate are expressed in model-period (quarterly) units; the retirement hazard is chosen to reproduce a working life from age 25 to mean retirement age 64. The productivity-process parameters and decreasing-returns parameter reflect the capital-adjusted mapping described in the text and numerical appendix. The frontier-entry probability is measured from the IAB Establishment Panel digitalization supplement.

the model's entry age.[30] Using age-specific labor-force counts from the Destatis Microcensus, we compute the ratio of individuals aged 25–34 to all individuals aged 25 and older in the labor force.[31] This share averages 22.3% over the calibration period.

Finally, we set the probability that an entrant starts with the frontier technology to $p_F = 0.332$, the share of start-ups that use at least one digital technology in the Digitalization Survey supplement to the 2017 IAB Establishment Panel.[32] Table A3 in the Appendix reports usage rates for the three digital technologies covered by the survey, which range from 6 to 23 percent. Overall, 32 percent of establishments use at least one of them, a slightly lower share than among start-ups.[33] Consistent with the model, establishments that use digital technologies are larger, employ a higher share of young workers, and are more productive, as measured by value added per worker (Figure A1).

4.3 Internally Calibrated Parameters

4.3.1 Technology and Skills

We calibrate the parameters governing technology and skills to three sources of evidence: establishments' responses to technology adoption, the distribution of wages and transitions early in workers' careers, and establishment-level measures of training and technology use. Although all parameters are calibrated jointly, we discuss each in connection with the moments that are most informative about it.

Technology adoption and workforce composition. The event-study estimates of Section 2 are central to the calibration. The employment response to adoption is informative about the productivity

[30] Defining young workers ar under 35 y.o. is standard in the literature. See, for instance, Bianchi and Paradisi (2022).

[31] The official data from Destatis can be found at Federal Statistical Office (Destatis) (2026).

[32] We define start-ups as establishments less than five years old.

[33] Behrens et al. (2017) find a similar diffusion of digital technologies among German establishments using a different data source.

premium of the frontier technology, A_F, and the shift in workforce age composition about its relative intensity in technological skills, $\Delta\alpha$. We also target the differential wage response of young and old incumbents, which is informative about younger workers' comparative advantage in technological skills, governed by Γ_T. Appendix D.4 develops the intuition for this identification in a static version of the model.

Constructing the model counterparts of these moments requires workers' age, which is not a state variable in the model. Fixing the age at labor-market entry to 25, however, the equilibrium transition rates determine the distribution of worker types and employer states at each subsequent age. Aggregating across entry cohorts, weighted by their size given retirement and labor-force growth, yields the probability that a worker is young (under 35) or old (over 45), conditional on her type and her employer's state. Weighting these probabilities by the firm's employment of each type gives its expected age composition. Appendix B.3 details the construction.

We compute the model counterparts of the event-study estimates by replicating the empirical design on the stationary equilibrium. At baseline, $k = -3$, we select firms operating the standard technology and follow them over six annual dates, $k \in \{-3, \dots, 2\}$.[34] We keep firms that remain standard through $k = -1$. Adopters are firms that upgrade to the frontier between $k = -1$ and $k = 0$, and controls are firms that do not. As in the data, we match adopters and controls on bins of employment and workforce age composition at $k = -1$. We do not match on productivity z, which is unobserved in the data. Because adoption in the model responds to productivity shocks, adopters and their matched controls may differ in productivity, and the model counterparts embed the same selection on unobservables as the empirical estimates. This is how the quantitative analysis accounts for the endogeneity concern discussed in Section 2.3.

Let $\bar{Y}_k^1$ and $\bar{Y}_k^0$ denote the weighted means of outcome Y among adopters and controls at event time k. The model counterpart of the event-study coefficient at $k = 2$ is

$$\hat{\beta}_2^{\text{model}} = \left(\bar{Y}_2^1 - \bar{Y}_2^0\right) - \left(\bar{Y}_{-1}^1 - \bar{Y}_{-1}^0\right). \tag{11}$$

We compute it for log employment, the share of young workers, and the log wage gap between young and old workers, $\log w_y - \log w_o$, and target the corresponding estimates at $k = 2$: β_2 from (1) for the first two outcomes, and τ_2 from (2) for the wage gap.

Finally, we calibrate the technological-skill intensity of standard firms, α_S, to match the average firm's share of young workers.[35]

Worker skills and early-career wage dynamics. We complement the adoption moments with moments of early-career wage dispersion and mobility from the Sample of Integrated Labour Market

[34] To mimic the size selection of the LPP, which covers establishments with at least 50 employees, we restrict the baseline sample to the largest 6 percent of firms in the model, the approximate percentile of the 50-employee threshold in the empirical establishment-size distribution.

[35] To be consistent with the age of entry in the model, we compute the share of young workers employed in an establishment in year t as $\text{Young}_t = \frac{N_{25-34,t}}{N_{25_+,t}}$, where $N_{25-34,t}$ is the number of full-time employees between 25 and 34 years old, and $N_{25_+,t}$ is the number of full-time employees older than 25.

Biographies (SIAB).[36] For each year between 2005 and 2015, we form a cohort of new entrants into full-time employment aged 25–29, and measure their initial wage distribution and their wage outcomes five and ten years later.[37] We compute each moment cohort by cohort and average across cohorts.

Because skill types are unobserved, these moments inform the parameters through the structure of the model. All entrants are inexperienced, and a share $1-\omega$ also has low technological skills (Li). Lacking both experience and technological skills, Li entrants tend to start at the bottom of the wage distribution. The share of each cohort earning below the 30th percentile of the contemporaneous wage distribution of all workers is therefore informative about $1-\omega$.

Subsequent wage mobility disciplines skill accumulation and obsolescence. Among workers who start in the bottom 30 percent, mostly Li entrants, the probability of leaving that group within five years is informative about the rate of experience accumulation, λ_E, and their average five-year wage growth about the productivity of experience, Γ_E. The evolution of within-cohort wage dispersion is informative about obsolescence. Workers at both ends of the initial distribution gain from experience, but those at the top, disproportionately Hi entrants, also lose technological skills over time. A higher obsolescence rate, λ_T, therefore slows wage growth at the top relative to the bottom and compresses the cohort's wage distribution. We target the ten-year change in the within-cohort interquartile range of log wages.

On-the-job training. We calibrate the training-cost shifter κ_{tr} to match the six-month incidence of employer-supported ICT training. The IAB Establishment Panel records the number of employees who participated in continuing training during the first half of the year. Restricting attention to training whose costs are paid at least partly by the employer, the aggregate participation rate is

$$\hat{\lambda}_{tr}^{6m} = \frac{\sum_j \varpi_j T_j}{\sum_j \varpi_j N_j} = 0.312,$$

where T_j is the number of participating employees in establishment j, N_j is its employment, and ϖ_j is its cross-sectional sample weight.

We make two adjustments to map this participation rate into the upgrades of technological skills in the model. First, training in the model restores the technological skills needed to operate digital and ICT technologies, whereas the IAB Establishment Panel covers continuing training of any content, most of which is unrelated to ICT. We therefore retain only the share of training devoted to information and communication technology. Janssen and Wölfel (2017) find that about one in ten non-formal training courses in the German National Educational Panel Study concern ICT, so we set this share to $\pi_{\text{ICT}} = 0.10$.[38]

[36] We use daily wages, including base pay and one-time bonuses, in 2012 prices. All wage moments are computed on wages residualized with respect to gender, state, and two-digit NACE industry.

[37] In the model, workers enter the labor force at age 25 as inexperienced; the age window accommodates variation in the timing of entry into full-time employment.

[38] This value is consistent with the German Adult Education Survey tabulations published by Eurostat: among individuals aged 25–64 in Germany in 2016, ICT accounted for 9.5 percent of job-related non-formal education and training activities. See Eurostat, *Distribution of education and training activities by field* (dataset `trng_aes_166`).

Second, training in the model always succeeds and fully restores the worker's technological skills. In practice, training may fail to generate new skills, and the skills acquired may be only partly transferable to the worker's job. We capture both margins with an effectiveness probability p_{tr}^{succ}. Evidence from the German Adult Education Survey shows that 62 percent of non-formal training activities led participants to report performing tasks at work or in everyday life differently afterward (Bilger et al., 2017). We therefore set $p_{tr}^{\text{succ}} = 0.60$.

The resulting calibration target is

$$\lambda_{tr,\text{target}}^{6m} = \hat{\lambda}_{tr}^{6m}\, \pi_{\text{ICT}}\, p_{tr}^{\text{succ}} = 0.312 \times 0.10 \times 0.60 \approx 0.019.$$

Technology upgrading and downgrading. Finally, we use the prevalence and persistence of frontier status to discipline the investment sensitivities ψ_{SF} and ψ_{FS}. The frontier technology in the model represents current best practice rather than a specific technology. During our calibration period, digital and ICT technologies were at the frontier.[39] We therefore calibrate the upgrading parameter, ψ_{SF}, to match the share of establishments using at least one digital technology in the 2017 IAB Establishment Panel, 32 percent.

The digitalization supplement is a single cross-section, so it cannot measure how often establishments lose frontier status. To discipline the downgrading parameter, ψ_{FS}, we instead use a question asked every year in the IAB-BP, in which establishments rate the technical state of their equipment relative to other establishments in their industry on a five-point scale, from "outdated" to "state of the art." Because this rating is relative to current best practice, it tracks an establishment's position with respect to a moving frontier. This matches the notion of obsolescence in the model, in which a firm that stops investing falls behind as technologies evolve. We classify establishments in the two highest categories as frontier and the remaining ones as standard. Among surviving establishments classified as frontier, about 16 percent are classified as standard one year later, and we target the corresponding annual transition probability in the model. Appendix A.5 details the construction of this measure and provides validation results.

4.3.2 Labor Market and Firm Dynamics

We discipline recruitment and exogenous worker separations using aggregate monthly transition rates between unemployment and employment. We target the average unemployment to employment (UE) and employment to unemployment (EU) rates over 2005–2014, calculated from the publicly available series of Hartung et al. (2026), who construct monthly transition rates for Germany over 2000–2014 using the SIAB. The UE rate primarily disciplines the recruitment-cost shifter, κ_v, which governs firms' incentives to recruit, while the EU rate helps discipline the exogenous separation rate, s_W, conditional on the other sources of job loss in the model.

Employer-initiated dismissals provide an additional moment to distinguish endogenous firing from other sources of separation. We calibrate the firing-cost shifter, κ_δ, to match the six-month dis-

[39] Using German linked employer–employee data for 2011–2016, Genz et al. (2021) similarly classify establishments adopting connected digital technologies as being at the technological frontier.

Table 2: Internally Calibrated Parameters and Targeted-Moment Fit

Technology and Skills						
Parameter		**Value**	**Target**	**Model**	**Data**	**Source**
A_F	TFP F	1.207	Log Employment DID	0.958	0.95	LPP
$\frac{\alpha_F}{\alpha_S}$	Relative skill intensity	1.52	Age Composition DID	0.0279	0.0275	LPP
α_S	Tech Intensity S	0.487	Average Young Workers (%)	0.214	0.225	LPP
Γ_T	High Tech-skill	1.876	Log wage gap DID	0.0185	0.020	LPP
ω	Entrant skill composition	0.64	Entrant (%) below P30	0.327	0.354	SIAB
λ_E	Experience acquisition	0.04	Pr. Above P30 (5 years)	0.38	0.41	SIAB
Γ_E	High Experience	2.27	Wage Growth Entrants below P30 (5 years)	0.35	0.38	SIAB
λ_T	Tech-skill obsolescence	0.036	Change sd log wage (5 years)	-0.067	-0.07	SIAB
κ_{tr}	Training cost	46.73	Six-month training rate	0.019	0.019	SIAB
ψ_{FS}	Tech Downgrade	2.12	$F \to S$ transition rate	0.158	0.162	IAB-BP
ψ_{SF}	Tech Upgrade	1.5	F firms (%)	0.312	0.32	IAB-BP
Labor Market and Firms						
Parameter		**Value**	**Target**	**Model**	**Data**	**Source**
κ_v	Recruitment cost shifter	0.3	Monthly UE	0.0562	0.0552	Hartung et al. (2026)
κ_δ	Dismissal cost	15.23	Dismissal rate	0.0148	0.0152	Bellmann et al. (2018)
s^F	Firm entry rate	0.013	Annual firm entry rate	0.0515	0.0506	Schröpf (2023)
κ	Firm cost parameter	7,002	Average log firm size	1.855	1.850	IAB-BP
b_{He}	Unemployment value (He)	1	$\frac{b_{He}}{C_{He}}$	0.967	1	Carrillo-Tudela et al. (2023)
b_{Hi}	Unemployment value (Hi)	0.88	$\frac{b_{Hi}}{C_{Hi}}$	0.974	1	Carrillo-Tudela et al. (2023)
b_{Le}	Unemployment value (Le)	0.78	$\frac{b_{Le}}{C_{Le}}$	0.982	1	Carrillo-Tudela et al. (2023)
b_{Li}	Unemployment value (Li)	0.65	$\frac{b_{Li}}{C_{Li}}$	0.976	1	Carrillo-Tudela et al. (2023)

Notes: The table reports internally calibrated structural parameters and the associated empirical and model moments. All parameters are estimated jointly by simulated method of moments with equal weight on each target. The first panel covers technology choice, skill endowments, skill transitions, and training; the second covers recruiting, separations, firm entry, and unemployment flow payoffs. Adoption-event targets come from the LPP matched event studies, early-career wage moments from SIAB, and additional labor-market and technology moments from the cited data sources.

missal rate of 1.52% reported by Bellmann et al. (2018). Using the IAB Establishment Panel over 1993–2014, they measure dismissals during the first six months of each year relative to average establishment employment. This moment directly informs the cost of adjusting employment through the firing margin.

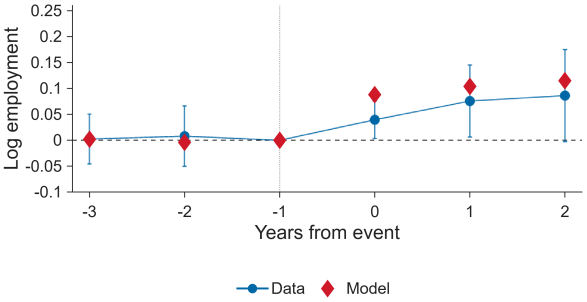
We calibrate the firm entry cost, κ, to match mean log firm employment. Through the free-entry condition, this parameter governs the equilibrium mass and size of firms. We also calibrate the exogenous firm exit rate to match an average annual establishment-entry rate of approximately 5.06% over 2005–2015, based on the IAB Establishment History Panel (BHP) evidence in Schröpf (2023). The entry rate is informative about firm exit because, along the balanced growth path, entry must replace exiting firms and accommodate growth in the firm population.

Finally, we set the flow value of unemployment separately for each worker type i according to

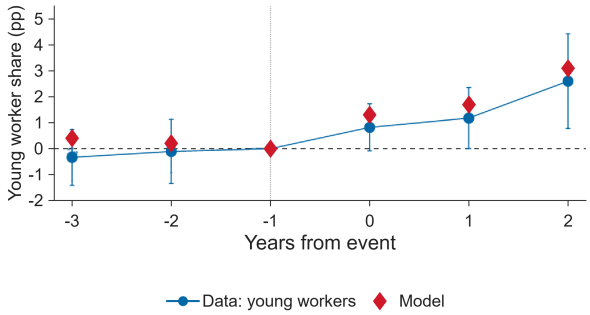
$$b_i = 0.46\bar{w}_i + 0.24\frac{Y}{E} \equiv C_i$$

where $\bar{w}_i$ is the average wage of employed workers of type i and Y/E is average labor productivity in the calibrated equilibrium. The first component captures unemployment cash income. Following

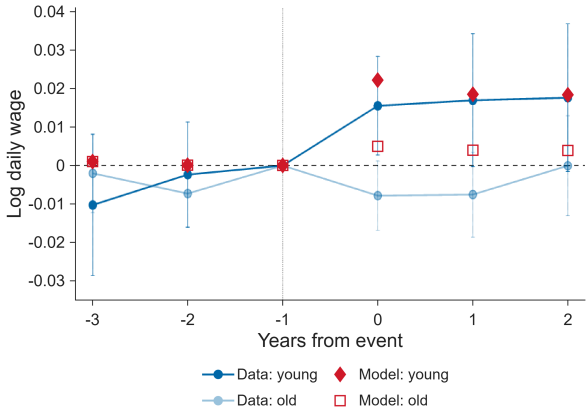
Figure 7: Technology-Adoption Event Studies: Model and Data



(a) Employment response to technology adoption



(b) Workforce age-composition response to technology adoption



(c) Wages response to technology adoption

Notes: Blue estimates are the LPP matched event-study coefficients, while red markers are their simulated counterparts from the calibrated model. The same event-study design is applied to model-generated firms: standard-technology firms are followed around an upgrade to the frontier, with controls matched on pre-adoption employment and workforce age composition. The model's latent age distribution is used to construct young- and old-worker outcomes. Panels report log employment, the young-worker employment share in percentage points, and the log-wage responses of young and old incumbent workers. Only the coefficients at $k = 2$ are targeted in calibration; the remaining event times are untargeted validation moments. The omitted event time is $k = -1$.

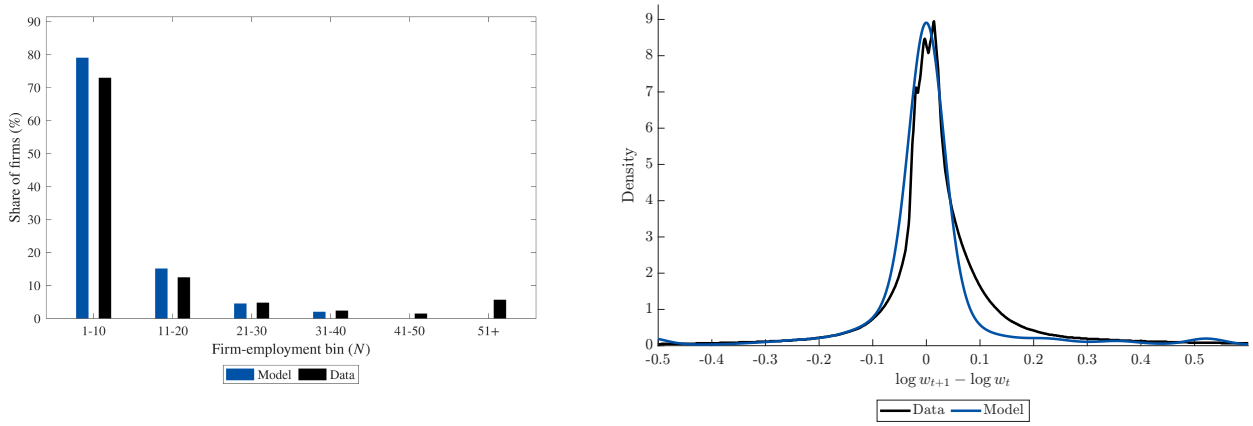
Carrillo-Tudela et al. (2023), we set it equal to 46% of the type-specific average wage.[40] The second component captures the value of leisure. We use the leisure value of 0.24 calibrated by Hartung et al. (2026) as a benchmark and scale it by average labor productivity in our model. All worker types therefore share the same replacement rate and leisure value, while differences in their average wages generate variation in b_i.

4.4 Moment Fits and Validation

We estimate the internally calibrated parameters through a Simulated Method of Moments, minimizing the mean squared deviations between simulated and empirical moments. We use a diagonal weighting matrix that assigns equal weight to each moment. Table (2) reports the calibrated parame-

[40] Carrillo-Tudela et al. (2023) adopt this benchmark from Krebs and Scheffel (2013), who use OECD tax-benefit data to compute a post-Hartz net replacement rate of 46% for long-term unemployment.

Figure 8: Firm-Size and Wage-Change Distributions: Model and Data



(a) Firm Size **(b)** $\log w_{t+1} - \log w_t$

Notes: The figure compares untargeted distributional moments generated by the calibrated model with their empirical counterparts. Panel (a) reports the fraction of establishments in employment-size bins, and panel (b) the distribution of annual changes in log daily wages, $\log w_{t+1} - \log w_t$. Firm-size moments are based on establishment data, while wage-change moments are computed from linked worker employment histories. The model distributions are computed from the stationary equilibrium and the implied worker transitions.

ter values and compares the targeted moments in the model and the data.

We validate the calibrated model against moments that we do not target. First, although we target only the event-study coefficients at $k = 2$, the model reproduces the full dynamic response to adoption, including the absence of pre-trends (Figure 7). Because these responses discipline the technology parameters, matching their entire path, and not only the targeted horizon, supports our calibration strategy. Second, the model generates a realistic firm size distribution (Figure 8a) and a realistic distribution of the share of young workers across firms (Table 3). Since firm size and workforce age composition are key determinants of the incentive to adopt in the model, matching these distributions suggests that the model also captures how adoption incentives vary across firms. Finally, the model reproduces the distribution of annual wage changes across all workers (Figure 8b). Since we discipline the skill parameters using early-career cohorts only, this suggests that the implied skill dynamics also generate realistic wage changes for the workforce as a whole.

5 Quantitative Results

This section quantifies how population aging affects aggregate productivity and output through changes in the skill composition of the labor force. We introduce aging by varying the demographic parameter g while holding the retirement rate λ_r fixed, thereby changing the inflow of new workers relative to the size of the labor force. Section 5.1 describes the construction of this demographic experiment. We first compare the calibrated baseline steady state with an alternative steady state calibrated to the demographic conditions projected for the 2050s. Section 5.2 examines the resulting changes in workforce skills, technology adoption, productivity, output, and consumption. Section 5.3

Table 3: Distribution of Young-Worker Shares: Model and Data

Percentile	Data (%)	Model (%)
Mean	22.53	21.38
p_{25}	8.78	5.60
p_{50}	20	16.80
p_{75}	33.33	34.23

Notes: The table compares the distribution of the establishment-level share of young employees in the data with the distribution generated by the stationary model. Young workers are aged 25–34, and shares are measured relative to employees aged 25 and older, consistent with the model's labor-market entry age. The model constructs chronological ages from the equilibrium worker-state transition process and then aggregates age probabilities across each firm's workforce. Percentiles are across establishments, not across individual workers; entries are percentage shares.

introduces an exact decomposition of the change in productivity, which isolates the contributions of changes in the skill composition of the workforce and in technology adoption.

5.1 Construction of the Demographic Experiment

This subsection explains how we introduce aging into the model. Appendix A.7 provides the details of the construction. In the model, workers enter at age 25 and leave the labor force at the constant retirement rate λ_r. Conditional on this rate and the initial age distribution, the sequence of new entrants determines the evolution of cohort sizes and, consequently, the age composition of the workforce. The law of motion of the labor force (eq. 3) implies that, for a given λ_r, the growth rate g_t maps one-to-one into the inflow rate of young workers. We therefore construct the path of g_t to match the projected changes in the inflow rate of young workers relative to the labor force.

We combine observed population age composition for 2015–2025 with projections from the Federal Statistical Office's 16th Coordinated Population Projection (Statistisches Bundesamt, Statistisches Bundesamt, 2025).[41] Let $P_{25,t}$ denote the population aged exactly 25 in year t, and let $P_{a,t}$ denote the population in age group a. We define the empirical entrant share as

$$\widehat{e}_t = \frac{\bar{\pi}_{25\text{–}29} P_{25,t}}{\sum_{a \in \mathcal{A}} \bar{\pi}_a P_{a,t}}, \qquad \mathcal{A} = \{25\text{–}29, 30\text{–}34, \ldots, 60\text{–}64, 65+\},$$

where $\bar{\pi}_a$ is the average labor-force participation rate of age group a over 2005–2015.[42] The numerator approximates the inflow of new labor-market participants, while the denominator measures the labor force implied by the population age distribution at fixed participation rates. Changes in $\widehat{e}_t$ therefore reflect changes in population age composition. We anchor the entrant-share series to the calibrated

[41] Observed population stocks by age are from Destatis GENESIS table 12411-0005. For subsequent years, we use Model Calculation 1 ($G2L2W0$) of the 16th Coordinated Population Projection, reported in GENESIS table 12421-0002. This scenario assumes moderate developments in fertility and life expectancy and zero net migration.

[42] Participation rates are computed from the Destatis Microcensus, GENESIS table 12211-9001 (Federal Statistical Office (Destatis), 2026), and held fixed throughout the experiment. We use the participation rate for ages 25–29 to approximate participation at age 25 because the historical data do not report a separate rate for that age.

steady state in 2016 and use its subsequent relative changes to recover the path of g_t from the model's labor-force law of motion. Averaging the implied quarterly rates over 2050–2059 gives an annual growth rate of approximately $g^{2050} \simeq -0.358\%$.

Given the greater uncertainty surrounding population projections at longer horizons, we assume that the demographic change eventually stabilizes at the level implied by the 2050s. We therefore hold g_t fixed at g^{2050} from 2060 onward. The long-run steady state we solve for retains all parameters of the calibrated economy except g, which is set to this value.

5.2 Long-Run Impact of Aging

We compare the calibrated baseline steady state with the long-run alternative steady state. Figure 9 summarizes the resulting changes in workforce skills, technology adoption (defined as the share of frontier firms), aggregate productivity, employment rate, and welfare. To account for decreasing returns to scale, we measure aggregate productivity as output per efficiency-adjusted unit of employment, $\mathcal{P} \equiv \frac{\mathcal{Y}}{\mathcal{E}^\eta}$, where η is the returns-to-scale parameter. We define GDP as aggregate output net of the resource costs incurred by firms, $GDP \equiv \mathcal{Y} - c$.[43] Finally, because workers have linear utility, aggregate welfare equals GDP plus the total flow value of unemployment, $\mathcal{W} \equiv GDP + \sum_{i \in \mathcal{G}} u_i b_i$, where u_i is the mass of unemployed workers of type i and b_i is their flow value of unemployment. Aggregate quantities are reported per member of the labor force. All reported changes are percentage deviations from the baseline.

The decline in cohort entry shifts the workforce toward workers who have had more time to accumulate experience. The share of experienced workers, $(N_{He} + N_{Le})/\mathcal{L}$, rises by 2.81% in the long-run equilibrium. This increase occurs in both technological-skill groups: the shares of He and Le workers each rise by roughly the same percentage over the transition (Figure A10). At the same time, lower entry reduces the inflow of workers endowed with high technological skills, while incumbent workers remain exposed to skill obsolescence. The share of high-tech workers, $(N_{He} + N_{Hi})/\mathcal{L}$, consequently falls by 1.71%. Firms respond to the technological skill shortage by increasing training investment: training incidence rises by approximately 1%, helping to replenish technological skills among incumbent workers.

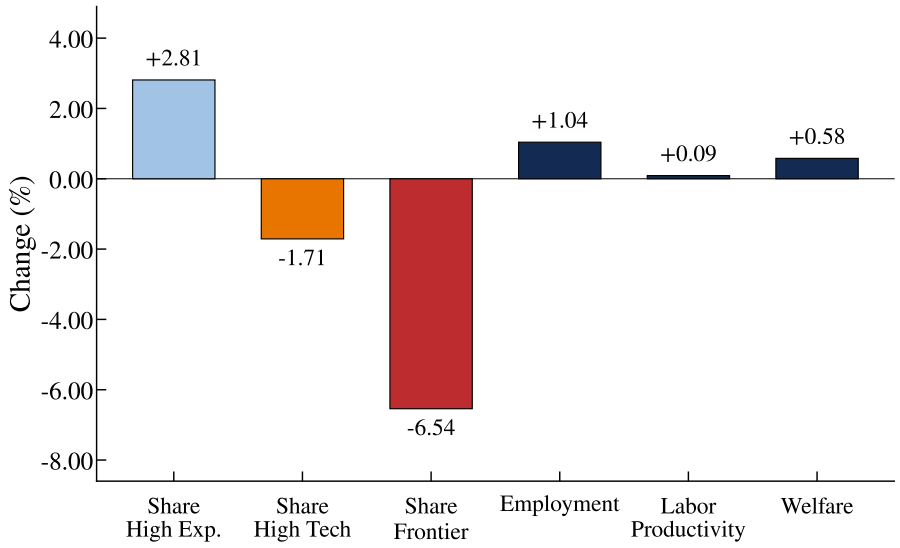
The change in skill composition affects firms' technology choices. The shift toward experience and away from technological skills reduces the relative profitability of the frontier technology. In equilibrium, technology adoption, measured as the share of active firms operating the frontier technology, falls by approximately 6.5%. Despite this decline in adoption, aggregate productivity is basically unchanged (+0.09%). Since the overall employment rate increases (+1.04%) while productivity is constant, the GDP per capita is larger, and aggregate welfare increases by 0.58%.[44] The employment gains, however, accrue entirely to experienced workers (+0.8%), as the expansion of experience-intensive standard firms raises the demand for their skills, while the employment rate of inexperienced workers falls (-1.47%).

Table A7 in the appendix reports additional statistics that summarize the long-run impact of aging

[43] Resource costs include expenditures on recruitment, firing, training, and technology investment.

[44] The increase in welfare also reflects changes in resource costs between the two steady state equilibria.

Figure 9: Long-Run Effects of Aging on Skills, Technology, and Welfare



Notes: Bars report percentage changes between the stationary equilibrium calibrated to Germany in 2005–2015 and the equilibrium with the labor-force entry rate implied by projected 2050s demographics. All other structural parameters, including the retirement hazard, are held fixed. High experience and high technological skill denote the shares of workers with the corresponding skill, and technology adoption denotes the share of active firms operating the frontier technology. Employment is the fraction of the labor force employed. Aggregate TFP is defined as output divided by employment raised to the returns-to-scale parameter, Y/E^{η}. Welfare is the stationary flow value of net output plus the unemployment flow payoff, per member of the labor force.

in the economy. Consistent with recent papers that link demographic aging to the slowdown of business dynamics (e.g., Karahan et al. (2024); Hopenhayn et al. (2022); Inokuma and Sanchez (2023)), our model predicts a sharp decrease in firm entry rate due to aging (-7.8%), as well as an increase in the average firm size (+0.5%).

5.3 Productivity Decomposition

This section answers the central question of the paper: how large is the productivity cost of the technological skill shortage generated by aging? Aging changes three determinants of aggregate labor productivity at once: it raises the stock of experienced workers, lowers the stock of tech-proficient workers, and reduces the share of firms operating the frontier technology. Because these changes push productivity in opposite directions, the small net effect documented in Section 5.2 does not reveal how large each of them is. We therefore develop an exact decomposition of the change in aggregate labor productivity that isolates the contribution of each of the three changes.

The decomposition compares the initial stationary equilibrium, calibrated to 2005–2015, with the final equilibrium implied by the labor-force growth rate projected for the 2050s. Let equilibrium 0 denote the initial equilibrium, with labor-force growth g_0, and equilibrium 1 the 2050s equilibrium, with slower growth $g_1 < g_0$. Let $\mathcal{P} = \frac{y}{\mathcal{E}^{\eta}}$ denote aggregate productivity. We decompose $\Delta\mathcal{P} \equiv \frac{\mathcal{P}^1 - \mathcal{P}^0}{\mathcal{P}^0}$, so that negative contributions are channels through which aging lowers aggregate productivity, and positive contributions are offsetting forces.

To understand the relevant forces at play, consider an individual firm with employment l, productivity z, technology τ, tech skills per worker $\bar{T} \equiv T/l$, and experience per worker $\bar{E} \equiv E/l$. The firm's output can be written as

$$y = \underbrace{z A_\tau \bar{T}^{\eta\alpha_\tau} \bar{E}^{\eta(1-\alpha_\tau)}}_{\text{TFP}} \times l^\eta$$

Thus, firm-level TFP depends on four objects: the exogenous productivity state, the technology used, and technological skills and experience per employed worker. Aggregate productivity $\mathcal{P}$ combines these firm-level objects across firms. Because of decreasing returns to scale, however, it is not simply an average of firm-level TFP: for given total employment, aggregate output also depends on how employment is spread across firms, and rises when it is distributed more evenly. Aggregate productivity therefore reflects the distribution of firms across productivity states and technologies, the economy-wide supply of technological skills and experience, how these skills are allocated across firms, and the distribution of employment across firms.

For any stationary equilibrium, let Λ denote the stationary per-capita measure of active firms across the state space. Define aggregate employment as $\mathcal{E}$. Also, denote the decreasing-returns term as $\Psi_{\text{DRS}} \equiv \frac{\int l(\hat{s})^\eta \, d\Lambda(\hat{s})}{\mathcal{E}^\eta}$, which captures how the number of active firms and the distribution of employment across them affect aggregate output under decreasing returns to scale: holding total employment fixed, it increases when employment is distributed more evenly across a larger number of firms. We also define the size-adjusted distribution of firms

$$\widetilde{\Lambda}(d\hat{s}) \equiv \frac{l(\hat{s})^\eta}{\int l(\hat{s})^\eta \, d\Lambda(\hat{s})} d\Lambda(\hat{s}).$$

This distribution weights firms according to the employment component of their contribution to aggregate output. Because firm output is proportional to l^η, larger firms receive greater weight, but, under decreasing returns to scale, less than proportionally to their employment. Let $\widetilde{\mathbb{E}}$ denote expectations under this distribution, and let $\widetilde{\pi}_z$ and $\widetilde{\pi}_{\tau|z}$ denote its implied marginal and conditional probabilities. Finally, define employment-weighted average technological skills and experience as

$$\bar{T} \equiv \frac{1}{\mathcal{E}} \int l(\hat{s}) T(\hat{s}) \, d\Lambda(\hat{s}), \qquad\qquad \bar{E} \equiv \frac{1}{\mathcal{E}} \int l(\hat{s}) E(\hat{s}) \, d\Lambda(\hat{s}).$$

The following proposition shows how to rewrite aggregate productivity to isolate economically meaningful channels.

Proposition 3 (Aggregate Productivity)**.** *Aggregate productivity can be written as*

$$\frac{\mathcal{Y}}{\mathcal{E}^\eta} \;=\; \Psi_{\text{DRS}} \underbrace{\sum_z z \, \widetilde{\pi}_z}_{\textit{Selection } z} \underbrace{\sum_\tau A_\tau \, \widetilde{\pi}_{\tau|z}}_{\textit{Technology Adoption}} \underbrace{\widetilde{\mathbb{E}}\left[T^{\eta\alpha_\tau} E^{\eta(1-\alpha_\tau)} \,\middle|\, \tau, z \right]}_{\textit{Skill Composition}}. \tag{12}$$

Equivalently,

$$\frac{\mathcal{Y}}{\mathcal{E}^{\eta}} = \Psi_{\mathrm{DRS}} \underbrace{\sum_{z} z\,\widetilde{\pi}_{z}}_{\text{Selection } z} \underbrace{\sum_{\tau} A_{\tau}\,\widetilde{\pi}_{\tau|z}}_{\text{Technology Adoption}} \underbrace{\bar{T}^{\eta\alpha_{\tau}}\bar{E}^{\eta(1-\alpha_{\tau})}}_{\text{Skill Supply}} \underbrace{\widetilde{\mathbb{E}}\left[\left(\frac{T}{\bar{T}}\right)^{\eta\alpha_{\tau}}\left(\frac{E}{\bar{E}}\right)^{\eta(1-\alpha_{\tau})}\middle|\,\tau, z\right]}_{\text{Skill Allocation}}. \tag{13}$$

The final conditional expectation captures the allocation of workforce skills across firms, relative to their economy-wide averages.

Equation (12) implies that aggregate output retains the same general structure as firm-level production:

$$\mathcal{Y} = \underbrace{\Psi_{\mathrm{DRS}}\widetilde{\mathbb{E}}\left[zA_{\tau}T^{\eta\alpha_{\tau}}E^{\eta(1-\alpha_{\tau})}\right]}_{\text{Aggregate Productivity}}\mathcal{E}^{\eta}.$$

Aggregate productivity therefore combines the size-adjusted average of firm-level productivity with the decreasing-returns term Ψ_{DRS}. Under constant returns to scale, $\eta = 1$, this term equals one and the size-adjusted firm distribution reduces to the standard employment-weighted distribution.

This representation provides the basis for our productivity decomposition. Moving from the initial to the new equilibrium changes each component of (13): the decreasing-returns term, the distribution of firms across productivity states, technology adoption, the average levels of technological skills and experience, and the allocation of those skills across firms. We use a Shapley decomposition to quantify the contribution of each channel to the total change in aggregate productivity (Shorrocks, 2013).[45]

Figure 10 decomposes the percentage change in aggregate productivity between the 2005–2015 stationary equilibrium and the 2050s one. Aging raises aggregate productivity by only 0.09%, but this small net effect masks large offsetting forces.

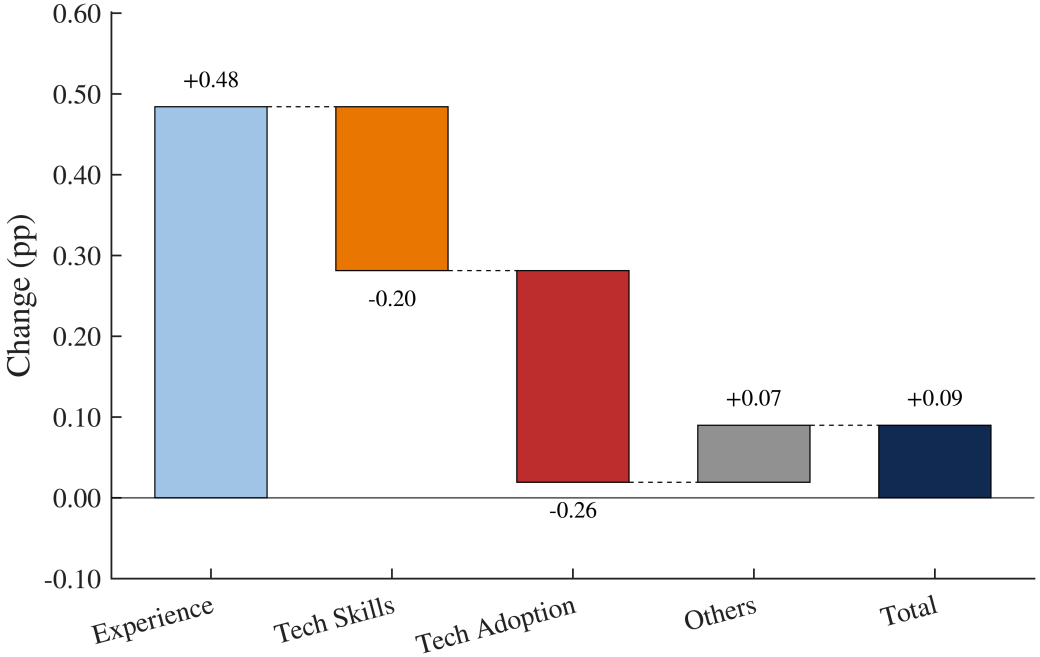
The larger stock of experienced workers raises aggregate productivity by 0.48%. This is the channel emphasized by standard theories of human capital accumulation, and on its own it would imply a sizable productivity gain from aging. A model without the interaction between technological skills and adoption would therefore erroneously predict that aging raises productivity significantly.

The technological skill shortage, the novel channel of this paper, entirely offsets this gain. The smaller stock of tech-proficient workers lowers aggregate productivity directly by 0.20%. Because frontier technologies rely heavily on these skills, their scarcity also makes the frontier technology less profitable, and the resulting decline in adoption lowers productivity by a further 0.26%. Together, these two effects reduce aggregate productivity by 0.46%, canceling the long-run productivity gains from experience.

The remaining channels, grouped as "Others" in Figure 10, account for the small net increase. They comprise the selection channel (-0.13%), which reflects a shift of the size-adjusted firm distri-

[45] For each channel, we compute the change in aggregate productivity obtained by replacing its equilibrium-(0) value with its equilibrium-(1) counterpart, holding a subset of the remaining objects at their equilibrium-(1) values and the others at their equilibrium-(0) values. We then average this marginal effect over all possible orderings of the changes. The resulting contributions are order-independent and sum exactly to the total change in aggregate productivity.

Figure 10: Decomposition of the Aggregate Productivity Effect of Aging



Notes: The bars are grouped Shapley contributions to the percentage change in aggregate TFP between the 2005–2015 and 2050s stationary equilibria. Experience and technological skills distinguish the two components of skill supply; technology adoption captures shifts between standard and frontier firms. The residual group labeled "Others" collects the remaining channels, including productivity selection, skill allocation across firms, and the firm-mass/size-distribution term arising from decreasing returns. The contributions sum to the total change.

bution toward lower idiosyncratic productivity states; the skill-allocation channel ($+0.03\%$); and the decreasing-returns channel ($+0.17\%$), which captures changes in the mass of firms and in the distribution of employment across them. Figure A8 in the appendix shows the decomposition results unpacking these residual channels.

5.4 Heterogeneous Impact of Aging

We conclude this section by examining how the effects of aging differ across workers. Because the decline in frontier adoption shifts labor demand across skill types, aggregate outcomes mask substantial heterogeneity in employment and welfare. As reported in figure 9, Table A7 in the Appendix shows that the aggregate unemployment rate increases as the economy ages, but this masks opposite effects across worker types. Unemployment declines among experienced workers, He and Le, whereas it rises among inexperienced workers, Hi and Li. This pattern reflects the endogenous response of technology adoption to aging. As the firm distribution shifts away from frontier technologies and toward standard technologies, which use experience more intensively, equilibrium labor demand tilts toward experienced workers and away from inexperienced workers. This is consistent with Bianchi et al. (2023), who argue that longer careers among older incumbents reduce employment opportunities for younger workers and slow their progression along the job ladder.

Table A8 shows that the welfare effects of aging are similarly heterogeneous. Aggregate welfare

rises slightly, but the gains are concentrated among inexperienced tech-proficient workers. Experienced workers, He and Le, lose by similar amounts regardless of their technological skills, because aging makes experience more abundant and lowers its return. Among inexperienced workers, technological skills determine the outcome: tech-proficient workers, Hi, gain the most, because they supply a skill that becomes scarcer and better rewarded, while workers without these skills, Li, experience small losses. The same contrast arises among new entrants. Because inexperienced workers are disproportionately young, aging increases welfare and consumption dispersion among younger workers, while reducing it among experienced workers. This pattern is consistent with evidence of rising dispersion among young workers in developed economies (Antonczyk et al., 2018).

6 Policy Evaluation

The mechanisms highlighted by our model are consistent with growing policy concerns that population aging exacerbates technological skill shortages and widen the mismatch between the skills supplied by workers and those required by firms. This is highlighted, for instance, by the European Commission's 2024 action plan, which identifies demographic change and evolving skill requirements as pressures that could hinder the digital transition (European Commission, 2024).

Expanding support for continuing training has been a prominent policy response to these concerns.[46] In Germany, the focus of our quantitative analysis, the *Qualifizierungschancengesetz* broadened support for training employed workers, explicitly seeking to help them adapt to demographic and technological change (Bundesregierung, 2018).[47]

The productivity decomposition in Section 5.3 clarifies how training policies may support productivity in an aging economy. The gains from a more experienced workforce are completely offset by a decline in technological skills and the subsequent foregone adoption of frontier technology. Training can alleviate both these negative forces. By updating the technological skills of incumbent workers, it can reduce the skill shortages associated with smaller entering cohorts while preserving the benefits of accumulated experience, and, as a result, also raise the profitability of adopting frontier technology.

We use the calibrated model to quantify the effect of Germany's expansion of training support on aggregate productivity. The model is well suited for this exercise because it jointly determines workers' skill composition, firms' training and adoption decisions, and the allocation of employment across firms. The policy's welfare implications, however, differ from its productivity effects. Because the decentralized equilibrium is constrained efficient, a utilitarian planner facing the same frictions cannot improve on it, so a training subsidy cannot raise aggregate welfare over the transition, even if

[46] Governments have also pursued measures that directly support digitalization. Japan's Society 5.0 strategy promotes digital technologies and robotics as part of its response to demographic change (Council for Science, Technology and Innovation, 2015), while Australia's 2023 Intergenerational Report discusses investment in digital infrastructure and support for technology diffusion as means of sustaining productivity growth (Australian Treasury, 2023).

[47] Other aging economies have expanded support for skill renewal through related instruments. Italy's *Fondo Nuove Competenze* finances labor costs associated with employee training for firms' digital transitions (Ministero del Lavoro e delle Politiche Sociali, 2024), while Sweden provides income support for experienced workers undertaking further education (CSN, 2026).

it raises productivity. The subsidy can nonetheless redistribute welfare across workers by changing skill acquisition, wages, and employment opportunities. We therefore identify its winners and losers across worker types and cohorts.

6.1 Training Subsidies in Germany

The *Qualifizierungschancengesetz*, effective January 2019, substantially broadened training support, extending it to employed workers of all ages, qualification levels, and firm sizes. Its stated objective was to prepare the incumbent workforce for future skill needs, addressing shortages associated with demographic change and helping workers adapt to technological change before their employment prospects deteriorate (German Federal Government, 2018).[48] Reforms in 2020 and 2024 further strengthened support and simplified access, making training subsidies a permanent feature of German labor-market policy.

The current scheme, under Section 82 of Social Code III, supports training within an existing employment relationship. It covers up to 100% of course fees and, under the standard schedule, up to 75% of wages paid during training-related absences.[49] Eligible courses must exceed 120 hours, be approved together with their providers, and develop skills beyond job- or firm-specific knowledge (Federal Employment Agency, nd). These features make the policy well suited to our framework: its focus on employed workers corresponds to the model's training margin, and its support for general skills matches our treatment of training as an investment in transferable technological skills.

6.2 Macroeconomic Implications of Training Subsidy

We model public support for employee training as an exogenous subsidy rate, $\sigma_{tr} \in [0, 1)$, that proportionally reduces firms' training costs. For a training intensity λ_{tr}, the firm pays $(1-\sigma_{tr})C_{tr}(\lambda_{tr})$, while the government covers the remaining share. Accordingly, we modify the training-cost term in the joint-surplus problem, equation (6), to reflect the subsidy. Firms continue to choose training intensity optimally, taking the subsidy rate as given.

The government finances the subsidy through a uniform, non-distortionary lump-sum tax, $\mathcal{T}_t$, paid by every member of the labor force. Let Λ_t denote the measure of active firms over firm states $s = (\vec{n}, z, \tau)$, and recall that $\mathcal{G}_L = \{Li, Le\}$ contains the worker types eligible for training. Total government expenditure per-capita on the subsidy at date t is

$$\mathcal{B}_{tr,t} = \sigma_{tr} \int \sum_{i \in \mathcal{G}_L} n_i C_{tr} \left(\lambda_{tr,i,t}(s)\right) d\Lambda_t(s)$$

The government balances its budget at every instant, so the tax paid by each worker satisfies $\mathcal{T}_t = \mathcal{B}_{tr,t}$.

[48] The original bill projected additional annual expenditure of up to EUR 1 billion in the medium run on the broader expansion of training support, including course costs, wage subsidies, and benefits for unemployed participants (German Federal Government, 2018, pp. 19–20).

[49] For firms with fewer than 50, 50–499, and at least 500 employees, standard course-fee subsidies are 100%, 50%, and 25%, and wage subsidies are 75%, 50%, and 25%. More generous provisions apply to some groups: course fees can be fully covered for workers aged 45 or older and for workers with severe disabilities in firms with fewer than 500 employees, and wage support can reach 100% for workers training toward a vocational qualification they lack.

We calibrate σ_{tr} using observed participation in publicly supported employee training. In 2025, approximately 70,000 employees began a subsidized training course, corresponding to approximately 0.166% of employment (Statistics of the Federal Employment Agency, 2026). We use this participation rate to set the scale of the policy experiment, targeting an increase of 0.166 percentage points in annual training incidence relative to the baseline steady state. In the calibrated economy, approximately 1.9% of employed workers are successfully retrained each year. The target therefore corresponds to an increase of approximately 9% in annual training incidence. Holding the other parameters at their baseline values, we choose the subsidy rate that generates this increase. As shown in Figure A11, the implied subsidy rate is approximately $\sigma_{tr} = 0.10$.

To evaluate the policy in an aging economy, we compare two equilibrium transitions starting from the same calibrated 2005–2015 allocation and following the demographic path described in Section 5.1. In the first, the economy undergoes the demographic transition absent the training subsidy. In the second, the calibrated subsidy is introduced at the beginning of the transition and remains in place permanently.[50] Both economies converge to equilibria with the demographic parameters implied by the 2050s projections, but under different policy regimes. Comparing these transitions allows us to quantify the aggregate and distributional effects of the subsidy while holding the demographic experiment fixed.

Figure (11) compares the aging transition without policy intervention with the transition under the permanent 10% training subsidy. The subsidy induces an immediate increase of approximately 9.5% in the annual training rate relative to the calibrated baseline, as shown in Figure A11. As additional training replenishes incumbent workers' technological skills, labor productivity rises and adopting frontier technology becomes more profitable. However, training is a costly investment that requires resources to materialize. Consequently, flow consumption initially falls by approximately 0.6% relative to the baseline and remains below the no-policy path during the early years of the transition.

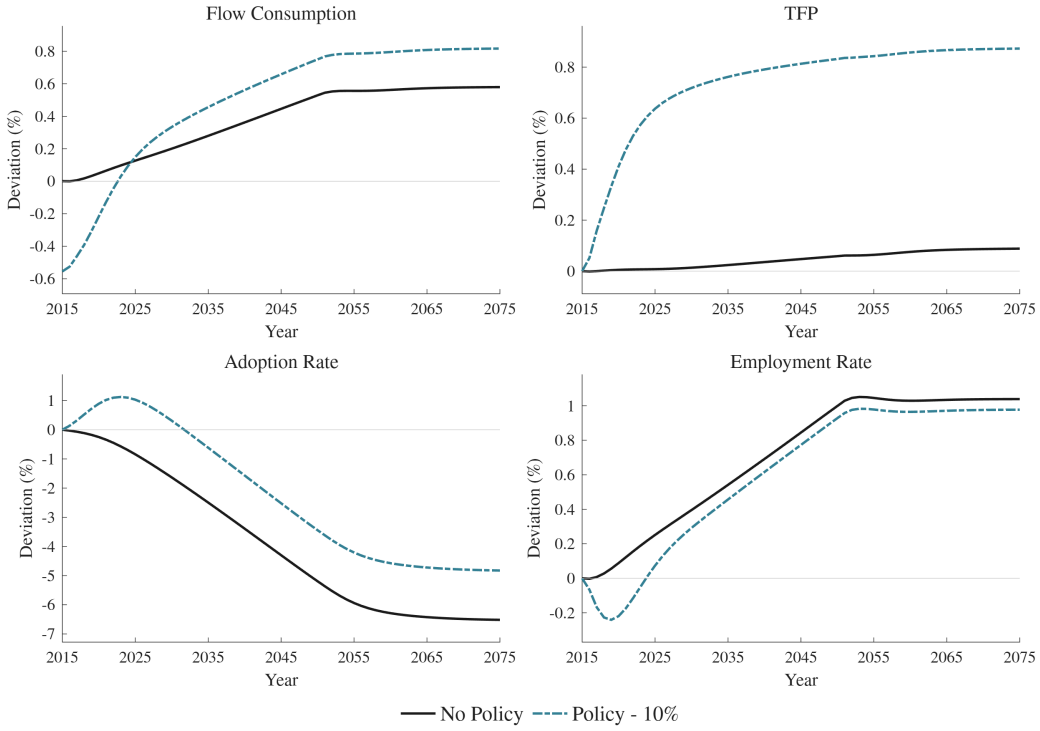
Over time, the accumulation of technological skills supports higher productivity and consumption while attenuating the decline in frontier-technology adoption associated with aging. Adoption initially rises above its baseline level before declining as the demographic transition progresses, but remains above the no-policy path throughout the transition. By the end of the plotted horizon, adoption is approximately 4.8% below the baseline, compared with 6.4% without the subsidy, while TFP is approximately 0.85% above the baseline, compared with 0.09% without intervention. Consumption eventually exceeds the no-policy path, although its improvement is smaller because maintaining higher training intensity continues to absorb resources.

6.3 Winners and Losers

We now examine how the welfare effects of training subsidies differ across workers. Let $t = 0$ denote the beginning of the aging transition and $\mathcal{W}_0$ aggregate discounted welfare evaluated at that date. Because the equilibrium without intervention is constrained efficient, a training subsidy cannot

[50] The reform took effect in 2019, three years after the start of our demographic transition. For simplicity, we assume that the subsidy is introduced at the beginning of the transition.

Figure 11: Aging Transitions With and Without a Training Subsidy



Notes: Each panel shows percentage deviations from the initial 2005–2015 stationary equilibrium along the demographic transition beginning in 2016. The solid path holds training policy unchanged; the dashed policy path introduces a permanent 10% subsidy to firms' training expenditures at the beginning of the transition, financed by an equal lump-sum tax on workers. The two economies face the same demographic path. Flow consumption is measured after the resources used for training and other firm adjustments; TFP is Y/E^{η}, adoption is the frontier-firm share, and employment is the employed fraction of the labor force. The horizontal axis is calendar time.

increase $\mathcal{W}_0$ under the utilitarian welfare criterion.[51] Figure A13 confirms that aggregate discounted welfare decreases with the subsidy rate. This aggregate loss, however, can coexist with gains for particular groups of workers. We therefore examine how the policy affects incumbent workers with different initial skills and successive cohorts of labor-market entrants.

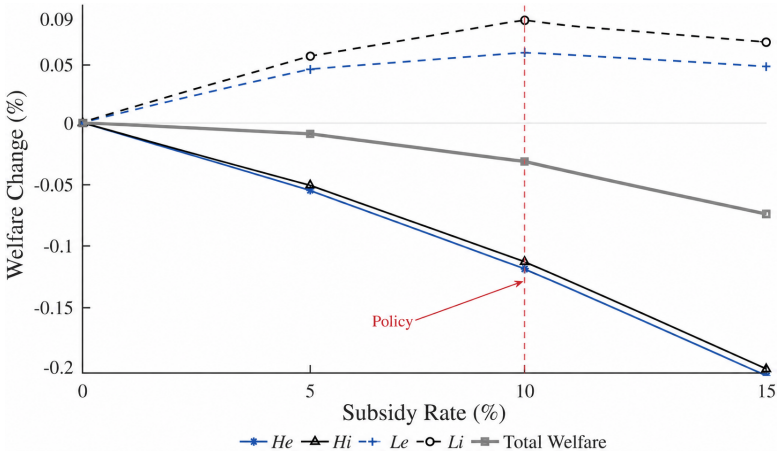
Denote the average life-discounted utility of workers of type i alive at instant t as $\bar{V}_{i,t}$. Appendix B.5 shows how we compute these welfare measures for each group of workers. We compute the percentage welfare changes relative to the aging transition without intervention. For incumbents

[51] Let υ_t denote aggregate flow utility per member of the labor force. Then

$$\frac{\mathcal{W}_0}{\mathcal{L}_0} = \int_0^{\infty} \exp\left(-\rho t + \int_0^t g_u \, du\right) \upsilon_t \, dt.$$

The population adjustment accounts for the changing size of the labor force along the demographic transition. With linear utility, υ_t sums consumption and any nonmarket unemployment payoffs; it need not coincide with market consumption alone.

Figure 12: Welfare Effects of Training Subsidies on Incumbent Workers



Notes: The figure shows percentage changes in the discounted welfare of workers present at the start of the demographic transition, comparing a permanent training subsidy with the otherwise identical no-subsidy demographic transition. Workers are grouped by their skill type at the initial date, $i \in \{Li, Le, Hi, He\}$; their subsequent employment and skill transitions enter their continuation values. Welfare includes the present value of lump-sum taxes financing the subsidy. The horizontal axis varies the subsidy rate, and the vertical marker identifies the calibrated 10% policy.

present at $t = 0$ and workers entering at date t, respectively, these changes are

$$\Delta_i^I(\sigma_{tr}) = 100\,\frac{\overline{V}_{i,0}(\sigma_{tr}) - \overline{V}_{i,0}(0)}{\overline{V}_{i,0}(0)},$$
$$\Delta_{i,t}^E(\sigma_{tr}) = 100\,\frac{V_{i,t}^e(\sigma_{tr}) - V_{i,t}^e(0)}{V_{i,t}^e(0)}.$$

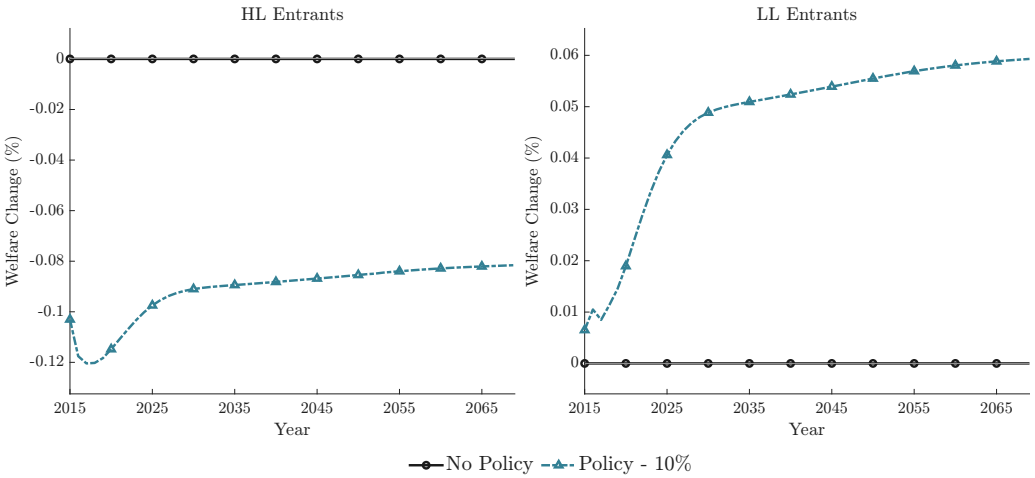
Incumbent comparisons use the same initial distribution of workers across employment states, while entrant comparisons evaluate each cohort's welfare at its own entry date.

Figure 12 shows that, among incumbents, workers without technological skills (Le and Li) gain the most. Greater access to training raises their probability of acquiring technological skills, and with it their value in the labor market. At the calibrated subsidy rate of 10%, their welfare rises by 0.074% and 0.098%, respectively. These gains are small because the subsidy raises equilibrium training only slightly, so few incumbents actually benefit.

For both groups, welfare is hump-shaped in the subsidy rate, reflecting two opposing forces. A higher subsidy raises their chances of acquiring technological skills and therefore their share of aggregate consumption, but the resource cost of additional training, financed through taxes, lowers aggregate consumption in discounted terms. These workers thus receive a larger share of a smaller pie: at modest subsidy rates the gain in their share dominates, while at high rates the contraction in aggregate resources prevails.

Tech-proficient incumbents (He and Hi), by contrast, lose at every positive subsidy rate. The expansion of technological skills erodes their scarcity and the earnings advantage it brings. They also do not benefit from training, yet contribute to financing it through taxes. At the calibrated subsidy rate, the welfare of both groups falls by about 0.10%.

Figure 13: Welfare Effects of Training Subsidies on Future Entrants



Notes: The figure reports the percentage effect of the permanent 10% training subsidy on lifetime welfare evaluated at labor-market entry, separately for cohorts arriving with high and low technological skills. Each cohort initially enters unemployment without experience; the high-tech and low-tech groups correspond to Hi and Li, respectively. Each plotted comparison is between the subsidized and unsubsidized demographic transitions at the same entry date. Expected future wages, employment and skill transitions, and the present value of subsidy-financing taxes are included in the welfare calculation. The horizontal axis gives the calendar year of entry.

Figure 13 reports the effects for successive cohorts of entrants under the calibrated subsidy. Tech-proficient entrants lose in every cohort, while entrants without technological skills gain. These gains grow for later cohorts, approaching 0.06% at the end of the horizon shown. Later cohorts enter after aggregate consumption has recovered from its initial decline and remains permanently above its level without the policy. They therefore avoid the adjustment costs borne by earlier cohorts, although they still help finance the policy throughout their working lives.

At the calibrated rate, the subsidy therefore redistributes welfare from current and future tech-proficient workers to workers without technological skills. Among incumbents, inexperienced workers without these skills (Li) gain the most, followed by experienced ones (Le). Despite its aggregate welfare cost, the subsidy may therefore be politically attractive in an aging economy. Its beneficiaries include experienced workers without technological skills, who are older on average. As the population ages, older workers account for a growing share of voters, so securing their support becomes increasingly important for winning elections. This mechanism is reminiscent of the political economy of pension systems, in which the rising electoral weight of older voters sustains transfers in their favor even when these transfers are inefficient (Galasso and Profeta, 2002). Training subsidies can be viewed through the same lens: they transfer resources toward older incumbents, at the expense of tech-proficient workers, who are younger on average. This view is also consistent with one of the policy's stated objectives: helping incumbent workers adapt to technological and structural change and preserve their employment prospects (Bundesregierung, 2018).

7 Conclusions

Workforces in advanced economies are aging rapidly. Because younger workers are more likely to have the up-to-date skills that new technologies require, this shift has raised a growing policy concern: an older workforce may leave firms short of the skills needed to operate at the technological frontier. Yet evidence on the size of the resulting productivity losses has been limited. This paper develops a methodology to measure them, combining firm-level evidence on technology adoption with an equilibrium model of the labor market.

Using German administrative data linked to surveys of technology adoption, we show that younger workers complement new technologies: establishments with younger workforces adopt them more often, and after adopting, they hire disproportionately more young workers, whose wages rise relative to those of older colleagues. We build an equilibrium model of technology adoption with a frictional labor market, in which younger workers are more technologically proficient and older workers more experienced, and discipline it with the estimated responses of employment, workforce composition, and wages to adoption.

The calibrated model predicts that Germany's projected aging reduces the share of frontier firms by 6.5%. The resulting shortage of technological skills lowers aggregate productivity by 0.5%, enough to completely offset the gains from a more experienced workforce predicted by standard theories of human capital. A model without the interaction between technological skills and adoption would therefore wrongly conclude that aging raises productivity.

These results point to a role for policies that sustain the supply of technological skills. Germany's recent expansion of training subsidies attenuates the decline in frontier adoption and raises long-run productivity by 0.8%. It does not raise aggregate welfare, but it redistributes welfare from tech-proficient workers, whose comparative advantage erodes, to workers who acquire technological skills through training. Because the beneficiaries include older workers, whose share of voters rises as the population ages, such policies may be politically attractive in aging societies.

Our analysis abstracts from channels that could amplify or offset these results. Endogenous growth, through which younger workers also contribute new ideas, would likely magnify the losses from technological skill shortages. Conversely, aging may encourage the adoption of labor-saving technologies, such as automation, that substitute for scarce workers. Our framework can also evaluate other responses to aging, such as raising the statutory retirement age, which expands labor supply but further ages the workforce and could therefore reinforce the decline in adoption we document. Studying these channels and policies is a promising direction for future research.

References

Acemoglu, D. and P. Restrepo (2017). Secular stagnation? the effect of aging on economic growth in the age of automation. *American Economic Review 107*(5), 174–179.

Acemoglu, D. and P. Restrepo (2020). Robots and jobs: Evidence from us labor markets. *Journal of political economy 128*(6), 2188–2244.

Acemoglu, D. and P. Restrepo (2022). Demographics and automation. *The Review of Economic Studies 89*(1), 1–44.

Adão, R., M. Beraja, and N. Pandalai-Nayar (2024). Fast and slow technological transitions. *Journal of Political Economy Macroeconomics 2*(2), 183–227.

Akerman, A., I. Gaarder, and M. Mogstad (2015). The skill complementarity of broadband internet. *The Quarterly Journal of Economics 130*(4), 1781–1824.

Aksoy, Y., H. S. Basso, R. P. Smith, and T. Grasl (2019). Demographic structure and macroeconomic trends. *American Economic Journal: Macroeconomics 11*(1), 193–222.

Angelini, D. (2023). Aging population and technology adoption. Technical report, Department of Economics, University of Konstanz.

Antonczyk, D., T. DeLeire, and B. Fitzenberger (2018). Polarization and rising wage inequality: comparing the us and germany. *Econometrics 6*(2), 20.

Atkeson, A. and P. J. Kehoe (2007). Modeling the transition to a new economy: lessons from two technological revolutions. *American Economic Review 97*(1), 64–88.

Aubert, P., E. Caroli, and M. Roger (2006). New technologies, organisation and age: firm-level evidence. *The Economic Journal 116*(509), F73–F93.

Australian Treasury (2023, August). 2023 Intergenerational Report. Technical report, Australian Government.

Barth, E., J. C. Davis, R. B. Freeman, and K. McElheran (2023). Twisting the demand curve: Digitalization and the older workforce. *Journal of Econometrics 233*(2), 443–467.

Behrens, V., M. Berger, M. Hud, P. Hünermund, Y. Iferd, B. Peters, C. Rammer, and T. Schubert (2017). Innovation activities of firms in germany-results of the german cis 2012 and 2014: Background report on the surveys of the mannheim innovation panel conducted in the years 2013 to 2016. Technical report, ZEW-Dokumentation.

Bellmann, L., H.-D. Gerner, and R. Upward (2018). Job and worker turnover in german establishments. *The Manchester School 86*(4), 417–445.

Ben-Porath, Y. (1967). The production of human capital and the life cycle of earnings. *Journal of political economy 75*(4, Part 1), 352–365.

Bianchi, N., G. Bovini, J. Li, M. Paradisi, and M. Powell (2023). Career spillovers in internal labour markets. *The Review of Economic Studies 90*(4), 1800–1831.

Bianchi, N. and M. Paradisi (2022). Countries for old men: An analysis of the age wage gap. *Available at SSRN 3880501*.

Bilger, F., F. Behringer, H. Kuper, and J. Schrader (Eds.) (2017). *Weiterbildungsverhalten in Deutschland 2016: Ergebnisse des Adult Education Survey (AES)*. Bielefeld: wbv Media.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The quarterly journal of economics 117*(1), 339–376.

Brücker, H., N. Rother, J. Schupp, C. Babka von Gostomski, A. Böhm, T. Fendel, M. Friedrich, M. Giesselmann, Y. Kosyakova, M. Kroh, et al. (2016). Flucht, ankunft in deutschland und erste schritte der integration. *DIW Wochenbericht 83*(46), 1103–1119.

Bundesregierung (2018, October). Entwurf eines gesetzes zur stärkung der chancen für qualifizierung und für mehr schutz in der arbeitslosenversicherung (Qualifizierungschancengesetz). Drucksache 19/4948, Deutscher Bundestag.

Carrillo-Tudela, C., H. Gartner, and L. Kaas (2023). Recruitment policies, job-filling rates, and matching efficiency. *Journal of the European Economic Association 21*(6), 2413–2459.

Chari, V. V. and H. Hopenhayn (1991). Vintage human capital, growth, and the diffusion of new technology. *Journal of political Economy 99*(6), 1142–1165.

Council for Science, Technology and Innovation (2015, December). Report on the 5th Science and Technology Basic Plan. Technical report, Cabinet Office, Government of Japan.

CSN (2026). Vad är omställningsstudiestöd? `https://www.csn.se/bidrag-och-lan/omstallningsstudiestod-for-dig-mitt-i-arbetslivet/vad-ar-omstallningsstudiestod.html`. Updated September 15, 2026. Accessed September 23, 2026.

Deming, D. J. and K. Noray (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics 135*(4), 1965–2005.

Deng, L., S. Müller, V. Plümpe, and J. Stegmaier (2024). Robots, occupations, and worker age: A production-unit analysis of employment. *European Economic Review 170*, 104881.

Deutsche Rentenversicherung Bund (2025, October). Rentenversicherung in Zeitreihen. DRV-Schriften 22, Deutsche Rentenversicherung Bund, Berlin. Ausgabe 2025.

Engbom, N. (2020). Misallocative growth. *Working Paper*.

Eurofound (2025). Keeping older workers in the labour force. Technical report, Publications Office of the European Union, Luxembourg.

European Commission (2024, March). Labour and skills shortages in the EU: An action plan. Communication COM(2024) 131 final, European Commission.

Federal Employment Agency (n.d.). Individuelle förderung von beschäftigten. Accessed September 23, 2026.

Federal Statistical Office (Destatis) (2026). Population, persons in employment, unemployed persons, economically active population, economically inactive population: Germany, reference month/year, age groups. Microcensus, GENESIS-Online, Table 12211-9001.

Galasso, V. and P. Profeta (2002). The political economy of social security: a survey. *European Journal of Political Economy 18*(1), 1–29.

Gathmann, C. and U. Schönberg (2010). How general is human capital? a task-based approach. *Journal of Labor Economics 28*(1), 1–49.

Genz, S., T. Gregory, M. Janser, F. Lehmer, and B. Matthes (2021). How do workers adjust when firms adopt new technologies? IZA Discussion Paper 14626, Institute of Labor Economics (IZA), Bonn.

German Federal Government (2018). Entwurf eines gesetzes zur stärkung der chancen für qualifizierung und für mehr schutz in der arbeitslosenversicherung (Qualifizierungschancengesetz). Bundestagsdrucksache 19/4948, Deutscher Bundestag.

Graf, T., S. Grießemer, M. Köhler, C. Lehnert, A. Moczall, M. Oertel, A. Schmucker, A. Schneider, S. Seth, U. Thomsen, and P. vom Berge (2023). Weakly anonymous version of the sample of integrated labour market biographies (SIAB)—version 7521 v1. Research Data Centre of the Federal Employment Agency (BA) at the Institute for Employment Research (IAB).

Hartung, B., P. Jung, and M. Kuhn (2026). Unemployment insurance reforms and labour market dynamics. *Review of Economic Studies 93*(1), 517–555.

Hopenhayn, H., J. Neira, and R. Singhania (2022). From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share. *Econometrica 90*(4), 1879–1914.

Inokuma, H. and J. M. Sanchez (2023). *From Population Growth to TFP Growth*. Federal Reserve Bank of St. Louis, Research Division.

Janssen, S. and O. Wölfel (2017). Weiterbildung in der informations- und kommunikationstechnologie: Jüngere belegen inhaltlich andere kurse als ältere. IAB-Kurzbericht 17/2017, Institut für Arbeitsmarkt- und Berufsforschung, Nuremberg.

Jeong, H., Y. Kim, and I. Manovskii (2015). The price of experience. *American Economic Review 105*(2), 784–815.

Jones, C. I. (2022). The end of economic growth? unintended consequences of a declining population. *American Economic Review 112*(11), 3489–3527.

Kaas, L. and P. Kircher (2015). Efficient firm dynamics in a frictional labor market. *American Economic Review 105*(10), 3030–3060.

Kalyani, A. (2024). The creativity decline: Evidence from us patents.

Karahan, F., B. Pugsley, and A. Şahin (2024). Demographic origins of the start-up deficit. *American Economic Review 114*(7), 1986–2023.

Katz, L. F. and K. M. Murphy (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *Quarterly Journal of Economics 107*(1), 35–78.

Krebs, T. and M. Scheffel (2013). Macroeconomic evaluation of labor market reform in germany. *IMF Economic Review 61*(4), 664–701.

Kredler, M. (2014). Experience vs. obsolescence: A vintage-human-capital model. *Journal of Economic Theory 150*, 709–739.

Liang, J., H. Wang, and E. P. Lazear (2018). Demographics and entrepreneurship. *Journal of Political Economy 126*(S1), S140–S196.

Lipowski, C. (2024). No teens, no tech: How shortages of young workers hinder firm technology investments. *ZEW-Centre for European Economic Research Discussion Paper* (24-062).

Lochner, B. and B. Schulz (2024). Firm productivity, wages, and sorting. *Journal of Labor Economics 42*(1), 85–119.

Mackeben, J., K. Ruf, S. Wolter, P. Grunau, T. Graf, S. Grießemer, M. Köhler, C. Lehnert, A. Moczall, M. Oertel, A. Schneider, and U. Thomsen (2023). LPP: Linked personnel panel survey data linked with administrative data of the IAB (LPP-ADIAB 7521) 1975–2021. Research Data Centre of the Federal Employment Agency (BA) at the Institute for Employment Research (IAB).

Maestas, N., K. J. Mullen, and D. Powell (2023). The effect of population aging on economic growth, the labor force, and productivity. *American Economic Journal: Macroeconomics 15*(2), 306–332.

Menzio, G. and S. Shi (2011). Efficient search on the job and the business cycle. *Journal of Political Economy 119*(3), 468–510.

Ministero del Lavoro e delle Politiche Sociali (2024, December). Competenze per le innovazioni, pubblicato l'avviso Fondo nuove competenze 3. `https://lavoro.gov.it/pn-giovani-donne-lavoro/comunicazione/notizie/pubblicato-lavviso-fondo-nuove-competenze-3-competenze-per-le-innovazioni`. Published December 5, 2024. Accessed September 23, 2026.

OECD (2025). Staying in the game: Skills and jobs of older workers in a changing labour market. In *OECD Employment Outlook 2025*. Paris: OECD Publishing.

Panahian Fard, D., A. Schmucker, S. Seth, M. Umkehrer, and F. Zimmermann (2024). Linked-employer-employee-data of the IAB: LIAB longitudinal model (LIAB LM) 1975–2021. FDZ-Datenreport 04/2024 (en), Institute for Employment Research (IAB), Nuremberg.

Peters, M. and C. Walsh (2021). Population growth and firm dynamics. Technical report, National Bureau of Economic Research.

Petrongolo, B. and C. A. Pissarides (2001). Looking into the black box: A survey of the matching function. *Journal of Economic literature 39*(2), 390–431.

Pijoan-Mas, J. and P. Roldan-Blanco (2024). Dual labor markets and the equilibrium distribution of firms.

Schaal, E. (2017). Uncertainty and unemployment. *Econometrica 85*(6), 1675–1721.

Schröpf, B. (2023). The dynamics of wage dispersion between firms: the role of firm entry and exit. *Journal for Labour Market Research 57*(1), 1.

Shorrocks, A. F. (2013). Decomposition procedures for distributional analysis: A unified framework based on the Shapley value. *Journal of Economic Inequality 11*(1), 99–126.

Statistics of the Federal Employment Agency (2026, August). Förderung der beruflichen weiterbildung: Beschäftigtenqualifizierung. Technical report, Bundesagentur für Arbeit, Nuremberg. Reference month: May 2026. Released August 25, 2026.

Statistisches Bundesamt. Population by age: Germany. GENESIS-Online, table 12411-0005. Resident population by single year of age and sex.

Statistisches Bundesamt (2025). 16th coordinated population projection. GENESIS-Online, table 12421-0002. Model Calculation 1 (G2L2W0), base date 31 December 2024.

Vandenbroucke, G. (2021). The baby boomers and the productivity slowdown. *European Economic Review 132*, 103609.

- SUPPLEMENTARY APPENDICES -

- **Appendix A**: Empirical Appendix ..... p. A2
- **Appendix B**: Model Appendix ..... p. A12
- **Appendix C**: Numerical Appendix ..... p. A23
- **Appendix D**: Static Model ..... p. A24
- **Appendix E**: Additional Figures and Tables ..... p. A30

A Empirical Appendix

A.1 Descriptive Statistics: LPP Survey

Table A1: Establishment Characteristics: LPP Versus IAB Establishment Panel

	LPP			IAB-BP		
	Mean	Median	SD	Mean	Median	SD
Log full-time employment	4.20	4.21	1.08	1.39	1.10	1.32
Log real VA	15.71	15.52	1.19	12.44	12.29	1.47
Log real VA per worker	10.99	10.96	0.82	10.48	10.52	0.89
Age of establishment	28.46	26.50	13.92	16.63	12.50	14.01
Average age of full-time employees	44.08	44.04	4.66	42.68	42.84	9.26
Qualified employees (%)	73.64	76.96	16.20	71.61	80.00	30.27
Employees under 35 (%)	28.86	27.46	14.43	29.09	25.00	28.22
Manufacturing (%)	30.60	—	—	9.13		
Services (%)	65.55	—	—	77.24		
Construction (%)	3.85	—	—	10.40		
Utilities & mining	Excluded	—	—	0.80		
Agriculture	Excluded	—	—	2.44		
N. establishments (average per wave)	737.0	—	—	17776.0		

Notes: The table compares establishments in the LPP-ADIAB survey waves with establishments in the IAB Establishment Panel, using the respective survey weights. Entries are arithmetic averages of the 2019 and 2021 wave-specific summary statistics, including the medians and standard deviations, rather than pooled sample statistics. Employment and value added are reported in logs; workforce shares and industry shares are percentages. The LPP samples establishments with at least 50 employees, whereas the IAB Establishment Panel represents the broader establishment population. Valid observation counts may differ by variable.

Table A2: Technology Use and Recent Adoption in the LPP

Technology	Use before (%)	Adopt (%)	Not use (%)	N valid
Digital				
Digital distribution channels	72.95	8.11	18.94	730.0
Big-data analytics	12.97	7.03	80.00	722.5
Internet of Things	9.64	4.65	85.71	721.0
Cyber-physical systems	14.56	4.12	81.31	724.5
Artificial intelligence (machine learning)	-	-	-	-
Automation				
Additive manufacturing	5.71	4.05	90.24	723.5
Robots	5.79	2.46	91.75	724.0
Automated transport systems	3.12	2.31	94.57	725.5
Survey N (average per wave)				737.0

Notes: The table describes use and recent adoption of technologies in the LPP establishment survey, averaged across its two technology waves. "Use before" indicates technologies already used before the two-year reference window; "Adopt" indicates introduction during the preceding two years; "Not use" indicates no current use. These three categories are mutually exclusive for each technology. Percentages use the LPP establishment survey weights, while valid-response counts are unweighted. Statistics are averages of wave-specific figures, not pooled estimates; unavailable or suppressed technology items do not have comparable two-wave statistics.

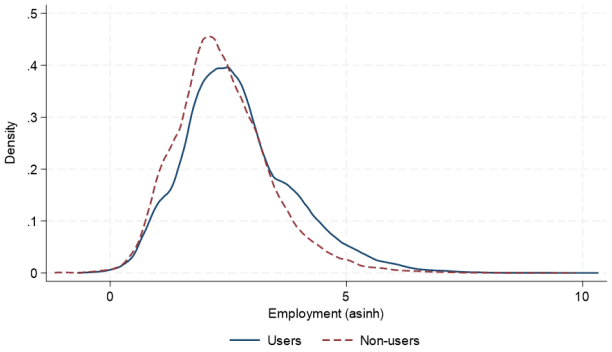
A.2 Digitalization Supplement: IAB-BP 2017

Table A3: Digital-Technology Use by Establishment Age

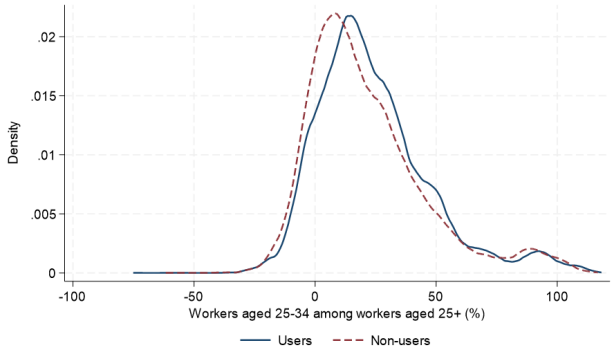
Technology	All establishments	Age 0–5	Age 6–15	Age 16+	N.Obs
Digital Procurement	19.29	19.52	18.55	18.72	10668
Digital Sales	22.52	24.01	21.62	21.31	10762
Internet of Things	6.86	5.90	6.03	7.36	12510
Digital (Any)	32.06	33.24	30.45	31.33	11414

Notes: The table reports establishment-weighted rates of digital-technology use by establishment age in the 2017 IAB Establishment Panel Digitalization Survey supplement. Technologies comprise digital procurement, digital sales, and internet-of-things/networked data exchange; "Digital (Any)" denotes use of digital procurement or sales. Entries are percentages obtained with establishment survey weights. Age categories aggregate the available establishment-age bins, and the all-establishments column includes observations without a valid establishment-age category. Sample counts vary with item response.

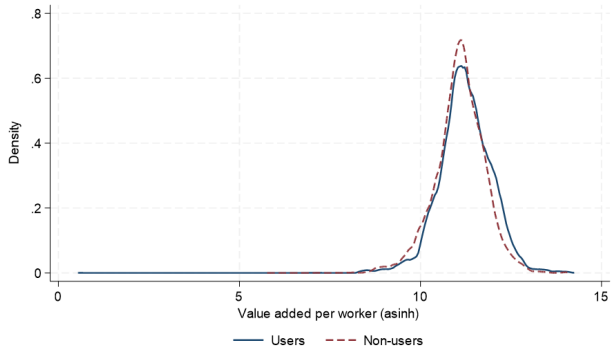
Figure A1: Characteristics of Digital-Technology Users and Nonusers



(a) Employment



(b) Share of Young



(c) Value added per worker

Notes: The figure compares the distributions of log establishment employment, the share of employees younger than 35, and log value added per worker between users and nonusers of digital technologies. Digital-use status is measured using the 2017 Digitalization Survey supplement to the IAB Establishment Panel, linked to establishment administrative outcomes. The curves are smoothed empirical distributions for the respective groups. The comparison is descriptive and does not condition on establishment characteristics; value added is expressed per worker.

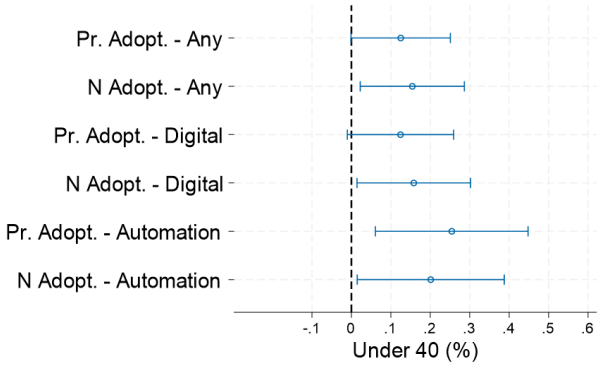
A.3 Cross Sectional Regressions

Table A4: Workforce Age Composition and Technology Adoption Regressions

	Any technology			Digital technology			Automation technology		
	(1) Adopt	(2) Adopt	(3) Adopt	(4) Adopt	(5) Adopt	(6) Adopt	(7) Adopt	(8) Adopt	(9) Adopt
Under-35 share (std.)	0.058***	0.048**	0.067***	0.051***	0.041**	0.055**	0.026**	0.028**	0.025
	(0.020)	(0.023)	(0.025)	(0.017)	(0.020)	(0.022)	(0.012)	(0.014)	(0.017)
asinh Full-time employment	0.053**	0.055**	0.049**	0.046**	0.048**	0.057**	0.034***	0.035***	0.027*
	(0.021)	(0.022)	(0.024)	(0.019)	(0.020)	(0.023)	(0.011)	(0.012)	(0.016)
asinh Avg full-time wage	0.228***	0.248***	0.335***	0.227***	0.240***	0.277***	0.053	0.058	0.131**
	(0.076)	(0.082)	(0.094)	(0.068)	(0.072)	(0.085)	(0.036)	(0.040)	(0.054)
asinh Firm age		0.017	−0.017		0.019	−0.015		0.006	−0.002
		(0.026)	(0.028)		(0.024)	(0.026)		(0.013)	(0.017)
Qualified-worker share (std.)		−0.030	−0.031		−0.024	−0.027		0.001	−0.003
		(0.022)	(0.028)		(0.021)	(0.026)		(0.011)	(0.015)
Inflow rate (std.)		0.021	−0.014		0.023	−0.010		0.003	−0.001
		(0.028)	(0.032)		(0.026)	(0.029)		(0.014)	(0.018)
asinh Investment per worker			0.034***			0.029***			0.008**
			(0.007)			(0.006)			(0.004)
asinh Sales per worker			−0.039			−0.008			−0.016
			(0.033)			(0.029)			(0.020)
Fixed effects									
Constant	✓	✓	✓	✓	✓	✓	✓	✓	✓
Year	✓	✓	✓	✓	✓	✓	✓	✓	✓
Sector 3-digit	✓	✓	✓	✓	✓	✓	✓	✓	✓
State	✓	✓	✓	✓	✓	✓	✓	✓	✓
Technologies used before	✓	✓	✓	✓	✓	✓	✓	✓	✓
Owner type		✓	✓		✓	✓		✓	✓
Change management (2 years)			✓			✓			✓
Mean Outcome (weighted)	0.254	0.254	0.268	0.216	0.216	0.227	0.076	0.077	0.081
N. Observations	1,387	1,365	1,068	1,388	1,366	1,069	1,381	1,359	1,065
Adjusted R^2	0.116	0.114	0.174	0.110	0.107	0.156	0.115	0.112	0.122

Notes: The table reports linear probability regressions of technology adoption during the two years preceding the LPP survey. The outcome is an adoption indicator, separately for any, digital, and automation technologies. The regressor of interest is the standardized lagged share of workers under 35; other controls and fixed effects vary by column as indicated. Regressions use LPP survey weights and standard errors clustered by establishment (in parentheses). Share variables and inflow rates are standardized; selected levels are transformed using the inverse hyperbolic sine. The weighted sample mean of the dependent variable and observations are shown below. Significance stars are defined by the conventional 10%, 5%, and 1% levels.

Figure A2: Technology Adoption and the Share of Workers Under 40



Notes: This robustness exercise repeats the cross-sectional adoption regressions while replacing the baseline share of employees under 35 with the share under 40. Outcomes distinguish the probability of adopting any, digital, and automation technologies and the number of technologies adopted among adopting establishments. Adoption is defined over the two years preceding the LPP survey. The regressions use lagged establishment controls and industry, state, survey-year, and existing-technology fixed effects, as in the baseline analysis. Horizontal bars depict 95% confidence intervals.

Establishments with younger workforce adopt technologies at higher rates

A.4 Event Study

Table A5: Pre-Adoption Balance Between Adopters and Matched Controls

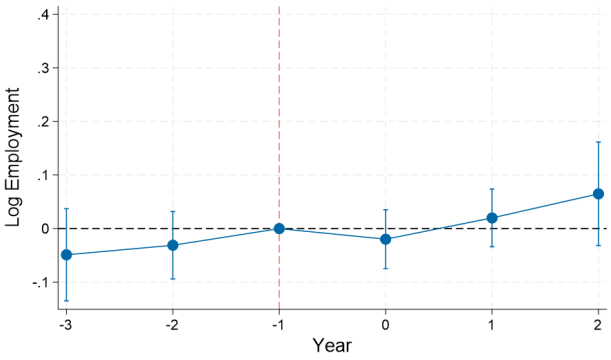
	N		Weighted mean		
Variable	Adopters	Non-adopters	Adopters	Non-adopters	SMD_W
asinh Employment	222	539	5.44	5.39	0.064
asinh Sales per worker	161	377	12.54	12.57	-0.028
Log real VA (t-1)	136	313	15.82	15.88	-0.053
Log real VA per worker (t-1)	136	313	11.06	11.20	-0.175
Age of establishment	222	539	25.98	26.29	-0.024
Avg age of employees	222	539	42.65	42.53	0.028
Qualified employees (%)	222	539	75.73	74.86	0.062
Employees under 35 (%)	222	539	31.61	31.48	0.010
Technologies used before	222	539	0.55	0.74	-0.396

Notes: The table compares pre-event characteristics of matched adopting and non-adopting establishments in the baseline LPP event-study sample at $k = -1$. Matching cells are formed by event year, two-digit industry, pre-event employment-size bins, and pre-event young-worker-share bins. Weighted means use the final event-study matching weights; controls may appear in more than one event-year cell. The reported standardized mean difference scales the weighted treated–control difference by the original unweighted pooled standard deviation. Counts are establishment–event observations, and outcome-specific sample sizes differ when responses are missing.

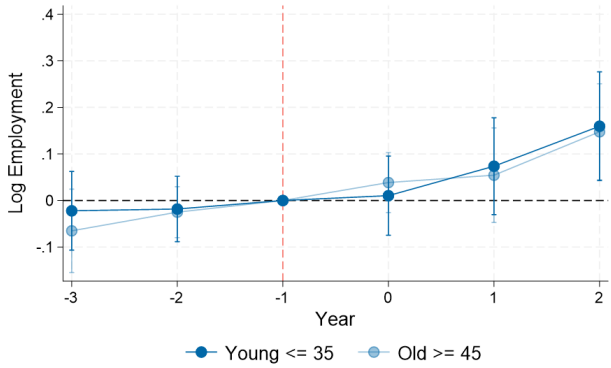
Alternative Employment Shock

We use the same sample of firms surveyed in the LPP-ADIAB, but restricting to years prior the technology adoption survey. First, we residualize annual establishment employment growth with

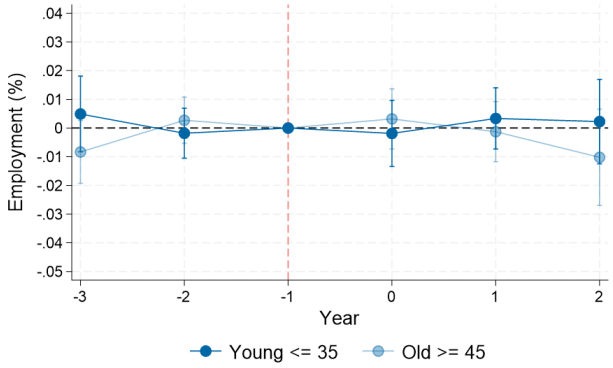
Figure A3: Additional Employment and Age-Composition Event Studies



(a) Employment response to technology adoption



(b) Employment responses to technology adoption, by worker age



(c) Workforce age-composition response to technology adoption

Notes: The three panels present establishment-level event-study responses for the adoption sample used in this appendix, using non-adopting establishments matched on industry, employment, and young-worker share before the event. The outcomes are total log employment, log employment by age group, and the employment shares of young and old workers. Coefficients are measured relative to $k = -1$, and vertical bars show 95% confidence intervals clustered at the establishment level.

respect to 3-digit industry and German-state trends. Then, we define a general large employment-expansion event event as a year in which an establishment's residualized employment growth is above the 75th percentile of the corresponding distribution. Subsequently, using the same matching procedure detailed above, we create a matched sample of establishment experiencing an employment expansion and establishments that do not.[1]

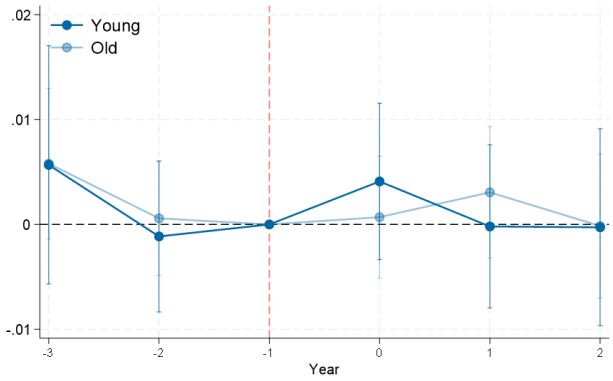
We estimate an event-study specification analogous to the one used around technology adoption:

$$y_{ft} = \alpha_f + \gamma_t + \sum_k \theta_k D^k_{ft} + \sum_k \beta_k (D^k_{ft} \times \text{Employment Expansion}_f) + \varepsilon_{ft},$$

where event time is now defined relative to the large employment-expansion event.

[1](As before, the matching variables are 2-digit industry, same bin of lagged employment and workforce age composition.

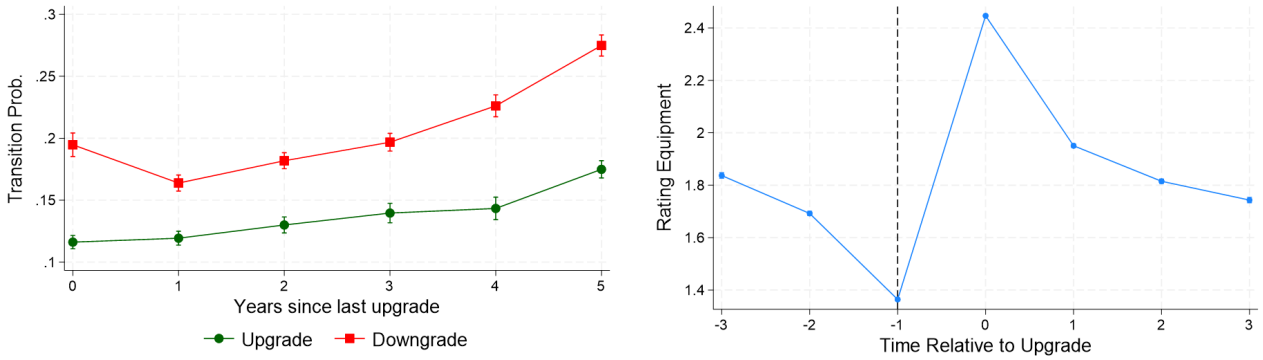
Figure A4: Incumbent Days Worked Around Technology Adoption by Age



Notes: The figure shows matched event-study estimates for incumbent workers' total days employed during the year, separately for younger and older workers. The sample consists of workers employed before their establishment's technology-adoption event, with age groups defined from pre-event age. The omitted reference period is $k = -1$ and event times are assigned using the survey adoption window. Vertical bars show 95% confidence intervals clustered by establishment.

A.5 State of the Equipment

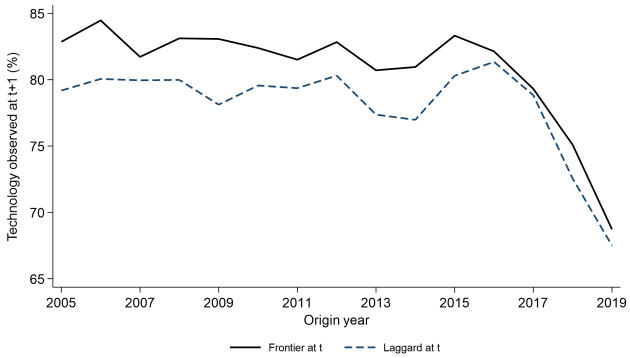
Figure A5: Technology-State Transitions and Equipment Ratings



(a) Hazard Transition **(b)** Cyclical Rating

Notes: The panels examine how technology-state transitions and equipment ratings evolve around an observed upgrade in the IAB Establishment Panel. Panel (a) plots the probability of an upgrade or downgrade by elapsed time since the latest upgrade. Panel (b) plots the reported equipment rating before and after an upgrade, where higher ratings indicate more technologically advanced equipment relative to industry peers. The technology classification follows the distinction between frontier ratings (categories 4–5) and standard ratings (categories 1–3).

Figure A6: Subsequent-Year Survey Response Rates by Technology State



Notes: The figure reports, by origin year, the fraction of establishments with a valid equipment-technology response in year $t + 1$ among establishments observed with a valid response in year t, separately by the technology category reported in t. The categories are frontier (equipment ratings 4–5) and standard/laggard (ratings 1–3) in the IAB Establishment Panel. The graph describes next-year response availability for the sample used to estimate technology-state transitions.

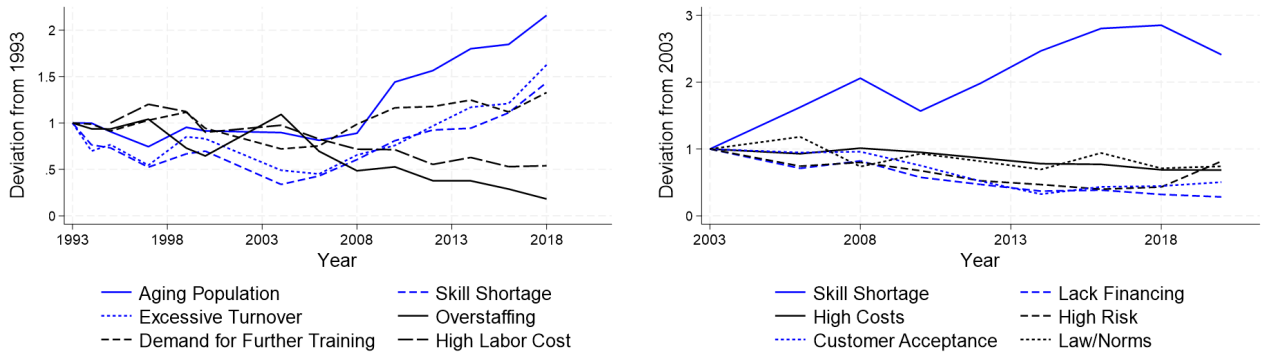
Table A6: Annual Transition Probabilities Between Technology States

	Origin state	
Destination	Frontier	Laggard
Frontier	83.7813	28.7311
	82.9874	29.6624
Laggard	16.2187	71.2689
	17.0126	70.3376

Notes: Entries are average annual transition probabilities between technology states inferred from the IAB Establishment Panel equipment-rating question. "Frontier" denotes ratings 4–5 and "Laggard" ratings 1–3. Origin states are columns and destination states rows. For each destination, the first line is the unweighted probability and the second is weighted using the origin-year establishment survey weight. Probabilities average transitions for origin years 2005–2019 among surviving establishments with valid responses in both years. Each column sums to 100% within each weighting scheme.

A.6 Motivational Trends

Figure A7: Trends in Staffing Constraints and Barriers to Investment



(a) Staffing-Related Issues

(b) Reasons Preventing Investment

Notes: The panels display survey-reported constraints related to staffing and technology investment over time. Panel (a) covers workforce aging, skill shortages, excessive turnover, overstaffing, demand for additional training, and high labor costs; panel (b) covers skill shortages, access to financing, costs, risk, customer acceptance, and legal or regulatory barriers. Series are presented relative to the initial year displayed in each panel.

A.7 Construction of the Demographic Path

This appendix describes how we anchor the demographic experiment to the calibrated economy and convert the empirical entrant-share series into the model's quarterly path of g_t.

Anchoring to the calibrated steady state. The transition begins in 2016 from the steady state calibrated to 2005–2015. Let g_0 denote the calibrated quarterly growth intensity and let e_0 denote the corresponding cumulative annual inflow of workers relative to the labor force at the beginning of the year. We normalize the initial labor force to $\mathcal{L}_{2016} = 1$. To preserve the calibrated starting point, we use the empirical series only to determine the relative evolution of entry:

$$e_t^{\text{target}} = e_0 \frac{\widehat{e}_t}{\widehat{e}_{2016}}. \tag{A1}$$

Thus, the entry target in 2016 equals e_0, and subsequent changes follow the empirical entrant-share series. The baseline value must satisfy $e_0 = \Phi(g_0)$, where Φ is the annual mapping defined below. If baseline growth is reported as an annual discrete rate, its quarterly continuous counterpart is $g_0 = \log(1 + g_0^{\text{ann}})/4$.

Construction of the empirical entrant share. We calculate $\widehat{e}_t$ using equation (5.1). The population aged exactly 25 at the end of each year serves as a proxy for the cohort reaching the model's entry age during that year. Multiplying this stock by the fixed participation rate for ages 25–29 gives the entry proxy. We construct the denominator by applying fixed participation rates to population stocks in the age groups $25\text{–}29, 30\text{–}34, \ldots, 60\text{–}64, 65+$. Each participation rate is the simple average of the corresponding annual rates over the calibration period:

$$\bar{\pi}_a = \frac{1}{11} \sum_{s=2005}^{2015} \frac{LF_{a,s}}{P_{a,s}}, \tag{A2}$$

where $LF_{a,s}$ denotes the labor force in age group a and year s.

Population stocks are observed for 2015–2025 and taken from Model Calculation 1 of the 16th Coordinated Population Projection thereafter. The data sources and projection assumptions are reported in Section 5.1. The zero-net-migration assumption applies to the projected period; observed population stocks retain the effects of historical migration. Moreover, the empirical denominator reflects the aging of existing cohorts. The construction therefore targets the evolution of entry relative to the labor force without imposing the projected aggregate growth rate on the model.

The empirical entry proxy combines population stocks and fixed participation rates and need not coincide in levels with the model's cumulative annual entry rate. The normalization in equation (A1) accommodates this difference while retaining the relative changes in the empirical series.

Recovering quarterly growth rates. One calendar year contains $Q = 4$ quarterly model periods. Let g_t^q denote the quarterly continuous growth rate applied throughout calendar year t; this is the parameter denoted by g_t in the model. For within-year time $s \in [0, Q]$, the labor force satisfies

$$\frac{d\mathcal{L}_t(s)}{ds} = \mu_t(s) - \lambda_r \mathcal{L}_t(s) = g_t^q \mathcal{L}_t(s). \tag{A3}$$

Writing $\mathcal{L}_t$ for the labor force at the beginning of the year gives

$$\mathcal{L}_t(s) = \mathcal{L}_t \exp(g_t^q s), \qquad \mu_t(s) = (g_t^q + \lambda_r)\mathcal{L}_t(s). \tag{A4}$$

The cumulative annual entrant flow relative to the beginning-of-year labor force is therefore

$$\frac{1}{\mathcal{L}_t}\int_0^Q \mu_t(s)\, ds = \Phi(g_t^q), \qquad \Phi(x) = \begin{cases} (x+\lambda_r)\dfrac{\exp(Qx)-1}{x}, & x \neq 0, \\ Q\lambda_r, & x = 0. \end{cases} \tag{A5}$$

This expression accounts for changes in the labor force within the year. For each year, we recover the quarterly growth rate as the solution to

$$\Phi(g_t^q) = e_t^{\text{target}}, \qquad g_t^q > -\lambda_r. \tag{A6}$$

The inequality ensures a positive entrant flow. Since Φ is strictly increasing on this domain, each positive target determines a unique quarterly rate. In particular, the normalization of the entry target guarantees $g_{2016}^q = g_0$.

The model labor force and the annual discrete growth rate then satisfy

$$\mathcal{L}_{t+1} = \mathcal{L}_t \exp(Q g_t^q), \qquad g_t^{\text{ann}} = \exp(Q g_t^q) - 1. \tag{A7}$$

Throughout this construction, the retirement intensity remains fixed at its calibrated value.

The long-run 2050s steady state. We define the terminal demographic parameter as the arithmetic mean of the quarterly continuous rates implied by the 2050s:

$$g^{2050} = \frac{1}{10}\sum_{t=2050}^{2059} g_t^q. \tag{A8}$$

Its annual discrete equivalent is

$$g^{2050,\text{ann}} = \exp(4g^{2050}) - 1. \tag{A9}$$

Using a baseline annual entrant share of $e_0 = 0.0218$ and a quarterly retirement intensity of $\lambda_r = 0.0064$, the empirical series implies an average target entrant share of approximately 1.881% during 2050–2059. Averaging the corresponding quarterly solutions gives $g^{2050} \simeq -0.0016818$, or approximately -0.6705% in annual discrete terms.

We set $g_t = g^{2050}$ from 2060 onward. This assumption fixes the demographic parameter while allowing the age distribution to converge endogenously. The long-run 2050s steady state is the balanced-growth equilibrium associated with this constant value of g and all other parameters fixed at their baseline calibration.

B Model Appendix

B.1 Proof of Proposition 1

Proof. Recall that the firm's physical state is $\hat{s} = (\vec{n}, z, \tau)$ and that $\vec{W} = \{W_i\}_{i \in G}$ denotes the outstanding promises to its incumbent workers. The full contract is

$$\mathcal{C} = \left\{ \left\{ w_i, v_i, \delta_i, \lambda_{tr,i}, \{W_i'(\hat{s}')\}_{\hat{s}' \in S} \right\}_{i \in G}, r \right\}.$$

For a given contract $\mathcal{C}$, let $W_i(\hat{s}; \mathcal{C})$ be the continuation value actually delivered to an incumbent worker of type i and let $J(\hat{s}, \vec{W}; \mathcal{C})$ be the firm's value. Feasibility requires promise keeping and participation:

$$W_i(\hat{s}; \mathcal{C}) \geq W_i, \qquad \text{for every incumbent type } i, \tag{A10}$$

$$W_j'(\hat{s}') \geq U_j, \qquad \text{for every worker of type } j \text{ remaining employed in } \hat{s}'. \tag{A11}$$

In particular, the second restriction applies to a worker's *new* type following a skill transition.

Binding promises. At an optimum, (A10) holds with equality. Otherwise, the firm could lower the current wage paid to workers of a type receiving strictly more than their outstanding promise, increasing firm value without changing any continuation promise or violating participation. Thus,

$$W_i(\hat{s}; \mathcal{C}) = W_i \qquad \text{whenever } n_i > 0. \tag{A12}$$

The worker's continuous-time Bellman equation is linear in the flow wage w_i. Hence, for any feasible outstanding promises and any chosen state-contingent continuation promises, equation (A12) uniquely determines the wage that implements the outstanding promise.

Aggregation of firm and worker values. Define the joint value associated with a feasible contract by

$$\widetilde{\Sigma}(\hat{s}, \vec{W}; \mathcal{C}) \equiv J(\hat{s}, \vec{W}; \mathcal{C}) + \sum_{i \in G} n_i W_i. \tag{A13}$$

Add the firm's Bellman equation to the Bellman equations of all its incumbent workers, with each type-i equation weighted by n_i. The current wage bill cancels with the sum of workers' wage receipts:

$$-\sum_{i \in G} n_i w_i + \sum_{i \in G} n_i w_i = 0.$$

Likewise, in every transition, the firm's post-transition value plus the continuation values of the workers who remain with it equals the joint value of the continuing coalition. We verify the terms

generated by each type of transition.

When the firm hires one worker of type i, it moves to $\hat{s}_i^+$. The promised continuation value $W_i'(\hat{s}_i^+)$ belongs to the new hire, who was not a member of the original coalition. Consequently, the hiring contribution to the joint Bellman equation is

$$\sum_{i \in G} v_i \eta\big(\theta_i(W_i'(\hat{s}_i^+))\big) \left[\widetilde{\Sigma}(\hat{s}_i^+) - \widetilde{\Sigma}(\hat{s}) - W_i'(\hat{s}_i^+)\right]. \tag{A14}$$

Here and below, the arguments $\vec{W}$ and $\mathcal{C}$ of joint values at successor states are suppressed to avoid clutter. Notice that no continuation promise to an incumbent worker remains in (A14).

If a type-i worker is fired or separates exogenously, that worker receives unemployment value U_i, and the remaining coalition moves to $\hat{s}_i^-$. The corresponding term is

$$\sum_{i \in G} n_i(\delta_i + s_i^W) \left[\widetilde{\Sigma}(\hat{s}_i^-) + U_i - \widetilde{\Sigma}(\hat{s})\right]. \tag{A15}$$

Retirement also moves the firm to $\hat{s}_i^-$, but the retiring worker's continuation value is zero, by assumption. Its contribution is therefore

$$\lambda_r \sum_{i \in G} n_i \left[\widetilde{\Sigma}(\hat{s}_i^-) - \widetilde{\Sigma}(\hat{s})\right]. \tag{A16}$$

Now consider a type transition $i \to j$, which changes the physical state to $\hat{s}_{i \to j}$. Conditional on this event, one of the n_i workers changes type and the other $n_i - 1$ retain type i. Their aggregate continuation value after the transition is

$$(n_i - 1)W_i'(\hat{s}_{i \to j}) + W_j'(\hat{s}_{i \to j}).$$

Adding the continuation values of all other incumbent workers and the firm gives precisely $\widetilde{\Sigma}(\hat{s}_{i \to j})$. In particular, experience accumulation, technological-skill obsolescence, and on-the-job training contribute, respectively,

$$\lambda_E \sum_{i \in \{Li, Hi\}} n_i \left[\widetilde{\Sigma}(\hat{s}_{i \to e(i)}) - \widetilde{\Sigma}(\hat{s})\right], \tag{A17}$$

$$\lambda_T \sum_{i \in \{Hi, He\}} n_i \left[\widetilde{\Sigma}(\hat{s}_{i \to \ell(i)}) - \widetilde{\Sigma}(\hat{s})\right], \tag{A18}$$

$$\sum_{i \in \{Li, Le\}} n_i \lambda_{tr,i} \left[\widetilde{\Sigma}(\hat{s}_{i \to h(i)}) - \widetilde{\Sigma}(\hat{s})\right]. \tag{A19}$$

The same cancellation of incumbent promises applies to technology transitions and to idiosyncratic productivity shocks. Their joint-value contributions are

$$\mathbb{1}\{\tau = S\}\, q_{SF}(r) \left[\widetilde{\Sigma}(\hat{s}^F) - \widetilde{\Sigma}(\hat{s})\right] + \mathbb{1}\{\tau = F\}\, q_{FS}(r) \left[\widetilde{\Sigma}(\hat{s}^S) - \widetilde{\Sigma}(\hat{s})\right], \tag{A20}$$

$$\sum_{z' \in Z,\, z' \neq z} \Gamma_z(z' \mid z) \left[\widetilde{\Sigma}(\hat{s}^{z'}) - \widetilde{\Sigma}(\hat{s})\right]. \tag{A21}$$

Finally, firm dissolution ends the coalition and sends all its workers to unemployment. Since the

firm's value after exit is zero, its post-dissolution joint value is

$$\Sigma^{\text{exit}}(\vec{n}) = \sum_{i \in G} n_i U_i,$$

so the dissolution term is

$$s_F \left[\Sigma^{\text{exit}}(\vec{n}) - \widetilde{\Sigma}(\hat{s}) \right]. \tag{A22}$$

For a separation that leaves the firm without workers, the value of the empty firm is similarly zero under the free-entry condition.

Combining (A14)–(A22) with the current net flow payoff

$$(1-r)y(\hat{s}) - \sum_{i \in G} \left[n_i C_\delta(\delta_i) + C_v(v_i) \right] - \sum_{i \in \{Li, Le\}} n_i C_{tr}(\lambda_{tr,i})$$

yields the joint-surplus Bellman equation in (6). Crucially, neither that equation nor the feasible set of its reduced controls contains the outstanding promises $\vec{W}$ or the current wages $\{w_i\}_{i \in G}$. The only continuation promises that remain are $\{W_i'(\hat{s}_i^+)\}_{i \in G}$, since they determine both the job-filling rates and the value delivered to newly hired workers. Thus the maximized joint value depends only on the physical state:

$$\widetilde{\Sigma}(\hat{s}, \vec{W}) = \Sigma(\hat{s}), \tag{A23}$$

and it can be obtained by maximizing over the reduced contract

$$\mathcal{C}^\Sigma = \left\{ \{v_i, \delta_i, \lambda_{tr,i}, W_i'(\hat{s}_i^+)\}_{i \in G}, r \right\}, \qquad W_i'(\hat{s}_i^+) \geq U_i.$$

Equivalence of the optimization problems. First, let $\mathcal{C}$ solve the firm's original problem, subject to promise keeping and participation. By (A12) and (A23), its value is

$$J(\hat{s}, \vec{W}) = \Sigma(\hat{s}) - \sum_{i \in G} n_i W_i. \tag{A24}$$

Since the outstanding promises are given, an improvement in joint surplus would raise firm value by exactly the same amount. Hence the reduced part $\mathcal{C}^\Sigma$ of an optimal full contract must maximize joint surplus.

Conversely, suppose $\mathcal{C}^\Sigma$ maximizes joint surplus. Choose the remaining continuation promises to satisfy the incumbent workers' participation constraints, including after changes in type. Keep the promises to new hires specified by $\mathcal{C}^\Sigma$. The worker's Bellman equation is affine in the current wage, so the promise-keeping equality uniquely determines the wage for each incumbent type. This completes a feasible full contract $\mathcal{C}$ without changing the reduced policies or joint value. By (A24), the completed contract attains the highest possible firm value conditional on $\vec{W}$; it therefore solves the original firm problem.

Thus, the firm and joint-surplus problems are equivalent, and the joint surplus is independent of outstanding promises. □

B.2 Computation of the Stationary Equilibrium

This appendix describes the numerical procedure used to compute a balanced-growth-path (BGP) equilibrium of the model. All firm and worker masses are expressed per unit of the labor force, so that their cross-sectional distributions are constant along the BGP even when the labor force grows at rate g. The equilibrium is computed on the state space that excludes chronological age. We then

describe how to recover the latent age distribution of workers, conditional on their skill type and the state of their employer.

State space and equilibrium unknowns

Recall that the four worker types are $G = \{Li, Le, Hi, He\}$, where the first letter indicates technological skills and the second experience. Write the physical state of an operating firm as $\hat{s} = (\vec{n}, z, \tau)$, where $\vec{n} = \{n_i\}_{i \in G}$, $z \in Z$, and $\tau \in \{S, F\}$. Let $\mathcal{S}$ denote the finite set of operating-firm states in the numerical grid. The state reached by hiring (respectively, losing) one worker of type i is denoted by $\hat{s}_i^+$ (respectively, $\hat{s}_i^-$); the states reached when a worker's type changes from i to j and when productivity changes to z' are denoted by $\hat{s}_{i \to j}$ and $\hat{s}^{z'}$. The technology-switching states are $\hat{s}^F$ and $\hat{s}^S$. When the final worker leaves a firm, the firm exits the operating-state space.

Let U_i be the value of unemployment for a worker of type i. We distinguish this value from the per-capita mass u_i of unemployed workers. Similarly, e_i denotes per-capita employed workers of type i. The stationary per-capita mass of firms in state $\hat{s}$ is $\Lambda(\hat{s})$, with $f = \sum_{\hat{s} \in \mathcal{S}} \Lambda(\hat{s})$ the mass of operating firms and f^e the mass of potential entrants per worker. Employment consistency requires

$$e_i = \sum_{\hat{s} \in \mathcal{S}} n_i(\hat{s}) \Lambda(\hat{s}), \qquad \sum_{i \in G} (u_i + e_i) = 1. \tag{A25}$$

The population law of motion implies that the flow of new workers per unit of the current labor force equals $g + \lambda_r$; a fraction ω arrives as type Hi and the remaining fraction as type Li. In particular, $g + \lambda_r \geq 0$.

The principal equilibrium unknowns are the four values $\{U_i\}_{i \in G}$. For any candidate vector of unemployment values we solve the incumbent-firm problem, compute entry values and the implied stationary distribution, and then evaluate unemployment-flow and free-entry residuals. A stationary equilibrium is a zero of these residuals.

Incumbent surplus and directed-search policies

Fix a candidate positive vector $\{U_i\}_{i \in G}$. The value of unemployment implies the job-finding rate in any active submarket (i, W). For convenience, define

$$D_i \equiv (\rho + \lambda_r) U_i - b_i - \mathbb{1}\{i \in \{Hi, He\}\} \lambda_T \big(U_{\ell(i)} - U_i\big), \tag{A26}$$

where $\ell(Hi) = Li$ and $\ell(He) = Le$. The unemployed worker's Bellman equation and the indifference condition across active submarkets give

$$f_i\big(\theta_i(W)\big)(W - U_i) = D_i, \qquad W > U_i. \tag{A27}$$

Here $f_i(\theta)$ is a job-finding rate, not the mass of firms f. Under the Cobb–Douglas matching function $M(V, U) = V^\gamma U^{1-\gamma}$, with $0 < \gamma < 1$, the worker job-finding rate is $f_i(\theta) = \theta^\gamma$ and the job-filling rate per unit of recruiting intensity is $\eta(\theta) = \theta^{\gamma - 1}$. Consequently,

$$\theta_i(W) = \left(\frac{D_i}{W - U_i}\right)^{1/\gamma}, \tag{A28}$$

$$\eta\big(\theta_i(W)\big) = \left(\frac{W - U_i}{D_i}\right)^{(1-\gamma)/\gamma}. \tag{A29}$$

These expressions are used on active submarkets, for which $D_i > 0$ and $W > U_i$. Inactive vacancies have zero recruiting intensity and do not require an inversion of market tightness.

We solve the joint-surplus Bellman equation in the main text (equation (6)) for $\Sigma(\hat{s})$ at every $\hat{s} \in \mathcal{S}$. The maximization is over the reduced contract $\mathcal{C}^\Sigma$, comprising recruiting intensities v_i, firing intensities δ_i, training intensities $\lambda_{\mathrm{tr},i}$ for $i \in \{Li, Le\}$, technology investment r, and the promises to

new hires $W_i'(\hat{s}_i^+)$. The associated output is

$$y(\hat{s}) = zA_\tau\big(T^{\alpha_\tau}E^{1-\alpha_\tau}\big)^\eta,$$

with T and E determined by the workforce composition as in the main text. Value-function iteration is used until the joint-surplus Bellman residual is sufficiently small.

For each feasible hire of type i, define its marginal joint surplus by $\Sigma(\hat{s}_i^+) - \Sigma(\hat{s})$. The optimal promise is

$$W_i'(\hat{s}_i^+) = \max\big\{\gamma U_i + (1-\gamma)[\Sigma(\hat{s}_i^+) - \Sigma(\hat{s})],\ U_i\big\}. \tag{A30}$$

When hiring is active, substitution into (A29) yields

$$\eta_i(\hat{s}_i^+) = \left[\frac{(1-\gamma)\big(\Sigma(\hat{s}_i^+) - \Sigma(\hat{s}) - U_i\big)}{D_i}\right]^{(1-\gamma)/\gamma}. \tag{A31}$$

For an inactive hire the corresponding recruiting intensity and effective hiring rate are set to zero. Given the optimal $v_i(\hat{s})$, the firm's hiring rate is

$$h_i(\hat{s}) \equiv v_i(\hat{s})\eta_i(\hat{s}_i^+). \tag{A32}$$

All other policies are obtained from the same joint-surplus maximization. In particular, the endogenous training intensity $\lambda_{tr,i}$ is the arrival rate of a successful $L \to H$ technological-skill transition in the model considered here. The firm-state transition rates are therefore $n_i\lambda_{tr,i}$ for $i \in \{Li, Le\}$, without a separate training-success probability. The remaining transitions are experience accumulation at rate λ_E, technological-skill obsolescence at rate λ_T, firm technology changes at rates $q_{SF}(r)$ and $q_{FS}(r)$, productivity transitions $\Gamma_z(z' \mid z)$, worker separations at rates $\delta_i + s_i^W$, retirement at rate λ_r, and firm dissolution at rate s_F. The maximization also includes the corresponding recruiting, firing, training, and technology-investment resource costs.

B.2.1 Entry values and the inflow of new firms

Potential entrants pay the entry cost κ, draw productivity according to $\Gamma_e(z)$, and start with technology F with probability p_F and with technology S with probability $1 - p_F$. Write $\hat{s}_i^e(z,\tau)$ for the state of a firm with productivity z, technology τ, and exactly one worker of type i. Entry entails one vacancy of each type, with unit recruiting intensity, and occurs when one of these vacancies is filled.

For the first worker, the predecessor is an empty firm with zero continuation surplus. Thus the entry promise and the firm's value after a successful hire are

$$W_i^e(z,\tau) = \max\big\{\gamma U_i + (1-\gamma)\Sigma(\hat{s}_i^e(z,\tau)),\ U_i\big\}, \tag{A33}$$

$$J_i^e(z,\tau) = \Sigma(\hat{s}_i^e(z,\tau)) - W_i^e(z,\tau). \tag{A34}$$

The associated job-filling rate $\eta_i^e(z,\tau)$ follows from (A29), replacing W with $W_i^e(z,\tau)$. For notational convenience in the next expressions, let $p_\tau = p_F$ if $\tau = F$ and $p_\tau = 1 - p_F$ if $\tau = S$. The ex-ante value of entry is

$$J^e = -\kappa + \sum_{\tau\in\{S,F\}} p_\tau \sum_{z\in Z} \Gamma_e(z) \sum_{i\in G} \eta_i^e(z,\tau) J_i^e(z,\tau). \tag{A35}$$

Free entry requires $J^e = 0$.

The same entry block produces the inflow of firms into each operating state, per unit mass of

potential entrants:

$$b^e(\hat{s}) = \sum_{\tau\in\{S,F\}} \sum_{z\in Z} \sum_{i\in G} p_\tau \Gamma_e(z) \eta_i^e(z,\tau)\, \mathbf{1}\{\hat{s} = \hat{s}_i^e(z,\tau)\}. \tag{A36}$$

Its level counterpart is $f^e b^e(\hat{s})$ per unit of the labor force.

B.2.2 Stationary distribution of operating firms

For fixed unemployment values and optimal policies, construct a sparse transition-rate matrix Q on $\mathbb{S}$. We use the convention that $Q_{\hat{s}',\hat{s}}$ is the rate of moving from $\hat{s}$ to $\hat{s}'$; thus its columns index source states. For distinct operating states, the nonzero entries include:

- Hiring at rate $h_i(\hat{s})$ from $\hat{s}$ to $\hat{s}_i^+$.
- Firing and exogenous worker separation at rate $n_i(\hat{s})[\delta_i(\hat{s}) + s_i^W]$ from $\hat{s}$ to $\hat{s}_i^-$, and retirement at rate $n_i(\hat{s})\lambda_r$ for the same workforce change.
- Experience accumulation at rate $n_i\lambda_E$ for $i \in \{Li, Hi\}$, obsolescence at rate $n_i\lambda_T$ for $i \in \{Hi, He\}$, and training at rate $n_i\lambda_{\mathrm{tr},i}$ for $i \in \{Li, Le\}$.
- Technology transitions $S \to F$ and $F \to S$ at rates $q_{SF}(r)$ and $q_{FS}(r)$, and productivity transitions at rates $\Gamma_z(z' \mid z)$.

The diagonal of Q is minus the total rate of leaving its source state. In particular, its diagonal includes the exogenous dissolution hazard s_F and all transitions that leave the operating-state space, including the loss of a firm's last worker. These exit events do not appear as off-diagonal entries of Q. At employment-grid boundaries, endogenous choices that would lead outside the grid are constrained. Exogenous skill and productivity transitions must still be accounted for under the explicitly specified boundary convention; they cannot simply be discarded without changing the model.

Let Λ_t be the column vector of per-capita operating-firm masses. Its law of motion is

$$\dot{\Lambda}_t = (Q - gI)\Lambda_t + f_t^e b^e, \tag{A37}$$

where I is the identity matrix and the term $-g\Lambda_t$ adjusts for labor-force growth. Stationarity therefore requires

$$(gI - Q)\Lambda = f^e b^e. \tag{A38}$$

Because the system is linear in the mass of entrants, we first solve it for one unit of potential entrants. Define $\widetilde{\Lambda} = \Lambda / f^e$. Then

$$(gI - Q)\widetilde{\Lambda} = b^e, \qquad \widetilde{\Lambda} = (gI - Q)^{-1} b^e. \tag{A39}$$

Numerically, we solve the sparse linear system in (A39) directly, rather than forming the matrix inverse. This calculation requires the corresponding resolvent to exist; in particular, the equilibrium must generate sufficient eventual exit from operating-firm states relative to population growth. The calculation recovers the relative distribution of firms but not yet its scale.

B.2.3 Stationary employment, unemployment, and market clearing

The standardized employment of type i is

$$\widetilde{e}_i = \sum_{\hat{s}\in\mathbb{S}} n_i(\hat{s})\widetilde{\Lambda}(\hat{s}). \tag{A40}$$

To compute unemployment, we impose market clearing in each directed-search submarket. Label a vacancy by the type i and its post-hiring firm state $\hat{s}'$. This is equivalent to indexing by its associated promise when multiple vacancy classes offering the same promise are aggregated. The standardized intensity-weighted mass of vacancies directed toward $(i, \hat{s}')$ is

$$\begin{aligned}\widetilde{V}_i(\hat{s}') = &\sum_{\hat{s}\in\mathcal{S}} v_i(\hat{s})\widetilde{\Lambda}(\hat{s})\mathbb{1}\{\hat{s}_i^+ = \hat{s}'\} \\ &+ \sum_{\tau\in\{S,F\}}\sum_{z\in Z} p_\tau \Gamma_e(z)\mathbb{1}\{\hat{s}_i^e(z,\tau) = \hat{s}'\}. \end{aligned} \tag{A41}$$

The second line accounts for vacancies posted by potential entrants. At an active submarket, the vacancy mass divided by market tightness gives the associated measure of job seekers. Consequently,

$$\widetilde{u}_i = \sum_{\hat{s}' \text{ active}} \frac{\widetilde{V}_i(\hat{s}')}{\theta_i(\hat{s}')} = \sum_{\hat{s}' \text{ active}} \widetilde{V}_i(\hat{s}')\big[\eta_i(\hat{s}')\big]^{1/(1-\gamma)}. \tag{A42}$$

The final equality uses $\eta(\theta) = \theta^{\gamma-1}$. Only vacancy classes that recruit workers contribute to this sum.

Finally, the labor-force normalization pins down the entrant mass:

$$f^e = \left[\sum_{i\in G}(\widetilde{u}_i + \widetilde{e}_i)\right]^{-1}. \tag{A43}$$

The stationary per-capita masses then follow without further fixed-point iterations:

$$\Lambda = f^e\widetilde{\Lambda}, \qquad u_i = f^e\widetilde{u}_i, \qquad e_i = f^e\widetilde{e}_i, \qquad f = \sum_{\hat{s}\in\mathcal{S}}\Lambda(\hat{s}). \tag{A44}$$

This two-step normalization—first determining the shape of the firm distribution per entrant, then determining the entry scale from the labor-force constraint—avoids including the full firm distribution and the mass of potential entrants in the nonlinear equilibrium solver.

B.2.4 Unemployment-flow residuals and the equilibrium root

The preceding steps can be carried out for any admissible candidate $\{U_i\}_{i\in G}$, but the resulting cross-sectional measures need not yet satisfy the unemployment laws of motion. To impose these laws, let the per-capita inflow of hires of type i be

$$H_i \equiv \sum_{\hat{s}\in\mathcal{S}}\Lambda(\hat{s})h_i(\hat{s}) + f^e\sum_{\tau\in\{S,F\}}\sum_{z\in Z} p_\tau\Gamma_e(z)\eta_i^e(z,\tau). \tag{A45}$$

The per-capita inflow into unemployment arising from ordinary separations and firm dissolution is

$$\mathcal{D}_i \equiv \sum_{\hat{s}\in\mathcal{S}}\Lambda(\hat{s})n_i(\hat{s})\big[\delta_i(\hat{s}) + s_i^W + s_F\big]. \tag{A46}$$

Retirement is not included in $\mathcal{D}_i$, since retirees leave the labor force rather than enter unemployment. The stationary unemployment law for each $i \in G$ can now be written as

$$\begin{aligned}0 = &(g+\lambda_r)\big[\omega\mathbb{1}\{i = Hi\} + (1-\omega)\mathbb{1}\{i = Li\}\big] + \mathcal{D}_i - H_i - (g+\lambda_r)u_i \\ &+ \lambda_T\big[\mathbb{1}\{i\in\{Li, Le\}\}u_{h(i)} - \mathbb{1}\{i\in\{Hi, He\}\}u_i\big], \end{aligned} \tag{A47}$$

where $h(Li) = Hi$ and $h(Le) = He$ in the final line. The last term captures technological-skill obsolescence during unemployment; experience accumulation and training take place only on the job.

There are four unemployment equations, but only three are independent once the stationary firm distribution, population law, matching feasibility, and total-worker normalization hold. We therefore solve for the four unemployment values using three unemployment-flow conditions and the free-entry condition:

$$\dot{u}_{He} = 0, \qquad \dot{u}_{Hi} = 0, \qquad \dot{u}_{Le} = 0, \qquad J^e = 0. \tag{A48}$$

The omitted equation $\dot{u}_{Li} = 0$ is checked independently after solving. An alternative selection of three unemployment equations gives an equivalent system whenever the accounting identities hold.

For numerical stability, the nonlinear root can be expressed in logarithms. Write

$$c = \log U_{Li}, \qquad \xi = \begin{pmatrix} \log U_{He} - c \\ \log U_{Hi} - c \\ \log U_{Le} - c \end{pmatrix}. \tag{A49}$$

The scalar c controls the common level of unemployment values, while ξ controls their relative magnitudes. At a fixed c, an inner nonlinear solver finds $\xi(c)$ by setting the three unemployment-flow residuals in (A48) to zero. The outer problem solves the single equation

$$J^e\big(c, \xi(c)\big) = 0. \tag{A50}$$

A sign-changing bracket for the outer residual allows a safeguarded scalar root method, with bisection as a fallback. The solution of a nearby inner problem can be used as a warm start. Starting values may also be improved by varying the common level of the unemployment values while temporarily holding their relative log differences fixed. These steps are numerical accelerations; the equilibrium conditions remain (A48).

Once the root has been found, we recompute the firm policies and all stationary measures at the converged unemployment values. The non-hiring continuation promises and incumbent wages, if needed, are recovered under the Markov-perfect contract restriction. In particular, for a state containing a type-i worker, the promised value $W_i(\hat{s})$ is consistent with the optimal hiring promise offered in predecessor state $\hat{s}_i^-$, with single-worker firms initialized from (A33). The firm value is $J(\hat{s}) = \Sigma(\hat{s}) - \sum_i n_i(\hat{s}) W_i(\hat{s})$. Given a complete, feasible schedule of continuation promises, each flow wage is obtained from the corresponding employed-worker Bellman equation by imposing promise keeping with equality. These wage-recovery calculations do not enter the stationary distribution or the nonlinear root for $\{U_i\}_{i \in G}$.

B.2.5 Numerical consistency checks

We verify that the numerical solution satisfies the following conditions:

- The joint-surplus Bellman and policy optimality residuals are small, including at boundaries of the employment grid.
- The standardized and level firm distributions are nonnegative, and $(gI - Q)\widetilde{\Lambda} = b^e$ as well as $(gI - Q)\Lambda = f^e b^e$ hold to numerical tolerance.
- Every active submarket satisfies $f_i(\theta) = \theta^\gamma$, $\eta(\theta) = \theta^{\gamma-1}$, $f_i(\theta) = \theta\eta(\theta)$, and $f_i(\theta_i(W))(W - U_i) = D_i$.
- Employment consistency, matching feasibility, and $\sum_i (u_i + e_i) = 1$ hold.

- The free-entry residual and all four unemployment-flow residuals, including the equation omitted from root finding, are small.

B.3 Recovery of the Latent Age Distribution

Chronological age is not a payoff-relevant state in the model. All workers with the same skills and employment state face the same continuation values and transition rates, independently of their age. Nevertheless, workers enter the labor market at a known age and acquire experience and lose or regain technological skills over time. We use these transitions to recover the conditional age distribution after solving the stationary equilibrium. The age calculation does not alter the equilibrium policies.

B.3.1 The tagged-worker transition process

Construct a state space for one randomly tagged worker,

$$\mathcal{W} = \{(\hat{s}, i) : \hat{s} \in \mathcal{S},\ n_i(\hat{s}) > 0\} \cup \{\mathfrak{u}_i : i \in G\}.$$

The state $(\hat{s}, i)$ means that the worker is employed at a firm in physical state $\hat{s}$ and has type i. The state $\mathfrak{u}_i$ means that the worker is unemployed with type i. The distinct notation u_i continues to denote the stationary unemployment mass.

Let K be the transition-rate matrix of this tagged-worker process, with rows indexing the source states and columns indexing destination states. Its entries are implied by the equilibrium policies, rather than chosen separately. They include the following events:

- When a firm hires a coworker of type j at rate $h_j(\hat{s})$, the tagged worker moves from $(\hat{s}, i)$ to $(\hat{s}_j^+, i)$.
- When a coworker of type j separates or retires, the rate is $[n_j(\hat{s}) - \mathbb{1}\{j = i\}][\delta_j(\hat{s}) + s_j^W + \lambda_r]$; the tagged worker remains employed in the corresponding post-separation firm state.
- The tagged worker moves to unemployment $\mathfrak{u}_i$ at rate $\delta_i(\hat{s}) + s_i^W + s_F$. Firm dissolution also sends all coworkers to unemployment. Retirement of the tagged worker instead removes the worker from $\mathcal{W}$ at rate λ_r.
- Experience accumulation, technological-skill obsolescence, and training change either the tagged worker's type or a coworker's type. For coworker transitions the appropriate multiplicity is $n_j(\hat{s}) - \mathbb{1}\{j = i\}$; for the tagged worker's own transition the multiplicity is one. Only unemployed high-tech workers can change type while unemployed, by obsolescence $Hi \to Li$ or $He \to Le$ at rate λ_T.
- Firm technology switches and idiosyncratic productivity shocks change $\hat{s}$ without directly changing the tagged worker's type.
- Matches from unemployment to firms occur at the equilibrium rates implied by incumbent and entrant vacancies. Specifically, if $H_i(\hat{s}')$ is the per-capita matching flow from type-i unemployment into a post-hiring firm state $\hat{s}'$, then the tagged-worker transition rate from $\mathfrak{u}_i$ to $(\hat{s}', i)$ is $H_i(\hat{s}')/u_i$ whenever $u_i > 0$, where $H_i(\hat{s}') = f^e \widetilde{V}_i(\hat{s}')\eta_i(\hat{s}')$ from the submarket-specific matching flows.

The off-diagonal entries of K are these transition intensities. The diagonal is minus the sum of all rates leaving a tagged-worker state, including the absorbing retirement hazard. Under the common

retirement hazard assumed in the model, $K\mathbf{1} = -\lambda_r \mathbf{1}$. Unsuccessful or nonexistent skill-change events do not create transitions.

The stationary mass of each tagged state is known from the age-free equilibrium:

$$x^*_{\hat{s},i} = \Lambda(\hat{s}) n_i(\hat{s}), \qquad x^*_{\mathrm{u}_i} = u_i. \tag{A51}$$

Let x^* collect these masses as a column vector. Since the labor force is normalized to one, $\mathbf{1}^\top x^* = 1$.

B.3.2 Cohort occupation measures

The flow of new workers per capita is $g + \lambda_r$. All new entrants are inexperienced: a fraction ω enters unemployment as type Hi, and a fraction $1 - \omega$ as type Li. We keep these two entry groups separate to identify their age composition. Let b^H and b^L be vectors defined on $\mathcal{W}$, with all entries equal to zero except

$$b^H_{\mathrm{u}_{Hi}} = \omega(g + \lambda_r), \qquad b^L_{\mathrm{u}_{Li}} = (1 - \omega)(g + \lambda_r). \tag{A52}$$

These are flows of workers, not the firm-entry vector b^e. For each entry group $c \in \{H, L\}$, let x^c be its stationary occupation-measure vector. It solves

$$0 = K^\top x^c - g x^c + b^c, \qquad x^c = (gI - K^\top)^{-1} b^c. \tag{A53}$$

Equivalently,

$$x^c = \int_0^\infty \exp\big[(K^\top - gI)t\big] b^c \, dt. \tag{A54}$$

The exponential discounts the past inflow of each cohort by labor-force growth while propagating its members forward through the tagged-worker transition process, including exits from the labor force. In numerical applications, we solve the sparse linear systems in (A53) directly.

The sum $x = x^H + x^L$ must reproduce the tagged-worker masses calculated from the stationary firm and unemployment distributions:

$$x = x^*. \tag{A55}$$

This equality is a substantive diagnostic of the consistency of the labor-market flows and the tagged-worker generator, not an imposed renormalization.

B.3.3 Occupation measures by age group

Let a_0 be the age at labor-market entry, and let a_Y and a_O be the two age cutoffs. In the baseline calibration, $a_0 = 25$, $a_Y = 35$, and $a_O = 45$. The corresponding groups are young ($a < a_Y$), middle-aged ($a_Y \le a < a_O$), and old ($a \ge a_O$). Write κ_{yr} for the number of model-time units in one calendar year (four under the quarterly calibration), and define the times between entry and the cutoffs by

$$T_Y = \max\{0, (a_Y - a_0)\kappa_{\mathrm{yr}}\}, \qquad T_O = \max\{0, (a_O - a_0)\kappa_{\mathrm{yr}}\}. \tag{A56}$$

Thus $T_Y = 40$ and $T_O = 80$ quarters in the baseline convention.

For any entry group $c \in \{H, L\}$, the stationary occupation measure of workers whose elapsed time since entry is at least T is

$$r^c(T) = \exp\big[(K^\top - gI)T\big] x^c. \tag{A57}$$

The exponential is a matrix exponential applied to a vector; it need not be formed as a dense matrix.

The young, middle-aged, and old occupation measures of group c are then

$$x^{c,Y} = x^c - r^c(T_Y), \tag{A58}$$

$$x^{c,M} = r^c(T_Y) - r^c(T_O), \tag{A59}$$

$$x^{c,O} = r^c(T_O). \tag{A60}$$

Summing across the two entrant groups gives

$$x^Y = x^{H,Y} + x^{L,Y}, \quad x^M = x^{H,M} + x^{L,M}, \quad x^O = x^{H,O} + x^{L,O}. \tag{A61}$$

These measures partition the stationary tagged-worker distribution:

$$x^Y + x^M + x^O = x^*. \tag{A62}$$

The same procedure permits other age cutoffs without resolving the equilibrium, provided the underlying equilibrium policies and demographic parameters are unchanged.

B.3.4 Conditional age probabilities and empirical moments

For any employed tagged-worker state $(\hat{s}, i)$ with positive mass, we recover the probabilities of belonging to each age group as

$$\Pr(Y \mid \hat{s}, i) = \frac{x^Y_{\hat{s},i}}{x^*_{\hat{s},i}}, \qquad \Pr(M \mid \hat{s}, i) = \frac{x^M_{\hat{s},i}}{x^*_{\hat{s},i}}, \qquad \Pr(O \mid \hat{s}, i) = \frac{x^O_{\hat{s},i}}{x^*_{\hat{s},i}}. \tag{A63}$$

For an unemployed worker of type i and $u_i > 0$, the analogous probabilities are

$$\Pr(k \mid \mathfrak{u}_i) = \frac{x^k_{\mathfrak{u}_i}}{u_i}, \qquad k \in \{Y, M, O\}. \tag{A64}$$

The probabilities are undefined in zero-mass states, which receive zero weight when aggregating economic outcomes.

For example, the expected fraction of young workers employed by a firm in physical state $\hat{s}$ is

$$\text{YoungShare}(\hat{s}) = \frac{\sum_{i \in G} n_i(\hat{s}) \Pr(Y \mid \hat{s}, i)}{\sum_{i \in G} n_i(\hat{s})}. \tag{A65}$$

The analogous expression provides the expected old-worker share. Age probabilities can also be used to compute aggregate employment, unemployment, and wage moments by age, weighting each tagged state by its corresponding age-specific mass x^k. For example, if $w_i(\hat{s})$ is the wage of a type-i worker in state $\hat{s}$, the mean wage among employed young workers is

$$\bar{w}_Y = \frac{\sum_{\hat{s} \in \mathcal{S}} \sum_{i \in G} x^Y_{\hat{s},i} w_i(\hat{s})}{\sum_{\hat{s} \in \mathcal{S}} \sum_{i \in G} x^Y_{\hat{s},i}}, \tag{A66}$$

provided the denominator is positive. These are age-conditioned moments, not additional equilibrium state variables.

As checks, we verify that the three age-group masses are nonnegative, that their sum reproduces x^* state by state, and that each conditional probability vector sums to one wherever the corresponding state has positive mass. Because worker skill transitions, firm policies, and matching rates depend on the equilibrium, the tagged-worker generator and the latent age distribution must be recomputed

for each new counterfactual stationary equilibrium.

B.4 Constraint Efficiency

B.5 Welfare: Training Policy

Workers have linear utility and are hand-to-mouth. Their labor-market continuation utility is summarized by the promised value $W_{i,t}(s)$ when employed as type i in firm state s, and by $U_{i,t}$ when unemployed. These values incorporate future employment and skill transitions, as well as retirement. At each instant t, the subsidy-financing taxes are

$$A_t^{\text{tax}} = \int_t^\infty e^{-(\rho+\lambda_r)(u-t)} \mathcal{T}_u \, du.$$

A worker's discounted welfare, conditional on their current labor-market state x, is therefore

$$V_t(x) = \begin{cases} W_{i,t}(s) - A_t^{\text{tax}}, & \text{if employed as type } i \text{ in firm state } s, \\ U_{i,t} - A_t^{\text{tax}}, & \text{if unemployed as type } i. \end{cases}$$

For each type $i \in \mathcal{G}$, we compute mean welfare by averaging over the distribution of employment states:

$$\overline{V}_{i,t} = \frac{\int n_i(s) W_{i,t}(s) \Lambda_t(ds) + U_{i,t} \mathcal{U}_{i,t}}{\mathcal{E}_{i,t} + \mathcal{U}_{i,t}} - A_t^{\text{tax}},$$

where $\mathcal{E}_{i,t}$ and $\mathcal{U}_{i,t}$ are employed and unemployed masses of type i, also expressed per member of the labor force. New workers enter unemployed, with initial type Hi with probability ω and Li otherwise. Their welfare, evaluated at entry, is

$$V_{i,t}^e = U_{i,t} - A_t^{\text{tax}}, \qquad i \in \{Hi, Li\}.$$

These measures condition on a worker's initial type while allowing their skills and employment status to change subsequently.

C Numerical Appendix

C.1 Idiosyncratic Productivity

Empirical targets. We discipline the productivity process using the estimates of Lochner and Schulz (2024). Their annual log-productivity component follows

$$\omega_{j,t+1} = \rho_\omega \omega_{jt} + \xi_{j,t+1}, \qquad \rho_\omega = 0.75, \qquad \text{Var}(\omega) = \sigma_\omega^2 = 0.074. \tag{A67}$$

The variance is reported in Table 1, column (1), and the annual autocorrelation in the accompanying discussion on p. 101. The same specification yields labor and capital elasticities $\beta_l = 0.7129$ and $\beta_k = 0.1706$. The variance target refers to the persistent log-productivity component, rather than its innovation or the transitory production residual.

Capital residualization. To map these estimates into the model's labor-only technology, suppose capital is chosen flexibly at a common rental price R. Suppressing controls and constant productivity shifters, gross value added is $Y = e^\omega L^{\beta_l} K^{\beta_k}$, where L denotes effective labor. The first-order

condition for $\max_K\{Y - RK\}$ implies $RK^* = \beta_k Y^*$, so output net of capital expenditure is

$$Y^* - RK^* = C_K e^{\omega/(1-\beta_k)} L^\eta, \qquad \eta = \frac{\beta_l}{1-\beta_k}, \tag{A68}$$

where $C_K = (1-\beta_k)(\beta_k/R)^{\beta_k/(1-\beta_k)}$ is absorbed into the model's technology shifters A_τ. Capital adjustment amplifies log-productivity fluctuations by $1/(1-\beta_k)$ and leaves their autocorrelation unchanged. Thus, the adjusted process has

$$\eta \simeq 0.85954, \qquad \sigma_z = \frac{\sqrt{0.074}}{1-0.1706} \simeq 0.32798, \qquad \rho_z = 0.75, \tag{A69}$$

where σ_z is the stationary standard deviation of log productivity and ρ_z its annual autocorrelation.

Discretization. Let x_t denote centered adjusted log productivity. We match these moments with an Ornstein–Uhlenbeck process,

$$dx_t = -\kappa_z x_t\, dt + \nu_z\, dB_t, \qquad \kappa_z = -\frac{\log \rho_z}{4}, \qquad \nu_z = \sigma_z\sqrt{2\kappa_z}, \tag{A70}$$

where time is measured in quarters and B_t is a standard Wiener process. The factor four converts annual persistence into quarterly units; stationary dispersion is unchanged.

We approximate this diffusion by a continuous-time Ehrenfest chain with $N_z = 6$ states. Set $M = N_z - 1$ and construct the equally spaced grid

$$x_k = \frac{(2k-M)\sigma_z}{\sqrt{M}}, \qquad k = 0, \ldots, M. \tag{A71}$$

Only adjacent states communicate directly, with generator entries

$$Q_{k,k+1} = (M-k)\frac{\kappa_z}{2}, \qquad Q_{k,k-1} = k\frac{\kappa_z}{2}, \qquad Q_{kk} = -\sum_{\ell \neq k} Q_{k\ell}. \tag{A72}$$

Transitions outside the grid are omitted and all other off-diagonal entries are zero. The stationary probabilities are $\pi_k = \binom{M}{k} 2^{-M}$. We exponentiate and normalize the grid as

$$z_{k+1} = \frac{e^{x_k}}{\sum_{\ell=0}^{M} \pi_\ell e^{x_\ell}}, \qquad k = 0, \ldots, M. \tag{A73}$$

This normalization fixes mean productivity under the chain's stationary distribution at one, while preserving the log variance σ_z^2 and annual log autocorrelation ρ_z exactly. It assigns the common productivity scale to A_τ; it does not impose a zero mean for log productivity. The model's transition intensities are $\Gamma_z(z_{\ell+1} \mid z_{k+1}) = Q_{k\ell}$ for $\ell \neq k$, and the transition matrix over h quarters is $P_z(h) = \exp(hQ)$.

D Static Model

We start by presenting a static version of the model to illustrate how population aging affects productivity through the supply of skills and firms' technology choices. Workers supply technological skills and experience in proportions that vary with age, while firms combine these inputs using either a standard technology or a frontier technology, which offers higher TFP but relies more heavily on

technological skills. Aging increases the relative abundance of experience, raising the relative cost of frontier production and weakening adoption incentives. The model separates the productivity effect of the change in skill supply from the effect operating through adoption. It also clarifies which parameters govern these two channels, motivating the empirical analysis and the richer quantitative model that follow.

D.1 Environment

The economy is static and comprises two sets of agents: workers, with total measure $\mathcal{L} = 1$, and firms, with total measure M.

Workers There are two types of workers: young (with share y) and old (with share $o = 1 - y$). Age types are labels: they characterize the worker's exogenous endowment of the two skills in terms of efficiency units of labor. Young workers possess high tech skills, but low experience. We denote their skill vector as $\gamma_y = (\Gamma_T, 1)$, where $\Gamma_T \geq 1$. In contrast, old workers feature low tech skills, but high experience. Their skill vector is given by $\gamma_o = (1, \Gamma_E)$, where $\Gamma_E > 1$.[2] Therefore, the aggregate supply of tech skills and experience is given by

$$\begin{aligned} T^s(o) &= (1-o)\Gamma_T + o \\ E^s(o) &= (1-o) + o\Gamma_E \end{aligned} \tag{A74}$$

Firms There is a measure M of homogeneous firms. They are perfectly competitive and can potentially operate either the frontier technology or the standard one.

Given their technology, they produce output using the following technology-specific decreasing-returns-to-scale production function:

$$y_\tau = A_\tau \left(T^{\alpha_\tau} E^{1-\alpha_\tau}\right)^\eta, \quad \eta < 1, \alpha_\tau \in [0,1].$$

Hence, total output only depends on which technology the firm operates, and on the skill allocation, but not on which types of workers are employed.

We make two assumptions about the structure of production.

Assumption 1. The frontier technology features higher TFP: $A_F > A_S = 1$.

Assumption 2. The frontier technology is more intensive in technological skills: $\alpha_F > \alpha_S$.

Adopting the frontier technology is a costly and risky process. In particular, we assume that all firms are initially endowed with the standard technology. Before production, a firm can invest a fraction r of the output the firm would produce if it remained at the standard technology to generate an adoption probability $\phi(r) \in [0,1]$. We assume that the function ϕ is strictly increasing and strictly concave in the investment, so that $\phi' > 0, \phi'' < 0$, and we normalize $\phi(0) = 0$. If the investment is successful, the firm operates the frontier technology; otherwise, it remains at the standard technology.

D.2 Competitive Equilibrium

Skill prices and wages Denote by p_T and p_E the equilibrium price of an efficiency unit of tech skills and experience, respectively. These are the fundamental prices in the economy. Because workers are

[2] Notice that for the main results to hold, we only need old workers to have a strict comparative advantage in experience ($\Gamma_E > 1$), while young workers are allowed to have a weak comparative advantage ($\Gamma_T \geq 1$).

fixed bundles of skills, and the labor market is perfectly competitive, the wage rates for young and old workers are given by

$$\begin{aligned} w_y(p_T, p_E) &= \Gamma_T p_T + p_E \\ w_o(p_T, p_E) &= p_T + \Gamma_E p_E. \end{aligned}$$

Hence, the wages w_y and w_o simply label the value of the skill bundles carried by young and old workers.

A competitive equilibrium is an allocation of tech skills and experience across frontier and standard firms, a technology investment r, and skill prices p_T, p_E, such that every firm type maximizes profits, firms choose which the technology investment optimally, and the markets for both skills clear.

Firm operating problem Consider a firm operating technology $\tau \in \{F, S\}$. If the firm hires n_y young workers and n_o old workers, its total skill composition would amount to

$$\begin{aligned} T &= \Gamma_T n_y + n_o \\ E &= n_y + \Gamma_E n_o, \end{aligned}$$

for a total cost of

$$w_y n_y + w_o n_o = p_T T + p_E E.$$

Hence, while the firm hires workers, its production and costs only depend on the amount of skills employed and their price.

It is convenient to rewrite the firm's problem in terms of total employment and workforce composition. Define total employment $l = n_y + n_o$ and the share of young workers $x = n_y/l$. It follows that the average tech skills and experience supplied by one employee are

$$\begin{aligned} \bar{T}(x) &= 1 + (\Gamma_T - 1)x \\ \bar{E}(x) &= \Gamma_E - (\Gamma_E - 1)x. \end{aligned}$$

Total skill inputs are $T = l\bar{T}(x),\ E = l\bar{E}(x)$.

Therefore, the problem of a firm of a given technology class τ can be stated as maximizing profits by choosing the optimal number of employees (l) and the share of young workers to hire x. In particular, each firm solves

$$\begin{aligned} \pi_\tau(p_T, p_E) \coloneqq \max_{l \in \mathbb{R}_+, x \in [0,1]} & A_\tau l^\eta \left(\bar{T}(x)^{\alpha_\tau} \bar{E}(x)^{1-\alpha_\tau}\right)^\eta - l\left[p_T \bar{T}(x) + p_E \bar{E}(x)\right] \\ \text{sub. to.} \quad \bar{T}(x) &= 1 + (\Gamma_T - 1)x \\ \bar{E}(x) &= \Gamma_E - (\Gamma_E - 1)x. \end{aligned}$$

Technology adoption problem Recall that a firm is initially endowed with the standard technology and therefore obtains operating profits π_S absent successful adoption. By investing r, it gives up a fraction of standard output ry_S and reaches the frontier with probability $\phi(r)$. The return to adopting the frontier technology is $\Delta\pi = \pi_F - \pi_S$. Hence, each firm chooses a technology investment r to solve

$$\max_r \ \phi(r)\Delta\pi - ry_S.$$

Denote by $a \equiv \phi(r)$ the equilibrium adoption probability chosen by each firm, and by $m \equiv p_T/p_E$ the equilibrium price ratio. Since all firms are homogeneous, by the law of large numbers, a is also

the share of firms that operate the frontier technology in equilibrium.

The following proposition holds.

Proposition 4. *Equilibrium aggregate technology adoption is a function of the price ratio only, i.e.,* $a = a(m)$. *An increase in the price ratio lowers aggregate technology adoption:* $\frac{da(m)}{dm} < 0$.

Proposition (4) establishes that equilibrium prices affect the share of firms that adopt the frontier technology only through their ratio m. This result arises because a proportional change in the level of p_T and p_E only alters the firms' optimal scale, but not the relative profitability of the two technologies. The response of adoption to the price ratio m hinges on the relative skill intensity of frontier production. When $\alpha_F > \alpha_S$, an increase in the relative price of tech skills erodes the profitability of frontier production and weakens adoption incentives.

D.3 Population Aging as a Skill-Supply Shock

We model population aging as an increase in the share of old workers, o, holding total labor supply fixed at $\mathcal{L} = 1$. Equation (A74) implies that aging increases the supply of experience relative to technological skills. This change in relative scarcity is reflected in equilibrium skill prices, as the following lemma establishes.

Lemma 5. *An increase in the share of old workers increases the equilibrium price ratio:* $\frac{dm}{do} > 0$.

The increase in m reduces the relative profitability of frontier production because it is more intensive in technological skills. The adoption response therefore rests on Assumption (2): the change in relative skill prices disproportionately raises the cost of operating the frontier technology. Combining this result with Proposition (4) yields the following corollary.

Corollary 6. *An increase in the share of old workers reduces the equilibrium aggregate technology adoption rate:* $\frac{da(m)}{do} = \frac{da(m)}{dm} \times \frac{dm}{do} < 0$.

To characterize the productivity consequences of these adjustments, define aggregate output as $\mathcal{Y} = M[ay_F + (1-a)y_S]$. Since all workers are employed and $\mathcal{L} = 1$, this quantity also measures gross output per worker. To relate productivity to the underlying parameters, define the frontier share of aggregate output, the output-weighted tech-skill intensity, and the frontier-to-standard output ratio as

$$\omega \equiv \frac{May_F}{\mathcal{Y}}, \qquad \bar{\alpha} \equiv \omega\alpha_F + (1-\omega)\alpha_S, \qquad \tilde{\pi} \equiv \frac{y_F}{y_S}.$$

Finally, let $\varepsilon_{a,\tilde{\pi}} \equiv d\log a / d\log \tilde{\pi}$ denote the elasticity of adoption with respect to relative profitability, and $\varepsilon_{m,o} \equiv d\log m / d\log o$ the elasticity of relative skill prices with respect to the old-worker share.

The following proposition separates the effect operating through skill supply from the effect operating through technology adoption.

Proposition 7. *The effect of population aging on aggregate labor productivity can be decomposed into a direct channel and an adoption channel:*

$$\frac{d\log \mathcal{Y}}{d\log o} = \underbrace{\eta\left[(1-\bar{\alpha})\frac{o(\Gamma_E - 1)}{E^s} - \bar{\alpha}\frac{o(\Gamma_T - 1)}{T^s}\right]}_{\textit{Direct channel}} \underbrace{- \eta\,\Delta\alpha\,\omega\left(1 - \frac{1}{\tilde{\pi}}\right)\varepsilon_{a,\tilde{\pi}}\,\varepsilon_{m,o}}_{\textit{Adoption loss}}. \tag{A75}$$

The direct channel captures the change in productive skills at a given technology mix. Replacing young workers with older workers adds experience but reduces technological skills, with the size

of these changes governed by Γ_E and Γ_T. Their contribution to productivity depends on the skill intensity of production: a higher $\bar{\alpha}$ places greater weight on the loss of technological skills and less on the gain in experience.

The same shift in skill supply also affects productivity by discouraging frontier adoption. All else equal, the adoption loss increases with the skill-intensity gap, $\Delta\alpha$, and the price and adoption elasticities. The gap determines how strongly rising relative tech-skill prices erode frontier profitability, while the elasticities govern how the initial shift in skill supply translates into lower adoption. Furthermore, for a given proportional decline in adoption, the output loss is larger when frontier firms account for a greater share of production (high ω) and have a larger output advantage, (high A_F).

Quantifying these channels therefore requires jointly determining workers' comparative advantages and the technological differences between frontier and standard production. The next subsection explains how employment, workforce composition, and wage outcomes provide information about these parameters.

D.4 Identification

We use the static model to illustrate how firm and worker outcomes inform the joint estimation of the parameters governing the productivity effects of aging. Consider two otherwise identical firms facing the same skill prices, one operating the frontier technology and the other the standard technology. Denote their differences in employment and workforce composition by $\Delta l = l_F - l_S$ and $\Delta x = x_F - x_S$.

Role of $\Delta\alpha$. At an interior optimum, workforce composition satisfies

$$\frac{\bar{T}(x_\tau)}{\bar{E}(x_\tau)} = \frac{\alpha_\tau}{(1-\alpha_\tau)m}.$$

Conditional on the skill price ratio m and workers' skill endowments, optimal composition depends on skill intensity but is independent of TFP.

Corollary 8. *At fixed price m, and conditional on skill endowments Γ_T, Γ_E the following holds:*

(i) $\Delta x > 0$ *if and only if* $\Delta\alpha > 0$

(ii) Holding α_S fixed, Δx is increasing in $\Delta\alpha$

Comparing the workforce composition of frontier and standard firms, therefore, helps separate $\Delta\alpha$ from A_F.

Role of A_F. Optimal employment, in turn, satisfies

$$l_\tau = A_\tau^{\frac{1}{1-\eta}} \left[\frac{\eta \left[\bar{T}(x_\tau)^{\alpha_\tau} \bar{E}(x_\tau)^{1-\alpha_\tau} \right]^\eta}{p_T \bar{T}(x_\tau) + p_E \bar{E}(x_\tau)} \right]^{\frac{1}{1-\eta}}.$$

Corollary 9. *At fixed price m and age composition x_τ, Δl is increasing in A_F.*

Conditional on $\Delta\alpha$, therefore, the employment gap between frontier and standard firms helps identify A_F.

Role of Γ_T, Γ_E. The optimality conditions shown so far describe a frictionless benchmark, whereas employment and workforce composition may adjust gradually following adoption. To illustrate the implications of incomplete adjustment, suppose that production switches immediately to the frontier technology, while employment and the young share remain fixed at the optimal level implied by the standard technology, i.e., l_S and x_S. The ratio of young to old workers' marginal products is

$$R_\tau(x) \equiv \frac{MP_{y,\tau}}{MP_{o,\tau}} = \frac{\Gamma_T \alpha_\tau \bar{E}(x) + (1-\alpha_\tau)\bar{T}(x)}{\alpha_\tau \bar{E}(x) + \Gamma_E(1-\alpha_\tau)\bar{T}(x)}.$$

Corollary 10. *Because frontier production places greater weight on tech skills ($\Delta\alpha > 0$), young workers gain in relative productivity upon the technology upgrade:*

$$\widehat{\Delta \log R} \equiv \log R_F(x_S) - \log R_S(x_S) > 0.$$

Conditional on α_F, α_S, the gain $\Delta \log R$ is increasing in both Γ_T and Γ_E.

This response motivates an additional empirical prediction. If proportional changes in marginal products pass through to wages with some positive elasticity,[3] adoption raises the young-to-old wage ratio within the establishment.

Mapping Model to Data. These results guide the empirical analysis in Section (2). Whereas the static model compares otherwise identical firms operating different technologies, in the data we estimate changes within German establishments around digital-technology adoption relative to matched non-adopting establishments. The responses of workforce composition and employment play distinct roles. Conditional on workers' skill endowments and the baseline skill intensity of production, the change in the young-worker share helps discipline the skill-intensity gap, $\Delta\alpha$. Given this composition response, the change in employment provides information about the frontier productivity advantage, A_F, because a more productive frontier technology supports a larger optimal scale.

Instead, wage outcomes provide information about workers' comparative advantages. The response of young-to-old wages among incumbent workers is the empirical proxy of $\widehat{\Delta \log R}$: if changes in marginal products pass through to wages, its magnitude is governed by Γ_T and Γ_E. Because this response depends on both parameters, we complement it with moments describing wage growth and dispersion during workers' early careers. At given skill prices, a larger experience advantage, Γ_E, raises older workers' earnings relative to those of younger workers, whereas a larger technological advantage, Γ_T, narrows this gap. In the quantitative model, these moments jointly discipline workers' initial skills and their subsequent evolution over the life cycle.

[3] A feature that will be present in the quantitative model.

E Additional Figures and Tables

E.1 Long-Run Impact of Aging

Table A7: Stationary Equilibria Under Baseline and 2050s Demographics

	Baseline (2005–2015)	Long-Run (2050s)	Change (%)
Share of Frontier Firms (Firm weighted), percent	27.594	25.789	-6.542
Share of Frontier Firms (Employment weighted), percent	26.390	24.158	-8.459
Firm Entry Rate (Annualized), percent	5.346	4.930	-7.784
Average Firm Size (Average Log)	1.822	1.832	0.500
Share of He, percent	22.456	22.931	2.114
Share of Hi, percent	7.310	6.326	-13.452
Share of Le, percent	58.468	60.267	3.078
Share of Li, percent	11.767	10.476	-10.971
Unemployment rate, percent	14.849	13.963	-5.968
within He, percent	9.096	8.568	-5.797
within Hi, percent	24.019	24.754	3.060
within Le, percent	11.657	10.873	-6.730
within Li, percent	35.995	37.035	2.890

Notes: The table compares the stationary equilibrium calibrated to the 2005–2015 German labor market with the long-run equilibrium under the cohort-entry rate implied by 2050s demographic projections. All other structural parameters are held fixed. Frontier firms are reported both as a share of establishments and weighted by their employment. The entry rate is annualized, firm size is the mean log of employment, and worker-type shares are fractions of the labor force for $i \in G = \{Li, Le, Hi, He\}$. Unemployment rates are expressed as fractions of the corresponding labor-force group. The final column gives percentage changes relative to baseline.

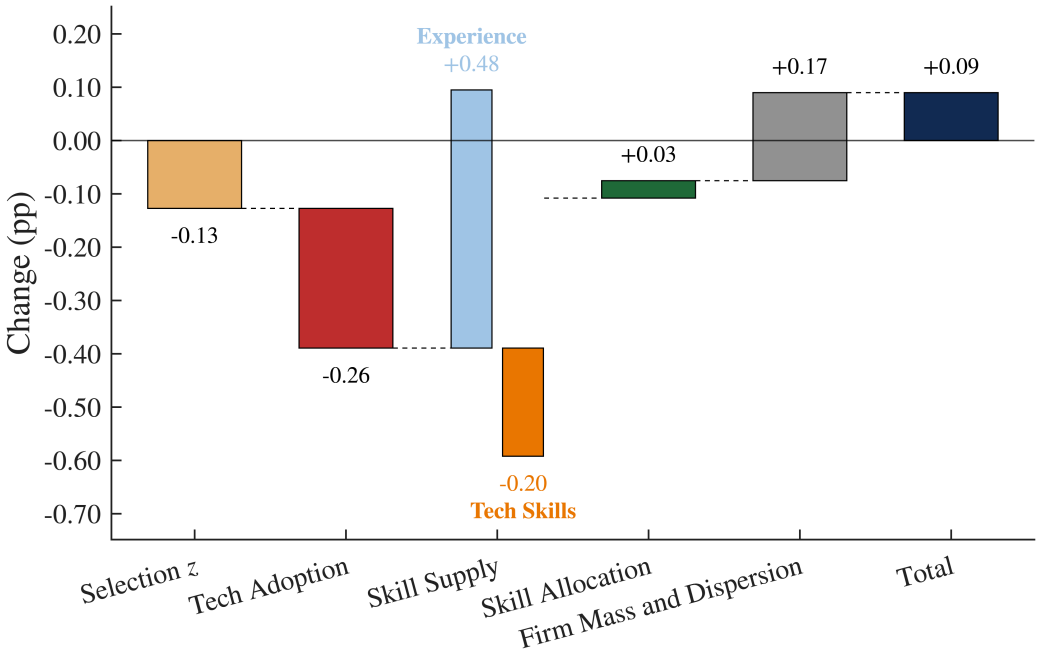
Table A8: Welfare and Consumption by Worker Type, Excluding Dividends

	Long-Run (2050s)
Change in Welfare excluding dividends, percent	**0.050**
Incumbents He	-0.089
Incumbents Hi	0.100
Incumbents Le	-0.077
Incumbents Li	-0.007
New entrants Hi	0.099
New entrants Li	-0.003
Change in SD log Welfare excluding dividends, percent	**-1.678**
He	-4.157
Hi	1.217
Le	-4.908
Li	0.029
Change in Consumption excluding dividends, percent	**0.172**
Incumbents He	-0.163
Incumbents Hi	0.772
Incumbents Le	-0.137
Incumbents Li	0.413
New entrants Hi	0.000
New entrants Li	0.000
Change in SD log Consumption excluding dividends, percent	**-1.928**
He	-3.004
Hi	2.082
Le	-3.395
Li	3.398

Notes: Current workers include employed and unemployed workers. Employed consumption is the wage; unemployed consumption is the flow payoff b_i. The common dividend entitlement and dividend flow are excluded.

Notes: The table shows percentage changes from the calibrated stationary equilibrium to the 2050s stationary equilibrium in welfare and consumption measures by initial/current worker type, before any common dividend entitlement. Incumbent averages include employed and unemployed workers in the initial type, with weights given by their stationary worker-state masses; entrant values are evaluated upon entry into unemployment as Hi or Li. Consumption is the current wage if employed and the unemployment flow payoff b_i if unemployed. Dispersion rows report changes in the standard deviation of log welfare or log consumption within the indicated group. All figures exclude dividend income, as indicated in the heading.

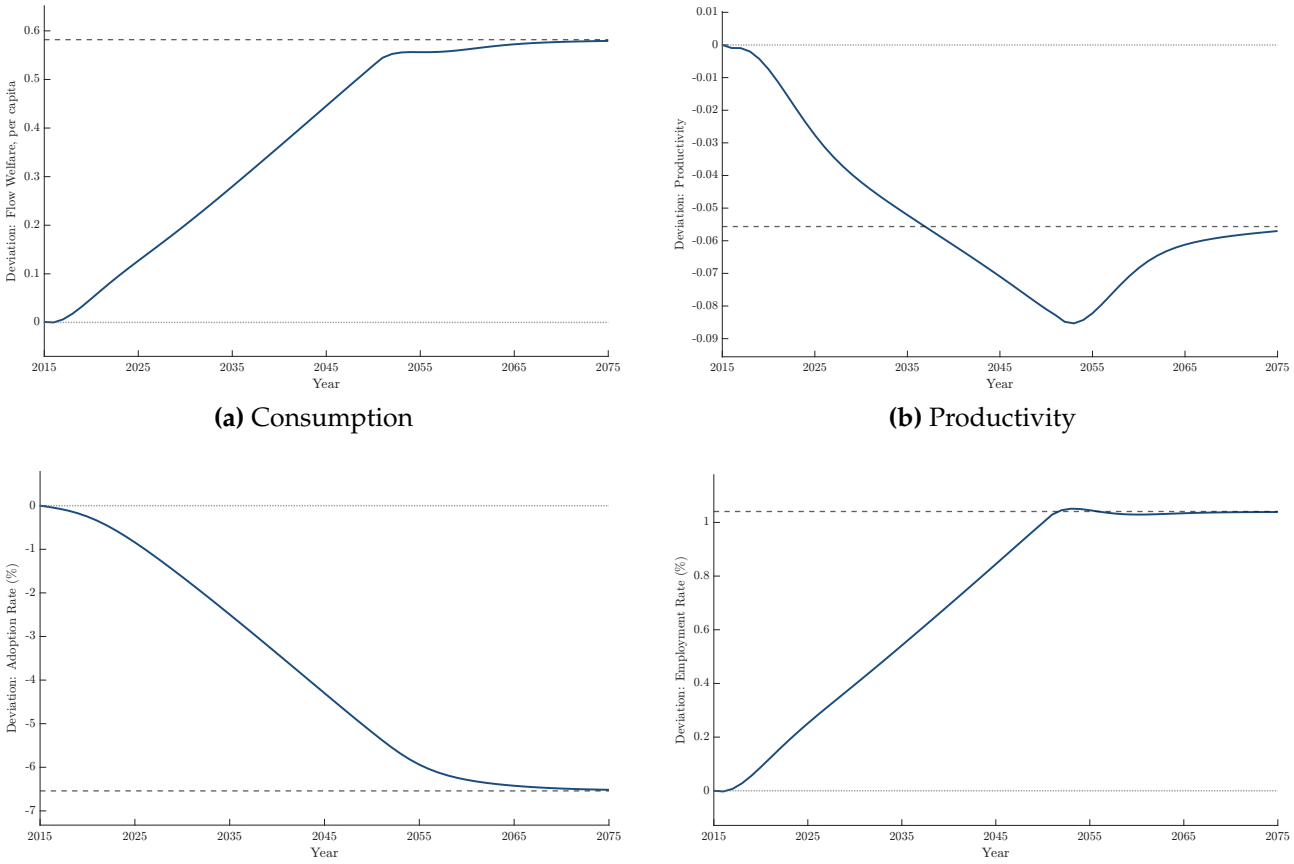
Figure A8: Decomposition of the Aggregate Productivity Effect of Aging



Notes: The bars are grouped Shapley contributions to the percentage change in aggregate TFP between the 2005–2015 and 2050s stationary equilibria. Experience and technological skills distinguish the two components of skill supply; technology adoption captures shifts between standard and frontier firms. The residual group labeled "Others" collects the remaining channels, including productivity selection, skill allocation across firms, and the firm-mass/size-distribution term arising from decreasing returns. The contributions sum to the total change.

E.2 Transitional Dynamics

Figure A9: Aggregate Outcomes Along the Aging Transition



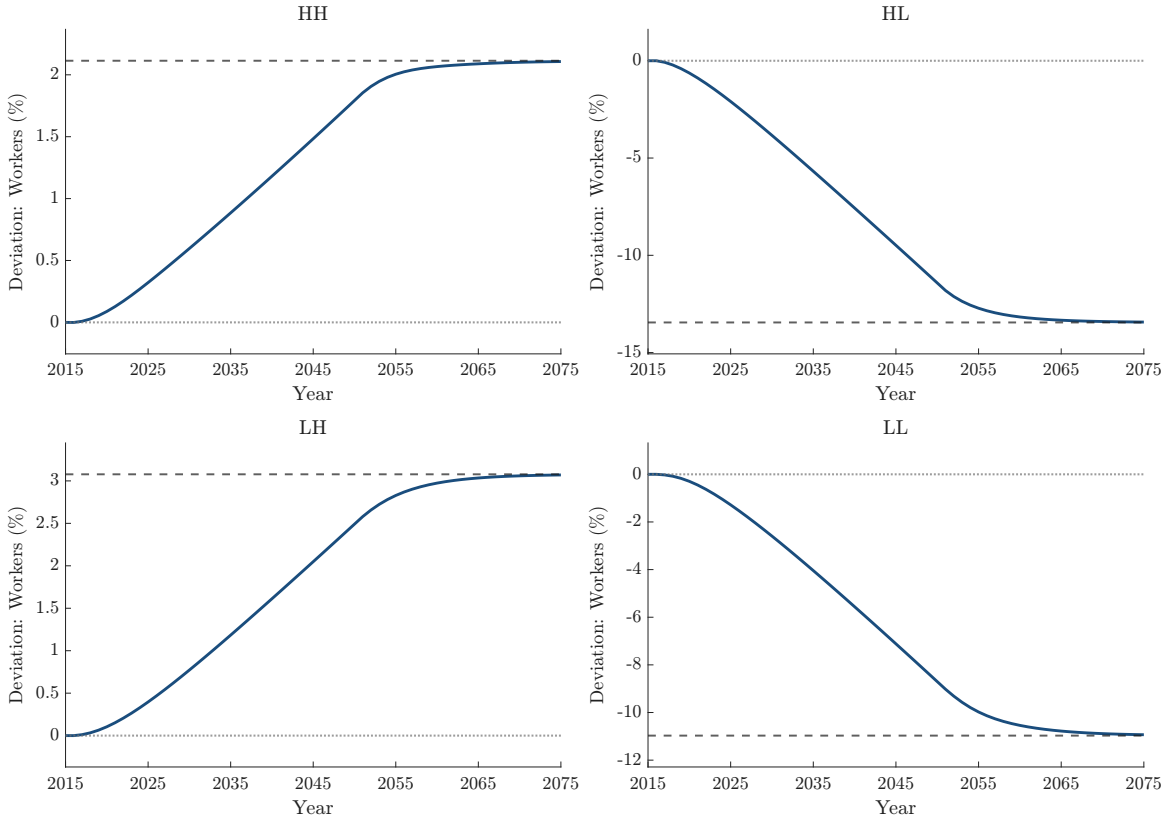
(a) Consumption **(b)** Productivity

(c) Adoption **(d)** Employment Rate

Notes: The figure traces the equilibrium transition from the 2005–2015 calibration to demographics implied by the projected 2050s labor force, with all other structural parameters held fixed. Each panel plots the deviation from the initial stationary equilibrium: flow consumption per labor-force member, aggregate productivity, the frontier-technology firm share, and the employment rate. The demographic entry rate changes according to the path described in the text and becomes constant at its long-run 2050s level. Horizontal reference lines indicate the initial and long-run outcomes.

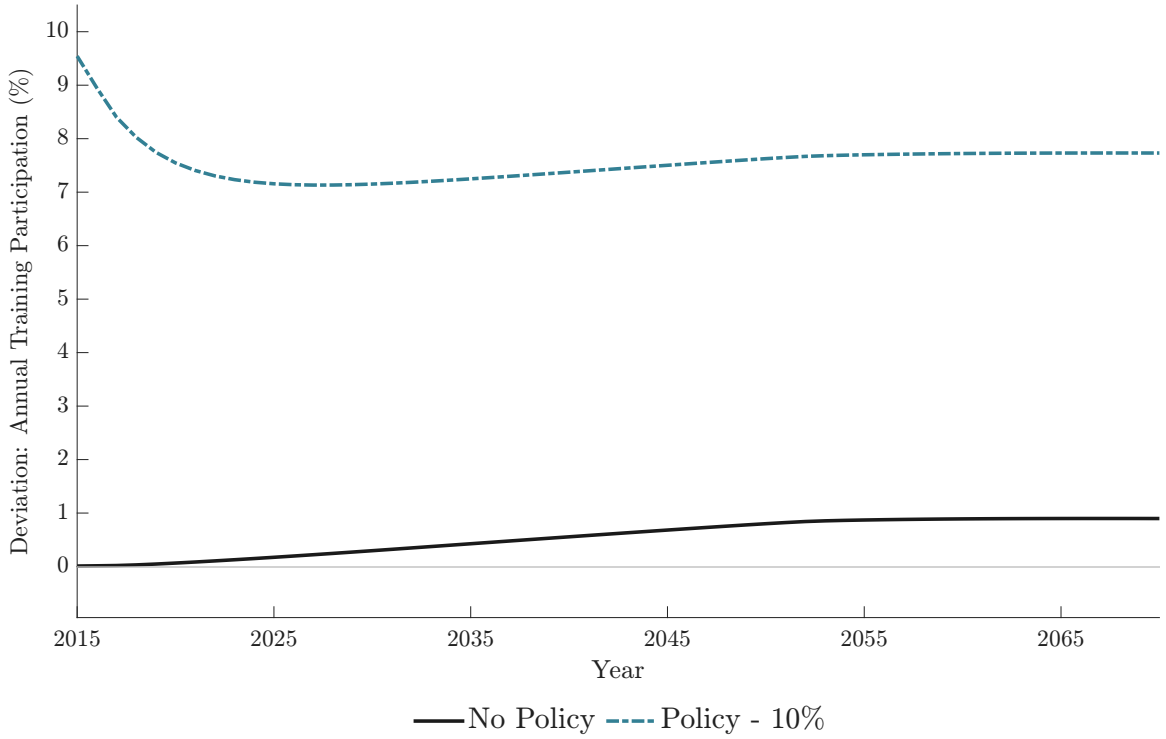
E.3 Macroeconomic Implications of Training Subsidies

Figure A10: Workforce Skill Composition Along the Aging Transition



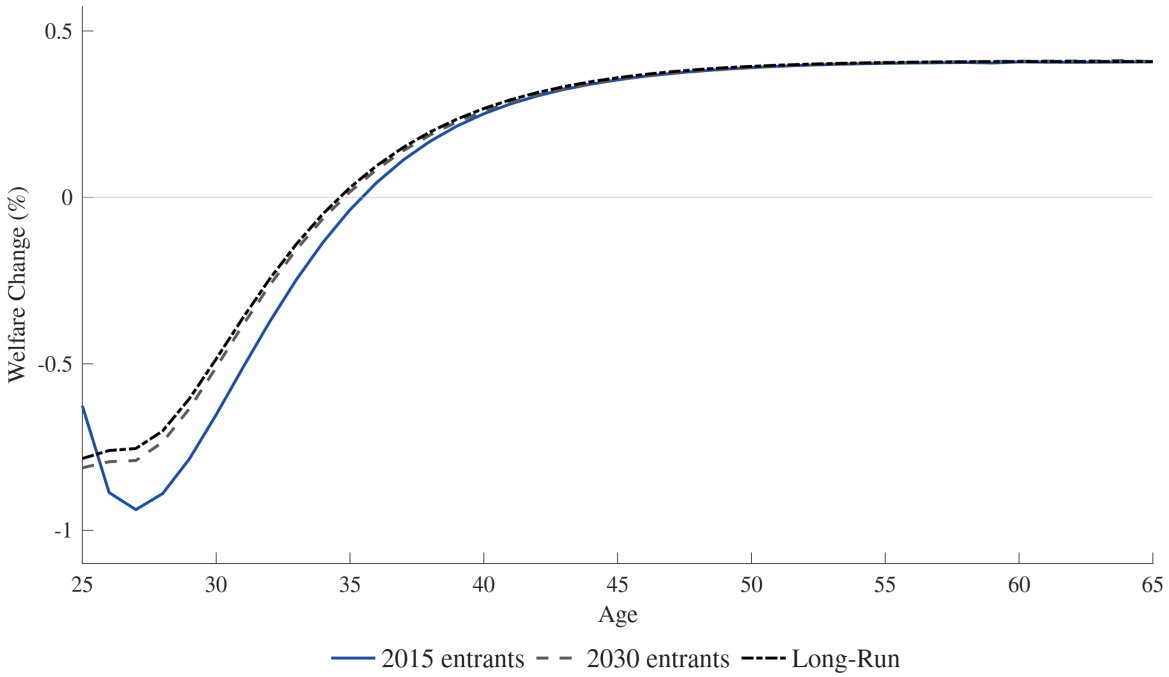
Notes: The panels report percentage deviations from the baseline in the population shares of the model's four worker types during the demographic transition. The corresponding types in the main paper are He, Hi, Le, and Li: high/low technological skills crossed with experienced/inexperienced status. The changing cohort-entry rate affects these shares through experience accumulation, technological-skill obsolescence, endogenous training, and retirement. Horizontal dashed lines mark long-run deviations. All shares are expressed relative to the total labor force.

Figure A11: Training Participation Along the Subsidized Aging Transition



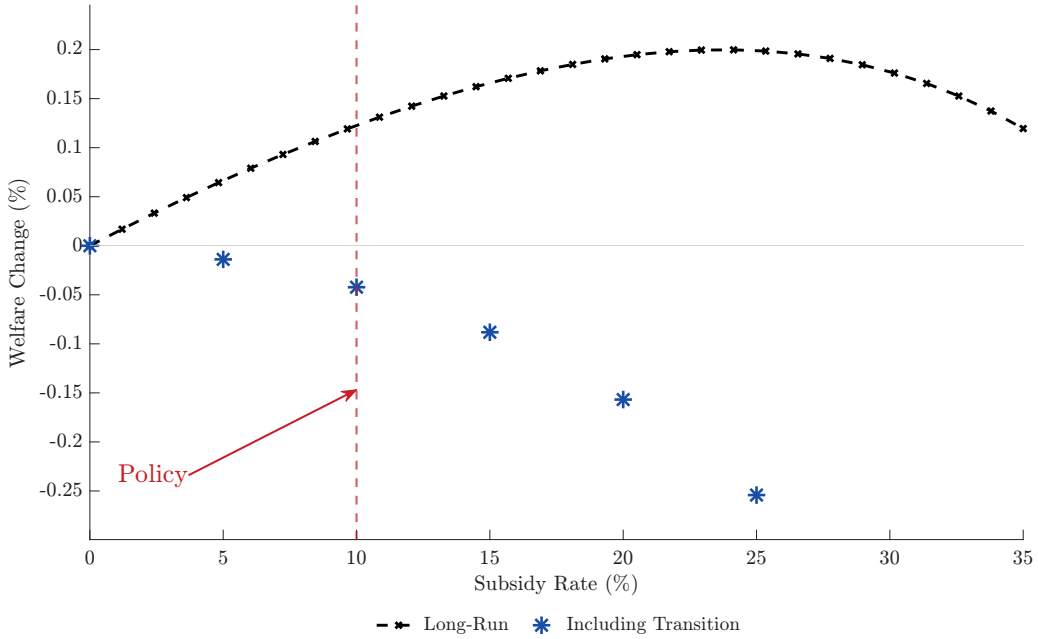
Notes: The figure compares the percentage deviation in annual successful on-the-job technological training from its initial stationary rate along the aging transition, with and without a permanent 10% training-cost subsidy. The subsidy is introduced at the beginning of the transition and financed through lump-sum taxes; both economies follow the same demographic path. Training is available to employed low-tech workers and converts them to high-tech types. The scale reports relative changes in training incidence, not percentage-point changes in the workforce share.

Figure A12: Life-Cycle Flow Consumption With and Without Training Subsidies



Notes: The curves compare the percentage effect of the 10% training subsidy on flow consumption at each age for workers entering the labor force in different calendar years, together with the corresponding long-run comparison. Labor-market entry occurs at age 25. The subsidized and unsubsidized economies follow the same demographic path, so the curves isolate the policy difference for each cohort. These are age-specific flow-consumption comparisons, not discounted lifetime welfare measures; the initially lower consumption reflects the resource cost of additional training.

Figure A13: Training Subsidies: Long-Run Versus Transition Welfare



Notes: The figure compares the welfare change from a permanent training subsidy under two evaluation criteria. The long-run curve compares welfare in the eventual stationary equilibria with and without the subsidy; the transition-inclusive series evaluates discounted welfare from the beginning of the demographic transition, allowing for its full adjustment path and subsidy-financing costs. Both use the same demographic scenario. The vertical marker identifies the calibrated 10% subsidy. Values are percentage changes relative to the corresponding no-policy allocation.