

Selective Hiring and Welfare Analysis in Labor Market Models

Christian Merkl

University of Erlangen-Nuremberg, Kiel Institute and IZA

Thijs van Rens

CREI, Universitat Pompeu Fabra, IZA and CEPR

August 2011

Abstract

Firms select not only how many, but also which workers to hire. Yet, in standard labor market models, all workers have the same probability of being hired. We argue that selective hiring crucially affects welfare analysis. Our model is isomorphic to a standard search and matching model under random hiring but allows for selective hiring. With selective hiring, the positive predictions of the model change very little, but the welfare costs of unemployment are much larger because unemployment risk is distributed unequally across workers. As a result, optimal unemployment insurance may be higher and welfare is lower if hiring is selective.

Keywords: labor market models, welfare, optimal unemployment insurance

JEL codes: E24, J65

Preliminary

Please do not cite without permission

1 Introduction

Standard labor market models are not suitable to evaluate the welfare costs of unemployment and business cycle fluctuations. To make these models tractable, the literature typically assumes that workers are homogeneous and markets are complete (Lucas 1987). Combined, these assumptions eliminate the most important source of welfare costs of unemployment: the fact that unemployment is unequally distributed across workers. Previous attempts to do welfare analysis of unemployment and labor market policies have focused on relaxing the complete markets assumption. If markets are incomplete, ex ante identical workers cannot share unemployment risk, making these workers heterogeneous ex post. Because we need to keep track of the entire distribution of asset holdings, this class of models is difficult to solve. Moreover, because workers are ex ante homogeneous, welfare costs of unemployment and business cycles are small (Krusell and Smith 1998).

We propose a framework, in which workers are ex-ante heterogeneous. In this model, some workers are more attractive to employers than others, because they have lower training costs. Individual-specific training costs are fully observable to workers and firms and determine how likely it is that a worker finds a job in a given period. We analyze two polar cases in this framework. If individual training costs are serially uncorrelated, in expectation each worker has the same probability of being hired in future periods. We call this version of the model the case of perfectly random hiring. If individual training costs are permanent, the probability of being hired in the future depends on current training costs, which will be the same in future periods. This is the case of perfectly selective hiring.

If training costs are persistent but not permanent, then hiring is partly random and partly selective. There is ample evidence that hiring decision in the real world are partly selective, and not all workers have the same probability of finding a job. We discuss this evidence in section 2. By analyzing the polar cases of perfectly random and perfectly selective hiring, we aim to understand how selectivity in hiring decision matters for our understanding of the labor market.

Our model is set up in such a way that, under perfectly random hiring, it is isomorphic to a standard search and matching model (Pissarides 2000, chapter 1) in terms of its predictions for labor market dynamics. Specifically, the model can be parameterized to generate the same aggregate job finding and unemployment rate, and the same elasticities of these variables with respect to changes in productivity. The distribution of idiosyncratic training costs plays the same role as the aggregate matching function in the standard models. Thus, we provide a framework that on the one hand maintains most of the insights from standard labor market models, and on the other hand allows us to compare the predictions of the model under selective versus random hiring. While most predictions of the model are very similar under selective hiring, the implications for welfare analysis are very different.

If hiring is selective, unemployment is costly because unemployment risk is spread unequally across workers. With perfectly selective hiring, some workers are always employed, while others are always unemployed. Thus, unemployment risk is uninsurable and the welfare costs of unemployment are much larger than under random hiring. As a result, there is a role for government intervention, insuring unborn workers against their unemployment risk.

As an application, we formally study the question of optimal unemployment insurance.

Under random hiring, the government can replicate the efficient allocation, using unemployment benefits and lump-sum taxes as instruments. In this case, unemployment benefits are set to make sure the level of job creation is efficient, similar to the Hosios (1990) condition in search and matching models. Under selective hiring there is an additional motive for unemployment insurance because workers cannot self-insure against their characteristics, which determine their individual-specific unemployment risk.¹ Thus, the government faces a trade-off between efficient job creation and efficient redistribution. We solve the Ramsey problem for the government in this case and find two results. First, the maximum welfare that can be reached under selective hiring is substantially lower than under random hiring. Second, to obtain a more equitable income distribution with selective hiring, it may be optimal to set unemployment benefits substantially higher than under random hiring.

The remainder of this paper is structured as follows. Section 2 describes some evidence from microeconomic labor market data that hiring decisions of firms in the real world are partially selective. Section 3 sets up the model and solves for the efficient allocation. We describe the equilibrium of the model over two sections. In section 4, we derive the equilibrium job creation condition and discuss conditions, under which equilibrium job creation is efficient. We also formally establish the equivalence of the random hiring model to a search and matching model in this section, and discuss the positive differences between the models with random and selective hiring. Section 5 concerns welfare analysis. Here, we show analytically how the welfare costs of unemployment differ starkly under selective versus random hiring. Section 6 establishes this point more formally in an application to optimal unemployment insurance. Section 7 concludes and briefly discusses the costs of business cycles as another possible application, in which selective hiring is likely to make a large difference.

2 Selective Hiring: Motivating Evidence

We start the paper with a review of the evidence that hiring decisions are at least partially selective. This evidence consists of facts that have been documented in other contexts, but that have not always been interpreted as evidence for selective hiring.

The most direct evidence comes from the distribution of job finding rates. If hiring is perfectly random, then all workers have the same probability of finding a job. If hiring is perfectly selective, then some ‘good’ workers find jobs immediately, whereas other ‘bad’ workers never find jobs. In the data, the job finding rate decreases with unemployment duration, both in the US (Abraham and Shimer 2002) and in Europe (Wilke 2005). This is consistent with bad workers being over-represented in the pool of long-term unemployed.

A similar picture emerges when we compare the aggregate job finding rate to the average unemployment duration. If all workers have the same job finding rate, then the average unemployment duration D must simply be the inverse of the aggregate job finding rate, $D = 1/f$. If hiring is selective, then bad workers (with low job finding probabilities) are over-represented in the average unemployment duration and under-represented in the average job finding rate, so we would expect $D > 1/f$. In the data, unemployment duration is indeed much longer than expected based on the aggregate job finding rate (Shimer 2007). By

¹This result derives from ex-ante heterogeneity, not from credit constraint, as in Landais, Michailat and Saez (2011).

a similar argument, selective hiring explain why the net job finding rate, which excludes workers with unemployment duration shorter than the period of observation, is smaller than the gross job finding rate (Shimer 2007).

A third piece of evidence comes from the composition of the pools of employed and unemployed workers. If hiring is selective, we would expect the quality of the employment pool to be countercyclical, because workers that are only hired in booms are relatively bad compared to workers that already had jobs in the recession. But by the same token, we would expect the quality of the unemployment pool to be countercyclical as well, because workers that are hired in booms are relatively good compared to workers that remain unemployed even in booms. It has long been known that there is a composition bias in the cyclicity of wages consistent with this story (Solon, Barsky and Parker 1994). In addition, Mueller (2010) recently documented that the average predicted wage of unemployed workers is countercyclical. Although Mueller gives a different interpretation to his finding, we interpret it as evidence that hiring is selective.

3 Model Environment

Our economy is populated by a continuum of worker-consumers i , characterized by ε_{it} . We model worker characteristics as training costs: a firm that hires worker i in period t needs to pay ε_{it} for this worker to become productive. Alternatively, we may think of ε_{it} as a measure of worker productivity or match-specific skills, which would lead to minor modifications to the model but would leave the results unchanged. Worker characteristics are fully observable to workers and firms, so that there is perfect information in the economy. In our model, training costs (or worker characteristics in general) determine how likely an individual worker is to be hired in a given period. Let G and g denote the distribution function and the probability density function of training costs, $\varepsilon_{it} \sim G$. The distribution G is assumed to be constant across individuals and time-invariant.² This modelling framework is inspired by Brown, Merkl and Snower (2010), although both the focus of the analysis and the details of the model are very different from that paper.

Whether hiring is selective or random depends on the persistence of individual worker characteristics. If ε_{it} is serially uncorrelated for each individual, then each worker expects to have the same probability of being hired in future periods, so that hiring decisions are independent of current worker characteristics and thus effectively random from today's perspective. If ε_{it} is permanent, i.e. $\varepsilon_{it} = \varepsilon_i$ is fixed for each worker over time, then current worker characteristics fully determine how likely an individual worker is to be hired in the future. This is what we call selective hiring. If ε_{it} has an autocorrelation higher than 0 but lower than 1, then hiring is partly random and partly selective.

3.1 Preferences

Workers may be employed or unemployed. Employed workers earn a wage w_t and unemployed workers receive unemployment benefits, including utility from leisure, b_t . Potentially, wages

²This does not mean, of course, that the average job finding rate is constant over time, since other factors than ε_{it} also affect the probability to be hired.

and unemployment benefits could depend on that worker characteristics ε_{it} , but we assume that this is not the case. This assumption, which we will maintain throughout the paper, is not very restrictive, because only the wage of the marginal hire is allocative, see section 4.2 for a more detailed discussion.

Worker-consumers are infinitely-lived, have time-separable utility and care about the expected net present value of utility from consumption c_{it} and leisure. We assume the parameters of the model are such that $b_t < w_t$ for all t , so that all unemployment is involuntary. For simplicity, we also assume the utility from leisure is included in the unemployment benefit b_t , so that the flow utility $\mathcal{U}(\cdot)$ depends only on consumption. Then, workers' objective function is given by,

$$E_0 \sum_{t=0}^{\infty} \beta^t \mathcal{U}(c_{it}) \quad (1)$$

where β is the discount factor and E_0 denotes rational expectations in period 0.

3.2 Production and Job Creation

Employed workers hold jobs, which produce output y_t in each period, including the period in which the worker was hired and trained. Given our assumption that worker characteristics take the form of training cost, the output of a job does not depend on the characteristics of the worker that holds it. Given a wage w_t , the firm's profits from a job equal $y_t - w_t$. The cost of creating a job is the cost of training a worker, which has a fixed component K and an idiosyncratic component ε_{it} . It is important to note that there are no search frictions in our model, so that jobs with positive value can be created immediately.

Since all jobs are identical after the worker has been trained, jobs created for workers with low training costs generate more output than jobs created for workers with high training costs. Thus, if it is profitable and/or efficient to create a job for a worker with training costs ε , then it must also be profitable and/or efficient to create a job for a worker with lower training costs $\varepsilon' < \varepsilon$. Thus, in the efficient as well as in the equilibrium allocation of this model, there exists a cutoff level $\tilde{\varepsilon}_t$, such that a worker seeking a job is hired if $\varepsilon_{it} < \tilde{\varepsilon}_t$ and not hired if $\varepsilon_{it} > \tilde{\varepsilon}_t$. Although the existence of this hiring threshold is clearly a property of the efficient allocation or equilibrium and not part of the environment, we will nevertheless impose it below in order to simplify the notation.

3.3 Markets

Worker-consumers and firms interact with each other on three markets. Firms hire workers on the labor market. The goods firms produce are sold to consumers on the goods market. Both firms and workers can trade on the capital market, which has a zero net supply of assets.

On the labor market, firms hire unemployed workers, generating jobs and employed workers. There is an exogenous probability λ that a job is destroyed, in which case the worker becomes unemployed again. We assume full commitment of both worker and firm, so that regardless of the worker's ε_{it} both the firm and the worker must continue the job unless it

is destroyed by a λ -shock, i.e. there is no endogenous job destruction. Let $f(\tilde{\varepsilon}_t)$ denote the aggregate job finding rate: the probability that an average job seeker finds a job in each period. The aggregate job finding rate depends on the hiring threshold defined in section 3.2 above: the higher the threshold, the larger the probability that any given job seeker is hired, everything else equal. Then, the number of employed workers in the economy evolves according to,

$$n_t = (1 - \lambda) n_{t-1} + f(\tilde{\varepsilon}_t) s_t = (1 - \lambda) (1 - f(\tilde{\varepsilon}_t)) n_{t-1} + f(\tilde{\varepsilon}_t) \quad (2)$$

where s_t is the number of workers that seek a job in a given period. We assume job destruction happens before job creation, so that the number of workers that are seeking jobs equals the number of workers that are unemployed since last period, $1 - n_{t-1}$, plus the number of workers that were employed last period but lost their job in this period, λn_{t-1} . Notice that the number of job seekers $s_t = 1 - (1 - \lambda) n_{t-1}$ does not equal the number of unemployed $u_t = 1 - n_t$, because some of the job seekers find new jobs within the period.

On the goods market, goods produced by firms are sold to workers for consumption. Goods market clearing requires that the amount of goods produced equals the amount of goods consumed by workers plus the amount of goods used to pay the training costs to create jobs. We assume that if firms make any profits in excess of the amount they need to pay the training costs, then these profits are distributed lump-sum to workers and then consumed. Thus, the aggregate resource constraint is given by,

$$\int_{-\infty}^{\infty} c_{it} dG = y_t n_t - [1 - (1 - \lambda) n_{t-1}] f(\tilde{\varepsilon}_t) (K + H(\tilde{\varepsilon}_t)) \quad (3)$$

where $K + H(\tilde{\varepsilon}_t)$ denotes the average training cost of all workers that were hired in period t , which depend on the hiring threshold defined in section 3.2 above.

Capital markets give both workers and firms access to borrowing and lending of a risk-free bond. The bond market allows workers to fully self-insure against any idiosyncratic fluctuations in their income. Since bonds are in zero net supply, aggregate risk is not insurable. And since the unborn do not have access to bond markets, workers cannot insure against their characteristics in period 0 either. Capital market clearing is implied by labor and goods market clearing, so we do not specify the capital market clearing condition separately.

3.4 Efficiency

The social welfare function aggregates the utility (1) of all workers in the economy. We assume a utilitarian welfare function, which weights the utility of all individuals equally. Thus, the social planner maximizes,

$$E_0 \sum_{t=0}^{\infty} \beta^t \int_{-\infty}^{\infty} \mathcal{U}(c_{it}) dG \quad (4)$$

subject to the law of motion for employment (2),

$$n_t = (1 - \lambda) (1 - f(\tilde{\varepsilon}_t)) n_{t-1} + f(\tilde{\varepsilon}_t) \quad (5)$$

and the aggregate resource constraint (3). The planner chooses how many workers to employ, which workers to employ, and how to distribute consumption over all employed and unemployed workers. As we argued in section 3.2, it is efficient to employ all workers with training costs ε_{it} below a threshold $\tilde{\varepsilon}_t$ and let workers with training costs above this threshold be unemployed. Imposing this property of the efficient allocation, the planner chooses the hiring threshold $\tilde{\varepsilon}_t$ and the consumption level of each worker $\{c_{it}\}_{i=-\infty}^{\infty}$ in each period t .

The solution of the social planner problem is straightforward. Details may be found in appendix A.1. The results may be summarized in two efficiency conditions, one about the efficient consumption allocation and the second one about efficient job creation.

In the efficient allocation, consumption is equal for all workers.

$$c_{it} = c_t \text{ for all } i \quad (6)$$

The level of consumption in period t can be found by substituting this result into the aggregate resource constraint, but is not of interest here. The important observation is that the social planner awards the same level of consumption to all workers, whether employed or unemployed and independent of their training costs ε_{it} . Of course this result depends to some degree on specific assumptions, in particular the additive separability of utility in consumption and leisure. The intuition for the result, however, is quite general. It is also important to note that any reasonable welfare function would deliver the same result. By equalizing consumption across workers, the planner minimizes the welfare loss from poor workers, who would have very steep marginal utility of consumption. This result will play a crucial role in the welfare analysis in section 5.

The efficient job creation equation determines the hiring threshold $\tilde{\varepsilon}_t$, i.e. the training cost of the marginal worker that is hired, which in turn pins down the efficient job finding rate $f(\tilde{\varepsilon}_t)$.

$$K + M(\tilde{\varepsilon}_t) = y_t + (1 - \lambda) E_t [Q_{t,t+1} \{K + (1 - f(\tilde{\varepsilon}_{t+1})) M(\tilde{\varepsilon}_{t+1}) + f(\tilde{\varepsilon}_{t+1}) H(\tilde{\varepsilon}_{t+1})\}] \quad (7)$$

where

$$M(\tilde{\varepsilon}_t) = H(\tilde{\varepsilon}_t) + \frac{H'(\tilde{\varepsilon}_t) f(\tilde{\varepsilon}_t)}{f'(\tilde{\varepsilon}_t)} \quad (8)$$

is a measure of marginal training costs and

$$Q_{t,t+1} = \frac{\beta \mathcal{U}'(c_{t+1})}{\mathcal{U}'(c_t)} \quad (9)$$

is the stochastic discount factor. For future reference, we also define $Q_{t,t+\tau} = 1$ for $\tau = 0$ and $Q_{t,t+\tau} = Q_{t,t+1} Q_{t,t+2} \dots Q_{t+\tau-1,t+\tau}$ for $\tau \geq 1$. In order to help with the intuition, it is useful to note that for perfectly random hiring $M(\tilde{\varepsilon}_t) = \tilde{\varepsilon}_t$, see appendix A.2 for details.

Condition (7) has the usual interpretation of a job creation condition, stating that the (social) cost of hiring the marginal worker must equal the expected net present value of the (social) value of having that worker employed. The cost of hiring the marginal worker equals the training costs of the marginal worker $K + M(\tilde{\varepsilon}_t)$, which gives the left-hand side of the condition. The right-hand side of the condition (7) represent the benefits of hiring this worker, which include the output produced by the worker y_t , plus the expected opportunity

cost of not having to hire a worker next period. The expected cost of hiring next period is discounted by a stochastic discount factor, which includes not only the rate of time-preference β and the ratio of the marginal utility of consumption between tomorrow and today, but also the separation rate λ , reflecting the possibility that the job that is created today may be destroyed tomorrow with probability λ .

Why is $K + M(\tilde{\varepsilon}_t)$ a measure of marginal training costs? With perfectly random hiring this is clear, as $K + M(\tilde{\varepsilon}_t) = K + \tilde{\varepsilon}_t$ equals the training costs of the marginal worker that is hired. In general, $K + M(\tilde{\varepsilon}_t)$ equals the average training costs $K + H(\tilde{\varepsilon}_t)$, plus a correction for the fact that the average cost function is increasing in the hiring threshold, so that the marginal hire is more expensive to train than the average worker. The expected costs of hiring next period, on the right-hand side of condition (7), are a weighted average between the expected marginal training costs $K + M(\tilde{\varepsilon}_{t+1})$ and the expected average training costs $K + H(\tilde{\varepsilon}_{t+1})$, where the weight on the average cost is the average job finding rate, as the social planner takes into account the effect of the hiring decision on the pool of job seekers in future periods.

3.5 Road Map

In the next two sections of the paper, we solve for the equilibrium allocation of the model, evaluate its properties and compare it to the efficient allocation. Section 4 deals with job creation and derives conditions under which efficient job creation can be supported as an equilibrium. Section 5 deals with the consumption allocation. This section, which contains the main result of the paper in its simplest form, shows that efficiency of the equilibrium depends crucially on the persistence of worker characteristics.

4 Equilibrium Unemployment

In this section, we derive the equilibrium job creation condition and compare it to the efficient job creation condition. We show that, under some conditions, equilibrium job creation is efficient. Then, we explore what are the aggregate job finding rate and unemployment rate implied by the job creation condition. To do this, we need to specify whether hiring decisions are random or selective. We explore both versions of the model and show that the predictions of our model for job creation with random hiring are very similar (and under some conditions identical) to the predictions of a standard Diamond-Mortensen-Pissarides model with search frictions. In that sense, we can think of our framework as micro-foundation for an aggregate matching function.³

The objective of this section is to show that the predictions of our model about the cyclical behavior of the labor market are very similar to those of standard search models of the labor market, and that it makes very little difference for those predictions whether hiring is random or selective. Welfare analysis, however, differs sharply with the selectivity of hiring. We postpone this issue to section 5 and in this section completely ignore the consumption allocation.

³This is consistent with the interpretation in Pissarides (2000, p.4) that “the matching function summarizes a trading technology between heterogeneous agents that is also not made explicit.”

4.1 Job Creation

In the decentralized equilibrium, firms decide whether or not to create jobs.⁴ Like the social planner, firms must choose how many workers as well as which workers to employ. As we argued in section 3.2, it is profit-maximizing to employ all workers with training costs ε_{it} below a threshold $\tilde{\varepsilon}_t$ and let workers with training costs above this threshold be unemployed. Imposing this property of equilibrium, firms directly choose the hiring threshold $\tilde{\varepsilon}_t$, which determines the total number of workers they employ.

The representative firm chooses $\tilde{\varepsilon}_t$ and n_t in each period, in order to maximize the expected net present value of its profits,

$$E_0 \sum_{t=0}^{\infty} Q_{0,t} [(y_t - w_t) n_t - f(\tilde{\varepsilon}_t) s_t (K + H(\tilde{\varepsilon}_t))] \quad (10)$$

where $Q_{0,t}$ is the stochastic discount factor as in equation (9), subject to the law of motion for its employment stocks (2),

$$n_t = (1 - \lambda) n_{t-1} + f(\tilde{\varepsilon}_t) s_t \quad (11)$$

where $f(\tilde{\varepsilon}_t) s_t$ the number of new hires in period t . Since each firm is small compared to the overall size of the economy, it takes wages w_t , the stochastic discount factor $Q_{t,t+1}$ and the total number of job seekers s_t as given.

Profit maximization gives rise to the following equilibrium job creation condition, see appendix A.3 for details on the derivation.

$$K + M(\tilde{\varepsilon}_t) = y_t - w_t + (1 - \lambda) E_t [Q_{t,t+1} \{K + M(\tilde{\varepsilon}_{t+1})\}] \quad (12)$$

To aid interpretation, remember that $M(\tilde{\varepsilon}_t) = \tilde{\varepsilon}_t$ in the case of perfectly random hiring. Solving forward, the job creation condition can be written as follows.

$$K + M(\tilde{\varepsilon}_t) = E_t \sum_{\tau=0}^{\infty} (1 - \lambda)^\tau Q_{t,t+\tau} (y_{t+\tau} - w_{t+\tau}) \quad (13)$$

In equilibrium, the hiring costs of the marginal hire, $K + M(\tilde{\varepsilon}_t)$, equals the expected net present value of profits generated from a job.

4.2 Wage Setting

The equilibrium job creation condition (12) may be compared to the efficient job creation condition (7). The two conditions are very similar, and can be made identical by choosing the appropriate wage setting rule. Thus, the only reason why equilibrium job creation is inefficient in this model, is that the wage may deviate from its efficient level. Comparing equations (12) and (7), we obtain the following efficient wage rule.

$$w_t = (1 - \lambda) E_t [Q_{t,t+1} f(\tilde{\varepsilon}_{t+1}) (M(\tilde{\varepsilon}_{t+1}) - H(\tilde{\varepsilon}_{t+1}))] \quad (14)$$

⁴We assume that workers always accept all job offers and there is full commitment of both firm and worker once a job has been created and a wage been agreed upon, see section 3.1.

In general, the efficient wage depends on the distribution and the persistence of training costs, which affect the average and marginal training cost functions $H(\tilde{\varepsilon})$ and $M(\tilde{\varepsilon})$.

In order for job creation to be efficient, it is not necessary that wages depend on individual worker characteristics ε_{it} , as we anticipated in section 3.1. The reason is that only the wage of the marginal worker, with training costs $\tilde{\varepsilon}_t$ is allocative, as long as it is more profitable to hire workers with lower training costs in equilibrium. Therefore, we limit our attention to wages that do not depend on ε_{it} for simplicity.

We argue the efficient amount of job creation can always be supported as an equilibrium with the right wage setting mechanism. Therefore, we must verify that there exist parameter values, for which the efficient wage rule (14) satisfies the full participation condition, $w_t > b_t$ for all t . To see that this is the case, notice that both the job finding probability f and the average training costs H must be strictly increasing in the hiring threshold $\tilde{\varepsilon}_t$, so that by (8) the average training costs of employed workers must be strictly lower than the training costs of the marginal worker if the distribution G is ‘well behaved’.⁵ Therefore, the efficient wage is always positive. Then, we can set unemployment benefits $0 < b_t < w_t$ such that full participation is satisfied.

What kind of wage determination mechanism would give rise to efficient wage rule (14)? Clearly, Nash bargaining would not.⁶ While potentially interesting, the relation between wage setting and labor market efficiency is not the purpose of the present study. Therefore, in the remainder of this section as well as in the next section we assume wages are set such that job creation is efficient, without explicitly specifying the wage determination mechanisms. In the numerical analysis in section 6 we use a simple wage rule of the form $w_t = w(y_t, b_t)$ and verify that this assumption does not affect our results.

4.3 Job Finding Rate

We now have a condition for the hiring threshold $\tilde{\varepsilon}_t$ in equilibrium (12), which under the efficient wage setting rule (14) reduces to the efficient job creation condition (7). The hiring threshold determines the aggregate job finding rate and unemployment rate. In this section, we formalize this link.

The first, and most important, observation is that in our framework, unlike in standard labor market models with search frictions, the job finding rate is not constant across workers. Since firms hire only workers with training costs below the hiring threshold $\tilde{\varepsilon}_t$, the job finding

⁵If G has zero mass at $\tilde{\varepsilon}_t$, then $f'(\tilde{\varepsilon}_{t+1}) = H'(\tilde{\varepsilon}_{t+1}) = 0$ and the efficient hiring threshold is not unique, so that the efficient wage is not well defined. If G has a mass point at $\tilde{\varepsilon}_t$, then the derivatives are infinity and all employed workers may have the same training costs so that the efficient wage equals zero. In all other cases, the efficient wage is positive. For perfectly random hiring, this point is easy to prove formally. Then, the distribution of job seekers with training costs below the hiring threshold $\tilde{\varepsilon}_t$ mirrors the truncated unconditional distribution of ε_{it} below the threshold. Thus, the average training costs and their derivative with respect to the training costs of the marginal hire $\tilde{\varepsilon}_t$, are given by

$$H(\tilde{\varepsilon}_t) = \frac{\int_{-\infty}^{\tilde{\varepsilon}_t} \varepsilon dG(\varepsilon)}{G(\tilde{\varepsilon}_t)} < \frac{\int_{-\infty}^{\tilde{\varepsilon}_t} \tilde{\varepsilon}_t dG(\varepsilon)}{G(\tilde{\varepsilon}_t)} = \tilde{\varepsilon}_t = M(\tilde{\varepsilon}_t)$$

⁶For a similar setup where Nash bargaining does not achieve efficiency, see Chugh and Merkl (2011).

probability of an individual worker f_{it} is either 1 or 0, depending on her training costs ε_{it} .

$$f_{it} = \begin{cases} 1 & \text{if } \varepsilon_{it} \leq \tilde{\varepsilon}_t \\ 0 & \text{if } \varepsilon_{it} > \tilde{\varepsilon}_t \end{cases} \quad (15)$$

The aggregate job finding rate $f(\tilde{\varepsilon}_t)$ is then given by the average of the individual job finding probabilities of all job seekers,

$$f(\tilde{\varepsilon}_t) = \frac{\int_{-\infty}^{\infty} f_{it} s_{it} dG}{\int_{-\infty}^{\infty} s_{it} dG} \quad (16)$$

where s_{it} is the fraction of type ε_{it} workers seeking a job. Notice that $f(\tilde{\varepsilon}_t)$ is the gross job finding rate, which includes workers who lost their job in the current period.

4.4 Random Hiring

In order to evaluate the integrals in (16), we need to know in a given period t how many workers of each type ε_{it} are looking for a job.⁷ The composition of the pool of job seekers depends crucially on the persistence of the individual training costs. If training costs are i.i.d. over time (as well as across workers), then each worker get a new draw for ε_{it} in each period, so that in any given period, the distribution of ε_{it} in the pool of job seekers mirrors the aggregate distribution G . In this case, the number of job seekers as a fraction of workers of each type equals the total number of job seekers as a fraction of the total labor force, $s_{it} = s_t$. In this case, the aggregate job finding rate equals the probability that training costs are below the hiring threshold.

$$f^{\text{RH}}(\tilde{\varepsilon}_t) = \frac{\int_{-\infty}^{\tilde{\varepsilon}_t} 1 \cdot s_t \cdot dG + \int_{\tilde{\varepsilon}_t}^{\infty} 0 \cdot s_t \cdot dG}{\int_{-\infty}^{\infty} s_t \cdot dG} = G(\tilde{\varepsilon}_t) \quad (17)$$

We refer to this case as perfectly random hiring, because at the beginning of the period, each unemployed workers has the same probability of getting a good draw for ε_{it} and therefore the same probability of finding a job, regardless of the current training costs. In other words, at the beginning of the period, it is random which workers will get hired and which will not.

Although we focus primarily on the job finding rate, for completeness we also calculate the steady state unemployment rate. The steady state unemployment rate equals $\bar{u}_t = 1 - \bar{n}_t$, where \bar{n}_t is the steady state fraction of workers that are employed implied by difference equation (2). The steady state unemployment rate for the model with random hiring equals

$$\bar{u}^{\text{RH}} = \frac{\lambda [1 - G(\tilde{\varepsilon})]}{\lambda [1 - G(\tilde{\varepsilon})] + G(\tilde{\varepsilon})} \quad (18)$$

Notice that the number of unemployed workers does *not* equal the number of workers with training costs above the hiring threshold, because many of these worker had lower training costs in the past and are currently still employed because they were hired then.

⁷We use the phrase “looking for a job” or “job seeker” loosely. There are some workers who do not have a job and who have zero (or very low) probability of being offered one, because their training costs are too high. We still include those workers in the pool of unemployed workers as well as job seekers, because at the current wage rate, they would accept a job if it were offered to them.

4.5 Comparison to Models with Search Frictions

Our model with random hiring can be made identical to a standard search and matching model in the tradition of Diamond (1982), Mortensen (1982) and Pissarides (1985) in terms of its predictions for unemployment fluctuations. We show that the job creation equation generated by our model with random hiring is the same as the job creation equation in a standard Pissarides model, if we choose the distribution of training costs G appropriately. The distribution function of worker heterogeneity plays the role of an aggregate matching function in search and matching models.

The job creation condition in a model with search frictions equates the expected net present value of firms' profits $y_t - w_t$ to the expected net present value of vacancy posting costs. Vacancy posting costs may include a fixed component K , which is paid only at the start of the vacancy, but also includes a flow cost k , the expected net present value of which depends on the probability the vacancy is filled in each period q_t .

$$K + \frac{k}{q_t} = y_t - w_t + (1 - \lambda) E_t \left[Q_{t,t+1} \left(K + \frac{k}{q_{t+1}} \right) \right] \quad (19)$$

See Petrosky-Nadeau and Van Rens (2011) for the derivation of this condition in a discrete time version of a standard search and matching model with capital and non-linear utility over consumption, and with the same timing assumptions as in this paper, and Pissarides (2009, section 5) for the role of fixed job creation costs in this type of model.

The job creation condition in the standard search model (19) equals the job creation condition in our model (12) with perfectly random hiring, i.e. with $M(\tilde{\varepsilon}_t) = \tilde{\varepsilon}_t$, if $k/q_t = \tilde{\varepsilon}_t$. The vacancy filling probability q_t in this model depends on labor market tightness θ_t , the ratio of vacancies v_t over the number of unemployed workers u_t , through the matching technology, which relates new matches m_t to the number of unemployed and the number of vacancies. With a standard constant returns to scale Cobb-Douglas matching function, $m_t = u_t^\mu v_t^{1-\mu}$, we get $q_t = m_t/v_t = \theta_t^{-\mu}$. The job finding rate in this model is also related to labor market tightness through the matching function, $f_t = m_t/u_t = \theta_t^{1-\mu}$. Thus, we can write the vacancy filling probability in terms of the job finding rate, $q_t = \theta_t^{-\mu} = f_t^{-\mu/(1-\mu)}$, so that $k/q_t = k f_t^{\mu/(1-\mu)}$. Thus, the job creation condition in our model equals the one from the standard search model if $k/q_t = k f_t^{\mu/(1-\mu)} = \tilde{\varepsilon}_t$ or

$$f_t = \left(\frac{\tilde{\varepsilon}_t}{k} \right)^{\frac{1-\mu}{\mu}} \quad (20)$$

Comparing expression (20) to (17), it is clear that we can choose a distribution function G such that the job creation condition is the same in both models. The distribution that makes job creation conditions identical is $G(\varepsilon) = (\varepsilon/k)^{(1-\mu)/\mu}$ for $0 \leq \varepsilon \leq k$, which means that $1/\varepsilon_{it}$ follows a Pareto distribution. Since the law of motion for employment (2) is also the same in both models, the predictions for (un)employment are identical as well. Thus, our model provides a framework to think about the selectivity of hiring, while maintaining all the insights about unemployment dynamics from standard labor market models.

In our model, worker heterogeneity plays the same role as the congestion externality, modelled through the aggregate matching function, in the standard model. In a boom,

when productivity is high, it becomes harder to hire in the search and matching model because the labor market gets ‘congested’ with vacancies. In our model, hiring is costlier in a boom because firms are forced to hire workers with larger training costs in order to increase employment.

4.6 Selective Hiring

Now consider the opposite polar case, in which individual training costs are fixed over time, $\varepsilon_{it} = \varepsilon_i$. In this case, there are two reasons why a worker may be seeking a job in period t . A worker with training costs above the hiring threshold was unemployed in period $t - 1$ and is therefore a job seeker in period t . Since this worker will have the same training costs in period t as she had in period $t - 1$, she will be very unlikely to be hired in period t . In fact, if the economy is in steady state, the individual job finding probability of these workers will be zero. A worker with training costs below the threshold in period $t - 1$ was employed in that period. However, such a worker may have been separated from her job in period t and consequently is a job seeker as well. Again assuming the economy is in steady state, if this worker was employed in period $t - 1$, she will again be offered a job in period t with probability one. The fraction of ‘good’ workers that are seeking jobs equals λ , the probability that any given existing job is destroyed, so that $s_i = \lambda$ if $\varepsilon_{it} \leq \tilde{\varepsilon}$. Since all ‘bad’ workers seek jobs, $s_i = 1$ if $\varepsilon_{it} > \tilde{\varepsilon}$. Thus, the (steady state) job finding rate in this case is given by the following expression.

$$f^{\text{SH}}(\tilde{\varepsilon}) = \frac{\int_{-\infty}^{\tilde{\varepsilon}} 1 \cdot \lambda \cdot dG + \int_{\tilde{\varepsilon}}^{\infty} 0 \cdot 1 \cdot dG}{\int_{-\infty}^{\tilde{\varepsilon}} \lambda \cdot dG + \int_{\tilde{\varepsilon}}^{\infty} 1 \cdot dG} = \frac{\lambda G(\tilde{\varepsilon})}{\lambda G(\tilde{\varepsilon}) + 1 - G(\tilde{\varepsilon})} \quad (21)$$

We refer to this case as perfectly selective hiring, because at the beginning of the period everyone knows which workers will be hired and which will remain unemployed. Firms pick out the ‘good’ workers, with low training costs, from the pool of job seekers and ignore the ‘bad’ workers.

The steady state unemployment rate for the model with selective hiring equals

$$\bar{u}^{\text{SH}} = 1 - G(\tilde{\varepsilon}) \quad (22)$$

Under selective hiring, the steady state unemployment rate equals the fraction of workers with training costs above the hiring threshold, because these workers will never be hired, whereas all other workers will always be immediately rehired in case they loose their job.

The differences between the model with random and selective hiring are driven by differences in the quality of the pool of job seekers between both models. If hiring is random, the pool of job seekers is a reflection of the overall distribution of workers. If hiring is selective on the other hand, workers with low training costs are unlikely to be unemployed, so that the pool of job seekers consists largely of lemons. How large this difference is depends on the separation rate λ . If $\lambda = 0$, the job-finding rate with selective hiring is equal to zero because all job seekers have training costs that are too high to be hired. If $\lambda = 1$, the job-finding rate is the same under selective and random hiring, because in both cases job seekers are representative for the distribution of all workers.

Comparing expressions (17) and (21) for the job finding rate and (18) and (22) for the unemployment rate, it seems that the models with random and selective hiring have very different predictions for labor market dynamics. This is not true. The difference between the job finding and unemployment rates under selective versus random hiring is mostly a level shift. If we were to use these models to generate a standard set of business cycle statistics for the volatility, persistence and comovement of labor market variables, we would calibrate the model parameters to match the steady state job finding or unemployment rate. The differences in calibration would offset the differences in the expressions, and the predictions of the models would be quite similar.⁸ This does not mean, of course, that the predictions of the two models are the same in all dimensions. In section 2, we discussed observable predictions that allow us to distinguish one model from the other in the data. In addition, the two models have very different implications for welfare analysis, to which we now turn.

5 Welfare Analysis

In this section, we derive the equilibrium consumption choices of workers and compare the resulting consumption allocation to the efficient allocation. In order to obtain simple, easily interpretable expressions, we evaluate the steady state of the model and make some additional simplifying assumptions. We show that, under these assumptions, the equilibrium consumption allocation with random hiring equals the allocation chosen by the social planner, but the equilibrium consumption allocation under selective hiring is far from efficient. The reason is that under selective hiring, unemployment risk is highly unequally distributed across workers. The objective of this section is to make this point in the simplest possible setting. Section 6 presents a numerical analysis to support the results of this section in a more general version of the model.

5.1 Equilibrium Consumption

Each worker i chooses her consumption in each period t in order to maximize the net present value of her utility (1), subject to a budget constraint. In order to smooth their consumption over time, workers may lend or borrow by buying or (short)selling bonds at the risk-free interest rate r . Let A_{it} denote bond holdings of worker i at time t . Then, workers have to satisfy the following dynamic budget constraint,

$$A_{it+1} = (1 + r)(A_{it} + m_{it} - c_{it}) \quad (23)$$

where m_{it} denotes period income of worker i , which equals either the wage w_t , if the worker is employed, or the unemployment benefit b_t , if she is unemployed.

Maximizing (1) subject to (23) as well as the usual no-Ponzi game condition, gives rise

⁸To see this, note that the elasticity of the job finding rate with respect to productivity y_t from equations (17) and (21) equals a constant times the elasticity of the hiring threshold \tilde{e}_t with respect to y_t , which is the same in both models. The proportionality factor is different in the two models, but depends only on the separation rate λ and the shape of the training costs distribution G . In a future version of this paper, we plan to include business cycle statistics for data simulated from a calibrated version of both models in order to make this point more formally, see section 7 for a brief description of this exercise. For now, we rely on intuition to convince the reader of this point.

to the usual Euler equation for consumption.

$$\mathcal{U}'(c_{it}) = \beta(1+r) E_t[\mathcal{U}'(c_{it+1})] \quad (24)$$

We simplify, as in Hall (1978), by assuming the utility function $\mathcal{U}(c)$ is quadratic in consumption and the interest rate equals the discount rate, $\beta(1+r) = 1$. Then, combining (24) with the budget constraint (23), we can write consumption directly as a function of assets and expected future income.

$$c_{it} = \frac{r}{1+r} \left[A_{it} + \sum_{\tau=0}^{\infty} \left(\frac{1}{1+r} \right)^{\tau} E_t m_{it+\tau} \right] \quad (25)$$

Consumption of each worker equals that worker's permanent income: the annuity value of her current assets plus the expected net present value of all her future income.

If the discount rate is sufficiently small, $r \rightarrow 0$, initial conditions do not matter. If in addition, we assume the economy is in steady state, the expression for consumption simplifies considerably further,

$$c_{it} = E_t m_{it+\tau} = u_i b + (1 - u_i) w \quad (26)$$

where u_i is the unemployment rate of type ε_{it} workers. Expected future income, and therefore consumption, depends exclusively on unconditional unemployment risk. With our simplifying assumptions, we have ruled out any welfare costs due to bad luck. We do this on purpose, in order to focus on the welfare loss deriving from the fact that unemployment risk is distributed unequally across workers. Comparing the equilibrium condition (26) to efficiency condition (6), we see that the consumption allocation is efficient, if and only if unemployment risk is distributed evenly across workers.

5.2 Random versus Selective Hiring

The worker-specific unemployment risk depends crucially on the persistence of training costs ε_{it} . In the case of perfectly random hiring, with ε_{it} uncorrelated over time, each worker gets a new draw for ε_{it} in each period, so that the unconditional unemployment risk of each worker equals the aggregate unemployment rate. In the case of perfectly selective hiring, with $\varepsilon_{it} = \varepsilon_i$ fixed over time for each worker, some workers, with low training costs, are always employed, whereas other, with training costs above the hiring threshold, are always unemployed. In this case, individual unemployment risk is highly unequally distributed.

$$u_i^{\text{SH}} = \begin{cases} 0 & \text{if } \varepsilon_i \leq \tilde{\varepsilon} \\ 1 & \text{if } \varepsilon_i > \tilde{\varepsilon} \end{cases} \quad (27)$$

Substituting (26) and individual unemployment risk (27) into the steady state version of welfare function (4), we get welfare under selective and random hiring.

$$\mathcal{W}^{\text{SH}} = u\mathcal{U}(b) + (1-u)\mathcal{U}(w) < \mathcal{U}(ub + (1-u)w) = \mathcal{W}^{\text{RH}} \quad (28)$$

Since wages and unemployment benefits are assumed to be the same random and selective hiring, the inequality follows directly from the concavity of utility \mathcal{U} by Jensen's inequality.

By assuming steady state and a sufficiently small rate of time discount, we have assumed that individual workers can completely self-insure against unemployment risk due to bad luck. However, the differences in unemployment risk between 'good' workers with low training costs and 'bad' workers with high training costs in the model with selective hiring, are uninsurable. Once a worker is born and enters the labor market, her type ε_{it} is observable to all market participants. At that point, for workers with high training costs the bad shock has already realized and they can no longer buy insurance against it. It is this unemployment risk across workers, rather than the unemployment risk over the life-time of a worker, that drives the difference in efficiency between the models with selective and random hiring. A different way to see the same point, is that while the two models are equally efficient in creating jobs, the distribution of job opportunities is more equitable with random hiring.

In the model with selective hiring, there is in some sense a 'missing market' for insurance against individual training costs. Therefore, there is a role for government intervention, insuring unborn workers against a bad draw for their training costs. We analyze this issue formally in the next section, using unemployment insurance policy as an example.

6 Application: Optimal Unemployment Insurance

In the previous sections, we showed that although the predictions of our model for unemployment fluctuations are very similar under random and selective hiring, welfare analysis is very different in the two versions of our model. As a concrete application of this general result, in this section we explore how optimal unemployment insurance differs under (perfectly) random and (perfectly) selective hiring. We assume the government does not observe individual workers characteristics ε_{it} and can only redistribute income based on employment status as a proxy for individual characteristics. By providing unemployment benefits, the government tries to insure workers against a bad draw for their training costs. This is a different motive for unemployment insurance from the intertemporal insurance motive typically considered in the literature, see e.g. Hopenhayn and Nicolini (2009). But of course the government faces a trade-off because unemployment benefits discourage job creation.

The objectives of this section are to illustrate the main result in a concrete application and to explore whether the result holds in a fully specified setup, in which wages are determined endogenously. In this application, we maintain the assumption from section 5 that the economy is in steady state (or, more precisely, we will analyze a static version of our model). In future work, we plan to work out a second application, on the costs of business cycles, in which we relax this assumption as well.

6.1 Ramsey Problem

To derive the optimal unemployment insurance policy, we specify the Ramsey problem for a government that sets its policy instruments, unemployment benefits and lump-sum taxes, subject to its budget constraint, such that the resulting competitive equilibrium is the best possible, in the sense that it maximizes social welfare. Thus, the government chooses b_t and

τ_t to maximize welfare (4). We assume the government needs to run a balanced budget, so that the government budget constraint is given by

$$(1 - n_t) b_t = \tau_t \quad (29)$$

In addition to its budget constraint, the government also takes the optimality conditions for job creation (12) and consumption allocation (24), the market clearing conditions for the labor market (2) and goods market (3), and an equilibrium wage setting rule as constraints on its optimization problem.

In its generality, the Ramsey problem is difficult to solve. Since dynamics do not play an important role for our result, we simplify the problem by considering a static version of our model: we assume that (i) there are no aggregate shocks, i.e. $y_t = y$ for all t , and (ii) employment is always in steady state, i.e. we replace equation (2) by the following equation.

$$n_t = \frac{f(\tilde{\varepsilon}_t) s_t}{\lambda} \quad (30)$$

where s_t is the number of job seekers, which the firm takes as given. Under these assumptions, the model becomes completely static. We will drop time subscripts from here onwards to reflect this.

In the static version of our model, the government chooses b and τ to maximize

$$\int_{-\infty}^{\infty} \mathcal{U}(c_{it}) dG \quad (31)$$

subject to the government budget constraint (29) and the equilibrium conditions of the static model. These equilibrium conditions include the static job creation equation,

$$K + \tilde{\varepsilon} = \frac{y - w}{\lambda} \quad (32)$$

the labor market clearing condition (30), which can be rewritten as,

$$n = \frac{f(\tilde{\varepsilon})}{\lambda + (1 - \lambda) f(\tilde{\varepsilon})} \quad (33)$$

and the consumption of employed and unemployed workers,

$$c_i = \begin{cases} c_u = b + \pi - \tau & \text{if } i \text{ unemployed} \\ c_n = w + \pi - \tau & \text{if } i \text{ employed} \end{cases} \quad (34)$$

under selective hiring, or $c_i = c = (1 - n) b + nw + \pi - \tau$ under random hiring, where π denotes profits, which we assume are redistributed lump-sum from firms to all workers, and which are pinned down by the aggregate resource constraint.

$$(1 - n) c_u + n c_n = yn - [1 - (1 - \lambda) n] f(\tilde{\varepsilon}) (K + H(\tilde{\varepsilon})) \quad (35)$$

The final constraint on the Ramsey problem of the government is a wage setting rule.

In sections 4 and 5, we assumed wages are set such that job creation is efficient. Here, we

deviate from that assumption because such a wage setting rule is probably not realistic and we want to explore whether our results hold up to more standard wage determination mechanism. More importantly, we want to think of the government facing a trade-off between efficient consumption redistribution and efficient job creation. If wages are always set efficiently and do not depend on the level of unemployment insurance, such a trade-off does not exist. Thus, we specify an ad-hoc wage rule, which is loosely inspired by the surplus sharing rule for wages in standard search and matching models.

$$w = \phi y + (1 - \phi) b \quad (36)$$

According to this rule, higher unemployment benefits improve workers' outside option, which drives up their wage and erodes profits, thus discouraging job creation.

6.2 Random Hiring

The first order conditions for the Ramsey problem are straightforward to derive, but messy and hard to interpret. Therefore, we present the results for the the optimal unemployment insurance policy numerically for a calibrated version of the model. This exercise is meant to be illustrative rather than quantitative. For the numerical results, we use logarithmic utility over consumption $\mathcal{U}(c) = \log c$, normalize productivity to $y = 1$, set the separation rate to $\lambda = 0.1$ per quarter and use a uniform distribution for the idiosyncratic component of training costs, $G = U[-4, 4]$. We calibrate the mean training costs $K = 1$ to target a job finding rate of 0.5, which corresponds to an unemployment rate of 9%. In our baseline calibration, we set the parameter of the wage setting rule to $\phi = 0.75$.

Under random hiring, the government can achieve the first best allocation. The reason is that there is no trade-off in this case. Since unemployment risk is equally distributed across all workers, no redistribution is necessary. Moreover, since all workers are unemployed an equal amount of time, unemployment benefits are not redistributive. Thus, unemployment benefits are irrelevant for the consumption allocation and the government can set b in order to implement the efficient level of job creation.

The solid line in figure 1 shows welfare as a function of unemployment benefits for the model with random hiring. The optimal amount of unemployment benefits, which implements the efficient amount of job creation, equals 0.6, which by equation (36) implies a wage of 0.9, so that the earnings of unemployed workers are two thirds of those of employed workers. Of course this result is sensitive to the parameterization. In figure 2 we explore an alternative calibration, setting the parameter of the wage setting rule to $\phi = 0.25$. In this case, efficient job creation is achieved by setting unemployment benefits to 0.87, which also implies a wage of 0.9 but a replacement ratio of 96%.

6.3 Selective Hiring

In the model with selective hiring there is a motive for redistribution, so that the government faces a trade-off: by raising unemployment benefits, the government redistributes income from unemployed to employed workers, but at the time discourages job creation. Therefore, in this case the Ramsey planner cannot replicate the first best allocation, and we would

than under random hiring.

This paper is very much a work in progress. All our results about welfare analysis are for steady state (or for a static version of the model) only. Yet, one of the most interesting applications of our framework is about the costs of business cycles. In order to analyze business cycles, we need to solve a dynamic version of our model with selective hiring, which is the next item on our to-do list. However, we can get an idea about the result of that exercise, by thinking of business cycles as comparative statics, i.e. changes in productivity y_t , which are unanticipated and believed to be permanent by workers and firms in our economy.

We expect the cost of business cycles to be (much) higher under selective hiring than under random hiring. The reason is again that unemployment risk under selective hiring is distributed unequally across workers. When the economy switches from a boom to a recession, the number of workers that are unemployed increases. With selective hiring these workers remain unemployed throughout the recession, so that their asset and consumption levels will decline much more than in models with random hiring, in which unemployment duration only increases marginally in recessions.

However, there are reasons why the cost of business cycles may be lower with selective hiring as well. Comparing the economy with business cycles to an alternative economy without aggregate shocks, there are winners as well as losers from business cycles. If the economy fluctuates, there is a group of workers that are sometimes employed and sometimes unemployed. In the economy without shocks, some of these workers would be employed at all times, but others would be unemployed at all times. Since utility is concave, the welfare gain for the workers that would otherwise always be unemployed should outweigh the welfare loss for the workers that would always be employed in the economy without business cycles. Whether this effect will dominate the previous one, depends on the persistence in individual worker characteristics, i.e. the degree of selectivity in hiring, compared to the degree of persistence in aggregate shocks. We leave this interesting issue for future work.

A Appendices

A.1 Social Planner Problem

The value function and the Bellman equation of the social planner problem (4) are given by

$$V(n_{t-1}; y_t) = \max_{\{\tilde{\varepsilon}_{t+\tau}, \{c_{it+\tau}\}_{i=-\infty}^{\infty}\}_{\tau=t}} E_t \sum_{\tau=0}^{\infty} \beta^{\tau} \int_{-\infty}^{\infty} \mathcal{U}(c_{it+\tau}) dG \quad (37)$$

$$= \max_{\tilde{\varepsilon}_t, \{c_{it}\}_{i=-\infty}^{\infty}} \left\{ \int_{-\infty}^{\infty} \mathcal{U}(c_{it}) dG + \beta E_t V(n_t; y_{t+1}) \right\} \quad (38)$$

where y_t is an exogenous state variable and n_{t-1} an endogenous state variable, with law of motion as in equation (2),

$$n_t = (1 - \lambda) (1 - f(\tilde{\varepsilon}_t)) n_{t-1} + f(\tilde{\varepsilon}_t) \quad (39)$$

Notice that hiring is instantaneous (there are no search frictions in this economy), so that n_{t-1} , not n_t , is the state variable. The hiring threshold and consumption in period t are

chosen subject to the aggregate resource constraint (3).

$$\int_{-\infty}^{\infty} c_{it} dG = y_t n_t - [1 - (1 - \lambda) n_{t-1}] f(\tilde{\varepsilon}_t) (K + H(\tilde{\varepsilon}_t)) \quad (40)$$

Let μ_t denote the multiplier associated with the aggregate resource constraint in period t .

The efficiency conditions resulting from this optimization problem are a set of first order conditions for c_{it}

$$\mathcal{U}'(c_{it}) = \mu_t \quad (41)$$

a first order condition for $\tilde{\varepsilon}_t$

$$y_t - K - H(\tilde{\varepsilon}_t) - \frac{H'(\tilde{\varepsilon}_t) f(\tilde{\varepsilon}_t)}{f'(\tilde{\varepsilon}_t)} + \frac{\beta E_t[V'(n_t; \Omega_{t+1})]}{\mu_t} = 0 \quad (42)$$

an envelope condition for n_{t-1}

$$\begin{aligned} V'(n_{t-1}; y_t) &= (1 - \lambda) \mu_t \{ (1 - f(\tilde{\varepsilon}_t)) y_t + f(\tilde{\varepsilon}_t) (K + H(\tilde{\varepsilon}_t)) \} \\ &\quad + (1 - \lambda) (1 - f(\tilde{\varepsilon}_t)) \beta E_t V'(n_t; y_{t+1}) \end{aligned} \quad (43)$$

and the aggregate resource constraint itself. These conditions define a system of $3 + M$ first-order expectational difference equations in the variables c_{it} , $\tilde{\varepsilon}_t$, μ_t and $V'(n_{t-1}; y_t)$, where $M \rightarrow \infty$ is the number of consumers i in the economy.

The first order condition for c_{it} immediately implies that $c_{it} = c_t$. The level of consumption in period t can be found by substituting this into the aggregate resource constraint, noting that G is a CDF, so that $\int_{-\infty}^{\infty} c_{it} dG = c_t \int_{-\infty}^{\infty} dG = c_t$.

Using the first order condition for $\tilde{\varepsilon}_t$ to substitute out for $E_t[V'(n_t; \Omega_{t+1})]$ in the envelope condition for n_{t-1} , we get an expression for $V'(n_{t-1}; y_t)$.

$$V'(n_{t-1}; y_t) = (1 - \lambda) \mu_t \left[K + H(\tilde{\varepsilon}_t) + (1 - f(\tilde{\varepsilon}_t)) \frac{H'(\tilde{\varepsilon}_t) f(\tilde{\varepsilon}_t)}{f'(\tilde{\varepsilon}_t)} \right] \quad (44)$$

Substituting this expression back into the envelope condition for n_{t-1} , we get an Euler equation for the hiring threshold $\tilde{\varepsilon}_t$.

$$K + M(\tilde{\varepsilon}_t) = y_t + \beta (1 - \lambda) E_t \left[\frac{\mu_{t+1}}{\mu_t} \{ K + (1 - f(\tilde{\varepsilon}_{t+1})) M(\tilde{\varepsilon}_t) + f(\tilde{\varepsilon}_{t+1}) H(\tilde{\varepsilon}_{t+1}) \} \right] \quad (45)$$

Substituting $\mu_t = \mathcal{U}'(c_t)$ gives the efficient job creation equation in the main text.

A.2 Marginal Training Costs under Random Hiring

Under perfectly random hiring, i.e. if ε_{it} is uncorrelated over time for each individual, we can calculate explicitly how the average training cost $H(\tilde{\varepsilon}_t)$ depends on the training cost of the marginal hire. With random hiring, the distribution of job seekers with training costs below the hiring threshold $\tilde{\varepsilon}_t$ mirrors the truncated unconditional distribution of ε_{it} below the threshold. Thus, the average training costs and their derivative with respect to the training

costs of the marginal hire $\tilde{\varepsilon}_t$, are given by the following expression.

$$H(\tilde{\varepsilon}_t) = \frac{\int_{-\infty}^{\tilde{\varepsilon}_t} \varepsilon dG(\varepsilon)}{G(\tilde{\varepsilon}_t)} \quad (46)$$

This implies

$$H'(\tilde{\varepsilon}_t) = \frac{\tilde{\varepsilon}_t g(\tilde{\varepsilon}_t)}{G(\tilde{\varepsilon}_t)} - \frac{\int_{-\infty}^{\tilde{\varepsilon}_t} \varepsilon dG(\varepsilon)}{G(\tilde{\varepsilon}_t)^2} g(\tilde{\varepsilon}_t) = \frac{\tilde{\varepsilon}_t g(\tilde{\varepsilon}_t)}{G(\tilde{\varepsilon}_t)} - \frac{H(\tilde{\varepsilon}_t) g(\tilde{\varepsilon}_t)}{G(\tilde{\varepsilon}_t)} \quad (47)$$

In addition, in section 4 we show that with perfectly random hiring $f(\tilde{\varepsilon}_t) = G(\tilde{\varepsilon}_t) \Rightarrow f'(\tilde{\varepsilon}_t) = g(\tilde{\varepsilon}_t)$. Substituting both results into equation (8), gives $M(\tilde{\varepsilon}_t) = \tilde{\varepsilon}_t$.

A.3 Equilibrium Job Creation

From the Bellman equation

$$V(n_{t-1}; y_t) = (y_t - w_t) n_t - f(\tilde{\varepsilon}_t) s_t (K + H(\tilde{\varepsilon}_t)) + E_t [Q_{t,t+1} V(n_t; y_{t+1})] \quad (48)$$

where

$$n_t = (1 - \lambda) n_{t-1} + f(\tilde{\varepsilon}_t) s_t \quad (49)$$

we get the first order condition for $\tilde{\varepsilon}_t$

$$K + H(\tilde{\varepsilon}_t) + \frac{f(\tilde{\varepsilon}_t) H'(\tilde{\varepsilon}_t)}{f'(\tilde{\varepsilon}_t)} = y_t - w_t + E_t [Q_{t,t+1} V'(n_t; y_{t+1})] \quad (50)$$

and the envelope condition for n_{t-1}

$$V'(n_{t-1}; y_t) = (1 - \lambda) \{y_t - w_t + E_t [Q_{t,t+1} V'(n_t; y_{t+1})]\} \quad (51)$$

Substituting the first order condition into the envelope condition

$$V'(n_{t-1}; y_t) = (1 - \lambda) \left\{ K + H(\tilde{\varepsilon}_t) + \frac{f(\tilde{\varepsilon}_t) H'(\tilde{\varepsilon}_t)}{f'(\tilde{\varepsilon}_t)} \right\} \quad (52)$$

and substituting back into the envelope condition, gives an Euler equation for the equilibrium hiring threshold, which is equation (12) in the main text.

References

- [1] Abraham, K. and R. Shimer (2002), Changes in Unemployment Duration and Labor-Force Attachment, in: Krueger, Alan and Robert Solow (eds.), *The Roaring Nineties: Can Full Employment Be Sustained?* Russell Sage Foundation, 367-420.
- [2] Brown, A.J.G., C. Merkl and D.J. Snower (2010). An Incentive Theory of Matching, Kiel Institute for the World Economy working paper.
- [3] Diamond, P.A. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy* 90 (5), 881-94.

- [4] Hall, R.E. (1978). Stochastic Implications of the Life-Cycle Permanent-Income Hypothesis, *Journal of Political Economy* 86(6).
- [5] Hopenhayn, H.A. and J.P. Nicolini (2009). Optimal Unemployment Insurance and Employment History, *Review of Economic Studies* 76(3), pp.1049-1070.
- [6] Hosios, A.J. (1990). Factor Market Search and the Structure of Simple General Equilibrium Models, *Journal of Political Economy* 98(2), pp.325-355.
- [7] Krusell, P. and A. Smith (1998), Income and Wealth Heterogeneity in the Macroeconomy, *Journal of Political Economy*, Vol. 5, 867-896.
- [8] Landais C., P. Michaillat and Emmanuel Saez (2010). Optimal Unemployment Insurance over the Business Cycle, NBER Working Paper 16526.
- [9] Lucas, R.E. (1987), *Models of Business Cycles*. New York: Basil Blackwell, 1987
- [10] Mortensen, D.T. (1982). Property rights and efficiency in mating, racing, and related games. *American Economic Review* 72 (5), 968–79.
- [11] Mueller, A. (2010). Separations, Sorting and Cyclical Unemployment, unpublished manuscript, Columbia University.
- [12] Petrosky-Nadeau, N. and T. van Rens (2011). A Search Model with Diminishing Returns to Labor, CREI working paper.
- [13] Pissarides, C.A. (1985). Short-run equilibrium dynamics of unemployment vacancies, and real wages. *American Economic Review* 75 (4), 676–90.
- [14] Pissarides, C.A. (2000). *Equilibrium Unemployment Theory* (2nd ed.). Cambridge: MIT Press.
- [15] Pissarides, C.A. (2009). The unemployment volatility puzzle: Is wage stickiness the answer? Walras-Bowley lecture, North American summer meetings of the Econometric Society, Duke University. *Econometrica* 77 (5), 1339-1369.
- [16] Shimer, R. (2007), Reassessing the Ins and Outs of Unemployment. NBER Working Paper, No. 13421.
- [17] Solon, G., R. Barsky and J. Parker (1994), Measuring the Cyclical of Real Wages: How Important is Composition Bias? *Quarterly Journal of Economics*, Vol. 109, 1-25.
- [18] Wilke, R. (2005), New Estimates of the Duration and Risk of Unemployment for West-Germany, *Schmollers Jahrbuch*, Vol. 125, 207-237.

Figure 1. Welfare as a function of unemployment benefits, $\phi = 0.75$




Figure 2. Welfare as a function of unemployment benefits, $\phi = 0.25$

