Worker protection and heterogeneous match quality

Saman Darougheh and Gustaf Lundgren*

This version: October 2, 2019

Abstract

We study the impact of worker protection in an environment with heterogeneous match productivity and a constrained wage setting. Firms can either employ costly screening to determine the match quality or hire workers out of their applicant pool at random, learn about the match quality, and disengage from bad matches. Thus, layoff protections intervene with a firm's ability to screen matches. In our calibration, a policy that prevents layoffs reduces unemployment and increases consumption in the new steady state. However, the economy becomes more susceptible to productivity shocks. Two additional channels transmit productivity shocks when layoffs are regulated. First, the value of hiring at random is more volatile when separating bad matches is no longer an option. Second, additional screening in recessions worsens the composition of the pool of unemployed workers. Consequently, recessions will be long lasting and hiring is lower even after the TFP shock has receded. We conclude that economies potentially have a higher average output under layoff regulations, but that this comes at the cost of higher volatility and jobless recoveries. (JEL D83, J63, J64)

*Saman Darougheh: PhD Candidate, Institute for International Economic Studies, Stockholm University. Gustaf Lundgren: Sveriges Riksbank. This chapter builds on joint work with Gustaf Lundgren. Parts of this work has been previously published as two separate chapters in Gustaf Lundgren's thesis "Essays on job market screening, in-group bias and school competition". The first of these chapters is entitled "A search model with multiple applications" and considers the impact of heterogeneous match quality. The second chapter is entitled "Ranking, unemployment duration and unemployment volatility" and considers the interaction between screening and business cycle fluctuations. We are indebted to advice from Per Krusell, Lars Ljungqvist, Kurt Mitman, Oskar Nordström Skans and seminar participants at the Stockholm School of Economics, Stockholm University, and the Oslo-Bi-NHH Workshop in Macroeconomics.

1 Introduction

In 2008, many countries were affected by the so-called Great Recession when a financial crisis and a collapse in housing prices led to a global productive slowdown and a worsening of the labor markets. Among them, southern European countries were hit particularly hard. While it is true that they were particularly exposed to the initial triggers of the recession, their labor markets are also highly regulated (Karamessini, 2008; Hassel, 2014), and some have faulted these rigid labor markets for the slow recoveries.

Since the seminal work of Lazear (1990), an extensive literature has analyzed the theoretical and empirical links between employment protection and the functioning of labor markets. In this paper, we revisit the impact of worker protection on aggregate output and unemployment in the presence of worker-firm pairs that differ in their match productivities. Workers and firms are both risk-neutral. Given our assumptions, aggregate consumption is an appropriate measure of steady-state welfare. The model uncovers a potential mean-volatility trade-off of protecting employment: in our simulation, the policy will improve aggregate consumption but render the economy more volatile to aggregate shocks.

Worker protection has some bite whenever it is not undone by Coasean transfers (Lazear, 1990) and layoffs are desired by the firm either in equilibrium, or as an off-equilibrium threat. To satisfy the second requirement, most of the literature focuses on shocks to a firm's profitability. For example, worker protection tampers with a firm's ability to respond to a recession by reducing size. Instead, we analyze the extent to which worker protection interferes with the screening potential of layoffs. Specifically, firm-worker matches differ in productivity: "good" matches are more productive than "bad" matches. Moreover, workers persistently differ in the probability of drawing a good match: workers of the "high" type are more likely to have a good draw than the "low" type. Wage contracts are constrained: they cannot discriminate by match-specific productivity. Here, wages do not offset the differences in productivity, and bad matches are not profitable for the firm. Firms observe match quality in the hiring stage only if they engage in the costly screening. After hiring, the match productivity is revealed instantly. Without employment protection, firms could forego costly screening of candidates and instead fire them once their match quality is revealed. Therefore, we analyze worker protection in a novel context where it interacts with an average match quality. After the introduction of employment protection, firms that hire without screening their applicant pool can no longer fire workers with a bad match quality. This affects aggregate consumption through two main channels. First, some firms will still hire without screening. Consequently, the average match quality in the economy decreases as non-screening firms employ workers but can no longer disengage from bad matches. Second, the policy decreases the value of a match. Consequently, the hiring intensity falls initially when the policy is introduced. Under our benchmark calibration, the welfare gains associated with such a policy are still positive: aggregate consumption is higher at the new steady state and during the transition thereto.

Next, we compare the economy's response to productivity shocks in the steady states with and without worker protection. Importantly, the ability to lay off workers puts a floor on the value of randomly hiring. The employment protection removes that floor and increases the pro-cyclicality of the value of hiring. Consequently, aggregate unemployment and consumption fall more in recessions when firms' ability to fire workers is reduced. In the recession, fewer firms hire without screening. Over the course of the recession, low-type workers have a relatively harder time when becoming unemployed, and the quality of the pool of unemployed workers deteriorates. When aggregate productivity recovers, the quality of the applicant pool is persistently lower, thus inducing a lower hiring rate: the recession leads to a jobless recovery.

Literature This paper builds on and is motivated by a large microeconomic literature that analyzes statistical discrimination and screening in the hiring process. Kroft, Lange, and Notowidigdo (2013) and Ghayad (2013) show that employers use unemployment duration to screen for persistent ability. Motivated by these empirical findings, Jarosch and Pilossoph (2019) argue that discrimination at the interview stage is irrelevant if it only arises for candidates that would ultimately not have been hired anyway. In our model, these candidates would have been laid off right after employment. Masters (2014) finds that statistical discrimination in the screening process can be self-fulfilling: discrimination against a group can worsen their unemployment pool and thereby solidify that discrimination. Josephson and Shapiro (2016) analyze the impact of screening in an environment where individuals have private information on their own type.

When placed in the context of aggregate fluctuations, Bertola (1999) argues that employment protection reduces layoffs but also hirings, with an ambiguous effect on total unemployment. Another theoretical prediction is that such protections will dampen the layoffs in recessions, but also employment in booms. In line with our model, Lindbeck (1993) argues that employment protections can increase the volatility of the value associated with hiring: in a recession, increased uncertainty could induce firms to choose not to hire at all, thus leading to long-lasting recessions. Lindbeck and Snower (2001) argue that worker protections strengthen the position of insiders within the firm and thus lead to wage gains at the cost of outsiders. O. Blanchard and Landier (2002) and Cahuc and Postel-Vinay (2002) analyze the impact of weakening temporary employment. Autor, Kerr, and Kugler (2007) measure the impact of employment protection on productivity in the context of the "wrongful discharge protection" in the United States. Dal Bianco, Bruno, and Signorelli (2015) estimate the role of labor market institutions in the European unemployment response.

Modeling-wise, our theory has the potential to generate large and persistent unemployment responses to a TFP shock. Here, we contribute to a large literature which has studied the volatility of unemployment since Merz (1995), Andolfatto (1996), and Shimer (2005), often referring to it as the "unemployment-volatility puzzle". The persistence of unemployment has received new attention since the so-called Great Recession of 2008, when the recovery did not coincide with a proportional reduction in unemployment. With respect to explaining jobless recoveries, Acharya and Wee (2019) performs a similar exercise where he shows that increased screening in a recession leads to large and persistent unemployment responses. The two papers vary in their source for additional screening: in our model, it stems from the number of applicants and a convex return to screening, while their model centers around rational inattention.

The remainder of this paper is as follows. Section 2 introduces our model. We analyze the introduction of employment protection in section 3, where we observe the economy's transition to a new steady state. In section 4, we compare the aggregate shocks in both steady states. Section 5 concludes the paper.

2 Model

In the model, time is continuous. The economy is populated by workers, firms, and vacancies. Firms post vacancies and collect applications from unemployed workers. These unemployed workers differ in the value they provide to the firm. Firms can employ costly screening for a vacancy to find out the value of all applicants to the firm. It can then decide to hire at most one of them. In the remainder of this section, we will lay out these parts in detail.

2.1 Matching

Vacancies on the labor market are opened subsequently and these remain open for a fixed duration τ. Each unemployed worker sends out a applications in each time window of length τ. Each application randomly arrives at one vacancy that is open during the time of submission. When a vacancy closes, it inspects all applications it received during the opening window of the vacancy of length τ. The appendix shows that the number of applications at each vacancy is Poisson distributed with mean $\lambda \equiv \frac{au}{v}$, where u is the mass of unemployed workers, and v the mass of vacancies. If u and v do not change during the window of length τ, the distribution of applications over vacancies is exact. This will be true in a steady state. We use this expression to approximate the distribution also outside of steady states.

2.2 Match-specific productivity

The canonical urn-ball model described so far can only match empirical job-finding rates when the unemployed workers apply to vacancies at a very low rate. That low application rate then implies that the vacancies receive very few applications, which limits the scope of screening and discrimination in the hiring process. From the view point of the vacancy: an empirically measured *monthly* vacancy filling rate around $0.6 - 0.8$ in a model with homogeneous vacancies and urn-ball matching implies that a sizable fraction of vacancies receives no applicant. The applicants are Poisson-distributed over the vacancies: a sizable fraction of vacancies not matching with an applicant is only possible if vacancies on average receive around one applicant.

This small number of applications per vacancy is not only inconsistent with the empirical estimates, but also affects the incentives for firms to screen and discriminate. For example, O. J. Blanchard and Diamond (1994)discuss the implications of discrimination by unemployment duration for wage dispersion and job-finding rates. In their calibration, the average number of applicants per vacancy is small. Moreover, business-cycle changes in market tightness lead to small variations in the number of applications per vacancy. They compare an economy in normal times with an unemployment rate of 5% to a recession with an unemployment rate of 10%. Despite calibrating to such a strong recession, their model only generates an increase in the average applications per vacancy from 0.7 to 1.22: the resulting impact of screening and discrimination does not vary a great deal with the business cycle.

We extend the urn-ball matching model to address this calibration dilemma. We envision the productivity of a worker at a firm to be a component of a multitude of factors, one of which is specific to the worker-firm pair. Such a catch-all productivity-term could contain different amenities at the firm, and how important they are for the worker: some components could be flexible time vs fixed schedules, private offices vs. shared office space, strong top-down management vs. horizontal structures, or whether the worker and the manager match on a personal level. We think about the match-specific productivity as a continuous random variable. For modeling purposes, we will approximate that distribution with two values, good and bad: $A_g > A_b$. In this paper, we consider calibrations where hiring an employee with a bad match value is not profitable for the firm. The calibrated model can then match many applications per vacancy if a large share of the applications draws a low match-specific productivity that potentially does not lead to employment.

We nest this match-specific productivity with a worker-specific component: some workers have a higher probability of drawing a good match-specific productivity than others. We will call these workers of the "high type". This feature of the model matches the fact that in reality, some workers are quite adaptable to many environments: they could be productive in a multitude of settings, they are very socially skilled or contain more general human capital. These factors enable them to be more likely to be a good match at any particular firm. Denote by $\mathcal{P}_i$ the probability of a worker of type i to draw a *low* match-specific probability. Consequently, the high types are less likely to draw a low match value: $\mathcal{P}_h < \mathcal{P}_\ell$. We normalize the size of the labor force to 1, and denote the (fixed) mass of low-type workers as L_ℓ. The endogenous share of low-type workers among the unemployed will be denoted by P_u.

2.3 Wages

If the vacancy decides to hire an unemployed individual, the two form a one-worker firm. Total productivity in each match is a composite of the match-specific productivity A_i and an aggregate component $A(t)$.

Key in the model is that the heterogeneity in match-specific productivity is not offset by wages: firms pay the workers in good or bad matches the same. Such fixed task-specific wage schedules could be motivated by unions that prevent heterogeneous compensation within occupation-firm. We assume that the wages for all workers are given by

$$w(t) = \beta A_g A(t) + (1-\beta)b \tag{1}$$

where b denotes the consumption of the unemployed workers. This wage schedule can be rationalized for workers in good matches using Nash bargaining with workers' bargaining power β, where in case of (off-equilibrium) disagreement, both sides make a pause in the negotiation instead of disbanding the match. During that pause, the worker receives home production b, and the firm receives 0.

The worker-firm matches separate at an exogenous rate δ. The firms discount the future at rate ρ. We denote by J_i the value of a firm that has match-specific productivity i:

$$\rho J_i(t) = A_i A(t) - w(t) - \delta J_i(t) + \dot{J}_i(t) \tag{2}$$

As emphasized before, the cutoff A_b ensures that bad matches are not profitable to the firm.

Figure 1: Payouts of the screening decision

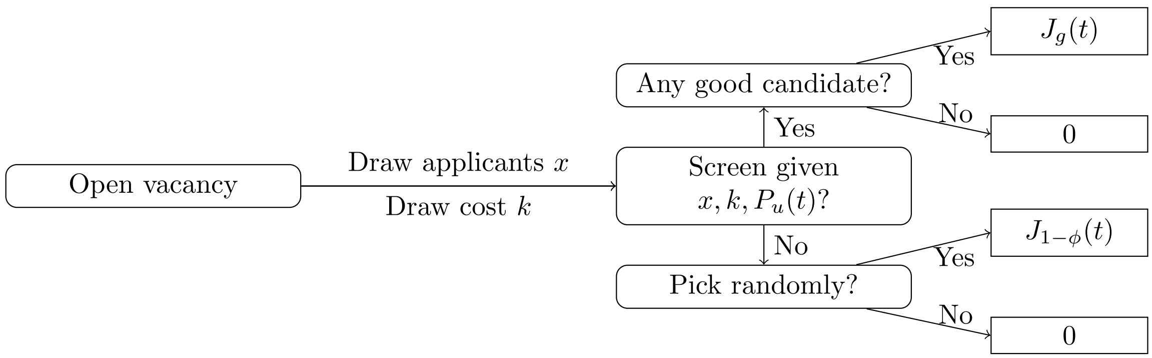

Assumption 1. *Match productivities are such that bad matches are not profitable.*

$$A_b A - \beta A_g A - (1-\beta) b \leq 0$$

We will hereinafter work with Assumption 1, which ensures that $J_b \leq 0$. There are many ways of approximating the productivity distribution such that J_b is negative. In the benchmark calibration, we will select A_b such that firms make exactly zero profits.

2.4 Job openings and screening

Now we present the details of the vacancy's problem. After having received x applications, the vacancy has to decide whether to screen the applications. Here, the cost of screening amounts to a fixed cost k.[1] After paying k, the vacancy learns about the match-specific productivity of every applicant. If one of the applicants is a good match, the firm will hire that person. If none of the applicants are a good match, the firm will refrain from hiring since bad matches are not profitable. If the firm posting the vacancy decides not to screen, it can either hire one applicant at random - the value of which we denote as $J_{1-\phi}$ - or not hire all. Figure 1 provides an overview over these choices and the associated outcomes. This model will contrast two policy regimes which differ in whether firms can lay off workers at will or not. $\mathcal{L}$ is a binary variable that takes the value of 1 whenever firms can lay off workers. Before setting up the vacancy's problem, it will be useful to establish two auxiliary variables.

[1] An earlier version of this paper incorporated a variable cost on top of the fixed cost. Variable costs add a great deal of complexity to the problem, but the results remain qualitatively similar.

Figure 2: Optimal screening choice and maximum profits

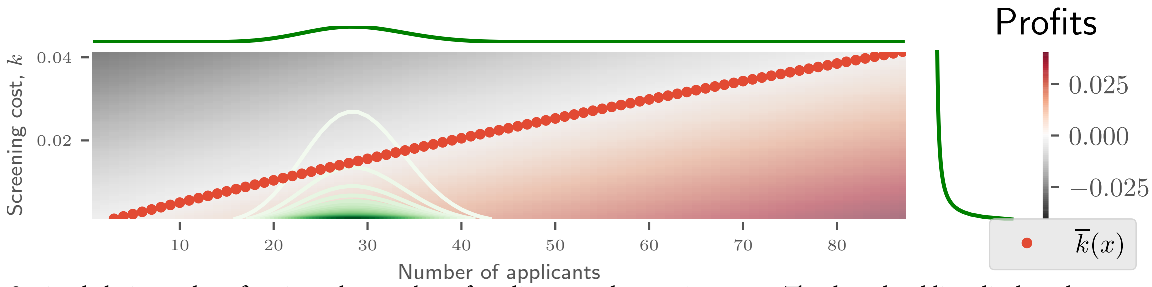

Optimal choice and profits given the number of applicants and screening costs. The dotted red line displays the curve where firms are indifferent between screening and not screening. Green contour lines display the joint-density of the distribution of firms across fixed-costs and applicants, with partial distributions displayed on the margins.

$J_{1-phi}(t)$ denotes the value of hiring at random in time t. With probability $\tilde{P}(t)$, the random draw has a low match value. When firms can fire at will, the downside risk of such a draw is limited to 0.

$$J_{1-\phi}(t) = \tilde{P}(t)\left[\mathcal{L}\cdot 0 + (1-\mathcal{L})J_b(t)\right] + (1-\tilde{P}(t))J_g(t)$$
$$\tilde{P}(t) = P_u(t)\mathcal{P}_\ell + (1-P_u(t))\mathcal{P}_h$$

With these in place, equations (3)-(5) compute the value of a vacancy with screening decision ϕ, x-number of applications and the screening cost k. If the vacancy decides to screen, it will hire if at least one application is of a good match quality. The probability of that happening is the complement to all of x applications being a bad match - which arises with probability $\tilde{P}^x$. If the vacancy does not screen, it can decide to either hire an unemployed individual at random - yielding $J_{1-\phi}(t)$ - or not hire at all, yielding 0. Figure 1 provides an overview of these choices and the associated outcomes.

$$\tilde{\pi}(\phi, x, k, t) = \phi\left[(1-\tilde{P}(t)^x)J_g(t) - k\right] + (1-\phi)\max\{J_{1-\phi}(t), 0\} \tag{3}$$
$$\phi(x, k, t) = \arg\max_{\phi\in[0,1]} \tilde{\pi}(\phi, x, k, t) \tag{4}$$
$$\pi(x, k, t) = \max_{\phi\in[0,1]} \tilde{\pi}(\phi, x, k, t) \tag{5}$$

The probability of finding at least one good match increases in the number of applications, while the cost of screening is invariant to the number of applications. Therefore, a vacancy with a given cost k will always screen if it receives a sufficiently large number of applications. Vacancies have heterogeneous screening costs: when opening a vacancy, the firm first pays a fixed vacancy cost c, and then draws the screening cost from a distribution with CDF $G^{\mathbf{k}}$. The number of applications that each vacancy receives is a random draw with density $g^{\mathbf{x}}$. This allows us to write the expected value of opening a vacancy V as in (6).

$$V(t) = -c + \int \sum_{x=1}^{\infty} \pi(x,k,t) g^{\mathbf{x}}(x) dG^{\mathbf{k}}(k) \tag{6}$$

It will be useful to define by $\overline{k}(x,t)$ the screening cost that makes firms indifferent between screening or not screening:

$$\overline{k}(x,t) \equiv \{k : \pi(0,x,k,t) = \pi(1,x,k,t)\}$$

It satisfies $\overline{k}(x,t) > 0 \quad \forall x > 0$. It is computed in (7) by collecting the ϕ terms in (3).

$$\overline{k}(x,t) = (1 - \tilde{P}(t)^x) - J_{1-\phi}(t) \tag{7}$$

We assume that the screening costs k are Gamma distributed with shape and scale parameters μ_k and σ_k. The number of applications are Poisson-distributed. The appendix uses this to show that we can write the value of a vacancy V as in (8).

$$\begin{aligned} V(t) &= -c + (1 - e^{-\lambda(t)}) J_{1-\phi}(t) \\ &\quad + \sum_{x=1}^{\infty} [J_h(t)(1 - \tilde{P}^x) - J_{1-\phi}(t)] g^{\mathbf{k}}(x) G^{\mathbf{k}}(\overline{k}(x,t)) - \mathcal{K}(t) \\ \mathcal{K}(t) &= \sum_{x=1}^{\infty} g^{\mathbf{x}}(x) \frac{\beta}{\Gamma(\alpha)} [\Gamma(1+\alpha, 0) - \Gamma(1+\alpha, \frac{\overline{k}(x,t)}{\beta})] \end{aligned} \tag{8}$$

Every vacancy has to pay the fixed entering cost c, and receives the value of a random draw $J_{1-\phi}$ if it draws more than zero applicants. Given x applications, screening yields the net value of $J_h(t)(1 - \tilde{P}^x) - J_{1-\phi}$ on top of that. The third term adds up this additional value, weighting it by the probability of drawing x applications $g^{\mathbf{x}}$, and the probability of drawing a screening cost k that is below $\overline{k}(x,t)$. Finally, $\mathcal{K}(t)$ computes the expected screening cost for the outcomes under which the firm posting the vacancy optimally decides to screen.

2.5 Flows

Now we will compute the job-finding rates for both high and low types of unemployed workers. We start by computing the probability that a firm posting a vacancy hires an unemployed individual of type i, denoted $o_i(t)$. If that firm screens, it finds a good match with probability $1 - \tilde{P}^x$, and ψ_h denotes the conditional probability of that hire being high type. The offer probability from a non-screening vacancy varies under the two policy regimes. A firm with a non-screening vacancy selects a high type with probability $1 - P_u$. If layoffs are not permitted, it will then hire that applicant if the value of randomly hiring is positive. If layoffs are permitted, the firm posting the vacancy will hire that applicant and keep him if he turns out to be a good match with probability $\mathcal{P}_h$. For individuals from the low-type pool, the offer probability is analogous.

$$
\begin{aligned}
o_h(t) &= \int \sum_{x=1}^{\infty} g(x)\Big[\phi(x,k,t)(1-\tilde{P}(t)^x)\psi_h(t) \\
&\quad + (1-\phi(x,k,t))(1-P_u(t)) \cdot \left[(1-\mathcal{L})\mathbf{1}_{J_{1-\phi}(t)\geq 0} + \mathcal{L}(1-\mathcal{P}_h)\right] dG^{\mathbf{k}}(k) \\
o_\ell(t) &= \int \sum_{x=1}^{\infty} g(x)\Big[\phi(x,k,t)(1-\tilde{P}(t)^x)\psi_\ell(t) \\
&\quad + (1-\phi(x,k,t))P_u(t) \cdot \left[(1-\mathcal{L})\mathbf{1}_{J_{1-\phi}(t)\geq 0} + \mathcal{L}(1-\mathcal{P}_\ell)\right] dG^{\mathbf{k}}(k) \\
\psi_h(t) &= \frac{(1-\mathcal{P}_h)(1-P_u(t))}{1-\tilde{P}(t)} \\
\psi_\ell(t) &= \frac{(1-\mathcal{P}_\ell)P_u(t)}{1-\tilde{P}(t)}
\end{aligned}
$$

We use these offer probabilities to compute the probability that any individual application of a type-i unemployed worker results in a job offer, which we denote μ_i. This is accomplished by employing an accounting identity of the matching function: the number of filled vacancies has to equal the number of hired unemployed workers.

$$
\begin{aligned}
v(t)o_\ell(t) &= a\mu_\ell(t)u(t)P_u(t) \\
v(t)o_h(t) &= a\mu_h(t)u(t)(1-P_u(t))
\end{aligned}
$$

Once we know the success probability of any individual application, we can compute the job-finding rates within each application cycle as the complement to all applications failing. We scale these rates up by τ to transform the job-finding rates from application-cycle length to period length.

$$f_i(t) = \tau(1 - (1 - \mu_i(t))^a)$$

The laws of motion for the unemployment rates can then be written as follows, where δ denotes the exogenous separation rates δ.

$$\dot{u_h}(t) = \delta e_h(t) - u_h(t) f_h(t)$$
$$\dot{u_\ell}(t) = \delta e_\ell(t) - u_\ell(t) f_\ell(t)$$

2.6 Equilibrium

Given regime status $\mathcal{L}$, initial $\{u_\ell(0), u_h(0)\}$, we can describe an equilibrium as a path $\{u_\ell(t), u_h(t)\}_{t>0}, \{J_i(t)\}_{i\in\{\ell,h\}}, v(t)\}_{t\geq 0}$ such that

1. $J_i(t)$ satisfies the Bellman equation (2) for all $t \geq 0$
2. Screening decisions solve the vacancy's problem for all x, k, t given $\mathcal{L}$, $\{u_\ell(t), u_h(t), J_i(t)\}_{t>0}$
3. Free-entry: $v(t)$ is such that $V(t) = 0$ in all periods, given $\{J_i(t), u_i(t)\}_{t\geq 0}$
4. $u_i(t)$ is consistent with the screening decisions and $v(t)$ for $i \in \{b, g\}$, $t > 0$

A steady state is an equilibrium where the unemployment rates, the vacancy value and the Bellman equation $\{u_\ell, u_h, v, J\}$ are all time-independent.

2.7 Parameter selection

Table 1 lists the chosen parameters. This paper lays down theory ahead of measurement: many of the moments in the model lack empirical counterparts, but we hope that the theory put forward motivates future empirical work. When the empirical counterparts are unclear, we choose parameter values that are either conservative estimates, or ease the illustration of the mechanisms at play. We will perform robustness checks for these parameters. We include job-to-job transitions when computing firms' effective discount rates ρ. The separation rates are taken from Bjelland et al. (2011). The values of home production b and bargaining power β are taken from Hagedorn and Manovskii (2008). We

Table 1: Parameter selection

Parameter	Value	Description	Source
General			
ρ	0.112	Discount rate	Includes J2J transitions
δ	0.039	Separation rate	Fallick and Fleischman (2016)
b	0.955	Home production	Hagedorn and Manovskii (2008)
β	0.520	Bargaining power	Hagedorn and Manovskii (2008)
Vacancies and matching			
c	0.015	Vacancy opening cost	Vacancy measure of 0.025
a	15.957	Average # applications per cycle	Benchmark
μ_k	0.014	Screening cost: Shape parameter	Screening share of 95%
σ_k	1.000	Screening cost: Scale parameter	Benchmark
τ	10.870	Inverse of vacancy duration	Vacancy duration of one week
Worker heterogeneity			
L_ℓ	0.500	Share of low types in labor force	Benchmark
$\mathcal{P}_\ell$	0.999	Probability of low match (low type)	$\Delta_{\mathcal{P}} = 10$
$\mathcal{P}_h$	0.985	Probability of low match (high type)	Unemployment rate of 0.04
A_b	0.978	Productivity of low match	Bad matches have zero value
A_g	1.000	Productivity of high match	Normalization

Chosen parameter values and targets. All rates are quarterly.

will discuss the importance of this calibration for the results and estimate alternative specifications for b.

The vacancy cost c targets an empirical average vacancy rate of 0.025. We calibrate a such that vacancies on average have 30 applications. We calibrate one screening cost parameter, μ_k, such that 95% of the vacancies are screening their applicant pool. The scale parameter of the screening cost distribution σ_k is chosen to be 1: this pins down the dispersion and skewness of the fixed-costs. For the chosen value of σ_k, the distribution of screening costs has a fat left tail, and many vacancies will draw zero or negligible screening costs. We will later test the robustness of the results under less skewed distributions $G^{\mathbf{k}}$.

To ease the exposition, we pick half of the labor force to be of the low type. $\mathcal{P}_\ell$ and $\mathcal{P}_h$ jointly satisfy an unemployment rate of 4%, and a chosen differential in success rates between the two types $\Delta_{\mathcal{P}} \equiv \frac{1-\mathcal{P}_h}{1-\mathcal{P}_\ell}$. A_g is normalized to one. A_b is set such that low types have zero value for the firm.

3 Introducing worker protection

We use the model to understand the impact of the firing restrictions on the aggregate economy in steady state and over the business cycle. To that end, we have calibrated the economy to a historic long-run steady state in which firings are not regulated ($\mathcal{R}^F = 0$). Then, we introduce the policy change ($\mathcal{R}^F = 1$) and observe the transition to a new steady state. We observe that regulating firings actually improves welfare, since firms were previously destroying matches that were generating a surplus, but were not beneficial to the firm.

The economy is at a steady state with $\mathcal{R}^F(0) = 0$, and at period 0 we introduce $\mathcal{R}^F(t) = 1, t \geq 0$. This policy change was unexpected to the agents up to $t = 0$, but the whole forward sequence of $\mathcal{R}^F$ from then onward is known to them.

Figure 3 shows that the change of policy does not directly affect the value of good and bad matches, J_g and J_b. Forbidding firms to fire badly matched workers does not directly affect the value of randomly hiring, $J_{1-\phi}$. Over time, we observe that $J_{1-\phi}$ improves together with the pool quality P_u. As $J_{1-\phi}$ remains positive, firms with non-screening vacancies still hire. However, now they cannot separate from bad matches: this improves the effective job-finding rate of both low and high type unemployed workers.

Figure 4 plots these job-finding rates on the same scale: the low-type unemployed worker's job-finding rate appreciates more as a response to the policy change. Consequently, the quality of the pool of the unemployed, P_u, appreciates, which makes vacancies more profitable. Market tightness increases as a response.

Figure 3: Job values and screening

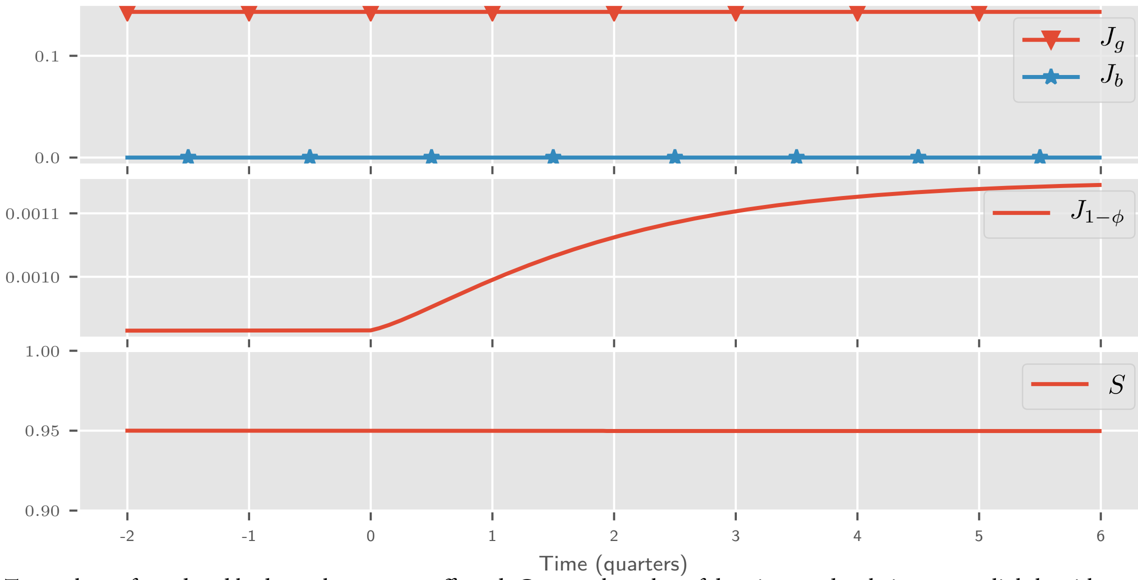

Top: values of good and bad matches are not affected. Center: the value of drawing randomly improves slightly with pool quality. Bottom: the share of screening vacancies is unaffected.

Figure 4: The rise in job-finding rates improves pool quality and raises market tightness

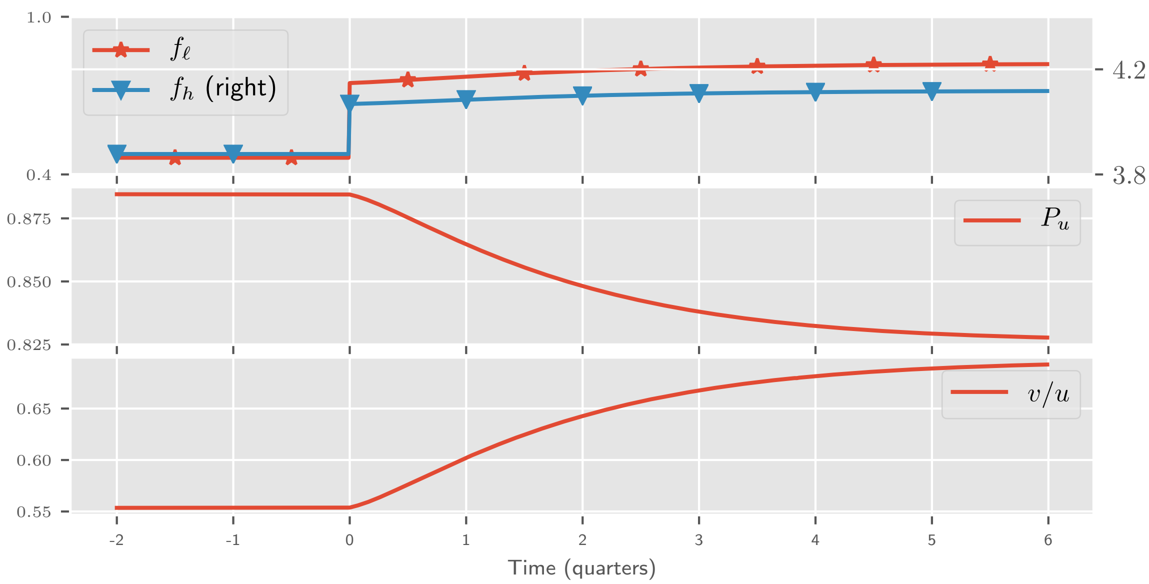

Transition of finding-rates, unemployment rates and vacancies to a steady state where layoffs are regulated.

Table 2: The regulation increases output

Moment	P_u	u	v	SS	Y	Vac c	Scr c	C	Consumption, net
$\mathcal{L} = 1$	0.885	0.044	0.024	0.95	0.9980	0.0006	8.2×10^{-6}	0.9975	0.9555
$\mathcal{L} = 0$	0.832	0.029	0.019	0.95	0.9987	0.0004	6.5×10^{-6}	0.9983	0.9710

Comparison of the steady state where firings are allowed ($\mathcal{L} = 1$) with the steady state where firings are forbidden ($\mathcal{L} = 0$). C: total consumption. *CN*: consumption, excluding home production of the unemployed. *SS*: share of screening vacancies, Vac c: total vacancy cost, Scr c: total screening cost.

Table 2 summarizes the two steady states in more detail. One key finding is that the improved job-finding rates lead to a reduction of unemployment. This reduction of unemployment reduces the number of total vacancies. To evaluate welfare, we measure output. The reform affects output via two channels. First, a reduction in unemployment mechanically increases the number of firms and hence output. However, a positive share of these pairs is now low quality matches that produce less output than a high quality match. The table shows that the latter effect is not dominating: total output is higher in the new steady state. In this model with homogeneous and risk-neutral households, total consumption is an appropriate measure of welfare. To compute consumption, we subtract recruitment costs from aggregate output. A lower vacancy rate in the new steady state together with a constant screening share imply that the total recruitment costs actually decreased over the period. Reducing unemployment here leads to small consumption gains as the unemployed consume a high amount of home production under the Hagedorn and Manovskii (2008) calibration. Instead of representing home production, b could represent unemployment benefits, financed with lump-sum taxation. In that interpretation, the correct consumption measure is net of unemployment benefits, which we display in the last column. When taking into account the transition to the new steady state, the present-discounted value of the consumption gain associated with the reform amounts to 1.7%.

In the appendix, table A.1 provides an overview of many alternative model specifications. The productivity of the bad matches, A_b, deserves special attention as it has the largest potential impact on the welfare effects. In the benchmark calibration, we assumed that bad matches still generate a positive surplus: workers in bad matches are more productive than at home. One of the parameters varied in table 2 is A_b. First, we calibrate $A_b = b$ - which implies that bad matches generate zero surplus. In that scenario, the extent to which unemployment is affected by worker protection is negligible, but still positive. Unemployment drops much more because $J_{1-\phi}$ - the value of randomly hiring - is now rendered negative by the policy: non-screening vacancies do no longer hire. In this case, the consumption gains are also positive, but negligible: the introduction of the policy is still welfare-improving in the steady state. As an extreme case, we calibrate workers to be completely unproductive in bad matches. The table shows that the welfare gains associated with worker protection are now negative, but also negligible. Notice that at $A_b = 0$, all firms are screening and only hiring good matches: lowering A_b below 0 is not going to have any additional effect. We conclude that under our preferred calibration where bad matches have a positive surplus, the policy improves aggregate consumption by around one percentage point – and a negative but negligible consumption loss is a reasonable lower bound for the introduction of such a policy.

Figure 5: TFP shock in unregulated environment, screening

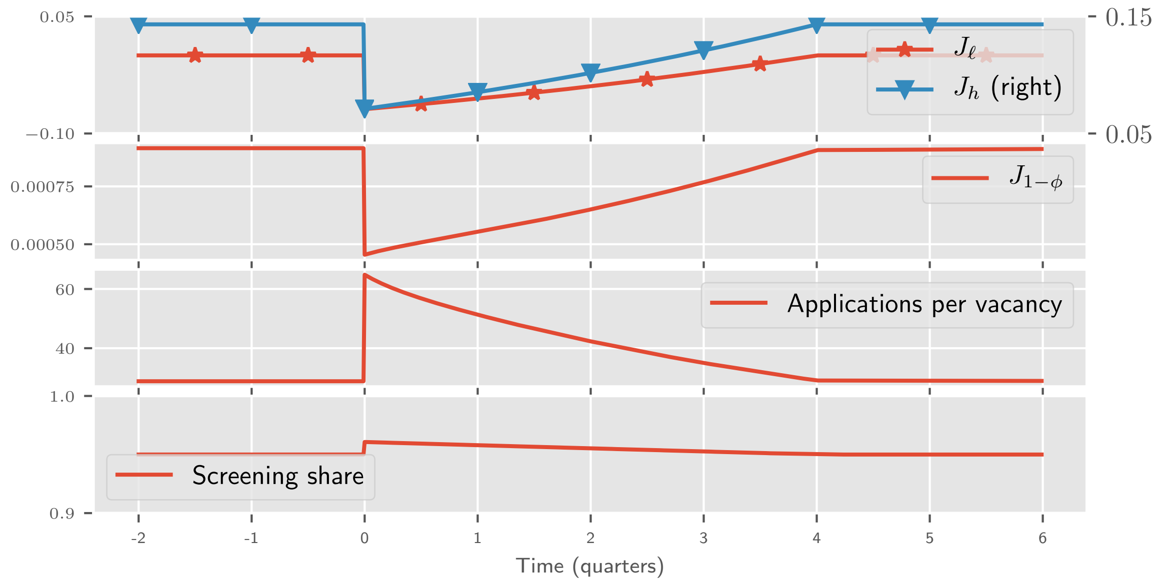

The job values, probabilities and resulting screening strategies in response to the TFP shock.

4 Business cycle fluctuations

Judging by the transition to the new steady state, the reform was a success: output has been increased and recruitment expenditure has been reduced. We now emphasize that this comes at a significant cost: the introduced policy renders the aggregate economy more volatile.

To be specific, we analyze the impact of an unexpected aggregate productivity shock in both steady states with and without the layoff constraint. In the simulations, we decrease $A(t)$ by 10% for 4 quarters in period 0. The shock is unexpected prior to period 0. From that period onward, the whole forward path of $A(t)$ is known by the agents.

4.1 Unconstrained equilibrium

First, we will analyze the impact of the shock when layoffs are unconstrained. As Figure 5 shows, the initial productivity shock reduces the value of both J_g and J_b. The value of random hiring, $J_{1-\phi}$, drops by half. The loss of job values leads to a drop in vacancies: the average number of applications received for vacancies doubles and it becomes profitable for more vacancies to screen their applicants.

Figure 6 displays the impact of this change in policy on the unemployed's job-finding and unem-

Figure 6: TFP shock in unregulated environment, flows

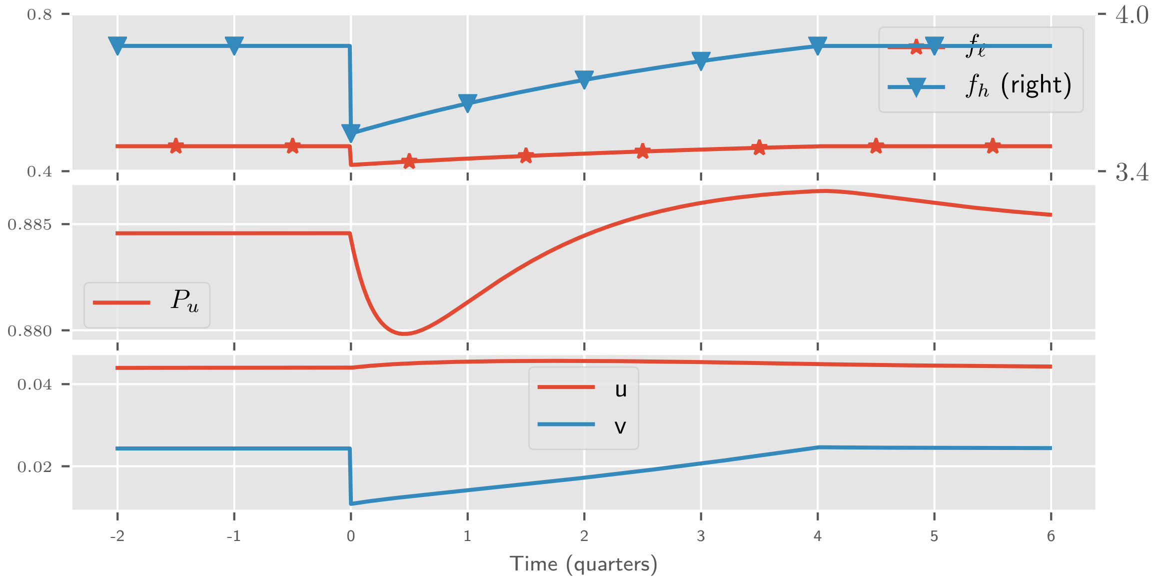

The job-finding rates, unemployment rates and vacancies in response to the TFP shock.

ployment rates. A sharp drop of vacancies at period 0 results in a drop in the job-finding rates for both types. The axes in the top panel have the same scale: the job-finding rates drop slightly more for the high type than for the low type. This is because the drop in v leads to a proportional drop in job-finding rates, and the high-type unemployed had a higher initial job-finding rate. Consequently, the share of low types among the unemployed, P_u, starts to drop slightly at the onset of the productivity shock. However, these changes in P_u and the increase in the unemployment rate are of negligible size.

Here, reasonable fluctuations in the vacancy rate do not generate any measurable changes in the aggregate unemployment rate. The model's failure to generate empirically measured unemployment fluctuations can be reproduced in the most basic version of the urn-ball model. Essentially, the Poisson distribution of the unemployed over vacancies implies that when the measure of vacancies drops, the remainder of the vacancies are much more likely to match with unemployed workers that they will then hire. This increased vacancy filling rate then offsets much of the impact of the vacancy rate on the job-finding rates.

Figure 7: TFP shock in regulated environment, screening

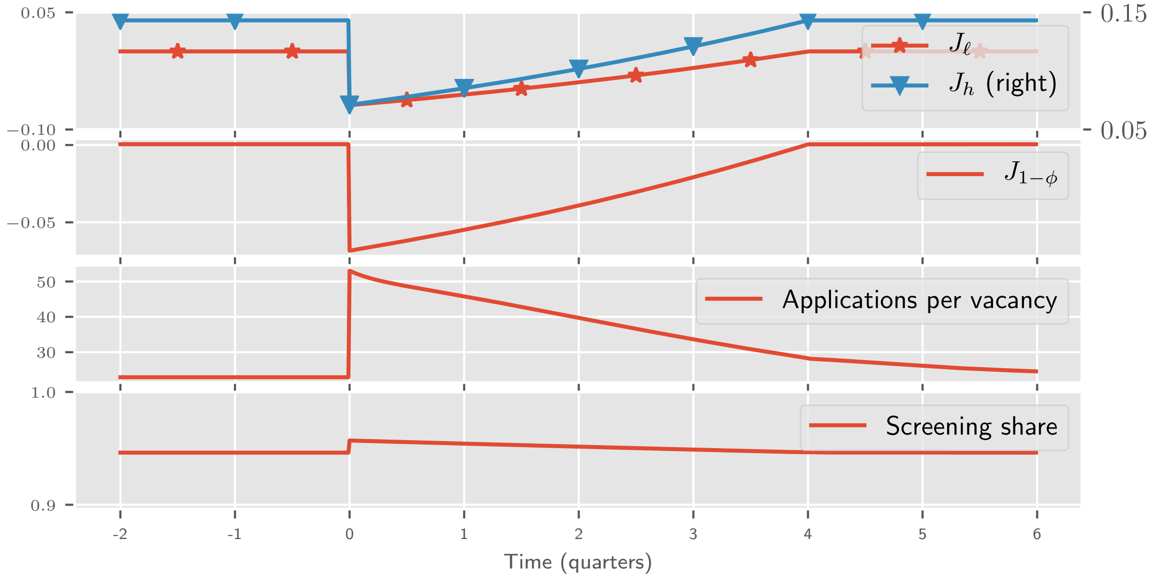

The job values, probabilities and resulting screening strategies in response to the TFP shock.

4.2 Constrained equilibrium

We introduce the same aggregate productivity shock in the steady state where layoffs are constrained, and observe that - unlike before - a TFP shock leads to a strong increase in the unemployment response.

Figure 7 shows that the impact of the shock on job-values $\{J_g, J_b\}$ is the same as before. Yet, the value of random firing, $J_{1-\phi}$ is much more affected under the new regime: firms have to keep bad matches, and the downside risk of randomly hiring is no longer floored at zero. Therefore, the productivity shock implies a much larger drop in $J_{1-\phi}$. In fact, the value of randomly hiring workers becomes negative, and firms with vacancies that decide not to screen their applicants will completely cease hiring. The drop in the vacancy rate and the increased share of screening vacancies are comparable to the previously analyzed recession.

Figure 8 shows that the unwillingness of non-screeners to hire leads to a sharp drop in the job-finding rates for both types, as they will no longer be hired if matched with a non-screening vacancy. Since low types are more dependent on non-screening vacancies, their job-finding rates drop relatively more: the unemployment rate increases for both low and high types, but more so for low types: P_u increases in response to the productivity shock. The job-finding rates start to recover

Figure 8: TFP shock in regulated environment, flows

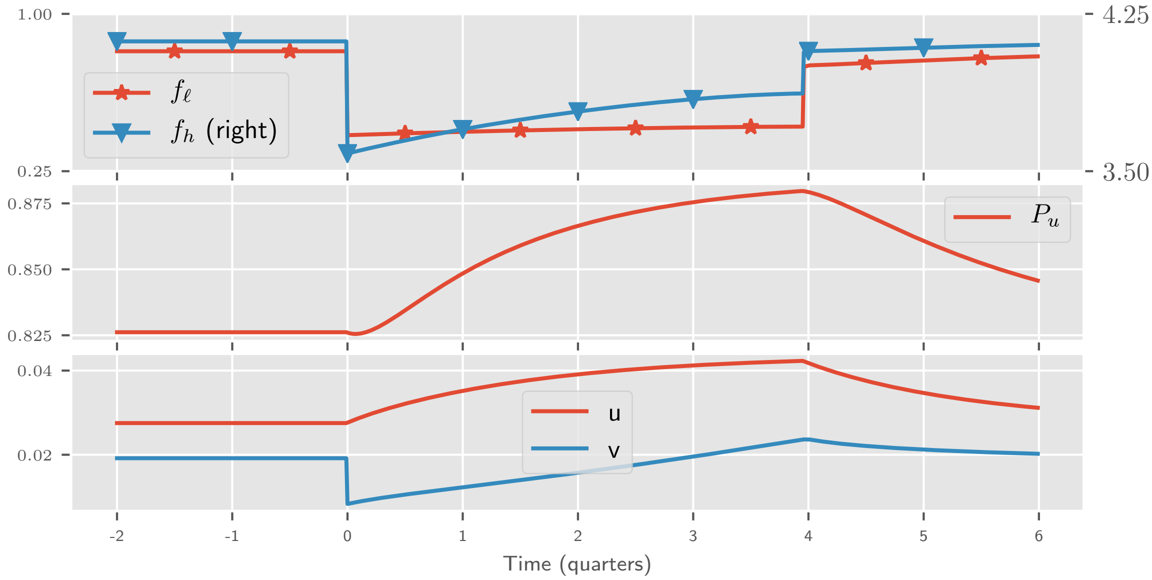

The job-finding rates, unemployment rates and vacancies in response to the TFP shock.

slightly as market tightness returns. However, the continuous recovery of the job-finding rate is much smaller than the discrete jumps of the job-finding rate at the beginning and end of the recession, when firms with non-screening vacancies stop and restart hiring at random. These large changes in the job-finding rate translate into a much larger increase in the unemployment rate: under the new policy regime, the unemployment rate increases by 2 percentage points. After 4 quarters, the aggregate productivity shock recedes. The pool quality P_u is persistent and remains lower for several additional quarters, which disincentivizes hiring. As a consequence, the economy suffers from a "jobless recovery", where the average job-finding rate in the recovery is persistently lower than prior to the recession.

Quantitatively, the unemployment response is smaller than empirically observed. The goal of this model was to emphasize the mechanisms underlying the relationship between worker protection and unemployment: many features that would increase the fluctuations in the unemployment rate are missing here. The key take-away is that the introduction of worker protection increases the sensitivity of aggregate consumption to TFP shocks. In what follows, we will shed some light on the mechanisms that lead to this response.

4.2.1 Screening and unemployment fluctuations: a tale of three mechanisms

In this model, the unemployment response comes from three channels: a direct effect of productivity on vacancy values, a screening externality, and a reduction in random hiring.

The direct effect The first effect is standard in the literature: a decrease in productivity directly reduces the value of all matches and disincentivizes vacancy opening and hiring. In urn-ball models, this direct effect will not lead to a large decrease in the job-finding rates since the effect of losing vacancies is to a large extent offset by a higher share of vacancies that are now receiving more applications, good applicants, and thereby increase their hiring rate. One way of showing this is by varying b: in the context of Cobb-Douglas matching functions, Hagedorn and Manovskii (2008) show that this direct effect is stronger when b is calibrated to a higher value, as this will reduce the firm's surplus and increase its sensitivity to any fundamental. Table A.2 shows that variations in the calibration of b do not significantly change the unemployment response to the TFP shock. Therefore, we argue that the direct effect does not play an important role in explaining the observed changes in the unemployment rate.

The screening externality The new mechanisms both involve the screening decision. First, an increase in the share of screening vacancies in the recession reduces the share of high-type applicants. An increase in P_u then increases the likelihood that none of the applicants for a vacancy will be of a high match, and the vacancy will close without hiring. In equilibrium, this lower quality pool reduces the incentives to open vacancies and leads to an even higher number of applications per vacancy, increases screening and hence reinforces this mechanism.

The strength of this channel depends on two factors: (i) how many marginal vacancies that start screening in the recession, and (ii) the difference in the high-match likelihood between good and bad type.

The first factor depends on the shape of the fixed-cost distribution. The Gamma distribution has two parameters, and we have targeted one moment of that distribution: screening costs are distributed such that 95% of the vacancies are subject to screening in equilibrium. This leaves one parameter free that controls the variance and the skewness. Figure 9 plots different screening cost distributions which all lead to the same share of screening vacancies. In our baseline simulation, we fix $\sigma_k = 1$ which implies a left-skewed distribution of vacancies over fixed costs, and only few marginal vacancies that respond to a change in applications per vacancy. Increasing the mass of marginal vacancies increases P_u between 0.059 and 0.065 but does not significantly affect the unemployment response.

Figure 9: Distribution of screening costs

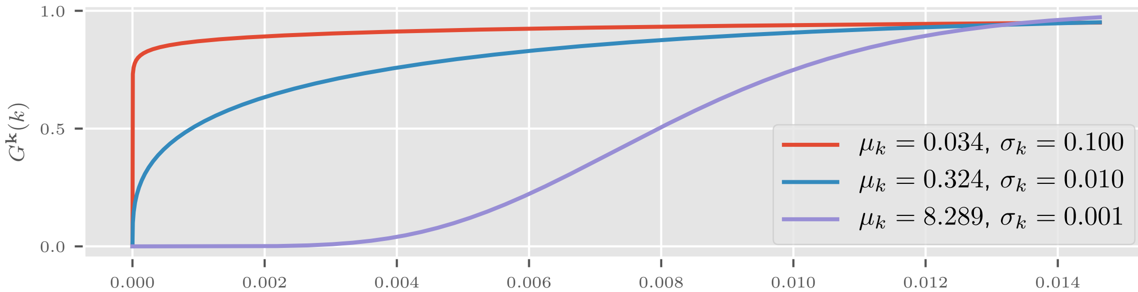

Different screening cost distributions that all imply the same average share of screening vacancies. Distributions with a higher σ_k have a larger mass around 0 and fewer vacancies

The second factor depends on our choice of $\mathcal{P}_\ell$ and $\mathcal{P}_h$. Figure 10 varies the calibrated differences in the probabilities. The pool composition P_u has a stronger impact on u when the difference in these likelihoods is higher: as we increase the difference in these likelihoods, the unemployment response increases. Note that $\frac{1-\mathcal{P}_h}{1-\mathcal{P}_\ell} = 1$ implies that there is no difference in the high-match likelihood across types, and the screening externality is without effect. In the benchmark calibration, the unemployment response from the other two mechanisms adds up to about 1.9 percentage points.

Reduced hiring for non-screeners The last mechanism involves the share of vacancies that do not screen. In this model, some vacancies do not screen because the combination of the fixed-cost draw and the number of applications does not justify screening. In the recession, the quality pool of the unemployed drops and thereby affects the value of randomly hiring a worker. When layoffs are forbidden, the value of randomly hiring - $J_{1-\phi}$ becomes negative under sufficiently negative productivity shocks. at is the case here and the firms posting non-screening vacancies stop hiring. The appendix performs a comprehensive robustness exercise. Table A.2 shows that the unemployment response roughly doubles when we double the calibrated (steady state) share of non-screening vacancies from 5% to 10%.

This channel is very strong because all vacancies have the same decision threshold: whenever the value of $J_{1-\phi}(t)$ crosses 0, all hiring decisions concerning non-screening vacancies change. In the real world, these thresholds probably vary by firm. For example, firms could be heterogeneous in the extent to which match-specific productivity matters in their production. This might be because some jobs - such as being a cashier - are very well regulated and standardized, and the extent to which match-specific productivity might affect firm profits would be limited. To capture that heterogeneity, job values could be modeled as follows:

Figure 10: Unemployment response increases in pool heterogeneity

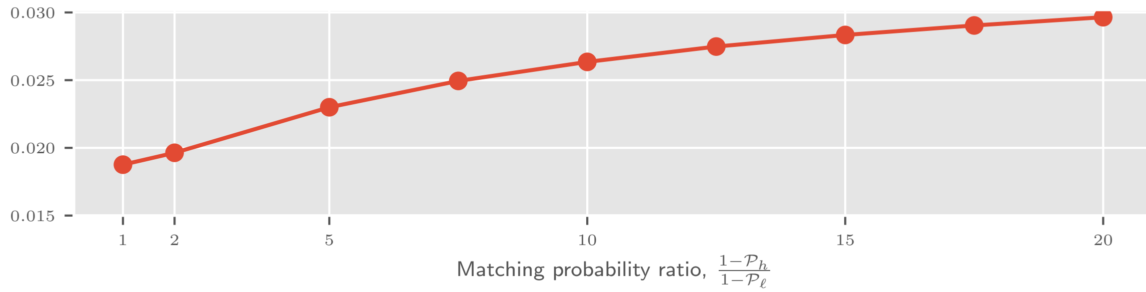

$$\rho J_{i,\gamma}(t) = A_i^{\gamma} A(t) - w(t) - \delta J_{i,\gamma}(t) + \dot{J}_{i,\gamma}(t)$$

where $\gamma \geq 0$ controls the extent to which the job depends on match-specific productivity. In an extension of this model, vacancies draw both k and γ when opened, and screening cost thresholds are now a function of both the number of applications x and γ: $\overline{k} = \overline{k}(x, \gamma)$. Such an extension would complicate the exposition of the model at little gain and is therefore not demonstrated here.

The remaining variations in Table A.2 concern the share of low types in the economy, L_ℓ, and the average number of applications per vacancy, au/v. Neither parameter has a sizable impact on the unemployment response.

5 Conclusion

We have built a model that adds to the literature on worker protection by analyzing its impact on firms' screening decision in an environment with heterogeneous match productivity. The model uncovers a new trade-off: the introduction of worker protections potentially increases aggregate consumption in the new steady state, and the transition thereto. However, this comes at the cost of higher volatility, as it renders the economy more susceptible to business cycle fluctuations. The presented model is very stylized and we do not provide a complete welfare assessment.

The model is a cautionary tale for welfare analysis: it is not always sufficient to evaluate a policy along the transition to the steady state, as it may affect other dynamics of the model.

The consumption gains in the steady state primarily stem from the inefficient wage schedule: firms make losses from bad matches that actually generate positive surpluses. When layoffs are outlawed, firms with bad matches have to keep these individuals and thereby contribute to aggregate consumption. Worker protection reduces the economy's resilience to negative productivity shocks as they increase the risk of hiring without screening. Without layoffs as a safeguard, hiring unscreened applicants becomes a very risky proposal, and recruitment becomes more sensitive to the business cycle. This takes place through two channels. First, firms for which screening is too costly decide to completely withdraw from the market. Second, some firms intensify their screening efforts to prevent the risk of bad matches. Through their more discriminatory hiring practices, they worsen the pool of applicants and thereby magnify the impact of the original productivity shock.

References

Acharya, Sushant and Shu Lin Wee (2019). "Rational inattention in hiring decisions". In: *SSRN Electronic Journal*, pp. 1–25. ISSN: 1556-5068.

Andolfatto, David (Mar. 1996). "Business cycles and labor-market search". In: *The American Economic Review* 86, pp. 112–132.

Autor, David H., William R. Kerr, and Adriana D. Kugler (June 2007). "Does employment protection reduce productivity? Evidence from US states". In: *The Economic Journal* 117.521, F189–F217. ISSN: 0013-0133.

Bertola, Giuseppe (1999). "Microeconomic perspectives on aggregate labor markets". In: *Handbook of Labor Economics*. Chap. 45, pp. 2985–3028.

Bjelland, Melissa et al. (Oct. 2011). "Employer-to-employer flows in the united states: estimates using linked employer-employee data". In: *Journal of Business & Economic Statistics* 29.4, pp. 493–505. ISSN: 0735-0015.

Blanchard, Olivier Jean and Peter Diamond (July 1994). "Ranking, unemployment duration, and wages". In: *The Review of Economic Studies* 61.3, pp. 417–434. ISSN: 0034-6527.

Blanchard, Olivier and Augustin Landier (June 2002). "The perverse effects of partial labour market reform: fixed-term contracts in france". In: *The Economic Journal* 112.480, F214–F244. ISSN: 0013-0133.

Cahuc, Pierre and Fabien Postel-Vinay (Feb. 2002). "Temporary jobs, employment protection and labor market performance". In: *Labour Economics* 9.1, pp. 63–91. ISSN: 09275371.

Dal Bianco, Silvia, Randolph L. Bruno, and Marcello Signorelli (Mar. 2015). "The joint impact of labour policies and the "great recession" on unemployment in europe". In: *Economic Systems* 39.1, pp. 3–26. ISSN: 09393625.

Ghayad, Rand (2013). "The jobless trap".

Hagedorn, Marcus and Iourii Manovskii (Aug. 2008). "The cyclical behavior of equilibrium unemployment and vacancies revisited". In: *American Economic Review* 98.4, pp. 1692–1706. ISSN: 0002-8282.

Hassel, Anke (2014). "Adjustments in the eurozone: varieties of capitalism and the crisis in southern europe". In: *SSRN Electronic Journal* 76. ISSN: 1556-5068.

Jarosch, Gregor and Laura Pilossoph (July 2019). "Statistical discrimination and duration dependence in the job finding rate". In: *The Review of Economic Studies* 86.4, pp. 1631–1665. ISSN: 0034-6527.

Josephson, Jens and Joel Shapiro (Mar. 2016). "Costly interviews". In: *International Journal of Industrial Organization* 45, pp. 10–15. ISSN: 01677187.

Karamessini, Maria (Mar. 2008). “Continuity and change in the Southern European social model”. In: *International Labour Review* 147.1, pp. 43–70. ISSN: 0020-7780.

Kroft, Kory, Fabian Lange, and Matthew J Notowidigdo (Aug. 2013). “Duration dependence and labor market conditions: evidence from a field experiment”. In: *The Quarterly Journal of Economics* 128.3, pp. 1123–1167. ISSN: 0033-5533.

Lazear, Edward P. (Aug. 1990). “Job security provisions and employment”. In: *The Quarterly Journal of Economics* 105.3, p. 699. ISSN: 00335533.

Lindbeck, Assar (1993). *Unemployment and Macroeconomics*. The MIT Press. ISBN: 9780262121750.

Lindbeck, Assar and Dennis J. Snower (Feb. 2001). “Insiders versus outsiders”. In: *Journal of Economic Perspectives* 15.1, pp. 165–188. ISSN: 0895-3309.

Masters, Adrian (Jan. 2014). “Statistical discrimination from composition effects in the market for low-skilled workers”. In: *Labour Economics* 26, pp. 72–80. ISSN: 09275371.

Merz, Monika (Nov. 1995). “Search in the labor market and the real business cycle”. In: *Journal of Monetary Economics* 36.2, pp. 269–300. ISSN: 03043932.

Shimer, Robert (Feb. 2005). “The cyclical behavior of equilibrium unemployment and vacancies”. In: *American Economic Review* 95.1, pp. 25–49. ISSN: 0002-8282.

A Tables

Table A.1: Robustness: transition to new steady state

change	Moment value	P_u	u	v	SS	C	CN
baseline		-0.05	-0.015	-0.005	-0.0003	0.0008	0.015
L_ℓ	0.25	-0.13	-0.017	-0.005	-0.0003	0.0009	0.017
	0.4	-0.08	-0.016	-0.005	-0.0002	0.0008	0.016
	0.6	-0.03	-0.014	-0.005	-0.0002	0.0008	0.014
	0.75	-0.02	-0.013	-0.005	-0.0001	0.0007	0.013
σ_k	0.1	-0.05	-0.015	-0.005	-0.0006	0.0008	0.016
	0.01	-0.05	-0.016	-0.005	-0.0028	0.0008	0.016
	0.001	-0.07	-0.019	-0.006	-0.0205	0.0010	0.019
b	0.3	-0.05	-0.015	-0.005	-0.0007	0.0125	0.017
	0.5	-0.05	-0.015	-0.005	-0.0006	0.0089	0.017
	0.8	-0.05	-0.015	-0.005	-0.0004	0.0036	0.016
$\frac{1-\mathcal{P}_h}{1-\mathcal{P}_\ell}$	1	0.00	-0.010	-0.006	0.0000	0.0006	0.010
	2	-0.03	-0.011	-0.005	-0.0000	0.0006	0.011
	10	-0.05	-0.015	-0.005	-0.0003	0.0008	0.015
au/v	10	-0.05	-0.016	-0.005	-0.0018	0.0008	0.016
	50	-0.05	-0.015	-0.005	-0.0002	0.0008	0.015
SS	0.9	-0.11	-0.024	-0.008	-0.0010	0.0013	0.025
	0.925	-0.08	-0.021	-0.007	-0.0005	0.0011	0.021
	0.975	-0.02	-0.009	-0.003	-0.0001	0.0004	0.009
A_b	b	-0.00	-0.001	-0.002	0.0292	0.0000	0.001
	0	-0.00	0.000	-0.010	0.0499	-0.0000	-0.000

Change in moments (new steady state minus old steady state) under different robustness checks. *SS*: Share of screening vacancies. *C*: total consumption. *CN*: consumption, excluding home production of the unemployed. *ss*: steady state value. *max*: maximum value in transition. *change*: difference between maximum deviation and steady state.

Table A.2: Robustness: Impact of negative TFP shock

change	value	u (diff)	θ (diff)	P_u (ss)	P_u (diff)	SS (max)	Y (% diff)	C (loss)	CN (loss)
baseline		0.026	-0.308	0.807	0.059	0.960	-0.006	-0.353	-8.352
L_ℓ	0.25	0.025	-0.306	0.506	0.105	0.960	-0.004	-0.347	-8.179
	0.4	0.027	-0.309	0.714	0.082	0.960	-0.005	-0.367	-8.683
	0.6	0.025	-0.306	0.872	0.040	0.960	-0.006	-0.329	-7.788
	0.75	0.023	-0.302	0.937	0.019	0.960	-0.007	-0.290	-6.848
σ_k	0.1	0.026	-0.313	0.807	0.059	0.971	-0.006	-0.348	-8.229
	0.01	0.025	-0.305	0.807	0.060	0.993	-0.006	-0.346	-8.161
	0.001	0.026	-0.303	0.804	0.064	1.000	-0.006	-0.387	-9.029
b	0.3	0.022	-0.097	0.807	0.050	0.951	-0.094	-4.622	-6.539
	0.5	0.024	-0.106	0.807	0.056	0.951	-0.066	-3.935	-7.748
	0.8	0.026	-0.111	0.807	0.058	0.953	-0.026	-1.653	-8.140
au/v	10	0.027	-0.308	0.808	0.059	0.961	-0.006	-0.358	-8.458
	50	0.026	-0.308	0.807	0.059	0.960	-0.006	-0.353	-8.346
Target SS	0.9	0.063	-0.319	0.649	0.112	0.919	-0.009	-0.819	-18.545
	0.925	0.043	-0.311	0.735	0.095	0.940	-0.008	-0.591	-13.510
	0.975	0.012	-0.306	0.855	0.027	0.980	-0.003	-0.141	-3.898
A_b	b	0.001	-0.242	0.887	-0.003	0.953	-0.019	0.026	-0.502
	0.5	0.001	-0.198	0.887	-0.003	0.951	-0.195	0.036	-0.484

Comparison of moments under different robustness checks. *SS*: share of screening vacancies. *C*: total consumption. *CN*: consumption, excluding home production of the unemployed. *ss*: steady state value. *max*: maximum value in transition. *diff*: difference between maximum deviation and steady state.