Moral Hazard vs. Liquidity in Unemployment Insurance

Raj Chetty*

UC-Berkeley and NBER

July 2007

Abstract

It is well known that higher unemployment benefits lead to longer unemployment durations. This result has been interpreted as evidence of "moral hazard" – a behavioral response to distorted marginal incentives to search. This paper shows that unemployment benefits also raise durations through a "liquidity" effect for households who cannot smooth consumption perfectly. The empirical importance of the liquidity effect is documented in two ways. First, state-level increases in unemployment benefits have larger effects on durations for households that are likely to be liquidity constrained (e.g., those with low assets). Second, lump-sum severance payments significantly increase durations among constrained households. Together, the empirical estimates imply that 60% of the effect of unemployment benefits on durations is due to the liquidity effect. To evaluate the welfare implications of this finding, I derive a new formula for the optimal level of unemployment benefits in terms of the ratio of the liquidity effect relative to moral hazard. Implementing this formula using the empirical estimates implies that the optimal unemployment benefit level exceeds 50% of the previous wage.

Keywords: unemployment durations, optimal benefits, wealth effects, borrowing constraints

*E-mail: chetty@econ.berkeley.edu. I have benefited from discussions with Alan Auerbach, Olivier Blanchard, David Card, Martin Feldstein, Amy Finkelstein, Jon Gruber, Jerry Hausman, Caroline Hoxby, Juan Jimeno, Larry Katz, Bruce Meyer, Emmanuel Saez, Jesse Shapiro, Robert Shimer, Adam Szeidl, Ivan Werning, anonymous referees, and numerous seminar participants. Philippe Bouzaglou, Greg Bruich, David Lee, Ity Shurtz, Jim Sly, and Philippe Wingender provided excellent research assistance. I am very grateful to Julie Cullen and Jon Gruber for sharing their unemployment benefit calculator, and to Suzanne Simonetta and Loryn Lancaster at the Dept. of Labor for assistance with state UI laws. Funding from the National Science Foundation and NBER is gratefully acknowledged. The data and code used for the empirical analysis are available on the author's website. An earlier version of this paper with additional results is available as NBER working paper 11760.

1 Introduction

One of the best known empirical results in public finance is that social insurance programs such as unemployment insurance (UI) reduce labor supply. For example, Moffitt (1985), Meyer (1990), and others have shown that a 10% increase in unemployment benefits raises average unemployment durations by 4-8% in the U.S. This finding has traditionally been interpreted as evidence of moral hazard caused by a *substitution* effect: UI distorts the relative price of leisure and consumption, reducing the marginal incentive to work. For instance, Krueger and Meyer (2002) remark that behavioral responses to UI and other social insurance programs are large because they "lead to short-run variation in wages with mostly a substitution effect." Similarly, Gruber (2005) notes that "UI has a significant moral hazard cost in terms of subsidizing unproductive leisure."

This paper questions whether the link between unemployment benefits and durations is purely due to moral hazard. The analysis is motivated by evidence that many unemployed individuals have limited liquidity and exhibit excess sensitivity of consumption to cash-on-hand (Gruber 1997, Browning and Crossley 2001, Bloemen and Stancanelli 2005). Indeed, nearly half of job losers in the United States report zero liquid wealth at the time of job loss, suggesting that many households may be unable to smooth transitory income shocks relative to permanent income.

Using a simple job search model, I show that when an individual cannot smooth consumption perfectly, unemployment benefits affect search intensity through a "liquidity" effect in addition to the moral hazard channel emphasized in earlier work. Intuitively, UI benefits increase cash-on-hand and consumption while unemployed for an agent who cannot smooth perfectly. Such an agent faces less pressure to find a new job quickly, leading to a longer duration. Hence, unemployment benefits raise durations purely through moral hazard for unconstrained individuals, but through both liquidity and moral hazard effects for constrained individuals.

The distinction between liquidity and moral hazard is of interest because the two effects lead to divergent views about optimal policy. The substitution effect is a socially suboptimal response to the creation of a wedge between private and social marginal costs. In contrast, the liquidity effect is a response to the correction of a market failure (incomplete credit and risk-sharing markets). Building on this logic, I develop a new test for the optimal unemployment benefit level – similar in spirit to the recent tests proposed by Chetty (2006a) and Shimer and Werning (2007) – based on the ratio of the liquidity and moral hazard effects. The formula uses revealed preference to calculate the value of insurance: if an agent chooses a longer duration primarily because he has more cash-on-hand (as opposed to distorted incentives), we infer that UI benefits bring the agent closer to the social optimum. Relative to existing approaches to optimal benefits, the formula proposed here has two advantages: (1) it requires data only on unemployment durations and (2) it does not rely on state-independence or a specific parametric form of utility.

I estimate the importance of moral hazard vs. liquidity in UI empirically using two complementary strategies. I first estimate the effect of UI benefits on durations separately for liquidity constrained and unconstrained households, exploiting differential changes in UI benefit levels across states in the U.S. for identification. Since households' ability to smooth consumption is unobserved, I proxy for it using three measures: asset holdings, single vs. dual-earner status, and an indicator for having to make a mortgage payment. I find that a 10% increase in UI benefits raises unemployment durations by 7-10% in the constrained groups. In contrast, changes in UI benefits have much smaller effects on durations in the unconstrained groups, indicating that the moral hazard effect is relatively small among these groups. These results suggest that liquidity effects could be quite important in the benefits-duration link. However, they do not directly establish that benefits raise the durations of constrained agents by increasing liquidity unless one assumes that the substitution elasticities are similar across constrained and unconstrained groups.

This limitation leads to the second empirical strategy, in which I estimate the fraction of the total benefit elasticity accounted for by the liquidity effect. I exploit variation across job losers in the receipt of lump-sum severance payments, which have a liquidity effect but no moral hazard effect. Using a survey of job losers from Mathematica that has information on severance pay and unemployment durations, I find that individuals who received severance pay (worth about $4000 on average) have substantially longer durations. An obvious concern is that this finding may reflect correlation rather than causality because severance pay is not randomly assigned. Two pieces of evidence support the causality of severance pay. First, the estimated effect of severance pay is not affected by controls for demographics, income, job tenure, industry, and occupation in a Cox hazard model. Second, severance payments have a large effect on durations among constrained (low asset) households, but have no effect on durations among unconstrained households. These findings, though not conclusive given the lack of randomized variation in cash grants, suggest that UI has a substantial liquidity effect.

Combining the point estimates from the two empirical approaches, I find that roughly 60% of the effect of benefits on durations is due to the liquidity effect. Using this estimate, I implement the test for optimal benefits derived from the search model, and find that the marginal welfare gain of raising the unemployment benefit level is small but positive. This implies that the optimal wage replacement rate for UI exceeds current benefit levels, which are approximately 50% of pre-unemployment wages. Although this calculation ignores important factors such as temporary layoffs and general equilibrium effects, it at least points toward a substantial role for policies that enhance households ability to smooth income shocks.

In addition to the empirical literature on unemployment insurance, this paper relates to and builds on several other strands of the literature in macroeconomics and public finance. First, several studies have used consumption data to investigate the importance of liquidity constraints (see e.g., Zeldes 1989; Johnson, Parker, and Souleles 2006). This paper presents analogous evidence from the labor market, showing that labor supply is "excessively sensitive" to transitory income because of limited liquidity. Second, several studies have explored the effects of incomplete insurance and credit markets for job search behavior and UI using simulations of calibrated search models (Hansen and Imrohoglu 1992; Acemoglu and Shimer 2000). The analysis here can be viewed as the empirical counterpart of such studies, in which the extent to which agents can smooth shocks is estimated empirically rather than simulated from a calibrated model. The estimates could be used to better calibrate dynamic models of household behavior in subsequent work.

Finally, the conceptual distinction between moral hazard and liquidity effects applies to other private and social insurance programs. The "revealed preference" approach to calculating the value of insurance proposed here can be adapted to other contexts. For example, one can calculate the value of a health insurance program by estimating the extent to which an agent's medical expenditures would differ if he were paid a lump-sum cash benefit rather than an indemnity benefit that covers health expenses. A key feature of this approach is that it does not require data on the *outcomes* of insurance provision, such as consumption, job match quality, or health. Though this approach relies on the assumption that agents' choices reveal their true preferences, it provides a parsimonious method of identifying policies that are optimal from a libertarian perspective.

The remainder of the paper proceeds as follows. The search model and test for optimal benefits are presented in the next section. Section 3 discusses the evidence on heterogeneous effects of unemployment benefits on durations. Section 4 examines the effect of severance payments on durations. The test for optimal benefits is implemented in section 5. Section 6 concludes.

2 Theory

I analyze a simple job search model that is closely related to the models studied by Lentz and Tranaes (2005) and Card, Chetty, and Weber (2007). I use this framework to distinguish the moral hazard and liquidity effects of UI and derive a test for the optimal benefit level in terms of these effects. I begin by characterizing the agent's job search behavior, and then turn to the government's problem of choosing the optimal benefit level.

2.1 A Job Search Model

Consider a discrete-time setting where the agent lives for T periods $\{0, ..., T-1\}$. Let δ denote the agent's time discount rate and r denote the interest rate. Suppose the agent becomes unemployed at $t = 0$. An agent who enters a period t without a job first chooses search intensity s_t. Normalize s_t to equal the probability of finding a job in the current period. Let $\psi(s_t)$ denote the cost of search effort, which is strictly increasing and convex. If search is successful, the agent begins working immediately in period t itself at a fixed pre-tax wage w_t.[1] Assume that all jobs last indefinitely once found, and that the wage schedule $\{w_t\}$ is deterministic, eliminating reservation-wage choices.

If the worker is unemployed in period t, he receives an unemployment benefit $b_t < w_t$. If the worker is employed in period t, he pays a tax τ that is used to finance the unemployment benefit. Let c_t^e denote the agent's consumption in period t if a job is found in that period. If the agent fails to find a job in period t, he sets consumption to c_t^u. The agent then enters period $t+1$ unemployed and the problem repeats.

Let $v(c_t)$ denote flow utility if employed in period t and $u(c_t) - \psi(s_t)$ denote flow utility if unemployed. Assume that u and v are strictly concave. Note that this state-dependent utility specification permits arbitrary complementarities between consumption and labor.

Optimal Search. The value function for an individual who finds a job at the beginning of period t, conditional on beginning the period with assets A_t is

$$V_t(A_t) = \max_{A_{t+1} \geq L} v(A_t - A_{t+1}/(1+r) + w_t - \tau) + \frac{1}{1+\delta} V_{t+1}(A_{t+1}), \tag{1}$$

where L is a lower bound on assets that may or may not be binding. The value function for an

[1] A more conventional timing assumption in search models without savings is that search in period t leads to a job that begins in period $t+1$. Assuming that search in period t leads to a job in period t itself simplifies the analytic expressions for $\frac{\partial s_t}{\partial A_t}$, as in Lentz and Tranaes (2005).

individual who fails to find a job at the beginning of period t and remains unemployed is:

$$U_t(A_t) = \max_{A_{t+1} \geq L} u(A_t - A_{t+1}/(1+r) + b_t) + \frac{1}{1+\delta} J_{t+1}(A_{t+1}) \tag{2}$$

where

$$J_t(A_t) = \max_{s_t} s_t V_t(A_t) + (1 - s_t) U_t(A_t) - \psi(s_t) \tag{3}$$

is the value of entering period t without a job with assets A_t. It is easy to show that V_t is concave because there is no uncertainty following re-employment; however, the function U_t could be convex. Lentz and Tranaes (2005) show that this problem can be addressed by introducing a wealth lottery prior to the choice of s_t, although they note that in simulations of the model with plausible parameters, non-concavity never arises. I simply assume here that U_t is concave in the parameter space of interest.

An unemployed agent chooses s_t to maximize expected utility at the beginning of period t, given by (3). Optimal search intensity is determined by the first-order condition

$$\psi'(s_t^*) = V_t(A_t) - U_t(A_t). \tag{4}$$

Intuitively, s_t is chosen to equate the marginal cost of search effort with its marginal value, which is given by the difference between the optimized values of employment and unemployment.

2.2 Moral Hazard and Liquidity Effects

To understand how changes in the unemployment benefit level affect search intensity, observe that

$$\partial s_t^* / \partial b_t = -u'(c_t^u) / \psi''(s_t^*)$$

Next, consider the effects of a \$1 increase in assets A_t and a \$1 increase in the period t wage w_t:

$$\partial s_t^* / \partial A_t = \{v'(c_t^e) - u'(c_t^u)\} / \psi''(s_t^*) \leq 0 \tag{5}$$

$$\partial s_t^* / \partial w_t = v'(c_t^e) / \psi''(s_t^*) > 0 \tag{6}$$

The effect of a cash grant on search intensity depends on the difference in marginal utilities between employed and unemployed states. If agents can smooth perfectly, they set $v'(c_t^e) = u'(c_t^u)$, and $\partial s_t^* / \partial A_t = 0$ because a cash grant raises $V_t(A_t)$ and $U_t(A_t)$ by the same amount. In contrast, if individuals face borrowing constraints, have incomplete insurance contracts, or voluntarily reduce c_t^u to maintain a buffer stock of savings, $v'(c_t^e) - u'(c_t^u) < 0$ and $\partial s_t^*/\partial A_t$ can be large.[2] The effect of an increase in w_t is proportional to $v'(c_t^e)$ because a higher wage increases the marginal return to search to the extent that it raises the value of being employed. Putting together (5) and (6) yields the decomposition

$$\partial s_t^*/\partial b_t = \partial s_t^*/\partial A_t - \partial s_t^*/\partial w_t. \tag{7}$$

Equation (7) shows that an increase in the unemployment benefit level lowers search intensity through two separate channels. The first channel is the *liquidity* effect ($\partial s_t^*/\partial A_t$): a higher benefit increases the agent's cash-on-hand, allowing the agent to maintain a higher level of consumption while unemployed and reducing the pressure to find a new job quickly.[3] The second channel is the *moral hazard* effect ($-\partial w_t^*/\partial A_t$): a higher benefit effectively lowers the agent's net wage ($w_t - \tau - b$), reducing the incentive to search though a substitution effect.[4]

As discussed in the introduction, the prevailing view has been that individuals take longer to find a job when receiving higher UI benefits solely because they perceive a lower private return to work. Equation (7) shows that this interpretation is valid only if agents are able to smooth perfectly. The fraction of the UI-duration link due to moral hazard is thus an open empirical question.

2.3 A New Test for the Optimal Benefit Level

The magnitudes of the moral hazard and liquidity effects are of interest because of their normative implications. In particular, the welfare gain of raising the UI benefit level can be written in terms of these effects. To simplify the exposition, I begin by establishing this result for the case where $T = 1$, a static search model.

Static Case. When $T = 1$, the social planner's problem is to maximize the agent's expected

[2] An optimizing agent will never have $v'(c_t^e) - u'(c_t^u) > 0$ because $b_t < w_t$ and re-employment is an absorbing state.

[3] The liquidity effect is the effect of a wealth grant *while unemployed*. The liquidity effect does not equal the conventional wealth effect (i.e., an increase permanent income) if the agent cannot smooth consumption perfectly. Indeed, there are models where the wealth effect is zero but the liquidity effect is positive because of liquidity constraints (see e.g. Shimer and Werning 2007).

[4] This decomposition parallels the Slutsky decomposition of an uncompensated elasticity into a price effect (the "moral hazard" effect here) and the income effect (the "liquidity" effect). I show below that the moral hazard effect reduces the welfare gain from UI, much as a distortionary tax creates a deadweight burden proportional to the price effect.

utility in period 0 subject to the balanced-budget constraint for the UI system:

$$\max_{b_0} \widetilde{W}(b_0) = \{(1 - s_0^*(b_0))u(A_0 + b_0) + s_0^*(b_0)v(A_0 + w_0 - \tau) - \psi(s_0^*(b_0))$$
$$\text{s.t. } b_0(1 - s_0^*(b_0)) = s_0^*(b_0)\tau$$

The welfare gain from increasing b_0 by \$1 is

$$\frac{\partial \widetilde{W}}{\partial b_0} = (1 - s_0^*)u'(c_0^u) - s_0^* v'(c_0^e)\frac{d\tau}{db_0}$$

Noting that $\frac{d\tau}{db_0} = \frac{1-s_0^*}{s_0^*} - \frac{1}{(s_0^*)^2} b_0 \frac{\partial s_0^*}{\partial b_0}$, it follows that

$$\frac{\partial \widetilde{W}}{\partial b_0} = (1 - s_0^*)[u'(c_0^u) - v'(c_0^e)] - \frac{\partial s_0^*}{\partial b_0}\frac{b_0}{s_0^*}v'(c_0^e).$$

To obtain a money metric for the welfare gain, define $\frac{\partial W}{\partial b_0}$ as the ratio of the welfare gain from raising benefits to the welfare gain of increasing the wage rate by \$1:

$$\frac{\partial W}{\partial b_0} = \frac{\partial \widetilde{W}}{\partial b_0}/s_0^* v'(c_0^e) = \frac{1 - s_0^*}{s_0^*}\{\frac{u'(c_0^u) - v'(c_0^e)}{v'(c_0^e)} - \frac{\varepsilon_{1-s,b}}{s_0^*}\}$$

where $\varepsilon_{1-s,b} = \frac{b_0}{s_0^*}\frac{\partial s_0^*}{\partial b_0}$ is the elasticity of the probability of being unemployed with respect to the benefit level. Using equations (5) and (6), this expression can be written as

$$\frac{\partial W}{\partial b_0} = \frac{1 - s_0^*}{s_0^*}\{\frac{-\partial s_0^*/\partial A_0}{\partial s_0^*/\partial w_0} - \frac{\varepsilon_{1-s,b}}{s_0^*}\} \tag{8}$$

This formula shows that the marginal welfare gain from provision of UI benefits can be calculated by comparing the liquidity and moral hazard effects of UI. I provide intuition for this result after showing that it holds in the general case under some approximations.

General Case. When $T > 1$, one can choose a different benefit level b_t in each period. I restrict attention to the optimal policy among the set of policies that have constant benefits: $b_t = b$ $\forall t$. I also assume for simplicity that the wage rate is constant over time: $w_t = w$ $\forall t$. The social planner's problem is to choose the UI benefit level that maximizes the agent's expected utility while balancing the budget:

$$\max J_0(b, \tau) \text{ s.t. } Db = (T - D)\tau \tag{9}$$

where $D = \sum_{t=0}^{T-1} \prod_{j=0}^{t} (1 - s_j)$ is the agent's expected unemployment duration.[5] As above, one can define a money-metric for the welfare gain from UI by comparing the effect of a \$1 increase in b with a \$1 increase in the wage rate w on the agent's expected utility: $\frac{\partial W}{\partial b} \equiv \frac{\partial J_0}{\partial b} / \frac{\partial J_0}{\partial w}$. However, there are two complications in deriving an expression for $\frac{\partial W}{\partial b}$ analogous to (8) when $T > 1$.

First, the benefit level b affects income in subsequent periods, and not just the first period. As a result, the effect of a cash grant in the first period cannot be directly compared with the effect of a benefit increase to identify the liquidity effect, since the timing of receipt of income may matter when agents cannot smooth consumption. To address this issue, suppose the agent receives an annuity payment of \$$a$ in every period $t = 0, ..., T-1$. With the introduction of this annuity, it is straightforward to obtain an analog of the decomposition in (7), as shown in Appendix A:

$$\frac{\partial s_0^*}{\partial b} = \frac{\partial s_0^*}{\partial a} - \frac{\partial s_0^*}{\partial w}. \tag{10}$$

Equation (10) shows that for a permanent increase in the benefit level, the liquidity effect corresponds to the effect of an increase in annuity income in all periods, while the moral hazard effect is given by the effect of a wage increase in all periods.

The second complication when $T > 1$ is that the wage and asset effects needed to compute the effect of the tax increase on welfare depend on the entire path of marginal utilities after re-employment. In particular, let $\mathrm{E}v'(c_{t,j}^e)$ denote the unconditional average marginal utility of consumption while employed over the agent's lifetime (see the appendix for a formal definition). Let $\mathrm{E}v'(c_{0,j}^e)$ denote the average marginal utility of consumption while employed conditional on finding a job in period 0. In addition to the components of interest for calculating the welfare gain, the expressions for $\frac{\partial s_0}{\partial w}$ and $\frac{\partial s_0}{\partial a}$ are affected by a new parameter $\rho = \frac{\mathrm{E}v'(c_{t,j}^e) - \mathrm{E}v'(c_{0,j}^e)}{\mathrm{E}v'(c_{t,j}^e)}$. The parameter ρ is the percentage difference in the marginal utility of an extra dollar while working if one finds a job after a spell of unemployment rather than immediately in period 0. If one assumes $\rho = 0$, one overestimates $\frac{\partial W}{\partial b}$. This upward bias arises because part of the liquidity effect is due to the difference in marginal utilities upon re-employment at different dates rather than the difference

[5] Note that benefits and taxes are not discounted by the interest rate in the budget constraint in (9). This is intended to capture the fact that UI budgets are balanced period-by-period rather than intertemporally in practice. In the stylized single-agent model considered here, period-by-period balancing is impossible because taxes are always collected after benefits are paid out. However, in a more plausible model with agents who have heterogeneous, uniformly distributed layoff dates but are otherwise identical, the planner can meet the constraint of equating benefits paid and taxes collected in each period. The problem specified in (9) can be viewed as a characterization of the optimal policy in this environment.

in marginal utilities between the employed and unemployed states, which is the transfer of interest for the UI system.

In the appendix, I provide a simple calibration argument which shows that ρ is likely to be small in practice. Intuitively, ρ is determined by the extent to which the length of the preceding unemployment spell affects an agent's average level of subsequent consumption *after* starting a new job. Since lost income during unemployment is typically small relative to lifetime wealth, ρ is negligible for plausible values of the coefficient of relative risk aversion. I therefore proceed with the approximation that $\rho = 0$.[6] This leads to the following simple expression for the welfare gain of raising UI benefits in the general case (see Appendix A):

$$\frac{\partial W}{\partial b} = \frac{1-\sigma}{\sigma}\left\{\frac{-\partial s_0^*/\partial a}{\partial s_0^*/\partial w} - \frac{\varepsilon_{D,b}}{\sigma}\right\} \tag{11}$$

where $\sigma = \frac{T-D}{T}$ denotes the fraction of his life that the agent spends employed and $\varepsilon_{D,b} = \frac{b}{D}\frac{\partial D}{\partial b}$ is the elasticity of unemployment durations with respect to benefits. The elasticity $\varepsilon_{d,b}$ includes both the moral hazard and liquidity effects of UI, and can be computed from an estimate of $\partial s_0^*/\partial b$. Equation (11) parallels (8) for the static case, except that the parameters reflect derivatives with respect to permanent changes in annuity income and the UI benefit level rather than changes in period 1 only. Note that the formula still depends on the effects of b and a on s_0^*, i.e. search intensity at the *beginning* of the spell, an issue taken into account in empirical implementation. The optimal benefit level b^* satisfies $\frac{\partial W}{\partial b}(b^*) = 0$.

Discussion. An important implication of (11) is that the optimal benefit level does not necessarily fall with $\varepsilon_{D,b}$, as is commonly thought. It matters whether a higher value of $\varepsilon_{D,b}$ comes from a larger liquidity $(-\partial s_0^*/\partial a)$ or moral hazard $(\partial s_0^*/\partial w)$ component. If the elasticity is large because of the liquidity effect, UI is essentially reducing the need for agents to rush back to work because they have insufficient ability to smooth consumption; if the elasticity is large because of the moral hazard effect, UI is subsidizing unproductive leisure. In this sense, the test provides a new method of quantifying the extent to which the full insurance benchmark is violated. The agent's capacity to smooth marginal utilities is assessed by examining the effect of transitory income shocks

[6] The key point is that consumption reverts to the annuitized value of permanent-income after re-employment. If human capital depreciates rapidly over the unemployment spell – a possibility ruled out by assumption in the present model – the resulting decline in permanent income could make ρ larger. In this case, one can apply equation (22) in Appendix A, which provides an exact formula for $\frac{\partial W}{\partial b}$ that includes the ρ term. An estimate of ρ could in principle be obtained using consumption data or data on the responsiveness of search intensity to benefits and cash grants that vary over the spell.

on the consumption of leisure instead of goods as in earlier studies (Cochrane 1991, Gruber 1997).

More generally, the key idea underlying the test is to measure the value of insurance using revealed preference. The effect of a lump-sum cash grant on the unemployment duration reveals the extent to which the insurance benefit permits the agent to attain a more socially desirable allocation. For example, suppose the agent spends an extra week searching when the benefit level is raised by 10%. If a lump-sum grant equivalent to the increase in the UI benefit has no effect on the duration of search, we infer that the agent is taking more time to find a job purely because of the price subsidy for doing so. In this case, UI does not provide any benefit and simply creates inefficiency by taxing work, leading to $\frac{\partial W}{\partial b} < 0$. In contrast, if the agent raises his duration substantially even when he receives a non-distortionary cash grant, we infer that the UI benefit permits him to make a more (socially) optimal choice, i.e. the choice he would make when incentives are not distorted and markets are more complete. In this case, $\frac{\partial W}{\partial b}$ is larger.

The test proposed here offers an alternative to the consumption-based test for optimal benefits of Baily (1978) and Chetty (2006a) and the reservation-wage test of Shimer and Werning (2007). The three tests each identify a different "sufficient statistic" that fully encodes the marginal value of insurance. One advantage of the liquidity vs. moral hazard test is that it only requires data on unemployment durations, which are typically more precise and widely available than data on consumption or reservation wages. In addition, this method does not rely on state-independence or a specific parametrization of the utility function.

The cost of this parsimonious approach is the assumption that an agent's actions reveal what maximizes his utility. To see why this may be a concern, suppose the provision of UI benefits has no effect on subsequent job match quality. Under the revealed preference test, this finding would have no bearing on the optimal benefit level – it does not matter if the worker rationally chooses to use the money to consume more leisure or search for a better match. However, one may believe that workers have time-inconsistent preferences that lead them to suboptimally reduce search effort when they have more liquidity. In such an environment, where the preferences revealed by choice are not those that the social planner wishes to maximize, data on outcomes are needed to calculate the benefits of insurance. In view of this limitation, the revealed preference test should be interpreted as a *libertarian* benchmark. The test identifies the benefit level that is best from the perspective of correcting market failures without intervening in individual choice.

In the remainder of the paper, I implement (11) by estimating $\frac{\partial s_0^*}{\partial b}$ and $\frac{\partial s_0^*}{\partial a}$ empirically and using the identity in (10) to calculate $\frac{\partial s_0^*}{\partial w}$.

3 Empirical Analysis I: The Role of Constraints

3.1 Estimation Strategy

The model suggests a natural first step in evaluating the empirical relevance of liquidity effects: compare the effect of UI benefits on durations for "unconstrained" individuals – who are able to smooth marginal utilities perfectly and have $\Delta = u'(c_t^u) - v'(c_t^e) = 0$ – and "constrained" individuals – who have $\Delta > 0$. For unconstrained individuals, the benefit-duration elasticity reflects the pure moral hazard effect of UI. For constrained individuals, the elasticity reflects the sum of the moral hazard and liquidity effects. Comparing the elasticity estimates gives an indication of the importance of liquidity relative to moral hazard. For instance, if the effects of benefits on durations were much stronger in the unconstrained group, it would be unlikely that liquidity effects are large.

To implement this heterogeneity analysis, I divide individuals into unconstrained and constrained groups and estimate benefit-duration elasticities for each group using cross-state and time variation in unemployment benefit levels. The ideal definition of the constrained group would be the set of households whose consumption is sensitive to transitory income shocks. Unfortunately, there is no panel dataset that contains high-frequency information on both household consumption and labor supply in the U.S. I therefore use proxies which identify households that are likely to be able to smooth intertemporally using savings and borrowing. Since unemployment shocks are small relative to lifetime wealth, such individuals should have $\Delta \simeq 0$.[7]

The proxy I focus on in this paper is liquid wealth net of unsecured debt, which I term "net wealth." Browning and Crossley (2001), Bloemen and Stancanelli (2005), and Sullivan (2007) report evidence from various panel datasets showing that households with little or no financial assets prior to job loss suffer consumption drops during unemployment that are mitigated by provision of UI benefits. In contrast, households with higher assets exhibit little sensitivity of consumption to unemployment or UI benefit levels. These findings suggest the use of net wealth as a proxy for being unconstrained.

I also consider two secondary proxies: spousal work status and mortgage status prior to job loss. Browning and Crossley find larger consumption drops and higher sensitivity to UI among single-earner households. Their interpretation of this finding is that those with a second income source are more likely to be able to borrow since at least one person is employed.[8] The mortgage

[7] An alternative strategy, which I do not pursue here because of data limitations, is to distinguish households by their ability to smooth consumption across states through risk-sharing mechanisms.

[8] A countervailing effect is that households with a single earner may be able to maintain their prior standard of

proxy is motivated by Gruber's (1998) finding that fewer than 5% of the unemployed sell their homes during a spell, whereas renters move much more frequently. Consequently, an individual making mortgage payments before job loss effectively has less ability to smooth the remainder of his consumption (Chetty and Szeidl 2007), and is more likely to be constrained than a renter.

Although these proxies predict being constrained on average, they are imperfect predictors for two reasons. First, some households classified as unconstrained are presumably misallocated to the constrained group and vice-versa. Second, no household truly has $\Delta = 0$ because insurance markets are likely to be incomplete, and intertemporal smoothing itself cannot fully eliminate consumption fluctuations. There is probably a small liquidity effect even among the groups classified as unconstrained. Since I attribute the entire response among the group classified as unconstrained to moral hazard, both of these misclassification errors lead to underestimation of the liquidity effect relative to moral hazard.

3.2 Data

I use data from the Survey of Income and Program Participation (SIPP) panels spanning 1985-2000. Each panel of the SIPP follows households for a period of two to four years. Relative to other widely used datasets for the U.S. such as the CPS and PSID, the main benefits of the SIPP are the availability of asset data, high-frequency information on employment status, and large sample size. At each interview, households are asked questions about their activities during the past four months, including weekly labor force status. Unemployed individuals are asked whether they received unemployment benefits in each month.

Starting from the universe of job separations in the pooled SIPP panels, I restrict attention to prime-age males who (a) report searching for a job, (b) are not on temporary layoff, (c) have at least three months of work history in the survey (so that pre-unemployment wages can be computed), and (d) took up UI benefits within one month after job loss.[9] Details on the sample construction and SIPP database are given in Appendix B. The restrictions leave 4,560 unemployment spells in the core analysis sample. Asset data are generally collected only once in each panel, so pre-

living more easily if the other earner can enter the labor force to make up for the lost income. Browning and Crossley's findings suggest that this effect is dominated by the added intertemporal smoothing capacity of dual earners, so that on net households with two earners are less constrained.

[9]Restricting the sample to those who take up UI could lead to selection bias because the takeup decision is endogenous to the benefit level (Anderson and Meyer 1997). I find that the elasticity of takeup with respect to the benefit level is similar across the constrained and unconstrained groups, suggesting that endogeneity is unlikely to be responsible for the heterogeneous effects estimated below.

unemployment asset data is available for approximately half of these observations.

The first column of Table 1 gives summary statistics for the core sample. Monetary values are in real 1990 dollars in this and all subsequent tables. The median UI recipient is a high school graduate and has pre-UI gross annual earnings of $20,726. Perhaps the most striking statistic is pre-unemployment wealth: median liquid wealth net of unsecured debt is only $128, suggesting that many unemployed individuals may not be in a position to smooth consumption while unemployed.

Information on UI laws was obtained from the Employment and Training Administration (various years) and supplemented with information directly from individual states. Unfortunately, measurement error and inadequate information about pre-unemployment wages for many claimants make it difficult to predict each claimant's benefit level precisely. I therefore use three independent methods to proxy for each claimant's (unobserved) actual UI benefits. First, I use average benefits for each state/year pair obtained from the Department of Labor in lieu of each individual's actual UI benefit amount. Second, I proxy for the actual benefit using maximum weekly benefit amounts, which are the primary source of variation in benefit levels across states, since most states replace 50% of a claimant's wages up to a maximum benefit level. Third, I simulate each individual's weekly UI benefit using a two-stage procedure. In the first stage, I predict each claimant's pre-unemployment annual income using education, age, occupation, and other demographics. In the second stage, I predict each claimant's unemployment benefits using a simulation program that assigns each claimant a benefit based on the predicted wage, state, and year of claim. See Appendix B for further details on the motivation for and implementation of this two-stage procedure.

3.3 Results

3.3.1 Graphical Evidence and Non-Parametric Tests

I begin by providing graphical evidence on the effect of unemployment benefits on durations in constrained and unconstrained groups. First consider the asset proxy for constraints. I divide households into four quartiles based on their net liquid wealth. Table 1 shows summary statistics for each of the four quartiles. Households in the lower net liquid wealth quartiles are poorer and less educated, but the differences between the four groups are not very large. As a result, UI benefit levels are fairly constant across the groups. In particular, the replacement rate – defined as each individual's simulated unemployment benefit divided by his predicted wage – is close to 50% on average in all four quartiles. This similarity of benefit and income levels suggests that differences in benefit-duration elasticities across the quartiles are unlikely to be driven purely by differences in the levels around which the elasticities are estimated.

Figures 1a-d show the effect of UI benefits on job-finding rates for households in the each of the four quartiles of the net wealth distribution. Since ex-post asset levels are endogenous to duration of unemployment, households for whom asset data are available only after job loss are excluded when constructing these figures. Including these households turns out to have little effect on the results, as we will see below in the regression analysis. I construct the figures by first dividing the full sample of UI claimants into two categories: those that are in (state, year) pairs that have average weekly benefit amounts above the sample median and those below the median. I then plot Kaplan-Meier survival curves for these two groups using the households in the relevant net wealth quartile. Note that the differences in average individual replacement rates between the low and high-benefit are fairly similar in the four quartiles.

These and all subsequent survival curves plotted using the SIPP data are adjusted for the "seam effect" in panel surveys. Individuals are interviewed at 4 month intervals in the SIPP and tend to repeat answers about weekly job status in the past four months. Consequently, a disproportionately large number of transitions in labor force status are reported on the "seam" between interviews, leading to artifical spikes in the hazard rate at 4 and 8 months. These spikes are smoothed out by fitting a Cox model with a time-varying indicator for being on a seam between interviews, and then recovering the baseline hazards to construct a seam-adjusted Kaplan-Meier curve. The resulting survival curves give the probability of remaining unemployed after t weeks for an individual who never crosses an interview seam. The results are similar if the raw data is used without adjusting for the seam effect.

Figure 1a shows that higher UI benefits lead to much lower job-finding rates for individuals in the lowest wealth quartile. For example, 15 weeks after job loss, 55% of individuals in low-benefit state/years are still unemployed, compared with 68% of individuals in high-benefit state/years. A nonparametric Wilcoxon test rejects the null hypothesis that the survival curves are identical with $p < 0.01$.[10] Figure 1b constructs the same survival curves for the second wealth quartile. UI benefits have a smaller effect on durations in this group. At 15 weeks, 63% of individuals in the low-benefit group are still unemployed, vs. 70% in the high benefit group. Equality of the survival curves is rejected with $p = 0.04$. Figures 1c and 1d show that effect of UI on durations

[10] The non-parametric test is conducted on the raw data because adjusting for the seam effect requires a parametric assumption about the hazard rate.

virtually disappears in the third and fourth quartiles of the wealth distribution. Not surprisingly, the equality of the survival curves is not rejected in these two groups. The fact that UI has little effect on durations in the unconstrained groups suggests that it induces little moral hazard among these households.

The secondary proxies confirm these results. Figure 2a shows that UI benefits have a clear, statistically significant effect on job finding rates among households that are paying off mortgages prior to job loss. In contrast, Figure 2b shows that the effect is smaller for households that are not paying off mortgages and are hence less constrained.[11] Results are similar for the spousal work proxy: UI benefits have a much larger effect on job finding hazards for single-earner families than dual-earner families (see Figure 2 in the working paper).

The preceding results show that the *interaction* effect of UI benefits and wealth (or other proxies) on durations is negative. An alternative approach to evaluating the importance of liquidity is to study the direct effect of the cross-sectional variation in wealth on durations, testing in particular if durations are an increasing and concave function of wealth. I focus on the variation in UI benefits because changes in UI laws are credibly exogenous to individuals' preferences. In contrast, conditional on demographics and income, cross-sectional variation in wealth holdings arises from heterogeneity in tastes for savings, confounding the effect of wealth on duration in the cross-section. For example, UI claimants with higher assets are also likely to have lower discount rates or higher anticipated expenses (e.g., college tuition payments), and hence may be reluctant to deplete their assets to finance a longer spell of unemployment.[12] In practice, I find no robust relationship between assets and unemployment durations in the cross-section (as indicated by the mean durations by quartile reported in Table 1), suggesting that the use of exogenous variation such as UI benefits is quite important. The same issue also motivates my focus on severance pay as a source of variation in wealth to identify the liquidity effect in section 4.

3.3.2 Hazard Model Estimates

I evaluate the robustness of the graphical results by estimating a set of Cox hazard models in Table 2. Let $h_{i,t}$ denote the unemployment exit hazard rate for individual i in week t of an unemployment

[11] In contrast with the other proxies, the constrained types in this specification (homeowners with mortgages) have *higher* income, education, and wealth than the unconstrained types, who are primarily renters. This makes it somewhat less likely that the differences in the benefit elasticity of duration across constrained and unconstrained groups is spuriously driven by other differences across the groups such as income or education.

[12] More generally, wealthier individuals may have unobserved characteristics (e.g. skills, job search technologies) that lead to different durations for reasons unrelated to their wealth.

spell, α_t the "baseline" hazard rate in week t, b_i the unemployment benefit level for individual i, and $X_{i,t}$ a set of controls. Throughout, I censor durations at 50 weeks to reduce the influence of outliers and focus on search behavior in the year after job loss.

Since the formula in (11) calls for estimates of the effect of UI benefits on search behavior at the beginning of the spell, I estimate hazard models of the following form:

$$\log h_{i,t} = \alpha_t + \beta_1 \log b_i + \beta_2 t \times \log b_i + \beta_3 X_{i,t} \tag{12}$$

Here, the coefficient β_1 gives the elasticity of the hazard rate with respect to UI benefits at the beginning of the spell ($t = 0$) because the interaction term $t \times \log b_i$ captures any time-varying effect of UI benefits on hazards. Note that the search model does not make a clear prediction about the sign of β_2. The effect of UI benefits could diminish over time ($\beta_2 < 0$) because the number of weeks for which benefits remain available is falling. But β_2 could also be positive because households are increasingly constrained and thus more sensitive to cash-on-hand late in the spell. In practice, there is no robust, statistically significant pattern in the β_2 coefficients across the quartiles, and I therefore do not report them in Table 2 in the interest of space.[13]

I first estimate (12) on the full sample to identify the unconditional effect of UI on the hazard rate. In this specification, as in most others, I use the average UI benefit level in the individual's (state,year) pair to proxy for b_i in light of the measurement-error issues discussed above. This specification includes the following controls: state, year, industry, and occupation fixed effects; a 10 piece log-linear spline for the claimant's pre-unemployment wage; linear controls for total (illiquid+liquid) wealth, age, education; and dummies for marital status, and being on the seam between interviews to adjust for the seam effect. Standard errors in this and all subsequent specifications are clustered by state. The estimate in column 1 of Table 2 indicates that a 10% increase in the UI benefit rate reduces the hazard rate by 5.3% in the pooled sample, consistent with the estimates of prior studies. The estimate of $\beta_2 = 0.001$ (s.e. $= 0.008$), indicating that there is no detectable variation in the UI benefit effect over the spell.

To examine the heterogeneity of the UI effect, I estimate separate coefficients for each of the four quartiles of the net wealth distribution. These specifications include all households for which asset data are available either before or after the spell. Consistent with the graphical evidence,

[13] The only stable pattern across the specifications is that β_2 is slightly negative in the highest wealth quartile (around -0.03). This could be because households that are intially unconstrained become increasingly liquidity constrained as they deplete their buffer stocks.

the estimates are similar (but less precise) if only households with ex-ante asset data are included.

Let $Q_{i,j}$ denote an indicator variable that is 1 if agent i belongs to quartile j of the wealth distribution. Let $\alpha_{t,j}$ denote the baseline job-finding hazard for individuals in quartile j in week s of the unemployment spell. Columns 2-5 of Table 2 report estimates of $\{\beta_1^j\}_{j=1,2,3,4}$ from the following stratified Cox model:

$$\log h_{itj} = \alpha_{t,j} + \beta_1^j Q_{i,j} \log b_i + \beta_2^j Q_{i,j}(t \times \log b_i) + \beta_3 X_{itj} \tag{13}$$

In this equation, β_1^j corresponds to the elasticity of the hazard rate w.r.t. UI benefits at $t = 0$ in quartile j of the net wealth distribution. Specification (2) of Table 2 reports estimates of (13) with no controls (no X). The effect of UI benefits declines monotonically with net wealth. Among households in the lowest quartile of net wealth, a 10% increase in UI benefits reduces the hazard rate by 7.2%. In contrast, there is a much weaker association between the level of UI benefits and the hazard among households in the third and fourth quartiles of net wealth. The null hypothesis that UI benefits have the same effect on hazard rates in the first and fourth quartiles is rejected with $p < 0.05$, as is the null hypothesis that the mean UI effect for below-median wealth households is the same as that for above-median wealth households.

Specification (3) replicates (2) with the full set of controls used in column (1), including state and year fixed effects so that the coefficients are identified from changes in UI laws within states rather than cross-state comparisons. This specification also includes interactions of the wage spline and industry/occupation dummies with the wealth quartile indicators, allowing these variables to have different effects across the quartiles. The pattern of the coefficients is unchanged, although the magnitudes of the coefficients in the first three quartiles is somewhat larger, perhaps because exogenous changes in UI laws are more effectively isolated when the controls are included.

In specifications (4) and (5), I explore robustness to changes in the definition of b_i. Both of these specifications include the control set used in (3). Column (4) uses the maximum UI benefit level in individual i's state/year and column (5) uses the simulated benefit for each individual i using the two-stage procedure described above. In the maximum benefit specification, the coefficient estimates are all smaller than their counterparts in (3), but the pattern is preserved: the effect of benefits is larger for low-wealth individuals and the hypothesis tests of equivalent effects in the lower and upper quartiles are both rejected with $p < 0.01$. In the individual simulated benefit specification, UI benefits are estimated to have little effect on durations in the highest wealth quartile, and the elasticity estimates are declining from quartiles 2 to 4. However, the estimate for the first quartile is smaller than that in the second quartile, breaking the monotonic declining pattern obtained with the other measures of benefits.

I have estimated a set of specifications analogous to (13) for the spousal work and mortgage proxies. An example is in column 6 of Table 2, which reports estimates of the effect of UI benefits on job-finding hazards for households with and without a mortgage prior to job loss. This specification includes the same controls as in column 3, except that the relevant covariates are interacted with the mortgage indicator rather than the asset quartiles. The estimates indicate that benefits have a considerably larger effect on durations among households that have mortgages. See Table 3 in the working paper for additional estimates using the spousal work and mortgage proxies.

I have also fit a variety of other specifications to further probe the robustness of the results (see Table 2b in the working paper). The estimates are similar when high income individuals are excluded or temporary layoffs are included. Results are also similar with controls for the average wage income in each state and year from the BLS, or when the wealth quartiles are defined in terms of wealth divided by wages. Unlike with liquid wealth, I find no systematic link between home equity and the benefit-duration elasticity. This is consistent with the importance of liquidity, since accessing home equity is difficult when one is unemployed (Hurst and Stafford 2004). Finally, I find no relationship between the level of benefits and durations for a control group of individuals who do not receive UI . This "placebo test" supports the identification assumption that the variation in UI benefits is orthogonal to unobservable determinants of durations.

In summary, the SIPP data indicate that the link between unemployment benefits and durations documented in earlier studies is driven by a subset of the population that has limited ability to smooth consumption. This pattern is suggestive of a substantial liquidity effect. If one were to assume that substitution effects are similar across unconstrained and constrained groups, this evidence would be sufficient to infer that liquidity effects are large. However, this assumption may be untenable: households with low liquidity might have different preferences that generate larger substitution effects than unconstrained households.[14] I therefore turn to a second empirical strategy to identify the magnitude of the liquidity effect.

[14] Even if constrained and unconstrained households have the same utility functions, their preferences could differ locally because their mean levels of consumption and assets differ. This is a further reason for caution in making inferences about the magnitude of the liquidity effect purely from the evidence documented thus far.

4 Empirical Analysis II: Severance Pay and Durations

4.1 Estimation Strategy

The ideal way to estimate the liquidity effect would be a randomized experiment where some job losers are given lump-sum grants while others are not. Since lump-sum grants do not distort marginal incentives to search, they yield an estimate of the pure liquidity effect. Lacking such an experiment, I exploit variation in severance pay policies across firms in the U.S.[15] Severance payments are lump-sum cash grants made at the time of job loss. The most common severance policy is one week of wages per year of service at the firm; however, some companies have flatter or steeper profiles with respect to job tenure, and others make no severance payments at all (Lee Hecht Harrisson 2001). Many companies have minimum job tenure thresholds to be eligible for severance pay, ranging from 3 to 5 years. For regular employees, there is little variation in severance packages within a given firm and tenure bracket. Hence, conditional on tenure, the variation in receipt of severance pay is driven primarily by cross-firm differences in policies.

I estimate the effect of severance pay on unemployment durations using hazard models similar to those above:

$$\log h_{i,t} = \alpha_t + \theta_1 \text{sev}_i + \theta_2 \text{sev}_i \times t + \gamma X_{i,t} \tag{14}$$

where sev_i is an indicator for receipt of severance pay. The coefficient θ_1 identifies the effect of lump sum grants on job finding hazards at the beginning of the spell if receipt of severance pay is orthogonal to other determinants of durations. After estimating the baseline model, I evaluate this key orthogonality condition.

4.2 Data

The data for this portion of the study come from two surveys conducted by Mathematica on behalf of the Department of Labor. The datasets contain information on unemployment durations, demographic characteristics, and most importantly for this study, an indicator for receipt of severance pay. The first dataset is a representative sample of job losers in Pennsylvania in 1991. The second dataset is a sample of unemployment durations in 25 states in 1998 that oversamples UI exhaustees.

[15] Receipt of severance pay intended to supplement UI benefits typically does not affect eligibility for UI, although some states can delay benefits if the claimant receives "wages in lieu of notice" (Kodryzcki 1998, McCulloch 1998). In Pennsylvania, the unemployment compensation law explicitly states that severance pay does not affect UI benefits (Pennsylvania Department of Labor and Industry 2007). Restricting the analysis to the Pennsylvania dataset below yields results similar to those obtained for the pooled sample.

Throughout the analysis, the data are reweighted using the sampling weights to obtain estimates for a representative sample of job losers.

For comparability to the preceding results, I make the same exclusions after pooling the two datasets to arrive at the final sample used in the analysis. In particular, I include only prime-age males and discard all individuals who expected a recall at the time of layoff (including temporary layoffs does not affect the results, as above). These exclusions leave 2,441 individuals in the sample, of whom 471 (18%) report receiving a severance payment at job loss. Details on the Mathematica datasets and sample construction are given in Appendix C.

Two measures of "unemployment duration" are available in these datasets: (1) an administrative record of the number of weeks for which UI benefits were paid ("compensated duration") and (2) the number of weeks from the end date of the individual's previous job to the (self-reported) start of the next job ("unemployment duration"). For consistency with the SIPP estimates, I focus on the second measure here. Results are similar, and more precisely estimated, using the administrative measure.

Table 3 shows summary statistics for severance pay recipients and non-recipients. The sample generally looks quite similar on observables to the SIPP sample used above. Given the minimum tenure eligibility requirement, it is not surprising that severance pay recipients have much higher median job tenures than non-recipients. Correspondingly, severance pay recipients are older and higher in observable characteristics than non-recipients. These differences underscore why one must be careful in drawing causal inferences from comparing severance pay recipients and non-recipients.

4.3 Results

I begin again with graphical evidence. Figure 3 shows Kaplan-Meier survival curves for two groups of individuals: those who received severance pay and those who did not. Since pre-unemployment job tenure is an important determinant of severance pay and is also highly positively correlated with durations, I control for it throughout the analysis. These survival curves have been adjusted for tenure by fitting a cox model with tenure as the only regressor and recovering the baseline hazards for each group. Severance pay recipients have significantly lower job finding hazards. As a result, 68% of individuals who received severance pay remain unemployed after 10 weeks, compared with 75% among those who received no severance payment.

An obvious concern in interpreting this result as evidence of a liquidity effect is that it may reflect correlation rather than causality because severance pay recipients differ from non-recipients. For instance, firms that offer severance packages might do so because their workers have accumulated more specific human capital and are likely to take a long time to find a suitable new job. This would induce a spurious correlation between severance pay and durations in the cross-section.

I use two approaches to examine the causality of severance pay. First, I investigate whether the effect of severance pay differs across constrained and unconstrained groups. The model in section 2 indicates that severance pay – which is a minor fraction of lifetime wealth – should causally affect durations only among households that cannot smooth consumption. In contrast, alternative explanations such as the one proposed above would not necessarily predict a differential effect of severance pay across constrained and unconstrained households. Hence, the heterogeneity of the estimated severance pay effect yields insight into the causality of severance pay.

Implementing this test requires division of households into constrained and unconstrained groups. Unfortunately, the Mathematica surveys do not contain data on assets and the other proxies for constraint status used in the SIPP data. To overcome this problem, I predict assets for each household with an equation estimated using OLS on the SIPP sample. The prediction equation is a linear function of age, wage, education, and marital status. I then divide households into two groups, above and below the median level of predicted assets. Note that results based on predicted assets (using the same prediction equation) and reported assets are similar in the SIPP data: the total elasticity of duration w.r.t. UI benefits is much larger among households with predicted assets below the median than for those above the median. Hence, the predicted asset measure succeeds in identifying the households whose search behavior is sensitive to UI benefits.

Figures 4a-b replicate Figure 3 for the two groups. Figure 4a shows that receipt of severance pay is associated with a large increase in survival probabilities for constrained (low asset) households. Figure 4b shows that severance pay has a much smaller effect on search behavior for households that are likely to be wealthier. Results are similar if households are split into constrained and unconstrained groups on the basis of age or income alone. Results are also unaffected by changes in the functional form of the asset prediction equation, prediction via quantile regression instead of OLS, or trimming of outliers. The fact that severance pay affects durations only in the group of households that are sensitive to UI benefits (those who are likely to be constrained) supports the claim that liquidity effects drive a substantial portion of the UI-duration link.

As a second approach to examining the causality of severance pay, I assess the sensitivity of the severance pay effect to controlling for observed heterogeneity. I estimate variants of the Cox model in (14), censoring all durations at 50 weeks as in the SIPP data. I first estimate a model with only a linear tenure control and a the time-varying interaction of severance pay with weeks unemployed. I then estimate the model with the following control set: ten piece linear splines for log pre-unemployment wage and job tenure; dummies for prior industry, occupation, and year; and controls for age, marital status, and education (using a dummy for dropout status and college graduation). The first two columns of Table 4 show that receipt of severance pay is estimated to lower the job-finding hazard at the beginning of the spell by $\theta_1 = -18\%$ in the tenure-control and $\theta_1 = -23\%$ in the full-control specification. The estimated value of $\theta_2 = 1.3\%$ (s.e. $= 0.2\%$) in both specifications. The effect of severance pay on search intensity diminishes over time, as one might expect if individuals deplete the grant over the course of the spell.

Specifications (3) and (4) estimate separate severance pay coefficients for constrained (below-median predicted assets) and unconstrained (above-median) households. The baseline hazards are stratified by predicted wealth group (above/below median) and the wage spline and industry/occupation dummies are interacted with the predicted wealth dummy, as in the SIPP specifications.[16] Consistent with Figure 4, the estimates indicate that severance pay reduces initial job finding hazards in the low-wealth group by 46-49%, but has little or no effect in the high-wealth group. The hypothesis that the effect of severance pay is the same in the low and high wealth groups is rejected with $p < 0.01$. Predicting whether the individual is above or below median wealth directly in the first stage, and interacting the predicted probability with the severance pay dummy when estimating the Cox model yields the same conclusion.

Although these findings all point toward a substantial liquidity effect, the evidence cannot be viewed as fully conclusive because the variation in cash grants is not randomized. One might, for instance, be concerned that low wealth workers with job-specific capital are especially likely to select into firms that offer severance payments, explaining the differential correlations in Figures 4a-b. Such concerns can be addressed only with a research design where the variation is severance pay is known to be exogenous. Reassuringly, in a followup study that exploits quasi-experimental variation in severance payments created by a discontinuity in the Austrian severance pay system, Card, Chetty, and Weber (2007) document a substantial effect of severance pay on durations, consistent with the evidence here. Further analysis along these lines is needed to obtain the most precise and compelling estimates of the liquidity effect in the U.S.

[16] Unlike in the SIPP specifications, I restrict the time interaction of severance pay with weeks unemployed to be the same across wealth groups here in order to increase power. Introducing an interaction yields similar point estimates but larger standard errors because the number of severance pay recipients is relatively small.

5 Calibration: Welfare Implications

I now combine the severance pay and UI benefit estimates to calculate how much of the benefit-duration link is due to the liquidity effect. Recall that the liquidity effect $\partial s_0/\partial a$ depends on the effect of a \$1 increase in an annuity payment. However, severance payments are typically one-time grants made at the start of the unemployment spell, i.e. variation in A_0. Mapping the estimate of $\partial s_0/\partial A_0$ in section 4 into $\partial s_0/\partial a$ requires some assumptions about the agent's intertemporal preferences. Assume that $r = \delta$, and let $k(\delta) = \sum_{t=0}^{T-1} \frac{1}{(1+\delta)^t}/T < 1$ represent the average discount factor applied to the annuity stream. In Appendix A, I derive the following lower bound for the magnitude of the liquidity effect:

$$|\frac{\partial s_0}{\partial a}| \ \geq \widehat{D}|\frac{\partial s_0}{\partial A_0}|k(\delta) \tag{15}$$

where $\widehat{D}$ is the mean compensated unemployment duration, the expected number of weeks for which benefits are received.[17] To simplify this expression further, consider the approximation that c_t^u and c_t^e do not vary with t, i.e. the consumption paths when unemployed and re-employed are flat. When $\delta = 0$, this leads to

$$\frac{\partial s_0}{\partial a} = \widehat{D}\frac{\partial s_0}{\partial A_0} \tag{16}$$

To understand the intuition for these formulas, first consider (16). When c_t^e does not vary with the date of re-employment, only the portion of the annuity received while unemployed affects search behavior, since the annuity adds to both V_t and U_t equally after reemployment. If $u'(c_t^u)$ is constant over the spell and $\delta = 0$, the date at which this payment is received does not matter. Thus the annuity has exactly the same effect on s_0 as a cash grant of equal expected value, \$$\widehat{D}$. When $\delta > 0$, (16) overstates the magnitude of $\frac{\partial s_0}{\partial a}$ because a \$$\widehat{D}$ grant is more valuable than a \$1 annuity payment. This is the source of the $k(\delta)$ shrinkage factor in (15). Finally, if one takes account of the fact that the consumption paths decline with spell length, the annuity becomes relatively more valuable because it provides income in states where the marginal utility is higher. This approximation makes (16) *understate* the true liquidity effect, and leads to the inequality in (15).

[17] In the model, there is no difference between the compensated and total duration because UI benefits are permanent. However, if UI benefits are provided at a constant level for T periods and the agent lives for $T' > T$ periods, the liquidity effect equals the effect of an annuity that lasts for T periods on s_0. This is why the mean compensated duration and the maximum duration of UI benefits are relevant for assessing the present value of the annuity.

Since the maximum duration of unemployment benefits is relatively short in the U.S. (six months), I assume that $\delta = r = 0$ over this horizon and use the approximation of flat consumption paths. I calculate the liquidity effect as a percentage of the total UI benefit effect using (10) and (16). I first rescale the estimates of the severance pay and UI benefit effects into comparable units. Based on a survey of firms by Lee Hecht Harrison (2001), I calculate that the mean severance payment equals 6.7 weeks of wages (see Appendix C). At the mean wage replacement rate of 50%, this corresponds to 13.4 weeks of UI benefits. The mean weeks of compensated unemployment in the Mathematica sample is $\widehat{D} = 15.8$. Hence, receipt of severance pay is equivalent to an annuity that pays $13.4/15.8 = 85\%$ of the UI benefit level each week.

To map the hazard model estimates into values of $\frac{\partial s_0}{\partial a}$ and $\frac{\partial s_0}{\partial b}$, let h_0 denote the baseline hazard rate in the first period for an agent who receives the mean UI benefit and does not receive a severance payment. Let h_1^s denote the hazard for an individual who receives the severance payment and h_1^b denote the hazard for an individual whose UI benefit level is doubled. Using the full-controls estimates from column 1 of Table 2 and column 2 of Table 4 yields $h_1^b/h_0 = \exp(-0.58)$ and $h_1^s/h_0 = \exp(-0.23)$. It follows that

$$\frac{\partial s_0}{\partial a}/\frac{\partial s_0}{\partial b} = \frac{[h_1^s - h_0]/h_0}{0.85 \times [h_1^b - h_0]/h_0} = \frac{\exp(-0.23) - 1}{\exp(-0.53) - 1} \times \frac{1}{0.85} = 0.6.$$

The point estimates imply that the liquidity effect accounts for roughly 60% of the effect of UI benefits on durations in the U.S., leaving a residual of 40% accounted for by moral hazard. While these numbers should be interpreted as rough estimates given sampling error and the approximations made above, the calculation indicates that the liquidity effect is substantial.

Welfare Gain. Using the preceding estimate, one can calculate $\frac{\partial W}{\partial b}$ by specifying $\varepsilon_{D,b}$ and the unemployment rate $1-\sigma$. To calculate $\varepsilon_{D,b}$, recall from specification 1 of Table 2 that the effect of benefits on the hazard rate does not vary significantly over the spell. Using the approximation that the benefit elasticity and baseline hazard rates are constant, $D = \frac{1}{h}$ and $\varepsilon_{D,b} = -\frac{\partial \log h_0}{\partial \log b} = 0.53$. Shimer and Werning (2007) report that there are 7.7 million unemployed individuals and 135 million workers in the U.S., i.e. $\sigma = \frac{135}{142.7} = 0.946$. Plugging these values into (11) yields

$$\frac{\partial W}{\partial b} == \frac{1-\sigma}{\sigma}\{\frac{-\partial s_0^*/\partial a}{\partial s_0^*/\partial b - \partial s_0^*/\partial a} - \frac{\varepsilon_{D,b}}{\sigma}\} = \frac{1-0.946}{0.946} \times (\frac{0.6}{0.4} - \frac{0.53}{0.946}) = 0.05$$

Starting from the mean benefit level in the U.S. over the sample period, a \$1 balanced-budget increase in the weekly benefit level would have raised each individual's welfare by the equivalent of a 5 cent increase in weekly wage income, or \$2.60 per year. Aggregating over the population, the net welfare gain of raising the UI benefit level by a dollar is equivalent to an increase in GDP of \$338 million. Starting from a mean benefit of \$200 per week, raising the benefit level by 10% would yield a gain equivalent to about \$6.7 billion, roughly 0.1 percent of GDP.[18] Since $\frac{\partial W}{\partial b} > 0$ at prevailing benefit levels, the optimal wage replacement rate for UI exceeds 50%. A replacement rate of 50% might be near optimal since $\frac{\partial W}{\partial b}$ is not very large.

This calculation should be viewed as a rough first pass because it abstracts from many factors. First, it focuses only on the duration margin. UI benefits can distort other margins of behavior such as the incidence of layoffs, which may make a lower benefit level desirable (Feldstein 1978, Topel 1983). Second, the analysis abstracts from the general equilibrium effects of UI, which as pointed out by Acemoglu and Shimer (1999) could make a higher benefit desirable. Finally, the calculation assumes a representative agent. Given the concentration of the incidence of unemployment, incorporating heterogeneity could affect the calculation.

Keeping these caveats in mind, it is interesting to compare the result of the revealed preference test with those of consumption and reservation-wage based studies. Baily (1978) and Gruber(1997) find that $\frac{\partial W}{\partial b}$ is substantially negative using plausible estimates of $\varepsilon_{d,b}$ and the consumption drop associated with job loss. Indeed, their preferred calibrations using log utility imply an optimal benefit rate close to zero. The reason for the sharp discrepancy is that the liquidity effect estimated here is inconsistent with the relatively low degrees of risk aversion considered by Baily and Gruber.[19] Shimer and Werning (2007) implement their formula using an estimate of the sensitivity of reservation wages to benefits from Feldstein and Poterba (1984). They find that a \$1 increase in the benefit level would yield a net welfare gain equivalent to raising GDP by \$2.4 billion, six times larger than the estimate obtained here. As Shimer and Werning emphasize, the credibility of existing reservation wage elasticity estimates is questionable, particularly in view of more recent evidence that UI benefit levels have little impact on subsequent wage rates.

The moral hazard vs. liquidity analysis also has some policy implications beyond the level of benefits. For example, a commonly held view is that policies which increase the marginal incentive to search – such as job-finding bonuses or more stringent search requirements – can

[18] This welfare gain calculation applies only *locally* at the level of benefits in the data. As b approaches the full insurance level, $\partial s_0^*/\partial a \to 0$ and $\frac{\partial W}{\partial b}(b) < 0$. Conversely, the marginal welfare gain of UI is likely to be larger at low benefit levels.

[19] A strength of the formula proposed here is that it does not require an assumption about the level of risk aversion, which can vary with context and may be difficult to estimate (Chetty and Szeidl 2007, Shimer and Werning 2007).

raise welfare significantly by reducing the moral hazard problem. The evidence here suggests that efforts to shorten durations through such reforms would yield welfare gains 60% smaller than suggested by studies that attribute the entire duration response to a substitution effect. The analysis also sheds light on the debate regarding means-testing of temporary income assistance programs. Recent evidence that UI does not smooth consumption for those who have high levels of pre-unemployment assets points in favor of asset-testing. However, UI does not appear to significantly affect unemployment durations for this group either. Since means-testing generates a distortionary incentive to save less, a universal benefit may maximize welfare.

6 Conclusion

This paper has developed a "revealed preference" approach to analyzing optimal unemployment insurance. UI benefits affect search behavior through both a welfare-enhancing "liquidity effect" (due to increased consumption-smoothing ability) and a welfare-reducing "moral hazard" effect (due to a distortion in marginal incentives). The ratio of the liquidity effect to the moral hazard effect reveals the extent to which UI permits the agent to attain a more desirable allocation, and encodes the marginal welfare gain of raising the benefit level. Using data on unemployment durations from the U.S., I estimate that 60% of the effect of UI benefits on durations is due to the liquidity effect rather than moral hazard. This estimate implies that increasing the UI benefit level from prevailing rates would yield a small welfare gain. Hence, a wage replacement rate of 50% is near optimal under the libertarian assumption that agents' choices reveal their true preferences.

While providing households' with resources to smooth income shocks appears to be valuable, one must be careful in drawing implications for government policy from this finding. Even if insuring unemployment is beneficial, government transfers may not be the best means of providing such insurance. Feldstein and Altman (1998) and Shimer and Werning (2006) argue that UI savings accounts or low-interest loans are a better means of providing liquidity than transfers.

Finally, the revealed preference approach to valuing insurance proposed here can be applied to a broad range of social and private insurance markets beyond unemployment. Several studies have documented large behavioral responses to health insurance, disability insurance, workers compensation, and social security. Investigating the relative importance of moral hazard vs. liquidity in these programs is a promising method of evaluating the welfare consequences of these programs.[20]

[20] Nyman (2003) and Autor and Duggan (2007) discuss related issues in health and disability insurance.

References

Acemoglu, Daron and Robert Shimer, "Efficient Unemployment Insurance," Journal of Political Economy 107: 893-928, 1999.

Acemoglu, Daron and Robert Shimer, "Productivity Gains from Unemployment Insurance," European Economic Review 44(7): 1195-1224, 2000.

Anderson, Patricia and Bruce Meyer, "Unemployment Insurance Takeup Rates and the After-Tax Value of Benefits,"Quarterly Journal of Economics 112(3): 913-937, 1997.

Atkinson, Anthony and John Micklewright, "Unemployment Compensation and Labor Transitions: A Critical Review," Journal of Economic Literature 29 (1991), 1679-1727.

Autor, David and Mark Duggan, "Income vs. Substitution Effects in Disability Insurance" American Economic Review Papers and Proceedings 97(2): 119-124, 2007.

Baily, Martin N., "Some aspects of optimal unemployment insurance," Journal of Public Economics, X (1978), 379-402.

Bloemen, Hans & Elena Stancanelli, "Financial Wealth, Consumption Smoothing and Income Shocks Arising from Job Loss," Economica 72(3), pp. 431-452, 2005.

Browning, Martin and Thomas Crossley, "Unemployment Insurance Levels and Consumption Changes," Journal of Public Economics 80 (2001), 1-23.

Card, David, Raj Chetty, and Andrea Weber, "Cash-on-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market," NBER Working Paper 12639, 2007.

Chetty, Raj, "Why do Unemployment Benefits Extend Unemployment Durations? Moral Hazard vs. Liquidity" NBER Working Paper 11760, 2005.

Chetty, Raj, "A General Formula for the Optimal Level of Social Insurance," Journal of Public Economics, 2006a.

Chetty, Raj, "A New Method of Estimating Risk Aversion." The American Economic Review XCVI (2006b), 1821-1834.

Chetty, Raj and Adam Szeidl, "Consumption Commitments and Risk Preferences," Quarterly Journal of Economics 122(2): 831-877, 2007.

Cochrane, John, "A Simple Test of Consumption Insurance," Journal of Political Economy 99(5): 957-76, 1991.

Cullen, Julie and Jonathan Gruber, "Does Unemployment Insurance Crowd Out Spousal Labor Supply?" Journal of Labor Economics 18 (2000), 546-572.

Employment and Training Administration, Significant Provisions of State Unemployment Insurance Laws, Washington D.C.: U.S. Government Printing Office (various years).

Feldstein, Martin, "The Effect of Unemployment Insurance on Temporary Layoffs," American Economic Review 68 (1978), 834-46.

Feldstein, Martin and Daniel Altman, "Unemployment Insurance Savings Accounts," National Bureau of Economic Research Working Paper No. 6860 (1998).

Feldstein, Martin and James Poterba, "Unemployment Insurance and Reservation Wages," Journal of Public Economics 23(1-2): 141-67, 1984.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," American Economic Review 87 (1997), 192-205.

Gruber, Jonathan, "Unemployment Insurance, Consumption Smoothing, and Private Insurance: Evidence from the PSID and CEX," Research in Employment Policy 1 (1998), 3-31.

Gruber, Jonathan, *Public Finance and Public Policy*, New York: Worth Publishers, 2005.

Hansen, Gary and Ayse Imrohoroglu, "The Role of Unemployment Insurance in an Economy with Liquidity Constraints and Moral Hazard," Journal of Political Economy 100(1): 118-42, 1992.

Hurst, Erik and Frank Stafford. "Home Is Where the Equity Is: Mortgage Refinancing and Household Consumption," Journal of Money, Credit, and Banking 36(6): 985-1014, 2004.

Johnson, David, Jonathan A. Parker and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," American Economic Review XCVI (December 2006), 1589-1610.

Kodrzycki, Yolanda, "The Effects of Employer-Provided Severance Benefits on Reemployment Outcomes," Federal Reserve Bank of Boston, New England Economic Review 0(0): 41-68, 1998.

Krueger, Alan and Bruce Meyer (2002) "Labor Supply Effects of Social Insurance" Handbook of Public Economics Vol 4. North-Holland: Amsterdam

Lee Hecht Harrison, "Severance and Separation Benefits: Bridges for Employees in Transition," monograph, 2001.

Lentz, Rasmus and and Torben Tranaes, "Job Search and Savings: Wealth Effects and Duration Dependence," Journal of Labor Economics XXIII (2005), 467-489.

McCulloch, Kenneth. *Termination of Employment: Employer and Employee Rights.* Boston: Warren, Gorham and Lamont, 1998.

Meyer, Bruce (1990), "Unemployment Insurance and Unemployment Spells," Econometrica, Volume 58, Issue 4, 757-782.

Moffitt, Robert (1985). "Unemployment insurance and the distribution of unemployment spells," Journal of Econometrics, vol. 28, no. 1, pp. 85–101.

Nyman, John. *The Theory of Demand for Health Insurance.* Stanford: Stanford University Press, 2003.

Pennsylvania Department of Labor and Industry (2007), "Unemployment Compensation Regulations." Available at http://www.dli.state.pa.us/landi/cwp/view.asp?a=185&q=170005

Shimer, Robert and Ivan Werning, "Liquidity and Insurance for the Unemployed," FRB Minneapolis Staff Report #366, December 2006.

Shimer, Robert and Ivan Werning, "Reservation Wages and Unemployment Insurance," forthcoming, Quarterly Journal of Economics (2007).

Sullivan, James. "Borrowing During Unemployment: Unsecured Debt as a Safety Net," forthcoming, Journal of Human Resources (2007).

Topel, Robert, "On Layoffs and Unemployment Insurance," American Economic Review 73 (1983), 541-549.

Zeldes, Stephen P, 1989. "Consumption and Liquidity Constraints: An Empirical Investigation," Journal of Political Economy 97(2), pp 305-46.

Appendix A: Optimal Benefit Level in the General Case

Derivation of (10). Let $c^e_{t,j}$ denote consumption in period j if a job was found in period $t \leq j$ and c^u_t consumption if unemployed in period t. For any variable $x \in \{a, b, w\}$,

$$\frac{\partial s_0}{\partial x} = \frac{1}{\psi''}[\frac{\partial V_0}{\partial x} - \frac{\partial U_0}{\partial x}] \tag{17}$$

Exploiting the envelope condition from agent optimization, we obtain the following derivatives for the value functions:

$$\begin{aligned}
\frac{\partial U_0}{\partial b} &= u'(c^u_0) + \sum_{t=1}^{T-1}\left(\frac{1}{1+\delta}\right)^t \prod_{i=1}^{t}(1-s_i)u'(c^u_t) \\
\frac{\partial V_0}{\partial b} &= 0 \\
\frac{\partial U_0}{\partial w} &= s_1 \sum_{j=1}^{T-1}\left(\frac{1}{1+\delta}\right)^j v'(c^e_{1,j}) + \sum_{t=2}^{T-1}[\prod_{i=2}^{t}(1-s_{i-1})]s_t \sum_{j=t}^{T-1}\left(\frac{1}{1+\delta}\right)^j v'(c^e_{t,j}) \\
\frac{\partial V_0}{\partial w} &= \sum_{j=0}^{T-1}\left(\frac{1}{1+\delta}\right)^j v'(c^e_{0,j}) \\
\frac{\partial U_0}{\partial a} &= \frac{\partial U_0}{\partial b} + \frac{\partial U_0}{\partial w} \\
\frac{\partial V_0}{\partial a} &= \frac{\partial V_0}{\partial w}
\end{aligned}$$

Using these expressions and equation (17), it follows that

$$\frac{\partial s_0}{\partial b} = \frac{\partial s_0}{\partial a} - \frac{\partial s_0}{\partial w}.$$

Derivation of (11) for $\frac{\partial W}{\partial b}$. Note that the unconditional average marginal utility of consumption while employed is

$$\begin{aligned}
\mathrm{E}v'(c^e_{t,j}) &= \frac{1}{T-D}[s_0 \sum_{j=0}^{T-1}\left(\frac{1}{1+\delta}\right)^j v'(c^e_{0,j}) + \sum_{t=1}^{T-1}[\prod_{i=1}^{t}(1-s_{i-1})]s_t \sum_{j=t}^{T-1}\left(\frac{1}{1+\delta}\right)^j v'(c^e_{t,j})] \\
&= \frac{1}{T-D}[(1-s_0)\frac{\partial U_0}{\partial w} + s_0\frac{\partial V_0}{\partial w}]
\end{aligned}$$

The average marginal utility of consumption while employed conditional on finding a job in period 0 is given by

$$\mathrm{E}v'(c^e_{0,j}) = \frac{1}{T}\sum_{j=0}^{T-1}\left(\frac{1}{1+\delta}\right)^j v'(c^e_{0,j}) = \frac{1}{T}\frac{\partial V_0}{\partial w}$$

To derive (11), recall that $J_0 = (1-s_0)U_0(b,\tau) + sV_0(b,\tau) + \psi(s_0)$. Therefore

$$\begin{aligned}
\frac{dJ_0}{db} &= (1-s_0)\{\frac{\partial U_0}{\partial b} - \frac{\partial U_0}{\partial w}\frac{d\tau}{db}\} - s_0\frac{\partial V_0}{\partial w}\frac{d\tau}{db} \\
&= (1-s_0)\frac{\partial U_0}{\partial b} - \frac{d\tau}{db}((1-s_0)\frac{\partial U_0}{\partial w} + s_0\frac{\partial V_0}{\partial w}) \\
&= -(1-s_0)\psi''\frac{\partial s_0}{\partial b} - \frac{d\tau}{db}((1-s_0)\frac{\partial U_0}{\partial w} + s_0\frac{\partial V_0}{\partial w})
\end{aligned} \tag{18}$$

Note that

$$\frac{\partial s_0}{\partial w} = \frac{1}{\psi''}\frac{1}{1-s_0}\{T\mathrm{E}v'(c^e_{0,j}) - T\mathrm{E}v'(c^e_{t,j}) + D\mathrm{E}v'(c^e_{t,j})\} \tag{19}$$
$$= \frac{1}{\psi''}\frac{1}{1-s_0}\mathrm{E}v'(c^e_{t,j})\{D - \rho T\}$$

where $\rho = \frac{\mathrm{E}v'(c^e_{t,j}) - \mathrm{E}v'(c^e_{0,j})}{\mathrm{E}v'(c^e_{t,j})}$ as in the text. Since $\tau = \frac{bD}{T-D}$, it follows that $\frac{\partial \tau}{\partial b} = \frac{D}{T-D} + \frac{bT}{(T-D)^2}\frac{\partial D}{\partial b}$. Using (10) and plugging these expressions into (18) yields

$$\frac{dJ_0}{db} = -(1-s_0)\frac{\partial s_0}{\partial a}\psi'' - \rho T\mathrm{E}v'(c^e_{t,j}) - \frac{bT}{T-D}\frac{\partial D}{\partial b}\mathrm{E}v'(c^e_{t,j}) \tag{20}$$

Next, I normalize the expression for $\frac{dJ_0}{db}$ by the expected welfare gain from increasing the wage by \$1, $\frac{dJ_0}{dw} = (T-D)\mathrm{E}v'(c^e_{t,j})$.

$$\frac{\partial W}{\partial b} = \frac{dJ_0}{db}/\frac{dJ_0}{dw} = \frac{1}{T-D}\{-(1-s_0)\frac{\partial s_0}{\partial a}\frac{\psi''}{\mathrm{E}v'(c^e_{t,j})} - \rho T - \frac{bT}{T-D}\frac{\partial D}{\partial b}\} \tag{21}$$
$$= \frac{1}{T-D}\{-\{D-\rho T\}\frac{\partial s_0^*}{\partial a}/\frac{\partial s_0^*}{\partial w} - \rho T - \frac{bT}{T-D}\frac{\partial D}{\partial b}\}$$
$$= \frac{D}{T-D}\{-\frac{\partial s_0^*}{\partial a}/\frac{\partial s_0^*}{\partial w} - \rho\frac{T}{D}(1 - \frac{\partial s_0^*}{\partial a}/\frac{\partial s_0^*}{\partial w}) - \frac{T}{T-D}\varepsilon_{D,b}\}$$

Finally, defining $\sigma = \frac{T-D}{T}$, it follows that

$$\frac{\partial W}{\partial b} = \frac{1-\sigma}{\sigma}\{\frac{-\partial s_0^*/\partial a}{\partial s_0^*/\partial w} - \frac{\rho}{1-\sigma}(1 - \frac{\partial s_0^*}{\partial a}/\frac{\partial s_0^*}{\partial w}) - \frac{\varepsilon_{D,b}}{\sigma}\}. \tag{22}$$

Plugging in $\rho = 0$ into this formula, we obtain (11).

Calibration of ρ. Assume $r = \delta$, so that the consumption path is flat after re-employment. Taking a quadratic approximation to the utility function, some algebra yields the bound $\rho \leq 1 - \frac{v'(c^e_{0,j})}{v'(\mathrm{E}c^e_{t,j})}$ where $\mathrm{E}c^e_{t,j}$ is the average level of consumption after re-employment, with strict equality iff $\delta = 0$. Intuitively, ρ depends on the difference in consumption if a new job is found immediately (in period 0) or after a spell length of average duration. An upper bound on $\frac{\mathrm{E}c^e_{t,j}}{c^e_{0,j}}$ can be obtained by recognizing that consumption if unemployed must be lower than consumption if employed ($c^u_t \leq c^e_{0,j}$). If $c^u_t = c^e_{0,j}$, the agent's wealth upon re-employment is lowered by the lost income during the unemployment spell.

At an average unemployment duration of 18.3 weeks with a UI replacement rate of 50%, the loss in wealth is $9.15w$, where w is the weekly wage. Since consumption upon re-employment is proportional to (lifetime) wealth, $\frac{\mathrm{E}c^e_{t,j}}{c^e_{0,j}} \leq \frac{Xw - 9.15w}{Xw}$ where X is the total weeks of work remaining after re-employment. In the SIPP data, the average age of the job losers is 37. Assuming a mean retirement age of 62 implies $X = 25 \times 52 = 1300$. Hence $\frac{\mathrm{E}c^e_{t,j}}{c^e_{0,j}} \leq 0.992$. Translating this value into an estimate of ρ requires an assumption about the curvature of the utility function. Under CRRA utility with coefficient of relative risk aversion γ, $\rho = 1 - (\frac{\mathrm{E}c^e_{t,j}}{c^e_{0,j}})^{\gamma}$. If $\gamma = 2$, a commonly used value in the literature (Chetty 2006b), $\rho \leq 0.015$. With other parameters chosen as in section 5 ($\sigma = 0.946$, $-\frac{\partial s_0^*}{\partial a}/\frac{\partial s_0^*}{\partial w} = 0.6$), it follows that $\frac{\partial W}{\partial b}(\rho = 0) - \frac{\partial W}{\partial b}(\rho = 0.015) \leq 0.02$. Hence, the

$\rho = 0$ approximation used in (11) leads to at most a modest upward bias in $\frac{\partial W}{\partial b}$, one that would not change the conclusion that $\frac{\partial W}{\partial b} \geq 0$ at current benefit rates.

Relationship between $\frac{\partial s_0}{\partial A}$ and $\frac{\partial s_0}{\partial a}$. The effect of a lump-sum grant at $t = 0$ on search effort is:

$$\frac{\partial s_0}{\partial A_0} = \frac{1}{\psi''}[v'(c^e_{0,0}) - u'((c^u_0)]$$

The effect of an increase in the annuity is

$$\frac{\partial s_0}{\partial a} = \frac{1}{\psi''}\frac{D}{1-s_0}[Ev'(c^e_{0,j})\frac{1-\rho-\sigma}{(1-\sigma)(1-\rho)} - Eu'((c^u_t)]$$

where $Eu'((c^u_t) = \frac{\partial U_0}{\partial b}/(\frac{D}{1-s_0})$ is the average marginal utility of consumption while unemployed. Note that $\widehat{D} = \frac{D}{1-s_0}$ is the expected unemployment duration conditional on having a positive duration. As noted in the text, in a model where UI benefits have finite duration, the liquidity effect would correspond to the effect of an annuity with the same duration as the benefits. In this case, $\widehat{D}$ would equal the expected compensated unemployment duration conditional on having a positive duration.

Given $r = \delta$, $v'(c^e_{0,0}) = v'(c^e_{0,j})$ for all j. Agent optimization implies $\rho > 0$ and $c^u_{t+1} \leq c^u_t \Rightarrow u'(c^u_t)$ is rising with t. Using these results, some algebra yields

$$\frac{\partial s_0}{\partial a} \leq \frac{1}{\psi''}\frac{D}{1-s_0}k(\delta)[v'(c^e_{0,0}) - u'((c^u_0)] = \widehat{D}\frac{\partial s_0}{\partial A_0}k(\delta)$$

where $k(\delta) = \sum_{t=0}^{T-1} \frac{1}{(1+\delta)^t}/T$ is a decreasing function of δ which satisfies $k(0) = 1$. When $\delta = 0$, $\frac{\partial s_0}{\partial a} \leq \widehat{D}\frac{\partial s_0}{\partial A_0}$, with strict equality if $v'(c^e_{0,0}) = Ev'(c^e_{t,j})$ and $u'((c^u_0) = Eu'((c^u_t)$. The last two conditions are met if c^u_t and c^e_t do not vary with t.

Appendix B: SIPP Sample and Variable Definitions

The data used in section 3 are from the 1985, 1986, 1987, 1990, 1991, 1992, 1993, and 1996 panels of the Survey of Income and Program Participation (SIPP). The SIPP collected information from a sample of approximately 13,000 households in 1985 that grew over time to over 36,000 households in 1996. Interviews were conducted every four months for a period of two to four years, so the data span the beginning of 1985 to the middle of 2000.

Pooling the eight panels yield a universe of 468,766 individuals from 149,286 households. 99,880 of these individuals experience at least one job separation (as defined below) during the sample period. Further restricting the sample to individuals between the ages of 18 and 65 who have at least three months of work history and have been included in the panel for at least three months leaves 78,168 individuals. Because of a problematic definition of unemployment status in the 1985 to 1987 versions of the SIPP, individuals sometimes report a job separation while also reporting unemployment duration equal to zero. Redefining unemployment status to only include those who report becoming unemployed and also a non-zero unemployment duration leaves 65,135 individuals.

I drop observations from Maine, Vermont, Iowa, North Dakota, South Dakota, Alaska, Idaho, Montana, and Wyoming because the SIPP does not provide unique state identifiers for individuals residing in these small states. This leaves me with a sample of 62,598 individuals and 86,921 unemployment spells. 33,149 of these spells are for women, whom I exclude. I also keep only those individuals who report actively searching for a job, as defined below, to eliminate those who have dropped out of the labor force. This leaves a sample of 16,784 individuals (3.6% of original sample) who experienced a total of 21,796 unemployment spells. Next, I drop temporarily layoffs, since these individuals may not have been actively searching for a new job, leaving 21,107 spells. I then exclude individuals who never received UI benefits, leaving 7,015 spells. Finally, I further limit the sample to individuals who take up benefits within the first month after job loss because it is unclear how UI should affect hazards for individuals who delay takeup. This last step produces a core sample consisting of 4,015 individuals (0.86% of the original sample) and 4,560 unemployment spells, of which 4,337 have asset and mortgage information.

Measurement of unemployment durations. The measurement of unemployment durations in the SIPP differs from conventional measures because it requires the tabulation of responses to questions about employment at the weekly level. In particular, the SIPP reports the employment status of every individual over 15 years old for every week that they are in the sample. Weekly employment status (ES) can take the following values: 1. With a job this week; 2. With a job, absent without pay, no time on layoff this week; 3. With a job, absent without pay, spent time on layoff this week; 4. Looking for a job this week; 5. Without a job, not looking for a job, not on layoff. A job separation is defined as a change in ES from 1 or 2 to 3, 4, or 5. Following Cullen and Gruber (2000), I compute the duration of unemployment by summing the number of consecutive weeks that ES >= 3, starting at the date of job separation and stopping when the individual finds a job that lasts for at least one month (i.e., reports a string of four consecutive ES=1 or ES =2). Individuals are defined as being on temporary layoff if they report ES = 3 at any point in the spell. They are defined as "searching" if they report ES = 4 at any point during their spell.

Prediction of Individual-Level Unemployment Benefits. I estimate a first-stage equation for earnings using OLS on the full sample of individuals who report a job loss at some point during the sample period. I regress nominal log wages in the year before job loss on years of education, age at job loss, years of tenure on the last job, a dummy for left-censoring of this job tenure variable, industry, occupation, month, and year dummies, and the unemployment rate in the relevant state/year. Since many individuals in the sample do not have a full year's earning's history before a job separation, I define the annual income of these individuals by assuming that they earned the average wage they report before they began participating in the SIPP. For example, individuals with one quarter of wage history are assumed to have an annual income of four times that quarter's income. Using the coefficient estimates, I predict log wages for each job loser, and recover the predicted wage in levels. I then use this predicted wage to simulate the claimant's unemployment benefit using the UI benefit calculator under the assumption that wages are constant over the "base period" (typically the five quarters before job loss).

Note that simulating benefits using individuals' actual reported wage histories rather than the predicted wages yields a distribution of unemployment benefit levels that has much higher variance and a much weaker correlation with durations in the full sample. The predicted wage measure smooths the reported income fluctuations by isolating permanent differences in income correlated with stable characteristics such as education. Reassuringly, when the sample is restricted to observations in which the deviation between the predicted wage and actual reported wage is small (e.g. <25%), the point estimates for the specifications in Table 2 obtained are similar to those obtained in the full sample.

Appendix C: Mathematica Sample and Variable Definitions

The data for the analysis in section 5 from two surveys conducted by Mathematica on behalf of the Department of Labor, matched with administrative data from state UI records. The datasets are publicly available through the Upjohn Institute. The first dataset is the "Pennsylvania Reemployment Bonus Demonstration," which contains information on 5,678 durations for a representative sample of job losers in Pennsylvania in 1991. This dataset contains information on prior wages, weeks of UI paid, as well as demographic characteristics, household income, job characteristics (tenure, occupation, industry), and receipt of severance pay. The second dataset is the "Study of Unemployment Insurance Exhaustees," which contains data on the unemployment durations of 3,907 individuals who claimed UI benefits in 1998. This dataset is a sample of unemployment durations in 25 states of the United States, with oversampling of individuals who exhausted UI benefits. The information in the dataset is similar to that in the Pennsylvania study. Combining the two datasets yields a pooled sample contains 9,585 individuals. Note that Pennsylvania is not included in the Exhaustees study, and hence there is only one year of data for each state in the sample.

For comparability, I make the same exclusions as in the SIPP. First, I include only prime-age males, dropping 44.7% of original sample. Second, I exclude temporary layoffs by discarding all individuals who expected a recall at the time of layoff, dropping an additional 24.8% of the original sample. Finally, I drop all individuals with missing data either on severance payments, years of job tenure, reported survey durations, or the variables used to predict net liquid wealth, losing another 5% of the original sample. These exclusions leave 2,441 individuals in the sample.

Consistent with the SIPP definition, I measure unemployment durations as the number of weeks elapsed from the end of the individual's prior job to the start of his next job as reported on the Mathematica survey.

Calculation of Mean Severance Amount. Lee Hecht Harrison (2001) report results from a survey of severance pay policies of human resource executives at 925 corporations in the U.S. in 2001. At companies where the severance amount is based on years of service (which is the most common practice), the report provides a tabulation of the number of weeks of severance provided per week of service (see Table 1 in section II, page 4). Using the percentages reported in this table, I compute the mean severance amount for "exempts" (salaried workers) and "non-exempts" (hourly workers), coding the < 1 week category as 0.5 weeks. I then take an unweighted average of the mean severance amount for these two categories to arrive at a mean of 1.49 weeks of severance pay per year of service. Finally, multiplying by the mean job tenure reported in Table 3 implies that the average individual in the Mathematica sample would receive $4.5 \times 1.49 = 6.7$ weeks of severance pay if he were eligible for it.

TABLE 1

Summary Statistics by Wealth Quartile for SIPP Sample

		Net Liquid Wealth Quartile			
		1	2	3	4
	Pooled	(< -$1,115)	(-$1,115-$128)	($128-$13,430)	(>$13,430)
Prior to or at job loss:					
Mean Annual Wage	$20,711	$19,638	$15,971	$20,950	$26,726
Median Annual Wage	$17,780	$17,188	$14,346	$18,584	$23,866
Age	37.0	35.5	35.2	36.7	41.7
Years of Education	12.1	12.2	11.2	12.2	13.1
Percent Married	61%	64%	59%	60%	63%
Percent Spouse Working	37%	40%	28%	40%	44%
Post-layoff:					
Weekly indiv. unemp. benefits	$166	$163	$152	$167	$184
Indiv. replacement rate	49%	50%	50%	49%	47%
Mean unemp. duration (weeks)	18.3	18.0	19.1	17.6	19.4
Median unemp. duration	15.0	15.0	17.0	14.0	17.0
Assets and Liabilities:					
Mean Liq. Wealth	$22,701	$1,536	$502	$5,898	$87,912
Median Liq. Wealth	$1,763	$466	$0	$4,273	$53,009
Mean Unsecured Debt	$3,964	$10,008	$697	$1,752	$3,171
Median Unsecured Debt	$960	$5,659	$0	$353	$835
Mean Home Equity	$31,053	$19,768	$12,866	$30,441	$62,663
Median Home Equity	$8,143	$2,510	$0	$11,794	$48,900
Percent with Mortgage	45%	46%	27%	49%	59%
Percent Renters	39%	43%	61%	35%	16%

NOTE--Table entries are means unless otherwise noted. Data source is 1985-87, 1990-93, and 1996 SIPP panels. Sample includes prime-age males who (a) report searching for a job, (b) are not on temporary layoff, (c) take up UI benefits within one month of layoff, and (d) have at least 3 months of work history in the dataset. Pooled sample size is 4,560 observations. See Appendix B for further details on construction of sample. Indiv. unemp. benefit is simulated individual-level benefit based on two stage procedure described in text. Replacement rate is individual benefit divided by weekly pre-unemployment predicted wage. Unemployment duration is defined as time elapsed from end of last job to start of next job. Asset and liability data is collected once per panel, prior to job loss for approximately half the sample and after job loss for the remainder. Liquid wealth is defined as total wealth minus home, business, and vehicle equity. Net liquid wealth is liquid wealth minus unsecured debt. All monetary values are in real 1990 dollars.

TABLE 2

Effect of UI Benefits: Cox Hazard Model Estimates

	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled	Stratified	Stratified with Full Controls			Mortgage
	Full cntrls	No cntrls	Avg WBA	Max WBA	Ind. WBA	Full cntrls
log UI ben	-0.527					
	(0.267)					
Q1 x log UI ben		-0.721	-0.978	-0.727	-0.642	
		(0.304)	(0.398)	(0.302)	(0.241)	
Q2 x log UI ben		-0.699	-0.725	-0.388	-0.765	
		(0.484)	(0.420)	(0.303)	(0.219)	
Q3 x log UI ben		-0.368	-0.476	-0.091	-0.561	
		(0.309)	(0.358)	(0.370)	(0.156)	
Q4 x log UI ben		0.234	0.103	0.304	0.016	
		(0.369)	(0.470)	(0.339)	(0.259)	
mortg x log UI ben						-1.181
						(0.491)
no mortg x log UI ben						0.079
						(0.477)
log UI ben x spell wk.	x					
log UI ben x spell wk. int. with netliq or mortg		x	x	x	x	x
state, year, ind., occ. fixed effects	x		x	x	x	x
wage spline	x					
ind, occ interactions with netliq or mortg			x	x	x	x
wage spline interaction			x	x	x	x
Q1=Q4 p-val		0.039	0.013	0.001	0.090	
Q1+Q2=Q3+Q4 p-val		0.012	0.008	0.002	0.062	
mortg = no mortg p-val						0.005
Number of Spells	4529	4337	4054	4054	4054	2052

NOTE-Coefficients reported are elasticities of hazard rate w.r.t. UI bens. Standard errors clustered by state in parentheses. See note to Table 1 for sample definition. Sample in column 6 includes those with pre-unemp. mortgage data. Bottom rows of table report p-values from F-test for equality of reported coefficients across quartiles or mortgage groups. Columns 2-6 are cox models stratified by net liquid wealth quartile or mortgage status. All specifications except 2 include following additional controls: age, years of education, marital status, log total wealth, and a dummy for being on seam between interviews to adjust for "seam effect." Columns 3-6 include in addition interactions of occupation+industry dummies and wage spline with stratification variable (netliq quartile or mortgage). All columns include time-varying interaction between log UI benefit and weeks elapsed of spell; in columns 2-6, this time-varying effect is interacted with the stratification variable. In columns 1-3 and 6, UI ben is defined as average UI benefit in claimant's state/year pair.

TABLE 3

Summary Statistics by Severance Receipt for Mathematica Sample

	Pooled	No Severance	Severance
		(0.81)	(0.19)
Prior to or at job loss:			
Mean annual wage	$28,149	$26,213	$37,174
Median annual wage	$20,848	$19,347	$30,693
Age	36.2	35.2	40.6
Percent dropouts	14%	15%	6%
Percent college grads	17%	13%	34%
Percent married	58%	56%	68%
Mean job tenure (years)	4.5	3.8	8.1
Median job tenure (years)	1.9	1.5	4.8
Post-layoff:			
Weekly unemployment benefits	$198	$190	$236
Replacement Rate	49%	51%	43%
Mean unemployment duration	24.3	24.0	25.6
Median unemployment duration	20.0	20.0	22.0
Mean compensated duration	15.8	15.3	18.2
Median compensated duration	16.0	16.0	20.0

NOTE--Table entries are means unless otherwise noted. Data are from surveys of job losers conducted by Mathematica. Sample includes prime-age male UI claimants who are not on temporary layoff. Pooled sample size is 2,441 observations. See Appendix C for details. Data is reweighted using sampling probabilities to yield estimates for a representative sample of job losers. Pre-unemp job tenure is number of years spent working at firm from which worker was laid off. Weekly unemployment benefit is actual individual benefit based on UI records. Replacement rate is weekly benefit times 52 divided by annual wage. Unemployment duration is time elapsed from end of last job to start of next job. Compensated duration is weeks of UI collected. All monetary values are in real 1990 dollars.

TABLE 4

Effect of Severance Pay: Cox Hazard Model Estimates

	Pooled		By Net Liquid Wealth	
	(1) Tenure Control	(2) Full Controls	(3) Tenure Control	(4) Full Controls
Severance Pay Dummy	-0.179 (0.050)	-0.233 (0.071)		
(Netliq < Median) x Sev Pay			-0.493 (0.086)	-0.457 (0.099)
(Netliq > Median) x Sev Pay			0.030 (0.058)	-0.088 (0.081)
tenure spline		x		x
state, ind., occ. fixed effects		x		x
wage spline		x		x
ind., occ., and wage-spline interactions with netliq dummy				x
Equality of coeffs p-val			<0.001	<0.001
Number of spells	2441	2428	2441	2428

NOTE-Coefficients reported can be interpreted as percentage change in hazard rate associated with receipt of severance pay. Standard errors clustered by state in parentheses. See note to Table 3 for sample definition. Bottom row of specs 3 and 4 reports p-values from an F-test for equality of coefficients across low and high-asset groups. All specifications include a sevpay x spell week interaction variable to capture time-varying effects of severance pay. Columns 1 and 3 include only a linear control for tenure at pre-job loss employer in addition. Columns 2 and 4 include the following controls in addition to those listed in the table: age, marital status, dummies for dropout and college graduate, and log individual weekly UI benefit. Baseline hazards in specs 3-4 are stratified by Netliq < Median. Netliq < Median is an indicator variable for whether the household's predicted assets are below the sample median. Netliq > Median is defined analogously. Assets are predicted using the SIPP data as described in text.

Figure 1a

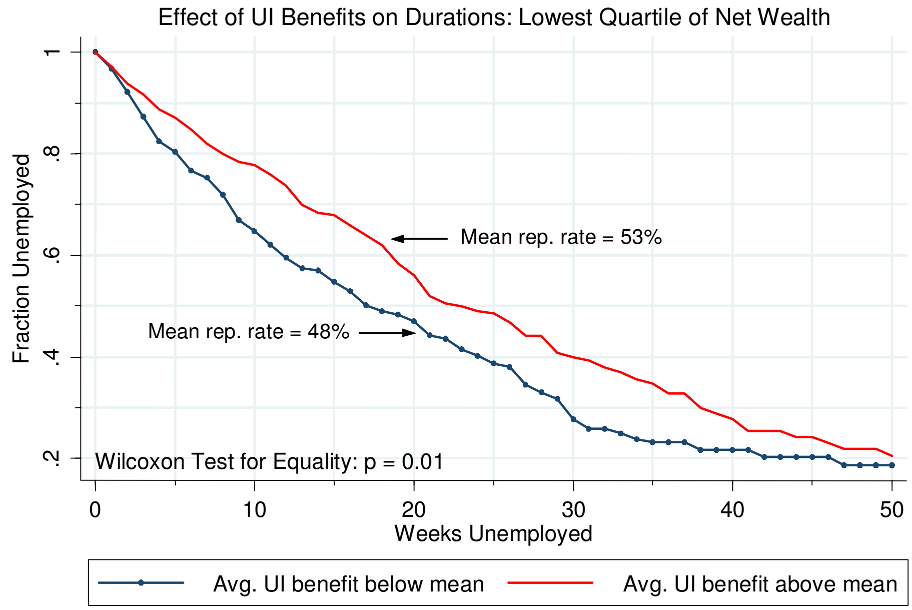

Figure 1b

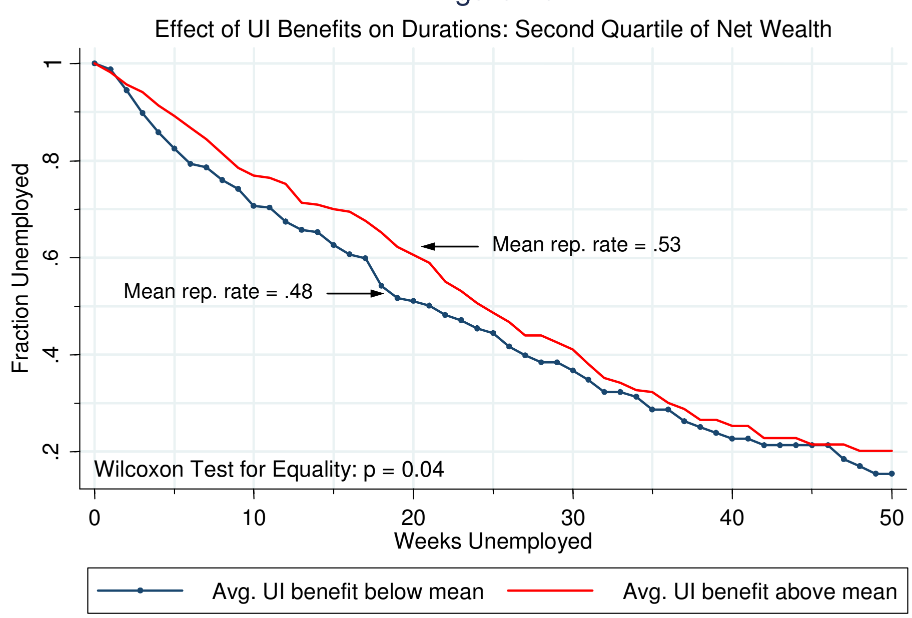

NOTE–Sample for both figures consists of observations in the core SIPP sample for which pre-unemployment wealth data are available. See Table 1 for definition of core sample and definition of net liquid wealth. Figure 1a includes households in lowest quartile of real net liquid wealth. Figure 1b includes those in second quartile. Each figure plots Kaplan-Meier survival curves for two groups of individuals: those in state/year pairs with average weekly benefit amounts (WBA) below the sample mean and those in state/year pairs with WBAs above the mean. The mean replacement rate is the average individual-level predicted benefit divided by wage for observations in the relevant group. Survival curves are adjusted for seam effect by fitting a Cox model with seam dummy and recovering baseline hazards.

Figure 1c

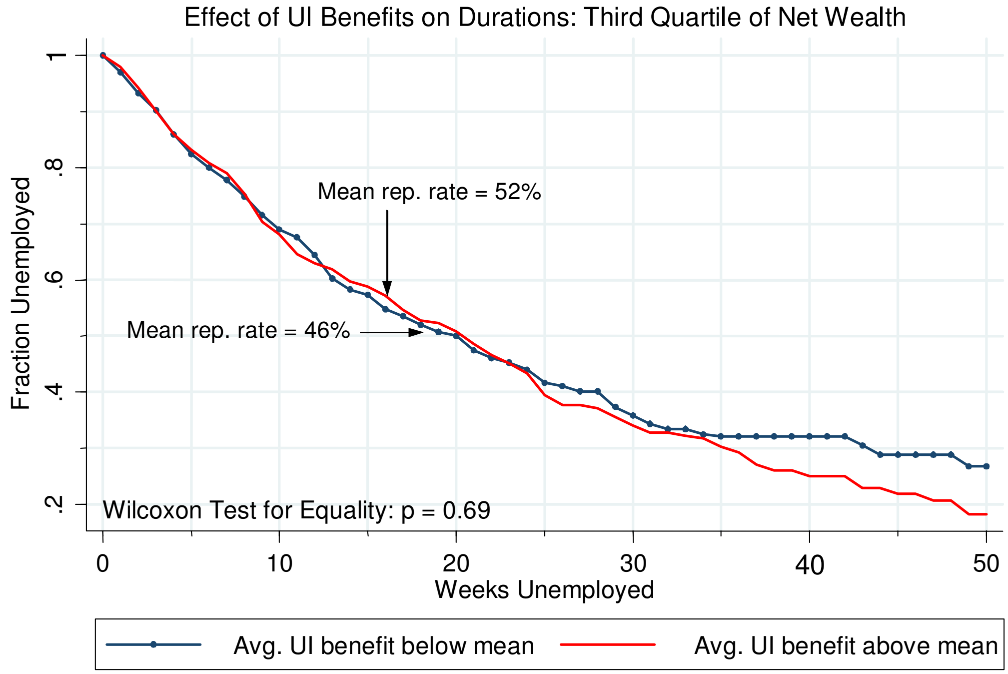

Figure 1d

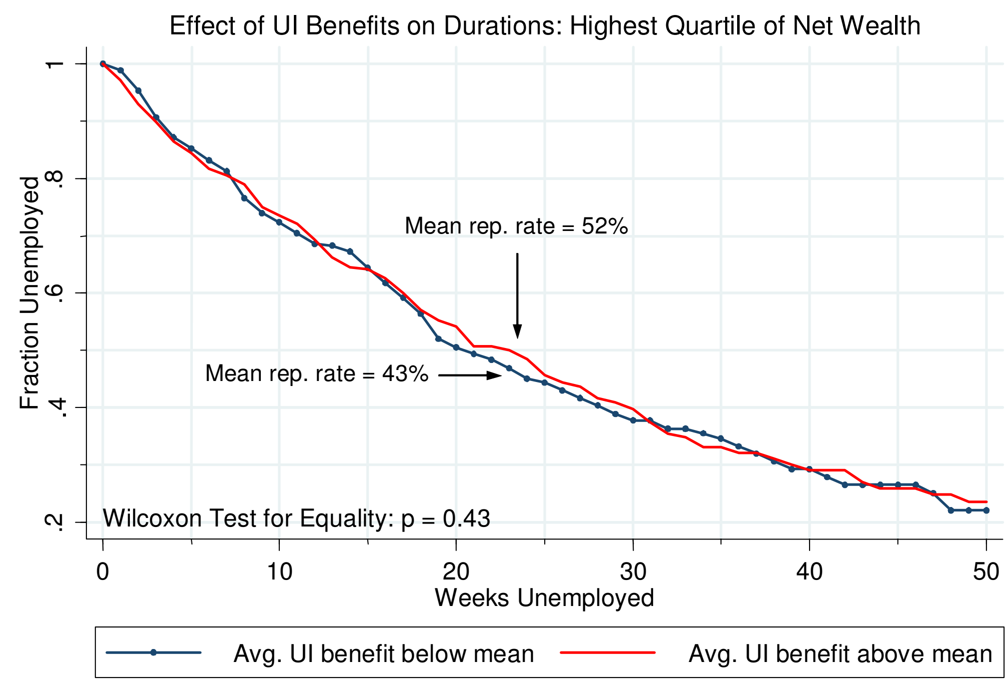

NOTE–These figures are constructed in the same way as Figures 1a-b using observations in the third and fourth quartiles of net wealth. See notes to Figures 1a-b for details.

Figure 2a

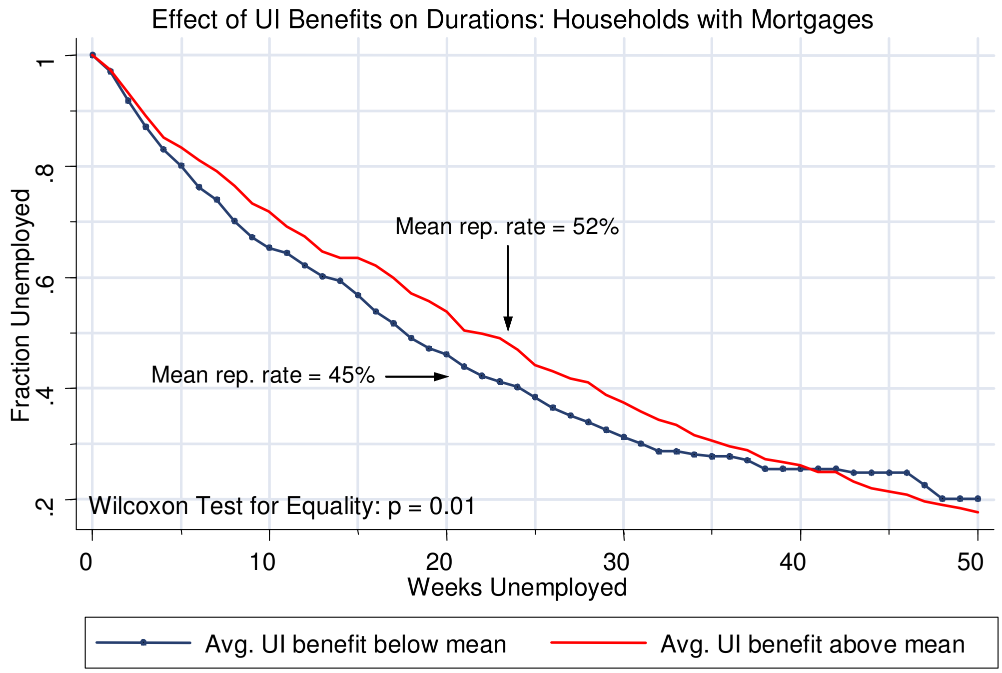

Figure 2b

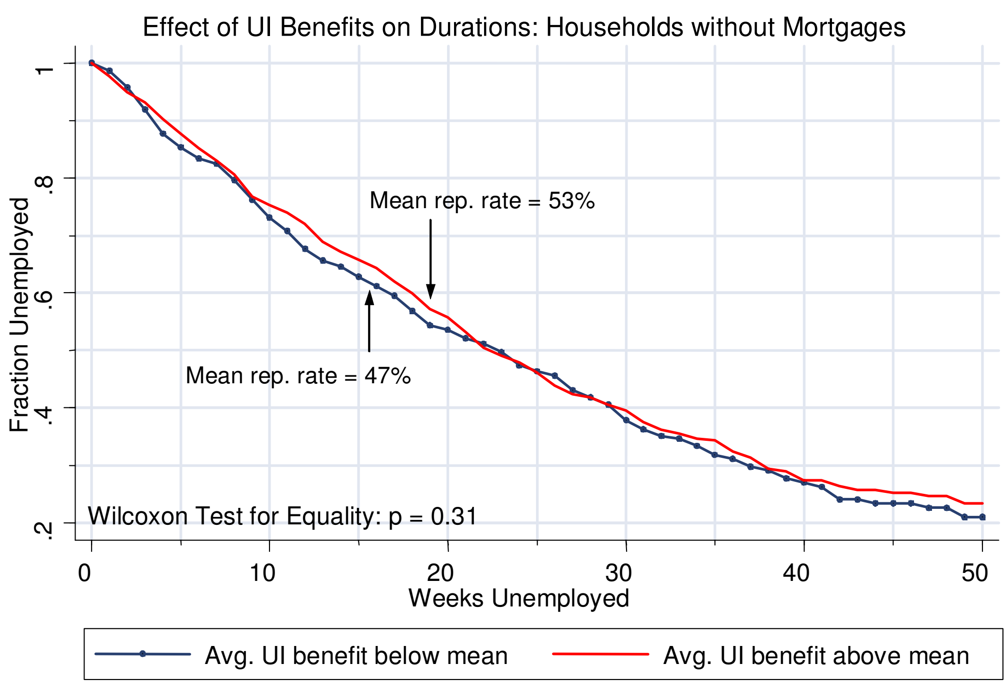

NOTE–These figures are constructed in the same way as Figures 1a-b; see notes to Figures 1a-b for additional details. Figure 2a includes households who make mortgage payments; 2b includes all others. Only observations with mortgage data prior to job loss are included.

Figure 3

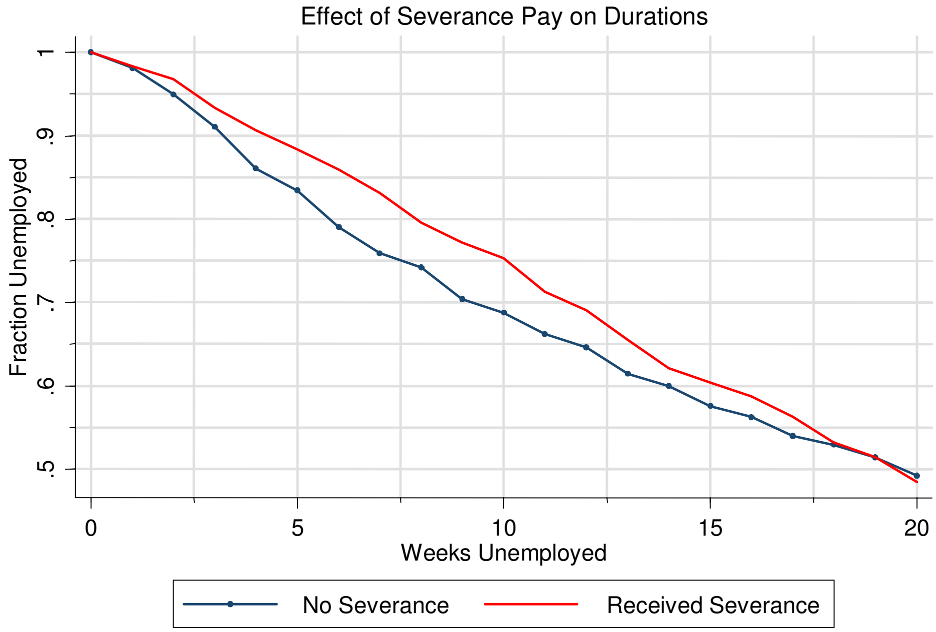

NOTE–Data are from Mathematica surveys; see note to Table 3 for additional details on data and sample definition. Data is reweighted using sampling probabilities to yield estimates for a representative sample of job losers. Kaplan-Meier survival curves are plotted for two groups of individuals: Those who received a severance payment at the time of job loss and those who did not. Survival curves are adjusted for the effect of pre-unemployment job tenure on durations by fitting a Cox model and recovering baseline hazards as described in text.

Figure 4a

Effect of Severance Pay on Durations: Below Median Net Wealth

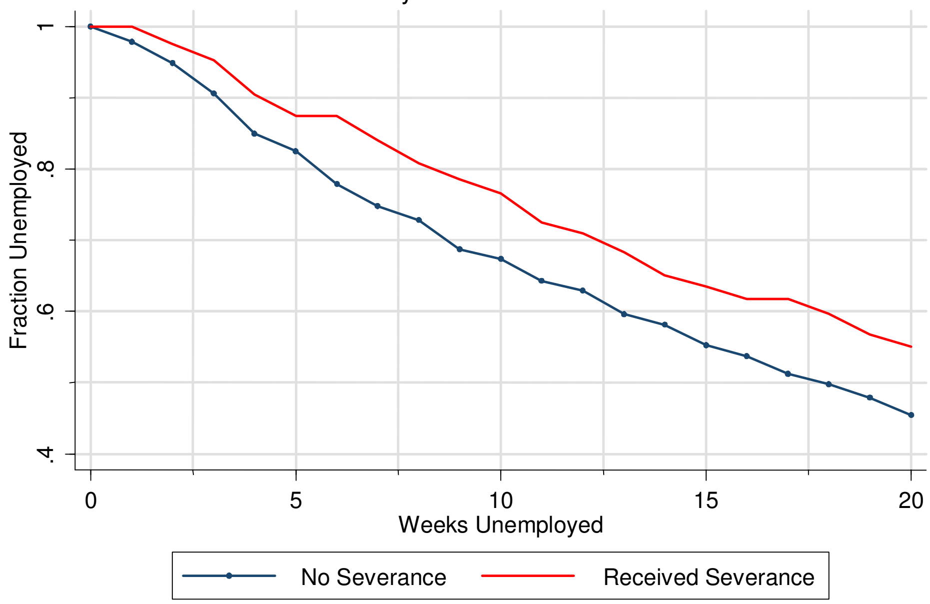

Figure 4b

Effect of Severance Pay on Durations: Above Median Net Wealth

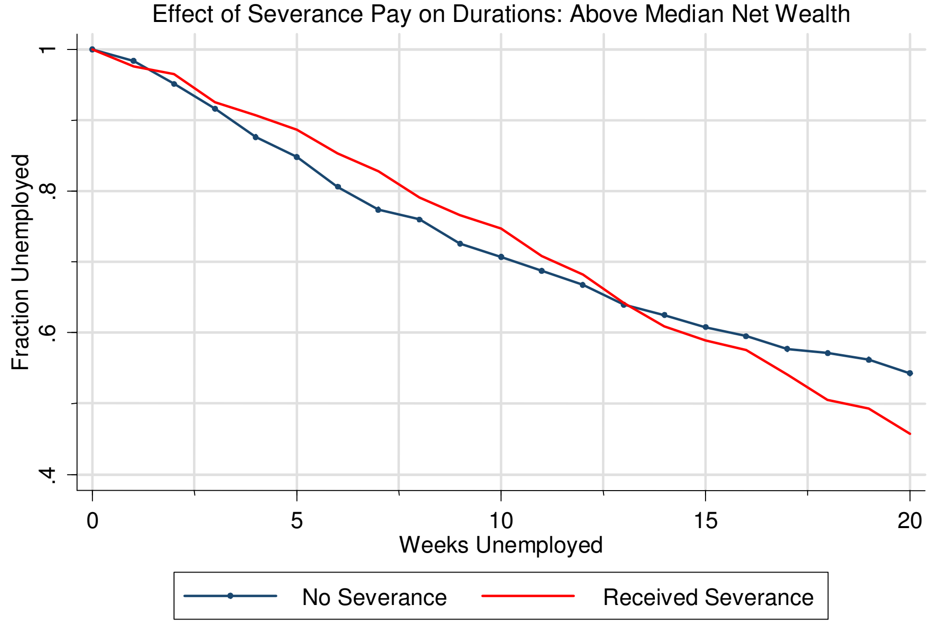

NOTE–See Figure 3 for sample definition. Each of these figures is constructed in exactly the same way as Figure 3. Figure 4a includes observations where predicted net wealth is below the sample median; Figure 4b includes those above the median. Net wealth is predicted using a linear function of age, wage, education, and marital status that is estimated on the core SIPP sample as described in text.