Employment Protection, Investment in Job-Specific Skills, and Inequality Trends in the United States and Europe*

Matthias Doepke
Northwestern University

Ruben Gaetani
University of Toronto

October 2017

Abstract

Since the 1980s, the United States has experienced a sharp rise in the college wage premium. In contrast, in most European economies the return to education rose only a little, and in Germany, Italy, and Spain the college premium actually fell. In this paper, we argue that differences in employment protection can account for a substantial part of these diverging trends. In our model, firms and workers can invest in relationship-specific capital: firms can create jobs that are complementary to experienced workers with long tenure, and workers can make corresponding investments in firm-specific skills. The incentives to undertake such investments interact with employment protection. Incentives are strong if employment protection favors older workers with long tenure, as is the case in the European countries where the college premium fell. Firms and workers also invest in relationship-specific capital in a calm economic environment where firm-specific shocks are small. The diverging inequality patterns between the United States and Europe emerge from different levels of employment protection combined with an increase in "turbulence" (Ljungqvist and Sargent 1998) in the economy starting in the 1980s.

*Preliminary and incomplete. Comments are welcome. Financial support from the National Science Foundation (grant SES-1260961) is gratefully acknowledged. Department of Economics, Northwestern University, 2211 Campus Drive, Evanston, IL 60208 (e-mail: doepke@northwestern.edu). Gaetani: Institute for Management and Innovation and Rotman School of Management, University of Toronto, 105 St. George Street, Toronto, ON M5S 3E6, Canada(e-mail: ruben.gaetani@utoronto.ca).

1 Introduction

Since the 1980s, the United States economy has experienced a sharp rise in education premia in the labor market, with the college premium going up by more than 30 percent. A similar trend is observed in other Anglo-Saxon economies, including Canada and the United Kingdom. During the same period, the relative supply of highly educated labor has also risen in these countries, suggesting a long-term upward shift in the demand for skills. Other industrialized countries also experienced a rise in the supply of educated workers, but education premia did not follow the Anglo-Saxon trend. Specifically, most European economies witnessed only a small rise in the return to education since the 1980s, and in Germany, Italy, and Spain the college premium actually fell.

The different trends in education premia suggests that it is not just global technological trends, but also country-specific factors that drive the return to education. In this paper, we argue that differences in employment protection can account for a substantial part of these diverging trends. In particular, we argue that firing restrictions affect the incentives to invest in job-specific skills. If there is a systematic difference in the extent to which workers with different education can transfer experience across different jobs, firing restrictions and between-group inequality can be tightly related. We also argue that restrictions for firing older workers are particularly relevant, which is where differences between Europe and the United States are the largest.

We start by documenting a strong negative correlation between the increase in the education premium since the 1980s and strength of employment protection legislation in a sample of OECD countries. We further show that the recent wave of labor-market deregulation that has taken place in many European countries (notably, Germany and Italy) in the last decades has mostly affected temporary employees, while leaving restrictions for regular employees substantially unaltered.

Using data from the Panel Study of Income Dynamics for the United States and the Socio-Economic Panel for Germany, we show that the share of old, non-college educated workers with long-term tenure has declined substantially in

the U.S. since the 1980s, but has remained unchanged in Germany. For college-educated workers, in contrast, the share of employees with long-term tenure has declined significantly both in the U.S. and Germany. In the spirit of Ljungqvist and Sargent (1998), we interpret the overall decline in the share of workers with long-term tenure as being due to increased "turbulence" in the economy, i.e., a rise in the variance of firm-specific shocks that can lead to separations. We take the view that firing restrictions in Germany have provided unskilled workers with some insurance against higher turbulence. We also show that returns to long-term tenure increase significantly in Germany and decline in the U.S. for non-college educated workers. This is consistent with the idea that the incentives to invest in job-specific skills are negatively related to the probability of separation.

We develop a model of the labor market in which workers can invest in skills while employed. Workers are subject to separation shocks, and firing restrictions—that determine the probability of separation—depend on age and tenure. The central assumption that we maintain is that transferability of skills depends on education. Skills acquired on the job by college-educated workers are transferrable across different employment spells. By contrast, non-college educated workers can only acquire job-specific skills that are lost upon separation. The model offers a novel perspective on changes in the skill premium. Namely, if college education gives an advantage in transferring skills across jobs, an increase in the rate of separation will have a first-order effect on the education premium.

We calibrate the model to match salient labor market facts in the United States and Germany. In our quantitative exercise, we infer the turbulence shock from the U.S. labor market variables, but assume that it only affects workers in Germany as long as they have low seniority and do not enjoy labor protection. Separation rates for senior workers in Germany—the ones who are protected by firing restrictions—are inferred directly from the German data. In low-turbulence times there is low probability of separation even in the absence of firing restrictions. As a result, all workers invest in skills regardless of regulation. In turbulent times, investment choices crucially depend on firing restrictions. In the absence of firing restrictions (i.e. in the U.S.) only educated workers invest in skills, generating

a high wage premium. However, in an environment with firing restrictions for experienced workers (i.e. in Europe) all workers continue to invest. This produces a lower education premium compared to the unconstrained economy. Our quantitative analysis implies that the mechanism explains about one-third of the overall increase in the education premium in the United States and provides a novel explanation for why the German labor market did not experience the same increase in the returns to education.

The idea that differences in labor market regulation lead to diverging labor market outcomes in Europe and the United States has a long history. Blanchard and Summers (1986) argue that European hysteresis can be due to labor market institutions that favor "insiders"versus "outsiders."This fact generates persistent, long-term shifts in the rate of unemployment. They document the following empirical facts: (1) Unemployment increased substantially in Europe (Germany, France and the UK) since the 1970s, less so in the United States and (2) Unemployment is highly persistent both in Europe and the US –slightly more so in Europe. The theoretical mechanism they suggest is that only insiders matter for setting the wage: If there is a shock that decreases the number of insiders, the remaining ones will set the wage in a way that keeps them employed, but will not care about outsiders re-entering employment. Hence, a shock to the unemployment rate can have high persistence. The theory is consistent with evidence on unionization in Europe and the US. Bentolila and Bertola (1990) study a partial equilibrium setting in which a firm takes hiring and firing decision under linear adjustment costs. Bertola and Ichino (1995) put forth the idea that shocks to the returns to skills (including trade shocks with developing countries) translate into higher wage dispersion in the US and higher unemployment rate in Europe. Nickell (1997) observes that there is large variation in labor market outcomes not only between Europe and the US, but also within Europe. He reports several cross-country correlations between various measures of unemployment and labor market participation and labor market institutions (e.g. unemployment benefits and labor protection). He finds no effect of labor protection or "well-designed"unemployment benefits and large effects of "badly-designed"unemployment benefits, unionization and labor taxes.

This paper contributes to the literature on the relationship between skills, technological adoption and labor market rigidities. Aghion, Howitt, and Violante (2002) argue that information technology is "general purpose": for educated workers, it increases the transferability of skills across sectors, and it makes physical capital less sector-specific and hence allows for a faster reallocation of capital to growing sectors. Both effects lead to an increase in inequality in response to the introduction of information technology. Our paper also builds on the notion that educated workers are characterized by higher transferability of human capital (here across employment spells), but focuses on other driving forces for changing inequality. On the empirical side, Gathmann and Schönberg (2010) estimate to what degree human capital acquired on the job can be thought of as task-specific. They use the German Employee Panel to test the following predictions: (1) tasks required in the "source" occupation are "close"to the tasks in the landing occupation; (2) the "distance"between source and landing occupation is negatively related to labor market experience; (3) the wage in the landing occupation is higher if the source occupation was "close"; (4) the current wage is positively related to "tenure"in tasks that are close to the current occupation. They also show that these patterns are typically strongest for the high-skilled, suggesting that task-specific skills are especially important for this education group. Wasmer (2006) develops a model to connect labor market rigidities with the incentives to invest in job-specific skills. High search frictions and firing costs increase the value of an existing match compared to separation. This leads to a lower rate of separation and increases the incentives to invest in match-specific skills compared to generic skills.

Acemoglu (2003) offers a direct analysis of the different experiences in terms of inequality between the United States and Europe that emphasizes the demand, rather than the supply, of skills. He decomposes the change in the skill premium in a given country as the change in technology and change in supply of skills. He then infers the change in technology by applying the decomposition to the United States, and uses the resulting estimate to compare the "actual"and the "predicted"change in the skill premium for a variety of countries. He finds significant differences between predicted and realized skill premium for several countries. He then formulates the following hypothesis: demand for skills changes

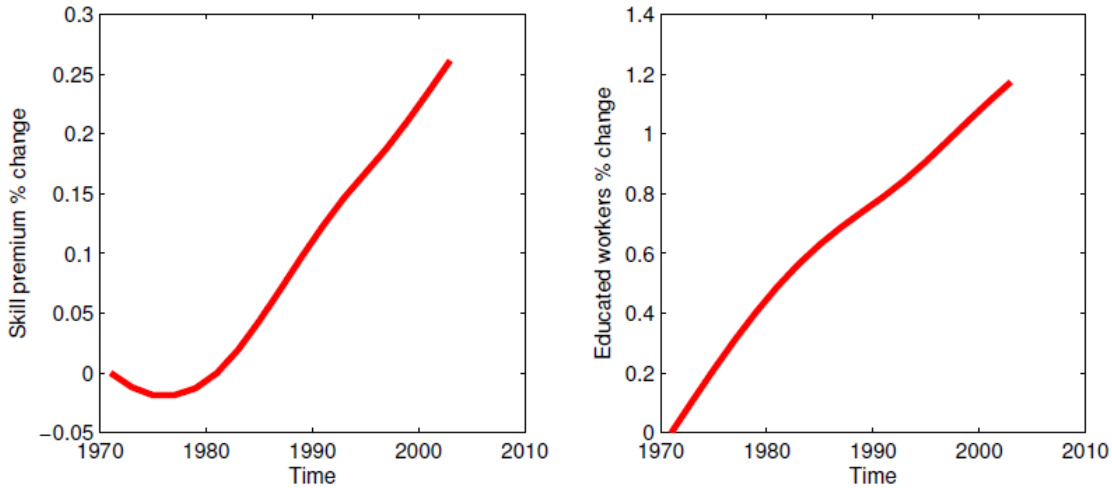

Figure 1: College wage premium and share of college educated workers in the U.S.

endogenously and, if institutions are such that relative wage of skilled relative to unskilled workers is compressed, firms will have an incentive to invest in technologies complementary to low-skill workers. The reason is that an increase in the productivity of a worker whose wage is above his marginal product will not require a wage increase even if his marginal product increases.

The following section describes the main empirical patterns that motivate this study. In Section 3 we illustrate the mechanism that we are proposing using a two-period model that can be solved analytically. Our quantitative analysis is described in Section 4, and Section 5 concludes.

2 Empirical Patterns

Figure 1 displays the contemporaneous increase in the supply of college educated workers (right panel) and the college wage premium (left panel) experienced in the United States in the last few decades. The fact that the two variables trend in the same direction has long been interpreted as evidence of a change in the demand for skills on the labor market. For example, Acemoglu (2002) argues that an increase in the supply of skills led firms to invest and adopt technologies that complement those skills, leading to skill-biased technical change.

An important limitation of the skill-biased technical change explanation is that

Figure 2: Labor protection and change in education premium. Data from Krueger et al. (2010) and Nickell (1997).

the supply of college educated workers has risen in all industrialized countries, whereas there is no uniform pattern in changes in the college wage premium. This observation suggests that there are additional country-specific factors that have a substantial impact on the skill premium. Our study is motivated by the observation that changes in the skill premium across countries are empirically related to measures of labor protection. Figure 2 displays the percentage change in the college premium in the years 1980–2006 in a sample of OECD countries, plotted against the OECD labor protection index. The data on education premia are taken from Krueger et al. (2010), while the OECD employment protection index is taken from Nickell (1997). The picture shows a strong negative correlation between the two. The United States, the United Kingdom and Canada have experienced a substantial increase in the return to education, with the figure for the United States being over 40%. In contrast, Germany, Italy and Spain—economies characterized by significantly stronger employment protection legislation—have experienced a flat or declining return to college.

In this paper, we analyze a specific mechanism that can account for the empirical relationship between employment protection and changes in the skill pre-

mium. In particular, we argue that employment protection can increase the incentives for workers and firms to make relationship-specific investments that pay off if workers stay with the same firm for a long time. These investments take the form of investment in firm-specific skills for workers, and the creation of jobs that employ such firm-specific skills on the part of firms. In addition to employment protection, the incentive to make relationship-specific investments depends also on the size and frequency of shocks to the productivity of a worker-firm match. We argue that employment protection is especially important when there are frequent shocks. This suggests that when there is an increase in economic turbulence, relationship-specific investments should decline in countries with little employment protection, but not in countries where employment protection creates an expectation of worker-firm matches of long duration. If high and low-education workers differ in their ability to transfer skills across employment spells, these differences will be reflected in the evolution of skill premium in response to changes in economic turbulence.

At first sight, our argument may appear to be contradicted by the fact that in recent decades many European countries deregulated labor markets at least to some extent. However, the kind of employment protection that is particularly relevant is one that becomes more stringent the longer the worker has been employed by a firm, so that there is an expectation that long-run relationship-specific investments will pay off. From this perspective, it is important to note that recent changes in employment protection in Europe have been mostly aimed at new entrants to the labor market, while keeping protections for workers with high tenure in place. In Figure 3 and 4 we plot the OECD index of labor protection for the U.S. and a subset of European countries over the last decades. We notice a striking difference between the evolution of the index for regular employees (Figure 3) which changes hardly at all, and the strongly declining protection for temporary employees (Figure 4) in Europe. In our model, we will make this distinction explicit and argue that what matters the most in determining the propensity to invest in skills on the job is the probability of separation for workers with long-term tenure.

Our mechanism linking employment regulation and inequality relies on the no-

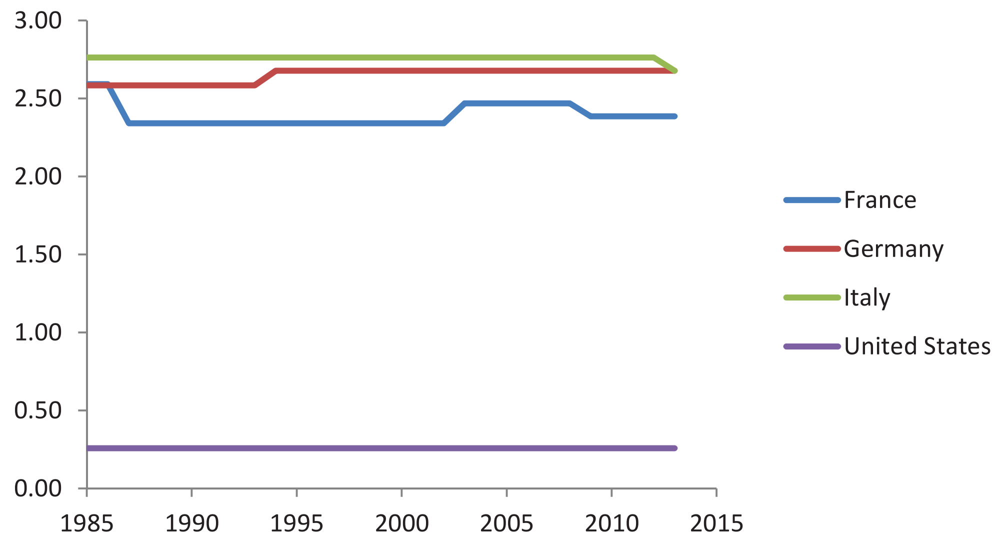

Figure 3: OECD index of labor protection for regular workers

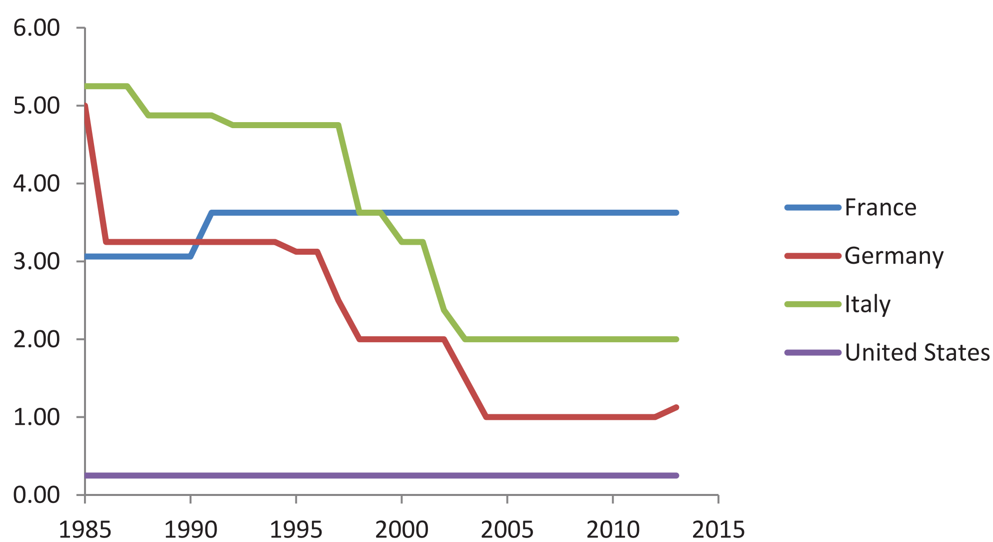

Figure 4: OECD index of labor protection for temporary workers

tion that in economies with more employment protection workers experience less job turnover and hence are more likely to achieve long tenure with a given employer. To verify whether this holds in the data, we use data from the Panel Study of Income Dynamics (for the United States) and the Socio-Economic Panel (for Germany) and compute the shares of workers with high tenure for the last decades. Figure 5 shows the evolution of the share of non-college educated workers of age 45–60 with long-term tenure (15 years or more with the current employer) for the United States and Germany. This share displays a clear downward trend in the United States: it declines from about 55% in 1984 to about 40% in 2013. The pattern for Germany is strikingly different. While in 1984 the share is comparable in magnitude to the U.S. share (about 57%) this figure does not display any clear trend and in 2013 is slightly above its value 1984. In contrast, these different patterns are not observed for college-educated workers (Figure 6), where the share of workers with long-term tenure declines by similar amounts in the United States and Germany.

The different patterns for more and less educated workers are what our theory predicts. In our model, college-educated workers are more likely to accumulate general skills that can be transferred across employers, whereas less educated workers accumulate mostly firm-specific skills. An increase in turbulence then leads to more turnover for educated workers, who respond to relationship-specific shocks by taking their skills to another employer.

Our theory also makes predictions for the returns to tenure. We posit that in the United States, firms and workers responded to increased turbulence by reducing relationship-specific investments, whereas such investments change less (or not at all) when there is a lot of employment protection for high-tenure workers as in Germany. Given that empirical measures of the returns to tenure capture the returns to relationship-specific investment, the theory then implies that the returns to tenure should have fallen in the United States, but not in Germany. Figure 7 shows that the data is consistent with this implications. The returns to long-term tenure decline in the U.S. for non-college educated workers, and in fact increase in Germany.

Finally, our mechanism relies on the assumption that the transferability of skills

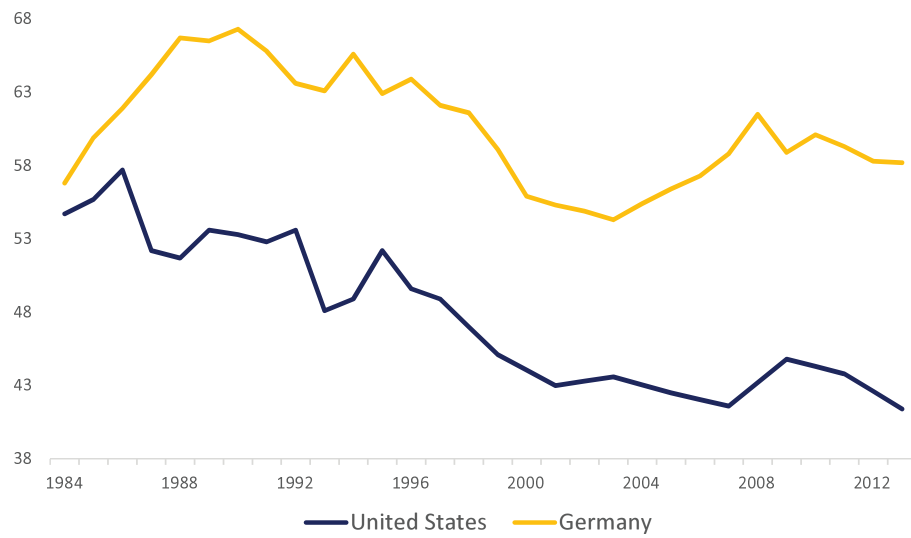

Figure 5: Share of non-college educated workers of ages 45–60 with long-term tenure (at least 15 years) in United States versus Germany

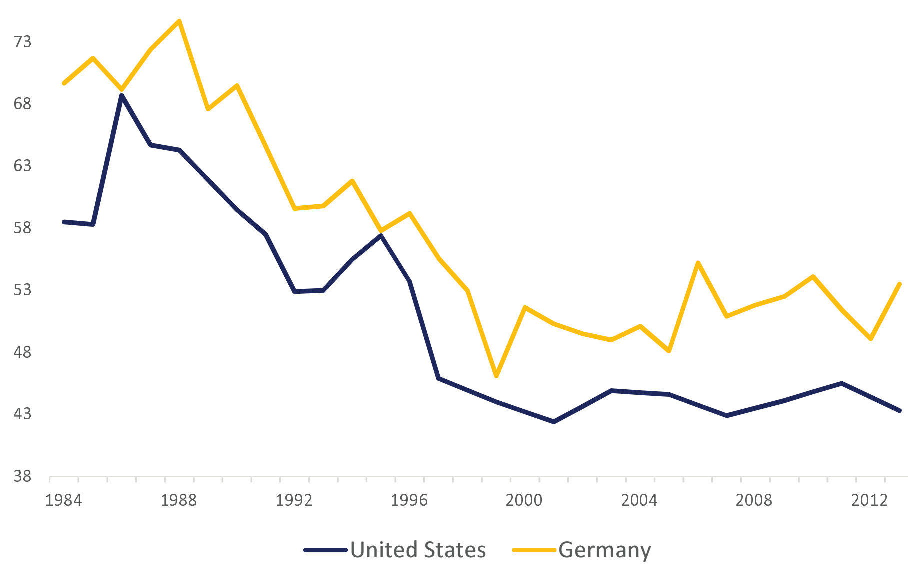

Figure 6: Share of college-educated workers of ages 45–60 with long-term tenure (at least 15 years) in United States versus Germany

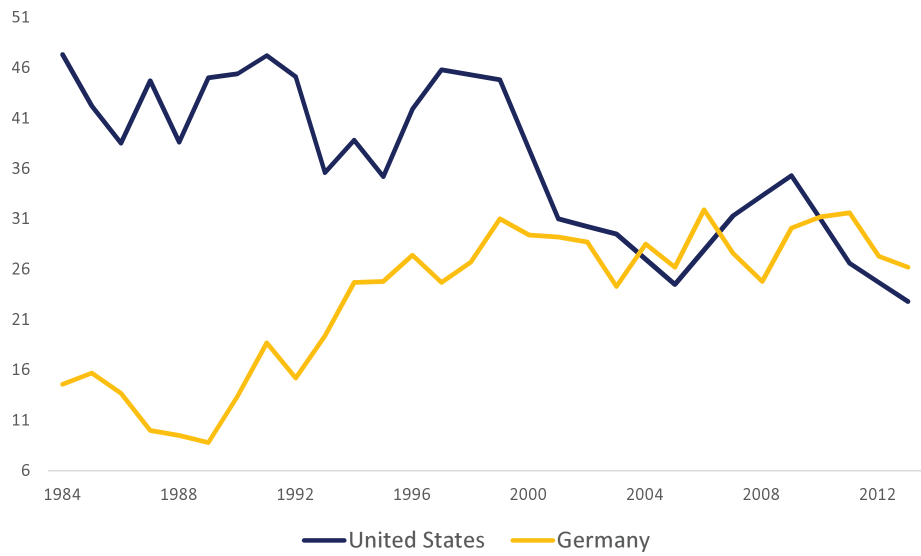

Figure 7: Returns to long-term tenure for non-college educated

acquired on the job depends on education. In particular, in our model college-educated workers can transfer their work experience to the next job. By contrast, non-college educated workers acquire job-specific skills and lose those skills upon separation. There is an existing literature (see previous section) in labor economics that is consistent with this view. We can also examine this assumption using the same PSID and SOEP data that we use to establish other facts above. If skills are fully transferable across employers, a worker's productivity should not depend on years of tenure with the current employer, after controlling for overall labor market experience. On the other hand, if transferability is not perfect, years of tenure with the current employer should positively affect a worker's salary, even after controlling for overall labor market experience. To test this hypothesis, we exploit the panel dimension of the PSID and the SOEP, restricting attention only to male individuals that are continuously employed (namely, who have labor earnings at each observation in the sample).

We run the following regression separately for college-educated and non-college educated workers:

$$y_{it} = \beta_t + \alpha_i + f\left(exp_{it}\right) + \gamma\, tenure_{it} + \epsilon_{it},$$

	Log of hourly wage			
	Germany (SOEP)		USA (PSID)	
Tenure, non-college	.006***		.004***	
	(.001)		(.001)	
Tenure, college		.001		.002
		(.001)		(.002)
Exper. 3rd degree pol.	yes	yes	yes	yes
Individual FE	yes	yes	yes	yes
Year FE	yes	yes	yes	yes

Table 1: Panel regression: returns to tenure for college- and non-college-educated workers in Germany and the United States.

where y_{it} is the logarithm of the hourly wage, β_t is a time fixed effect, α_i is an individual fixed effect, $f\left(exp_{it}\right)$ is a third-degree polynomial of potential labor market experience, and $tenure_{it}$ denotes years of tenure with the current employer. The estimate of γ denotes how much of his hourly wage a worker loses (in percentage terms) from the previous to the next employment spell for each year of continuous employment in his previous job.

The results can be found in Table 1. Consistent with our hypothesis, the estimate of γ is positive and significant for non-college educated workers. The magnitude of the coefficient implies that after a 10-year employment spell, a non-college educated worker loses on average 6% of his salary in Germany and 4% in the United States. By contrast, the estimate of γ is close to zero—and statistically insignificant—for college-educated workers, both in Germany and in the United States.

3 The Mechanism in a One-Period Model

We start by developing a simple static model that illustrates the mechanism behind our quantitative result. The purpose of the model is to explain how the rise of turbulence, which will be modeled as an increase in the probability of separation between employer and employee, is connected to the increase in the skill

premium, the decline in the tenure premium, and the prevalence of jobs for low-educated workers that do not allow for on-the-job skill accumulation. Also, we will highlight a potential role for- labor protection legislation in moderating these effects.

One challenge for modeling is that if both firm and worker can make relationship-specific investments, complicated strategic interactions may arise. We will address this by considering a setting where the firm decides ex-ante what type of vacancy to open: only some vacancies allow for high returns to effort. The model will generate a shift towards vacancies that do not reward investment in job-specific skills when there is an increase in the separation rate.

3.1 Setting

Consider an economy with unit mass of both college educated workers (H) and less educated workers (L). The labor markets for the two types of workers are separated, and there is a representative firm in each market. In the market for educated workers, all jobs allow for the accumulation of skills, and the representative firm opens unit mass of such vacancies (overall vacancy creation will be endogenized in the dynamic model below). In the market for less educated workers, the firm has a choice between opening vacancies that allow for skill accumulation and those that do not. Specifically, the firm in market L opens a mass v_T of vacancies at cost $\kappa_T v_T^2$ that allow for skill accumulation (T), and a mass $v_N = 1 - v_T$ of vacancies (at zero cost) that do not allow for accumulation of skills (N). Each worker in each market is randomly assigned to a vacancy, so that less educated workers are matched with a T job with probability v_T and with an N job with probability v_N.

Upon matching, workers with a job that allows for the accumulation of skill (all H workers and fraction v_T of L workers) can exert effort $e \in (0, 1)$ on acquiring skills on the job. In the dynamic model below, this effort will be endogenous and costly, but here we simply take e as given. Endogenizing e would only amplify the effect of a changing composition of vacancies that we focus on here, but not change the qualitative features of the model. Effort takes place before a wage is negotiated, so that it is not part of the match surplus and does not appear in the

wage negotiation (we can think of effort as being exerted in an early stage of the relationship, such as an internship, before a full employment contract is signed).

Effort is successful with probability e (in which case we say the worker becomes skilled) and is unsuccessful with probability $1 - e$. If effort is successful, output within the match will be ay, with $a > 1$, whereas in case effort is unsuccessful output is simply y. For N-type jobs, no skills can be accumulated, so that output is always y.

At the same time when the success of effort is realized, a negative productivity shock can hit the match, which occurs with probability $\gamma \in (0, 1)$. We will refer to this shock as a "turbulence shock," and to the parameter γ as the degree of labor market turbulence. If the negative productivity shock occurs, output in the match is reduced by a factor ϵ, where ϵ is a random variable that is uniformly distributed on the interval $(0, 1)$.

The outside option for workers is to engage in home production. For workers who did not become skilled the output in home production is given by $b > 0$, regardless of education. For workers whose effort was successful and who are therefore skilled, output in home production depends on the transferability of the accumulated skills. We assume that educated (H-type) workers acquire general skills that also apply to home production. Hence, an educated worker whose effort in accumulating skill was successful has productivity ab in home production (recall that $a > 1$). For less educated (L-type) workers, skills are more job specific and hence less transferable. Specifically, skilled L-type workers can transfer only a fraction $\tau < 1$ of the accumulated skills, so that home production is τab, where $a > a\tau \geq 1$.

After uncertainty has been resolved, if the surplus from the match is negative, separation occurs and the worker engages in home production. If there is positive surplus, a wage is negotiated via Nash bargaining, where the worker gets fraction β of the surplus.

Note that H and L workers differ along two dimensions. First, agents of type H only work in jobs that allow for on-the-job skill accumulation, whereas only a fraction v_T of agents of type L work in vacancies with returns to experience.

Second, H and L workers differ in their degree of skill transferability from employment to home production: H workers can fully transfer their skills, whereas L workers can only transfer a fraction τ of the skills acquired on the job.

3.2 Solving the Model

The key tradeoff in this model is the choice of the firm in the L market between posting vacancies with and without the possibility of skill accumulation. The firm's problem is to choose v_T and $v_N = 1 - v_T$ in such a way as to maximize its overall value, subject to the quadratic cost of posting the T-type vacancy. The optimal v_T is given by:

$$v_T^* = \frac{J_T - J_N}{2\kappa_T}, \tag{1}$$

where J_T and J_N are the firm's profit from posting each type of vacancy. We assume that the cost parameter κ_T is sufficiently large so that $v_T^* < 1$.

To determine J_T and J_N, we can solve the model backwards, starting from the wage determination after all uncertainty has been resolved, and then solving for the firm's optimal choice of v_T and v_N. We first characterize the equilibrium wage and separation decisions for T-type vacancies. The solution for N-type vacancies is equivalent to the case of T-type vacancies with unsuccessful effort.

Consider firm-worker match where the worker is of type L and the job is of type T, i.e., the job allows for accumulation of skill and the worker exerts effort e. If effort is successful and no negative productivity shock occurs, total surplus is:

$$S_T^1 = ay - a\tau b,$$

which implies that the wage of the worker is $w_T^1 = \beta ay + (1-\beta)\, a\tau b$ and the firm gets $q_T^1 = (1-\beta)\,(ay - a\tau b)$. Here the superscript 1 denotes successful effort.

If effort is successful and the negative shock does occur (this happens with probability γ) the realization of ϵ will determine whether there will be separation. Surplus is now given by

$$\hat{S}_T^1(\epsilon) = a\epsilon y - a\tau b,$$

where the hat denotes an equilibrium variable under the negative productivity

shock. We assume efficient separation decisions, so that separation will occur if $\epsilon < \bar{\epsilon}_T^1 = \frac{\tau b}{y}$. If separation does not occur (conditional on the turbulence shock, this happens with probability $\frac{y-\tau b}{y}$) the expected surplus is:

$$E\left[\hat{S}_T^1\right] = a\left(\frac{\tau b + y}{2} - \tau b\right)$$

and the expected wage for the worker is: $E\left[\hat{w}_T^1\right] = \beta a \frac{\tau b + y}{2} + (1-\beta)\, a\tau b$.

Finally, if separation occurs (conditional on the turbulence shock, this happens with probability $\frac{\tau b}{y}$), the worker gets $a\tau b$ and the firm gets 0.

Taking into account each possible outcome, in case of successful effort (i.e. with probability e) the expected profit of the firm is given by:

$$J_T^1 = (1-\gamma)(1-\beta)\, a\,(y - \tau b) + \gamma(1-\beta)\, a\left(\frac{y^2 + \tau^2 b^2 - 2y\tau b}{2y}\right). \tag{2}$$

The expected value of the vacancy in case of unsuccessful effort is also given by (2) once we set $a = \tau = 1$:

$$J_T^0 = (1-\gamma)(1-\beta)(y - b) + \gamma(1-\beta)\left(\frac{y^2 + b^2 - 2yb}{2y}\right). \tag{3}$$

Note that J_T^0 is also equal to the value of a vacancy of type N, J_N.

From equation (1), we can see that the composition of vacancies only depends on the difference between the returns from opening either type of vacancy:

$$\Delta J = J_T - J_N = eJ_T^1 + (1-e)\, J_T^0 - J_N = e\left(J_T^1 - J_T^0\right).$$

Plugging (2) and (3) into the expression for ΔJ, we obtain:

$$\frac{\Delta J}{e} = (1-\gamma)(1-\beta)\, y\,(a-1) - (1-\beta)\, b\,(\tau a - 1) + \frac{\gamma(1-\beta)}{2y}\left[y^2(a-1) + b^2\left(a\tau^2 - 1\right)\right]. \tag{4}$$

This expression, combined with equation (1), determines the optimal choice of

the composition of vacancies for workers of type L.

From a direct inspection of equations (1) and (4), the following result is immediate:

Proposition 1. *In equilibrium, the share of jobs that allow for on-the-job skill accumulation is a decreasing function of the degree of labor market turbulence* γ*.*

Intuitively, higher turbulence lowers the return to creating vacancies that allow for skill accumulation, because these skills will not pay off if a productivity shock leads to a separation. As a consequence, an increase in turbulence also lowers the average wage of L-type workers and increases the wage gap between H and L type workers. Before considering how these effects are modified if there is labor protection legislation, we discuss the welfare properties of the equilibrium.

3.3 Planner's Solution

The equilibrium choice of vacancies is generally constrained inefficient. Specifically, a planner who can control the composition of vacancies and takes everything else (effort choice, wage setting, and separation decisions) as given will generally choose a value of v_T that is higher than the equilibrium value. The reason behind this inefficiency is that, while the marginal private cost of posting a vacancy is equal to the marginal social cost, the marginal private and social benefit differ. There are two sources of the gap between private and social returns. First, in the case of no separation, only a fraction $1 - \beta$ of the resulting surplus is captured by the firm. Second, in case of separation, only the worker will benefit from the skills acquired in the job. This second source of inefficiency is more severe the higher the degree of skill transferability τ, and the higher the frequency of job destruction, which depends on labor market turbulence γ.

We now characterize the solution of the constrained planning problem, and show how the inefficiency is related to the labor market characteristics. Below, we will argue that because of the inefficiency, a moderate amount of labor protection legislation (LPL) can both improve welfare and lower equality. Consider a planner who can control the composition of vacancies in the labor market for L workers, while leaving all other outcomes decentralized. Since agents are risk neutral and

utility is transferrable, the planner's problem reduces to maximizing output net of the cost of posting vacancies.

The planner's value of a T vacancy if effort is successful (1) and unsuccessful (0) is, respectively:

$$V_T^1 = (1-\gamma)\, ay + \gamma a \frac{y^2 + \tau^2 b^2}{2y},$$
$$V_T^0 = (1-\gamma)\, y + \gamma \frac{y^2 + b^2}{2y}.$$

The planner's solution is of the form (1), where the value of a vacancy for the planner V is used in place of the value of the same vacancy for the firm J. Hence, the choice of vacancies for the planner only depends on:

$$\frac{\Delta V}{e} = (1-\gamma)\, y\,(a-1) + \gamma \frac{y^2\,(a-1) + b^2\,(a\tau^2 - 1)}{2y}. \tag{5}$$

To evaluate the size of the inefficiency, it is instructive to look at the difference between the solutions to the planner's and the firm's problems. Although this difference is not directly interpretable as a welfare loss, there is a one-to-one mapping between the two:

$$\text{Loss} \approx \frac{\Delta V}{e} - \frac{\Delta J}{e} = (1-\gamma)\,\beta y\,(a-1) + (1-\beta)\, b\,(\tau a - 1) + \frac{\gamma\beta}{2y}\left[y^2\,(a-1) + b^2\left(a\tau^2 - 1\right)\right]. \tag{6}$$

Note that in general the welfare loss is positive, meaning that the equilibrium is supplying fewer vacancies of type T than the planner would.

The first term of the welfare loss corresponds to the loss of value from the fact that only a fraction $1-\beta$ of the surplus is captured by the firm, specifically when the turbulence shock does not occur (with probability $1-\gamma$). The second term originates from the fact that as long as there is some skill transferability, skill acquisition improves the bargaining position of the worker within the match. This improvement in the bargaining position of the worker reduces the firm's incentives to invest in the worker's skills. The third and final term combines two

effects: first, the firm does not internalize the value of skills outside the match (i.e. if separation occurs), and second, only a fraction $1-\beta$ of the surplus is appropriated by the firm if the turbulence shock occurs, but separation does not happen.

3.4 Introducing Firing Costs

We now would like to consider how labor protection legislation affects education premia in our model. We consider a simple form of labor protection, namely a firing cost $f > 0$ that the firm has to pay in case of separation. We focus on the case of a small firing cost, i.e., we will initially set $f = 0$ and then consider the effect of marginally increasing the firing cost.

If there is a separation, the profit of the firm will now be $-f$, which implies that the surplus when there is no separation is shifted up by f. Hence, if the turbulence shock does not occur and if effort is successful, match surplus is:

$$S_T^{1,f} = ay - a\tau b + f,$$

while in the case of unsuccessful effort we have:

$$S_T^{0,f} = S_N^f = y - b + f.$$

When the match experiences a turbulence shock, we can again define two threshold values $\bar{\epsilon}_T^{1,f}$ and $\bar{\epsilon}_T^{0,f}$ such that in case of successful effort and unsuccessful effort, respectively, separation occurs if and only if ϵ is below the threshold:

$$\bar{\epsilon}_T^{1,f} = \frac{a\tau b - f}{ay},$$
$$\bar{\epsilon}_T^{0,f} = \bar{\epsilon}_N = \frac{b - f}{y}.$$

The resulting expected value of the surplus conditional on the negative productivity shock is:

$$E\left[\hat{S}_T^{1,f}\right] = \frac{ay - a\tau b + f}{2},$$

$$E\left[\hat{S}_T^{0,f}\right] = E\left[\hat{S}_N^f\right] = \frac{y-b+f}{2}.$$

Letting J denote the value of the vacancy to the firm minus the firing cost, we obtain:

$$J_T^{i,f} + f = (1-\beta)\, S_T^{i,f}, \;\; i \in (0,1).$$

In case of separation, $J_T^{i,f} = 0$ and the firm pays the firing cost f. Writing down the expression for the expected firm value we obtain:

$$J_T^{1,f} = (1-\beta)(ay - a\tau b + f) - \gamma(1-\beta)\frac{a^2y^2 - (a\tau b - f)^2}{2ay} - f,$$
$$J_T^{0,f} = J_N^f = (1-\beta)(y-b+f) - \gamma(1-\beta)\frac{y^2-(b-f)^2}{2y} - f.$$

Note that although the cost f is paid only in case of separation, it still enters positively in the share of surplus accrued by the firm. As a result, when evaluating (1), f will appear as a constant in the expected value of posting either type of vacancy, and will affect the firm's decision only by altering the values of $E\left[S_T^f\right]$ and $E\left[S_N^f\right]$. In other words, only the difference ΔJ^f of the expected values (without the firing cost) are relevant for the firm's decision (recall that v_T^* only depends on the difference in the expected value of the two vacancy types):

$$\Delta J^f = eJ_T^{1,f} + (1-e)\, J_T^{0,f} - J_N^f = e\left(J_T^{1,f} - J_T^{0,f}\right). \tag{7}$$

Plugging the expressions for $J_T^{1,f}$ and $J_T^{0,f}$ into (7), we obtain the following equilibrium expression for ΔJ^f:

$$\frac{\Delta J^f}{e} = C_1 + \gamma(1-\beta)\frac{(a\tau b - f)^2 - a(b-f)^2}{2ay},$$

where C_1 is a constant that does not depend on f. Differentiating $\frac{\Delta J^f}{e}$ with respect to f yields:

$$\frac{d\Delta J^f}{df} \propto \gamma(2f - 2a\tau b - 2af + 2ab). \tag{8}$$

We are interested in the case of a small firing cost. We therefore evaluate (8) at

$f = 0$ to obtain:

$$\left. \frac{d\Delta J^f}{df} \right|_{f=0} \propto 2\gamma ab\,(1-\tau) > 0.$$

Also, it is straightforward to show that the cross derivative with respect to the degree of labor market turbulence is positive as long as $\tau < 1$ (i.e. skill transferability is not perfect):

$$\left. \frac{d^2\Delta J^f}{df\,d\gamma} \right|_{f=0} > 0.$$

In other words, the gain from adding moderate firing costs is larger when γ (turbulence) is higher. We can summarize these results as follows:

Proposition 2. *A small firing cost increases the share of vacancies of type T in the labor market for workers of type L. This marginal effect is larger, the larger the degree of labor market turbulence γ.*

The result also implies that the education wage premium will react less strongly to a change in turbulence if the firing cost is large. For educated workers, all vacancies allow for the accumulation of skills. An increase in turbulence affects educated workers through the direct impact of negative productivity shocks, but there are no add-on effects on investment in relationship-specific capital. By contrast, for workers with less education an increase in turbulence implies that fewer opportunities for accumulating skills are available, which lowers average wages and hence increases the education premium. The downward effect on unskilled wages and upward effect on the education premium are less pronounced when there are firing costs.

To illustrate this point, in Figure 8 we display the relationship between the degree of turbulence and the education premium for the cases of no firing cost ($f = 0$) and a moderate firing cost ($f = 0.3$).[1] In the case of the higher firing cost, the relationship is noticeably flatter. The same qualitative relationship also arises in the quantitative model discussed in the next section. There, these effects are even more pronounced, because worker's investment in skill is also endogenous, and investment in relationship-specific capital takes place over many periods.

[1] We set $y = 1$, $a = 2$, $\tau = 0.75$, $e = 0.8$, $b = 0.65$, and $\beta = 0.65$.

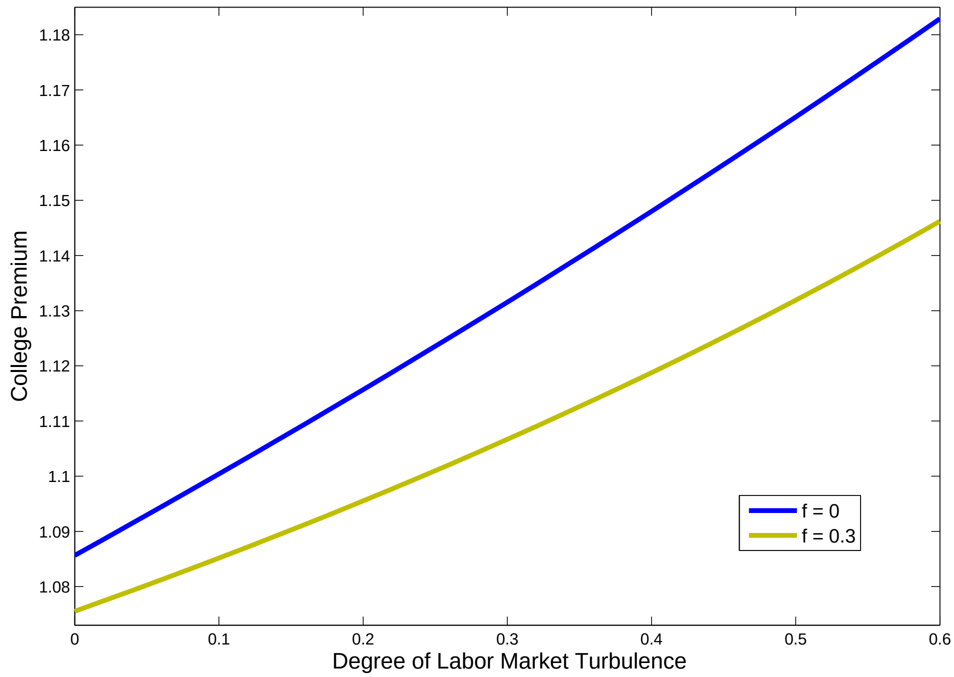

Figure 8: Effect of firing cost on college premium for different degrees of labor market turbulence

Returning to the simple model, introducing the firing cost also affects welfare. In principle, introducing a firing cost can reduce welfare, since it alters separation decisions taken by workers and firms. We assume that revenues from firing costs are redistributed as lump sum transfers, so that there is no direct welfare loss from the firing cost. We can however compute the loss associated with the distortion in the separation decisions. As we show in the Appendix, we find that the distortion in separations has no first order impact on output.

Given risk neutrality, welfare is defined as output minus the cost of vacancy posting. We focus here on the L market (the case for the H labor market is analogous, once we impose $\tau = 1$ and $v_T = 1$). Welfare is given by:

$$W_L^f = v_T e_L Y_{L,T}^{1,f} + \left[v_T\left(1 - e_L\right) + \left(1 - v_T\right)\right] Y_{L,T}^{0,f} - \kappa_T v_T^2, \tag{9}$$

where $Y_{L,T}^{1,f}$ and $Y_{L,T}^{0,f}$ is the expected output in case of successful effort and unsuccessful effort (or N vacancies), respectively. Plugging the equilibrium values of v_T and Y into (9), differentiating with respect to f, and evaluating at $f = 0$ yields

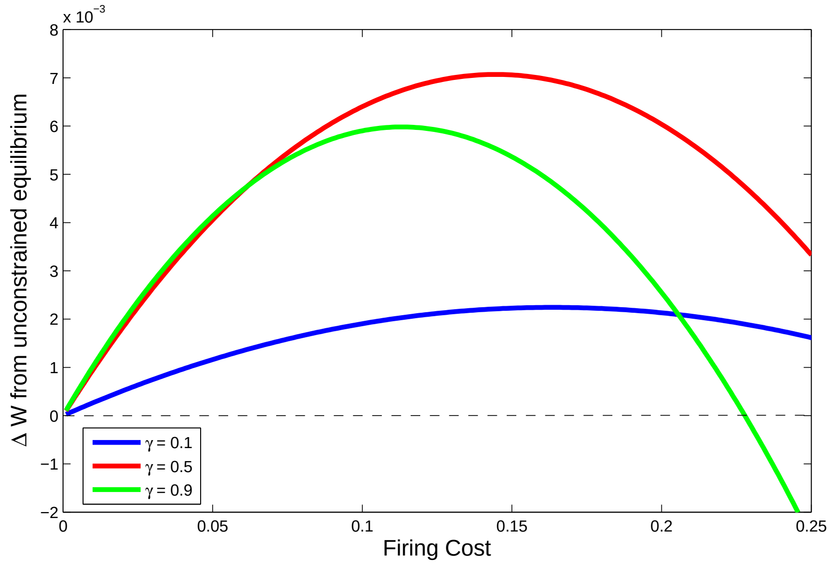

Figure 9: Welfare effect of firing costs for different degrees of labor market turbulence

the following result:

Proposition 3. *A small firing cost has a positive effect on welfare.*

To illustrate the overall effect of firing costs on welfare (i.e. for higher levels of the firing cost), we consider a simple numerical example. We set $y = 1$, $a = 2$, $\tau = 0.75$, $e = 0.8$, $b = 0.65$, and, consistently with our parametrization of the dynamic model in Section 4, $\beta = 0.65$.

Figure 9 shows the impact of firing costs on welfare for different values of the turbulence parameter γ. For low values of γ, the distortion both in the vacancy composition and in the separation decision is low, and the impact of the firing cost on welfare is small. For high values of γ, the welfare loss in the vacancy composition is high enough to generate large welfare gains for small values of f. However, the distortion in the separation decision ultimately kicks in: welfare is maximized at a firing cost equal to 12% of output. The impact of the firing cost on welfare is the highest for intermediate values of γ, where the introduction of firing costs improves the efficiency of vacancy allocations with a moderate negative effect on the efficiency of separations.

4 Dynamic Model and Quantitative Findings

In this section, we examine the quantitative importance of the mechanism developed in the previous section by calibrating a dynamic model to match salient features of the U.S. labor market between 1980 and 2010. We use the calibrated model to quantify the extent to which the mechanism can account for the evolution of average worker tenure by education and tenure and education wage premia. We also explore how labor protection regulation can help explain the different reactions of the U.S. and European economies to rising turbulence starting in the 1980s.

4.1 Setting

Consider an environment with mass M_L of low-education workers and mass M_H of college-educated workers. Agents start working at the age of 30, become old with probability π_o in each period and, once old, retire with probability π_r. We set π_o and π_r so that agents become old on average at age 45 and retire at age 60. We normalize the utility from retirement to zero. This perpetual youth structure yields a computationally tractable model while preserving a life-cycle structure.

Utility is linear and workers are risk neutral. Workers discount the future with the probabilities of aging and retiring. Workers' instantaneous utility is given by their wage or home production, minus the cost of effort exerted on acquiring skills on the job.

Each worker is born with a productivity level $a \in \{a_1, \ldots, a_n\}$ drawn from a probability distribution $F(\cdot)$. Unemployed workers find the labor market with probability λ. Upon finding the labor market, they are randomly matched with a vacancy. Workers of type H are only matched with jobs that allow for on-the-job skill accumulation. Workers of type L are matched with jobs that allow for skill accumulation with probability v_T, and with jobs that do not allow for skill accumulation with probability $v_N = 1 - v_T$. We model vacancy posting in a reduced form way, by postulating that type-T vacancies are a linear function of the difference in the expected returns from T and N vacancies, $\Delta E[J]$:

$$v_T = \bar{v_T} + c_1 \left(\Delta E[J] - c_0\right) \tag{10}$$

where c_0 and c_1 are constants.

Upon matching, if the surplus of the match is positive, the firm and the worker negotiate a wage, with the worker retaining a share β of the resulting surplus. The wage can be renegotiated in any period without additional costs. If at any point in the employment relationship the surplus becomes negative, separation occurs. Upon separation, the firm pays a firing cost $f \geq 0$. We set $f = 0$ in the benchmark model.

After separation, a worker spends one period as unemployed before attempting again to find the labor market (which happens with probability λ). An unemployed worker with low education consumes b_L, and a college-educated unemployed worker consumes $b_H = A_H b_L$ with $A_H > 0$.

4.2 Production and Turbulence Shocks

A match between a firm and worker of skill level a produces output $y_L(a) = a$ if the worker is of type L, and $y_H(a) = A_H a$ if the worker is of type H.

Before production takes place, the match is hit by a turbulence shock with a probability that depends on the education of the worker. If the shock hits, productivity of the match is reduced by a factor $\epsilon \sim$ Uniform $(0, 1)$. This low-state productivity lasts until separation, or until a positive shock hits the match, bringing productivity back to normal. This shock occurs with probability $p(\epsilon)$, where $p(\cdot)$ is an increasing function of ϵ.

Turbulence shocks are the main source of separations in this model. In order to generate quantities that are comparable with the ones in the data, we also keep track of the level of seniority of a worker within his current employment spell. In particular, we assume that a worker starts a job as a junior (jr) and transitions to seniority (sr) with probability π_{sr}. We set π_{sr} so that a worker is with the same employer for 15 years on average before acquiring seniority.

4.3 Skill Accumulation and Skill Loss

Workers in jobs that allow for skill accumulation choose how much effort to exert in accumulating skills. Exerting effort e costs e^2 and improves a worker's skills

with probability $\frac{e}{e+1}$. If effort is successful, a worker with skill level a_i upgrades her skills to a_{i+1}.

Upon separation from any job, workers experience a downgrading of skills, so that a worker of education $edu \in \{H, L\}$ and skill level i transitions to skill level j with probability $Q_{edu}(i, j)$. We assume that Q is such that $Q(i, j) = 0$ for each $j > i$. We also assume that for each i, Q_H first order stochastically dominates Q_L, so that the skill loss upon separation is more severe for L workers. This assumption captures the idea that part of the advantage of high education in the labor market is giving access to skills that are more transferrable across different employment spells, as argued in Section 2.

4.4 Functional Forms and Parameter Values

We would like explore the quantitative implications of changing labor market turbulence in our model. To do this, we let the turbulence parameters γ_{edu}^{year} vary over education levels, $edu \in \{H, L\}$, and time, $year \in \{1980, 2010\}$. The parameters are chosen to match the observed decline in long-term seniority observed in the United States over the same period for each education group. The other model parameters are held constant.

The distribution of productivities at birth is given by a Pareto distribution with a minimum at $a_{min} = 1$ and shape parameter $\eta = 5$. The job finding rate λ is fixed at 0.8 to obtain an average duration of unemployment of three months. Home production is set to 40% of the minimum of the support of the skill distribution. The share of surplus captured by the worker, β, is set equal to 0.65.

For the transition probabilities controlling skill loss upon separation, we assume a functional form that captures the idea that skill transferability across employment spells is higher for H than for L workers. The following functions make the ex-post distribution of productivities following separation for a worker of skill a_i more compressed towards a_i for H workers and more dispersed towards lower values of a for L workers:

$$Q_H(i, j) = 0.5^{(i-j)+1} 1(j \leq i),$$
$$Q_L(i, j) = \frac{2j}{(i+1)(i)} 1(j \leq i).$$

We also set the probability $p(\epsilon)$ of match productivity recovering following a turbulence shock to be increasing in ϵ:

$$p(\epsilon) = \epsilon^2.$$

This functional form captures the idea that more severe productivity shocks are also more likely to be persistent.

4.5 Quantitative Exploration

In our quantitative exercise, we explore the ability of the model to account for the empirical trends observed in the US economy between 1980 and 2010. We then analyze the consequences of introducing a firing cost for inequality and welfare.

The key parameters to be calibrated are the degree of labor market turbulence in 1980 and 2010 for both H and L labor markets; the parameters controlling the policy function for vacancy posting (10); and the intrinsic productivity advantage of H workers A_H.

We set parameters such that in 1980 we have $v_T^{80} = 1$, i.e., all vacancies in labor market L allow for on-the-job skill accumulation. We also set $\bar{v}_T = v_T^{80}$, which implies $c_0 = \Delta E[J^{80}]$. [2] Then, we calibrate γ_L^{80} to match the share of non-college educated workers at ages 45–60 with seniority of 15 years or more. This share is 54% in 1980, which implies a value of $\gamma_L^{80} = 0.10$. We jointly calibrate A_H and γ_H^{80} to match the skill premium (44%) and the share of college-educated workers with seniority of 15 years more (60%) in 1980. This procedure yields $\gamma_H^{80} = 0.045$ and $A_H = 1.11$.

In 2010, we jointly calibrate γ_L^{10} and c_1 to match the share of senior workers in 2010 and the change in the seniority premium between 1980 and 2010. This implies $\gamma_L^{10} = 0.26$ and $c_0 = 0.61$, yielding a share of vacancies with returns to experience of $v_T^{10} = 0.55$. This sharp increase in the degree of labor market turbulence and

[2] The reason for this choice is that, under the current parametrization, the model generates the maximum seniority premium at $v_T = 1$, but this premium is still lower than the one observed in the data (36% in the model, 44% in the data). Hence, we impose $v_T^{80} = 1$ and calibrate c_1 in equation 10 to match the change, rather than the level, of the seniority premium between 1980 and 2010.

	Data	Model	Parameter	Value
1980 College premium	44%	44%	A_H	1.11
1980 L Senior Share	54%	54%	γ_L^{80}	0.10
1980 H Senior Share	60%	60%	γ_H^{80}	0.045
2010 L Senior Share	43%	43%	γ_L^{10}	0.26
2010 H Senior Share	45%	45%	γ_H^{10}	0.135
1980-2010 L Seniority premium change	-10%	-10%	c_1	0.61
2010 College premium	96%	94%	Non-targeted	

Table 2: Model calibration

decrease in the share of vacancies with returns to experience rationalize the steep decline in the share of workers with high seniority (from 54% to 43%) and the seniority premium (from 44% to 30%) recorded in the data between 1980 and 2010. Analogously, we set the degree of labor market turbulence for college graduates in 2010 to match the corresponding share of senior workers (44%), which delivers $\gamma_H^{10} = 0.135$. Table 2 summarizes the calibrated parameters and the targeted moments.

4.6 Model Experiments and Decompositions

The ability of the model to capture the response of the college premium to a change in turbulence can be gauged from Table 2. Notably, the model accounts almost fully for the increase in the college premium from 44% to 96%, even when the intrinsic productivity shifter, A_H, is kept constant at its calibrated 1980 value. In the calibration, this moment was not directly targeted; rather, the rise in the college premium follows from matching the returns to tenure together with matching changes in average tenure by education in the data.

Next, we provide a decomposition to illustrate the three main components behind the increase in the college wage premium. First, the increase in the degree of labor market turbulence, inferred by the decline in the share of workers with long term seniority, weakens workers' incentives to invest in skill accumulation,

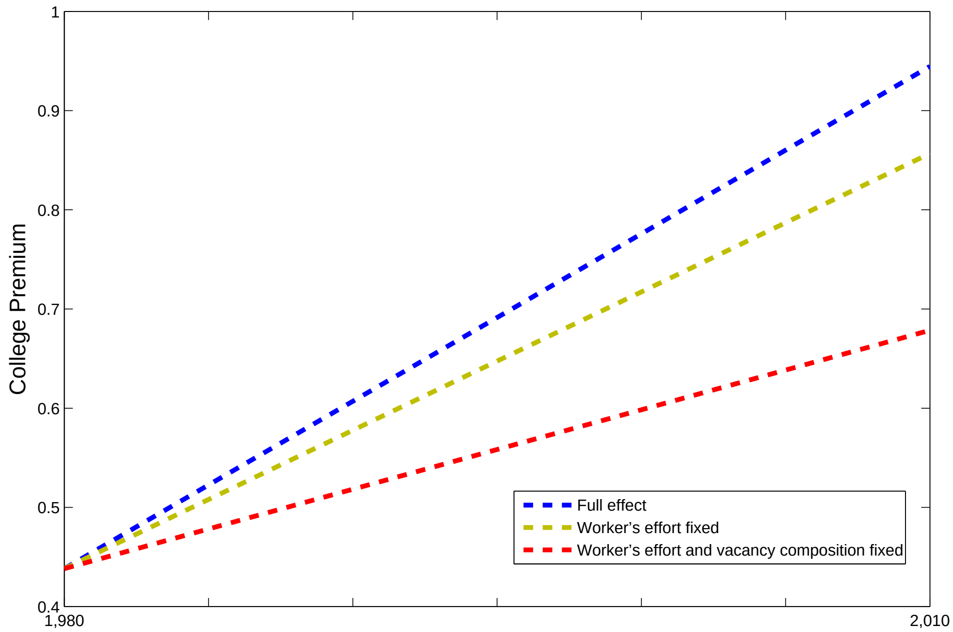

Figure 10: Decomposition of effect of labor market turbulence on college wage premium

because of imperfect transferability. This negative effect is more pronounced for L- than for H-type workers, because of the different degrees of skill transferability for less- and more-educated workers. Second, the labor market for L workers experiences a shift toward vacancies that do not allow for on-the-job skill accumulation. Third, even keeping vacancies and worker's effort constant, an increase in the rate of job destruction will damage L workers more, since low skill transferability and more frequent job transitions imply a more severe loss of skills.

To evaluate the contribution of each component, we first simulate the model holding workers' effort fixed at the 1980 level, while changing the composition of vacancies and labor market turbulence to the calibrated 2010 levels. The results are shown in Figure 10: The college premium increases to 85%, suggesting that roughly 18% of the overall effect is due to a change in worker's incentives to invest in skills. Next, we shut down jointly the vacancy composition and the workers' effort channels. The college premium in 2010 is now 67%, suggesting

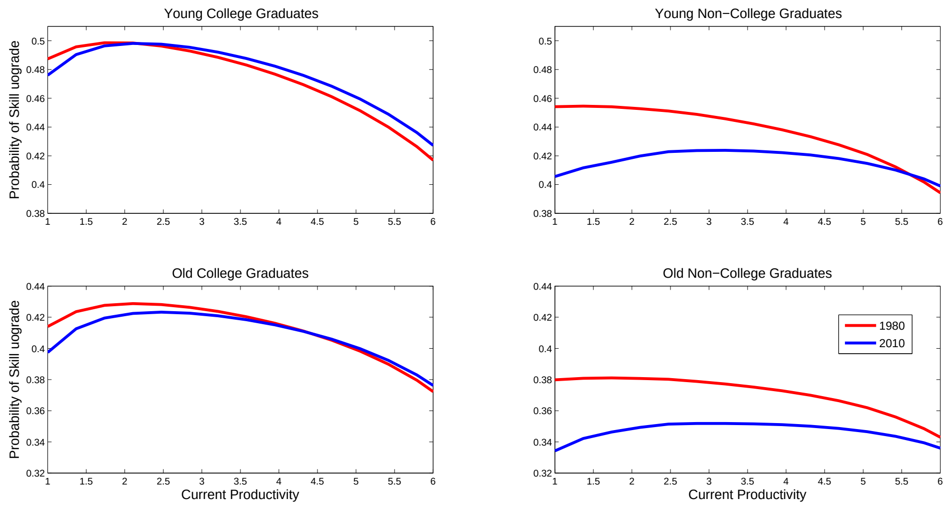

Figure 11: Effect of labor market turbulence on investment in skills

that roughly 36% of the overall effect is due to the change in vacancy composition. The remainder, roughly 23 percentage points, or 46% of the overall effect, is due to the direct effect of more frequent skill loss for L workers after the increase in turbulence.

To shed light on the underlying mechanism, Figure 11 shows the impact of the increase in labor market turbulence on the effort exerted by each type of worker in accumulating skills. In the left panels, we show the probability of skill upgrading for young (top) and old (bottom) H workers for a given current level of skill, both in 1980 and 2010. The two lines almost perfectly overlap: because of their high skill transferability, an increase in labor market turbulence has a marginal impact on their investment decisions. By contrast, the right panels show how the probability of skill upgrading for L workers in T vacancies changes between 1980 and 2010. The increase in labor market turbulence has a substantial negative impact on the probability of skill upgrading for L workers with low current skills, which (given the steady state skill distribution) affects most L workers. The average probability that an L worker in a job of type T upgrades skills in a given period drops from 44.2% to 41.5% for young workers and from 36.7% to 34.6% for old workers.

	College Premium	Output
No LPL	94%	0
LPL in both markets	91%	+0.34%
LPL in L market	88%	+0.34%

Table 3: Effect of labor protection legislation on skill premium and output

4.7 The Role of Labor Protection Legislation

We now evaluate the effect of introducing a firing cost, in the form of a tax f the firm has to pay in case of separation. A common argument in the policy debate is that labor market turbulence should make labor protection legislation less desirable from a welfare viewpoint, since separations can be severely distorted. As already discussed in the static model, this argument potentially neglects firms' and workers' incentives to invest in relationship-specific capital. The inefficiencies in this process can in fact deteriorate when labor market turbulence increases.

In our experiment, we introduce a firing cost equal to an L worker's average annual wage in the high-turbulence economy (i.e., in 2010). We consider two cases: in the first case, the firing cost applies to both labor markets, while in the second case, it only applies to the labor market for L workers.

Table 3 shows the result of this policy experiment. In both cases, output increases by 0.34% compared to the economy with no firing costs. This happens because the effect of introducing the firing cost in the H labor market is negligible: the improvement in the investment decisions and the distortion on separations offset each other almost perfectly. Hence, the output improvement originates entirely from the L labor market. However, as we would expect, between-group inequality responds differently in the two cases: the college premium drops to 91% when the firing cost applies to both labor markets, and to 88% when it only applies to L workers.

It is important to ask whether the rise in output is due to the decrease in the rate of separations and subsequent skill loss, or to the improvement in the incentives of firms and workers to invest in accumulation of relationship-specific capital.

Vacancies of type T in the labor market for low education workers increase from 55% to 57.2%. Moreover, when vacancies and worker's effort are kept at the 2010 level in the unconstrained case (when $f = 0$) output declines by 0.03%, suggesting that the rise in output is driven entirely by the change in firms' and workers' incentives.

5 Conclusion

There are large differences in employment protection across countries. In Europe, insider-outsider labor markets and protection of senior workers at the expenses of junior and temporary workers are common. In this paper, we argue that this dimension has important implications for investment in job-specific skills. This fact can help explain the diverging inequality trends between countries with tight labor market regulation and countries with low levels of employment protection.

In our analysis, we have focused on match-specific investments that improve the productivity of a given worker-firm pair. An interesting extension would be to also consider more general investment by firms in technologies that are complementary to workers' accumulated firm-specific skills. That is, some firms' production technology relies on having experienced workers, whereas other technologies work equally well with inexperienced workers. In a model of directed technological change along the lines of Acemoglu (2002), the incentive to develop technologies that work well with experienced workers would be higher if (because of labor regulation) the firm is more likely to have many such workers in the future. The direction of technical change in the context of a search model has previously been considered by Michelacci and Lopez-Salido (2007), but not from the perspective of the skill premium in the labor market. Considering the role of labor market regulations in models of endogenous directed technological change is an important challenge for future research.

References

Acemoglu, Daron. 2002. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature* 40 (1): 7–72.

———. 2003. "Cross-Country Inequality Trends." *Economic Journal* 113:F121–F149.

Aghion, Philippe, Peter Howitt, and Giovanni L. Violante. 2002. "General Purpose Technology and Wage Inequality." *Journal of Economic Growth* 607:315–345.

Alon, Titan. 2017. "Earning More by Doing Less: Human Capital Specialization and the College Wage Premium." Unpublished Manuscript, Northwestern University.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4): 1279–1333.

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

Bertola, Giuseppe, and Andrea Ichino. 1995. "Wage Inequality and Unemployment: United States versus Europe." In *NBER Macroeconomics Annual*, edited by Ben S. Bernanke and Julio J. Rotemberg, 13–66. MIT Press.

Blanchard, Olivier J., and Lawrence H. Summers. 1986. "Hysteresis and the European Unemployment Problem." In *NBER Macroeconomics Annual 1986*, edited by Stanley Fischer, 15–90. MIT Press.

Gathmann, Christina, and Uta Schönberg. 2010. "How General Is Human Capital? A Task-Based Approach." *Journal of Labor Economics* 28 (1): 1–49.

Krueger, Dirk, Fabrizio Perri, Luigi Pistaferri, and Giovanni L. Violante. 2010. "Cross-sectional Facts for Macroeconomists." *Review of Economic Dynamics* 13 (1): 1–14. Special issue: Cross-Sectional Facts for Macroeconomists.

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

Michelacci, Claudio, and David Lopez-Salido. 2007. "Technology Shocks and Job Flows." *Review of Economic Studies* 74 (4): 1195–1227.

Nickell, Stephen. 1997. "Unemployment and Labor Market Rigidities: Europe versus North America." *The Journal of Economic Perspectives* 11 (3): 55–74.

Wasmer, Etienne. 2006. "General versus Specific Skills in Labor Markets with Search Frictions and Firing Costs." *American Economic Review* 96 (3): 811–831.

A Proof of Proposition 3

In this Appendix, we derive analytically the effect on welfare of introducing a firing cost $f \geq 0$, and show that this effect is always positive for small f. We will look at welfare in labor market H and L separately.

In labor market H, welfare is equal to output:

$$W_H^f = eY_H^{1,f} + (1-e)Y_L^{0,f}. \tag{11}$$

Output in case of successful effort can be written explicitely as:

$$Y_H^{1,f} = (1-\gamma)ay + \gamma\left(\frac{ay-ab+f}{ay}\right)\frac{ab-f+ay}{2} + \gamma\left(\frac{ab-f}{ay}\right).$$

This can be rewritten as:

$$Y_H^{1,f} = C_{Y,H}^{1,f} - \gamma\frac{f^2}{2ay},$$

where $C_{Y,H}^{1,f}$ is a term that does not depend on f, and $Y_H^{0,f}$ is analogous, once we impose $a = 1$. It is immediate to see that the firing cost has no first-order effect on output in the labor market for workers of type H.

Welfare in the L labor market depends on the share of vacancies of type T, and on the associated cost of posting them:

$$W_L^f = v_T eY_{L,T}^{1,f} + \left[v_T\left(1-e\right) + \left(1-v_T\right)\right]Y_{L,T}^{0,f} - \kappa_T v_T^2, \tag{12}$$

where:

$$v_T = \frac{e\Delta J^f}{2\kappa_T}.$$

Letting $\Delta Y_L^f = Y_{L,T}^{1,f} - Y_{L,T}^{0,f}$, differentiating (12) with respect to f yields:

$$\frac{dW_L^f}{df} = \frac{e^2}{w\kappa_T}\left\{\frac{d\Delta J^f}{df}\Delta Y_L^f + \Delta J^f\frac{d\Delta Y_L^f}{df}\right\} - \frac{\gamma f}{y} - \frac{e^2}{2\kappa_T}\Delta J^f\frac{d\Delta J^f}{df}. \tag{13}$$

The partial derivaties of ΔJ^f and ΔY_L^f with respect to f are:

$$\frac{d\Delta J^f}{df} = \gamma\frac{ab(1-\tau) - f(a-1)}{ay},$$
$$\frac{d\Delta Y_L^f}{df} = \gamma f(a-1)ay.$$

Plugging these expression into 13 and evaluating at $f = 0$ gives:

$$\frac{dW_L^f}{df}|_{f=0} = \frac{e^2\gamma b(1-\tau)}{2\kappa_T y}\left[\Delta Y_L^f - \Delta J^f\right]$$

which is always positive since, at $f = 0$:

$$\Delta Y_L^f - \Delta J^f = (1-\gamma)\beta y(a-1) + (1-\beta(1-\gamma))b(a\tau - 1).$$

This proves Proposition 3.