Labor Market Fluidity and Human Capital Accumulation *

Niklas Engbom
NYU, CEPR & NBER

April, 2026

Abstract

Using panel data from 14 advanced economies, I document that life-cycle wage growth is steeper in countries with higher rates of job-to-job mobility. This pattern is not simply the mechanical consequence of wage gains at job changes: wages also grow faster within continuing employment spells, and young workers in more fluid labor markets are more likely to work at large, training-intensive firms. I interpret these facts through an estimated equilibrium model of careers in which firms differ in both productivity and the quality of their learning environment. Lower fluidity increases mismatch and depresses human capital accumulation through two mechanisms: it slows young workers' movement into high-learning firms and lowers the return to skill by making future mismatch more likely. The estimated model accounts for the cross-country relationship between fluidity and life-cycle wage growth. Moving from the least- to the most-fluid labor market raises life-cycle wage growth by more than 30 percentage points, while the aggregate stock of human capital is more than 20 percent higher.

*Niklas Engbom: ne466@nyu.edu. This paper incorporate new cross-country data on the allocation of workers across heterogeneous firms and proposes a richer theory of sorting of workers across firms that differ in multiple attributes relative to a much earlier draft.

1 Introduction

A large literature studies how labor market institutions affect the allocation of workers across firms and the resulting level of aggregate productivity (Bentolila and Bertola, 1990; Hopenhayn and Rogerson, 1993; Ljungqvist and Sargent, 1998). This work emphasizes firms' job creation and destruction decisions and the static misallocation of labor across producers. Yet labor markets do more than allocate a given stock of skills across firms. They also allocate workers, especially young workers, to jobs that shape the work experiences they accumulate over their careers. Through this channel, frictions and regulations that impede mobility may reduce productivity not only by putting workers in the wrong firms at a point in time, but also by changing the human capital workers carry into the future.

This paper quantifies the dynamic consequences of cross-country differences in *labor market fluidity*—the rate at which workers switch employers without an intervening spell of unemployment. I combine individual-level panel data from 14 advanced economies with an equilibrium model of careers in which firms differ in both contemporaneous productivity and the quality of their learning environment. The analysis yields three main empirical findings. First, more fluid labor markets exhibit greater life-cycle wage growth. Second, this relationship is not accounted for solely by wage gains at job changes: wages also grow faster within continuing employment spells. Third, young workers in more fluid labor markets sort into large, training-intensive firms, precisely the firms in which young workers' wages grow faster. Quantitatively, I find that moving from a low-fluidity to a high-fluidity market raises life-cycle wage growth by more than 30 percentage points, matching the data closely. The human-capital response is central to this result, more than doubling the static misallocation effect of differences in labor market fluidity.

I begin by establishing three empirical facts on cross-country differences in career outcomes using individual-level panel data for 14 advanced economies in Western Europe and the United States between 1993 and 2024. I measure labor market fluidity as the share of employed workers who were hired directly from another employer in the past year. This measure can be computed systematically across countries and over time, and it varies sharply across countries. In the least-fluid countries, fewer than five percent of employed workers are hired directly from another employer in a given year; in the most-fluid countries, the share is close to 15 percent.

The first fact is that life-cycle wage growth is steeper in more fluid labor markets. I estimate wage profiles using the panel dimension of the data, controlling for individual fixed effects and country-year fixed effects, under a restriction on wage growth late in life. Between ages 20–29 and 50–59, real wages grow by about 15 percent in the least-fluid countries and by more than 50 percent in the most-fluid countries. This relationship remains after allowing wage profiles to differ by education, gender, and occupation, and it is especially pronounced for college-educated workers. Entry wages, by contrast, do not vary systematically with fluidity. Workers in high-fluidity countries therefore do not simply start from a lower base and catch up. Instead, they experience faster wage growth over their careers from a similar starting point.

The second fact is that the relationship between fluidity and life-cycle wage growth does not merely reflect the cumulative effect of wage gains at job-to-job transitions. Job-to-job transitions are associated with wage gains in the data, and this direct job-ladder channel is non-trivial. Yet young workers' wages also grow faster in more fluid countries in years in which they remain with the same employer. The cross-country relationship between labor market fluidity and life-cycle wage growth therefore reflects more than the mechanical effect of more frequent wage jumps at job changes.

The third fact is that young workers in more fluid labor markets sort more often into firms that appear to provide better learning opportunities. Consistent with recent evidence that firms differ importantly in their *learning environment* (Gregory, 2026), I document that young workers experience faster wage growth when they work at larger firms and at firms that offer training. Moreover, young workers in more fluid labor markets are more likely to work at large and training-intensive firms. Taken together, these patterns suggest that fluid labor markets facilitate the allocation of young workers to high-learning environment jobs.

Guided by this evidence, I develop and estimate an equilibrium model of careers in which jobs differ along two dimensions—contemporaneous productivity and the quality of the learning environment—extending the work of Gregory (2026) to a general equilibrium setting. Workers enter the labor market unemployed, search for jobs while unemployed and employed, and accumulate human capital at a rate that depends on the learning environment of their employer. Wages are determined by bargaining and respond to outside offers as in Cahuc, Postel-Vinay and Robin (2006). Differences in fluidity affect wage growth through three channels: the speed of transition to more productive jobs, the allocation of workers across learning environments and hence the growth rate of human capital, and the number of outside offers workers receive to bargain up their wage.

A key force in the model is that the value of a good learning environment depends on expected future mismatch. By increasing the probability that workers will later find matches in which their accumulated human capital is well used, a more fluid labor market raises the return to skill. Fluidity therefore raises human capital accumulation in two ways. It allows young workers to find good learning environment firms more quickly, and it tilts young workers' preferences toward such firms. This shift in worker preferences, in turn, increases the recruiting intensity of firms with better learning environments.

I estimate the model to match a hypothetical medium-fluidity economy using my panel data on wage growth, job mobility, employment dynamics, firm size, and training. In particular, the dispersion in learning environments is disciplined by the excess wage growth young workers experience when working at training firms. Although the model is overidentified, it fits the main patterns of wage growth, mobility, and sorting across firm types by age. The estimates imply that human capital accumulation is a central source of wage growth over the career.

I then use the estimated model to quantify the consequences of policies and frictions that reduce labor market fluidity. To this end, I vary the cost of vacancy creation and employed search intensity to match cross-country differences in job-finding rates from unemployment and employment. I interpret vacancy creation costs as a reduced-form representation of policies and frictions that raise the cost of creating jobs. Employed search intensity captures less intense job search by employed workers when it is harder to find a job. The exercise compares a low-fluidity economy, a high-fluidity economy, and a very-fluid economy calibrated to span the range of fluidity observed in the data.

The model implies large effects of observed differences in labor market fluidity on both life-cycle and aggregate outcomes. Moving from the least to the most fluid economy raises wage growth between ages 20–29 and 50–59 by more than 30 percentage points, accounting for essentially all of the empirical cross-country relationship. Human capital is central to this result: it grows more than 10 percentage points more over the life cycle in the most-fluid economy than in the least-fluid economy. Better match productivity also contributes, but by less. In the aggregate, the stock of human capital is more than 20 percent higher, while average match productivity is roughly 10 percent higher.

These results change how we should interpret policies that reduce labor mobility. The standard concern is that such policies distort reallocation and lower the level of productivity by preventing workers from reaching the firms where they are most productive. This paper shows that the same policies can also distort the production of human capital. When young workers face a labor market in which future mismatch is likely, high-learning jobs become less attractive and are harder to find. The resulting loss is not merely a lower wage today, but a lower stock of skills tomorrow.

Related literature. This paper contributes first to the literature on labor market institutions, reallocation, and aggregate productivity. Classic work shows that firing costs, employment protection, and other labor-market distortions change firms' job creation and destruction decisions and generate allocative losses (Bentolila and Bertola, 1990; Hopenhayn and Rogerson, 1993; Ljungqvist and Sargent, 1998; Pries and Rogerson, 2005). Related empirical and structural work documents large cross-country differences in worker flows, unemployment dynamics, and job-to-job mobility (Jolivet, Postel-Vinay and Robin, 2006). My contribution is to show that these differences matter not only because they affect which workers are matched to which firms at a point in time, but also because they shape the path of human capital accumulation over the career. The paper therefore complements the misallocation view of labor-market regulation with a dynamic, career-based channel.

Second, the paper contributes to the literature on job mobility and wage dynamics. A long tradition studies how search frictions, employer competition, and job ladders generate wage growth and wage dispersion (Diamond, 1982; Mortensen and Pissarides, 1994; Burdett and Mortensen, 1998; Postel-Vinay and Robin, 2002; Cahuc, Postel-Vinay and Robin, 2006). Empirically, job changes are especially important early in the career and account for a substantial share of young workers' wage growth (Topel and Ward, 1992). Structural models with on-the-job search and human capital accumulation decompose wage growth into returns to experience, tenure, search, and employer heterogeneity (Bagger et al., 2014). Relative to this work, I emphasize a cross-country object: the extent to which differences in labor market fluidity explain differences in life-cycle wage growth. The key distinction is that mobility affects not only the level of wages through better matches and bargaining, but also the rate of wage growth by changing the firms at which workers accumulate skills.

Third, the paper relates to work on human capital and life-cycle wage profiles. The idea that workers invest in skills early in life and reap the returns later is central to human capital theory (Becker, 1962; Ben-Porath, 1967). Recent work uses differences in life-cycle wage growth across countries to measure differences in human capital accumulation and assess their aggregate implications (Lagakos et al., 2018). I build on this perspective but focus on advanced economies and on labor-market fluidity as a determinant of human capital accumulation. The analysis shows that differences in career wage growth across rich countries are closely linked to the functioning of labor markets, not only to schooling or aggregate technology.

Finally, the paper contributes to the growing literature on firm heterogeneity in wage growth and learning. Matched employer–employee studies document that firms differ systematically in pay, productivity, and the workers they employ (Abowd, Kramarz and Margolis, 1999; Lise, Meghir and Robin, 2016). More recent work emphasizes that firms also differ in their ability to promote worker learning and that these differences matter for life-cycle earnings inequality (Gregory, 2026; Arellano-Bover and Saltiel, 2026; Arellano-Bover, 2024). I use this insight to study cross-country differences in careers. The paper documents that young workers sort more often into high-learning firms in fluid labor markets and embeds this sorting margin in an equilibrium model of vacancy creation and on-the-job search. This allows me to quantify how labor-market institutions affect not only sorting and wages, but also the aggregate stock of human capital.

This paper is structured as follows. Section 2 establishes three motivating facts on cross-country differences in life-cycle career outcomes. Section 3 presents the theory, which Section 4 estimates. Section 5 uses the estimated model to understand the role of differences in labor market fluidity in generating the observed cross-country patterns, and to quantify their aggregate implications.

2 Three Facts on Cross-Country Differences in Career Outcomes

I document three facts about cross-country differences in worker career outcomes. The facts link labor market fluidity to life-cycle wage growth and to the allocation of young workers across firms with different learning environments.

2.1 Data

I use individual-level panel data from the European Community Household Panel (ECHP), the European Union Statistics on Income and Living Conditions (EU-SILC), and the Panel Study of Income Dynamics (PSID). Together, these data provide 21–30 years of information for advanced economies on employment histories, wages, and worker and firm characteristics. The analysis sample contains 14 countries after the access and data-quality restrictions described below.

ECHP. The ECHP ran from 1993 to 2001 and covered 15 Western European countries: Austria (starting in 1995), Belgium, Denmark, Finland (starting in 1994), France, Germany, Greece, Ireland, Italy, Luxembourg, the Netherlands, Portugal, Spain, Sweden (starting in 1997), and the U.K. Five features of the data are relevant for the analysis.

First, Germany participated in the ECHP only from 1994 to 1996, but partly overlapping data from the German Socio-Economic Panel (SOEP) for 1994–2001 are included in the ECHP files. Because the German component of the original ECHP suppresses the interview month for confidentiality reasons, I cannot compute labor market fluidity in those data. Since the ECHP and SOEP samples cover different respondents, I use both sources for the other outcomes in order to maximize sample size.

Second, Luxembourg participated in the ECHP only from 1994 to 1996. For 1995–2001, the ECHP includes complementary data from the Panel Socio-Economique Liewen zu Lëtzebuerg (PSELL). These data do not report the start date of the current job for 1995–1997 and do not record the survey month throughout. I therefore cannot compute labor market fluidity in the PSELL, and hence cannot compute it for Luxembourg in 1997–2001. Since the ECHP and PSELL samples cover different respondents, I use both sources for the other outcomes in order to maximize sample size.

Third, each country's statistical agency must grant researchers access to its national components of the ECHP and EU-SILC. The Netherlands has not granted me access to its ECHP or EU-SILC files, so I exclude the Netherlands from the analysis.

Fourth, Sweden did not directly participate in the ECHP, instead providing complementary data from a national survey. Because this survey lacks a panel dimension, it cannot be used for the main analyses below. I therefore exclude Sweden from the ECHP sample.

Fifth, the U.K. participated in the ECHP only from 1994 to 1996, with overlapping data from the British Household Panel Survey (BHPS) for 1994–2001. Although the BHPS contains the variables needed to construct the main outcomes, the implied fluidity series displays implausibly large year-to-year fluctuations. I therefore exclude the BHPS when computing labor market fluidity.

EU-SILC. The EU-SILC was designed as the successor to the ECHP. The files used here cover 2004–2024, with several exceptions. Denmark and Luxembourg also provide data for 2003, while

Denmark and Ireland end in 2023. Data for Greece start in 2006, data for Germany start in 2015 and are unavailable in 2020, and data for the U.K. end in 2018.

Eurostat releases the EU-SILC in separate cross-sectional and longitudinal files, even though the underlying data are collected as a panel. The two files contain different variables; for example, firm size and sector are included only in the cross-sectional file. I use the cross-sectional file to impute sector and firm size in the longitudinal file based on demographic characteristics that are available in both files.

The question needed to compute labor market fluidity is not asked in the EU-SILC after 2020. My EU-SILC-based measure of fluidity therefore ends in 2020.

PSID. The PSID provides annual data from 1984 to 1997 and biennial data thereafter. To compute labor market fluidity, I rely on retrospective questions about employment status in each month of the calendar year before the survey year. Specifically, I use the subsequent survey to recover employment status in the calendar months around the previous survey. Starting in 2001, the PSID asks about employment status in each month of the previous two calendar years, which allows me to compute labor market fluidity. The 1999 and 2001 surveys, however, record employment status only for the previous calendar year, so I cannot compute labor market fluidity for 1997–2000.

The PSID records firm size only from 2005 onward, apart from additional measures in 1993 and 1999. Because the survey's earnings questions have changed over time, I harmonize earnings as total wage, salary, and self-employment income in order to maintain consistency with the European data.

2.2 Sample Selection and Variable Definitions

In most of the analysis, I focus on men and women aged 20–59 who are older than the imputed completion age for their highest degree. I set the completion age to 17 for workers with less than an upper secondary degree, 19 for workers with an upper secondary degree, and 23 for workers with a tertiary degree.[1] The cross-country share of young workers enrolled in school is not systematically correlated with labor market fluidity.

Because the data do not consistently distinguish private from public employment, or wage employment from self-employment, I classify private-sector wage employees, public-sector wage employees, and self-employed workers as employed. The non-employed are all other individuals. Although I cannot systematically distinguish public from private employment across all data sets, I can do so in the ECHP. In those data, public employees are concentrated in a small number of

[1] I do not consistently observe the age at which a respondent completed schooling. I do, however, observe the share of workers in each education group who are enrolled in school by age. Among workers with less than an upper secondary degree, XX percent are recorded as students at age XX. Among workers with an upper secondary degree, XX percent are recorded as students at age XX. Among workers with a tertiary degree, XX percent are recorded as students at age XX.

one-digit occupations, primarily public administrators and teachers. Since occupation is observed consistently across surveys and over time, I can verify that the results below are similar within occupations that consist primarily of private-sector employees. Excluding the self-employed also produces similar patterns.

I recode education into three categories: less than upper secondary, upper secondary, and tertiary. Occupation is standardized into one-digit ISCO codes. Income is total wage and salary income, inclusive of overtime pay and bonuses, plus self-employment income, all measured over the previous calendar year. Hourly wages are total income divided by total hours worked, where total hours are the product of weeks worked and usual hours per week. Nominal wages are first converted to real 2023 local-currency wages and then to real U.S. dollars using 2023 PPP-adjusted exchange rates.

Labor market fluidity. The available data do not allow me to construct monthly job-to-job mobility consistently across countries and years. I instead measure the annual *poaching rate*: the share of currently employed workers who started working for their current employer in the past 12 months and were never non-employed during that period. This measure is available in the European data through 2020 and in the U.S. data through 2023, subject to the survey-specific gaps described above.

Life-cycle wage growth. I estimate life-cycle wage growth from the panel dimension of the data. Specifically, I project log hourly wages on individual fixed effects, country-year fixed effects, and country-age effects. The country-age effects are restricted to be flat after age 45, which resolves the perfect collinearity among individual, year, and age fixed effects:

$$\log w_{it} \quad = \quad \alpha_i + \gamma_{ct} + \xi_{ca} + \varepsilon_{it}. \tag{1}$$

The objects of interest are the country-specific age effects, ξ_{ca}.

2.3 Three Motivating Facts

I now present three empirical facts that motivate the model.

Wages Grow More Over the Life Cycle in More Fluid Labor Markets. Figure 1 plots wage growth between ages 20–29 and 50–59, based on the estimates from (1), against the share of workers who were poached in the past year. Three observations stand out. First, countries differ substantially in labor market fluidity: in the most fluid labor market, workers are more than twice as likely to be poached in a given year as in the least fluid labor market. Second, countries also differ substantially in life-cycle wage growth. Between ages 20–29 and 50–59, real wages grow by less than 15 percent in Greece, holding fixed aggregate growth over calendar time, but by more than 50 percent in Sweden. Third, life-cycle wage growth is steeper in countries with more frequent job-to-job mobility.

Figure 1: Life-Cycle Wage Growth and Labor Market Fluidity

Notes: Each marker is a country. Wage growth is the estimated increase in log hourly wages between ages 20–29 and 50–59 after removing individual and country-year fixed effects as in (1). Labor market fluidity is the annual poaching rate. *Source:* ECHP 1993–2001; EU-SILC 2003–2024; PSID 1984–2023.

A natural concern is that these patterns reflect differences in workforce composition. To assess this possibility, I project wages on the interaction between aggregate labor market fluidity and the restricted life-cycle age profile, while controlling for individual fixed effects, country-year fixed effects, and restricted age effects interacted with education (college versus non-college), gender, and one-digit occupation:

$$\log w_{it} \quad = \quad \beta\left(\text{Fluidity}_c \times \text{Age}_{it}\right) + \alpha_i + \gamma_{ct} + \xi_{ae} + \xi_{ag} + \xi_{ao} + \varepsilon_{it}.$$

Table 1 summarizes the results. College-educated workers, men, and workers in some occupations tend to experience more wage growth over the life cycle. However, workforce composition along these dimensions does not covary systematically with labor market fluidity. Consequently, differences in education, gender, and occupational composition do not account for the pattern in Figure 1.

Table 1: Life-Cycle Wage Growth, Entry Wages, and Labor Market Fluidity

	(1)	(2)	(3)	(4)	(5)	(6)
	Life-cycle wage growth		By education		Entry wages	
	Raw	Residual	Non-college	College	Raw	Residual
Fluidity	0.010***	0.010***	0.006***	0.014***	0.063	0.035
	(0.002)	(0.002)	(0.002)	(0.004)	(0.130)	(0.117)
Individual-years	1,730,856	1,730,856	1,049,439	681,417	238,512	238,512
Individuals	509,886	509,886	312,167	197,719	122,710	122,710
Years	29	29	29	29	29	29
Countries	14	14	14	14	14	14

Although workforce composition does not drive the result, heterogeneity across subpopulations helps discipline potential explanations. Columns (3) and (4) of Table 1 show that life-cycle wage growth is especially steep in high-fluidity countries for college graduates. Since college-educated workers are less likely to be directly affected by minimum wages or collectively bargained pay scales, this pattern points away from explanations based primarily on wage floors or centralized pay-setting institutions, although it does not rule them out.

The final two columns of Table 1 relate wages at ages 20–29 to labor market fluidity, expressing entry wages in real U.S. dollars using the PPP adjustment described above. Entry wages are modestly higher in more fluid labor markets, but the relationship is not statistically significant. Thus, workers in high-fluidity countries do not start at lower wages and then catch up. Rather, workers across these advanced economies enter the labor market at broadly similar pay levels, but wages subsequently grow faster in more fluid labor markets.

Wages Also Grow More Within Jobs. Across all countries in the sample, job-to-job transitions are associated with wage gains. It is therefore natural to ask whether the steeper life-cycle wage profiles in high-fluidity countries are simply the mechanical consequence of workers making more such transitions. To assess the importance of this direct effect, I project year-on-year wage growth on an indicator for whether the worker is younger than 40 interacted with labor market fluidity, controlling for country-year, age-year, education-year, gender-year, and occupation-year effects. I then add an indicator for whether the worker made a job-to-job transition in the past 12 months:

$$\Delta \log w_{it} \quad = \quad \beta \left(\text{Young}_{it} \times \text{Fluidity}_c\right) + \alpha \, \text{JJ}_{it} + \xi_{cy} + \xi_{ay} + \xi_{ey} + \xi_{gy} + \xi_{oy} + \varepsilon_{it}. \qquad (2)$$

Equation (2) relates the excess wage growth of young workers relative to older workers within a country to aggregate labor market fluidity.

Table 2 summarizes the results. Consistent with Figure 1, young workers experience faster

Table 2: Wage Growth Within and Between Jobs and Labor Market Fluidity

	(1)	(2)
	Baseline	With JJ control
Age 20–39 × fluidity	0.007***	0.006**
	(0.003)	(0.003)
JJ transition in last 12 months		0.028***
		(0.007)
Individual-years	625,936	625,936
Individuals	292,608	292,608
Years	24	24
Countries	14	14

wage growth in more fluid countries. Part of this relationship reflects the direct wage gains associated with job-to-job transitions. Yet column (2) shows that young workers' wages grow faster in more fluid labor markets even in years in which they do not make a job-to-job transition. Put differently, wages also grow more within employment spells in more fluid labor markets.

Young Workers Sort Toward Larger, Training-Intensive Firms. Why do wages grow more on the job in more fluid labor markets? A recent literature emphasizes heterogeneity across firms in their *learning environments* (Gregory, 2026): some firms allow workers to accumulate skills more rapidly than others. If young workers in more fluid labor markets are more likely to match with such firms, this allocation channel can help explain the steeper life-cycle wage growth in those countries.

I begin by regressing wage growth on an indicator for whether the respondent works at a large firm (50 or more employees) as a proxy for learning environment, interacted with an indicator for being younger than 40. The specification controls for individual fixed effects, country-year fixed effects, and the restricted country-age effects used above:

$$\Delta \log w_{it} \quad = \quad \beta \left(\text{Young}_{it} \times \text{Large}_{it}\right) + \xi_i + \xi_{cy} + \xi_{ca} + \varepsilon_{it}. \tag{3}$$

I then replace firm size with an indicator for whether the respondent's firm offers on-the-job training. The training measure is available only in the 1993–2001 ECHP.

Table 3 shows that young workers experience faster wage growth while working at large firms or at firms that offer training. Because the specification includes individual fixed effects, the estimates compare wage growth for the same worker across different firm environments. Older workers, by contrast, experience little differential wage growth across large versus small firms or training versus non-training firms. The wage-growth gains from working at a large firm are especially large for young workers in more fluid labor markets; the gains associated with training,

Table 3: Wage Growth and Wages at Large and Training Firms

	(1)	(2)	(3)	(4)	(5)	(6)
	Baseline		Fluidity interactions		Wage levels	
	Large	Training	Large	Training	Large	Training
Age 20–39	0.018***	0.014***	0.020***	0.012**	0.025***	0.001
	(0.004)	(0.004)	(0.004)	(0.005)	(0.003)	(0.003)
× hired					-0.040***	-0.026***
					(0.004)	(0.005)
× fluidity			0.035***	-0.017		
			(0.012)	(0.020)		
Age 40–59	0.001	-0.002	0.002	-0.004	0.017***	0.004
	(0.004)	(0.005)	(0.004)	(0.005)	(0.003)	(0.003)
× hired					-0.015***	-0.002
					(0.005)	(0.008)
× fluidity			0.020*	-0.019		
			(0.012)	(0.021)		
Individual-years	598,575	172,873	598,575	172,873	633,211	177,259
Individuals	204,354	39,633	204,354	39,633	217,164	41,022
Years	24	7	24	7	23	7
Countries	14	12	14	12	14	12

however, are not systematically related to aggregate fluidity.

The final two columns analyze wage levels. Young and older workers earn more at larger firms, consistent with large firms paying more because they are more productive. At the point of hire, however, young workers earn less at larger firms, while older workers earn similar wages across firms. Young and older workers earn similar wages at training and non-training firms, but young hires earn less at training firms. This pattern is consistent with young workers partly paying for training through lower initial wages.

I next ask whether the allocation of workers across firm types differs across countries. Figure 2 shows that young workers in more fluid labor markets are more likely to work at large firms and at firms that offer training. The evidence is consistent with two related mechanisms: greater fluidity may make it easier for young workers to locate firms where they can accumulate skills, and greater fluidity may increase young workers' willingness to trade off current pay for future wage growth.

3 Model

This section develops a stationary search model in which firms differ both in productivity and in the learning opportunities they provide. The model links aggregate labor-market fluidity to

Figure 2: Employment Shares at Large and Training Firms and Labor Market Fluidity

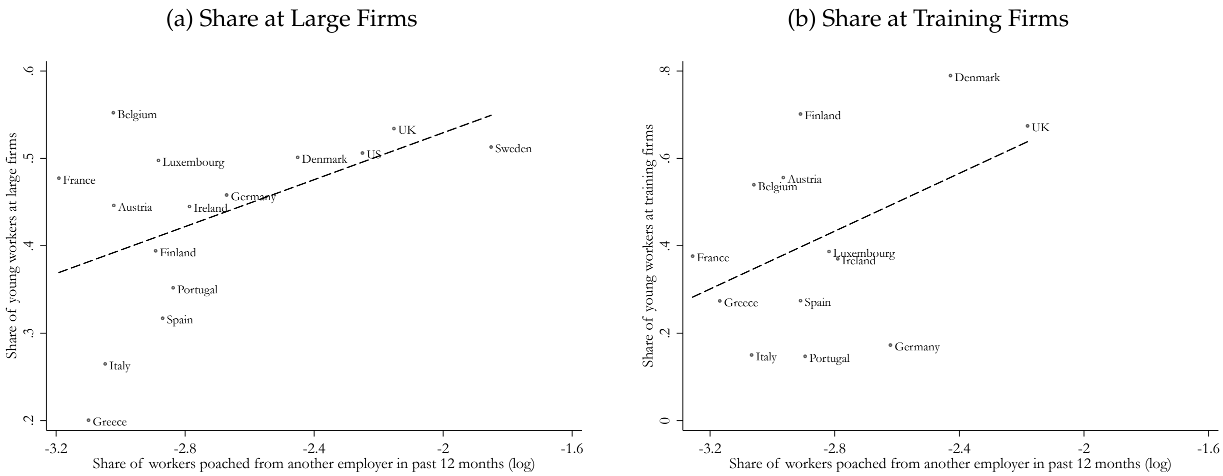

the types of jobs workers accept early in their careers and, through those choices, to subsequent human-capital accumulation and wages. Apart from the different focus, a key difference to Gregory (2026) is that the model is set in general equilibrium.

3.1 Environment

The economy is populated by a unit mass of workers and a mass M of firms.

Demographics. Workers enter the economy at age $a = 0$ as unemployed with normalized human capital $h = 0$. Age is continuous. Workers age deterministically at unit speed until retirement at age A, at which point they receive zero continuation value and are replaced by new entrants. The stationary age distribution is therefore uniform on $[0, A]$ with density $1/A$.

Preferences. Workers and firms are risk neutral and discount the final good at rate ρ. The final-good market is competitive, and the final good is the numeraire. An unemployed worker of age a with human capital h receives the consumption-equivalent flow value $e^{b(a)+h}$.

Technology. Firms differ along two permanent dimensions: productivity z and learning environment x. Let $\Gamma(z,x)$ denote the distribution of firm types and $\gamma(z,x)$ its density. A firm of type (z,x) that employs a worker with human capital h produces

$$y = e^{z+h}.$$

While employed at a firm with learning environment x, a worker of age a accumulates human capital according to

$$dh = \psi(a) x \, dt.$$

Human capital does not depreciate and remains fixed during unemployment.

Search. The labor market is frictional. Unemployed workers search randomly with age-specific intensity $\phi^u(a)$, normalized so that $A^{-1} \int_0^A \phi^u(a) da = 1$. Employed workers also search randomly for better jobs, with efficiency $\phi^e \geq 0$ relative to unemployed workers of the same age. Their voluntary job-finding rate is therefore $\phi^e \phi^u(a) p$, where p is the equilibrium job-finding rate per unit of search intensity.

Employed workers are additionally hit by a job-termination shock at rate $\lambda(a) = \delta(a) + \phi^f \phi^u(a) p$. This shock has two components. At rate $\delta(a)$, the worker separates to unemployment. At rate $\phi^f \phi^u(a) p$, the worker receives a forced outside offer that must be accepted whenever it is preferable to unemployment; ϕ^f governs the relative efficiency of these forced offers. The case $\phi^f = 0$ can generate job-to-job wage cuts, for example when a worker moves to a better learning environment at a less productive firm, but this channel alone is too weak to match the large share of job-to-job moves with wage cuts in the data. Allowing $\phi^f > 0$ gives the model an additional source of such transitions. Workers may also quit to unemployment whenever doing so is optimal. In the continuous-age formulation, this quit option is imposed as an obstacle condition rather than through a discrete aging transition.

Let U denote the aggregate search effort supplied by unemployed workers:

$$U = \int_0^A \int \phi^u(a) u(a,h) \, dh \, da, \tag{4}$$

where $u(a,h)$ is the density of unemployed workers of age a with human capital h. Let $m(a,h)$ be the corresponding density of all workers. Aggregate search intensity is

$$\begin{aligned} S &= \int_0^A \int \phi^u(a) u(a,h) + (\phi^f + \phi^e) \phi^u(a) \big(m(a,h) - u(a,h)\big) \, dh \, da, & (5) \\ &= U + (\phi^f + \phi^e)(1 - U), & (6) \end{aligned}$$

where the second line uses the normalization of age-specific search intensities and the stationary age distribution.

A firm of type (z,x) posts $v(z,x)$ vacancies at cost $cv(z,x)^{1+\eta}/(1+\eta)$. Total vacancies are

$$V = M \int v(z,x) \, d\Gamma(z,x).$$

The job-finding rate for workers and the worker-finding rate for vacancies are generated by the matching function:

$$
\begin{aligned}
p &= \chi\left(\frac{V}{S}\right)^{\alpha}, \\
q &= \chi\left(\frac{V}{S}\right)^{\alpha-1}.
\end{aligned}
$$

Conditional on meeting a firm, a worker draws a type from the vacancy-weighted distribution

$$
f(z,x) = \frac{Mv(z,x)\gamma(z,x)}{V}.
$$

Let F denote the distribution induced by f.

Wages are determined through the bargaining protocol of Cahuc, Postel-Vinay and Robin (2006), with worker bargaining weight β. A worker who receives a forced outside offer bargains against unemployment as the fallback option. If, at the current wage, the worker would prefer to quit to unemployment, the incumbent firm makes a take-it-or-leave-it offer.[2]

3.2 Equilibrium characterization

Because human capital enters all flow payoffs multiplicatively through e^h and the law of motion $dh = \psi(a)x\,dt$ is independent of the level of h, unemployment and match values are homogeneous of degree one in e^h. I therefore write the value of unemployment as $e^h W(a)$ and the joint value of a match as $e^h J(a,z,x)$.

Value functions. The value of unemployment satisfies the HJB equation

$$
\rho W(a) = e^{b(a)} + \frac{\partial W(a)}{\partial a} + \phi^u(a)p\beta \int_{(z,x)\in\mathcal{A}(a)} \left(J(a,z,x) - W(a)\right) dF(z,x), \tag{7}
$$

with terminal condition $W(A) = 0$. Unemployed workers receive flow value $e^{b(a)}$ in the normalized economy and meet potential employers at rate $\phi^u(a)p$. They accept offers that generate positive surplus, where the acceptance set is

$$
\mathcal{A}(a) = \{(z,x) : J(a,z,x) > W(a)\}.
$$

[2] Because pay is a piece rate of output that never exceeds one, the firm always earns nonnegative flow profits at the current wage. It therefore has no credible threat to lay off the worker.

For accepted matches, the joint value solves

$$\begin{aligned}\big(\rho+\lambda(a)-\psi(a)x\big)J(a,z,x) &= e^z+\frac{\partial J(a,z,x)}{\partial a}+\lambda(a)W(a)\\ &+ \phi^f\phi^u(a)p\beta\int_{(\tilde z,\tilde x)\in\mathcal{A}(a)}\big(J(a,\tilde z,\tilde x)-W(a)\big)\,dF(\tilde z,\tilde x)\\ &+ \phi^e\phi^u(a)p\beta\int_{(\tilde z,\tilde x)\in\mathcal{M}(a,z,x)}\big(J(a,\tilde z,\tilde x)-J(a,z,x)\big)\,dF(\tilde z,\tilde x), \qquad (8)\end{aligned}$$

with terminal condition $J(A,z,x)=0$ for all (z,x). This equation is imposed on the continuation region in which the worker prefers the current match to unemployment, $J(a,z,x)>W(a)$. If the match value falls below $W(a)$, the worker quits to unemployment. In the normalized economy, a match produces e^z. At rate $\lambda(a)$, the match is terminated either by a separation to unemployment or by a forced outside offer. In the latter event, the worker accepts any offer that is preferable to unemployment and receives share β of the surplus over unemployment. At rate $\phi^e\phi^u(a)p$, the worker receives a voluntary outside offer, accepts it only if it improves on the current match, and receives share β of the surplus from moving. Offers below the current match value affect only the wage renegotiation problem and therefore do not enter the joint value. The mobility set is

$$\mathcal{M}(a,z,x) = \{(\tilde z,\tilde x):J(a,\tilde z,\tilde x)>J(a,z,x)\}.$$

I choose the flow value of leisure $b(a)$ so that, for each age, workers are indifferent between unemployment and employment at the reference firm $(\underline{z},\underline{x})$. Evaluating (8) at this reference match gives

$$\big(\rho-\psi(a)\underline{x}\big)W(a) = e^{\underline{z}}+\frac{\partial W(a)}{\partial a}+(\phi^f+\phi^e)\phi^u(a)p\beta\int_{(z,x)\in\mathcal{A}(a)}\big(J(a,z,x)-W(a)\big)\,dF(z,x).$$

Combining this condition with (7) yields

$$e^{\underline{z}} = e^{b(a)}-\psi(a)\underline{x}W(a)+p\phi^u(a)\beta(1-\phi^f-\phi^e)\int_{(z,x)\in\mathcal{A}(a)}\big(J(a,z,x)-W(a)\big)\,dF(z,x). \qquad (9)$$

Equation (9) illustrates two forces behind the reservation threshold. The standard search option value raises the productivity needed to make an offer acceptable, provided $\phi^f+\phi^e<1$. Learning on the job lowers the threshold when $\psi(a)\underline{x}>0$, because even a low-productivity job can increase future human capital.

Substituting (9) into (7) gives the unemployment value in terms of the reference match:

$$\big(\rho-\psi(a)\underline{x}\big)W(a) = e^{\underline{z}}+\frac{\partial W(a)}{\partial a}+p\beta\phi^u(a)(\phi^f+\phi^e)\int_{(z,x)\in\mathcal{A}(a)}\big(J(a,z,x)-W(a)\big)\,dF(z,x) \qquad (10)$$

In estimation, I impose the indifference condition at $(\underline{z}, \underline{x})$ for each age. Equations (8) and (10), together with the terminal conditions at retirement, then determine $J(a,z,x)$ and $W(a)$ without directly solving for $e^{b(a)}$; the latter is recovered ex post from (9).

Laws of motion. Let the total density of employed workers of age a with human capital h be

$$e(a,h) \;=\; \int g(a,z,x,h)\,dz\,dx,$$

where $g(a,z,x,h) = 0$ for $(z,x) \notin \mathcal{A}(a)$. Let $\mathcal{Q}(a,h)$ denote the deterministic quit flow into unemployment generated by matches that cease to be acceptable as workers age. Equivalently,

$$\mathcal{Q}(a,h) \;=\; \lim_{\Delta a \downarrow 0} \frac{1}{\Delta a} \int_{\{(z,x)\in\mathcal{A}(a):\,(z,x)\notin\mathcal{A}(a+\Delta a)\}} g(a,z,x,h)\,dz\,dx,$$

whenever this limit exists, with the corresponding object interpreted as a boundary flux in weak form otherwise.

The density of unemployed workers satisfies the stationary Kolmogorov equation

$$0 \;=\; -\frac{\partial u(a,h)}{\partial a} + \Big(\delta(a) + \phi^f \phi^u(a) p F(\mathcal{R}(a))\Big)\, e(a,h) - p\phi^u(a) F(\mathcal{A}(a)) u(a,h) + \mathcal{Q}(a,h), \tag{11}$$

where $\mathcal{R}(a) = \{(z,x) : (z,x) \notin \mathcal{A}(a)\}$ is the rejection region. Entry is imposed as the boundary condition

$$\begin{aligned} u(0,h) &= A^{-1}\Psi_0(h), \\ g(0,z,x,h) &= 0, \end{aligned}$$

where Ψ_0 is degenerate at the normalized entrant human-capital level $h = 0$. Workers retire at age A, so no distributional boundary condition is imposed beyond the terminal outflow at A.

For accepted jobs, $(z,x) \in \mathcal{A}(a)$, the density of employed workers satisfies

$$\begin{aligned} 0 \;=\;& -\frac{\partial g(a,z,x,h)}{\partial a} - \frac{\partial}{\partial h}\left(\psi(a) x g(a,z,x,h)\right) \\ &- g(a,z,x,h)\left(\lambda(a) + \phi^e \phi^u(a) p F(\mathcal{M}(a,z,x))\right) \\ &+ f(z,x) p \phi^u(a) \left[u(a,h) + \phi^f e(a,h) + \phi^e \int \mathbf{1}\{(z,x)\in\mathcal{M}(a,\tilde{z},\tilde{x})\} g(a,\tilde{z},\tilde{x},h)\,d\tilde{z}\,d\tilde{x} \right]. \end{aligned} \tag{12}$$

This equation is understood on the age-dependent continuation domain $\mathcal{A}(a)$. The flux of employed mass that reaches the moving boundary of this domain is transferred to unemployment through $\mathcal{Q}(a,h)$.

Vacancy creation. A firm that posts v vacancies pays $cv^{1+\eta}/(1+\eta)$. A posted vacancy meets a worker at rate q and yields expected return $R(z,x)$ upon contact. The firm's first-order condition is

$$cv(z,x)^{\eta} \quad = \quad qR(z,x). \tag{13}$$

The return $R(z,x)$ has two components. If the vacancy contacts an unemployed worker or a worker subject to a forced outside offer, the firm receives share $1-\beta$ of the surplus over unemployment. If it contacts an employed worker through voluntary on-the-job search, the firm receives share $1-\beta$ of the surplus over the worker's current match. Therefore,

$$\begin{aligned} R(z,x) \quad = \quad & \frac{1-\beta}{S} \int_0^A \phi^u(a) \int e^h \Big[\mathbf{1}\{(z,x) \in \mathcal{A}(a)\} \big(u(a,h) + \phi^f e(a,h)\big) \\ & \times \big(J(a,z,x) - W(a)\big) + \phi^e \int \mathbf{1}\{(z,x) \in \mathcal{M}(a,\tilde{z},\tilde{x})\} \big(J(a,z,x) - J(a,\tilde{z},\tilde{x})\big) \\ & \times g(a,\tilde{z},\tilde{x},h)\, d\tilde{z}\, d\tilde{x} \Big] dh\, da. \end{aligned}$$

Here $R(z,x)$ is the expected return to a contact for a firm of type (z,x), integrating over the distribution of workers by continuous age, employment status, current job, and human capital.

Using (13), the matching function, and the definition of f, the equilibrium job-finding rate and offer distribution satisfy

$$p \quad = \quad \theta \left(\frac{1}{S(p,f)} \right)^{\frac{\alpha\eta}{1+\eta-\alpha}} \left(\int R(z,x;p,f)^{1/\eta}\, d\Gamma(z,x) \right)^{\frac{\alpha\eta}{1+\eta-\alpha}}, \tag{14}$$

$$f(z,x) \quad = \quad \frac{R(z,x;p,f)^{1/\eta} \gamma(z,x)}{\int R(\tilde{z},\tilde{x};p,f)^{1/\eta}\, d\Gamma(\tilde{z},\tilde{x})}. \tag{15}$$

The composite parameter is

$$\theta \quad = \quad \left(\frac{\chi^{1+\eta} M^{\alpha\eta}}{c^{\alpha}} \right)^{\frac{1}{1+\eta-\alpha}}.$$

The return $R(z,x;p,f)$ depends on p and f because these objects affect match values, unemployment values, and the stationary allocation of workers across unemployment and firm types.

Equilibrium. A stationary equilibrium consists of value functions $W(a)$ and $J(a,z,x)$ for $a \in [0,A]$, acceptance and mobility sets $\mathcal{A}(a)$ and $\mathcal{M}(a,z,x)$, a job-finding rate p, aggregate search intensity S, an offer distribution $f(z,x)$, and worker densities $u(a,h)$ and $g(a,z,x,h)$ such that:

1. Given p and f, the value functions and policies solve (8) and (10), with terminal conditions

at retirement and the quit obstacle imposed continuously in age.

2. Given match values, unemployment values, and the stationary distribution of workers, p and f satisfy firm profit maximization, (14)–(15).

3. Given policies, p, and f, the worker densities solve (11)–(12), with entry boundary conditions at age 0, retirement at age A, and aggregate search intensity given by (6).

The counterfactual exercises vary the vacancy-posting cost c to study how labor-market fluidity affects life-cycle outcomes. An increase in c lowers firms' incentive to post vacancies and reduces the equilibrium job-finding rate p. Before turning to the quantitative exercises, it is useful to isolate the partial-equilibrium effect of a lower p on the jobs accepted by young unemployed workers.

Fix age a and let $\overline{z}(a,x)$ be the minimum productivity required for a job with learning environment x to be acceptable. A worker is indifferent between unemployment and the job $(\overline{z}(a,x),x)$ when

$$\begin{aligned}(\rho-\psi(a)x)W(a) &= e^{\overline{z}(a,x)}+\frac{\partial J\big(a,\overline{z}(a,x),x\big)}{\partial a}\\ &+ \beta(\phi^f+\phi^e)\phi^u(a)p\int_{(\tilde{z},\tilde{x})\in\mathcal{A}(a)}\big(J(a,\tilde{z},\tilde{x})-W(a)\big)\,dF(\tilde{z},\tilde{x}).\end{aligned}$$

Subtracting (10) yields

$$e^{\overline{z}(a,x)}-e^{\underline{z}} = -\psi(a)W(a)(x-\underline{x})+\frac{\partial W(a)}{\partial a}-\frac{\partial J\big(a,\overline{z}(a,x),x\big)}{\partial a}.$$

The indifference curve is

$$\overline{z}(a,x) = \log(e^{\underline{z}}-\psi(a)W(a)(x-\underline{x})+\mathcal{C}(a,x)), \qquad \mathcal{C}(a,x)\equiv\frac{\partial W(a)}{\partial a}-\frac{\partial J\big(a,\overline{z}(a,x),x\big)}{\partial a}.$$

Differentiating this with respect to x using the fact that $\frac{\partial^2 J\left(a,\overline{z}(a,x),x\right)}{\partial a\partial x}=0$ since people are indifferent on the boundary,

$$\frac{\partial\overline{z}(a,x)}{\partial x} = -\frac{\psi(a)W(a)}{e^{\underline{z}}-\psi(a)W(a)(x-\underline{x})} = -\frac{\psi(a)W(a)}{e^{\overline{z}(a,x)}} < 0.$$

Thus the indifference curve is downward sloping. The derivative with respect to the job-finding rate is

$$\frac{\partial}{\partial p}\left(\frac{\partial\overline{z}(a,x)}{\partial x}\right) = -\frac{\psi(a)e^{\underline{z}}}{(e^{\underline{z}}-\psi(a)W(a)(x-\underline{x}))^2}\frac{\partial W(a)}{\partial p},$$

The value of unemployment $W(a)$ is increasing in p. Workers can never be made worse off from receiving more offers, since they can always discard additional offers. Using a version of this logic, it is easy to show that they are strictly better off (as long as some of the offers are accepted). Thus, a higher p steepens the indifference curve: for a given deterioration in the learning environment, workers require larger productivity compensation in a more fluid economy. The intuition is twofold. First, young workers in fluid labor markets expect to be less mismatched later in life, which raises the value of human capital and makes high-learning jobs more attractive today. Second, fluid markets make it easier to move later from high-learning firms to high-productivity firms as workers' comparative advantage shifts toward exploiting accumulated human capital. Anticipating this option, young workers are more willing to start at high-learning firms.

3.3 Wage determination

After solving for the equilibrium allocation, wages are determined from the bargaining protocol. Pay is delivered as a piece rate w of output, so take-home pay is we^{z+h}. By homogeneity, the worker's value under piece rate w can be written as

$$\widetilde{V}(a,z,x,h,w) \quad = \quad e^h V(a,z,x,w).$$

For $a \in [0, A]$, this value satisfies, on the continuation region in which the worker does not quit,

$$\begin{aligned}
&\big(\rho + \lambda(a) - \psi(a)x\big)V(a,z,x,w) \\
&= \quad we^z + \frac{\partial V(a,z,x,w)}{\partial a} + \lambda(a)W(a) \\
&+ \quad \phi^f \phi^u(a) p\beta \int_{(\tilde{z},\tilde{x})\in\mathcal{A}(a)} \big(J(a,\tilde{z},\tilde{x}) - W(a)\big)\, dF(\tilde{z},\tilde{x}) \\
&+ \quad \phi^e \phi^u(a) p \int_{(\tilde{z},\tilde{x})\in\mathcal{N}(a,z,x,w)} \big(\beta J(a,z,x) + (1-\beta)J(a,\tilde{z},\tilde{x}) - V(a,z,x,w)\big)\, dF(\tilde{z},\tilde{x}) \\
&+ \quad \phi^e \phi^u(a) p \int_{(\tilde{z},\tilde{x})\in\mathcal{M}(a,z,x)} \big((1-\beta)J(a,z,x) + \beta J(a,\tilde{z},\tilde{x}) - V(a,z,x,w)\big)\, dF(\tilde{z},\tilde{x}),
\end{aligned}$$

with terminal condition $V(A,z,x,w) = 0$. If the continuation value falls below $W(a)$, the worker quits to unemployment. The renegotiation set is

$$\begin{aligned}
\mathcal{N}(a,z,x,w) \quad = \quad &\Big\{(\tilde{z},\tilde{x}) : (\tilde{z},\tilde{x}) \notin \mathcal{M}(a,z,x) \\
&\quad \text{and } \beta J(a,z,x) + (1-\beta)J(a,\tilde{z},\tilde{x}) > V(a,z,x,w)\Big\}.
\end{aligned}$$

The equilibrium allocation requires only the joint match values and unemployment values.

Given $J(a,z,x)$, $W(a)$, $\mathcal{A}(a)$, $\mathcal{M}(a,z,x)$, p, and F, define

$$
\begin{aligned}
T(a,z,x) &= \rho + \lambda(a) + \phi^e \phi^u(a) p F(\mathcal{M}(a,z,x)) - \psi(a) x, \\
B(a,z,x) &= \lambda(a) W(a) + \phi^f \phi^u(a) p \beta \int_{(\tilde{z},\tilde{x}) \in \mathcal{A}(a)} \big(J(a,\tilde{z},\tilde{x}) - W(a) \big) \, dF(\tilde{z},\tilde{x}) \\
&+ \phi^e \phi^u(a) p \int_{(\tilde{z},\tilde{x}) \in \mathcal{M}(a,z,x)} \left((1-\beta) J(a,z,x) + \beta J(a,\tilde{z},\tilde{x}) \right) \, dF(\tilde{z},\tilde{x}).
\end{aligned}
$$

Then the worker value can be written compactly as the continuous-age equation

$$
\begin{aligned}
& \left(T(a,z,x) + \phi^e \phi^u(a) p F(\mathcal{N}(a,z,x,w)) \right) V(a,z,x,w) \\
&= w e^z + B(a,z,x) + \frac{\partial V(a,z,x,w)}{\partial a} \\
&+ \phi^e \phi^u(a) p \int_{(\tilde{z},\tilde{x}) \in \mathcal{N}(a,z,x,w)} \left(\beta J(a,z,x) + (1-\beta) J(a,\tilde{z},\tilde{x}) \right) \, dF(\tilde{z},\tilde{x}).
\end{aligned}
$$

The remaining objects, $\mathcal{N}(a,z,x,w)$ and $V(a,z,x,w)$, are recovered numerically after the equilibrium allocation has been solved.

The piece rate paid to a worker of age a hired from unemployment, or to a worker who receives a forced outside offer, is defined by

$$
V\big(a,z,x,w(a,z,x,u)\big) = W(a) + \beta \big(J(a,z,x) - W(a) \big).
$$

If the incumbent worker would otherwise quit to unemployment, the piece rate satisfies

$$
V\big(a,z,x,w(a,z,x,q)\big) = W(a).
$$

Finally, if a worker employed at (z,x) receives a less attractive outside offer $(\tilde{z},\tilde{x})$, the renegotiated piece rate is defined by

$$
V\big(a,z,x,w(a,z,x,\tilde{z},\tilde{x})\big) = \beta J(a,z,x) + (1-\beta) J(a,\tilde{z},\tilde{x}).
$$

4 Calibration

This section calibrates the model to a hypothetical average-fluidity benchmark economy. The calibration proceeds in two steps. I first set or normalize parameters that are standard, weakly identified by the available data, or directly pinned down by individual moments. I then estimate the remaining parameters by matching moments that discipline job mobility, wage growth, firm size, and worker sorting.

4.1 Calibration Strategy

I proxy the continuous time model for age with four 10-year bins: 20–30, 30–40, 40–50 and 50–60.

Externally set and normalized parameters. Panel A of Table 4 reports the parameters set outside the internal calibration. These include a monthly discount rate $\rho = 0.003$, corresponding to an annual real interest rate of approximately four percent, worker bargaining power $\beta = 0.5$, and matching-function elasticity $\alpha = 0.5$. To simplify the estimation, I also fix three objects that map closely into individual moments.

First, the worker-level data identify employment shares in broad firm-size bins but do not identify average firm size. Comparable cross-country data on average firm size are also limited. Moreover, in the quantitative exercises below, M mainly scales the economy and has little effect on the relative changes of interest. I hence set the mass of firms per worker to $M = 0.05$, which given an employment rate of about 75 percent in the average-fluidity country for workers aged 20–59, implies an average firm size of 15 employees.

Second, I set the exogenous separation hazards $\delta(a)$ to match employment-to-unemployment (EU) rates by age. The model also generates a small number of endogenous separations when workers age into the separation region, but these separations are negligible in the calibrated economy.

Third, I set the unemployment job-arrival rates $\phi^u(a)p$ to match unemployment-to-employment (UE) rates by age. Age variation in workers' reservation thresholds implies that not all posted jobs are accepted in equilibrium, but the share of rejected offers is modest in the calibrated economy. Moreover, as for the number of firms, $\phi^u(a)p$ mostly affect the scale of the employment pool, with second-order effects on the relative changes I consider below. Let UE_a denote the UE rate for age group a. Under the normalization $A^{-1}\int \phi^u(a)da = 1$, the average job-finding rate is $p = A^{-1}\sum_a UE(a)$, and the age-specific search efficiencies are $\phi^u(a) = UE(a)/p$. Given p and $[\phi^u(a)]_{a=1}^{A}$, I solve for $J(z,x,a)$, $W(a)$, $\mathcal{A}(a)$, $\mathcal{M}(z,x,a)$, $u(h,a)$, $g(z,x,h,a)$, and $f(z,x)$. Ex post, I recover the scalar vacancy-cost parameter c from (14) so that p is consistent with equilibrium. Without vacancy data, the matching efficiency χ and the vacancy-cost scale c are not separately identified; I therefore normalize $\chi = 1$.[3]

Figure 3 compares the model-implied EU and UE rates with their empirical counterparts. Because endogenous separations are rare and few offers are rejected, the model matches pretty well these transition rates.

Internally calibrated parameters. Panel B of Table 4 reports the eight internally calibrated parameters. The calibration minimizes a weighted sum of squared deviations between model and

[3] Because χ is a free parameter, the implied value of c, even if expressed relative to output, has no independent economic interpretation.

Figure 3: EU and UE Rates

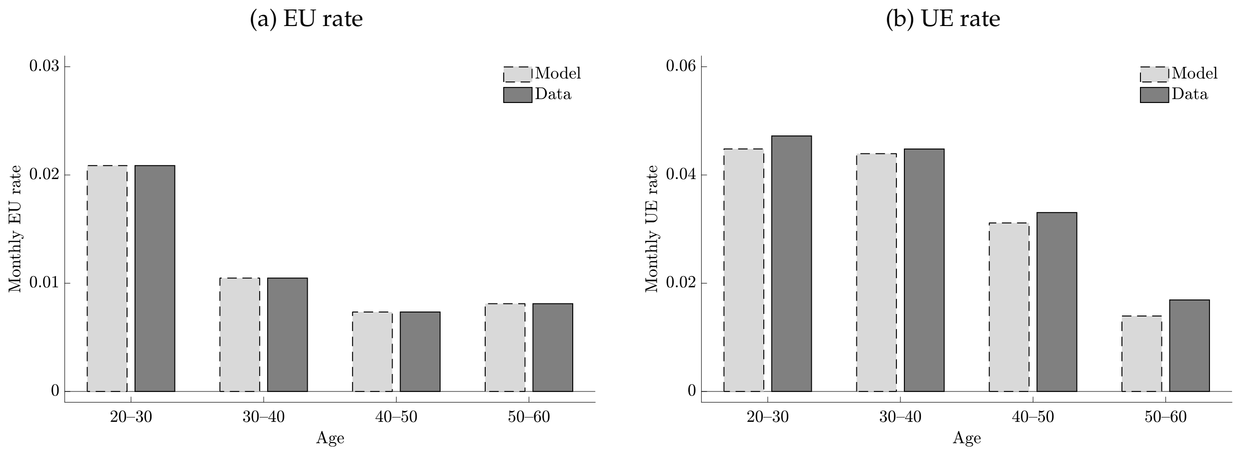

data moments. Since the targeted moments are either rates or log differences, the objective is comparable across moments. Although the parameters are estimated jointly, the discussion below highlights the empirical variation that most directly disciplines each parameter.

On-the-job search efficiency for higher-paying jobs, ϕ^e, is disciplined by the job-to-job (JJ) mobility rate. To mirror the empirical measure, the model JJ rate is the share of currently employed workers who started a new job in the previous 12 months without experiencing unemployment during that period. I construct the measure by recording employment status at monthly intervals and classifying workers as employed or unemployed at each observation. Because the age distribution in the model differs from that in the data, I compute the rate within age bins 20–29, 30–39, 40–49, and 50–59 in both the model and the data, and then aggregate by weighting each age group equally. I target both the aggregate JJ rate and the age-specific rates. Together, these moments receive total weight five because job-to-job mobility is central to the analysis.

Forced-reallocation search efficiency, ϕ^f, is disciplined by the wage growth of JJ movers. To construct the corresponding model moment, I condition on workers employed in month 17, paralleling the timing of the survey, which is typically fielded in May. A JJ mover is a worker who changes employer at some point between months 6 and 17 while remaining employed in every month over that interval. I compute total income and total months worked separately in months 1–12 and 13–24, and define wages as income per month worked. In both the data and the model, I drop workers without a valid wage in either year t or year $t+1$. I compute the average wage in year t and year $t-1$ (in levels) at the age group level, then compute the log difference in wages, and aggregate weighing all age groups equally.[4] I target both the aggregate and the age-specific moment, weighing equally so that they get total weight one.

[4] Wages in the data are deflated to have mean zero in each country-year to take out any aggregate time trend. Aggregate wages by construction do not grow in the model.

I assume that age-specific learning ability takes the form

$$\psi(a) = \mu \left(\frac{A-a}{A-1} \right)^{\psi}.$$

Thus, $\psi(1) = \mu$ and $\psi(A) = 0$, with μ governing average learning ability and ψ governing the speed at which learning ability declines with age. Because mean learning ability is not separately identified from the mean learning environment, I normalize $\mu = 1$. The curvature parameter ψ is disciplined by two sets of moments. First, I target annual wage growth at training firms by age group. A high value of ψ implies that the human-capital gain from working at training firms declines sharply with age, and therefore that the wage-growth premium from such firms also declines sharply with age. Second, I target wages at ages 30–39, 40–49, and 50–59 relative to ages 20–29. A high value of ψ generates rapid human-capital accumulation early in the life cycle and slower accumulation later, producing a more concave wage profile.

The underlying distribution of firm productivity z and learning environment x is jointly normal. Since the overall productivity scale is indeterminate, I normalize mean productivity to $\mu_z = 0$. I estimate the dispersion of productivity, σ_z; the mean learning environment, μ_x; the dispersion of learning environments, σ_x; and the correlation between productivity and the learning environment, ϱ_{zx}.

The dispersion of productivity, σ_z, is disciplined by log average wages at firms with at least 50 employees relative to firms with fewer than 50 employees. To mimic the data, I record firm size in month 17 and define the wage as total income in months 1–12 divided by total months worked in months 1–12. I compute average wages by firm-size category and age group, take logs, difference large relative to small firms, and aggregate across age groups using equal weights, applying the same procedure in the data and the model. This moment receives total weight one.

For any candidate parameter vector, I choose the age-specific flow value of leisure b_a so that workers of each age group are indifferent between unemployment and employment in a $(z, x) = (0, 0)$ match: $J(0, 0, a) = W(a)$ for all a. Because the mean of the distribution of firms with $x > 0$ depends on both the untruncated mean and σ_x, I let μ_x denote the mean of the distribution conditional on $x > 0$. Equivalently, the untruncated mean m_x solves

$$\mu_x = m_x + \sigma_x \lambda \left(\frac{m_x}{\sigma_x} \right), \quad \lambda(t) \equiv \frac{\varphi(t)}{\Phi(t)},$$

where φ and Φ are the standard normal density and distribution function. This parametrization separates the mean accepted learning environment from dispersion in x and improves identification of σ_x. I discipline μ_x using wages at age 50–59 relative to age 20–29. This moment receives a weight of one.

The dispersion of learning environments, σ_x, is disciplined by wage growth at training firms relative to non-training firms. I define training firms as the employment-weighted top 41 percent of firms ranked by x, matching the employment share at training firms in the data. As above, the employer is recorded in month 17 in the model. I compute average wages between months 1–12 and 13–24, compute the mean (in levels) at the age group level, take the log difference, and aggregate across age groups using equal weights, applying the same procedure in the data and the model. This moment receives total weight one.

The correlation between productivity and learning environment, ϱ_{zx}, is disciplined by the share of training firms within firm-size categories 0–4, 5–19, 20–49, and 50 or more employees, as well as the share of workers at training firms by age. A higher ϱ_{zx} implies a higher training share among large, high-productivity firms. Young workers value more learning environment, leading them to gradually relocate away from training firms. If ϱ_{zx} is high, however, the strong correlation between z and x implies that older workers are found at training firms, even though they only value their high z. As with the other moments, I compute training shares first by firm-size category and age group, and then aggregate to the firm-size level by weighting all age groups equally in both the model and the data. These moments receive total weight one.

The curvature of the vacancy cost, η, governs how costly it is for firms to grow large. Conditional on productivity dispersion σ_z, I set η to match the employment share at firms with at least 50 employees. I target this share separately for workers aged 20–39 and 40–59, with the two moments receiving total weight one.

Conditional on the calibrated parameter vector, I infer the flow value of leisure $b(a)$ at each age so that workers are indifferent between unemployment and a $(z, x) = (0, 0)$ match. Panel C of Table 4 reports the implied values of c, m_x, and $b(a)$.

4.2 Parameter Estimates and Model Fit

Table 4 summarizes the parameter values, and Table 5 reports the fit to the targeted moments. Employed workers receive directed outside offers at 0.20 times the average offer rate of unemployed workers, while forced-reallocation offers arrive at 0.13 times that rate. The estimate of ψ implies that workers aged 30–39 learn at approximately 13 percent of the speed of workers aged 20–29, so learning ability declines rapidly with age.

The estimated standard deviation of firm productivity is 0.20. Although this value is below the empirical dispersion in value added across firms, value-added dispersion is larger in the model because older workers with more human capital tend to work at more productive firms. In the data, value-added dispersion also reflects differences in capital intensity, pricing power, and other firm-level factors that the model abstracts from. Because $\sigma_x \times 100 = 0.10$, a one-standard-deviation improvement in the learning environment raises monthly human-capital growth by 0.10 percentage points. Productivity and the learning environment are modestly positively correlated, with $\varrho_{zx} = 0.05$. Finally, the estimated vacancy-cost curvature, $\eta = 0.80$, implies only moderate convexity in hiring costs.

Table 4: Parameter Values

Parameter	Description	Value
Panel A. Preset or Normalized Parameters		
ρ	Discount rate	0.003
β	Worker bargaining weight	0.500
α	Matching elasticity	0.500
M	Mass of firms	0.050
$\delta(a)$	Separation hazard by age	
20–30		0.021
30–40		0.010
40–50		0.007
50–60		0.008
$\phi^u(a)$	Relative efficiency of unemployed search	
20–30		1.330
30–40		1.262
40–50		0.931
50–60		0.476
p	Job finding rate per unit of search intensity	0.035
χ	Matching efficiency	1.000
μ	Mean learning ability	1.000
Panel B. Internally Calibrated Parameters		
ϕ^e	Relative efficiency of directed on-the-job search efficiency	0.150
ϕ^f	Relative efficiency of undirected on-the-job search efficiency	0.150
ψ	Speed of decline in learning ability with age	2.500
σ_z	Std. dev. productivity	0.250
μ_x	Mean learning environment ($\times$100)	0.170
σ_x	Std. dev. learning environment ($\times$100)	0.200
ϱ_{zx}	Correlation between z and x	0.050
η	Vacancy cost curvature	1.000
Panel C. Implied Parameters		
c	Vacancy cost scale	9733
m_x	Mean of underlying (untruncated) x distribution ($\times$100)	0.028
$e^{b(a)}$	Flow value of leisure	
20–30		0.604
30–40		0.671
40–50		0.771
50–60		0.908

Table 5: Targeted Moments

Parameter	Moment	Data	Model
ϕ^e	JJ mobility rate (aggregate; annual)	0.070	0.070
	Life-cycle JJ mobility rate (annual)	See Figure 4a	
ϕ^f	Wage growth of JJ movers (aggregate; annual)	0.020	0.047
	Life-cycle wage growth of JJ movers (annual)	See Figure 4b	
ψ	Wage growth at training firms by age (annual)	See Figure 4c	
	Life-cycle wage profile	See Figure 5	
σ_z	Wages at large (50+ employees) relative to small ($<$ 50) firms	0.199	0.200
μ_x	Wage growth between ages 20–29 and 50–59	0.362	0.389
σ_x	Wage growth at training firms (aggregate; annual)	0.023	0.015
ϱ_{zx}	Share of employment at training firms by firm size	See Figure 4e	
	Share of workers at training firms by age	See Figure 4f	
η	Share of workers at firms with 50+ employees by age	See Figure 4d	

As shown in Table 5, the model matches reasonably well the aggregate JJ mobility rate, the aggregate wage gain associated with JJ mobility, the wage premium at large firms, life-cycle wage growth, and wage growth at training firms.

Figure 4 reports the fit for the remaining graphical moments. The model reproduces the age profile of JJ mobility and the wage gains from JJ moves. It also captures the faster wage growth of young workers at training firms relative to older workers, as well as the higher incidence of training firms among large firms. The age pattern in training-firm wage growth and the life-cycle wage profile discipline ψ; the firm-size pattern in training incidence disciplines ϱ_{zx}.

4.3 Model Mechanics

Having established the model fit, I use the calibrated model to analyze life-cycle wage and employment dynamics. Figure 5 decomposes life-cycle wage growth into growth in human capital, match productivity, and the piece-rate component of wages. The estimates imply that growth in the piece-rate component is the largest source of life-cycle wage growth, as workers use outside offers to renegotiate wages upward. Human-capital accumulation is the second-largest source, while growth in match productivity contributes the least.

Figure 6 compares the underlying firm distribution with the vacancy-weighted distribution. Recall that the flow value of leisure b_a is set so that workers are indifferent between unemployment and employment in a $(z, x) = (0, 0)$ match. Some low-productivity, low-learning-environment firms in the lower-left tail cannot profitably hire workers and therefore do not post vacancies. The vacancy-weighted distribution is consequently tilted toward firms with higher productivity and stronger learning environments.

Figure 7 plots the distribution of workers over productivity and learning environments by age,

Figure 4: Additional Targeted Moments

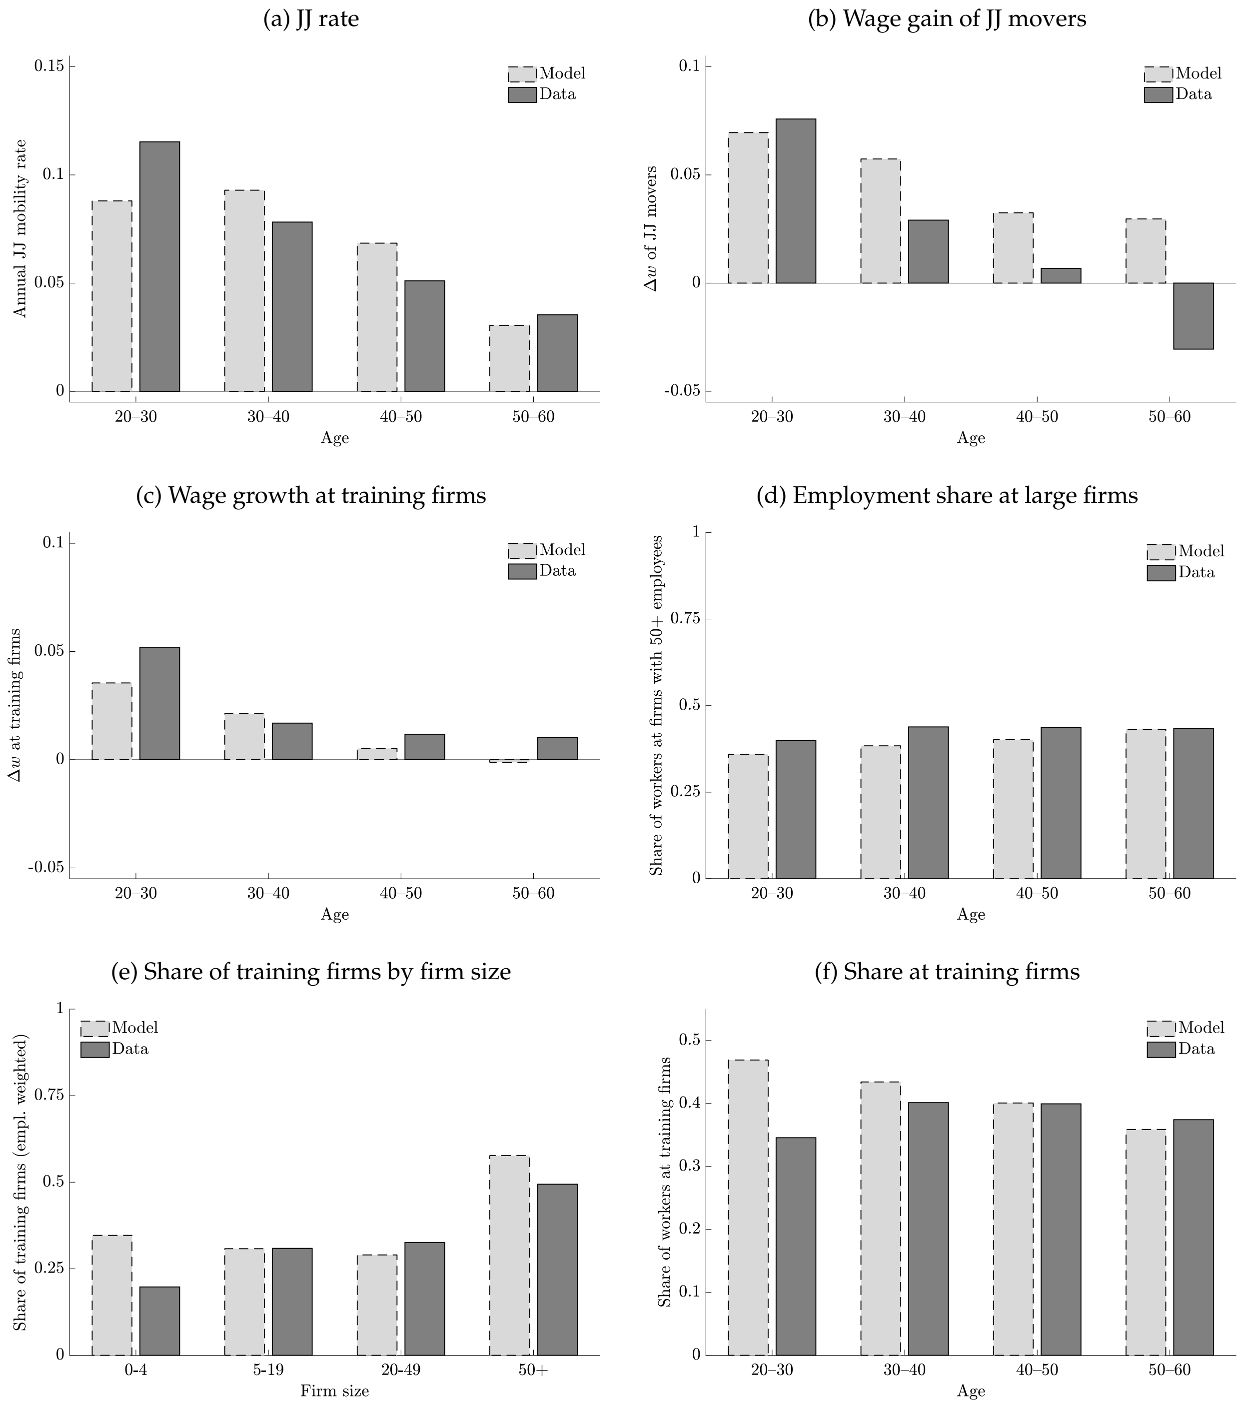

together with age-specific indifference curves. Young workers in panel (a) place greater value on learning environments than older workers in panel (b), who receive lower returns from human-capital accumulation. Young workers initially sort toward firms with stronger learning environments and higher productivity. Later in the life cycle, as the return to learning declines, workers

Figure 5: The Sources of Life-Cycle Wage Growth

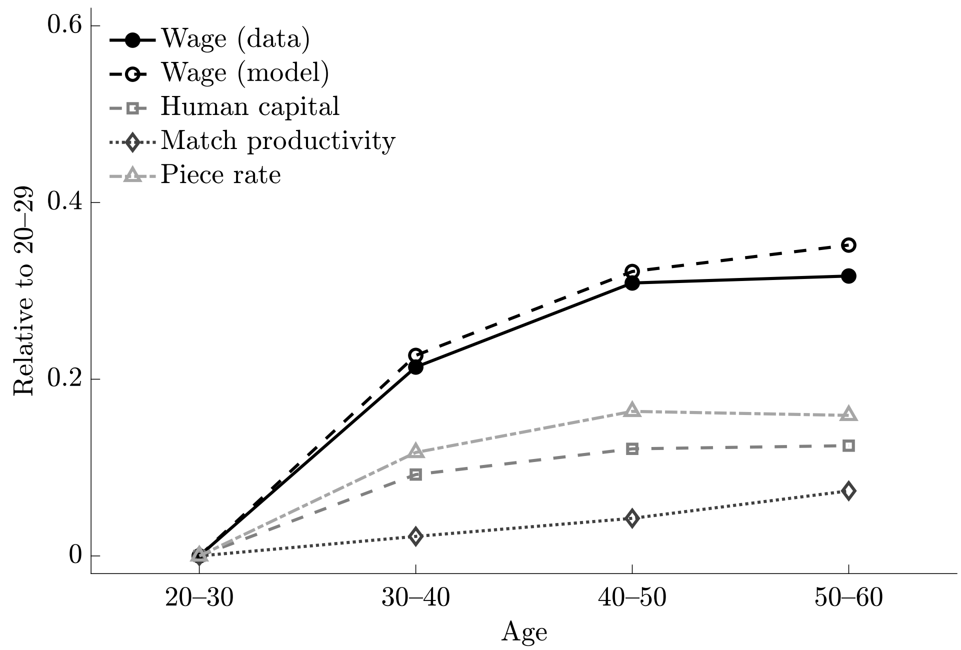

Figure 6: Firm Distribution and Vacancy Posting

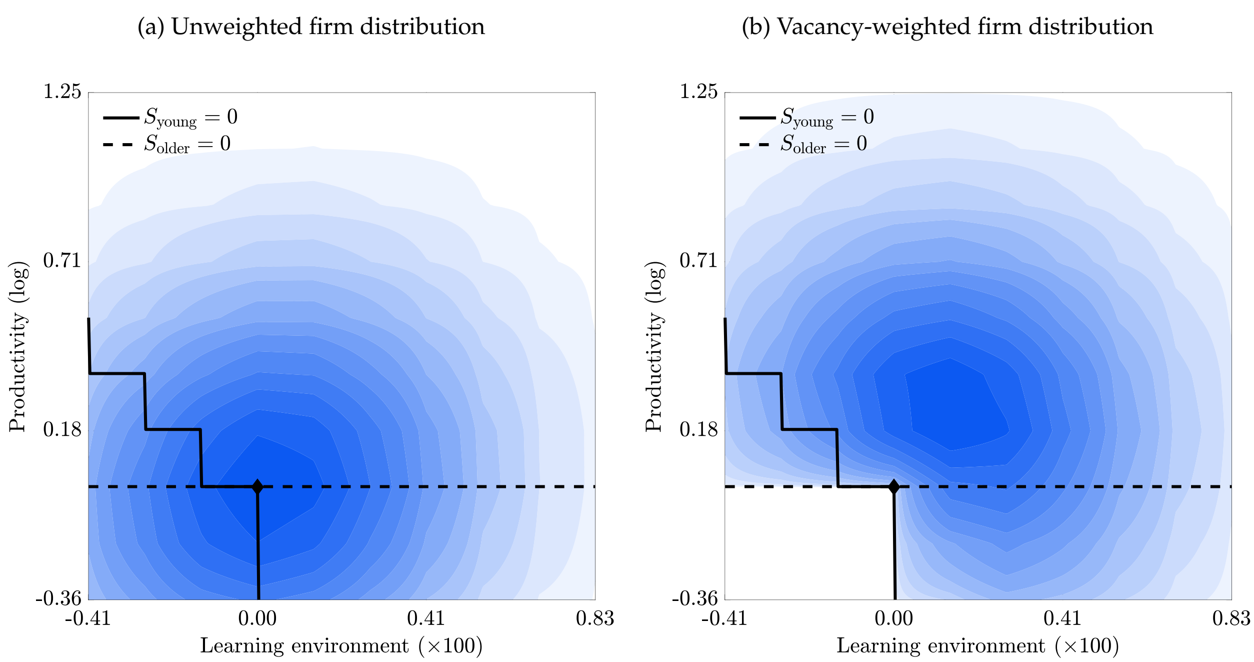

increasingly trade off learning opportunities for higher current productivity and wages.

Figure 7: Worker Sorting and Indifference Curves

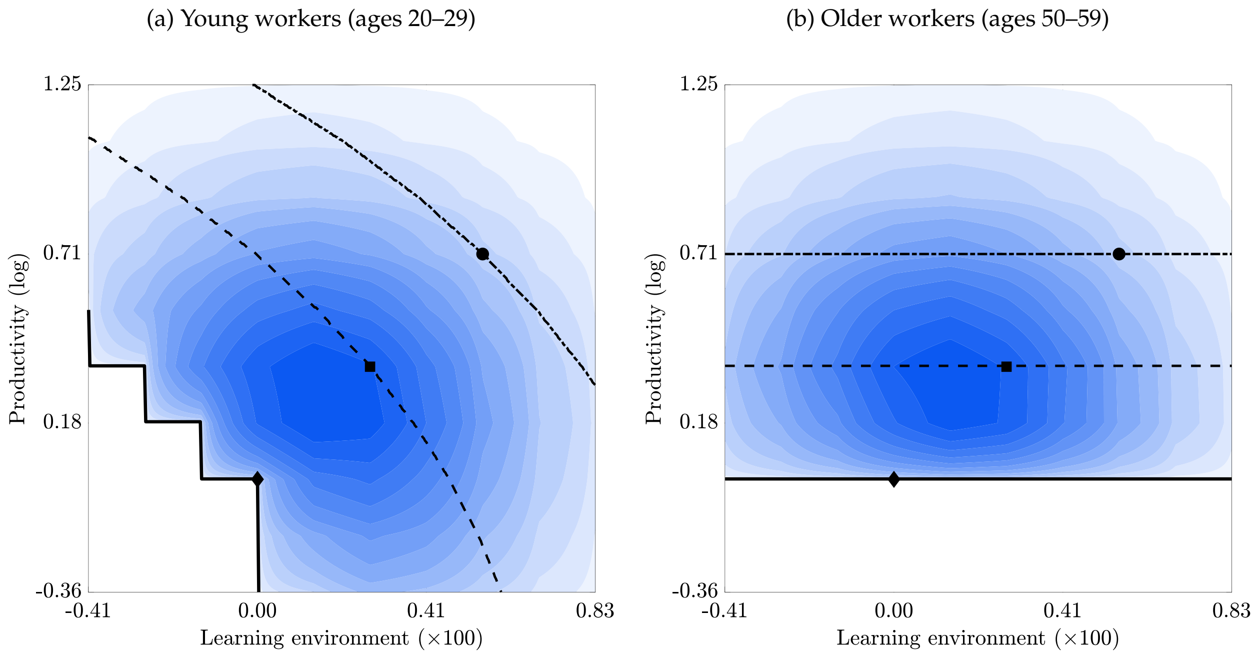

5 Results

I now turn to the my main quantitative exercise: understanding the role of labor market frictions in driving the observed cross-country differences in career outcomes.

5.1 Cross-Country Calibration

I recalibrate two parameters—p and ϕ^e—to match the level of fluidity and UE rate in a hypothetical low fluidity, high fluidity and super fluid country, respectively. The low fluidity country has log fluidity one standard deviation below the average, the high fluidity country has log fluidity two standard deviations above the average, and the super fluid country two standard deviations above the average. I obtain the corresponding UE rate from a linear projection on log fluidity. I set p to the UE rate and find ϕ^e to match associated labor market fluidity. Finally, I recover the implied cost of vacancy creation c that rationalizes p as an equilibrium outcome. All other parameters are held fixed in my analysis below.

Table 8 summarizes the calibrated cross-country parameters. The job finding rate p goes from three percent per month in the low fluidity country to over five percent per month in the super fluid country, which is required to match the UE rate shown in Figure 8. I assume that the total separation rate—the sum of separations to unemployment and reallocation shocks—is constant across countries. Consequently, the higher job finding rate p in high fluidity countries implies a (modestly) lower separation rate to unemployment in these countries. Although I cannot verify this assumption, Figure 8 shows that the calibrated model is broadly consistent with the EU rate (panel (b)). In any case, my results are little affected if I instead recalibrate δ to match the EU rate. The model is also consistent with the lack of a correlation between the wage gain of JJ movers and fluidity. The higher UE rate in more fluid labor markets is associated with a higher employment rate, consistent with the data. Directed search efficiency on-the-job ϕ^e is essentially zero in the low fluidity country, and over 0.4 in the super fluid country. One microfoundation for this pattern is that employed workers in more fluid labor markets respond to the higher job finding rate by searching harder for alternative jobs. The implied variation in the cost of creating jobs is large.

Table 6: Cross-Country Differences in Labor Market Structure

	Low	Middle	High	Super
Fluidity	0.0451	0.0690	0.1055	0.1620
p	0.0295	0.0355	0.0473	0.0561
$\lambda \equiv \delta + \phi^f p$ (mean)	0.0178	0.0182	0.0182	0.0182
δ (mean)	0.0128	0.0117	0.0095	0.0079
φ^f	0.1685	0.1828	0.1828	0.1828
φ^e	0.0000	0.0731	0.1576	0.4179
$\log c$ (rel)	0.5777	0.0000	-0.8483	-1.9471

Figure 8: Cross-Country UE and EU Rates

(a) UE Rate

(b) EU Rate

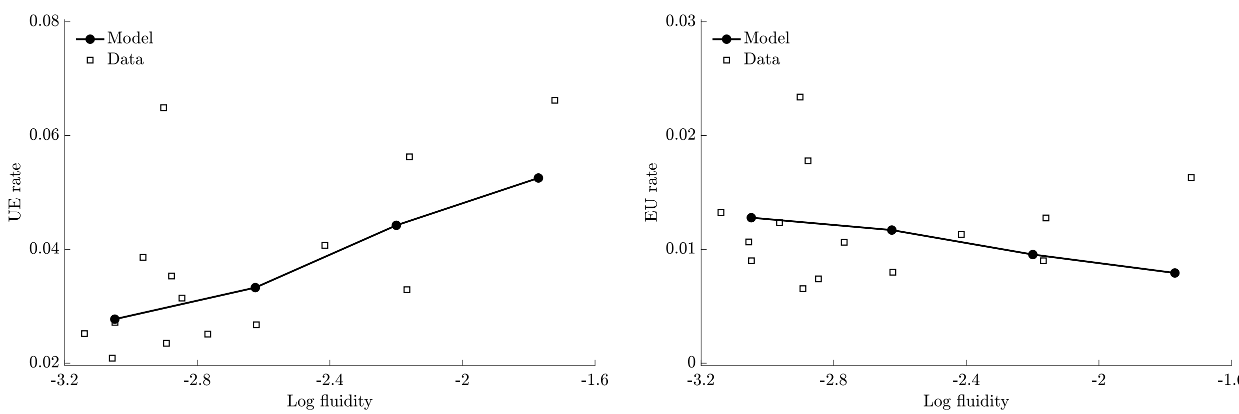

5.2 The Impact of Differences in the Functioning on the Labor Market

I start by quantifying the impact of the functioning of the labor market on life-cycle outcomes. Figure 9a plots wage growth between age 20–29 and 50–59 in the model and data. Wages grow by over 20 log points more in the most relative to the least fluid labor market, in both the model and data. In contrast, wages among workers aged 20–30 in panel (b) do not noticeably co-vary with fluidity, in either the model or data.

Panel (c) shows that young workers (aged 20–39) in more fluid labor markets tend to work for larger firms, with the model somewhat overstating the empirical relationship. Panel (d) shows that young workers are more likely to work for firms that offer training, with the model understating the empirical relationship. We construct this by tagging firms characterized by (z, x) as training or not in the baseline economy such that 41% of employment is at training firms—consistent with the data—and subsequently holding a firm's training type fixed across countries.

Figure 9: Cross-Country Life-Cycle Outcomes

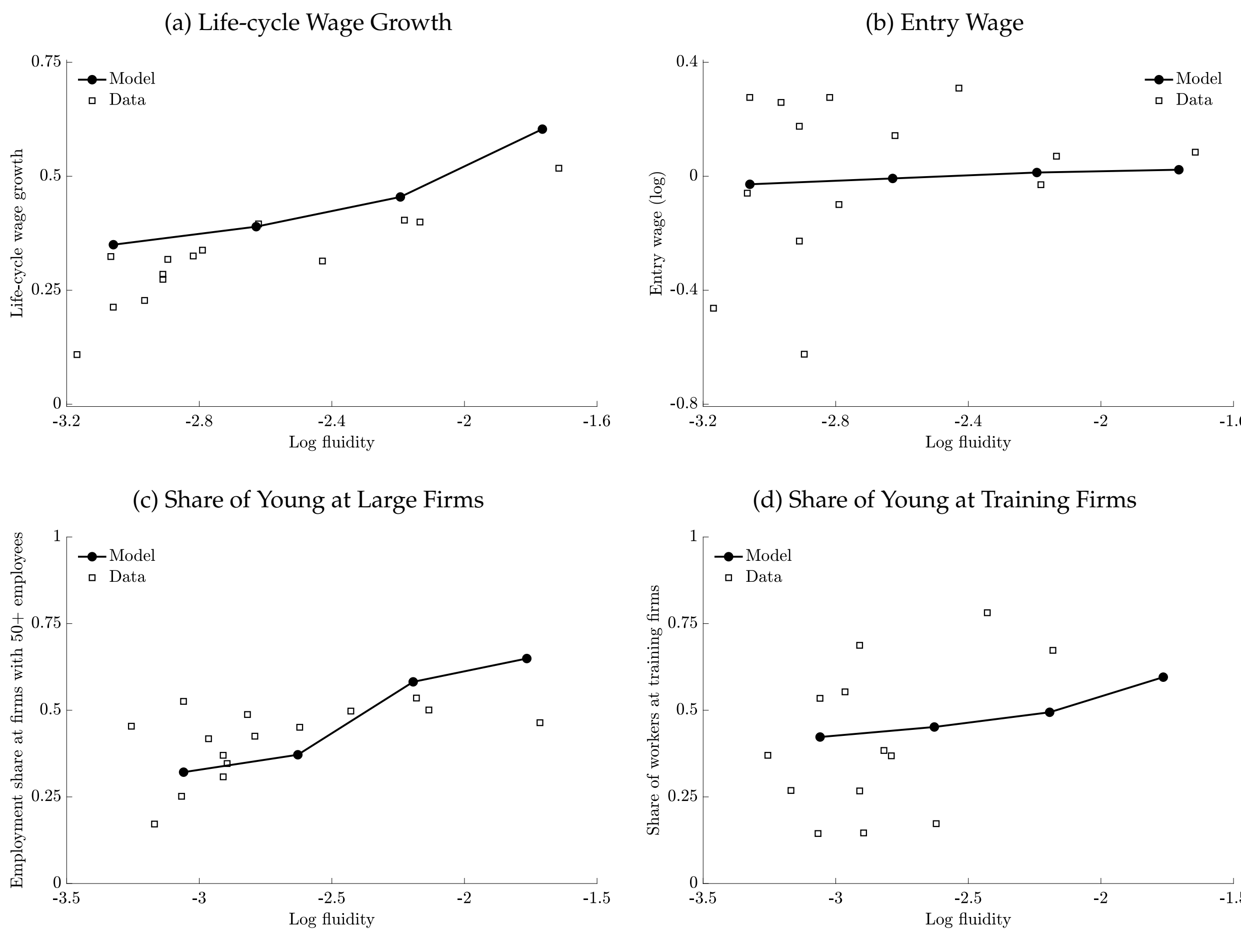

5.3 Understanding the Effects

Figure 10 illustrates the effect of these changes on the vacancy distribution and young workers' indifference curves. I contrast the super fluid to the low fluidity country, with blue intensity meaning a larger mass in the super fluid country and red a larger mass in the low fluidity country. In a more fluid economy, young workers anticipate to be less mismatched later in life—including spending less time as unemployed. Moreover, they know that they will have an easier time later gravitating away from high learning environment firms toward high productivity firms, as their learning ability and time horizon falls and they hence value learning environment less. Both forces lead young workers to tilt their preferences more toward high learning environment firms in high fluidity markets. Moreover, firms respond to workers' changing preferences, leading high-learning environment firms to disproportionately raise vacancy creation. For both reasons, younger workers are more likely to work in high-learning environment firms in high-fluidity countries.

Figure 10: Changes in Vacancy Distribution and Indifference Curves

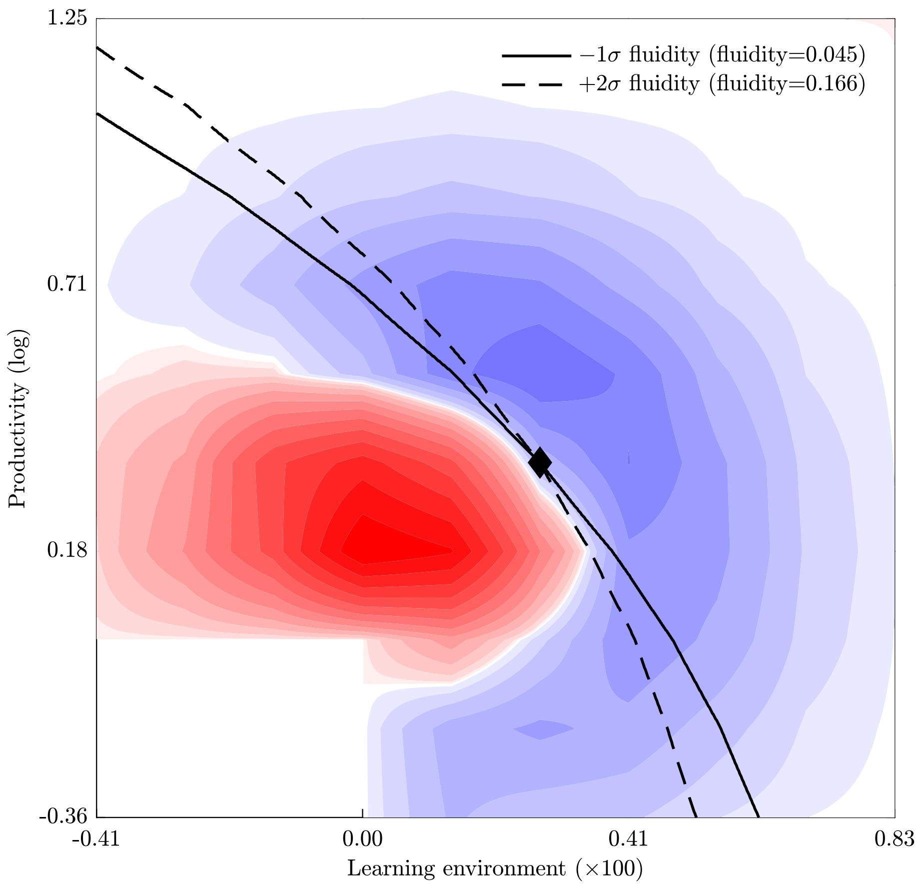

Figure 11 decomposes the cross-country variation in growth in various outcomes between ages 20–30 and 50–60. Recall that the wage is log additive in match productivity, human capital and the piece rate, so this decomposition is exact. According to panel (a), wages grow more over the life-cycle in roughly equal part due to faster growth in match productivity, human capital and the piece rate. Match productivity grows more as workers move up the job ladder more. Human capital grows more because workers have an easier time finding a good learning environment when they are young. Moreover, anticipating having greater use for their skills later in life as they will be less mismatched, young workers tilt their preferences toward finding a better learning environment. Finally, the piece rate grows faster in more fluid labor markets, as workers accumulate more outside offers that they can use to bargain up their wage.

Panel (b) shows the same outcomes among workers aged 20–30, relative to the corresponding outcomes in the targeted baseline economy. In a more fluid labor market, match productivity and human capital is already higher at age 20–30. In contrast, young workers have a lower piece rate in more fluid economies. The value of accumulating skills mostly accrue to the worker after leaving the firm, benefiting the current employer little directly. The worker hence compensates the employer for this learning through a lower piece rate initially. As the value of human capital grows, so does the initial compensation offered by workers to firms.

Figure 11: The sources of steeper life-cycle wage growth

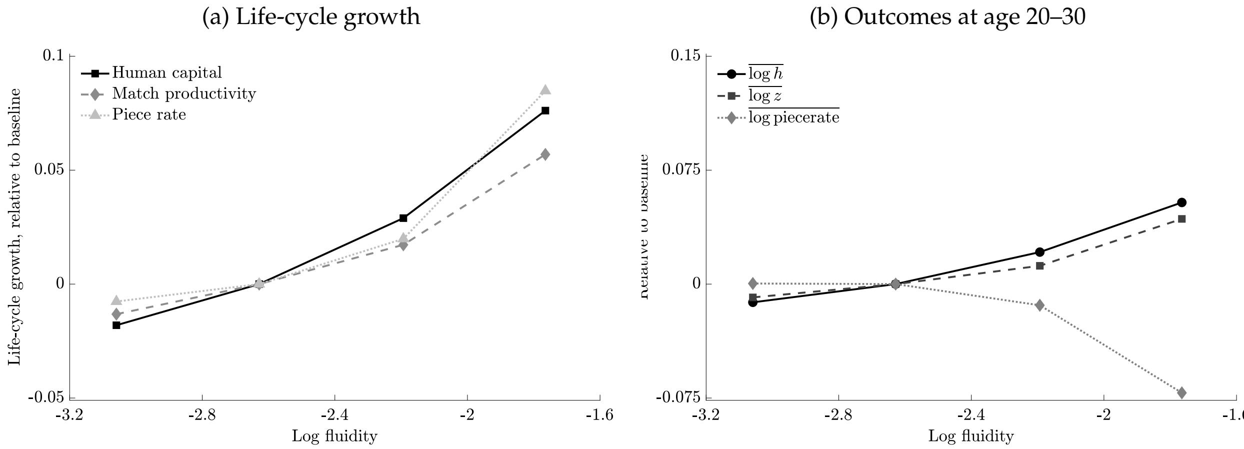

Figure 12 plots aggregate output, match productivity and human capital against fluidity. Aggregate output is roughly 40 percent higher in the super fluid country relative to the low fluidity country. Greater human capital accounts for roughly two-thirds of this pattern. Panel (b) shows an alternative counterfactual in which p and ϕ^e vary as estimated, but policies—including workers' acceptance and mobility decisions and firms' vacancy policies—are held fixed at the baseline policies. The mechanical effect of workers more quickly finding a good learning environment is important, with changes in workers' preferences across jobs and firms' vacancy creation decisions alone generating roughly a quarter of the aggregate increase in human capital.

Figure 12: Aggregate Implications

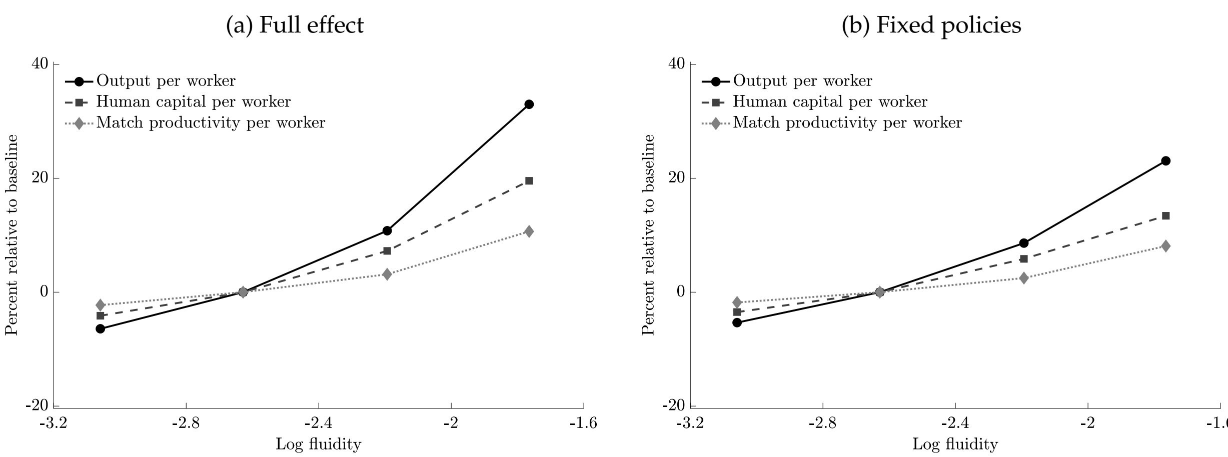

References

Abowd, John M., Francis Kramarz, and David N. Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica*, 67(2): 251–333.

Arellano-Bover, Jaime. 2024. "Career Consequences of Firm Heterogeneity for Young Workers: First Job and Firm Size." *Journal of Labor Economics*, 42(2): 549–589.

Arellano-Bover, Jaime, and Fernando Saltiel. 2026. "Differences in On-the-Job Learning across Firms." *Journal of Labor Economics*, 44(1): 149–188.

Bagger, Jesper, François Fontaine, Fabien Postel-Vinay, and Jean-Marc Robin. 2014. "Tenure, Experience, Human Capital, and Wages: A Tractable Equilibrium Search Model of Wage Dynamics." *American Economic Review*, 104(6): 1551–1596.

Becker, Gary S. 1962. "Investment in Human Capital: A Theoretical Analysis." *Journal of Political Economy*, 70(5, Part 2): 9–49.

Ben-Porath, Yoram. 1967. "The Production of Human Capital and the Life Cycle of Earnings." *Journal of Political Economy*, 75(4, Part 1): 352–365.

Bentolila, Samuel, and Giuseppe Bertola. 1990. "Firing Costs and Labour Demand: How Bad Is Eurosclerosis?" *Review of Economic Studies*, 57(3): 381–402.

Burdett, Kenneth, and Dale T Mortensen. 1998. "Wage differentials, employer size, and unemployment." *International Economic Review*, 257–273.

Cahuc, Pierre, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. "Wage Bargaining with On-the-Job Search: Theory and Evidence." *Econometrica*, 74(2): 323–364.

Diamond, Peter A. 1982. "Wage Determination and Efficiency in Search Equilibrium." *Review of Economic Studies*, 49: 217–227.

Gregory, Victoria. 2026. "Firms as Learning Environments." Working paper.

Hopenhayn, Hugo, and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis." *Journal of Political Economy*, 101(5): 915–938.

Jolivet, Gregory, Fabien Postel-Vinay, and Jean-Marc Robin. 2006. "The empirical content of the job search model: Labor mobility and wage distributions in Europe and the US." *Contributions to Economic Analysis*, 275: 269–308.

Lagakos, David, Benjamin Moll, Tommaso Porzio, Nancy Qian, and Todd Schoellman. 2018. "Life Cycle Wage Growth across Countries." *Journal of Political Economy*, 126(2): 797–849.

Lise, Jeremy, Costas Meghir, and Jean-Marc Robin. 2016. "Matching, Sorting and Wages." *Review of Economic Dynamics*, 19: 63–87.

Ljungqvist, Lars, and Thomas J. Sargent. 1998. "The European Unemployment Dilemma." *Journal of Political Economy*, 106(3): 514–550.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *The Review of Economic Studies*, 61: 397–415.

Postel-Vinay, Fabien, and Jean-Marc Robin. 2002. "Equilibrium Wage Dispersion with Worker and Employer Heterogeneity." *Econometrica*, 70(6): 2295–2350.

Pries, Michael, and Richard Rogerson. 2005. "Hiring Policies, Labor Market Institutions, and Labor Market Flows." *Journal of Political Economy*, 113(4): 811–839.

Topel, Robert H, and Michael P Ward. 1992. "Job mobility and the careers of young men." *Quarterly Journal of Economics*, 107(2): 439–479.