Population Aging, Firms, and Worker Composition

Santiago Caicedo, Cheng Chen, Aspen Gorry and Takahiro Hattori*

October 2026

Abstract

Population aging changes the *composition of workers* in the labor force. We develop a framework to understand how worker composition affects productivity in an economy with heterogeneous establishments and apply it to Japan. Reallocating workers to the optimal composition of experience and skill increases output by 8.5%. Strong complementarity between workers' skill and establishment productivity implies differential hiring strategies across establishments and heterogeneous effects of population aging on wages and market concentration. We find evidence consistent with these predictions using Japanese data. Despite large potential gains from reallocation, we find that further aging in Japan has only modest impacts on worker productivity.

Keywords: Population Aging, Labor Force Composition, Comparative Advantage of Productive Establishments

JEL-Codes: D21, D22, E23, E24, J11, J23

*Caicedo: Northeastern University, sa.caicedo@northeastern.edu; Chen: Clemson University, cchen9@clemson.edu; Gorry: Clemson University, cgorry@clemson.edu; Hattori: The University of Tokyo, hattori0819@gmail.com. We thank Sadhika Bagga, Zhang Chen, Seula Kim, Victoria Gregory, Christian vom Lehn, Oksana Leukhina, Xiao Ma, Makoto Nakajima, Elena Pastorino, Ben Pugsley, Michael Wong for insightful suggestions and the audiences at the 2023 SEA Meeting, the University of Hong Kong, the 2025 ASSA annual meeting, the Barcelona Summer Forum, Bentley Applied Macro Workshop, the Cleveland Fed., the 2026 SED annual meeting, and Northeastern University for their helpful comments. This work was supported by the funding from the University of Tokyo, Nomura Foundation (for international collaboration), and Japan Society for the Promotion of Science (JSPS) KAKENHI Grant-in-Aid (21K13321, 24K16394). All errors are ours. We thank Isabella Becerra for her excellent research assistance.

1 Introduction

The composition of workers by observable demographic groups, such as experience and skill, is a crucial factor in determining aggregate economic outcomes. For example, increases in cohort size of the baby boom generation reduced wages and wage growth for these larger cohorts, highlighting that workers with different experience are imperfect substitutes in the labor market (Welch, 1979; Berger, 1985). Moreover, if heterogeneous establishments employ different portfolios of workers in production, compositional changes in the workforce can have differential impacts across establishments. Recently, the composition of workers has been evolving rapidly due to population aging, declining fertility, and immigration (Toossi and Elka, 2017; Cardona, 2023; Frey, 2025; Qi, 2025; Geruso and Spears, 2026). Given large differences in demographic patterns across countries and over time, it is important to understand how the composition of workers influences the allocation of workers across establishments and aggregate productivity.

Demographic change influences aggregate output and productivity through multiple channels. Generally, a smaller working-age population reduces output and innovation by reducing labor input and firm entry.[1] Beyond these aggregate effects, demographic shifts change the composition of the labor force, potentially altering the allocation of workers across production units. This paper focuses on this latter channel. Holding labor supply constant, we examine how changes in *workforce composition* affect worker allocation and productive efficiency. In doing so, we abstract from changes in the size of the working-age population and from its consequences for business formation and innovation, which have been examined in previous studies. Understanding the optimal workforce composition across establishments sheds light on how workers with different characteristics are efficiently allocated across production units and informs policies that influence the demographic makeup of the labor force, such as immigration, education, fertility, and retirement policies.

This paper seeks to characterize the optimal allocation of workers to establishments to understand the optimal worker composition and the impacts of aging, using Japan as a case study. Japan is an interesting case because it is at the forefront of demographic aging.[2] Japan also

[1] Recent papers highlighting these channels include Maestas et al. (2016), Engbom et al. (2019), Hopenhayn et al. (2022), Hayashi (2023), Bianchi and Paradisi (2024), Karahan et al. (2024), and Weil (2026).

[2] The median age of the Japanese population was 48.4 years in 2020 and has been increasing over time (IMF, 2020). Moreover, the World Bank reports that the growth rate of the working-age population has declined from 0.24% in 2001 to -0.45% in 2021. For recent work studying the impact of rapid aging in Japan, see Braun et al. (2009), Braun and Joines (2015), and Kitao (2015).

has establishment data with rich demographic information at the worker level, including worker tenure. Using these data, we document that larger (more productive) establishments employ younger workers on average while also exhibiting higher average tenure. They display steeper tenure-experience profiles, indicating that tenure rises more rapidly with experience in larger establishments. In addition, larger establishments disproportionately hire younger workers and separate from older workers relative to smaller establishments.

Motivated by these facts, we develop a framework to characterize the efficient composition of workers across heterogeneous establishments. We start with a static framework that highlights how labor composition influences labor allocation and aggregate output. The economy consists of heterogeneous establishments that differ in productivity and employ workers from discrete groups defined by two characteristics: work experience and skill. Workers from different demographic groups are imperfect substitutes in production, and labor markets clear separately for each group, implying that all establishments face the same wage for a given type of worker. As a result, establishments employ a *portfolio* of workers from multiple groups, with diminishing marginal products within each group. This framework allows us to characterize both the efficient allocation of workers across establishments and the optimal composition of workers that maximizes aggregate output.

The optimal allocation from the static model accurately describes worker allocation when there are no costs of reallocating workers across establishments. In the presence of reallocation frictions, it is necessary to model both the cost of worker reallocation and the transitions workers make across demographic groups. We extend the model to include these dynamic features to understand worker transitions and the importance of reallocation costs. The dynamic model features workers whose experience evolves deterministically and skills that evolve stochastically over their life cycle. Heterogeneous establishments hire a portfolio of workers in each experience-skill bin and must plan for how that portfolio evolves over time. In this environment, reallocation frictions distinguish between internal and external labor markets, thereby providing the model with a meaningful definition of tenure.

We use Japanese establishment-employee matched data to parameterize the static model. The group-specific labor productivity parameters are crucial for determining the optimal composition of the labor force. These parameters are chosen to match observed moments on wages and the percentage of workers in each experience/tenure bin. Moreover, the parameter for the complementarity between establishment productivity and skill is informed by matching differential experience-tenure and wage profiles across establishments. The calibrated

parameters imply stronger complementarity (i.e., log-supermodularity) between establishment productivity and worker skill than between establishment productivity and worker experience. As a result, more productive establishments have a comparative advantage in employing skilled workers. Consistent with the model's predictions, worker-level hiring and separation data reveal systematic life-cycle patterns in worker flows across establishments.

Using these group-level productivity parameters, we compute the optimal composition of workers across groups. Our main result is that Japanese GDP per worker would be 8.5% higher with the optimal experience and skill composition. The optimal composition has more high-skill workers and a larger share of experienced workers, those with 21-40 years of experience or ages 41-60, than Japan's current population. Most of the gain comes from skill: moving only to the optimal skill composition would raise output by 5.6%, against 2.1% for moving only to the optimal experience composition. Additionally, the calibration implies a relative scarcity of skilled workers among experienced worker groups, which disproportionately constrains the most productive establishments.

We re-estimate the results using a dynamic model with reallocation frictions. Reallocation frictions imply that high-productivity establishments have weaker incentives to hire young workers and stronger incentives to retain older, low-skill workers. As a result, young and high-skill workers earn lower wages than in the baseline model. The dynamic model thus exhibits greater wage inequality across experience groups but lower wage inequality across skill levels. Moreover, the presence of reallocation frictions amplifies the estimated productivity differentials between high- and low-skilled workers as experience increases. Relative to the benchmark counterfactual, GDP per worker increases by a larger amount when the economy is endowed with the optimal worker composition (9.1% versus 8.5%). Importantly, nearly all of this gain is driven by the optimal composition of worker skills and experience, as eliminating reallocation frictions alone increases GDP per worker by only 0.16%. The larger gain instead reflects the change in establishment–worker group-specific productivity parameters, which become more dispersed across skill groups at large establishments.[3]

Finally, we apply the model to understand how population aging in Japan influences aggregate output, wage inequality, and establishment concentration. We simulate the evolution of Japan's demographic structure over the next 40 years, assuming the fertility rate remains

[3]The social planner does not face reallocation frictions when choosing the optimal composition of workers, since these costs are not technological constraints. As a result, the gain from the decentralized equilibrium to the optimal distribution of worker characteristics includes both the gain from eliminating frictions and the gain from moving to the optimal composition of worker skills and experience in the absence of frictions.

at 1.3. To isolate composition effects, we hold the total workforce constant. Despite the potentially large gains from reallocating workers toward the optimal workforce composition, we find that population aging has only modest effects on aggregate output. It stays within about 0.6% of its initial level. This result reflects two offsetting forces. On the one hand, population aging reduces the supply of young workers, who are more likely to be high-skill. On the other hand, it increases the supply of highly productive experienced workers. As these effects largely offset, aging has only a limited impact on aggregate productivity.

Population aging also affects wage inequality and market concentration. First, the wage premium of high-skill workers increases substantially over time, as workers mainly learn when they are young (with less than 10 years of experience), and aging reduces the relative supply of high-skill workers. Aging hurts high-productivity establishments, as their comparative advantage lies in utilizing workers' skills rather than experience. A shift in the distribution of workers toward older workers raises the skill premium, hurting high-productivity establishments. As a result, their employment and output shares shrink. Using evidence from Japanese prefectures, we show that the employment and sales shares of the largest establishments increase more in regions experiencing faster growth in the working-age population and declines in average age.

Related Literature Understanding how establishments choose a portfolio of workers contributes to several important areas of research. First, we relate to a literature studying the impact of population aging on economic outcomes. Our paper is closely related to papers that examine the effects of population aging on firm dynamics, including Peters and Walsh (2021), Hopenhayn et al. (2022), and Karahan et al. (2024). These papers examine how the growth rate of the labor force affects firm entry and exit (i.e., dynamism), innovation, and the labor share.[4] We complement this work by abstracting from changes in the overall size of the working-age population and focusing on how shifts in the composition of the labor force affect aggregate outcomes for workers and establishments. Specifically, we show how an increase in the relative supply of experienced workers affects the wage premiums of skilled and experienced workers, product market concentration, and aggregate output, with the latter crucially depending on the allocation of heterogeneous workers across establishments.

[4]Although related, population aging, old-age dependency, and fertility decline capture distinct dimensions of demographic change and need not move together. For example, in 2024, the Republic of Korea had a higher median age (45.5) than Sweden and Belgium (41.1 and 42.0), yet a lower old-age dependency ratio (0.294 versus 0.386 and 0.381), reflecting differences in cohort structure and longevity.

Second, we relate to the literature on firm dynamics that features the interplay between firms and workers. Recent papers in this literature, including Engbom et al. (2019), Gregory (2020), Bagga (2022), and Kim (2023), focus on workers who are either homogeneous or heterogeneous in a single dimension: skill. We emphasize the importance of worker heterogeneity along both work experience and skill. In particular, work experience is of special interest because it links establishments' workforce portfolios to the demographic changes associated with population aging. Our paper complements this literature by showing that establishments have heterogeneous exposure to worker skills and experience with implications for how population aging and reallocation frictions affect aggregate productivity, product market concentration, and wage inequality.

Finally, our paper contributes to the literature on how firing costs affect employer outcomes. Prior studies show that firing costs influence employment levels and distort hiring and separation decisions, reducing aggregate output and productivity (Hopenhayn and Rogerson, 1993). Extensions incorporate capital accumulation, search frictions, and firm dynamics (Veracierto, 2001; Alvarez and Veracierto, 2001; Poschke, 2009). We contribute by introducing heterogeneity in establishments' employment decisions and showing that reallocation frictions have important distributional effects across firms and workers.

2 Worker Composition in Japanese Establishments

In this section, we describe the Japanese establishment dataset and document how workforce composition varies across establishments. We show that larger establishments employ younger workers but have higher average tenure. Consistent with this pattern, they disproportionately hire younger workers and separate from older, longer-tenured workers. Japan provides an ideal setting for this analysis because of its aging population and shrinking working-age workforce, while the establishment survey contains detailed tenure information that allows us to measure worker tenure accurately.

2.1 Japanese Establishment Data

We use two datasets from Japan's Ministry of Health, Labour and Welfare. The first is the Basic Survey of Wage Structure (BSWS), an annual establishment survey covering roughly

Table 1: Summary statistics of the Basic Survey of Wage Structure (2005–2011)

Panel (a): Establishment-level Statistics								
	mean	count	sd	p10	p25	p50	p75	p90
Employment	90.14	377,740	293.2	8	13	26	63	164
Share regular workers	0.671	377,740	0.339	0.100	0.400	0.800	1	1
Panel (b): Worker-level Statistics								
	mean	count	sd	p10	p25	p50	p75	p90
Age (years)	41.12	8,505,380	13.35	23	30	41	52	59
Experience (years)	21.58	6,436,465	12.90	5	11	21	32	40
Tenure (years)	10.22	8,313,284	10.38	0	2	6	16	27
Education (years)	13.32	6,436,465	2.025	12	12	12	16	16
Annual Wage	31007.4	8,505,380	20286.1	8148	16308	27972	41796	57420
Annual Bonus	6662.2	8,505,380	9289.9	0	0	3170	10357	18230
Total Compensation	37669.5	8,505,380	27872.3	8280	16956	32266	51396	74532

Notes: Summary statistics are calculated as follows. In Panel (a), each establishment-year observation is treated as a unit of observation, and summary statistics are computed across establishment-year observations. In Panel (b), each worker-year observation is treated as a unit of observation, and summary statistics are computed across worker-year observations. Experience is defined as age minus years of education minus six. Annual wages, annual bonuses, and total compensation are measured in units of ten thousand yen (approximately USD 100 at 2005-2011 exchange rates). Non-regular workers include part-time, temporary, and contract workers.

55,000 establishments and 1.28 million workers per year. We use data from 1998-2011, focusing primarily on 2005-2011, when the survey was redesigned to improve coverage of non-regular and fixed-term workers. The BSWS provides detailed information on wages, hours, experience, education, tenure, and occupation, with no wage top- or bottom-coding.

We link establishments over time using identifiers from the Economic Census of Japan. Because establishments and workers rotate across survey waves, the data are effectively repeated cross-sections. We observe establishment age for 1998-2001 and 2004-2011 and parent-firm employment size categories. Table 1 reports summary statistics. The average establishment employs 90 workers, 67% of whom are regular employees. The average worker is 41.1 years old, with 21.6 years of experience and 10.2 years of tenure, substantially longer than in the United States. We also report annual salary, defined as the sum of wages and bonuses.

Our second dataset is the Employment Trend Survey (ETS), covering 2005-2007. The ETS follows a rotating sample of about 11,000 establishments and 45,000 hires and separations

Table 2: Summary Statistics of the Employment Trend Survey (2005–2007)

Panel (a): Establishment-Level Statistics								
	Mean	Count	SD	P10	P25	P50	P75	P90
Employment	351.1	64,161	888.3	9	22	81	471	915
Hired	24.05	64,161	117.5	0	0	3	15	56
Internal Transfers In	7.024	64,161	47.55	0	0	0	2	14
Separated	22.43	64,161	110.1	0	1	4	17	52
Internal Transfers Out	7.612	64,161	40.56	0	0	0	3	15
Panel (b): Hired Workers								
	Mean	Count	SD	P10	P25	P50	P75	P90
Age (years)	33.56	250,544	13.17	22	22	32	42	57
Job-market entrant or not	0.226	250,544	0.418	0	0	0	0	1
Education (years)	13.55	250,544	2.045	12	12	14	16	16
Panel (c): Separated Workers								
	Mean	Count	SD	P10	P25	P50	P75	P90
Age (years)	41.43	266,894	14.94	22	27	37	57	62
Education (years)	13.06	266,894	2.084	12	12	12	16	16
Tenure with previous employer (years)	8.167	266,403	8.299	0.25	0.75	3.50	20	20

Notes: In Panel (a), each establishment-half-year observation is treated as a unit of observation. In Panels (b) and (c), each worker-half-year observation is treated as a unit of observation. Summary statistics are computed across all observations within each panel. In the survey, age and tenure with the previous employer are reported as categorical variables. Age is grouped into intervals (below 19, 20–24, 25–29, ..., 60–64, and 65 or above), while prior-employer tenure is grouped into intervals (below 0.5 years, 0.5–1 year, 1–2 years, 2–5 years, 5–10 years, and 10 or more years). We convert these categories into continuous measures using the midpoint of each interval, with appropriate assumptions for the open-ended categories.

semiannually. It provides establishment-level information on employment, hiring, and separations, as well as worker-level demographic and employment histories. Table 2 shows that hiring and separation rates are substantial, with the average establishment employing 351 workers, hiring 24 workers, and separating from 22 workers over a six-month period. Hired workers are younger and more educated than separated workers, while separated workers display substantial variation in tenure.[5]

[5]Because the establishments covered by the ETS are disproportionately large, the average characteristics of newly hired and separated workers are not directly comparable to those in the overall workforce.

Table 3: Average Worker Characteristics by Two Size Bins (2005–2011)

Panel (a): All workers and all establishments					
Size Bin	Age	Experience	Tenure	Wage	Education
1	41.14	21.93	11.19	35435.56	13.20
2	40.41	20.83	13.46	42481.81	13.58
Panel (b): Workers younger than 61 and establishments before 1985					
Size Bin	Age	Experience	Tenure	Wage	Education
1	40.60	21.35	12.46	36602.95	13.25
2	40.12	20.59	14.61	43473.37	13.53

Note: Time period: 2005–2011. Unit of average wage: one hundred yen (or one USD). We divide establishments into two size bins. Size bin 1: employment ≤ 115; size bin 2: employment ≥ 116. Each worker-year observation is treated as a unit of observation.

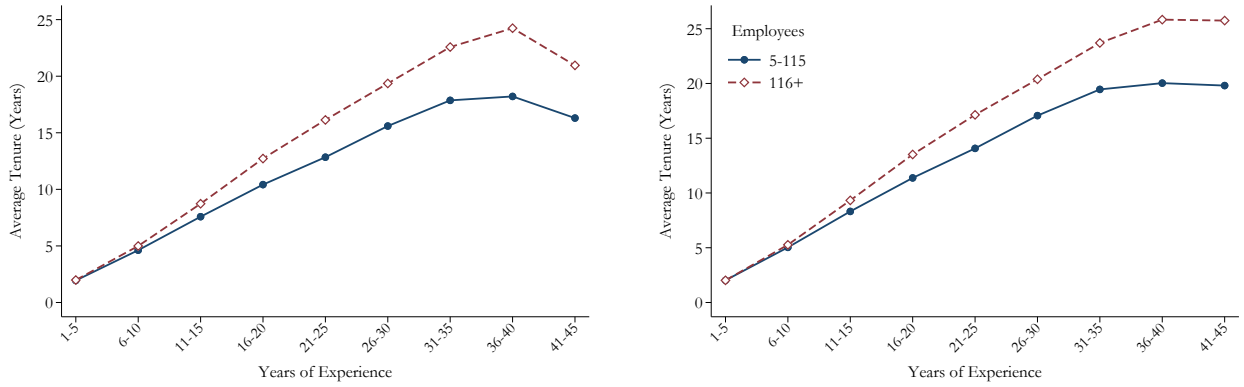
2.2 Large Establishments: Younger Workers with Longer Tenure

We next examine how these workers are distributed across establishments of varying sizes using the BSWS. We first divide establishments into two groups: those with employment of less than 116 and those with employment of 116 or more. This cutoff is chosen to keep both a sizable fraction of establishments and workers in each group. With this cutoff, approximately 20% of establishments employ more than 115 workers, while roughly 20% of workers are employed by establishments with fewer than 115 workers. Panel (a) of Table 3 documents that larger establishments employ workers who are about 0.7 years younger and have about 1.1 years less experience than those in smaller establishments. But these workers have roughly two more years of tenure with their current employer. This is surprising, as younger workers typically have shorter tenure.

Panel (a) of Figure 1 plots experience-tenure profiles for the two establishment size groups, illustrating that this pattern generates differential experience–tenure profiles across establishments. The experience–tenure profile is much steeper for workers in larger establishments. While young workers have similar tenure across establishments, middle-aged and older workers in larger establishments have substantially longer tenure. For example, workers with 36–40 years of experience in the larger establishments have about 24 years of tenure, compared with just 17 in the smaller establishments.

These findings relate to evidence that workers in large firms have lower separation rates

Figure 1: Experience-Tenure Profiles (2005–2011)



(a) Full Sample

(b) Exclude Old Workers and Young Establishments

Notes: Experience is defined as age minus (length of education + 6). Experience bins are defined as: bin 1 (1–5 years), bin 2 (6–10 years), ..., bin 8 (36–40 years), bin 9 (41–45 years). Panel (a) shows the full sample. Panel (b) includes only workers younger than 61 years old and establishments established before 1985.

than those in small firms in the U.S., reflecting the greater variety of jobs within large firms (Papageorgiou, 2018). However, differences in separation rates alone cannot account for our finding that large establishments have younger workers, as even if establishments hire from the same age distribution, longer tenure at large establishments would instead imply an older workforce. Bagga (2022) provides further support for our mechanism by showing that larger firms have stronger internal job ladders, consistent with young, high-skilled workers being more likely to enter and remain in large firms.

The pattern documented in Table 3 and Figure 1 are robust to different measurement choices. First, the results are robust to restricting the sample by time and worker type. Tables A.1 and A.2 and Figure A.1 in the appendix substantiate that the tenure–experience patterns become more pronounced when restricting the sample to the 1998–2004 period or to permanent workers. Second, Tables A.3 and A.4 and Figure A.2 in the appendix show that the tenure–experience patterns hold when we divide establishments into five equal-sized bins. Third, Figure A.3 in the appendix shows that the differential experience tenure profiles hold when using age rather than imputed experience.

Second, we show that our results are not driven by the age of establishments in the sample or the retirement age of workers in Japan. Establishment age is an important confounding factor, both because smaller establishments tend to be younger, so their workers naturally

have shorter tenure and because younger establishments are more likely to employ younger workers (Ouimet and Zarutskie, 2014). In Japan, substantial worker churn occurs after 60, when the lifetime employment system ends with retirement. To account for the influence of establishment age and worker retirement, we recalculate average worker experience and tenure in Panel (b) of Table 3 by restricting our sample to workers younger than 61 and establishments established before 1985 (excluding young establishments). The experience-tenure profiles with this sample are shown in Panel (b) of Figure 1. The documented patterns continue to hold in this sample, which we use to calibrate the model.

We conduct three additional checks to substantiate that our empirical findings are not driven by establishment age or workers' retirement age. One, regression results in Appendix Table A.5 show that larger establishments employ younger workers with longer tenure even after controlling for establishment age and various fixed effects. Two, the experience-tenure patterns hold within establishment-age bins across different sizes (see Appendix Figures A.4). Three, the differential experience-tenure profiles across establishments persist within these establishment age bins (see Appendix Figure A.5).

Third, given that modern economies are dominated by multi-establishment firms, we examine whether these patterns hold at the *firm* level.[6] Using firm-level employment data, we group establishments by the size groups of their parent firms and replicate the analysis, excluding young establishments and limiting workers to age 60 or younger. Appendix Table A.6 shows that larger firms employ younger workers with longer tenure on average, and Figure A.6 confirms that the steeper experience–tenure profile persists at the firm level.

The empirical findings suggest that large and small establishments follow different workforce adjustment strategies. Large establishments disproportionately hire young workers and separate older workers, while small establishments hire relatively fewer young workers and are more likely to recruit middle-aged and older workers. As a result, experienced workers leaving large establishments often migrate to smaller establishments, generating systematic differences in workforce age composition across establishments.

We provide evidence on this mechanism using the ETS. First, because the ETS reports employment by age group, we confirm that larger establishments employ younger workers. Panel (a) of Table 4 shows that workers in establishments with more than 500 employees are, on average, 2.7 years younger than those in establishments with 5–29 employees, con-

[6]Tenure information in the dataset is reported at the firm level, making it natural to aggregate establishments to firms to study how employer size affects average tenure.

Table 4: Worker Age and Establishment Size in the Employment Trend Survey (2005–2007)

Panel (a): Incumbent Workers		
Size Bin	Age (unweighted)	Age (weighted)
1	42.37	42.18
2	42.13	42.01
3	41.55	41.41
4	39.70	39.38

Panel (b): Newly Hired Workers			
Size Bin	Age	Work Experience	Share of Job-market Entrants
1	34.93	22.12	12.3%
2	35.04	22.03	16.7%
3	34.65	21.25	21.1%
4	31.07	16.88	29.8%

Panel (c): Newly Separated Workers			
Size Bin	Age	Work Experience	Tenure with Former Employer
1	39.39	26.78	5.56
2	40.83	28.14	6.17
3	41.81	28.85	7.96
4	41.67	28.19	10.05

Note: In Panel (a), each establishment-half-year observation is treated as a unit of observation. In Panels (b) and (c), each worker-half-year observation is treated as a unit of observation. Experience is defined as age minus years of education minus six. Size bin 1: employment $\in [5, 29]$; Size bin 2: employment $\in [30, 99]$; Size bin 3: employment $\in [100, 499]$; Size bin 4: employment ≥ 500. We use establishment employment as weights when calculating the weighted average age.

sistent with the patterns documented in the BSWS.[7] Next, Panel (b) of Table 4 confirms that newly hired workers are younger and have less prior work experience in larger establishments. Moreover, larger establishments hire a disproportionately greater share of labor market entrants. Finally, Panel (c) shows that newly separated workers in larger establishments are older and have accumulated more labor market experience and longer tenure with their previous employers. These findings corroborate our main results and support the proposed hiring and separation mechanisms underlying these patterns.

[7]The correlation between log employment and average worker age is -0.111 and highly statistically significant. Larger establishments also have lower employment shares of workers aged 50 and above and higher employment shares of workers younger than 30.

3 Model: Efficient Worker Composition

This section presents a framework for understanding the efficient composition of workers and their allocation across establishments. We begin with a static model to analyze how shifts in the labor supply composition across worker groups affect labor allocations and aggregate output. We then extend the framework to a dynamic setting, incorporating the costs of reallocating labor across establishments, and examine how these costs influence both allocations and efficiency.

3.1 Static Framework

Consider an economy with a continuum of establishments indexed by productivity $z \sim \Phi(z)$. Each establishment produces a distinct good using a portfolio of workers drawn from a finite set of groups $g \in \mathcal{G}$, where a group collects workers with the same skill and experience. Workers in different groups are imperfect substitutes, so an establishment employs several groups rather than specializing in one. Labor markets clear group by group, so every establishment faces the same wage w_g for a worker of type g. The economy contains a mass L of workers who each inelastically supply one unit of labor, with L_g workers in group g.

An establishment with productivity z takes the wages $\{w_g\}$ as given and chooses its portfolio to maximize profits,

$$\max_{\{l_g \geq 0\}} \; F(\{l_g\}; z) - \sum_g w_g l_g.$$

When F is increasing and concave, the solution is unique and satisfies the familiar condition that each group is hired until its marginal product equals its wage,

$$\frac{\partial F(\{l_g\}; z)}{\partial l_g} = w_g, \qquad \forall z, g. \tag{1}$$

Let $l_g^c(z)$ denote the resulting demand. A competitive equilibrium is a set of wages and allocations such that every establishment satisfies (1) and each labor market clears,

$$\int l_g^c(z)\, d\Phi(z) = L_g, \qquad \forall g. \tag{2}$$

Equilibrium wages equalize the marginal product of a given group across establishments, but not across groups: a scarce or highly productive group commands a higher wage. This allocation is efficient once we take the supply of each group L_g as fixed. The central question of the paper is what happens when that supply is not fixed. Demographic policy, together with the slow-moving forces of population aging and declining fertility, changes the relative sizes of worker groups. We therefore ask how output would change if the economy could reshape the composition of its workforce.

Optimal Workforce Composition

To answer this question, we consider a planner who holds total labor L fixed but is free to change how it is divided among groups. The planner maximizes output subject only to the aggregate labor constraint,

$$\max_{\{l_g(z)\geq 0\}} \int F(\{l_g\}; z)\, d\Phi(z) \quad \text{s.t.} \quad \int \sum_g l_g(z)\, d\Phi(z) = L.$$

The first-order condition now equates the marginal product of labor across every establishment *and* every group to a single shadow value λ. The contrast with competitive equilibrium is precisely the contrast between (1), where marginal products are equalized only within each group, and a problem in which they are equalized across groups as well.

Two results follow, both proved in Appendix B.1. First, optimal composition can only increase output.

Proposition 1. *The planner's output is at least as large as competitive output, $Y^p \geq Y^c$, with strict inequality whenever two groups earn different wages.*

The argument is one of feasibility: the competitive allocation is always available to the planner, so the planner does at least as well. It does strictly better whenever groups earn different wages, since a wage gap signals unequal marginal products, and output rises by shifting labor toward the group valued more highly at the margin. This same logic fixes the direction of the reallocation.

Proposition 2. *Let g and g' be two groups with $w_g > w_{g'}$. Starting from the competitive allocation, shifting labor from g' to g raises output. If the planner may reallocate only between these two groups, the relative supply of the higher-wage group rises: $L_g^p/L_{g'}^p > L_g/L_{g'}$.*

The size of the gain depends on how far apart wages are and on how poorly groups substitute for one another. If all groups earned the same wage, or if they were perfect substitutes, there would be no gain from optimal composition.

The competitive economy and the fully flexible planner are the two ends of a spectrum. In between lie policies that reshuffle labor among some groups but not others, for instance, shifting workers across experience levels while holding the skill distribution fixed. Appendix B.1 treats these intermediate cases.

Corollary 1. *Reallocating labor freely across all groups yields the largest output gain relative to the competitive equilibrium. Restricting reallocation to a subset of groups yields a weakly smaller gain.*

The gain from full flexibility is therefore an upper bound on what any partial recomposition could achieve.

CES Production

For the quantitative analysis, we take F to be CES with decreasing returns,

$$Y = \int y(z)\, d\Phi(z) = \int \left(\sum_g A_g(z)^{\frac{1}{\sigma}} l_g(z)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\alpha\sigma}{\sigma-1}} d\Phi(z),$$

where $A_g(z)$ is the productivity of a group-g worker at an establishment of type z, $\sigma \in (0, \infty)$ governs substitutability across groups, and $\alpha \in (0, 1)$ controls returns to scale. The match productivities $A_g(z)$ are the objects we recover from the data, and their variation across establishments is what makes worker composition matter.

Profit maximization gives group-g employment and its wage,

$$l_g(z) = \frac{y(z)^{\sigma - \frac{\sigma-1}{\alpha}} A_g(z)}{\int y(z')^{\sigma - \frac{\sigma-1}{\alpha}} A_g(z')\, d\Phi(z')} L_g, \tag{3}$$

$$w_g = \alpha \left(\frac{\int y(z)^{\sigma - \frac{\sigma-1}{\alpha}} A_g(z)\, d\Phi(z)}{L_g} \right)^{1/\sigma}. \tag{4}$$

Employment of a group at an establishment rises with the group's total supply, the group's relative productivity at the establishment, and the establishment's scale. Wages rise with a

group's average productivity and fall with its supply.

The planner's allocation takes the same form,

$$l_g^*(z) = \frac{A_g(z)\, y^*(z)^{\sigma - \frac{\sigma-1}{\alpha}}}{\sum_{g'} \int y^*(z')^{\sigma - \frac{\sigma-1}{\alpha}} A_{g'}(z')\, d\Phi(z')} L. \tag{5}$$

While the competitive equilibrium allocates each group's fixed supply L_g across establishments in (3), the planner allocates the total supply L jointly across groups and establishments. This added flexibility is the source of the output gains.

The CES also sets the direction of reallocation. If $w_g > w_{g'}$, the ratio $l_g(z)/l_{g'}(z)$ is higher under the planner at every establishment, so establishments differ in how far they shift toward the higher-wage group, not in whether they do (Appendix B.1).

The size of the gain depends on how the productivity ratios $A_g(z)/A_{g'}(z)$ compare with the relative supplies $L_g/L_{g'}$. When technology values one group far more than another, but the workforce is not stocked accordingly, moving workers between groups raises output. The gains vanish in two cases: when groups are close substitutes ($\sigma \to \infty$), so marginal products cannot differ, and when relative supplies already track relative productivities, so no group is scarce at the margin. Appendix B.2 makes this precise. The productivities estimated in Section 4 determine how far Japan is from that second case, and the magnitude of the gains.

3.2 Dynamic Model

In this section, we extend the model to examine how dynamic considerations influence optimal labor allocation. This is important as the composition of groups can evolve over time. We consider two dimensions along which worker groups differ: experience and skill. Experience evolves deterministically as workers age, while skill captures the quality of a worker's match to their job and can evolve stochastically over time. Establishments in our model do not differ in industry characteristics, so skill in the model does not capture occupation or industry characteristics or the worker's education; rather, it reflects the worker's ability in their particular job, which is transferable across establishments. The model includes a reallocation friction in the form of a firing cost to illustrate the impact of these dynamic factors on the optimal labor composition.

In the model, heterogeneous establishments employ a portfolio of workers in different expe-

rience and skill groups. This portfolio evolves within establishments over time. The model allows us to characterize these optimal portfolios and the hiring strategies that support them.

Workers The model economy consists of overlapping generations of workers characterized by experience levels $X := (x_1, x_2, \ldots, x_J)$. Workers with the highest experience level, x_J, retire and are replaced by a new cohort entering at experience level $j = 1$. In addition, workers face random labor-market participation shocks that capture unmodeled entry and exit, such as death or illness, marriage, and re-entry into employment (from non-participation or unemployment). As a result, the total mass of workers with experience level j is given by $\mu_j = \delta_{j-1}\mu_{j-1} = \Pi_{i=1}^{j-1}\delta_i$, which determines the size of each cohort relative to the mass of entrants, normalized by $\mu_1 = 1$.[8]

Workers accumulate both experience and skill that are valuable to all establishments. We assume that work experience is acquired *deterministically*, while skill evolves *stochastically*. Specifically, a worker's experience evolves with age, so older workers are necessarily more experienced than younger ones. By contrast, skill evolves through stochastic transitions. We represent skill as $h \in H := \{h_1, h_2, ..., h_I\}$. Let Ψ_j be the $I \times I$ transition matrix, where $\psi^j_{i,i'}$ is the probability of moving from skill level h_i to $h_{i'}$ for an individual with experience x_j. The probability of skill transitions depends solely on experience and is independent of the employing establishment. At birth, workers' skills are distributed according to $g_0(h_i)$, the distribution of individuals with level h_i as they enter the economy.

Establishments Each period, establishments hire a portfolio of workers defined by their skill h_i and experience level x_j. Let L_{ij} denote the number of workers in the (i, j) bin at the beginning of the period (an establishment-level state variable), and l_{ij} denote its employment in that bin at the end of the period, after hiring and firing (an establishment-level choice variable). The model features two types of job separations. First, the establishment may choose to dismiss workers, incurring a reallocation cost equal to a fraction $\tau > 0$ of the worker's wage in that period. Second, a randomly selected fraction $1-\eta$ of incumbent workers separate from the establishment for exogenous reasons, such as job-to-job transitions, and incur no reallocation cost.

The presence of reallocation costs introduces a non-trivial dynamic trade-off: establishments

[8] We allow for the possibility that $\mu_{j'} > \mu_j$ for $j' > j$. This is a shortcut to match the relative cohort size in the data and could arise from fluctuations in labor-force participation by age.

must weigh the benefits of moving toward their frictionless optimal employment level against the costs of adjustment. This creates an *inaction region*, a range of employment levels where the establishment optimally maintains its current workforce. Two factors influence the characterization of these regions. First, skill evolves stochastically, creating uncertainty about the future composition of the workforce regardless of current choices. Second, worker groups are interdependent: employment decisions for one group alter the establishment's future state and influence the optimal choices of other groups through the production function. We now formalize this problem and derive the optimal policy.

An establishment with productivity z chooses labor l_{ij} to maximize the value function,

$$V(z, \{L_{ij}\}) = \max_{\{l_{ij}\}} F(\{l_{ij}\}, z) - \sum_{ij} w_{ij} l_{ij} - \sum_{ij} \eta\tau w_{ij} (L_{ij} - l_{ij})^+ + \beta V(z, \{\tilde{L}_{i'j+1}\}), \tag{6}$$

where the third term is the reallocation friction. The operator $(\cdot)^+ = \max\{\cdot, 0\}$ is active only when the establishment shrinks the (i, j) bin: if $l_{ij} < L_{ij}$ it dismisses $L_{ij} - l_{ij}$ workers and pays τw_{ij} for the η fraction of incumbent workers that remained, and if $l_{ij} \geq L_{ij}$ it pays nothing. This kink at $l_{ij} = L_{ij}$ is what generates the inaction region below, since the marginal cost of a worker rises from $w_{ij}(1 - \eta\tau)$ to w_{ij} as employment crosses the existing stock.

The employment stock evolves stochastically (depending on the current employment choices):

$$L_{i'j+1} = \delta_j \sum_i \psi^j_{i,i'} l_{ij}; \quad \forall\ j \in \{1, 2...J-1\}.$$

Three points are worth mentioning here. First, the state variables of employment stock, L_{ij}, form a *two-dimensional* vector of employment stocks and therefore have high dimensionality. Because of the curse of dimensionality, we do not solve the value function using standard iteration methods. Instead, we characterize analytical properties of the choice functions and use them to solve the problem. Second, although individual workers face uncertainty about future skill levels, the establishment knows its next-period employment stock once it has made its current-period employment choices, because employment stocks are treated as continuous variables. This is why the continuation value in equation (6) is deterministic. Third, the establishment still faces uncertainty when choosing l_{ij} in the current period because skill levels evolve stochastically, and different skill subgroups of workers in experience bin $j + 1$ may face different firing rates in the next period, some of which may be zero.

The first-order conditions (FOCs) of the optimization problem (6) are:

$$[l_{ij}]: \ \frac{\partial F}{\partial l_{ij}} = w_{ij}\left(1 - \eta\tau \mathbb{1}\left[L_{ij} > l_{ij}\right]\right) - \beta\delta_j \sum_{i'} \psi^j_{ii'} \frac{\partial V(z, \tilde{L}_{i'j+1})}{\partial L_{i'j+1}} \quad \forall i, j.$$

The Envelope Condition in the current and next periods implies that

$$\frac{\partial V}{\partial L_{ij}} = -\eta\tau w_{ij} \mathbb{1}\left[L_{ij} > l_{ij}\right].$$

Together, the FOCs imply

$$\frac{\partial F}{\partial l_{ij}} = w_{ij}\left(1 - \eta\tau \mathbb{1}\left[L_{ij} > l_{ij}\right]\right) + \beta\eta\delta_j\tau \sum_{i'} \psi^j_{ii'} w_{i'j+1} \mathbb{1}\left[\tilde{L}_{i'j+1} > l_{i'j+1}\right]. \tag{7}$$

Comparing (7) with (1) reveals three differences in hiring when reallocation costs are present. First, if employment l_{ij} falls below the existing stock L_{ij}, an extra hire spares the establishment a dismissal and lowers the effective marginal cost of labor by $\eta\tau w_{ij}$. Second, if next period's optimal employment $l_{i,j+1}$ is below the projected stock $\tilde{L}_{i,j+1}$, the establishment faces a potential reallocation cost. Third, the current choice l_{ij} feeds into skill subgroups next period, and whether each of those workers ends up in the firing region depends on choices the establishment has not yet made. This interdependence, rather than uncertainty, is what makes the problem non-trivial.

We now derive the establishment's optimal hiring and firing strategies, which imply an inaction region for employment given the current stock L_{ij}. Let FC_{ij} denote the expected continuation cost associated with choosing l_{ij} workers,

$$FC_{ij} := \beta\delta_j\eta\tau \sum_{i'} \psi^j_{ii'} w_{i'j+1} \mathbb{1}\left[\tilde{L}_{i'j+1} > l_{i'j+1}\right].$$

If the establishment hires workers ($l_{ij} \geq L_{ij}$), the FOC for the (i, j) worker group becomes

$$\frac{\partial F}{\partial l_{ij}} = w_{ij} + FC_{ij}, \tag{8}$$

which yields a unique solution for l_{ij}, denoted by l^{lb}_{ij}. Next, if the establishment fires workers

$(l_{ij} < L_{ij})$, the FOC is

$$\frac{\partial F}{\partial l_{ij}} = w_{ij}(1 - \eta\tau) + FC_{ij} \tag{9}$$

which yields a unique value l_{ij}^{ub}. Since $l_{ij}^{lb} < l_{ij}^{ub}$, the inaction region is defined by these bounds; if $L_{ij} \in [l_{ij}^{lb}, l_{ij}^{ub}]$, the establishment chooses neither to hire nor fire workers.

The establishment's labor choice, governed by equations (8) and (9), depends on employment decisions from both the previous and the next period. Moreover, the choice of l_{ij} depends not only on the current wage w_{ij} but also on future wages for the same and adjacent worker groups, since these influence expected firing costs. Focusing on the steady state, where wages are constant over time, the establishment only needs to form correct expectations of the current period's prevailing wages and the expected future firing probabilities.

Labor Markets Labor markets in this economy are indexed by (i, j), where i denotes the skill bin h_i and j the experience bin x_j. In equilibrium, labor demand must equal labor supply in each (i, j) market:

$$\int l_{ij}(z) d\Phi(z) = g(h_i, x_j); \quad \forall\ (i, j), \tag{10}$$

where $l_{ij}(z)$ denotes the number of workers with skill h_i and experience x_j employed by an establishment with productivity draw z, and $g(h_i, x_j)$ is the labor supply of workers with skill h_i and experience x_j. The wage schedule enters (10) through demand, since $l_{ij}(z)$ solves (8) or (9) given $\{w_{ij}\}$. These $I \times J$ conditions jointly determine the $I \times J$ wages.

In Appendix B.5 we develop an algorithm to solve for the equilibrium. The difficulty is that FC_{ij} depends on whether the establishment will fire in bin $j+1$ next period, which, in turn, depends on the choice $l_{i'j+1}$ it has not yet made. The algorithm, therefore, guesses the firing indicators $\mathbb{1}[\tilde{L}_{i'j+1} > l_{i'j+1}]$, solves the model forward from the youngest bin to the oldest, given that guess, and updates until the indicators implied by the solution reproduce the ones assumed. Two bins anchor the recursion: entrants carry no existing stock and so no firing cost, and workers in x_J retire, so their continuation cost is zero.

4 Model Calibration and Fit

In this section, we apply our model to Japanese worker–establishment data. We begin by describing the calibration procedure and the moments we target. We use the static model with no reallocation costs as our baseline case for calibration and then study how dynamic features influence the optimal composition of workers across establishments.

4.1 Calibration Specification

For our quantitative exercises, we parameterize both the production function and the stochastic learning process. We assume there are K types of establishments that differ in productivity. Each establishment of type $k \in \{1, ..., K\}$ operates a CES production function with decreasing returns to scale given by:

$$F\left(\{l_{ij}, z_k\}\right) = \left(\sum_{ij} A_{ijk}^{\frac{1}{\sigma}} l_{ijk}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\alpha\sigma}{\sigma-1}},$$

where l_{ijk} denotes the labor input from worker type (i, j) employed by a type-k establishment. The parameter $\sigma \in (0, \infty)$ represents the elasticity of substitution across worker types, capturing imperfect substitutability in production. The parameter $\alpha \in (0, 1)$ governs the degree of decreasing returns to scale. This specification also incorporates establishment-worker-specific productivity terms, A_{ijk}, that capture match-specific complementarities between worker and establishment types. Identifying these productivity parameters from observed data moments is a central component of our quantitative analysis.

We specify the skill transition process Ψ^j as a learning–depreciation chain, in which workers move up or down the skill ladder by at most one step per period. This parsimonious formulation captures two forces: skill acquisition through experience-dependent learning, with probability p_j, and skill depreciation through knowledge obsolescence, with constant probability q. We impose reflecting conditions at the two ends of the ladder, so workers at the lowest level cannot depreciate further, and workers at the highest level cannot advance.[9]

[9]For interior levels $1 < i < I$, learning and depreciation are independent, so $\psi^j_{i,i+1} = p_j(1-q)$, $\psi^j_{i,i-1} = q(1-p_j)$, and $\psi^j_{i,i} = 1 - p_j - q + 2p_jq$, where a worker stays put if both events occur or neither does. At the boundaries, the same probabilities apply, and a move off the ladder leaves the worker in place: $\psi^j_{1,2} = p_j(1-q)$ and $\psi^j_{1,1} = 1 - p_j(1-q)$ at the bottom, $\psi^j_{I,I-1} = q(1-p_j)$ and $\psi^j_{I,I} = 1 - q(1-p_j)$ at the top.

4.2 Target Moments and Estimation

We estimate the model using Japanese establishment data from 2005–2011. Workers are classified into $J = 9$ experience bins: $x_j \in \{1\text{-}5, 6\text{-}10, ..., 36\text{-}40, 40+\}$ years, and establishments are divided into two size categories: small (fewer than 116 employees) and large (116 or more employees), indexed by $k \in \{\mathrm{S}, \mathrm{L}\}$. Skill is unobserved and takes two values, $h_i \in \{\mathrm{L}, \mathrm{H}\}$, representing low-skill and high-skill workers, respectively. With $I = 2$, both skill levels are boundary states, so the transition matrix reduces to a single pair of probabilities: a low-skill worker with experience x_j becomes high-skill with probability $p_j(1-q)$, when she learns and does not lose knowledge, and a high-skill worker becomes low-skill with probability $q(1-p_j)$. All wages are normalized by the economy-wide average wage.

We first set the values of some parameters externally from our calibration approach. We set the discount factor to $\beta = 0.815$, consistent with a five-year period and an annual discount rate of 0.96. The elasticity of substitution is set to $\sigma = 2.5$, implying moderate substitutability across worker types.[10] Following Hopenhayn et al. (2022) and Karahan et al. (2024), we set the returns-to-scale parameter to $\alpha = 0.6$, reflecting the labor share in production. The initial skill distribution, $g_0 = [0.52, 0.48]$, for low- and high-skill workers matches the share of college-educated workers entering the labor force during 2005–2011. Because learning and depreciation both move workers between skill levels, we estimate the learning rates and calibrate the depreciation rate q externally, reporting results for alternative values in Section 5.3. Given an annual human capital depreciation rate of approximately 6% documented in Blundell et al. (2021) and Ma et al. (2024), we set $q = 26\%$, consistent with a model period of five years. For the baseline model, we set reallocation costs to $\tau = 0$, though an infinitesimally small reallocation cost allows us to measure job separations of different types. Section 5.2 assesses the sensitivity of our results to this assumption.

After setting the values for external calibration, we have 45 remaining parameters: 36 match-specific productivity terms A_{ijk} (for $i \in \{L, H\}$, $j \in \{1, ..., 9\}$, $k \in \{S, L\}$), 8 experience-specific learning rates p_j for $j \in \{1, ..., 8\}$ (with $p_9 = 0$ for the most experienced workers), and the exogenous separation rate η. The employment and wage moments impose only 27 independent restrictions on the 36 productivity parameters. In each experience bin they leave

[10] Papers such as Card and Lemieux (2001), Borjas (2003), and Ottaviano and Peri (2012) estimate that the elasticity of substitution across workers in different experience groups is between 2 and 5. The elasticity of substitution across workers with different skill levels is estimated to be between 1.2 and 1.5 in Katz and Murphy (1992) and Card and Lemieux (2001).

one degree of freedom, which governs how high-skill workers sort across the two establishment types. We therefore estimate 18 parameters: one sorting parameter for each experience bin, the eight learning rates, and η. Given these, the productivities follow in closed form, and the model matches the employment and wage moments exactly. As we discuss below, the tenure-experience profiles and the tenure distribution identify the sorting, the learning rates, and η. We restrict the learning probabilities, p_j, and each profile of log productivities, $\log A_{ijk}$, to be single-peaked over the experience distribution.

We target the following four sets of moments. First, we target employment shares l_{jk} in aggregate labor demand, that is, the fraction of type-j workers employed by type-k establishments as a share of the total number of workers. Second, we target the mean wages of type-j workers in type-k establishments, w_{jk}. Third, we target experience-tenure profiles defined as the average tenure by experience bin-j within each establishment type-k. Finally, we target the tenure distribution measured as the distribution of employment across establishment type-k for each worker tenure bin-t.

These moments yield 63 targets for estimation: $J \times K = 18$ employment shares, 18 average wages, 18 tenure-experience observations, and $J = 9$ tenure distribution moments.[11]

We assume that wages increase monotonically with skill: $w_{Hj} > w_{Lj} \quad \forall j \in \{1, ..., 9\}$. This ensures that high-skill workers earn strictly more than low-skill workers with identical experience and is justified on the grounds that high-skill workers could also do low-skill jobs. We also rule out extreme sorting across establishment types and bound the wage dispersion that skill differences can generate within an experience bin and establishment type by the dispersion observed in the data (see details in Appendix B.4).

Denoting the 18 parameters collectively by Θ, we estimate them by minimizing:

$$\hat{\Theta} = \arg\min_{\Theta} \sum_{m} \omega_m \frac{(d_m - m_m(\Theta))^2}{d_m^2 + m_m(\Theta)^2} \tag{11}$$

where d_m and $m_m(\Theta)$ denote data and model moments, respectively. The denominator normalization prevents moments of different scales from dominating the objective function. The weights ω_m ensure equal importance across the four moment categories, with each category receiving a weight of $1/4$. For example, the 18 employment share moments each receive weight $\omega_m = 1/(4 \times 18)$, and the nine tenure distribution moments each receive weight

[11] Within each experience bin, the employment shares sum to that bin's share of the workforce, which is redundant, so only nine of the 18 employment moments are informative and the 63 targets amount to 54.

Table 5: Parameters

Parameter	Description	Value
Simulated Method of Moments		
A_{ijk}	Worker-Establishment Productivities	See Figure 3
p_j	Learning probability	$[0.61, 0.42, 0.14, 0.078, 0.071, 0.068, 0.068, 0.068]$
$1-\eta$	Exogenous Separation Rate	0.0984
Externally Calibrated		
β	Discount Rate (5 years)	0.815
σ	Elasticity of Substitution between Groups	2.5
α	Degree of Returns to Scale	0.6
g_0	Initial Distribution of Skills	[0.52, 0.48]
q	Skill Depreciation/Obsolescence	0.26 (Blundell et al., 2021)
τ	Firing Costs	$\epsilon \to 0$

Note: The table lists the externally calibrated parameters and the ones estimated using the simulated method of moments.

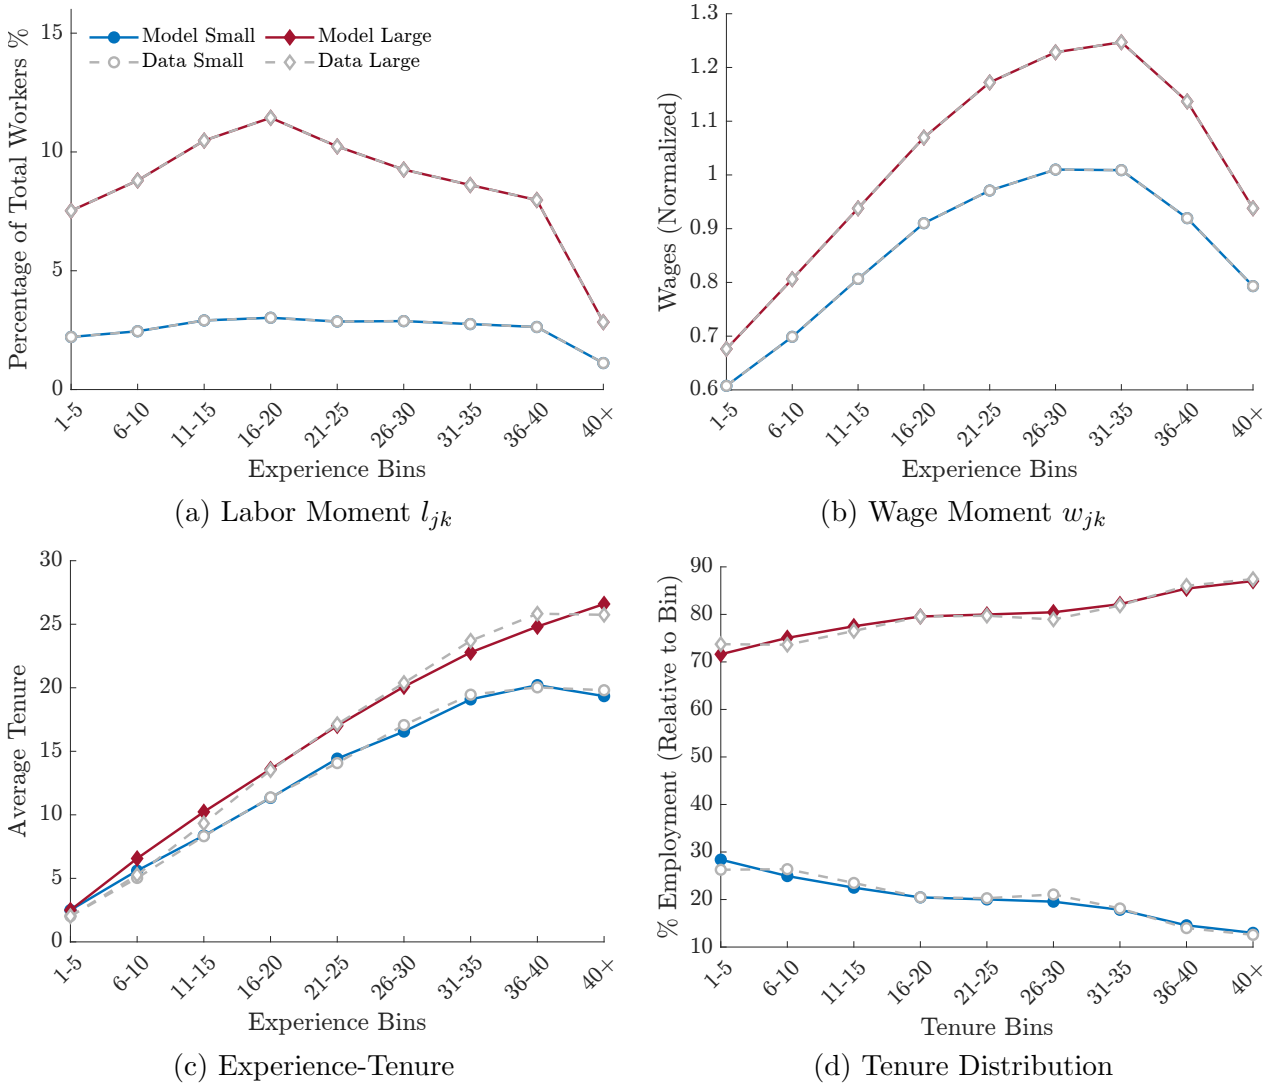
$\omega_m = 1/(4 \times 9)$. Because employment and wage moments are matched exactly, the two tenure categories determine the estimates. Table 5 reports the resulting parameter values.

Moments and Identification Figure 2 compares model predictions with the data. The model matches the employment shares and wages exactly and reproduces the heterogeneous experience-tenure profiles across establishment types. The employment shares by experience and establishment type (Figure 2a) and the wage–experience profiles (Figure 2b) determine the productivities A_{ijk} once we know how skills sort across establishment types. They also pin down the direction of sorting. Because every establishment pays the same wage for the same worker, large establishments can pay more on average within an experience bin only by employing a larger share of high-skill workers. They do not pin down the strength of sorting: strong sorting with a small skill premium and weak sorting with a large premium fit these moments equally well. Appendix B.4 formalizes this argument.

The tenure moments identify the strength of sorting, the learning rates p_j, and the exogenous separation rate η. When skill mixes differ across establishment types, workers whose skill changes must move between them, so sorting shapes the experience-tenure profiles by establishment type documented in Section 2 (Figure 2c). The tenure distribution (Figure 2d) separates learning from exogenous separations: η shifts tenure uniformly, whereas p_j moves the two establishment types in opposite directions where the skill composition changes fastest.[12] At the estimates, the bound on wage dispersion does not bind, so tenure pins down

[12]The calibrated p_j declines with experience while q is constant, so depreciation eventually dominates

Figure 2: Moments Fit



(a) Labor Moment l_{jk}

(b) Wage Moment w_{jk}

(c) Experience-Tenure

(d) Tenure Distribution

Notes: This figure displays the model fit for the first four sets of targeted moments. Panels (a) and (b) show employment shares and average wages by experience group and establishment type. Panel (c) presents average tenure by experience within each establishment type. Panel (d) shows the employment shares of large and small establishments for workers in each tenure bin (i.e., 1-5, 6-10, ...,40+). Grey dashed lines represent data moments; solid lines represent model predictions. Small establishments have fewer than 116 employees; large establishments have 116 or more employees.

sorting among entrants as well.

and skill shares fall among middle-aged and older workers. This helps the model match the differential experience-tenure profiles: large establishments shed these workers, and small establishments absorb them.

4.3 Estimation Results and Hiring/Separation Dynamics

Before proceeding to study the gains from optimal composition, it is informative to document the implications of the model parameterization. This section documents the implications of the estimated parameters for group-level productivities, worker skill dynamics, and the hiring and firing strategies of different establishments.

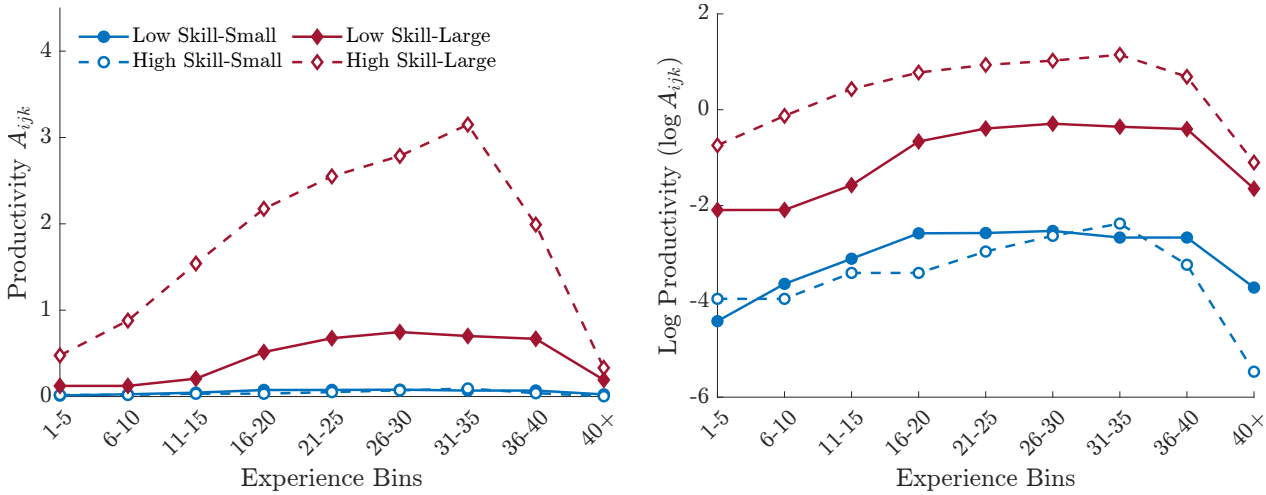
Log Supermodularity in Productivity Parameters Figure 3 shows the estimated match-specific productivities A_{ijk}, a key output of the structural estimation. They follow in closed form from the employment shares and wages, given the estimated sorting of skills across establishment types.

The results reveal two patterns in worker–establishment complementarities. First, productivity levels are higher in large establishments for both skill types, consistent with standard theories of firm heterogeneity. Second, in small establishments, low-skill workers have higher productivity parameters than high-skill workers at most experience levels. Combined with the productivity advantage of high-skill workers in large establishments, this pattern implies strong skill–size complementarity or *log supermodularity* in aggregate production. The sign of this complementarity is a feature of the data rather than of our estimates: large establishments pay more within every experience bin, which, with common wages, requires them to employ relatively more high-skill workers. Its strength reflects the estimated sorting.

Skill and Wages A feature of our approach is that skill is unobserved and inferred from the data. Figure 4 presents the skill composition and wage profiles resulting from our identification strategy. Figure 4a shows a pronounced life-cycle pattern in skill composition: the share of high-skill workers rises from 48% at labor market entry to 67% at 6–15 years of experience, then declines steadily, falling below 50% at 21–25 years and to 24% after 40 years. This inverted U reflects the interplay of learning and depreciation. Early-career workers accumulate skill rapidly, as learning probabilities exceed 50% in the first experience bin, but the constant depreciation rate q eventually dominates, leading to net skill deterioration. Skill in the model should be interpreted as the worker's fit for their given task. These patterns imply a process in which young workers match with tasks early in their career, with the fit then slowly degrading over the life cycle.

Figure 4b shows that the skill wage gap is not monotone over the life cycle. It is 64% at

Figure 3: Estimated Worker-Establishment Productivities across Experience A_{ijk}



(a) Productivity A_{ijk} (b) Log Productivity $\log A_{ijk}$

Notes: Estimated match-specific productivities for skill type $i \in \{L, H\}$, experience group $j \in \{1, ..., 9\}$, and establishment type $k \in \{S, L\}$. Panel (a) shows productivity levels; panel (b) shows log productivities. Each line represents a skill-establishment pair, with productivity varying across experience groups.

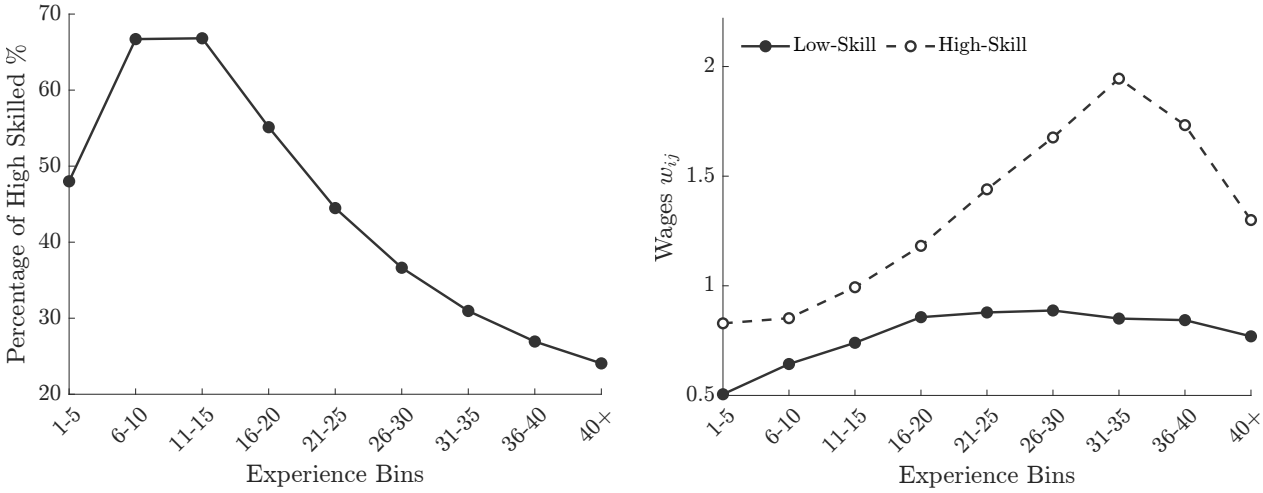
entry and narrows to about a third at 6–15 years of experience, when high-skill workers are most abundant. It then widens to 129% at 31–35 years, as high-skill wages continue to rise while low-skill wages plateau, before falling back to 69% after 40 years. This pattern in life-cycle wage inequality emerges endogenously from the interaction between skill dynamics and match-specific productivity differences.

Hiring and Separation Strategies We use our calibrated model to study the hiring and separation strategies of establishments of different sizes, showing that large, high-productivity establishments disproportionately employ young workers. We then provide evidence for those predictions using the ETS dataset.

Figure 5 highlights the patterns of hiring and firing by worker group across establishments.[13] Panel (a) shows that nearly 80% of newly hired workers in large establishments are job market entrants (experience 1–5), compared with about 50% in small establishments. Small establishments hire a larger share of middle-aged and older workers and employ a dispropor-

[13]We do not include exogenous separations and the associated replacement hiring when calculating the number of fired and hired workers. Because the exogenous separation rate is constant across establishments and worker groups, including them would not change the relative patterns of hiring and firing across establishments. ETS data indicate that between-group worker flows account for approximately 40% of total worker flows and that most worker separations are voluntary (i.e., do not result from dismissals or layoffs).

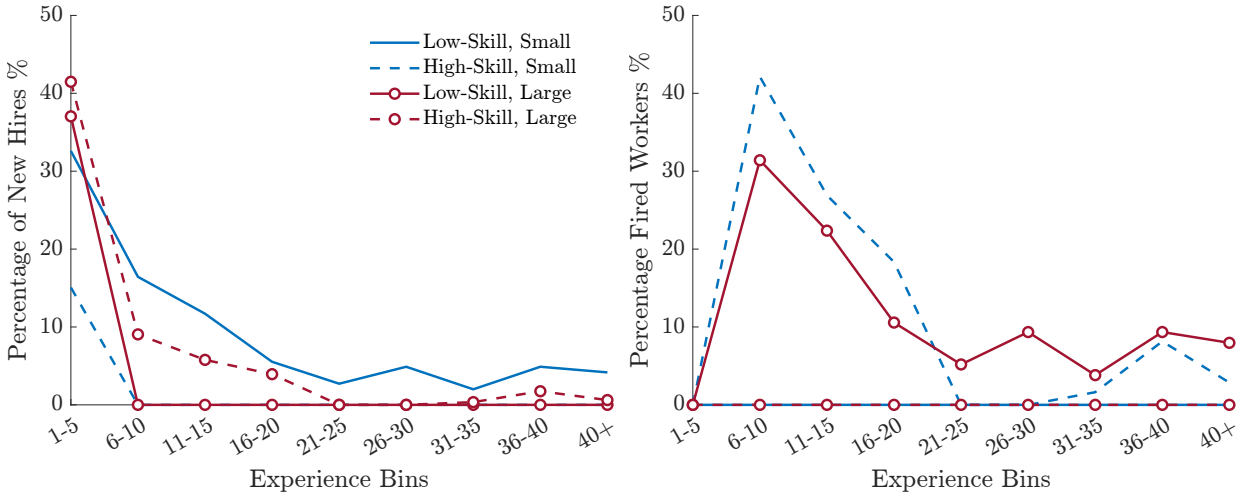
Figure 4: Percentage of High-Skill and Wages $\alpha = 0.6$



(a) Percentage of High Skill (b) Wages by Skill

Notes: Panel (a) shows the share of high-skill workers by experience group. Panel (b) shows wage profiles by skill type across experience groups.

Figure 5: Hiring and Firing Strategies: Up-or-Out Dynamics in Large Establishments



(a) Hiring (b) Firing

Notes: This figure shows the percentage of new hires and fired workers by establishment size and skill level for the case $\tau = 0$. Percentages are calculated relative to total hiring and firing within each firm.

tionate share of unskilled workers.

Panel (b) shows that small establishments fire disproportionately more young, skilled work-

Table 6: Hiring Strategies of Establishments (2005–2007)

	(1)	(2)	(3)	(4)	(5)	(6)
	Age		Experience		Entrant or Not	
Log(Employment)	-0.917***	-0.456***	-1.120***	-0.574***	0.0337***	0.0179***
	(0.0206)	(0.0277)	(0.0213)	(0.0287)	(0.000635)	(0.000865)
Prefecture × time FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × time FE	Yes	Yes	Yes	Yes	Yes	Yes
Size of parent firm FE	No	Yes	No	Yes	No	Yes
Observations	250,535	250,535	250,535	250,535	250,535	250,535
adj. R^2	0.132	0.135	0.144	0.148	0.203	0.206

Notes: Each hired worker-semi-year observation is treated as a unit of observation. Experience is defined as age minus years of education minus six. Robust standard errors are reported. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

ers, whereas large establishments tend to dismiss more experienced, unskilled workers. These differences in hiring and firing behavior stem from calibrated establishment–worker group-specific productivities and result in heterogeneous exposure to changes in worker skills and experience across establishments.

We assess whether the hiring and separation strategies implied by the model are consistent with the ETS data. We first regress the age, prior work experience, and labor-market entrant status of newly hired workers on the log employment of the hiring establishment, controlling for prefecture-by-time, industry-by-time, and parent-firm-size fixed effects.[14] Table 6 shows that larger establishments hire workers who are younger, less experienced, and more likely to be labor-market entrants. We then regress the age, prior work experience, and tenure with the former employer of newly separated workers on the log employment of the separating establishment. As shown in Table 7, workers separating from larger establishments are older, more experienced, and have longer tenure. Overall, these findings are consistent with the model's predictions regarding establishments' hiring and separation behavior.

[14]Establishment fixed effects cannot be included because the ETS is essentially cross-sectional.

Table 7: Separation Strategies of Establishments (2005–2007)

	(1)	(2)	(3)	(4)	(5)	(6)
	Age		Experience		Tenure	
Log(Employment)	0.127***	0.125***	0.0151	0.0652**	0.759***	0.296***
	(0.0228)	(0.0300)	(0.0235)	(0.0310)	(0.0123)	(0.0161)
Prefecture × time FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × time FE	Yes	Yes	Yes	Yes	Yes	Yes
Size of parent firm FE	No	Yes	No	Yes	No	Yes
Observations	266,632	266,632	266,632	266,632	266,142	266,142
adj. R^2	0.089	0.090	0.090	0.090	0.163	0.170

Notes: Each separated worker-semi-year observation is treated as a unit of observation. Experience is defined as age minus years of education minus six. Tenure refers to the worker's tenure with his/her former employer. Robust standard errors are reported. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

5 Results: Gains from Optimal Labor Composition

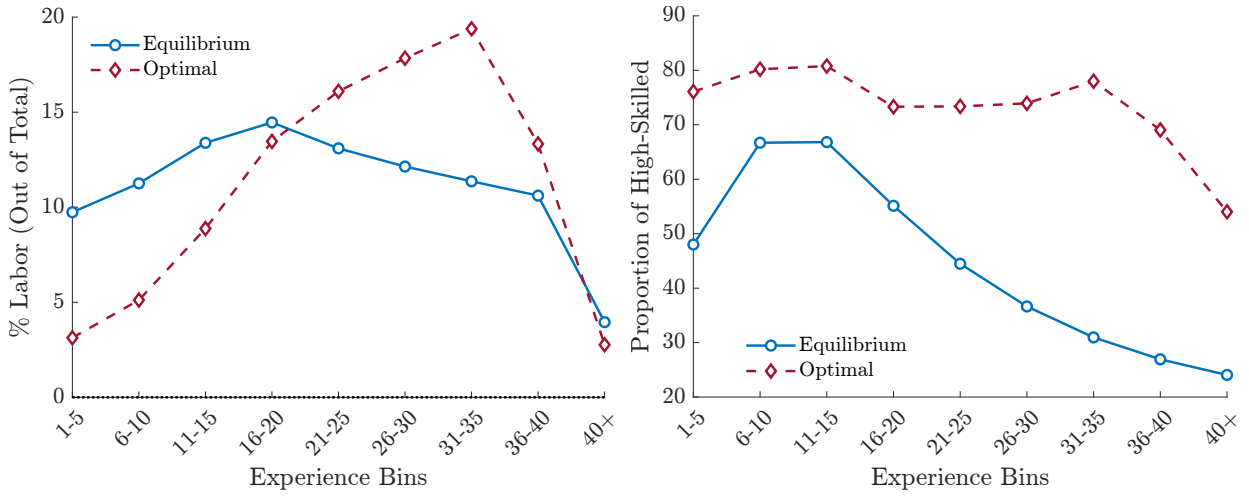
This section quantifies the efficiency gains from optimal labor composition by solving the social planner's problem that takes the estimated productivities A_{ijk} as given and reallocates workers across experience and skill groups to maximize aggregate output, taking estimated productivities as given. By allowing unrestricted reallocation subject only to aggregate labor supply constraints, the exercise provides an upper bound on potential gains and highlights the gap between the decentralized equilibrium and the first-best allocation. It also illustrates the potential scope for policies that influence workforce composition, such as education, fertility, and retirement reforms.

5.1 Implementation and Policy Implications

In the baseline model without reallocation costs, the competitive equilibrium shows no within-group misallocation: labor markets clear at the group level and establishments face no adjustment frictions. The efficiency gap arises solely from the economy's inability to move to the optimal composition of workers in response to productivity differences.

Figure 6 illustrates how the social planner would optimally restructure Japan's workforce to maximize aggregate output. The optimal composition reflects underlying productivity

Figure 6: Optimal Labor Composition: Experience and Skill



(a) Optimal Experience (b) Optimal Skills

Notes: Comparison of labor composition under the competitive equilibrium (solid line with circles) and the planner's allocation (dashed line with diamonds). Panel (a) shows the distribution of workers across experience groups, holding the aggregate skill composition fixed. Panel (b) shows the share of high-skill workers by experience group, holding the experience distribution fixed.

patterns and highlights the demographic shifts that would be most beneficial. Panel 6a compares the optimal experience composition with the current equilibrium. The planner would substantially reduce the share of early-career workers (under 16 years of experience) while increasing the proportion of prime-age workers (16–40 years).

Panel 6b shows that the planner would increase the share of high-skill workers at all experience levels. The gap between the optimal and equilibrium allocations is smallest for workers with 6–15 years of experience and largest for those with 26–40 years, where the optimal share is more than twice the equilibrium level. This pattern suggests potentially larger gains from skill upgrading among older workers, although the model does not account for training costs.

Table 8 quantifies the potential output gains from optimal labor composition in Japan. By jointly optimizing the experience and skill distribution of workers across establishments, the social planner could raise aggregate output by 8.53% (first column). Most of this gain comes from skill: moving to the optimal skill composition alone yields 5.61%, more than twice the 2.10% from moving to the optimal experience composition alone. The joint gain exceeds the sum of the two separate gains, so the two margins are complements.

Table 8: Output Gains from Optimal Labor Composition

Composition Scenario	**Output Gain**			
	$\tau = 0$	$\tau = 0.1$	$\tau = 0.2$	$\tau = 0.3$
No Friction ($\tau = 0$)	–	0.04%	0.16%	0.37%
Optimal Experience Structure Only	2.10%	1.75%	1.55%	1.51%
Optimal Skill Structure Only	5.61%	6.22%	6.94%	7.73%
Optimal Joint Composition	8.53%	8.73%	9.10%	9.59%

Notes: Gains relative to the competitive equilibrium. For $\tau > 0$, only the productivities are re-estimated, and the gains include the gain from removing the reallocation cost (first row).

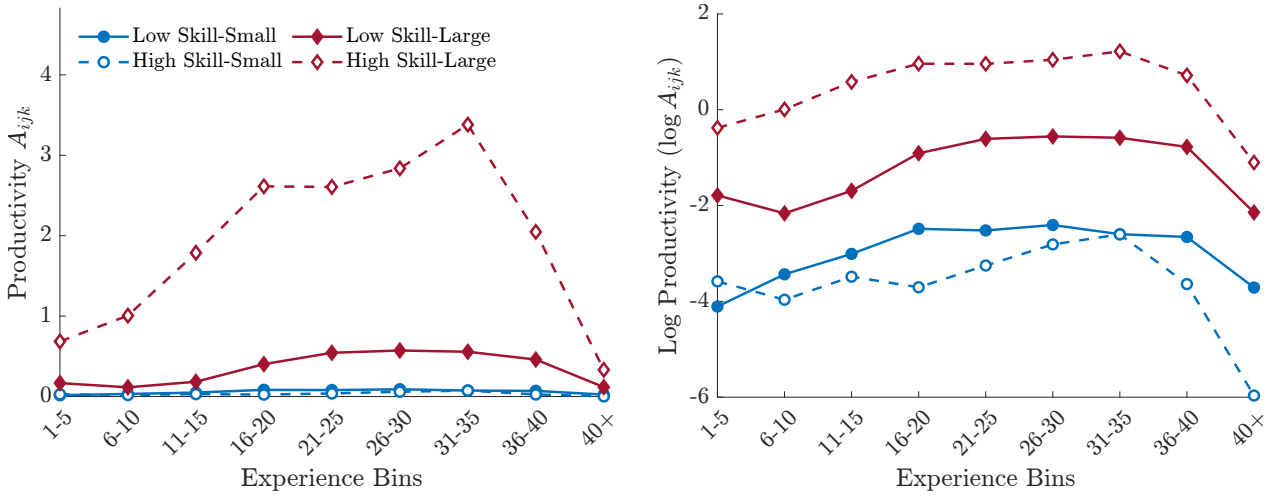
5.2 Dynamic Gains to Optimal Worker Composition

Our baseline model abstracts from labor reallocation frictions. To assess their importance, we recalibrate the model with reallocation costs up to $\tau = 0.3$, re-estimating only the match-specific productivities while holding sorting, learning rates, and η fixed. The model continues to match employment shares, wages, experience-tenure profiles, tenure distributions, and firing rates, but reallocation costs alter the mapping from wages to productivities because workers' marginal products reflect both the current and expected future dismissal costs(equations (8) and (9)). Comparing this dynamic model with the frictionless benchmark quantifies the economic costs of reallocation frictions.

Figure 7 reports the re-estimated productivities. The qualitative patterns remain unchanged: productivity is complementary with establishment size and worker skill, and follows a hump-shaped profile over experience. Quantitatively, reallocation costs widen the productivity gap between high- and low-skill workers at large establishments, particularly among experienced workers. The intuition is that large establishments hire high-skill workers despite the risk that some will lose skill and need to be dismissed in the future, raising the productivity required to justify these hires. Conversely, retaining low-skill workers helps avoid future dismissal costs, reducing the productivity needed to support their employment. Small establishments, which are more likely to separate from high-skill workers, exhibit the opposite pattern.

The last three columns of Table 8 report the gains from optimal workforce composition. At $\tau = 0.2$, the total gain rises to 9.10%, from 8.53% in the frictionless benchmark, with 6.94% attributable to skill and 1.55% to experience. Most of the increase reflects wider productivity differences across skill groups rather than the direct effect of reallocation frictions, as removing these frictions while holding workforce composition fixed raises output by only

Figure 7: Estimated Worker-Establishment Productivities across Experience A_{ijk} for $\tau = 0.2$



(a) Productivity A_{ijk} (b) Log Productivity $\log A_{ijk}$

Notes: Estimated match-specific productivities for skill type $i \in \{L, H\}$, experience group $j \in \{1, ..., 9\}$, and establishment type $k \in \{S, L\}$. Panel (a) shows productivity levels; panel (b) shows log productivities. Each line represents a skill-establishment pair, with productivity varying across experience groups.

0.16%. As τ increases to 0.3, the total gain reaches 9.59%, with the gains shifting further toward the skill margin.

5.3 Robustness to Externally Calibrated Parameters

We assess the sensitivity of our results to the externally calibrated parameters α, σ, and q. We move each parameter 20 percent below and above its benchmark value, holding the others fixed, and re-estimate the model under the benchmark's identification strategy and restrictions, starting from the benchmark estimates. Table 9 reports the gains in each case. Changes in α or σ affect reallocation gains through two channels: a direct effect on how output responds to a given reallocation and an indirect effect through the productivities inferred from the data. For both parameters, these two effects work in opposite directions.

Returns to Scale The decreasing-returns parameter α governs the curvature of the establishment production function. We consider $\alpha \in \{0.48, 0.72\}$ around the baseline $\alpha = 0.6$, holding $\sigma = 2.5$. A lower α implies larger estimated productivity gaps across establishments, as stronger diminishing returns require greater productivity differences to match observed em-

Table 9: Output Gains from Optimal Composition Varying α, σ and q

		α		σ		q	
Composition Scenario	Benchmark	0.48	0.72	2	3	0.208	0.312
Optimal Experience Structure Only	2.10%	1.67%	2.52%	1.73%	2.44%	2.10%	2.10%
Optimal Skill Structure Only	5.61%	4.38%	6.97%	4.85%	6.58%	4.51%	7.32%
Optimal Joint Composition	8.53%	6.68%	10.55%	7.05%	10.34%	7.15%	10.59%

Notes: Output gains from the optimal composition with one parameter at the value in the column header and the others at their benchmark values. Each column re-estimates the model with that parameter changed, under the benchmark's identification.

ployment shares. The log-supermodularity between establishment productivity and worker skill remains robust, since its sign comes from the wage gap between establishment types.

The joint gain rises with α, from 6.68% at $\alpha = 0.48$ to 10.55% at $\alpha = 0.72$, and both margins move in the same direction: the skill gain rises from 4.38% to 6.97%, and the experience gain from 1.67% to 2.52%. Increasing α reduces the estimated productivity gaps across establishments, lowering reallocation gains through the indirect channel, but it also increases the return to reallocation through the direct channel, which dominates.[15]

Elasticity of Substitution The parameter σ controls the substitutability of worker types in production. We re-estimate the model for $\sigma \in \{2, 3\}$ around the baseline $\sigma = 2.5$, holding $\alpha = 0.6$. Lower substitutability directly increases the value of workforce reallocation, but it also compresses the estimated productivity differences across worker types, since smaller gaps are needed to rationalize observed employment shares (see Appendix Figure C.2).

The indirect effect dominates. The joint gain rises with σ, from 7.05% at $\sigma = 2$ to 8.53% at the benchmark and 10.34% at $\sigma = 3$. Most of this increase comes from the skill margin, which rises from 4.85% to 6.58%, while the experience margin rises from 1.73% to 2.44%. Because the value of skill upgrading depends on productivity differences between high- and low-skill workers, this margin is particularly sensitive to changes in productivity dispersion.

Skill Depreciation The depreciation rate q determines the high-skill share at each experience, which the data do not reveal directly. We consider $q \in \{0.208, 0.312\}$, corresponding to annual depreciation of about 4.6 and 7.2 percent. The experience gain is 2.10% in both cases:

[15] This relates to the misallocation literature, where productivity losses increase as the revenue function becomes less concave (e.g., Hsieh and Klenow (2009)). That result captures only the direct channel. Because we re-estimate the model at each α, fundamentals also change, but not by enough to reverse the ranking.

with each experience bin's skill mix held fixed, this gain depends only on each bin's size and average wage, which the model matches exactly. The skill gain is sensitive to q. Faster depreciation implies fewer high-skill experienced workers, requiring a larger skill premium to match observed wage differences. Thus, the joint gain rises from 7.15% to 10.59%.

Taking Stock Three findings are robust across specifications. First, worker-establishment complementarity remains log-supermodular in skill. Second, the optimal workforce contains more workers with 21–40 years of experience and a larger share of high-skill workers than observed in Japan. Third, the consistently large output gains across parameterizations suggest that Japan's workforce remains far from its productivity-maximizing allocation.

The decomposition of these gains varies somewhat across parameter values, but their ranking is robust: the skill gain always exceeds the experience gain, while the latter increases with α and σ.[16]Because both margins contribute materially, neither education policies nor retirement and immigration reforms alone can achieve the full gains. Moreover, robust log-supermodularity and the scarcity of older high-skill workers imply that population aging disproportionately reduces productivity at the most productive establishments, which rely more heavily on young, high-skill workers.

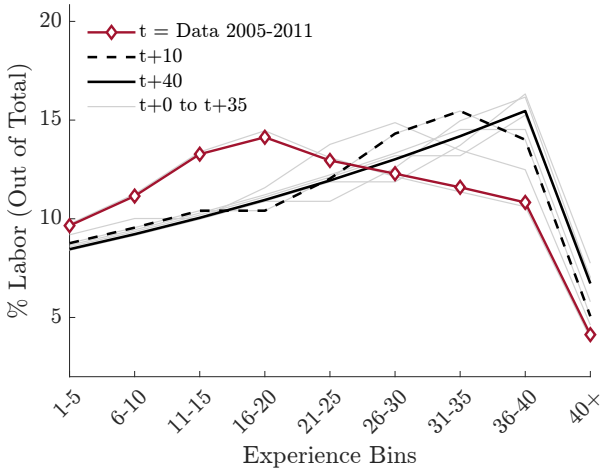
6 Application to Population Aging in Japan

The preceding analysis of the optimal composition highlights substantial potential gains from moving toward it. However, given that workforce composition evolves through slow-moving demographic processes, it is not clear that the optimal allocation is feasible. Specifically, the fact that workers accumulate experience mechanically with age implies that the optimal allocation is unlikely to be stationary. The inherent dynamics of population aging and skill accumulation impose feasibility constraints on compositional policies that the static reallocation exercise fails to capture.

To quantify the effects of population aging, we simulate the model over a 40-year horizon starting from the 2005-2011 baseline. We hold the total workforce fixed and allow only its age

[16]Because productivity differences between experience groups do not vary substantially across establishment types, gains from optimizing workforce composition along the experience dimension account for only a small share of the overall gains.

Figure 8: Evolution of Workforce Experience Distribution under Population Aging



Notes: Labor share by experience bin. The initial distribution is in red with diamonds (2005–2011), 10 years forward is the dashed black line, and 40 years forward is the solid black line. Grey lines show the remaining five-year steps. Based on the fertility rate of 1.3 and observed retirement patterns.

composition to evolve. The simulation incorporates Japan's low fertility rate, implying that each successive cohort is about 8% smaller than its predecessor, and assumes that 40% of workers remain employed after reaching 36-40 years of experience.[17] To maintain a constant labor force, entering cohorts are rescaled over time while preserving the demographic shift toward an older workforce. This setup captures Japan's rapid aging and allows us to isolate the effects of workforce composition on aggregate outcomes. This parameterization captures Japan's rapid demographic transition while preserving tractability in the dynamic analysis.[18]

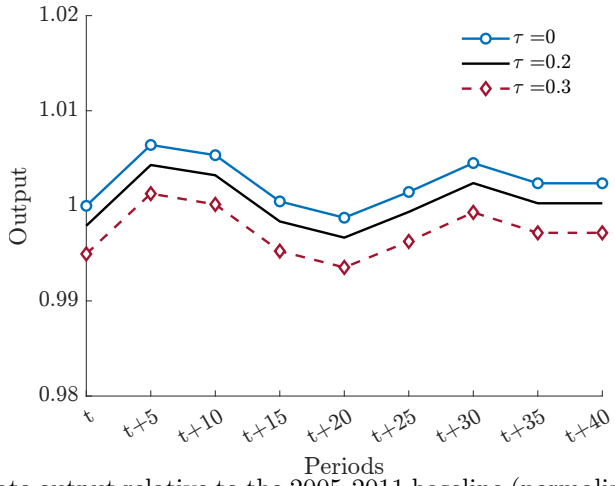
Figure 8 shows the projected evolution of Japan's workforce composition under our demographic assumptions. Starting from the 2005-2011 baseline, we simulate the experience distribution when each cohort is 8% smaller than the previous one. Over time, the distribution shifts toward more experienced workers, with the share of workers in the 36-40 experience group reaching 15.6% after 40 years, while the share of workers with less than 15 years of experience falls to about 10% or less.

We solve for equilibrium over the 40-year horizon by recomputing the steady state in each period using the contemporaneous experience and skill composition of the workforce. Thus,

[17] under the assumption that the average age of childbirth is 25, each successive cohort is 8% smaller than its predecessor as $L_j/L_{j+1} = (1.3/2)^{1/5} \approx 0.92$.

[18] Since the labor force must remain constant over time, we proportionately adjust the sizes of the current and previous entering cohorts (starting from the initial period) every five years, while maintaining the reduction in population of 8% between two consecutive entering cohorts.

Figure 9: Effect of Population Aging on Aggregate Output



Notes: Evolution of aggregate output relative to the 2005-2011 baseline (normalized to 1.0) under population aging. Lines show different reallocation cost scenarios: $\tau = 0$ (no reallocation cost), $\tau = 0.2$ (20% of wages), and $\tau = 0.3$ (30% of wages). The x-axis represents years forward from the initial period.

establishments behave as if the current labor force composition will persist, and labor and product markets clear in every period. This assumption is reasonable given the model's five-year period length and the slow pace of demographic change. Moreover, establishments' beliefs are irrelevant in the baseline static model, so they do not affect the counterfactual results based on that framework.

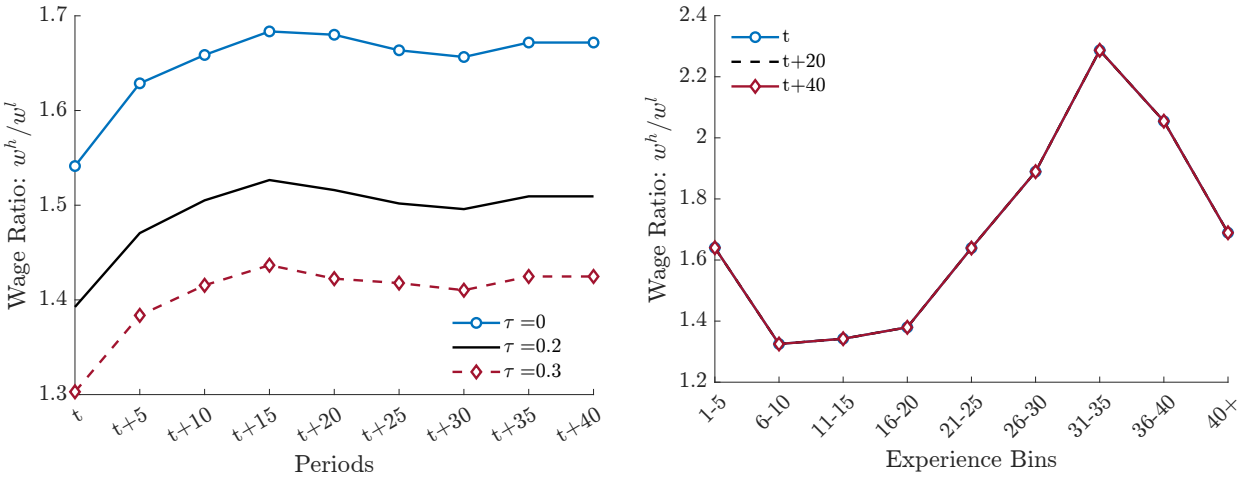
Throughout these counterfactuals, we use the benchmark parameters from Section 4, with no reallocation frictions. This frictionless benchmark allows us to isolate the effects of demographic composition on aggregate outcomes. We also explore the sensitivity of our results to reallocation frictions by conducting robustness checks with $\tau > 0$, holding all other parameters fixed at their baseline values. These exercises reveal how labor market frictions interact with demographic change to shape the economy's response to aging.

6.1 Population Aging, Productivity, and Concentration

We examine two dimensions of population aging's impact: its effects on aggregate productivity and its consequences for the distribution of workers across establishments. The analysis reveals distinct patterns at both the aggregate and distributional levels.

Figure 9 shows that population aging has only a modest effect on aggregate output over

Figure 10: Skill Premium Evolution under Population Aging



(a) Skill Premium Over Time

(b) Skill Premium by Experience ($\tau = 0$)

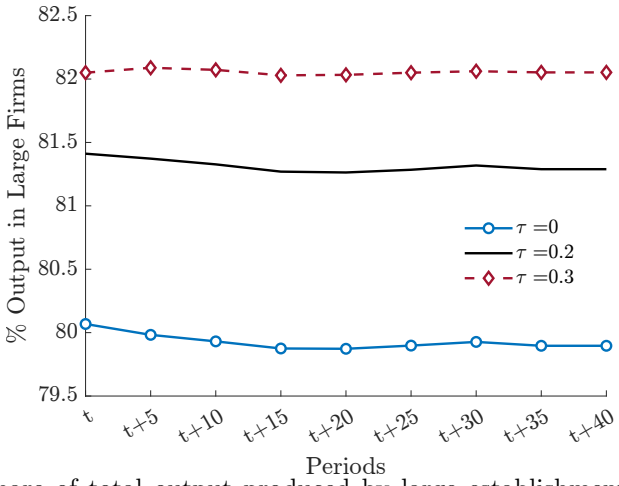
Notes: Panel (a) shows the evolution of the aggregate skill premium (high-skill/low-skill wage ratio) under different firing cost scenarios. Panel (b) displays skill premiums by experience group at three time points: initial (blue, solid), 20 years forward (black, dashed), and 40 years forward (maroon, with diamonds). The three are nearly indistinguishable. Experience bins are in 5-year intervals.

the 40-year horizon. Output initially rises by about 0.6% as large cohorts enter their most productive years, then declines slightly before ending about 0.2% above its initial level. Aging shifts the workforce through different experience groups but does not move it close to the optimal composition, as workers with 21–35 years of experience remain underrepresented. Reallocation costs lower output by 0.2% at $\tau = 0.2$ and 0.5% at $\tau = 0.3$, but have little effect on its overall trajectory, suggesting that labor market frictions play a limited role in mediating the aggregate productivity effects of aging.

Turning to the distributional effects, Figure 10 illustrates how population aging influences wage inequality across skill groups. Panel (a) traces the skill premium and finds that it rises from 54% to 67% over the projection period when $\tau = 0$. This 13-percentage-point increase corresponds to a 24% rise in relative wage inequality. This increase is driven by the growing share of more experienced workers, who command a higher skill premium, rather than by equilibrium wage adjustments.

Panel (b) decomposes this pattern by experience groups, highlighting the underlying compositional mechanism. We focus on the frictionless benchmark case with $\tau = 0$. The skill premium falls from 64% at entry to about 33% at 6–15 years of experience, rises to 129% at 31–35 years, and falls back to 69% for the most experienced workers. Within each bin

Figure 11: Large Establishment Output Share under Population Aging



Notes: Evolution of the share of total output produced by large establishments (116+ employees) under population aging. Lines show different firing cost scenarios.

the wage gap is constant over time (as indicated by the overlapping lines), so relative wages within a bin do not adjust.[19] The increase in the aggregate skill premium, therefore, reflects changes in the workforce composition rather than shifts in the returns to skill conditional on experience: aging moves workers away from 6 to 15 years of experience, where the premium is lowest, toward 26 to 40 years, where it is highest.

Reallocation costs primarily affect the level of the skill premium, which is about 15 percentage points lower at $\tau = 0.2$ and about 24 points lower at $\tau = 0.3$ than in the benchmark. As τ rises, it is more costly for large establishments to dismiss unskilled workers. Hence, more skilled workers stay in small establishments, and more unskilled workers remain employed in large ones. This mismatch between workers' skill and establishment productivity compresses the skill premium. Notably, these costs shift the level of wage dispersion, but not its trajectory.

Finally, we examine how population aging and reallocation frictions affect the establishment-size distribution. As the population ages, the output share of large establishments declines because productive establishments rely disproportionately on young and skilled workers. As the workforce becomes older and less skilled on average, their comparative advantage erodes, shifting output toward less productive establishments. This prediction is consistent with the

[19]This follows from the Cobb-Douglas technology. Equation (B.15) in the appendix shows that the skill wage ratio within each experience bin depends only on the ratio of productivities and the ratio of high- to low-skill supply. Since A_{ijk} and the experience-specific learning and depreciation probabilities do not change over time, the skill wage ratio within each bin is constant. Under a general CES technology, the changing supply of workers in other experience bins would also affect the skill wage ratio within a bin.

evidence in Section 6.2. Quantitatively, however, the effect is modest: the output share of large establishments falls by only 0.2 percentage points over 40 years.

Figure 11 shows that higher reallocation costs increase the concentration ratio of top establishments, from about 80% at $\tau = 0$ to 81.4% at $\tau = 0.2$ and 82% at $\tau = 0.3$. This effect arises through equilibrium adjustments in wages. As reallocation costs increase, the skill premium declines because productive establishments retain more low-skill workers and hire fewer high-skill workers. Since the most productive establishments employ disproportionately more high-skill workers, they benefit more than other establishments from the lower cost of skilled labor. Their profitability rises relative to that of less productive establishments, increasing their market share.

These results do not depend on the externally calibrated parameters. Appendix Table C.1 repeats the counterfactual for each re-estimation in Section 5.3. In every case, output stays close to its initial level, the aggregate skill premium rises, and the output share of large establishments falls. Levels differ across cases because each parameter maps the same wages to different productivities, but the response to aging does not.

6.2 Evidence: Aging and Market Concentration

In this subsection, we assess whether population aging reduces market concentration. We exploit variation in the growth rate of the working-age population across 47 Japanese prefectures and study its effect on the market shares of the top 5% and 10% of establishments at both the prefecture–year and prefecture–industry–year levels.[20] For the prefecture–year (prefecture–industry–year) analysis, we compute the sales and employment shares of the top 5% and 10% of establishments within each unit and year.[21]

We use a confidential establishment-level dataset and official statistics on prefecture-level demographic variables from Japan to conduct the analysis. The establishment-level dataset,

[20] We focus on within-country, across-region variation for two reasons. First, there is substantial heterogeneity in working-age population growth across Japanese prefectures. The average prefectural growth rate fell from 0.7% in 1985 to −1.3% in 2018, while the standard deviation declined only slightly, from 0.8% to 0.6%. Second, this design mitigates concerns that arise in cross-country studies, as regions in Japan are subject to broadly similar policies, institutions, and culture.

[21] The prefecture–industry–year specification helps control for compositional changes, but some cells contain too few establishments to compute market shares. We therefore report results at both levels and restrict the prefecture–industry–year sample to cells with at least 100 establishments. Results are similar when we use different thresholds such as 60 and 200.

the Census of Manufacturers (1985–2018), is mandatory for all manufacturing establishments and is provided by the Japanese Ministry of Economy, Trade, and Industry. It surveys all establishments with more than three employees annually, generating a panel spanning more than 30 years and covering approximately 300,000–350,000 establishments per year.

For prefecture-level demographic variables, we obtain data on the growth rate of the working-age population and the average age from the National Population Census (1985–2018), conducted by the Japanese Ministry of Internal Affairs and Communications.[22] We then merge the two datasets and estimate the following panel regression:

$$concentration_{p,y}^{top,est} = \beta_0 + \beta_1 working\ gr_{p,y} + \delta_p + \delta_y + \epsilon_{p,y}, \tag{12}$$

where $concentration_{p,y}^{top\ est}$ is the employment shares of top 5% (or 10%) establishments in prefecture p and year y at either the prefecture-year level or the prefecture-industry-year level. The variable of $working\ gr_{p,y}$ is the annual growth rate of the working-age population at the prefecture-year or prefecture-industry-year level. We include prefecture fixed effects δ_p (industry-prefecture fixed effects δ_{ip}) and year fixed effects into the prefecture-year (prefecture-industry-year) regression. As a result, we are comparing prefecture- or industry-prefecture pairs with different growth rates of the working-age population in the same period.

Results in Table 10 are consistent with our model's prediction: β_1 is positive and significant across all specifications, indicating that the employment share of top establishments declines in prefectures with slower growth in the working-age population. This result is robust to different top-establishment thresholds and concentration measures (see Table A.7). Using average population age as a proxy for the working-age population yields similar results: prefectures with more rapidly aging populations experience larger declines in the market share of the largest establishments over time (Table A.8). Moreover, the regression results are qualitatively similar and quantitatively stronger when we use the prefecture-year birth rate lagged by 20 years as an instrument for the annual growth rate of the working-age population.[23] The regressions capture all channels through which a slower-growing working-age population affects concentration, including entry and labor force size, whereas the model isolates workforce composition. The evidence therefore supports the sign of the prediction rather than its magnitude.

[22] The Population Census is conducted every five years (e.g., 1970, 1975, and 1980). We therefore linearly interpolate the average age and working-age population of each prefecture between census years.

[23] For details, see Appendix A.3.

Table 10: Working-age population growth and employment concentration

	Prefecture-Year		Industry-Prefecture-Year	
	(1) Top 5%	(2) Top 10%	(3) Top 5%	(4) Top 10%
Working-age pop. growth	1.458***	1.211***	0.464*	0.391*
	(0.427)	(0.372)	(0.232)	(0.225)
Prefecture FE	Yes	Yes	No	No
Industry × Prefecture FE	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	1,598	1,598	19,913	19,913
Adjusted R^2	0.894	0.902	0.730	0.733

Notes: The dependent variable is the employment share of the largest establishments (top 5% or top 10% by employment) in each prefecture-year (columns 1–2) or industry-prefecture-year (columns 3–4). Working-age population growth is measured at the prefecture-year level. Standard errors are two-way clustered by prefecture and year. Columns (3)–(4) include only industry-prefecture-year cells with at least 100 establishments. Data from Census of Manufacturing Plants, 1985–2018. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

7 Conclusion

As populations age, understanding how demographic change affects workforce composition and economic performance has become increasingly important. We develop a framework to characterize the optimal allocation of workers across heterogeneous establishments and quantify the effects of population aging using Japanese matched employer-worker data.

We show that larger, more productive establishments employ younger workers with longer tenure, reflecting complementarities between productivity and worker skill. Reallocating workers toward the optimal composition would raise Japan's GDP by 8.5%, primarily through skill composition. Incorporating dynamics and labor market frictions, we find that population aging increases wage inequality across skill groups while reducing market concentration as high-productivity establishments lose their comparative advantage. Empirical evidence is consistent with these predictions.

Our results imply that education, retraining, immigration, and retirement policies can generate sizable gains in aging economies. We also show that employment protection policies have heterogeneous effects across establishments and worker groups. An important direction for future research is to endogenize technological adaptation, such as automation and age-complementary innovations, which may amplify or offset the effects of population aging.

References

Alvarez, Fernando and Marcelo Veracierto, "Severance payments in an economy with frictions," *Journal of monetary Economics*, 2001, *47* (3), 477–498.

Bagga, Sadhika, "Internal and External Labor Markets and Declining Dynamism," 2022.

Berger, Mark C, "The Effect of Cohort Size on Earnings Growth: A Reexamination of the Evidence," *Journal of Political Economy*, June 1985, *93* (3), 561–573.

Bianchi, Nicola and Matteo Paradisi, "Countries for Old Men: An Analysis of the Age Pay Gap," Technical Report, National Bureau of Economic Research 2024.

Blundell, Richard, Monica Costa-Dias, David Goll, and Costas Meghir, "Wages, experience, and training of women over the life cycle," *Journal of Labor Economics*, 2021, *39* (S1), S275–S315.

Borjas, George J, "The labor demand curve is downward sloping: Reexamining the impact of immigration on the labor market," *The quarterly journal of economics*, 2003, *118* (4), 1335–1374.

Braun, R. Anton and Douglas H. Joines, "The implications of a graying Japan for government policy," *Journal of Economic Dynamics and Control*, None 2015, *57* (C), 1–23.

_ , Daisuke Ikeda, and Douglas H. Joines, "The Saving Rate In Japan: Why It Has Fallen And Why It Will Remain Low," *International Economic Review*, February 2009, *50* (1), 291–321.

Card, David and Thomas Lemieux, "Can falling supply explain the rising return to college for younger men? A cohort-based analysis," *The quarterly journal of economics*, 2001, *116* (2), 705–746.

Cardona, Carolina, "How a Vast Demographic Shift Will Reshape the World," *The New York Times*, 2023.

Engbom, Niklas et al., "Firm and worker dynamics in an aging labor market," Technical Report 2019.

Frey, William H., "Immigration drives the nation's healthy post-pandemic population growth, new census data show," Technical Report, Brookings Insitution 2025.

Geruso, Michael and Dean Spears, "The Likelihood of Persistently Low Global Fertility," *Journal of Economic Perspectives*, Winter 2026, *40* (1), 3–26.

Gregory, Victoria, "Firms as learning environments: Implications for earnings dynamics and job search," *FRB St. Louis Working Paper*, 2020, (2020-036).

Hayashi, Fumio, "The Root Cause of Japan's 30-year Stagnation: Implications for the US, Europe, and East Asia," *Europe, and East Asia (August 28, 2023)*, 2023.

Hopenhayn, Hugo and Richard Rogerson, "Job turnover and policy evaluation: A general equilibrium analysis," *Journal of political Economy*, 1993, *101* (5), 915–938.

_ , Julian Neira, and Rish Singhania, "From population growth to firm demographics: Implications for concentration, entrepreneurship and the labor share," *Econometrica*, 2022, *90* (4), 1879–1914.

Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly journal of economics*, 2009, *124* (4), 1403–1448.

IMF, "Japan: Demographic Shift Opens Door to Reforms," Technical Report, IMF Country Focus 2020.

Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin, "Demographic Origins of the start-up Deficit," *American Economic Review*, 2024, *114* (7), 1986–2023.

Katz, Lawrence F and Kevin M Murphy, "Changes in relative wages, 1963–1987: supply and demand factors," *The quarterly journal of economics*, 1992, *107* (1), 35–78.

Kim, Seula, "Workers' Job Prospects and Young Firm Dynamics," Technical Report, Working Paper 2023.

Kitao, Sagiri, "Fiscal cost of demographic transition in Japan," *Journal of Economic Dynamics and Control*, None 2015, *54* (C), 37–58.

Ma, Xiao, Alejandro Nakab, and Daniela Vidart, "Human capital investment and development: The role of on-the-job training," *Journal of Political Economy Macroeconomics*, 2024, *2* (1), 107–148.

Maestas, Nicole, Kathleen J Mullen, and David Powell, "The effect of population aging on economic growth, the labor force and productivity," Technical Report, National Bureau of Economic Research 2016.

Ottaviano, Gianmarco IP and Giovanni Peri, "Rethinking the effect of immigration on wages," *Journal of the European economic association*, 2012, *10* (1), 152–197.

Ouimet, Paige and Rebecca Zarutskie, "Who works for startups? The relation between firm age, employee age, and growth," *Journal of financial Economics*, 2014, *112* (3), 386–407.

Papageorgiou, Theodore, "Large firms and within firm occupational reallocation," *Journal of Economic Theory*, None 2018, *174* (C), 184–223.

Peters, Michael and Conor Walsh, "Population growth and firm dynamics," Technical Report, National Bureau of Economic Research 2021.

Poschke, Markus, "Employment protection, firm selection, and growth," *Journal of Monetary Economics*, 2009, *56* (8), 1074–1085.

Qi, Liyan, "China Sees a Fresh Decline in Population, Despite a Rise in Births," *The Wall Street Journal*, 2025.

Toossi, Mitra and Torpey Elka, "Older workers: Labor force trends and career options," Technical Report, Bureau of Labor Statistics 2017.

Veracierto, Marcelo, "Employment flows, capital mobility, and policy analysis," *International Economic Review*, 2001, *42* (3), 571–596.

Weil, David N., "How Much Would Continued Low Fertility Affect the US Standard of Living?," *Journal of Economic Perspectives*, Winter 2026, *40* (1), 27–46.

Welch, Finis, "Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust," *Journal of Political Economy*, 1979, *87* (5), S65–97.

Online Appendix for " Population Aging, Firms, and Worker Composition"

Santiago Caicedo, Cheng Chen, Aspen Gorry and Takahiro Hattori

This appendix provides supplementary empirical evidence, theoretical foundations, and computational details for the main analysis. Section A presents additional empirical patterns from using the Japanese establishment-employee matched data. Section B contains formal proofs of theoretical propositions, efficiency analysis, and identification arguments.

A Additional Empirical Evidence

This section presents additional empirical patterns that complement the main text, including robustness checks across different time periods, sample restrictions, and alternative establishment-size classifications, along with supporting summary statistics and regressions.

A.1 Worker Composition across Establishments

Table A.1: Average Worker Characteristics by Two Size Bins (1998-2004): full sample

Size Bin	Age	Experience	Tenure	Wage	Education (years)
1	40.45	21.69	11.62	36713.10	12.76
2	39.12	19.80	14.51	44632.34	13.32

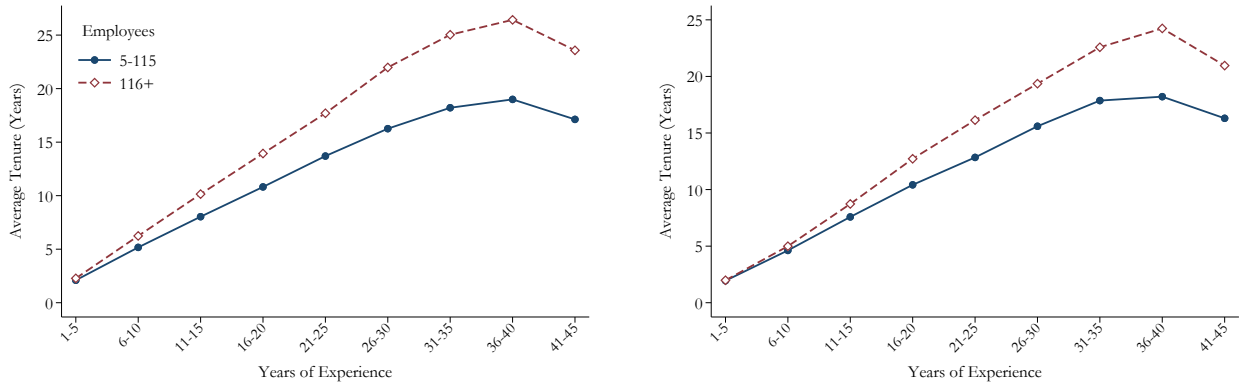
Note: Time period: 1998-2004. Unit of average wage: one hundred yen (or one USD). We divide establishments to two sized bins. Size bin 1: employment $\leq$ 115; Size bin 2: employment $\geq$ 116. Each worker-year observation is treated as a unit of observation.

Table A.2: Average Worker Characteristics by Two Size Bins (2005-2011): permanent workers only

Size Bin	Age	Experience	Tenure	Wage	Education (years)
1	40.67	21.35	12.00	37436.63	13.32
2	39.99	20.27	14.77	45464.95	13.72

Note: Time period: 2005-2011. Unit of average wage: one hundred yen (or one USD). We divide establishments to two sized bins. Size bin 1: employment ≤ 115; Size bin 2: employment ≥ 116. Each worker-year observation is treated as a unit of observation.

Figure A.1: Experience-Tenure Profiles by Period and Worker Type



(a) 1998–2004 (b) 2005–2011: Permanent Workers

Notes: Experience is defined as age minus (length of education + 6). Experience bins are defined as: bin 1 (1–5 years), bin 2 (6–10 years), ..., bin 8 (36–40 years), bin 9 (41–45 years). The experience bins are capped at 45 years. Small establishments have fewer than 116 employees; large establishments have 116 or more employees.

Table A.3: Average Worker Characteristics by Five Size Bins (2005-2011): full sample

Size Bin	Age	Experience	Tenure	Wage	Education (years)
1	42.58	23.81	10.57	31221.04	12.77
2	40.77	21.56	10.68	34566.54	13.21
3	40.92	21.68	11.10	35674.05	13.24
4	41.10	21.79	11.81	37151.72	13.31
5	40.41	20.83	13.46	42481.81	13.58

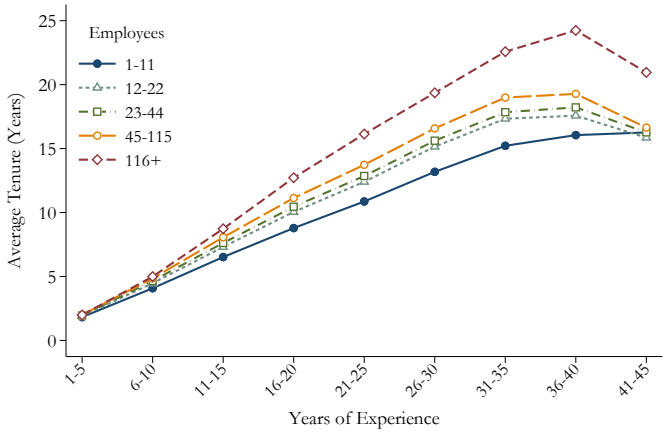
Note: Time period: 2005-2011. Unit of average wage: one hundred yen (or one USD). We divide establishments to five size equal-sized bins. Size bin 1: employment ≤ 11; Size bin 2: employment $\in [12, 22]$; Size bin 3: employment $\in [23, 44]$; Size bin 4: employment $\in [45, 115]$; Size bin 5: employment ≥ 116. Each worker-year observation is treated as a unit of observation.

Table A.4: Average Worker Characteristics by Five Size Bins (1998-2004): full sample

Size Bin	Age	Experience	Tenure	Wage	Education (years)
1	42.15	23.81	10.55	32056.02	12.34
2	40.61	21.92	10.95	35683.63	12.69
3	40.32	21.55	11.51	36838.49	12.77
4	40.06	21.17	12.28	38261.13	12.89
5	39.12	19.80	14.51	44632.34	13.32

Note: Time period: 1998-2004. Unit of average wage: one hundred yen (or one USD). We divide establishments to five size equal-sized bins. Size bin 1: employment ≤ 11; Size bin 2: employment $\in [12, 22]$; Size bin 3: employment $\in [23, 44]$; Size bin 4: employment $\in [45, 115]$; Size bin 5: employment ≥ 116. Each worker-year observation is treated as a unit of observation.

Figure A.2: Experience-tenure profile (2005-2011): Five Size Bins



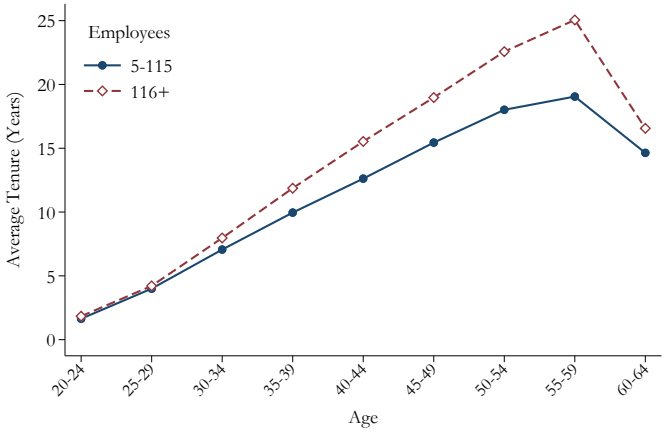
Note: Experience = age - (length of education + 6). Experience bin 1: 1-5 years; ... Experience bin 8: 36-40 years; Experience bin 9: 41-45 years. We cut the experience bin at 45 years.

Table A.5: Cross-sectional regression of experience, age and tenure: 2004-2011

Dep.Var:	(1) ave. experience	(2) ave. age	(3) ave. tenure
log(employment)	-0.355***	-0.310***	0.531***
	(0.011)	(0.011)	(0.011)
establishment age	0.056***	0.056***	0.082***
	(0.001)	(0.001)	(0.001)
constant	49.275***	53.854***	8.152***
	(0.167)	(0.151)	(0.119)
City FE	Yes	Yes	Yes
Industry-year FE	Yes	Yes	Yes
Prefecture-year FE	Yes	Yes	Yes
N	365721	365721	365721
R^2	0.400	0.398	0.454

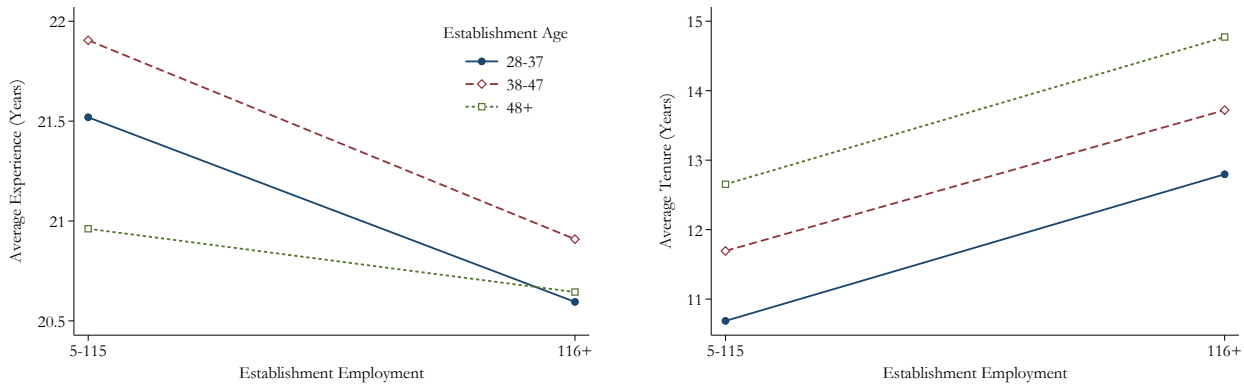
* 0.10 ** 0.05 *** 0.01. Standard errors are clustered at the establishment level.

Figure A.3: Age-tenure profile (2005-2011)



Note: Age bin 1: 20-24 years old; Age bin 2: 25-29 years old; ... Age bin 9: 60-64 years old. Small establishments have fewer than 116 employees; large establishments have 116 or more employees.

Figure A.4: Average Experience and Tenure by Establishment Age and Size (2005–2011)



(a) Average Experience (b) Average Tenure

Notes: Average experience (panel a) and tenure (panel b) by establishment age and size. Sample restricted to workers before retirement age. Small establishments (the first employment bin) have fewer than 116 employees; large establishments (the second employment bin) have 116 or more employees.

Table A.6: Worker Characteristics and Firm Size: 2005-2011

Size Bin	Experience	Tenure	Wage
1	22.94	10.04	27686.5
2	21.61	11.28	31746.4
3	20.22	12.81	36279.1
4	19.91	13.98	42235.4

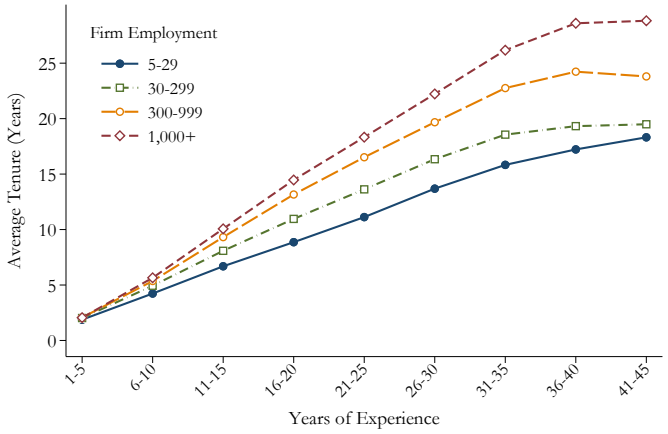
Time period: 2005-2011. *exp*: experience (years of working). Unit of *wage*: one hundred yen (or one USD). Establishments founded after 1985 are excluded, and workers older than 60 are excluded. We divide firms to four size bins. Size bin 1: emp $\in$ [5,29]; Size bin 2: emp $\in [30, 299]$; Size bin 3: emp $\in [300, 999]$; Size bin 4: emp $\geq$ 1000. Each worker-year observation is treated as a unit of observation.

Figure A.5: Experience-Tenure Profile by Establishment Age (2005-2011): Workers Before Retirement



Note: Experience = age - (length of education + 6). Experience bin 1: 1-5 years; ... Experience bin 9: 41-45 years. We cut the experience bin at 45 years.

Figure A.6: Experience-Tenure Profile at the Firm Level: 2005-2011



Note: Establishments founded after 1985 are excluded, and workers older than 60 are excluded. Experience = age - (length of education + 6). Experience bin 1: 1-5 years; ... Experience bin 8: 36-40 years; Experience bin 9: 41-45 years.

A.2 Population Growth, Aging, and Concentration

This subsection reports additional results for the concentration regressions in Section 6.1: alternative concentration measures, average population age in place of population growth, and instrumental variable estimates.

Table A.7: Working-age population growth and sales concentration

	Prefecture-Year		Industry-Prefecture-Year	
	(1) Top 5%	(2) Top 10%	(3) Top 5%	(4) Top 10%
Working-age pop. growth	1.722***	1.229***	0.309	0.250
	(0.522)	(0.381)	(0.369)	(0.286)
Prefecture FE	Yes	Yes	No	No
Industry × Prefecture FE	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	1,598	1,598	19,913	19,913
Adjusted R^2	0.871	0.876	0.691	0.702

Notes: The dependent variable is the sales share of the largest establishments (top 5% or top 10% by sales) in each prefecture-year (columns 1–2) or industry-prefecture-year (columns 3–4). Working-age population growth is measured at the prefecture-year level. Standard errors are two-way clustered by prefecture and year. Columns (3)–(4) include only industry-prefecture-year cells with at least 100 establishments. Data from Census of Manufacturing Plants, 1985–2018. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.8: Average population age and employment concentration

	Sales Share		Employment Share	
	(1) Top 5%	(2) Top 10%	(3) Top 5%	(4) Top 10%
Average age	−0.00723***	−0.00607***	−0.00484***	−0.00511***
	(0.00217)	(0.00201)	(0.00140)	(0.00131)
Industry × Prefecture FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	19,913	19,913	19,913	19,913
Adjusted R^2	0.692	0.702	0.733	0.733

Notes: The dependent variable is the sales share (columns 1–2) or employment share (columns 3–4) of the largest establishments (top 5% or top 10%) in each industry-prefecture-year cell. Average age is measured at the prefecture-year level. Standard errors are clustered at the prefecture-year level. All columns include only industry-prefecture-year cells with at least 100 establishments. Data from Census of Manufacturing Plants, 1985–2018. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

A.3 IV Regressions: Establishment Age and Worker Migration

Establishment (and firm) age is positively correlated with market share, as surviving establishments tend to grow. Thus, average age is a potential confounder in our regressions. Population aging or slower working-age population growth leads to older establishments, which would suggest higher market concentration, the opposite of our empirical finding. Thus, this confounding factor works *against* our results. Controlling for it would likely make the positive effect of population growth (and the negative effect of aging) on market concentration even stronger.

Worker mobility may induce endogeneity, as shocks that expand large establishments could also attract young workers, creating reverse causality. To address this, we follow Karahan et al. (2024) and use the 20-year-lagged birth rate at the prefecture-year level as an instrument for the annual growth rate of the working-age population. Since many young adults become part of the working-age population in their prefecture of birth, the prefecture's lagged birth rate will influence the future growth of its working-age population. The exclusion restriction is that, conditional on prefecture and time fixed effects, the determinants of a higher birth rate have no effect on future establishments except indirectly through their impact on the growth of the working-age population.[24]

[24]This assumption would be violated, for instance, if people in a given prefecture had a higher fertility rate

Table A.9: Population growth and employment concentration: IV estimates using lagged birth rates

	First Stage	Prefecture-Year		Industry-Prefecture-Year	
	(1)	(2)	(3)	(4)	(5)
	Working-Age Pop. Growth	Employment Share Top 5%	Top 10%	Employment Share Top 5%	Top 10%
Working-age pop. growth		5.007***	4.325***	1.607*	1.645**
		(1.754)	(1.492)	(0.803)	(0.721)
Birth rate$_{t-20}$	0.124***				
	(0.041)				
F-statistic (first stage)	9.15				
Prefecture FE	Yes	Yes	Yes	No	No
Industry × Prefecture FE	No	No	No	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	1,588	1,588	1,588	19,873	19,873
R^2	0.908	0.866	0.876	0.744	0.747
Adjusted R^2	0.903	0.859	0.869	0.729	0.733

Notes: Column (1) presents the first-stage regression of working-age population growth on the 20-year lagged birth rate at the prefecture level. Columns (2)–(5) present second-stage IV estimates where the dependent variable is the employment share of the largest establishments (top 5% or top 10% by employment) in each prefecture-year (columns 2–3) or industry-prefecture-year (columns 4–5). Working-age population growth is instrumented using the 20-year lagged prefecture birth rate. Standard errors in parentheses are clustered at the prefecture level. Columns (4)–(5) include only industry-prefecture-year cells with at least 100 establishments with non-missing sales and employment. Data from Census of Manufacturing Plants, 1985–2018. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table A.9 confirms the results reported in Table 10.[25] The effect is quantitatively sizable: a 0.86% decline in working-age population growth (one standard deviation) reduces the employment share of the top 10% of establishments by 3.72% at the prefecture-year level and 1.41% at the industry-prefecture-year level (standard deviations: 5.21% and 13.0%, respectively). In summary, both OLS and IV results are consistent with our model's key prediction that higher average worker age reduces the employment share of top establishments.

20 years ago in anticipation of stronger labor market conditions (relative to other prefectures). However, as local economies in Japan (and the national economy) experienced significant fluctuations between 1975 and 2018, and given the long 20-year time lag, it is unlikely that individual households had such perfect foresight when making their fertility decisions.

[25]The correlation between the working-age population growth rate and the 20-year lagged birth rate is 0.87 in our sample.

B Theory Appendix

This section provides formal proofs of propositions from the static model, efficiency characterizations under alternative production technologies, and identification arguments for the structural parameters.

B.1 Proofs of Propositions from the Static Model

This subsection provides proofs for all propositions stated in Section 3.1 of the main text. We establish conditions under which the planner's allocation dominates the competitive equilibrium and characterize the optimal reallocation of labor across worker groups.

Setup: Partitions and the Planner's Problem

The text presents the two endpoints of the planner's problem: competitive equilibrium, in which each group's supply is fixed, and full flexibility, in which the planner reallocates labor freely across all groups. The proofs below cover the full range of intermediate cases, indexed by the number of groups the planner may reshuffle jointly. The statements in the text are the leading special cases of the general results established here.

Let $\mathcal{P}$ be a partition of $\mathcal{G}$. Each element $P_s \in \mathcal{P}$ collects the groups whose labor the planner may reallocate jointly, with $L_s = \sum_{g \in P_s} L_g$ and $\sum_s L_s = L$. Taking $\mathcal{P}$ as given, the planner solves

$$\max_{l_g(z) \geq 0} \int F(\{l_g\}; z)\, d\Phi(z) \quad \text{s.t.} \quad \int \left(\sum_{g \in P_s} l_g(z) \right) d\Phi(z) = L_s, \quad \forall s,$$

with associated multipliers μ_s. The first-order conditions equalize the marginal product of labor within each element of the partition,

$$[l_g(z)] : \ \frac{\partial F(\{l_g\}; z)}{\partial l_g} \leq \mu_s \qquad \forall z,\ g \in P_s,\ \forall s. \tag{B.13}$$

Two special cases bracket the problem. The coarsest partition $\mathcal{P}^{max} = \{\mathcal{G}\}$ places all groups in a single element, so (B.13) collapses to a single shadow value and recovers the fully flexible planner of the text. The finest partition $\mathcal{P}^{min} = \{\{g\} : g \in \mathcal{G}\}$ imposes one constraint per

group, so $\mu_g = w_g$ and (B.13) reproduces the competitive conditions (1) together with market clearing (2).

We rank partitions by how much flexibility they grant.

Definition 1. *Let $\mathcal{P}$ and $\mathcal{P}'$ be partitions of $\mathcal{G}$. We say that $\mathcal{P}$ is a refinement of $\mathcal{P}'$ if every element of $\mathcal{P}$ is a subset of some element of $\mathcal{P}'$. In this case $\mathcal{P}$ is finer than $\mathcal{P}'$, written $\mathcal{P} \leq \mathcal{P}'$.*

The finer-than relation is a partial order on the partitions of $\mathcal{G}$. Let $Y^p(\mathcal{P})$ denote output under partition $\mathcal{P}$. A coarser partition gives the planner more room to reallocate, and so cannot lower output.

Proposition 3. *If $\mathcal{P}$ is a refinement of $\mathcal{P}'$, then $Y^p(\mathcal{P}') \geq Y^p(\mathcal{P})$.*

Proposition 1 in the text is the statement that $Y^p(\mathcal{P}) \geq Y^c = Y^p(\mathcal{P}^{min})$ for any $\mathcal{P}$ with fewer elements than groups, and Corollary 1 is the special case $\mathcal{P} = \mathcal{P}^{max}$, which by Proposition 3 attains the largest gain. Proposition 2 concerns any two groups placed in a common element P_s. The proofs of all four results follow.

Preliminary Lemmas

Lemma 1. *If $\mathcal{P}$ is a refinement of $\mathcal{P}'$, then any allocation feasible under $\mathcal{P}$ is also feasible under $\mathcal{P}'$.*

Proof. Since $\mathcal{P}$ refines $\mathcal{P}'$, every element $P_s \in \mathcal{P}$ is contained in some element $P'_t \in \mathcal{P}'$. Consider an allocation $\{l_g(z)\}$ feasible under $\mathcal{P}$:

$$\int \sum_{g \in P_s} l_g(z) d\Phi(z) = L_s \quad \forall P_s \in \mathcal{P}.$$

For any $P'_t \in \mathcal{P}'$, let $\mathcal{S}_t = \{P_s \in \mathcal{P} : P_s \subseteq P'_t\}$. Since $\mathcal{P}$ partitions each P'_t:

$$\int \sum_{g \in P'_t} l_g(z) d\Phi(z) = \sum_{P_s \in \mathcal{S}_t} \int \sum_{g \in P_s} l_g(z) d\Phi(z) = \sum_{P_s \in \mathcal{S}_t} L_s = L'_t.$$

Hence, $\{l_g(z)\}$ is feasible under $\mathcal{P}'$. □

Proofs of Main Results

Proof of Proposition 1. We show that $Y^p(\mathcal{P}) \geq Y^c$ whenever $|\mathcal{P}| = S < G$.

The competitive allocation $\{l_g^c(z)\}$ satisfies $\int l_g^c(z)\, d\Phi(z) = L_g$ for every g, so it is feasible under $\mathcal{P}^{\min}$. Since $\mathcal{P}^{\min}$ refines every partition, Lemma 1 makes it feasible under $\mathcal{P}$ as well. The planner maximizes over a set containing this allocation, so

$$Y^p(\mathcal{P}) = \max_{\{l_g(z)\} \in \mathcal{F}(\mathcal{P})} \int F(\{l_g\}; z) d\Phi(z) \geq \int F(\{l_g^c\}; z) d\Phi(z) = Y^c,$$

where $\mathcal{F}(\mathcal{P})$ denotes the feasible set under partition $\mathcal{P}$.

Now we show the strict inequality. Suppose there exist $g, g' \in P_s$ for some s such that $w_g \neq w_{g'}$. In a competitive equilibrium, the first-order conditions imply:

$$\frac{\partial F(\{l_g^c\}; z)}{\partial l_g} = w_g \neq w_{g'} = \frac{\partial F(\{l_g^c\}; z)}{\partial l_{g'}} \quad \text{for all } z.$$

However, the planner's optimality conditions require:

$$\frac{\partial F(\{l_g^p\}; z)}{\partial l_g} = \mu_s = \frac{\partial F(\{l_g^p\}; z)}{\partial l_{g'}} \quad \text{for all } g, g' \in P_s.$$

Since marginal products are not equalized at the competitive allocation, while they must be equalized at the planner's optimum, we have $\{l_g^p(z)\} \neq \{l_g^c(z)\}$. By strict concavity of F, the optimum is unique, implying $Y^p(\mathcal{P}) > Y^c$. □

Proof of Proposition 2. Let $g, g' \in P_s$ with $w_g > w_{g'}$ in competitive equilibrium. We show that $\frac{L_g^p}{L_{g'}^p} > \frac{L_g^c}{L_{g'}^c} = \frac{L_g}{L_{g'}}$.

In competitive equilibrium, the first-order conditions for establishment optimization yield:

$$\frac{\partial F(\{l_g^c\}; z)}{\partial l_g} = w_g > w_{g'} = \frac{\partial F(\{l_g^c\}; z)}{\partial l_{g'}} \quad \text{for all } z.$$

Now we show that the competitive allocation is suboptimal for the planner. Consider a perturbation from the competitive allocation: for each z, increase $l_g(z)$ by ϵ and decrease $l_{g'}(z)$ by ϵ, where $\epsilon > 0$ is small. This perturbation satisfies the planner's constraint since

aggregate labor in P_s is unchanged:

$$\int \left[(l_g^c(z) + \epsilon) + (l_{g'}^c(z) - \epsilon) + \sum_{h \in P_s \setminus \{g,g'\}} l_h^c(z) \right] d\Phi(z) = L_s.$$

The first-order change in output from this perturbation is:

$$\int \left[\frac{\partial F}{\partial l_g}(\{l_g^c\}; z) - \frac{\partial F}{\partial l_{g'}}(\{l_g^c\}; z) \right] \epsilon \, d\Phi(z) = (w_g - w_{g'})\epsilon > 0.$$

Thus, output strictly increases when labor is shifted from group g' to group g.

Finally, we establish the inequality. For $\ell \in \mathbb{R}$, let $\varphi(\ell)$ denote the maximum of $\int F(\{l_g\}; z) d\Phi(z)$ subject to $\int l_g(z) d\Phi(z) = L_g + \ell$, $\int l_{g'}(z) d\Phi(z) = L_{g'} - \ell$, and $\int l_h(z) d\Phi(z) = L_h$ for every other group h. The constraint set is convex and F is concave, so φ is concave. At $\ell = 0$ the constraints are the market-clearing conditions (2), so the competitive allocation solves the problem with multipliers equal to the competitive wages, and the envelope theorem gives $\varphi'(0) = w_g - w_{g'} > 0$. Output therefore rises when labor is shifted from g' to g. A concave function with a strictly positive derivative at zero attains its maximum at some $\ell^* > 0$, so if the planner reallocates between g and g' alone, $L_g^p = L_g + \ell^* > L_g$ and $L_{g'}^p = L_{g'} - \ell^* < L_{g'}$, and the ratio $L_g^p / L_{g'}^p$ exceeds $L_g / L_{g'}$. □

Corollary 2. *Under the CES technology, the comparison extends to any set of groups the planner may reallocate among. If $g, g' \in P_s$ and $w_g > w_{g'}$, then $l_g^p(z)/l_{g'}^p(z) > l_g^c(z)/l_{g'}^c(z)$ at every establishment z. When $\sigma = 1/(1-\alpha)$ the same inequality holds in aggregate, $L_g^p/L_{g'}^p > L_g/L_{g'}$.*

Proof of Corollary 2. Under CES the marginal rate of technical substitution between g and g' at establishment z is $MRTS_{g,g'}(z) = (A_g(z)/A_{g'}(z))^{1/\sigma} (l_{g'}(z)/l_g(z))^{1/\sigma}$, which is strictly decreasing in $l_g(z)/l_{g'}(z)$ and does not depend on the inputs of any other group. Profit maximization sets it equal to $w_g/w_{g'}$, so $l_g^c(z)/l_{g'}^c(z) = (A_g(z)/A_{g'}(z)) (w_{g'}/w_g)^\sigma$, while the planner's first-order conditions for $g, g' \in P_s$ set it equal to one, so $l_g^p(z)/l_{g'}^p(z) = A_g(z)/A_{g'}(z)$. Since $w_g > w_{g'}$, the factor $(w_{g'}/w_g)^\sigma$ is smaller than one, and the ratio rises at every z. Because the pairwise condition holds establishment by establishment, the direction of reallocation is uniform across establishments. When $\sigma = 1/(1-\alpha)$ the exponent $\sigma - (\sigma - 1)/\alpha$ in (3) and (5) is zero, so $l_g^c(z) = A_g(z) L_g / \bar{A}_g$ and $l_g^p(z) = A_g(z) L / \bar{A}$, giving $L_g^p / L_{g'}^p = \bar{A}_g / \bar{A}_{g'}$. By (4), $w_g > w_{g'}$ is equivalent to $\bar{A}_g / L_g > \bar{A}_{g'} / L_{g'}$, which is the claim. □

Proof of Proposition 3. Let $\mathcal{P}$ be a refinement of $\mathcal{P}'$. Denote by $\{l_g^p(z;\mathcal{P})\}$ the planner's optimal allocation under partition $\mathcal{P}$.

By Lemma 1, $\{l_g^p(z;\mathcal{P})\}$ is feasible under $\mathcal{P}'$. Therefore:

$$Y^p(\mathcal{P}') = \max_{\{l_g(z)\}\in\mathcal{F}(\mathcal{P}')} \int F(\{l_g\};z)d\Phi(z) \geq \int F(\{l_g^p(\cdot;\mathcal{P})\};z)d\Phi(z) = Y^p(\mathcal{P}).$$

For strict inequality, suppose $\mathcal{P} \neq \mathcal{P}'$. Then there exists some $P_t' \in \mathcal{P}'$ that properly contains multiple elements of $\mathcal{P}$. If these elements have workers with different competitive wages, Proposition 1 implies $Y^p(\mathcal{P}') > Y^p(\mathcal{P})$. □

Proof of Corollary 1. The maximal partition $\mathcal{P}^{max} = \{\mathcal{G}\}$ is coarser than any proper partition $\mathcal{P} \neq \mathcal{P}^{max}$, meaning $\mathcal{P}$ refines $\mathcal{P}^{max}$. By Proposition 3, $Y^p(\mathcal{P}^{max}) \geq Y^p(\mathcal{P})$.

Strict inequality holds if there exist groups $g, g' \in \mathcal{G}$ with $w_g \neq w_{g'}$ that belong to different elements of $\mathcal{P}$. Since $\mathcal{P}^{max}$ places all groups in a single element while $\mathcal{P}$ separates some groups, the planner gains additional flexibility under $\mathcal{P}^{max}$.

Under $\mathcal{P}^{min} = \{\{g\} : g \in \mathcal{G}\}$, each partition element contains exactly one group. The planner's constraints become:

$$\int l_g(z)d\Phi(z) = L_g \quad \forall g \in \mathcal{G},$$

which are identical to the competitive market-clearing conditions (2). Since both problems maximize the same objective subject to the same constraints, $Y^p(\mathcal{P}^{min}) = Y^c$. □

B.2 Efficiency Analysis Under CES Production

This subsection analyzes how the efficiency gains from optimal labor reallocation depend on the elasticity of substitution between worker types and the dispersion of comparative advantages across establishments.

Proposition 4. *Under the CES technology, $Y^p \geq Y^c$, with strict inequality if and only if B_g/L_g varies across groups, where $B_g := \int y^c(z)^{\sigma-(\sigma-1)/\alpha} A_g(z)\, d\Phi(z)$ and $y^c(z)$ is competitive output. Equivalently, the gain is strict if and only if two groups earn different wages. The gain vanishes as $\sigma \to \infty$.*

Proof. The competitive allocation is feasible for the planner, so $Y^p \geq Y^c$. It is optimal for the planner only if the marginal product of labor is equalized across all pairs (g, z), whereas the competitive first-order conditions equalize it at w_g within each group. The two coincide if and only if w_g is constant across g, and by (4) $w_g = \alpha(B_g/L_g)^{1/\sigma}$, so this holds if and only if B_g/L_g is constant. Strict concavity of F makes the planner's optimum unique, so the inequality is strict otherwise. As $\sigma \to \infty$ the technology becomes $y(z) = (\sum_g l_g(z))^\alpha$, in which only total employment enters, and the two allocations deliver the same output. □

B.3 Analytical Results: Cobb-Douglas Case

This subsection derives closed-form solutions for the competitive equilibrium and the planner's problem under Cobb-Douglas production, providing analytical intuition for the efficiency gains from labor reallocation. This is the case $\sigma = 1/(1-\alpha)$, which is the specification we estimate.

Let $F(\cdot)$ be the sum of Cobb-Douglas functions in labor and match productivity $A_g(z)$, so that establishment z solves

$$\max_{l_g} \sum_g l_g^\alpha A_g(z)^{1-\alpha} - \sum_g w_g l_g.$$

The first-order condition gives $l_g^c(z) = (\alpha/w_g)^{1/(1-\alpha)} A_g(z)$. Letting $\bar{A}_g := \int A_g(z) d\Phi(z)$, market clearing $\int l_g^c(z)\, d\Phi(z) = L_g$ then yields

$$w_g = \alpha \left(\frac{\bar{A}_g}{L_g}\right)^{1-\alpha} \quad \text{and} \quad l_g^c(z) = L_g \frac{A_g(z)}{\bar{A}_g},$$

so each establishment's demand for group g is proportional to its productivity for that group relative to the economy-wide average. Substituting into the production function and integrating,

$$Y^c := \int y^c(z) d\Phi(z) = \int \sum_g \left(\frac{L_g}{\bar{A}_g}\right)^\alpha A_g(z) d\Phi(z) = \sum_g L_g^\alpha \bar{A}_g^{1-\alpha}. \tag{B.14}$$

Planner's Problem. Let the planner reallocate workers across all groups, which by Corollary

1 gives the maximal gain. The problem is

$$\max_{l_g(z)\geq 0} \int \sum_g l_g^{\alpha} A_g(z)^{1-\alpha} d\Phi(z) \quad s.t \quad \int \sum_g l_g(z) d\Phi(z) = L,$$

with first-order conditions $\alpha l_g^{\alpha-1} A_g(z)^{1-\alpha} = \lambda$ for all g and z. Marginal productivity is therefore equalized across all worker groups and establishments, which requires $l_{g'}(z') = l_g(z) A_{g'}(z')/A_g(z)$ for every g, g', z, z'. Substituting into the resource constraint gives

$$l_g^p(z) = \frac{A_g(z)}{\bar{A}} L, \qquad \bar{A} := \sum_g \bar{A}_g,$$

so group g receives a share $\bar{A}_g/\bar{A}$ of total labor. Aggregate output is $Y^p = L^{\alpha} \bar{A}^{1-\alpha}$, and the maximal gain from reallocating labor supply is

$$\Delta\% Y := \frac{Y^p}{Y^c} - 1 = \frac{L^{\alpha} \bar{A}^{1-\alpha}}{\sum_g L_g^{\alpha} \bar{A}_g^{1-\alpha}} - 1.$$

B.4 Identification Strategy and Theory

This subsection describes what the targeted moments identify. We work with the baseline calibration, in which $\sigma = 1/(1-\alpha)$ and the technology is the sum of Cobb–Douglas functions of Appendix B.3.

Fix an experience bin j and take the skill supplies $L_{ij} := g(h_i, x_j)$ as given. Firm optimization and market clearing give, for each skill $i \in \{L, H\}$ and establishment type $k \in \{S, L\}$,

$$l_{ijk} = L_{ij} \frac{A_{ijk}}{\bar{A}_{ij}}, \quad w_{ij} = \alpha \left(\frac{\bar{A}_{ij}}{L_{ij}} \right)^{1-\alpha}, \tag{B.15}$$

where $\bar{A}_{ij} = \sum_k A_{ijk} \phi_k$ and ϕ_k is the mass of type-k establishments.

Static moments. Let $\theta_{ij} := \phi_S A_{ijS}/\bar{A}_{ij}$ be the share of skill-i workers employed at small establishments, $W_{ij} := w_{ij} L_{ij}$ their wage bill, and s_{jk} the high-skill share at establishment type k. Market clearing implies $l_{jS} + l_{jL} = L_{Lj} + L_{Hj}$ for any $\{A_{ijk}\}$, so each bin yields three

restrictions on four productivities:

$$\begin{aligned} L_{Lj}\theta_{Lj} + L_{Hj}\theta_{Hj} &= l_{jS}, \\ \theta_{Lj}W_{Lj} + \theta_{Hj}W_{Hj} &= w_{jS}\, l_{jS}, \\ (1-\theta_{Lj})W_{Lj} + (1-\theta_{Hj})W_{Hj} &= w_{jL}\, l_{jL}. \end{aligned} \tag{B.16}$$

Given the split, (B.16) determines (W_{Lj}, W_{Hj}) and (B.15) determines A_{ijk}, so productivities are identified up to one parameter per bin, s_{jS}. The restriction $w_{Hj} > w_{Lj}$ selects the branch of solutions on which large establishments employ relatively more high-skill workers.

Objects invariant to the split. With $\sigma = 1/(1-\alpha)$, the output of each skill group in each bin is W_{ij}/α, and the joint gain is $\mathcal{M}_r(w_{ij})/\mathcal{M}_1(w_{ij}) - 1$, where $\mathcal{M}_r$ is the employment-weighted power mean of order $r = 1/(1-\alpha)$. The experience gain is the same ratio computed over bin-average wages, and at $\tau \to 0$ the paths of output and of the large-establishment output share under aging depend only on wage bills by bin and establishment type. None of these depends on s_{jS}. Since wages are common across establishments, $A_{HjL}/A_{LjL} > A_{HjS}/A_{LjS}$ if and only if $w_{jL} > w_{jS}$. The direction of sorting is therefore identified, but not its strength, which determines w_{Hj}/w_{Lj} and hence the joint and skill gains.

Tenure. At $\tau \to 0$ establishments reach their frictionless employment every period, so tenure depends on $\{\theta_{ij}\}$, $(\{p_j\}, q)$ and η, but not on wages given the allocation: when skill mixes differ across establishment types, workers whose skill changes must move between them. We estimate $\{s_{jS}\}_{j=1}^{J}$, $\{p_j\}$ and η on tenure by experience and establishment type and on its distribution across types, solving (B.16) in closed form so that the static moments hold exactly.

Entry. Skill is unobserved at every experience level, including entry: g_0 sets the share of high-skill entrants, and their split across establishment types is estimated like the split in any other bin. The entry size wage gap must then come from composition, $w_{1L}/w_{1S} = (1+\pi_1 s_{1L})/(1+\pi_1 s_{1S})$ with π_1 the entry skill premium, so weaker initial sorting requires a larger entry premium. At the estimates, tenure pins down the initial sorting on its own: the dispersion of entry wages implied by the two skill levels stays below the dispersion in the data, so restriction (iii) below does not bind.

Restrictions. Tenure is less informative about s_{jS} late in careers, where small changes in the split imply large changes in the premium. We impose (i) single-peaked $\log A_{i\cdot k}$ and p_j in experience; (ii) $[s_{jL}/(1-s_{jL})]/[s_{jS}/(1-s_{jS})] \leq 10$, well above the corresponding odds ratio for a college degree; and (iii) $\sqrt{s_{jk}(1-s_{jk})}\,\pi_j/(1+s_{jk}\pi_j) \leq CV_{jk}$, with $\pi_j = w_{Hj}/w_{Lj} - 1$ and CV_{jk} the coefficient of variation of wages in the data: the dispersion that two skill levels imply in a cell cannot exceed the dispersion in the data, which also reflects other sources of dispersion. At the estimates, restriction (ii) binds in two experience bins, 11–15 and over 40 years, and restriction (iii) binds nowhere: the largest ratio of the implied to the observed dispersion is 0.94, among entrants at small establishments. None of the restrictions affects the experience gain.

B.5 Computational Algorithm for Dynamic Equilibrium

This subsection describes the numerical algorithm used to solve for the dynamic competitive equilibrium. The key challenge is that establishments' current employment decisions depend on expected future reallocation costs, which themselves depend on future employment choices.

In our setting, two special experience bins do not require employment information in two periods when solving for the current employment choice. First, the establishment does not incur reallocation costs when employing workers with the least experience, as there is no employment stock for this group. Second, for bins with the most experience (i.e., the oldest workers), the establishment does not need to account for reallocation costs it would incur in the next period, as they all retire. Therefore, we can solve the employment choices for the youngest workers first and proceed to solve them *sequentially* from experience group x_j to x_{j+1}, where $j \in [1, J-1]$. Other than the wage schedule, key equilibrium objects needed for solving the employment choices are the expected firing probabilities of a worker with h_i and x_j hired this period (i.e., $\mathbb{1}[\tilde{L}_{i,j+1} > l_{i,j+1}]]$ and $\mathbb{1}[\tilde{L}_{i+1,j+1} > l_{i+1,j+1}]]$).

Practically, we solve the labor market equilibrium using the following steps.

1. First, we solve the model without the reallocation cost and obtain the wage schedule $w^{no}_{i,j}$ and establishments' employment choices, $l^{no}(z)_{ij}$ $(j \in [1,2,3,...,J])$, and employment, stocks $\tilde{L}^{no}_{ij}$ $(j \in [2,3,...,J])$. Based on these solved equilibrium objects, we construct the expected reallocation costs that show up in the right-hand side of equations

(8) and (9), denoted by FC^0_{ij}.

2. Next, we solve for $w_{i,1}$ that equalizes the labor supply in the submarket $(g(h_i, x_1))$ to the aggregated labor demand across all active establishments whose establishment-level labor demand is denoted by $l_{i,1}(z)$. For this step, we assume that each establishment uses the expected reallocation cost derived in the previous step. Having solved $w_{i,1}$, we calculate establishments' employment choices of the youngest workers, $l_{i,1}(z)$, and their employment stocks of the second youngest experience bins $\tilde{L}_{i,2}(z)$.

3. Third, we solve for $w_{i,j}$ (where $j \in [2, J]$) that equalizes the labor supply in the submarket $(g(h_i, x_j))$ to the aggregated labor demand across all active establishments *sequentially* (i.e., from experience group x_{j-1} to experience group x_j). Importantly, we use the *derived* $\tilde{L}_{i,j}(z)$ calculated from the employment choice made for the experience group x_{j-1}. For this step, we also assume that each establishment uses the expected reallocation cost derived in the first step. Note that when $j = J$, there is no need to calculate the expected reallocation cost as all workers retire at the end of experience x_J.

Using the above three steps, we derive the wage schedule, $w_{i,j}$, employment choices and stocks, $l_{i,j}(z)$ and $\tilde{L}_{i,j}(z)$, which lead to the newly calculated values of the expected reallocation cost:

$$\begin{aligned} FC^1_{ij}(z): \;=\; & \beta\eta\tau\delta_j\Big[(1-p_j-q+2p_jq)w_{ij+1}\mathbb{1}[\tilde{L}_{i,j+1} > l_{i,j+1}] \\ & +p_j(1-q)w_{i+1,j+1}\mathbb{1}[\tilde{L}_{i+1,j+1} > l_{i+1,j+1}] + q(1-p_j)w_{i-1,j+1}\mathbb{1}[\tilde{L}_{i-1,j+1} > l_{i-1,j+1}]\Big], \end{aligned}$$

where $w_{i,j}$, $l_{i,j}(z)$ and $\tilde{L}_{i,j}(z)$ are derived in the second and third steps. With $I = 2$, a low-skill worker becomes high-skill with probability $p_j(1-q)$ and a high-skill worker becomes low-skill with probability $q(1-p_j)$, so the equation simplifies to,

$$\begin{aligned} FC^1_{Lj}(z) &= \beta\eta\tau\delta_j\Big[\big(1-p_j(1-q)\big)\,w_{L,j+1}\,\mathbb{1}[\tilde{L}_{L,j+1} > l_{L,j+1}] + p_j(1-q)\,w_{H,j+1}\,\mathbb{1}[\tilde{L}_{H,j+1} > l_{H,j+1}]\Big], \\ FC^1_{Hj}(z) &= \beta\eta\tau\delta_j\Big[q(1-p_j)\,w_{L,j+1}\,\mathbb{1}[\tilde{L}_{L,j+1} > l_{L,j+1}] + \big(1-q(1-p_j)\big)\,w_{H,j+1}\,\mathbb{1}[\tilde{L}_{H,j+1} > l_{H,j+1}]\Big]. \end{aligned}$$

In the final step, we update the expected reallocation cost by taking a weighted average of

the assumed expected cost in previous steps and the newly calculated one,

$$FC_{ij}^{2} = \zeta FC_{ij}^{1} + (1-\zeta)FC_{ij}^{0},$$

where $\zeta \in (0,1)$ is the adjustment factor. We then iterate by using FC_{ij}^{2} in place of FC_{ij}^{0} and repeating steps two and three. We stop when the newly calculated wage schedule and employment choices coincide with those from the previous iteration.[26]

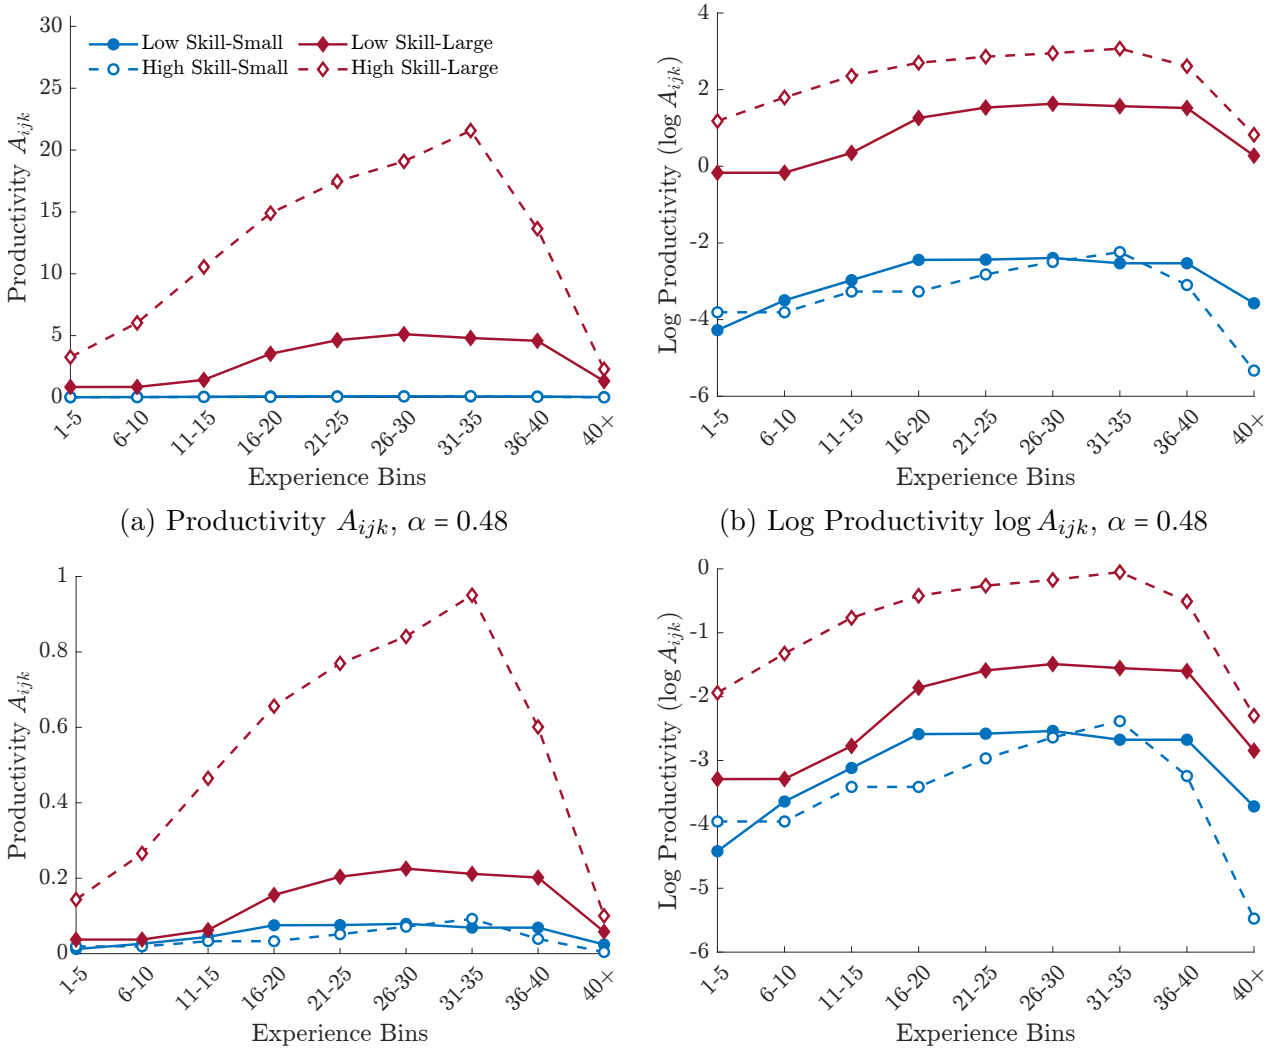
At the edge of the inaction region, the iteration can alternate: an establishment that expects to dismiss workers hires fewer of them and then has no one to dismiss. When $\tau > 0$ we therefore replace $\mathbb{1}[\tilde{L}_{i,j+1} > l_{i,j+1}]$ with the probability $1/(1+e^{-f/s})$, where f is the share of next period's stock the establishment would dismiss and $s = 0.05$, as if stocks were subject to a small idiosyncratic shock. Lowering s to 0.02 changes output by 0.04 percent.

C Additional Quantitative Results

This section presents a sensitivity analysis of key structural parameters α, σ, and q from Sections 5 and 6.

[26] In such a case, the newly calculated expected reallocation cost is also the same as the one calculated from the previous iteration.

Figure C.1: Estimated Worker-Establishment Productivities by Experience A_{ijk} Varying α



(a) Productivity A_{ijk}, $\alpha = 0.48$

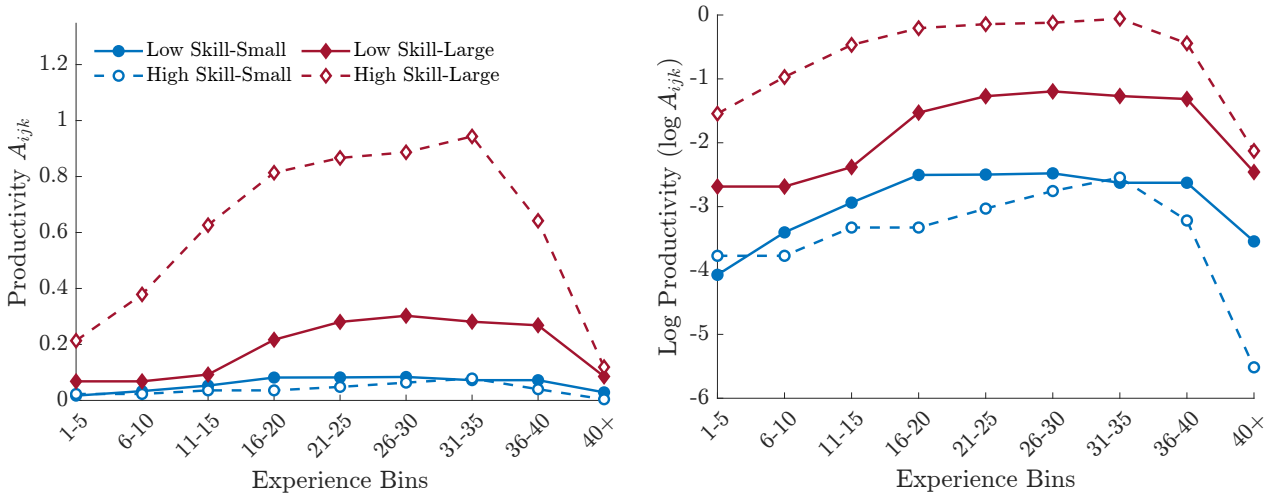
(b) Log Productivity $\log A_{ijk}$, $\alpha = 0.48$

(c) Productivity A_{ijk}, $\alpha = 0.72$

(d) Log Productivity $\log A_{ijk}$, $\alpha = 0.72$

Notes: Estimated match-specific productivities for skill type $i \in \{L, H\}$, experience group $j \in \{1, ..., 9\}$, and establishment type $k \in \{S, L\}$. Panel (a) shows productivity levels; panel (b) shows log productivities. Each line represents a skill-establishment pair, with productivity varying across experience groups.

Figure C.2: Estimated Worker-Establishment Productivities by Experience A_{ijk} Varying σ



(a) Productivity A_{ijk}, $\sigma = 2$

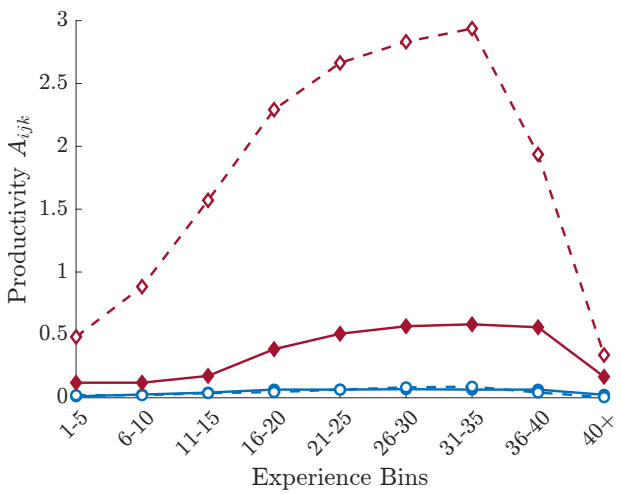
(b) Log Productivity $\log A_{ijk}$, $\sigma = 2$

(c) Productivity A_{ijk}, $\sigma = 3$

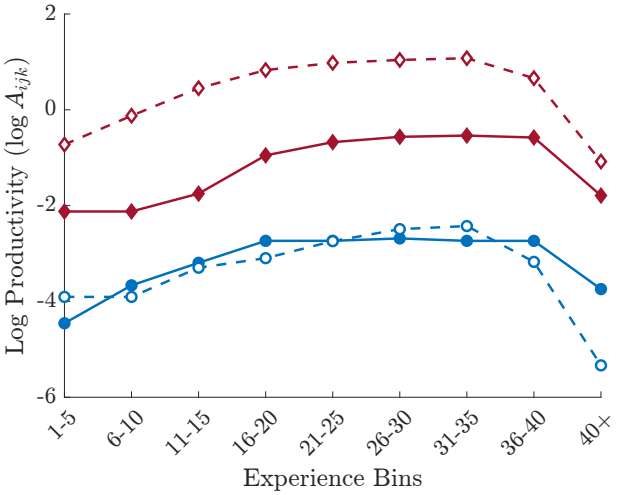
(d) Log Productivity $\log A_{ijk}$, $\sigma = 3$

Notes: Estimated match-specific productivities for skill type $i \in \{L, H\}$, experience group $j \in \{1, ..., 9\}$, and establishment type $k \in \{S, L\}$. Panel (a) shows productivity levels; panel (b) shows log productivities. Each line represents a skill-establishment pair, with productivity varying across experience groups.

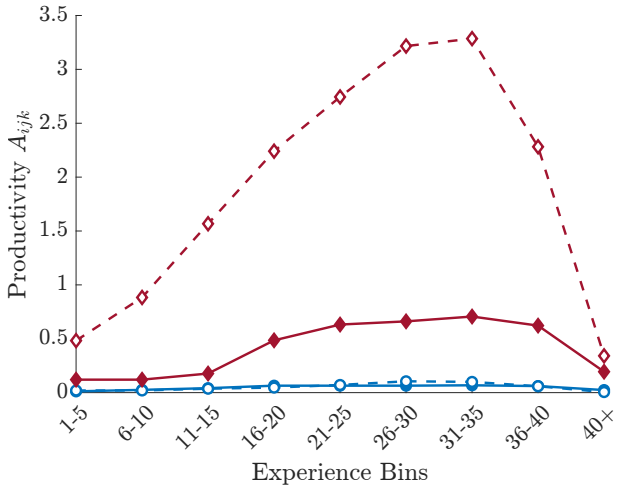
Figure C.3: Estimated Worker-Establishment Productivities by Experience A_{ijk} Varying q



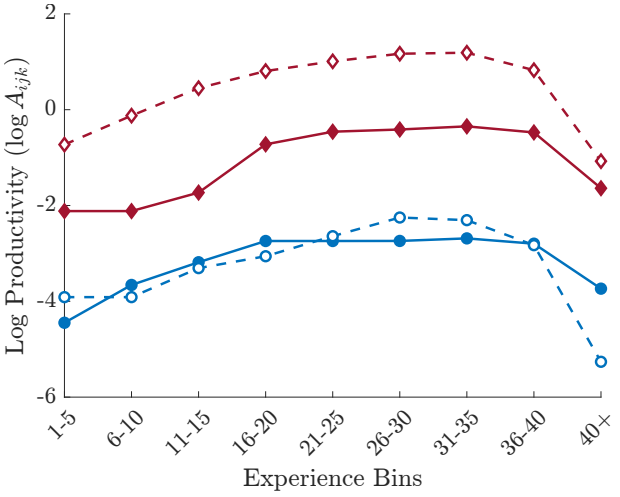
(a) Productivity A_{ijk}, $q = 0.208$



(b) Log Productivity $\log A_{ijk}$, $q = 0.208$



(c) Productivity A_{ijk}, $q = 0.312$



(d) Log Productivity $\log A_{ijk}$, $q = 0.312$

Notes: Estimated match-specific productivities for skill type $i \in \{L, H\}$, experience group $j \in \{1, ..., 9\}$, and establishment type $k \in \{S, L\}$. Panel (a) shows productivity levels; panel (b) shows log productivities. Each line represents a skill-establishment pair, with productivity varying across experience groups.

Table C.1: Population Aging under Alternative Calibrated Parameters

		α		σ		q	
	Benchmark	0.48	0.72	2	3	0.208	0.312
Output relative to its initial level							
After 5 years	+0.64%	+0.71%	+0.56%	+0.46%	+0.88%	+0.64%	+0.64%
After 20 years	−0.12%	−0.10%	−0.03%	−0.22%	+0.06%	−0.12%	−0.12%
After 40 years	+0.24%	+0.29%	+0.26%	+0.09%	+0.47%	+0.24%	+0.24%
Aggregate skill premium							
Initial	54%	104%	20%	27%	82%	55%	63%
After 40 years	67%	124%	28%	38%	98%	65%	79%
Output share of large establishments							
Initial	80.1%	96.0%	56.7%	61.7%	91.2%	80.1%	80.1%
After 40 years	79.9%	96.0%	56.1%	61.3%	91.1%	79.9%	79.9%

Notes: The population aging counterfactual of Section 6 at $\tau = 0$, run on the model re-estimated with one parameter at the value in the column header and the others at their benchmark values.