

Productivity Gains From Unemployment Insurance*

Daron Acemoglu
Department of Economics
M.I.T.

Robert Shimer[†]
Department of Economics
Princeton

November 1999

Abstract

This paper argues that unemployment insurance increases labor productivity by encouraging workers to seek higher productivity jobs, and by encouraging firms to create those jobs. We use a quantitative model to investigate whether this effect is comparable in magnitude to the standard moral hazard effects of unemployment insurance. Our model economy captures the behavior of the U.S. labor market for high school graduates quite well. With unemployment insurance more generous than the current U.S. level, unemployment would increase by a magnitude similar to the micro-estimates; but because the composition of jobs also changes, total output and welfare would increase as well.

Keywords: efficiency, risk-aversion, search, unemployment insurance, consumption smoothing

JEL Classification: D83, J64, J65.

*This paper was prepared for the International Seminar on Macroeconomics 1999. Thanks to seminar participants there, at the Federal Reserve Bank of Cleveland, the London Business School, the New York Area Macro Workshop, Pennsylvania, Princeton, and the NBER Summer Institute, and to Gadi Barlevy, Olivier Blanchard, Charles Bean, Robert Gordon, Jonathan Gruber, Per Krusell, Dale Mortensen, James Poterba, two anonymous referees, and the editor for comments.

[†]Corresponding Author: 204 Fisher Hall, Princeton University, Princeton, NJ 08544, United States. Phone: 609-258-4031. Fax: 609-258-6419. Email: shimer@princeton.edu

1 Introduction

This paper argues that unemployment insurance increases labor productivity by encouraging workers to seek higher productivity jobs, and by encouraging firms to create those jobs. We use a quantitative model to investigate whether this effect is comparable in magnitude to the standard moral hazard effects of unemployment insurance (UI). Our results indicate that a decrease in the generosity of UI from its current U.S. level would not only decrease welfare but also reduce the level of output.

Most analyses of unemployment insurance focus on its consumption-smoothing and risk-sharing roles. For example, Gruber (1997) finds that workers who receive more generous unemployment benefits experience a smaller drop in consumption following the loss of a job. Standard approaches to unemployment insurance compare this benefit with the adverse moral hazard effects, and compute the optimal amount of UI by equating marginal costs and benefits (Shavell and Weiss 1979, Hopenhayn and Nicolini 1997).

While this tradeoff is likely to be important, UI could also affect the type of jobs that workers look for and accept.¹ According to the theory we develop, in an economy without any UI, workers avoid the risk of unemployment by applying to low productivity jobs that are easier to obtain. Firms offer implicit insurance to workers by opening jobs with low unemployment risk, and charging an insurance premium in the form of lower wages. The resulting composition of jobs is inefficient and can be improved by a moderate level of unemployment insurance, which encourages workers to take on more risk, and increases both welfare and the level of output (see also Acemoglu and Shimer 1999a).

Although this effect is qualitatively reasonable, a major goal of this paper is to show that it is likely to be quantitatively important as well. In a realistic environment, it must outweigh two significant forces. First, unemployment insurance will encourage workers to reduce their search effort, lowering employment and output. Second, workers can self-insure by saving, considerably reducing the need for unemployment insurance.

To address these issues, we consider a quantitative dynamic economy. Workers are risk averse, with constant relative risk aversion (CRRA). They optimally choose their consumption, labor supply, and search effort while unemployed. Unfortunately, this model cannot be solved analytically,² because the workers' optimal policy (their consumption, labor supply,

¹In the data, Ehrenberg and Oaxaca (1976) find that workers who receive more UI find higher wage jobs. On the other hand, Meyer (1989) finds no evidence that more generous UI results in higher earnings for covered workers. His estimates have large standard errors, however, and cannot rule out substantial earnings effects.

²In Acemoglu and Shimer (1999a), we provided closed form solutions for the case with constant absolute

and search rules) depends on their wealth level, which is itself determined by the optimal policy. We therefore undertake a calibration exercise, anchoring our model to plausible preferences and to the unemployment rate and unemployment insurance system faced by U.S. workers with a high school degree. With a coefficient of relative risk aversion of four, a replacement rate a little below fifty percent for six months, and productivity differences between good and bad jobs on the order of 30%, our model generates levels of unemployment, the consumption drop upon job loss, and low-frequency income variability similar to those found in U.S. data.

Moderate increases in the replacement rate or the duration of unemployment benefits lead to increases in unemployment duration similar to those observed in the U.S. economy. For example, Meyer (1989) finds that a 10 percent increase in unemployment benefits raises unemployment duration by about a week, while Ehrenberg and Oaxaca (1976) and Atkinson and Micklewright (1991) estimate a slightly smaller response. At the same time, the policy change raises average wages by about 1.2 percent, which is a substantial effect, but still somewhat less than the gains reported by Ehrenberg and Oaxaca (1976). The overall effect of the policy change is to raise output and welfare by a little over half a percent.

The result that economies with moderate UI have higher output and welfare than those without social insurance is very robust. Increasing the value of leisure, reducing risk aversion, reducing the wage gap between different types of jobs, and allowing on-the-job search does not alter this conclusion, although it does sometimes affect whether the output-maximizing replacement rate is above or below U.S. levels.

Our paper is related to a number of previous studies. As mentioned above, we build on Acemoglu and Shimer (1999a). Other papers have also pointed out beneficial effects of UI, *inter alia* Diamond (1981), Acemoglu (1998), and Marimon and Zilibotti (1999). For example, Acemoglu shows that UI may improve welfare by encouraging workers to wait for higher capital jobs that pay higher wages because of holdup problems. Marimon and Zilibotti emphasize matching between workers and firms according to comparative advantage, and show that UI encourages workers to wait for jobs better suited to them. All these papers consider risk-neutral agents, however, so unemployment benefits are simply a subsidy to search. Our approach differs in explicitly modeling risk aversion and precautionary saving, and in contrasting the beneficial effects of UI with its conventional costs. We are not aware of any other study that has compared these costs and benefits in such a realistic setting.

Other studies, including Shavell and Weiss (1979), Hansen and Imrohroglu (1992), Atkinson and Lucas (1995), and Hopenhayn and Nicolini (1997) analyze optimal unemployment insurance with asymmetric information, but do not model labor market search. Costain

risk aversion. In this paper, we use the more conventional CRRA preferences, which imply that poorer workers may be unwilling to accept gambles that richer workers find attractive. The utility cost of low consumption is also much larger with CRRA preferences, increasing the role of insurance.

(1996), Valdivia (1996), and Gomes, Greenwood, and Rebelo (1997) examine labor market behavior in a quantitative general equilibrium framework, but do not look at the productivity gains from unemployment insurance.

Our emphasis on the importance of uninsured risk may appear to contradict the findings of Aiyagari (1994), Huggett (1993), and Krusell and Smith (1998). These papers show that wealth heterogeneity and risk are unlikely to have an important effect on the behavior of aggregate macroeconomic variables. They consider general equilibrium models with an endogenous interest rate. If all workers have the same rate of time preference and there are no exogenous borrowing constraints, the interest rate will be quite close to the discount rate. The precautionary savings motive then induces workers to build up a large buffer stock of assets, insulating them from risk. In contrast, we focus on the behavior of low skill workers (e.g., high school graduates) rather than the whole economy. We use a partial equilibrium model, with an exogenous interest rate, significantly less than the discount rate.³ Workers maintain low asset levels, which provides an accurate characterization of the behavior of preretirement, low skill workers (see the numbers given in footnote 11). Since self-insurance is imperfect, agents change their behavior to avoid risk.

Another significant difference between our model and the previous literature, is that in these other papers, agents only affect macroeconomic outcomes through their savings decisions.⁴ Since poor agents own very little of the aggregate capital stock, they only have a limited effect on aggregate outcomes. In contrast, in our economy poor agents' behavior determines the composition of jobs, and therefore has a large effect on aggregate income.

In the next section, we start with a static model that illustrates the main tradeoffs that we focus on. We outline our full dynamic model and characterize the solution in Section 3. The model is calibrated to features of the U.S. labor market for high school graduates and the U.S. unemployment insurance system in Section 4. We show that the model performs well and matches a variety of salient patterns of the labor market that we did not use in our choice of parameters. We also find in this section that moderate increases in UI generosity starting from U.S. levels could increase output, productivity, and welfare. In Section 5, we show that our general results are robust to introducing on-the-job search, to reducing the wage gap between different types of jobs, and to lower levels of risk aversion. Section 6 concludes.

³In a previous version of the paper, we endogenized the interest rate by introducing home production income and assuming that riskless bonds are in zero net supply. With small values of home production income, there was a large wedge between the discount rate and the equilibrium interest rate, so the results were identical to those presented here.

⁴Krusell and Smith (1998) consider labor supply choices in an extension to their basic model. They find, as we do, that labor supply varies little with wealth levels.

2 Static Model

To illustrate the qualitative benefits of UI, we start with a simple static model. There is a continuum 1 of identical workers, each with the von Neumann-Morgenstern utility function $U(c, h)$ over final consumption c and work hours, h . U satisfies standard assumptions: it is continuously differentiable, strictly increasing in c , and weakly concave. For now, we restrict $h = \widehat{h}$, so $u(c) = U(c, h)$, but variations in hours will play an important role in the quantitative dynamic model.

A worker's consumption is the sum of her initial assets A , plus her net labor income y . Depending on whether she finds a job, her income will either be equal to her wage w net of proportional UI taxes τ , or to her unemployment benefit b .

There is a larger continuum of potential firms, each with access to the same production technology. Each firm can open a job with 'specificity' $\alpha \in [0, 1]$ which then produces $g(\alpha)$ units of output when filled. A job with higher α produces more output, so g is an increasing function. However, a high α job is also harder to fill. In particular, it requires a better match between the firm and its employee, so the probability that a random worker possesses the skills and abilities required for a job of specificity α is given by the decreasing function $M(\alpha)$. Moreover, these specific skills are 'inspection' goods; workers do not know before applying for the job whether they will be a good fit.

Workers and firms come together via search. The timing of events is as follows. At the beginning of each period, each firm decides whether to open a job. Those that open a job incur a cost $\gamma > 0$. An active firm j then chooses its specificity α_j and posts a wage w_j . Each worker observes all the wage offers and associated specificities, and decides where to apply for a job. In a pure strategy equilibrium, each worker applies to a different job, and each job has one applicant.⁵ After the matching stage, the pair learns whether the worker has the requisite skills. If she doesn't, both remain unmatched. If she does, the pair produce $g(\alpha)$.

In equilibrium, each firm that creates a job chooses a specificity α and wage w that will attract a worker and maximize its profit. Also, profits are driven to zero by the free entry

⁵There is also a mixed strategy equilibrium in which there is more competition for some jobs than for others, see Peters (1991), Montgomery (1991), Burdett, Shi, and Wright (1997), or Acemoglu and Shimer (1999b). In such an equilibrium, matching probabilities are determined by the relative supply of workers seeking and jobs offering each specificity-wage combination. We simplify the analysis here by ignoring the mixed strategy equilibrium and the corresponding analysis of market tightness.

condition. Thus, an equilibrium specificity-wage combination (α, w) must solve⁶

$$\begin{aligned} \max_{\alpha, w} \quad & M(\alpha)u(A + (1 - \tau)w) + (1 - M(\alpha))u(A + b) \\ \text{subject to} \quad & M(\alpha)(g(\alpha) - w) \geq \gamma \end{aligned} \tag{1}$$

That is, an equilibrium (α, w) maximizes the expected utility of a worker, subject to the constraint that the firm makes non-negative profits. Otherwise, a firm could offer (α', w') that yields higher utility to workers, while ensuring positive expected profits for itself.⁷

This problem can be represented in a simple diagram like Figure 1. The objective is a downward sloping function in (α, w) space, as workers prefer higher wages and less specific jobs. On the other hand, firms prefer lower wages, and face a tradeoff regarding specificity: higher specificity makes it more difficult to hire a worker, but also raises productivity upon hiring a worker. The non-monotonic shape of the constraint highlights this tradeoff.

[Figure 1 about here.]

An alternative representation is obtained by solving the constraint for the lowest specificity consistent with a particular wage and zero profits, $\alpha(w)$. Substitute this into the matching function to obtain $m(w) \equiv M(\alpha(w))$. This simplifies the characterization of equilibrium to:

$$\max_w \quad m(w)u(A + (1 - \tau)w) + (1 - m(w))u(A + b). \tag{2}$$

This form of the problem is convenient, because it emphasizes the parallel between our formulation and that of Lucas and Prescott (1974), in which workers search across a set of exogenously given wages. In our problem, wages are determined by firms' profit maximization. Nevertheless, neither assets nor the level of unemployment insurance affects the wage-specificity relationship $\alpha(w)$. As a result, the reduced form matching function $m(w)$ is unaffected by policy, and can be taken as given for the purposes of this paper.

Comparative statics with respect to unemployment benefits and asset levels are now straightforward. In terms of the problem in (1), the asset level and unemployment benefit only affect the objective function. Hence comparative statics can be conducted only by

⁶The equilibrium of this model need not be unique, even with homogeneous workers and *ex ante* identical firms. Nevertheless, workers have the same utility in any equilibrium; and more importantly, the comparative statics given below hold across the *set* of equilibria. Acemoglu and Shimer (1999a) prove this result and the generic uniqueness of equilibrium in a more general model.

⁷A potential issue is why firms do not offer insurance by promising payments to applicants who are not well-suited to the job. In practice, there a number of reasons why this may not be feasible. First, such promises may not be credible if firms can declare bankruptcy. Second, they may create adverse selection problems, attracting workers who are of lower ability, hence unlikely to be suited to the job. Third, they may create moral hazard problems, as workers attempt to understate their suitability.

tracing the movements of the indifference contours. (In terms of (2), they do not affect the function $m(w)$.) An increase in the unemployment benefit makes workers more willing to seek higher wages and more specific jobs, flattening the indifference curves. Therefore, when unemployment income increases, the tangency point shifts up and to the right, as in Figure 2. Equilibrium wages increase, and workers seek jobs requiring more specific skills (see Acemoglu and Shimer (1999a) for more details). Intuitively, without UI, workers are unwilling to apply for specific jobs, because these jobs entail too much unemployment risk. UI makes workers more willing to endure the possibility of unemployment. This unambiguously increases labor productivity and unemployment.

[Figure 2 about here.]

Attitudes towards risk have a similar impact on behavior. Less risk-averse workers are more willing to accept a tradeoff of higher unemployment risk for a higher wage. That is, a less risk averse worker has a flatter indifference curve. Again, this shifts the tangency point up and to the right, raising wages and increasing specificity. Intuitively, risk-averse workers dislike the risk of unemployment, and so are willing to pay an insurance premium to firms, in the form of lower wages, in return for a higher employment probability. Firms provide this insurance by reducing job specificity.

A change in the level of assets affects the equilibrium by altering workers' attitude towards risk. With constant relative risk aversion, the most common preferences used in macro models and the one that we will use in our quantitative exercise, richer workers have lower absolute risk aversion. As a result, they are more willing to accept employment gambles, compared to poorer workers. They apply to higher wage jobs with higher unemployment risk. In the dynamic model, we show that this creates a natural source of persistence in individual wealth levels.

We can also show that in an economy without unemployment insurance, risk-averse workers will apply to jobs with too little specificity. Moderate levels of UI will not only improve risk-sharing, but also increase the level of output in the economy. The simplest way to see this is to consider an economy where all agents have the same level of assets A . As a method of proof, consider first an economy in which agents are risk-neutral and have no UI, $b = \tau = 0$. From (2), we see that the equilibrium of this economy maximizes $m(w)w$, total output. Since we know that an economy with more risk-averse agents has lower wages and less job specificity, it is necessarily 'inefficient' (i.e. fails to maximize output). Introducing UI raises wages and specificity. A moderate level of UI will return the productive sector back to the output maximizing allocation. At the same time, it raises workers' payoff while unemployed, reducing risk and raising utility. This beneficial effect of unemployment insurance obviously ignores the 'standard' moral hazard effect on search effort. The dynamic model introduces these adverse effects, and investigates whether, for plausible parameterizations,

the beneficial effects outlined here outweigh the standard costs.

3 Dynamic Model

We now consider our full dynamic model, which contains the benefits of UI demonstrated in the previous section, but also incorporates moral hazard, precautionary savings, and later, on-the-job search.

Workers are infinitely lived, and maximize expected utility, which depends on consumption $C > 0$ and hours $h \in [0, \bar{h}]$, spent either working or searching for a job:

$$E_t \sum_{s=0}^{\infty} \beta^s U(C_{t+s}, h_{t+s}) = E_t \sum_{s=0}^{\infty} \beta^s \frac{(C_{t+s}(\bar{h} - h_{t+s})^\eta)^{1-\theta}}{1-\theta}. \quad (3)$$

θ is the coefficient of relative risk aversion, β is the discount factor, \bar{h} is workers' endowment of time, and η is the relative value of leisure compared to consumption. When $\theta = 1$, preferences are logarithmic with $u(C, h) = \log C + \eta \log(\bar{h} - h)$.

These standard preferences ensure that workers' labor supply will be unchanged by aggregate income growth (i.e., along a balanced growth path). This restriction also implies that when the coefficient of relative risk aversion θ is greater than one, an unemployed worker who is forced to reduce her consumption will have a higher return to leisure, potentially exacerbating the moral hazard problem created by unemployment insurance.

To simplify the computation of equilibrium, we assume that the technology in the economy is such that there exists only two types of jobs. Also, since as shown above, the relationship between specificity and wages is pinned down by firms' zero profit condition and, crucially, is unaffected by policy and assets, we directly work with wages. In particular, there are two types of jobs; 'good' with hourly wage W_{gt} and 'bad' with wage $W_{bt} < W_{gt}$. At any point in time, a worker may be in one of four employment states; she may be employed in a good or bad job, or she may be unemployed and either eligible or ineligible for unemployment benefits B_t .⁸ If she is eligible, her labor income Y_t is equal to her unemployment benefits, $Y_t = B_t$. Otherwise, $Y_t = 0$. Employed workers pay taxes at the rate τ that are used to finance the unemployment insurance system, so their labor income is $Y_t = (1 - \tau)h_t W_{gt}$ or $(1 - \tau)h_t W_{bt}$, where h_t is their hours of work.

The worker then faces a sequence of intertemporal budget constraints:

$$A_{t+1} = R(A_t + Y_t - C_t). \quad (4)$$

⁸We assume that the level of unemployment benefits is not tied to a worker's earlier wage. This reflects the fact that in most countries, the replacement rate falls substantially as earnings increase. It also simplifies our analysis by reducing the number of state variables. Moreover, by restricting attention to a proportional labor income tax and a single unemployment benefit level, we stack the cards against finding benefits from unemployment insurance. Optimal public policy could do better than our calculations suggest.

A_t is the assets at the start of period t . These increase with her labor income Y_t and decrease with consumption C_t . R is the exogenous gross rate of return on risk-free bonds, the only asset in the economy.⁹ Workers may buy and sell bonds, subject only to the solvency ('no-Ponzi game') constraint that their debt cannot explode with positive probability. Since there is always a positive probability of being unemployed and ineligible for unemployment benefits for an arbitrarily long period of time, *any* positive debt may explode. Thus the solvency constraint prevents indebtedness, $A_t \geq 0$.

Next consider the transition of workers between employment states. An unemployed worker finds a job at a rate proportional to the number of hours she searches h_t .¹⁰ The proportionality constant depends on whether she looks for good jobs (m_g) or bad jobs (m_b). Conversely, good and bad jobs end exogenously according to a Poisson process with arrival rates δ_g and δ_b , respectively, leaving the worker unemployed but eligible for unemployment insurance. A worker in turn loses her eligibility with probability ϕ in each period that she is unemployed, so the expected unemployment benefit duration is $1/\phi$ periods. However, any employment spell restores her eligibility. Finally, a worker may voluntarily quit a job at any time, in which case she is ineligible for benefits.

We assume productivity grows at a constant rate $\Gamma > 1$, so $w_g \equiv \Gamma^{-t}W_{gt}$ and $w_b \equiv \Gamma^{-t}W_{bt}$. Unemployment benefits grow at the same rate, with $b \equiv \Gamma^{-t}B_t$. Under these conditions, there is a balanced growth path, in which a worker's consumption and leisure decisions only depend on her employment state and on her normalized asset level $a_t \equiv \Gamma^{-t}A_t$, and the government budget is balanced at each point in time.

We characterize the equilibrium using Bellman equations. Consider first the expected present value of a worker in a good job as a function of her assets A_t and the current time period:

$$\mathcal{J}_{gt}(A_t) = \max_{C_t, h_t} \frac{(C_t(\bar{h} - h_t)^\eta)^{1-\theta}}{1-\theta} + \beta((1-\delta_g)\mathcal{J}_{g,t+1}(A_{t+1}) + \delta_g\mathcal{J}_{i,t+1}(A_{t+1}))$$

where $A_{t+1} \equiv R(A_t + (1-\tau)h_tW_{gt} - C_t)$ is next period's assets. The value of a good job comes from the utility of current consumption and leisure plus the continuation value, which is the probability that the job is exogenously terminated, δ_g , times the value of an unemployed, insured worker $\mathcal{J}_{i,t+1}(A_{t+1})$, plus $(1-\delta_g)$ times the value of a good job next period.

Let $a_t \equiv \Gamma^{-t}A_t$, $c_t \equiv \Gamma^{-t}C_t$, and $J_x(a_t) \equiv \Gamma^{-(1-\theta)t}\mathcal{J}_{xt}(A_t)$, $x \in \{g, i\}$. Our balanced

⁹In an earlier version of this paper, we endogenized the return on bonds, and found that it had little effect on our results. In any case, we will parameterize the model to describe the behavior of workers with a high school diploma. Such workers own a tiny fraction of the aggregate capital stock, and so a change in the unemployment insurance system for these workers will have little effect on the real interest rate (Krusell and Smith 1998).

¹⁰For now, we assume employed workers cannot search. We introduce on-the-job search in Section 5.4, with little effect on the results.

growth assumption yield an autonomous Bellman equation:

$$J_g(a) = \max_{c,h} \frac{(c(\bar{h} - h)^\eta)^{1-\theta}}{1-\theta} + \beta\Gamma^{1-\theta}((1 - \delta_g)J_g(a_+) + \delta_g J_i(a_+)) \quad (5)$$

where $a_+ \equiv (R/\Gamma)(a + (1 - \tau)hw_g - c)$ is next period's normalized assets. Notice that the growth rate of productivity acts to reduce the discount factor (if $\theta > 1$) and the interest rate, thereby raising desired consumption today. The autonomous Bellman equation for bad jobs is identical, except that these workers may also choose to quit.

For employed workers, the relationship between consumption and hours is given by a static first order condition:

$$\frac{\eta c}{\bar{h} - h} = (1 - \tau)w \quad (6)$$

where w is the current wage rate. The left hand side is the marginal rate of substitution between consumption and leisure, while the right hand side is the relative price of leisure, the after-tax wage. Workers will work more hours when consumption is lower or the after-tax wage is higher. Thus we expect to see poor workers in good jobs working the hardest in this economy. In addition, if employed workers approximately consume their after-tax labor income, $c \approx (1 - \tau)hw$, they will supply about $h \approx \bar{h}/(1 + \eta)$ hours of labor. This explains why we interpret η as the value of leisure.

The autonomous Bellman equation for unemployed workers is similar. For those ineligible for unemployment benefits:

$$J_u(a) = \max_{c,h} \frac{(c(\bar{h} - h)^\eta)^{1-\theta}}{1-\theta} + \beta\Gamma^{1-\theta} \max \left\langle m_g h J_g(a_+) + (1 - m_g h) J_u(a_+), \right. \\ \left. m_b h J_b(a_+) + (1 - m_b h) J_u(a_+) \right\rangle \quad (7)$$

where $a_+ = (R/\Gamma)(a - c)$. The first term gives their current utility, while the second term is the continuation value, discounted to account both for impatience and growth. An unemployed worker must choose whether to seek a good job or a bad one. This affects the job finding rate as well as the continuation value upon finding a job. The Bellman equation for insured workers has the same structure, but adds a term to allow for the possibility of losing benefit eligibility, and adds benefits to the asset accumulation equation.

4 Computation of Benchmark Equilibrium

This section describes a calibrated version the model, choosing parameters in line with U.S. data for workers with exactly a high school degree, approximately forty percent of the U.S.

labor force. We focus on this education category because it is relatively homogeneous, and our model presumes that all workers have the same ‘skill level’. In addition, these workers save less,¹¹ yet suffer more unemployment than their more educated peers. This means that they are less likely to do a good job of self-insuring against labor income shocks, and more likely to benefit from unemployment insurance programs. In Section 5, we change key parameter values to verify the robustness of our results.

4.1 Parameterization

We interpret a period to be a week, in order to have some hope of capturing the multitude of very short duration unemployment spells; on average, over 40 percent of unemployed workers in the U.S. have been unemployed for less than five weeks. We normalize the time endowment to $\bar{h} = 1$ and the maximum weekly wage in a bad job to $w_b = 1$. We then must choose the following parameters:

- the interest rate R .
- Technology parameters: the productivity growth rate Γ ; the wage in good jobs w_g ; the destruction rates δ_b and δ_g ; and the matching rates m_b and m_g .
- Preference parameters: the discount rate β ; the coefficient of relative risk aversion θ ; and the relative value of leisure η .
- Policy parameters: the benefit level b ; and the exhaustion rate of benefits ϕ .

In addition, the tax rate τ must balance the government budget.¹²

We set the gross interest rate to $R = 1.0006$ per week, which corresponds to an annual interest rate of 3.2%, approximately the real return on 10-year treasury bonds over the past ten years (IMF 1995). This is the interest rate used in other calibrations of precautionary savings models (e.g. Gourinchas and Parker 1997). The growth rate is $\Gamma = 1.0003$ per week, or 1.6% per year, which matches the average growth rate of labor productivity from 1960–96 (OECD 1997).

We set the wage in a specialized job to be 30% higher than the wage in an unspecialized job, $w_g = 1.3$. This yields a wage differential between specialized and unspecialized jobs

¹¹High school graduates who are less than 40 years old have very low levels of liquid assets. According to the PSID, the median level of financial wealth (bank accounts, stocks, and bonds) excluding debt was \$700 for this group in 1984. It rose to \$1200 by 1989 and \$1325 in 1993. Older workers generally have more assets, in large part because of their retirement savings. Since our model does not have retirement, we focus on data for younger workers without a lifecycle savings motive. We thank Annette Vissing-Jorgensen for providing us with these numbers.

¹²The budget balancing tax rate is generically not unique, due to ‘Laffer curve’ reasoning. A very high tax rate chokes off most production, leading to a high unemployment balanced budget equilibrium. We always look for the lowest tax rate consistent with a balanced budget.

approximately equal to the difference between average wages in the manufacturing and retail sectors (See Tables 1 and 2 in Krueger and Summers 1988).¹³

We assume that all jobs are destroyed with 0.5% probability per week, $\delta_g = \delta_b = 0.005$. These numbers are in line with the flows reported by Poterba and Summers (1986), who find a monthly rate of 0.019 (see their Table 5).¹⁴

We set the matching rate m_g and m_b to be consistent with two facts. First, the equilibrium unemployment rate should be approximately equal to the unemployment rate among high school graduates over the age of 25 in the U.S., which averaged 6.4% from 1979 to 1997. This pins down the ‘average’ matching rate. Second, Krueger and Summers (1988) report that the *observed standard deviation* of log wages is about 0.13. In order to achieve this in our model, with $w_g = 1.3w_b$, about half of employment must be in specialized jobs, and half in unspecialized jobs. This will only happen if unspecialized jobs are sufficiently easier to get, that is the difference between m_g and m_b is relatively large, so that poor (hence more risk-averse) workers are willing to search for them. In our baseline specification, the fraction of specialized jobs is indeed about 50%.

Our choice of the preference parameters β and θ is motivated by two facts. First, high school graduates maintain very low levels of financial wealth (see footnote 11), despite the relatively large income risk that they face. This implies that they must be quite impatient, so we set the weekly discount factor at $\beta = 0.998$, about 0.9 at annual frequencies. This is lower than the typical discount factors used in business cycle analyses, but fairly conventional in the precautionary savings literature. For example, Deaton (1991) uses a 10% discount rate in his simulations, arguing as we do, that this is necessary to justify the number of workers with low asset levels in the presence of liquidity constraints. Carroll and Samwick (1997) argue that one requires a discount rate of about 11% in order to rationalize the fact that the size of agents’ precautionary wealth ‘buffer stocks’ are fairly insensitive to the extent of their income uncertainty. This level of impatience is enough to moderate the precautionary savings motive. Second, workers smooth their consumption quite well (Gruber 1997), despite the relatively low levels of assets. This implies that they must have a low intertemporal elasticity of substitution $1/\theta$. Thus we set the coefficient of relative risk aversion at $\theta = 4$, slightly higher than usual, but once again in line with parameters used in the precautionary savings literature.

¹³Krueger and Summers (1988) and Gibbons and Katz (1992) find that workers who change industry obtain approximately the same wage change as the cross-sectional difference (see, for example, Table 5 in Krueger and Summers 1988). Nevertheless, some of the interindustry wage differences are likely still due to unobserved worker heterogeneity. For this reason, we do not use the largest differentials in the data. The differential between sectors such as tobacco, petroleum, chemicals, on the one hand, and retail, on the other, is on the order of 50 percent.

¹⁴These rates are lower than the worker turnover rates reported by Anderson and Meyer (1994), because we are concerned only with movement from employment to unemployment.

We set the value of leisure to $\eta = 1/6$. According to equation (6), an employed worker who consumes her after-tax income, $(1 - \tau)hw$, will work about $\bar{h}/(1 + \eta) = 0.86$ hours. Interpreting the maximum feasible work week \bar{h} to be 45 hours, workers supply about 38 hours of labor each week. This elasticity is lower than conventionally used (e.g. Hansen 1985). With a higher elasticity of leisure, unemployed workers who are eligible for unemployment insurance will actually consume less than ineligible workers, but offset this with much more leisure. This implausible result follows from the fact that labor and leisure are substitutes, and is clearly counterfactual. Studies such as Hansen and Imrohoroglu (1992) do not encounter this problem because they assume that unemployed workers receive job offers at an exogenous rate, and allocate all of their time to leisure.

Finally, we parameterize unemployment insurance to match U.S. policy. To get an average benefit duration of six months, we set the probability that a worker’s benefits will be exhausted at 4% per week, $\phi = 0.04$. We set the unemployment benefit to $b = 0.4$ in our benchmark model. Since we find that employed workers set their hours to about 0.85, this is about 47% of the wage in bad jobs, and 36% of the wage in good jobs. This matches the evidence on 1987 post-income tax replacement rates in Table A1 of Gruber (1998), that the replacement rate is slightly less than fifty percent.

[Table 1 about here.]

Although we choose the technology and preference parameters to match a similar number of pieces of evidence, below we also report a number of other statistics that are implied by our model and can be compared with the data. Most of these results appear to be in line with the evidence, even though the calibration was not designed to match these facts. These can be interpreted as a test of the ‘over-identifying’ restrictions implied by our numerical exercise.

4.2 Numerical Methodology

We calculate the optimal policy using backward induction. Begin with a conjecture for the Bellman functions. Then at each step, calculate the optimal policy on a grid of asset levels for workers in each employment state.¹⁵ Note that workers are not constrained to choose asset levels from that grid. Instead, we can calculate new value functions using a

¹⁵We compute the optimal policy at about 400 asset levels, concentrating particularly on the more curved region of the policy function, with low asset levels. More precisely, we calculate the policy at points in the set $\{1.01^n - 1\}$ for $n \in \{0, 1, \dots, N\}$, where N is large enough that in equilibrium no one wants to save beyond that point. This grid is much more dense than is standard in many other calibration exercises. However, we found that with less dense grids, workers incorrectly become too concerned about low asset levels, and thus save too much.

linear interpolation. This process converges to yield the Bellman and policy functions on a balanced growth path.

In order to do this calculation, we must specify the tax rate. However, it is endogenously determined in equilibrium to balance the government budget. Whether this happens depends on the ergodic distribution of workers across asset and employment states, which can in turn be simulated from the policy functions calculated in the first stage. This makes the equilibrium tax rate the solution to a fixed point problem.

Moreover, calculating the government budget surplus is extremely time consuming, since it requires calculating the ergodic distribution of workers across states. Even though the cross-sectional dispersion of asset holdings is relatively small (for example, in our baseline parameterization, no one maintains more than 37 units of assets), asset holdings are very persistent (see Figure 6 and the accompanying discussion). This limits the number of significant figures in our reported results.

4.3 Benchmark Results

The policy functions in the benchmark parameterization are well-behaved. Wealthy workers look for good jobs and poor workers look for bad ones. For insured workers, the critical threshold for assets, above which they seek good jobs, is approximately 4 (times their weekly income in bad jobs). Uninsured workers start looking for bad jobs when their assets fall below 9.

[Figure 3 about here.]

Figure 3 shows that workers in good jobs consume more than those in bad ones, and employed workers consume more than the unemployed, conditional on their assets. However, uninsured workers sometimes consume more than insured workers at intermediate asset levels. This occurs because insured workers choose to enjoy a lot of leisure and low consumption (Figure 4), especially when they seek good jobs where the marginal value of an additional hour of search is relatively low. Similarly, unemployed workers who receive benefits raise their consumption when their assets fall so low that they begin to seek bad jobs.

[Figure 4 about here.]

More generally, Figure 4 shows that hours are a fairly complicated function of the worker's state. While employed, richer workers enjoy more consumption and more leisure. Workers in good jobs typically work longer hours as they take advantage of this temporarily high wage by intertemporal substitution. For unemployed workers, hours depend on the type of job that they seek, with the lowest search effort for unemployed workers near the threshold of

applying for bad jobs. In any case, there is little variability in labor supply, with a maximum range of 0.78–0.92.¹⁶

[Figure 5 about here.]

Workers' consumption and search patterns give rise to an ergodic distribution of asset holdings (Figure 5). This is both the cross-sectional distribution of asset holdings in a large economy, and the time-series distribution of asset holdings for one individual. The distribution is skewed, because at any point in time, about half the workers are in bad jobs, where they maintain very low asset holdings. Because it is easy to regain another bad job following the loss of one, asset levels cannot fall too far. The other half of the workers are in good jobs, where they have substantial precautionary savings. Many of these achieve quite high asset levels. However, upon losing a good job, it takes a long time to find another one, during which time the assets are run down rapidly.

[Table 2 about here.]

The first column in Table 2 summarizes other results from the benchmark model. A balanced budget requires a 2.2% labor income tax rate, and yields a mean asset level 16.5 times the wage in bad jobs. The next four lines show some results that the parameterization was chosen to match. About half the jobs are good, yielding a 6.3% unemployment rate, or conversely an average unemployment duration of a quarter year. The standard deviation of weekly labor income also matches the 0.13 number reported by Krueger and Summers (1988).

The following four lines give the average working and searching hours for workers in different employment states. As indicated above, uninsured workers have the least leisure, while insured workers enjoy the most. Employed workers fall somewhere in between. Again, there is very little variation in hours.

The next two lines report the first two over-identifying restrictions. First we calculate the immediate drop in consumption for an average worker following the loss of her job. We can see from Figure 3 that this varies considerably with asset level and job type. On average, it is 15.2% in the benchmark model. This is somewhat higher than Gruber's (1997) estimate of an eight percent decline in food consumption for high school graduates at this replacement rate. Since food consumption is more inelastic than overall consumption, however, our estimate is plausible.

Second, in order to check the implications of our model for low frequency income variability, we calculate the standard deviation of log annual income. Notice that this is different

¹⁶Since poorer workers are more likely to search harder and to look for bad jobs, the model predicts that the exit rate from unemployment is generally lower for richer workers, which appears counterfactual. In practice, this may be because richer workers have better job opportunities due to unobserved heterogeneity.

from the standard deviation of weekly income for two reasons: It includes workers who are unemployed for part or all of the year, and it aggregates income to an annual level. The second difference is quite important, for if income fluctuations were completely transitory, the standard deviation of annual income would be quite small. Our benchmark model delivers a standard deviation of 0.17, only slightly smaller than Carroll and Samwick’s (1997) estimate of 0.19 for all U.S. workers, or 0.20 for workers with a high school diploma (see their Table 1). This shows that our model generates individual income persistence at annual frequencies in line with the patterns in the data, despite the transitory nature of shocks.

[Figure 6 about here.]

To further understand this point, it helps to look at a typical 10,000 week (192 year) sample path for an individual’s assets (Figure 6). Asset holdings are extremely persistent. There are long periods during which the individual maintains low asset levels and takes bad jobs, and other periods when she is luckier and so is able to save more and work in good jobs. In this particular example, the individual spends the middle 2000 weeks (48 years) primarily in bad jobs. But in the last 4000 weeks, she never takes a bad job.

The last four rows of Table 2 summarize productivity and welfare in this economy. We give four measures: the average wage; total output (hence consumption); total leisure; and the welfare of an average worker at any point in time. Welfare is measured as the level of consumption that must be given to a worker to make her indifferent between consuming this without having to work at all, and participating in the labor market as in the model.

4.4 Changes in Unemployment Insurance Policy

The remaining columns of Table 2 give the results from three policy changes. These are useful both because they enable us to check whether our model generates plausible increases in unemployment duration and wages in response to more generous UI policy, and because they will show whether moderate levels of UI can raise productivity, output, and welfare.

In Column II, we raise the level of unemployment benefits by ten percent. This raises the mean unemployment duration by exactly one week, which agrees precisely with Meyer’s (1989) estimated response, and is somewhat larger than Atkinson and Micklewright’s (1991) conclusion from their literature survey. In Column III, we reduce the benefit exhaustion rate ϕ by ten percent, resulting in a 2.8 week increase in unemployment benefits. This leads to a half week increase in unemployment duration, which is consistent with Katz and Meyer’s (1990) estimate that a one week increase in the duration of benefits will raise unemployment duration by 0.16 to 0.20 weeks. Our model, therefore, performs quite well in these ‘tests of over-identifying restrictions’, which gives us some confidence in evaluating its implications regarding the impact of UI policy on output and welfare.

The table shows that in response to both more generous and longer benefits, the share of good jobs increases significantly (row 5). Consequently wages and consumption rise as well (see the last four rows). Interestingly, the response of average wages to the change in policy is quite reasonable: a ten percent increase in the replacement rate, for example, raises the average wage rate by 1.2%, which is smaller than some of Ehrenberg and Oaxaca's (1976) estimates. Leisure, on the other hand, is virtually unchanged, and as a result, per capita output rises sharply in either policy experiment. Utility also increases, because workers benefit both from the reduction in risk and from the increase in productivity.

Notice also that in both cases, the primary reason for the increase in the unemployment rate and unemployment duration is the change in the types of jobs that workers seek. In fact, there is very little change in search effort. Therefore, our analysis suggests that a large part of what is usually labeled moral hazard, may be driven by the willingness of better-insured workers to seek higher wage jobs.

We point out two other results from these experiments. First, the change in the UI scheme has very little effect on mean asset holdings, justifying our partial equilibrium analysis. In fact, more generous UI actually induces more savings, as more workers prepare to undertake the risky search for good jobs. Second, there is one dimension in which our model does not perform particularly well: Gruber (1997) estimates that an increase in the replacement rate results in a smaller decline in consumption following job loss. In our model, the consumption decline (surprisingly) increases. Again, this occurs because workers prepare for longer unemployment spells.

Finally, the fourth column of Table 2 shows the results from eliminating all UI in our benchmark economy. These results should be interpreted cautiously, since they are based on estimates that are well 'out-of-sample'. Nevertheless, the qualitative behavior of such an economy is instructive. No worker looks for a good job, resulting in very low unemployment and wage dispersion. However, the productivity loss implies a substantial reduction in output and welfare.

4.5 Discussion

How large are the productivity and welfare gains from UI in this economy? Our results indicate that an increase in unemployment benefits that raises UI receipts and expenditures by 17.6% will lead to a 0.7% increase in output and 0.6% increase in welfare. An increase in benefit duration that raises receipts and expenditures by 5.6% will lead to a 0.5% increase in output and 0.4% increase in welfare. According to this metric, the gains are plausibly small.

On another metric, however, the gains appear large. Increasing benefit generosity raises the UI system's expenditures by 0.0035 per worker, but leads to twice as large an increase in output. Increasing benefit duration raises expenditures by 0.0011, and leads to about four

times as large an increase in output. This suggests substantial returns to raising the size of the UI system.

These findings should be interpreted with caution, however. Our results are ‘local’, in the sense that we use U.S. observations to pin down parameters, and we have the most confidence in our results for small deviations from the U.S. benchmark.¹⁷ At least two sets of assumptions are suspect when one considers large deviations from the benchmark. First, the output of good and bad jobs may be imperfect substitutes, so a large increase in the number of good jobs may reduce productivity and wage differentials across sectors (e.g. Acemoglu 1998). Second, our assumption of two job types imposes a linearity on our model economy. Although this is a valid local approximation when we conduct small policy changes, it is problematic when there are large changes in UI. For example, an earlier version of the paper showed that with very generous UI, workers may become ‘overinsured’ and seek out excessively risky jobs.

There are two other important caveats to these results. First, there are distortions associated with the administration of UI, like the subsidy of temporary layoffs and volatile employment. This alters the pattern of production towards industries like construction, and is only partially mitigated by the experience rating of insurance premiums.¹⁸

Finally, the results in Table 2 compare two steady states, and ignore the costs of transition. In particular, in moving to a system with more generous UI, many workers in bad jobs must locate good ones, which is a time-consuming process. To obtain a bound on the transition costs, we set the initial distribution of workers at its ergodic distribution across employment and asset states in the benchmark economy. We then increase the generosity or duration of benefits *and the tax rate* to the values indicated in the second and third columns of Table 2. Utility now increases from 0.651 in the benchmark economy to 0.652 with either high or long UI, rather than 0.654 or 0.655 as in the calculations that ignored the transition costs. Although the welfare gains are smaller, the numbers reported here represent an upper bound on the transition costs. Low unemployment rates in the early years ensure that the government runs a budget surplus, which is never redistributed to the population.

5 Robustness

We consider four robustness checks on our model. First, we raise the relative value of leisure by a factor of six, to $\eta = 1$. This increases the importance of traditional moral hazard, as

¹⁷This is why we do not investigate the implications of increasing UI generosity from the U.S. to European levels.

¹⁸Increasing the experience rating of the UI system is not necessarily an improvement either. A perfectly experience rated UI system effectively taxes firing, similar to the practice in Europe, and as a result, distorts job flows and matching.

insured workers choose to cut back more on their job search effort. Then we cut the wage differential between good and bad jobs in half, to $w_g = 1.15$, reducing the productivity gains from workers looking for good jobs. Third, we lower the coefficient of relative risk aversion to $\theta = 1$, which makes utility linear in log consumption and log leisure. Finally, we extend the model to allow workers employed in bad jobs to look for good jobs. In each case, the model does not perform as well as the benchmark parameterization along some dimensions. Still, the robustness checks support our main conclusion, that moderate levels of UI, similar to those in the U.S., are welfare enhancing.

5.1 The Value of Leisure

[Table 3 about here.]

To investigate the implications of more severe moral hazard problems, we raise the value of leisure substantially, to $\eta = 1$. To keep hours, hence the replacement rate, at the same level as our benchmark, we also increase the available time allotment to $\bar{h} = 1.7$, but do not change the other parameters. Table 3 summarizes the results in this new economy.

The most significant change in results from the benchmark parameterization is that the search effort of insured workers declines sharply compared with their uninsured peers. While previously insured workers searched about six percent less than uninsured workers, they now search 24% less, and consume substantially more leisure.

This traditional moral hazard problem has both positive and normative implications. On the positive side, unemployment duration now responds too much to changes in UI policy. This is primarily due to the responsiveness of search effort to unemployment benefits. In fact, the share of good jobs is relatively constant. Additionally, consumption declines by about fifty percent more than in the benchmark model following the loss of a job. This is due to the substitutability between consumption and leisure, together with the moral hazard-induced increase in leisure. In our view, the magnitudes of the decline in consumption and increase in unemployment are implausible. These were the main motivations for choosing a smaller value of leisure as the benchmark.

The increase in moral hazard implies that, while the average product of a labor hour still increases with the generosity of the UI policy, output actually declines slightly in the neighborhood of the current U.S. system. Welfare, on the other hand, is locally constant, so that even with this extreme degree of moral hazard, the current system does not reduce economic welfare. Additionally, both output and welfare are much higher than they would be in the absence of any insurance.¹⁹

¹⁹Hansen and Imrohoroglu (1992) compute very small benefits of UI when $\eta = 2$. However, because unemployed workers do not search in their model, they do not obtain our conclusion that consumption is too variable when the relative value of leisure is high. Our model indicates why such high values of η are

5.2 Wage Dispersion

[Table 4 about here.]

Reducing the wage gap between good and bad jobs scarcely changes our results. We set $w_g = 1.15$, implying only a 15 percent gap between good and bad jobs. We also reduce the difference between the job finding rates to $m_b = 0.15$ and $m_g = 0.06$, which enables us to maintain a mix of good and bad jobs similar to the baseline. Table 4 shows that the results of this parameterization are very similar to the benchmark, except, of course, that the standard deviations of wages and income decline.

A small increase in unemployment benefits or benefit duration still leads to an increase in output and welfare, despite the lower benefits of obtaining a good job. However, one must interpret this result cautiously, as unemployment duration is not as responsive to policy in this case as in the data. Correcting this by increasing the value of leisure η would reduce the benefits of UI along the lines of the previous robustness check.

5.3 Risk Aversion

[Table 5 about here.]

Our third robustness check is to reduce the coefficient of relative risk aversion all the way to $\theta = 1$. This has two effects. First, consumption and leisure are no longer substitutes. And second, workers are far more willing to endure risk, which reduces savings considerably and increases workers' willingness to seek riskier opportunities. In fact, with this sharp decline in risk aversion, it is difficult to get workers to look for bad jobs. We therefore must also reduce the productivity differential to $w_g = 1.15$ and the gap in job finding rates to $m_b = 0.2$ and $m_g = 0.04$. The results are shown in Table 5.

With this change in parameterization, workers have much less desire to smooth consumption, which shows up in the low asset levels and large decline in consumption following job loss. Additionally, whether workers search for good or bad jobs is extremely sensitive to small changes in UI policy. The usual experiment of a ten percent increase in unemployment benefits causes a three week increase in unemployment duration, which seems implausible. Similarly, a 2.8 week increase in benefit duration causes a one week increase in unemployment duration, far more than the micro estimates. Surprisingly, despite this large increase in unemployment, both output and welfare are slightly higher with the more generous UI systems. Therefore, our main conclusion, that moderate UI increases output and welfare, is not dependent upon the high degree of risk aversion, although many of the other results are less plausible when workers have a coefficient of relative risk aversion much less than 4.

implausible.

5.4 On-the-Job Search

[Table 6 about here.]

Perhaps most surprisingly, including on-the-job search does not affect our main results either. We incorporate this feature in a natural way by allowing workers in bad jobs to search for good ones without first quitting.²⁰ These workers spend h_w hours working, yielding after-tax income $(1 - \tau)h_w w_b$; h_s hours searching for a good job, yielding transition probability $m_g h_s$; and they enjoy $\bar{h} - h_w - h_s$ hours of leisure. We keep all parameters as in the benchmark model. The results are presented in Table 6.

On-the-job search has very little effect on the benchmark equilibrium. Poor workers never search while employed. Instead, they supply a lot of labor, and gradually increase their asset holdings. When assets exceed a threshold value, they switch their activity from all-work to all-search.²¹ Assets immediately decline, so if their search was unsuccessful, they resume working in the bad job the following period. Workers in bad jobs spend about 2% of their time endowment, roughly 1 hour per week, searching. Effectively, on-the-job search allows workers to quit bad jobs and search, with a guarantee of being able to return to their old job should their search fail. Because bad jobs are easy to get, this option is not worth very much, so on-the-job search has little effect on the equilibrium of the model.

On-the-job search has a similarly small effect on the ‘High UI’ and ‘Long UI’ parameterizations, and does not affect our main results. However, it has a fairly large effect on the parameterization without any UI. The wealthiest workers in bad jobs look for good ones. When they find them, they save a lot, enough to be willing to look for a good job when they lose their job. As a result, over a third of jobs are good in the resulting equilibrium, and the cost of not having any UI is substantially smaller than in a world without on-the-job search. The behavior of the economy without UI seems more plausible here than in the benchmark model.

6 Concluding Comments

Conventional wisdom views unemployment insurance as a serious distortion that we have to live with in order to smooth income risk and consumption variability. In this paper, we have argued that moderate unemployment insurance may actually improve the allocation of

²⁰There is no reason in this model for workers in bad jobs to look for bad jobs, or workers in good jobs to search at all. Undoubtedly, in reality ‘sideways’ mobility represents an important component of job-to-job flows.

²¹This switching behavior may seem extreme, but would be more reasonable with a shorter period length. For example, a worker takes an afternoon off to look for a better job. Another way to obtain more reasonable conclusions, would be to introduce exogenous restrictions on minimum working hours, which would make the model behave more like the benchmark.

resources. Unemployment insurance enables workers to pursue riskier options, including jobs that are harder to get, but possibly also more productive. As a result, moderate UI may raise output by improving the composition of jobs. In practice, however, this effect of UI may be outweighed by the traditional moral hazard cost: insured workers search less hard for jobs, and therefore spend more time unemployed. Moreover, one might conjecture that the importance of uninsured risks is limited by workers' ability to self-insure by building up buffer stocks of assets. To investigate the quantitative importance of these opposing forces, we constructed a dynamic model in which workers make search effort, savings, and job application decisions. Although the decisions in question are complex, our model is sparse, enabling us to calibrate the few parameters to U.S. data, in particular to the labor market for high school graduates. The model performs well not only along the dimensions in which it is calibrated, but in a number of other dimensions as well, including the standard deviation of annual income, the decline in consumption following job loss, and the responsiveness of unemployment duration and wages to the magnitude and duration of unemployment benefits.

Armed with the confidence that this model captures some of the tradeoffs faced by workers in real labor markets, we investigated the implications of different UI policies on unemployment, output, labor productivity, and welfare. We find that reducing UI from its current U.S. level would reduce both risk sharing and total consumption. Conversely, moderate increases in UI raise output and improve risk sharing.

As this is a calibration exercise, our results are necessarily sensitive to parameter choices. To address this issue, we verified the robustness of our results to several key parameter changes. A complementary strategy would be to look for direct evidence that more generous unemployment insurance programs encourage the creation of more specialized and higher productivity jobs. At a more general level, we believe that more work needs to be done to understand the role and optimal design of UI programs. Labor market reform is a key issue in Europe. Social insurance programs, especially UI programs, are likely to be modified during the next decade. Similar reforms are underway in the U.S.. If we are correct that social insurance programs have a beneficial effect not only on welfare but also on output and productivity, the reforms may have very different consequences than currently envisioned.

Finally, we have assumed that UI is provided by the government. Why unemployment insurance is almost always publicly provided, in contrast to most other insurance contracts, remains an important, unresolved question.²² The answer will likely be relevant to the optimal design of unemployment insurance programs.

²²For recent work on this topic, see Chiu and Karni (1998).

References

- ACEMOGLU, DARON (1998): “Good Jobs versus Bad Jobs,” M.I.T. Working Paper.
- ACEMOGLU, DARON, AND ROBERT SHIMER (1999a): “Efficient Unemployment Insurance,” *Journal of Political Economy*, 107, 893–928.
- (1999b): “Wage and Technology Dispersion,” Mimeo.
- AIYAGARI, RAO (1994): “Uninsured Idiosyncratic Risk And Aggregate Saving,” *Quarterly Journal of Economics*, 109, 311–331.
- ANDERSON, PATRICIA, AND BRUCE MEYER (1994): “The Extent and Consequences of Job Turnover,” *Brookings Papers on Economic Activity*, Microeconomics, 137–236.
- ATKESON, ANDREW, AND ROBERT LUCAS (1995): “Efficiency and Equality in a Simple Model of Efficiency Unemployment Insurance,” *Journal of Economic Theory*, 66, 64–88.
- ATKINSON, ANTHONY, AND JOHN MICKLEWRIGHT (1991): “Unemployment and Compensation and Labor Market Transitions: A Critical Review,” *Journal of Economic Literature*, 58, 277–297.
- BURDETT, KENNETH, SHOUYONG SHI, AND RANDALL WRIGHT (1997): “Pricing with Frictions,” Mimeo.
- CARROLL, CHRISTOPHER, AND ANDREW SAMWICK (1997): “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, 40, 41–72.
- CHIU, W. HENRY, AND EDI KARNI (1998): “Endogenous Adverse Selection and Unemployment Insurance,” *Journal of Political Economy*, 106, 806–827.
- COSTAIN, JAMES (1996): “Unemployment Insurance in a General Equilibrium Model of Job Search and Precautionary Savings,” Mimeo.
- DEATON, ANGUS (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–48.
- DIAMOND, PETER (1981): “Mobility Costs, Frictional Unemployment and Efficiency,” *Journal of Political Economy*, 89, 798–812.
- EHRENBERG, RONALD, AND RONALD OAXACA (1976): “Unemployment Insurance, Duration of Unemployment, and Subsequent Wage Gain,” *American Economic Review*, 66, 754–766.

- GIBBONS, ROBERT, AND LAWRENCE KATZ (1992): “Does Unmeasured Ability Explain Inter-industry Wage Differentials,” *Review of Economic Studies*, 59, 515–535.
- GOMES, JOAO, JEREMY GREENWOOD, AND SERGIO REBELO (1997): “Equilibrium Unemployment,” Mimeo.
- GOURINCHAS, PIERRE OLIVIER, AND JONATHAN PARKER (1997): “Consumption over the Lifecycle,” Mimeo.
- GRUBER, JONATHAN (1997): “The Consumption Smoothing Benefits of Unemployment Insurance,” *American Economic Review*, 87, 192–205.
- (1998): “Unemployment Insurance, Consumption Smoothing and Private Insurance: Evidence from the PSID and CEX,” *Research in Employment Policy*, 1, 3–32.
- HANSEN, GARY (1985): “Indivisible Labor and the Business Cycle,” *Journal of Monetary Economics*, 16, 309–327.
- HANSEN, GARY, AND AYSE IMROHORGLU (1992): “The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard,” *Journal of Political Economy*, 100, 118–142.
- HOPENHAYN, HUGO, AND JUAN PABLO NICOLINI (1997): “Optimal Unemployment Insurance,” *Journal of Political Economy*, 105, 412–438.
- HUGGETT, MARK (1993): “The Risk-free Rate in Heterogeneous Agent Incomplete Insurance Economies,” *Journal Economic Dynamics and Control*, 17, 953–969.
- IMF (1995): “World Economic Outlook,” Washington, D.C., May.
- KATZ, LAWRENCE, AND BRUCE MEYER (1990): “The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment,” *Journal of Public Economics*, 41, 45–72.
- KRUEGER, ALAN, AND LAWRENCE SUMMERS (1988): “Efficiency Wages and the Inter-Industry Wage Structure,” *Econometrica*, 56, 259–293.
- KRUSELL, PER, AND ANTHONY SMITH (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106, 867–897.
- LUCAS, ROBERT, AND EDWARD PRESCOTT (1974): “Equilibrium Search and Unemployment,” *Journal of Economic Theory*, 7, 188–209.

- MARIMON, RAMON, AND FABRIZIO ZILIBOTTI (1999): “Unemployment versus Mismatch of Talent: Reconsidering Unemployment Benefits,” *Economic Journal*, 109, 266–291.
- MEYER, BRUCE (1989): “A Quasi-Experimental Approach to the Effects of Unemployment Insurance,” NBER Working Paper 3159.
- MONTGOMERY, JAMES (1991): “Equilibrium Wage Dispersion and Interindustry Wage Differentials,” *Quarterly Journal of Economics*, 106, 163–179.
- OECD (1997): “OECD Employment Outlook,” Paris, France, December.
- PETERS, MICHAEL (1991): “Ex Ante Price Offers in Matching Games: Non-Steady States,” *Econometrica*, 59, 1425–1454.
- POTERBA, JAMES, AND LAWRENCE SUMMERS (1986): “Reporting Errors and Labor Market Dynamics,” *Econometrica*, 54, 1319–38.
- SHAVELL, STEVEN, AND LAURENCE WEISS (1979): “The Optimal Payment of Unemployment Benefits Over Time,” *Journal of Political Economy*, 87, 1347–62.
- VALDIVIA, VICTOR (1996): “Policy Evaluation in Heterogeneous Agent Economies: The Welfare Impact of Unemployment Insurance,” Ph.D. thesis, Northwestern University.

Variable	Parameterization
<u>Gross interest rate</u>	$R = 1.0006$
<u>Technology Parameters</u>	
Gross growth rate	$\Gamma = 1.0003$
Bad job wage	$w_b = 1$
Bad job layoff rate	$\delta_b = 0.005$
Bad job matching rate	$m_b = 0.5$
Good job wage	$w_g = 1.3$
Good job layoff rate	$\delta_g = 0.005$
Good job matching rate	$m_g = 0.05$
<u>Preference Parameters</u>	
Discount factor	$\beta = 0.998$
Relative risk aversion	$\theta = 4$
Maximum hours	$\bar{h} = 1$
Relative value of leisure	$\eta = 0.167$
<u>UI policy</u>	
Unemployment benefits	$b = 0.4$
Benefit exhaustion rate	$\phi = 0.04$

Table 1: Baseline parameterization.

Variable	Column I Benchmark	Column II High UI	Column III Long UI	Column IV No UI
Unemployment Benefit b	0.4	0.44	0.4	0
Benefit Duration $1/\phi$	25	25	27.8	—
Tax Rate τ	0.022	0.025	0.023	0
Mean Assets	16.5	16.8	17.0	7.1
Fraction of Good Jobs	51.5%	56.1%	54.2%	0%
Unemployment Rate	6.3%	6.7%	6.5%	1.1%
Unemployment Duration	13.5	14.5	14.0	2.2
Std. Dev. of Log Weekly Income	0.133	0.132	0.132	0.002
Work Hours in Good Jobs	0.862	0.861	0.862	—
Work Hours in Bad Jobs	0.859	0.859	0.859	0.858
Search Hours by Insured	0.816	0.809	0.814	—
Search Hours by Uninsured	0.866	0.866	0.866	0.895
Consump. Drop Following Job Loss	15.2%	15.4%	15.4%	10.2%
Std. Dev. of Log Annual Income	0.169	0.168	0.171	0.020
Average Wage	1.155	1.169	1.163	1.000
Output	0.931	0.938	0.935	0.849
Leisure	0.142	0.142	0.142	0.142
Welfare: Consumption Metric	0.651	0.655	0.654	0.611

Table 2: Results in the benchmark parameterization, and for three alternative UI schemes.

Variable	Column I Benchmark	Column II High UI	Column III Long UI	Column IV No UI
Unemployment Benefit b	0.4	0.44	0.4	0
Benefit Duration $1/\phi$	25	25	27.8	—
Tax Rate τ	0.025	0.031	0.027	0
Mean Assets	16.2	15.5	16.1	7.1
Fraction of Good Jobs	53.0%	55.4%	53.8%	0%
Unemployment Rate	7.3%	8.0%	7.7%	1.0%
Unemployment Duration	15.8	17.5	16.6	2.0
Std. Dev. of Log Weekly Income	0.138	0.138	0.138	0.008
Work Hours in Good Jobs	0.866	0.865	0.866	—
Work Hours in Bad Jobs	0.855	0.854	0.855	0.853
Search Hours by Insured	0.672	0.624	0.663	—
Search Hours by Uninsured	0.886	0.893	0.888	0.971
Consump. Drop Following Job Loss	22.4%	24.2%	24.0%	4.2%
Std. Dev. of Log Annual Income	0.181	0.183	0.183	0.014
Average Wage	1.160	1.168	1.162	1
Output	0.926	0.923	0.924	0.844
Leisure	0.150	0.155	0.151	0.146
Welfare: Consumption Metric	0.446	0.446	0.446	0.419

Table 3: Results with a high value of leisure, $\eta = 1$.

Variable	Column I Benchmark	Column II High UI	Column III Long UI	Column IV No UI
Unemployment Benefit b	0.4	0.44	0.4	0
Benefit Duration $1/\phi$	25	25	27.8	—
Tax Rate τ	0.023	0.026	0.024	0
Mean Assets	27.1	26.9	27.1	27.1
Fraction of Good Jobs	46.5%	49.2%	48.8%	0%
Unemployment Rate	6.3%	6.5%	6.4%	3.7%
Unemployment Duration	13.5	14.0	13.7	7.6
Std. Dev. of Log Weekly Income	0.070	0.070	0.070	0.005
Work Hours in Good Jobs	0.860	0.860	0.860	—
Work Hours in Bad Jobs	0.859	0.859	0.859	0.860
Search Hours by Insured	0.822	0.816	0.820	—
Search Hours by Uninsured	0.875	0.878	0.875	0.873
Consump. Drop Following Job Loss	15.0%	15.3%	15.1%	12.2%
Std. Dev. of Log Annual Income	0.136	0.136	0.136	0.133
Average Wage	1.070	1.074	1.073	1
Output	0.862	0.863	0.863	0.828
Leisure	0.142	0.143	0.142	0.140
Welfare: Consumption Metric	0.608	0.609	0.610	0.588

Table 4: Results with low wage dispersion, $w_g = 1.15$.

Variable	Column I Benchmark	Column II High UI	Column III Long UI	Column IV No UI
Unemployment Benefit b	0.4	0.44	0.4	0
Benefit Duration $1/\phi$	25	25	27.8	—
Tax Rate τ	0.025	0.033	0.028	0
Mean Assets	5.8	6.0	5.9	5.8
Fraction of Good Jobs	33.8%	46.7%	38.3%	0%
Unemployment Rate	6.4%	7.7%	6.9%	2.6%
Unemployment Duration	13.7	16.8	14.7	5.4
Std. Dev. of Log Weekly Income	0.066	0.070	0.068	0.004
Work Hours in Good Jobs	0.859	0.859	0.859	—
Work Hours in Bad Jobs	0.859	0.859	0.859	0.859
Search Hours by Insured	0.828	0.821	0.827	—
Search Hours by Uninsured	0.921	0.921	0.922	0.910
Consump. Drop Following Job Loss	27.7%	27.9%	27.9%	28.6%
Std. Dev. of Log Annual Income	0.114	0.134	0.121	0.068
Average Wage	1.051	1.070	1.057	1
Output	0.845	0.848	0.846	0.836
Leisure	0.142	0.143	0.142	0.140
Welfare: Consumption Metric	0.603	0.604	0.604	0.598

Table 5: Results with low risk aversion, $\theta = 1$.

Variable	Column I Benchmark	Column II High UI	Column III Long UI	Column IV No UI
Unemployment Benefit b	0.4	0.44	0.4	0
Benefit Duration $1/\phi$	25	25	27.8	—
Tax Rate τ	0.021	0.025	0.022	0
Mean Assets	17.8	17.5	18.0	15.0
Fraction of Good Jobs	61.1%	63.1%	62.9%	37.4%
Unemployment Rate	6.4%	6.7%	6.6%	3.2%
Unemployment Duration	13.7	14.5	14.2	6.6
Std. Dev. of Log Weekly Income	0.212	0.208	0.168	0.199
Work Hours in Good Jobs	0.862	0.861	0.862	0.863
Work Hours in Bad Jobs	0.839	0.840	0.839	0.835
Search Hours in Bad Jobs	0.020	0.021	0.022	0.027
Search Hours by Insured	0.818	0.810	0.816	—
Search Hours by Uninsured	0.865	0.865	0.865	0.864
Consump. Drop Following Job Loss	15.6%	15.9%	16.7%	13.0%
Std. Dev. of Log Annual Income	0.172	0.173	0.173	0.176
Average Wage	1.185	1.191	1.190	1.115
Output	0.946	0.948	0.949	0.912
Leisure	0.142	0.141	0.138	0.862
Welfare: Consumption Metric	0.659	0.660	0.660	0.637

Table 6: Results with on-the-job search.




Figure 1: The thin curve is the worker's indifference curve. The thick curve is the firm's zero profit constraint.




Figure 2: An increase in unemployment income flattens workers' indifference curves, from the dashed line to the solid line, shifting the tangency point towards higher wages and specificity.




Figure 3: Consumption as a function of normalized assets in the four employment states, benchmark parameterization. Note that workers quit bad jobs when their assets exceed 19.



Figure 4: Hours of search or labor supply as a function of normalized assets in the four employment states, benchmark parameterization. Note that workers quit bad jobs when their assets exceed 19.




Figure 5: Ergodic distribution of asset holdings, benchmark parameterization.




Figure 6: Ergodic distribution of asset holdings, benchmark parameterization.