Moral Hazard vs. Liquidity in Unemployment Insurance

Raj Chetty*

UC-Berkeley and NBER

June 2007

Abstract

It is well known that unemployment benefits raise unemployment durations. This result has been interpreted as a moral hazard response to distorted incentives. 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 in documented in two ways. First, state-level increases in unemployment benefits have large effects on durations for liquidity constrained (e.g. low-asset) households, but small effects for unconstrained households. Second, lump-sum severance payments significantly increase durations among constrained households. These results suggest that unemployment benefits affect search behavior primarily through the provision of liquidity rather than moral hazard. Using a job search model, I derive a new test for the optimal level of unemployment benefits in terms of the magnitude of the liquidity effect relative to moral hazard. Implementing this test using the empirical estimates implies that the optimal wage replacement rate for UI is above 50%, considerably higher than the findings of some earlier studies.

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

*E-mail: chetty@econ.berkeley.edu. I have benefited from discussions with Alan Auerbach, David Card, Martin Feldstein, Amy Finkelstein, Jon Gruber, Jerry Hausman, Caroline Hoxby, Larry Katz, Bruce Meyer, Emmanuel Saez, Robert Shimer, Adam Szeidl, Ivan Werning, anonymous referees, and numerous seminar participants. Philippe Bouzaglou, David Lee, and Jim Sly provided excellent research assistance. Funding from the National Science Foundation and NBER is gratefully acknowledged.

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 arises purely from 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 standard moral hazard channel. Intuitively, UI benefits increase cash-on-hand and consumption while unemployed for an agent who cannot smooth perfectly. The agent therefore faces less pressure to find a new job quickly, leading to a longer duration. This liquidity effect disappears if the agent is unconstrained and can smooth consumption perfectly, because marginal utilities in the employed and unemployed states are equal in this case, making search behavior invariant to cash-on-hand. Hence, unemployment benefits raise durations purely through moral hazard for unconstrained individuals, but through both liquidity and moral hazard 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 derive a new test for the optimal unemployment benefit level – similar in spirit to the recent tests proposed by Shimer and Werning (2007) and Chetty (2006) – based on the ratio of the liquidity and moral hazard effects. The test 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.

Motivated by this theoretical analysis, 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 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 being constrained using three measures: asset holdings, single vs. dual-earner status, and an indicator for having to make a mortgage payment. I find clear, robust evidence that a 10% increase in UI benefits raises unemployment durations by 6-8% in all the constrained groups. In contrast, changes in UI benefits have much smaller, statistically insignificant effects on durations in the unconstrained groups – indicating that the moral hazard effect is small among these groups. Hence, the link between unemployment benefits and durations documented in earlier studies appears to be driven by a subset of the population that has limited ability to smooth consumption. This suggests that liquidity effects could be quite important in the benefits-duration link, but does not directly establish that benefits raise the durations of constrained agents by increasing cash-on-hand (unless one assumes that the substitution elasticities are similar across constrained and unconstrained groups).

This observation motivates the second portion of the empirical analysis, in which I decompose the benefit elasticity of unemployment durations into a liquidity and moral hazard effect. To do so, I exploit variation in cash-on-hand created by differential receipt of lump-sum severance payments across job losers. Using a new dataset that matches survey data from Mathematica on severance pay with administrative records on unemployment durations, I find that individuals who received a lump-sum severance payment at the time of job loss (worth about $2000 on average) have substantially lower unemployment exit hazards. 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 virtually unchanged with the inclusion of a large set of 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, point toward a substantial role for liquidity effects in unemployment insurance.

Using the point estimates from the two empirical approaches, a simple calculation indicates that 70% of the UI-duration link is due to the liquidity effect. I use this estimate to implement the test for optimal benefits derived from the search model, and find that the test for optimality is approximately satisfied at current benefit rates (which are approximately 50% of pre-unemployment wages). This calculation should be viewed as a rough first pass for several reasons. First, the analysis is based on a simple model that ignores factors such as temporary layoffs and general equilibrium effects (Feldstein 1978; Acemoglu and Shimer 1999). Second, the calculation assumes a constant benefit schedule, whereas benefits are in fact time-varying in the U.S. Finally, the empirical estimates on which the test is based should be interpreted cautiously given the lack of randomized variation in cash grants.

In addition to its implications for optimal benefits, this paper relates to some other literatures in macroeconomics and public finance. First, a large literature has used consumption data to investigate the extent to which households can smooth consumption (see e.g., Zeldes 1989; Johnson, Parker, and Souleles 2006). Analogously, the empirical analysis here shows that labor supply behavior is "excessively sensitive" to exogenous variation in disposable income because of limited liquidity. Second, several studies have explored the effects of incomplete insurance and credit markets for job search behavior and unemployment insurance using simulations of calibrated search models (Hansen and Imrohoglu 1992; Acemoglu and Shimer 2002). The analysis here can be viewed as the empirical counterpart of such studies, in which the extent to which agents can smooth shocks is evaluated empirically rather than simulated from a calibrated model. The estimated liquidity effects 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 a range of other social insurance programs. The "revealed preference" approach to calculating the value of insurance proposed here can be easily adapted to these 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 job match quality or health. While this approach relies on the assumption that agents' choices reveal their true preferences, it is a useful first step in assessing the set of policies that are optimal from a libertarian perspective.

The remainder of the paper proceeds as follows. The next section presents the search model, characterizes the moral hazard and liquidity effects, and derives the test for optimal benefits. Section 3 discusses the evidence on heterogeneous effects of UI benefits by consumption-smoothing capacity. 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). These studies analyze models with variable search intensity and savings decisions. I use this framework here to characterize 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 job search behavior given a UI benefit system, and then turn to the government's problem of choosing the optimal benefit level taking into account the agent's behavioral responses.

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 subjective time discount rate and r denote the interest rate. Flow utility in period t is given by $u(c_t) - \psi(s_t)$, where c_t represents consumption in the period, s_t denotes search effort, and the functions u and ψ are strictly concave and convex, respectively. Normalize s_t to equal the probability of finding a job in the current period.

Suppose the agent becomes unemployed at $t = 0$. If the worker is unemployed in period t, he receives an unemployment benefit b_t. If the worker is employed in period t, he pays a tax τ that is used to finance the unemployment benefit. An agent who enters a period t without a job first chooses search intensity s_t, and immediately learns if he has obtained a job. If search is successful, the agent begins working in that period at a fixed pre-tax wage w_t. Assume that all jobs last indefinitely once found, and that the wage schedule $\{w_t\}$ is exogenously fixed, eliminating reservation-wage choices. Let c_t^e denote the agent's optimal consumption choice 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.

Optimal Search Intensity. 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

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

where L is a lower bound on assets that may or may not be binding. The value function for an 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 expected 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 may in general 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 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 the marginal value of search effort, which is given by the difference between the optimized values of employment and unemployment. The analysis below follows from the comparative statics of (4).

2.2 Moral Hazard and Liquidity Effects

To understand how changes in the unemployment benefit level affect search intensity, first 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 on search intensity:

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

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

The effect of a cash grant on search intensity depends on the difference in marginal utilities between employed and unemployed states. In a model with perfect consumption smoothing ($c_t^u = c_t^e$), $\partial s_t^*/\partial A_t = 0$, because a cash grant raises $V_t(A_t)$ and $U_t(A_t)$ by the same amount. More generally, if c_t^u is close to c_t^e, as in a permanent income model with unrestricted borrowing or an environment with good insurance markets, the asset effect will be small. In contrast, if individuals face asset constraints, have incomplete insurance contracts, or voluntarily reduce c_t^u to maintain a buffer stock of savings, $\partial s_t^*/\partial A_t$ can be large.

The effect of an increase in the period t wage is proportional to the marginal utility of consumption while employed. Intuitively, a higher wage increases the marginal return to search to the extent that it raises the value of being employed, and thereby leads to higher search intensity through a substitution effect.

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 thereby making it desirable to take more time to find a new job. 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.

As discussed in the introduction, the existing literature has interpreted the effect of benefits on durations as being driven purely by the distortion in marginal incentives (moral hazard). The preceding analysis shows that this interpretation is valid for agents who are able to smooth consumption fully. However, for agents who cut consumption while unemployed, the liquidity effect could potentially be quite large. Hence, the fraction of the UI-duration link due to moral hazard in the population of job losers is ultimately an empirical question.

2.3 A New Test for the Optimal Benefit Level

The relative magnitudes of the moral hazard and liquidity effects are of interest because of their normative implications. To illustrate this, I derive a formula for the welfare gain of raising the UI benefit level in terms of these effects, assuming that benefits are constant over time. For expositional purposes, I begin with the case where $T = 1$, i.e. a static search model where the agents probability of employment is fully determined by s_0^*.

Static Case. When $T = 1$, the social planner's problem reduces to maximizing the agent's expected utility in period 0 subject to the balanced-budget constraint for the UI system:

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

The welfare gain from increasing b by \$1 is

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

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

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

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

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

Using equations (5) and (6), this expression can be rewritten as

$$\begin{aligned}\frac{\partial W}{\partial b_0} &= \frac{1 - s_0^*}{s_0^*}\{\frac{-\partial s_0^*/\partial A_0}{\partial s_0^*/\partial w_0} - \frac{\varepsilon_{1-s,b}}{s_0^*}\} \\ &= \frac{1 - s_0^*}{s_0^*}\{\frac{\partial s_0^*/\partial A_0}{\partial s_0^*/\partial b_0 - \partial s_0^*/\partial A_0} - \frac{\varepsilon_{1-s,b}}{s_0^*}\} \qquad (8)\end{aligned}$$

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 now show that this result holds in the general case under some approximations.

General Case. When $T > 1$, the policy space is much larger, because one can in principle choose a different benefit level b_t in each period. To simplify, I restrict attention to the optimal policy among the set of policies that have constant benefits: $b_t = b \ \forall t$ and also assume that $w_t = w \forall t$.[1]

There are two conceptual complications in deriving an empirically implementable expression for $\frac{\partial W}{\partial b}$ in terms of the liquidity and moral hazard effects when $T > 1$. First, an increase in b affects income in subsequent periods, and not just the first period. First, the effect of a cash grant in the first period cannot be directly compared with the effect of a benefit increase, since the timing of receipt of income matters when agents cannot smooth consumption. To address this issue, I characterize the liquidity effect of a benefit increase by comparing its effect to that of an annuity that pays \$$a$ in every period $t = 0, ..., T-1$. While this approach yields an analytically tractable formula, implementing it empirically requires an assumption about the cash-value of an annuity since the only variation in the data is in cash grants at $t = 0$. I revisit this issue when implementing the formula in section 5.

The second complication is that the UI tax τ also affects income in multiple periods. The effect of the tax thus depends on the marginal utility of consumption while employed in different periods. Relating this weighted average of marginal utilities to the effect of a change in wages on search intensity (the substitution effect), requires that the average marginal utility of consumption while employed $\mathrm{E}u'(c_t^e)$ can be approximated by the marginal utility of average consumption while employed $u'(\mathrm{E}c_t^e)$. The assumption that $\mathrm{E}u'(c_t^e) = u'(\mathrm{E}c_t^e)$ essentially requires that the third-order terms of the utility function are small given the variability of consumption while employed.[2]

Under this approximation, I show in the appendix that

$$\frac{\partial s_0}{\partial b} = \frac{\partial s_0}{\partial a} - \frac{\partial s_0}{\partial w}. \tag{9}$$

Equation (9) is the analog of (7). It shows that the moral hazard/liquidity decomposition of the UI benefit applies when the benefit level is simultaneously changed in all periods rather than a single period t. The liquidity effect is now determined by the effect of an increase in the annuity payment

[1] Shimer and Werning (2007) analyze the optimal path of benefits, and argue that constant benefits are likely to be optimal under certain conditions. In future work, it would be interesting to derive formulas for the optimal path or duration of benefits in terms of moral hazard and liquidity effects at different points of the spell.

[2] See Chetty (2006) for discussion of this condition and a simple calibration based on empirical estimates of the variance of c_t^e which suggests that this is likely to be a reasonable approximation for commonly used utilities.

a in all periods, while the moral hazard effect is determined by a wage increase in all periods.

Using these definitions of moral hazard and liquidity effects, one can obtain an expression for the welfare gain of raising UI benefits. In the general case, the government chooses b to maximize the agent's expected utility at time 0, $J_0(b, w, a)$. 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: $\frac{\partial W}{\partial b} \equiv \frac{\partial J_0}{\partial b} / \frac{\partial J_0}{\partial w}$. I show in the appendix that

$$\frac{\partial W}{\partial b} = \frac{1-\sigma}{\sigma}\{\frac{\partial s_0^*/\partial a}{\partial s_0^*/\partial b - \partial s_0^*/\partial a} - \frac{\varepsilon_{D,b}}{\sigma}\} \tag{10}$$

where D denotes the agents expected unemployment duration and $\sigma = \frac{T-D}{T}$ denotes the fraction of his life that the agent spends employed. This expression parallels (8) in the static case, except that the parameters reflect derivatives with respect to permanent changes in annuity income and UI benefits rather than changes in period 1 only.

Discussion. One can test whether a given level of UI benefits b' is optimal by estimating the moral hazard and liquidity effects of a benefit change and checking whether $\frac{\partial W}{\partial b}(b') = 0$. This test shows that a higher benefit elasticity does not necessarily imply a lower optimal benefit level, as is commonly thought. It matters whether this higher elasticity is coming from liquidity or moral hazard. 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.

The intuition for this test can be interpreted in several ways. First, the test provides an alternative method of estimating the extent to which the full insurance benchmark is violated. Existing studies estimate the effects of income shocks on consumption of goods to evaluate how far agents are from the full-insurance benchmark, and use this information to calculate the value of UI (e.g. Cochrane 1992, Gruber 1997). Here, I instead estimate the effect of income shocks on labor supply, i.e. consumption of leisure. The sensitivity of search to cash-on-hand ($\partial s_0^*/\partial a$) is determined by the difference between consumption when employed and unemployed. Hence, the effect of liquidity on search behavior provides an indirect means of assessing the extent to which agents can smooth consumption, and thus yields information on the value of insurance.

An alternative interpretation of the test comes from the literature on taxation. The substitution effect induced by a tax leads to a deadweight cost, whereas a response to cash transfers (the income effect) does not. The result in (10) parallels this intuition: the substitution effect of UI leads to an efficiency cost, and hence a lower optimal benefit level.

More abstractly, the idea underlying this test is to use revealed preference to measure the value of insurance to the agent. For example, suppose an agent extends his unemployment duration by 10% when b is raised by 10%. Now consider the effect of a lump-sum cash grant equivalent to the increase in the UI benefit amount. Suppose that the cash grant has no effect on the agent's duration of search. In this case, we can infer that the agent is "spending" his money on a longer duration only because of the subsidy for doing so. Hence, the provision of UI benefits does not facilitate choices that the agent would otherwise have been unable to make, and simply creates moral hazard. The welfare gain from UI is strictly negative, and the optimal benefit is 0. In contrast, suppose that the cash grant increases duration by 10%, the same as the UI benefit. In this case, we can infer that the agent would ideally "spend" his UI benefit in the same way as a lump sum grant. His choice reveals that the UI benefit permits him to make a more (socially) optimal choice, i.e. a choice when incentives are not distorted and markets are more complete. Here the marginal value of raising b is infinite, since there is no moral hazard effect.[3] More generally, the extent to which the agent raises duration in response to a cash grant reveals the extent to which the insurance benefit permits him to attain a more socially desirable allocation. The value of various private and social insurance programs can be assessed using this revealed preference approach; the present paper can be viewed as an application to unemployment insurance.

The test proposed here offers an alternative to the consumption-based test for optimal benefits of Baily (1978) and Chetty (2006) and the reservation-wage test of Shimer and Werning (2007). The three tests are essentially different "reduced form" representations of the condition that must be satisfied at the optimal benefit level, exploiting different data sources. An advantage of the test based on liquidity and substitution effects is that it requires only one source of data – unemployment durations – which are typically more precise and widely available than data on consumption or reservation wages.

The cost of this parsimonious approach is that is relies quite strongly on 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 test proposed here, this finding would have no effect on the optimal benefit level – it makes no difference if the

[3] It is important to note that the liquidity and moral hazard effects vary with the benefit level. In particular, the liquidity effect $\partial s_0^*/\partial a$ must approache zero as the benefit level approaches the full insurance level. This guarantees that the optimal wage replacement rate is below 100% if there is some non-zero moral hazard effect.

worker rationally chooses to use the money to consume more leisure rather than locate a better match. However, one may believe that workers are myopic or have time-inconsistent preferences that lead them to (suboptimally) consume cash-on-hand immediately through unproductive leisure rather than job search. In such an environment, where the preferences revealed by choice are not those that maximize lifetime utility, data on outcomes are needed to calculate the benefits of insurance.

In view of these issues, this revealed preference test should be viewed as a benchmark reflecting the choice of a libertarian social planner. In particular, the optimal benefit implied by the test is the one that corrects market incompleteness to the extent possible without interfering with individual choice. I turn now to implementing the libertarian test by estimating the moral hazard and liquidity effects of UI empirically.

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 consumption perfectly and have $\Delta c = c_t^e - c_t^u = 0$ – and "constrained" individuals – who have $\Delta c > 0$. If the UI-duration link is much stronger in the unconstrained group, it would be unlikely that liquidity effects are large since the liquidity effect arises only for constrained individuals. In contrast, if the link is driven by the constrained group, liquidity effects might matter.

I implement this heterogeneity analysis by replicating the identification strategy of Moffitt (1985) and Meyer (1990) on subsets of the data. In particular, I divide individuals into unconstrained and constrained groups (using proxies discussed below), and estimate equations of the following form using cross-state and time variation in UI benefit levels:

$$\log d_{it} = \beta_0 + \beta_1 \log b + \beta_2 X_{i,t} + \theta_{i,t} \tag{11}$$

where $X_{i,t}$ denotes a set of covariates and $\theta_{i,t}$ denotes an idiosyncratic error term. When (11) is estimated for unconstrained individuals, the estimated value of β_1 reflects the pure moral hazard effect of UI. When (11) is estimated for constrained individuals, the coefficient β_1 reflects the *sum* of the moral hazard and liquidity effects. Comparing the estimated β coefficients thus gives an indication of the importance of liquidity relative to moral hazard.

The formula in (10) calls for estimates of the effects of UI benefits on search behavior at the beginning of the spell. This issue, as well as censored spells, motivates estimation of a hazard model with time-varying covariates rather than estimation of (11) using OLS. Letting $h_{i,t}$ denote the unemployment exit hazard rate for household i in week t of an unemployment spell and $X_{i,t}$ denote a set of controls, the primary estimating equation used below is therefore

$$h_{i,t} = f(b_i, t \times b_i, X_{i,t}) \tag{12}$$

where the $t \times b_i$ term permits a time-varying effect of UI benefits on unemployment exit rates within the ex-ante constrained and unconstrained groups. This specification also addresses the concern that households may become constrained as an unemployment spell elapses, e.g. as they deplete their buffer stock. If liquidity effects are important, households that are initially unconstrained may exhibit greater sensitivity of search behavior to UI benefits late in the spell, a secondary hypothesis that can be evaluated to some extent by estimating (12).

Proxies for Constraint Status. The main difficulty in implementing the empirical strategy outlined above is in classifying households into unconstrained and constrained groups. As shown by equation (5), 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 several proxies for the ability to smooth consumption that have been identified by studies of consumption behavior. These proxies identify households that are likely to be able to smooth an unemployment shock intertemporally using savings and borrowing. Since unemployment shocks are small relative to lifetime wealth, individuals with adequate savings or borrowing capacity should have $\Delta c \simeq 0$.[4]

The primary proxy I use is liquid wealth net of unsecured debt, which I term "net wealth." Browning and Crossley (2001) and Sullivan (2005) find that households with little or no financial assets prior to job loss suffer consumption drops during unemployment that are mitigated by

[4] An alternative approach is to identify households that can smooth consumption through risk sharing (smoothing across states rather than intertemporally). In future work, it would be interesting to test whether the effect of UI benefits on durations varies with private insurance arrangements (e.g. social networks), as would be expected if liquidity effects are important.

provision of UI benefits. In contrast, households with higher assets do not show excess sensitivity of consumption to unemployment or UI benefit levels. These findings suggest the use of net wealth as a proxy for being unconstrained. With this proxy, the heterogeneity test amounts to asking whether an increase in UI benefits has a greater effect on durations when wealth is low.

Insofar as UI benefits have a liquidity effect by adding to wealth, this analysis can be interpreted as a test of whether the effect of wealth on unemployment duration is increasing and concave. This test could in principle be implemented by exploiting the cross-sectional variation in wealth itself, rather than focusing solely on the variation that arises from UI benefits. I focus on the variation in UI benefits because changes in UI laws are credibly exogenous to individuals' preferences, i.e. $Eb \times \theta_{i,t} = 0$ in (11). 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, individuals with higher assets are also likely to have lower discount rates or higher planned expenses such as college tuition payments, and hence may be reluctant to deplete their assets to finance a longer spell of unemployment. In practice, I find no robust association between assets and unemployment durations in the cross-section (consistent with Lentz (2003) and other studies), suggesting that the use of an exogenous source of variation in cash-on-hand such as UI benefits is quite important.

The second proxy I use for being unconstrained is whether the individual has a spouse who is also working 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 likely to be able to smooth shocks better intertemporally, e.g. because they are more likely to be able to borrow given that at least one person is employed.[5] The third proxy is an indicator variable for whether the individual was making mortgage payments prior to job loss. Gruber (1998) finds that fewer than 5% of the unemployed sell their homes during a spell, whereas renters move much more frequently. Consequently, if an individual must make a mortgage payment, he 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.[6]

[5] A countervailing effect is that households with a single earner may be able to maintain their prior standard of 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 countervailing effect is dominated by the intertemporal smoothing capacity, so that households with two earners are in fact less constrained in practice. I therefore use single-earner status as a proxy for being constrained.

[6] A potential concern with this proxy is that homeowners have more wealth than renters because of home equity. However, Hurst and Stafford (2004) point out that since most job losers have low levels of home equity, they must refinance to access this wealth. Hurst and Stafford find that while unemployment raises the probability of refinancing,

Although these proxies predict being constrained on average, they are imperfect predictors for two reasons. First, some households classified as unconstrained will be misallocated to the constrained group and vice-versa. This classification error pulls the estimated elasticities for the groups closer together, thereby causing us to *underestimate* the liquidity effect. Second, no household truly has $\Delta c = 0$ because insurance markets are likely to be incomplete, and intertemporal smoothing itself cannot fully eliminate consumption fluctuations. There is therefore likely to be a small liquidity effect even among the groups classified as unconstrained in our empirical analysis (e.g. high wealth households). Since I attribute the entire response among the group classified as unconstrained to moral hazard, this bias also leads to underestimation of the liquidity effect relative to moral hazard.

3.2 Data

The data used to estimate (12) are from 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. Other data about the demographic and economic characteristics of each household member are also collected. Further details on the database and sample definition are given in the appendix.

Starting from the universe of job separations in the pooled SIPP panels, I make five exclusions to arrive at the core analysis sample. First, following previous studies of UI, I restrict attention to prime-age males (over 18 and under 65). Second, I include only separations where the individual reported searching for a job at some point after losing their job, in order to eliminate individuals who have dropped out of the labor force. Third, I exclude individuals who report that they were on temporary layoff at any point during their spells, since they might not have been actively searching for a job.[7] Fourth, I exclude individuals who have less than three months work history within the

approximately 2/3 of homeowners who lost their jobs between 1991 and 1996 in the PSID did not refinance their mortgages over that five year period, perhaps because refinancing is a costly and slow process. This suggests that many homeowners with mortgages may be more liquidity constrained than renters, at least in the short run, despite having home equity wealth.

[7]Results are similar if temporary layoffs are included (see Table 2b of Chetty (2006)). This mitigates the concern that the estimated may be biased because of the endogeneity of temporay layoff status (Katz and Meyer 1990).

survey because there is insufficient information to estimate pre-unemployment wages for this group. Finally, I focus on individuals who take up UI within one month after losing their job because it is unclear how UI should affect hazards for individuals who delay takeup.

These exclusions leave 4,560 unemployment spells in the core sample. Asset data are generally collected only once in each panel, so pre-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.

Data on UI laws were obtained from the Employment and Training Administration (various years) and supplemented with information directly from individual states to construct a program that assigns unemployment benefit amounts to each individual in the sample.[8] Unfortunately, measurement error and inadequate information about pre-unemployment wages for many claimants make it difficult to simulate the potential UI benefit level for each agent precisely. I therefore use three independent approaches to proxy for each claimant's (unobserved) actual UI benefits, each of which yields qualitatively similar results. 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. The third method involves simulating each individual's weekly UI benefit using a two-stage procedure. In the first stage, I predict the claimant's pre-unemployment annual income using information on education, age, tenure, occupation, industry, and other demographics. The prediction equation for pre-UI annual earnings is estimated on the full sample of individuals who report a job loss at some point during the sample period.[9] In the second stage, I use the predicted wage as a proxy for the true wage, and calculate each claimant's unemployment benefits using the simulation program.

[8] I am grateful to Julie Cullen and Jon Gruber for sharing their simulation programs, and to Suzanne Simonetta and Loryn Lancaster in the Department of Labor for providing detailed information about state UI laws from 1984-2000.

[9] 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.

3.3 Results

3.3.1 Graphical Evidence and Non-Parametric Tests

I begin by providing graphical evidence on the benefit elasticity of unemployment 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 1a 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. Notably, quartiles 1 and 3 are similar in terms of income and education. Hence, UI benefits are similar both in levels and as a fraction of permanent income for all the groups.

Figures 1a-d show the effect of UI benefits on unemployment exit rates for households in the each of the four quartiles of the net wealth distribution. To construct Figure 1a, I first divide the observations 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. Kaplan-Meier survival curves are then plotted for these two groups using the households in the lowest quartile of the net wealth distribution. This procedure is repeated for the other three quartiles of the net wealth distribution to construct Figures 1b-d. Since ex-post asset levels are endogenous to the length of durations, 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.

These and all subsequent survival curves plotted using the SIPP data are adjusted for the "seam effect" common 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 (the "reference period"). As a result, they under-report transitions in labor force status within reference periods and overreport transitions on the "seam" between reference periods. Consequently, a disproportionately large number of spells appear to last for exactly 4 or 8 months in the data. These artificial spikes in the hazard rate are smoothed out by first fitting a Cox model with a time-varying indicator variable for being on a seam between interviews, and then recovering the (nonparametric) 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 qualitatively similar if the raw data is used without adjusting for the seam effect.

Figure 1a shows that higher UI benefits are associated with much lower unemployment exit rates for individuals in the lowest wealth quartile, who are most likely to be constrained ($\Delta c > 0$). 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 two survival curves are identical with $p < 0.01$. Figure 1b constructs the same survival curves for the second wealth quartile. UI benefits have a smaller, but still powerful 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. The Wilcoxon test again rejects equality of the survival curves in this group, with $p = 0.04$. Figures 1c and 1d show that effect of UI on durations virtually disappears in the third and fourth quartiles of the wealth distribution. Not surprisingly, the Wilcoxon test does not reject equality of the survival curves in these two cases. The fact that UI has little effect on durations in the unconstrained groups suggests that it induces little moral hazard among these households.[10]

I now replicate these graphs and nonparametric tests for the other two proxies of constraints. Table 1b shows summary statistics for the constrained and unconstrained groups based on spousal work and mortgage status. As with the asset cuts, there are differences across the constrained and unconstrained groups in income and education, but these are not extremely large. Figures 2a-b compare the effect of UI on unemployment exit rates for single and dual-earner households. Figure 2a shows that UI benefits significantly reduce exit rates for households who are more likely to be constrained at the time of job loss because they were relying on a single source of income. The Wilcoxon test rejects equality of the survival curves with $p < 0.01$. In contrast, UI benefits appear to have no effect on exit hazards for households with two earners (Figure 2b).

The results for the mortgage cut are similar. Figure 3a shows that UI benefits have a sharp effect on durations among households that have a mortgage to pay off at the time of job loss, and equality of the two survival curves is again rejected with $p < 0.01$. But among households without a mortgage pre-unemployment, the difference between the survival curves in the high-benefit and low-benefit groups is much smaller and statistically indistinguishable (Figure 3b). In contrast with the two other proxies, the constrained types in this cut (homeowners with mortgages) have *higher* income, education, and wealth than the unconstrained types, who are primarily renters (see Table 1b). This makes it somewhat less likely that the differences in the benefit elasticity of duration

[10] The similarity of the effect of UI benefits on hazard rates in the third and fourth quartiles is consistent with lthe model. Once households have sufficient assets to avoid hitting a constraint, income effects disappear, and further increases in assets should have no impact on the UI-duration link.

across constrained and unconstrained groups is spuriously driven by other differences across the groups such as income or education.

The identifying assumption underlying this analysis is that the variation in UI benefits is orthogonal to unobservable determinants of durations, i.e. that $Eb \times \theta_{i,t} = 0$ in (11) holds. To evaluate this assumption, Figure 4 shows the effect of UI benefits on durations for a "control group" of below-median net wealth individuals who do not receive UI benefits, either because of ineligibility or because they chose not to take up. The durations of these individuals are insensitive to the level of benefits, as are the durations of non-recipients who have net wealth levels above the median. The results of these placebo tests support the claim that UI *causes* longer durations among constrained UI recipients.

3.3.2 Hazard Model Estimates

I evaluate the robustness of the graphical results by estimating (12) using a Cox specification for the hazard function. The Cox model assumes a proportional form for the hazard rate:

$$\log h_{i,s} = \alpha_s + \beta_1 \log b_i + \beta_2 s \times \log b_i + \beta_3 X_{i,s} \tag{13}$$

where $X_{i,s}$ denotes a set of covariates and $\{\alpha_s\}$ are the set of baseline hazards.[11] The coefficient of interest is β_1, the elasticity of the hazard rate with respect to UI benefits. To control for the fact that the relationship between UI benefits and the hazard rate may vary over time, the model also includes an interaction of $\log(b_i)$ with s, the weeks elapsed since job loss. Note that this specification does not impose any functional form on the baseline hazards, so the β_1 coefficient is identified purely from variation in UI laws.

I first estimate (13) 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 in the data section. This specification includes a full set of 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, pre-unemployment spousal work status, and being on the seam between interviews to adjust for the seam effect.

[11] The Cox model is more commonly specified as $h_{i,s} = \alpha'_s \exp(\beta_1 \log b_i + ...)$, which is equivalent to (13). I use (13) because it clarifies why the β_1 coefficient on log UI benefits represents an elasticity.

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 5.8% in the pooled sample, consistent with the estimates of prior studies.

Heterogeneity by Net Liquid Wealth Quartiles. I now examine the heterogeneity of the UI effect by estimating separate coefficients for constrained and unconstrained groups as in the graphical analysis. Table 2 considers the asset proxy for constraints by dividing the data into four quartiles of the net wealth distribution. In these specifications, I include households for which asset data are available either before or after the spell. Consistent with the graphical evidence, results are similar (though less precise) if households with only ex-post asset data are excluded.

Let $Q_{i,j}$ denote an indicator variable that is 1 if agent i belongs to quartile j of the wealth distribution. Let $\alpha_{s,j}$ denote the baseline exit hazard for individuals in quartile j in week s of the unemployment spell. To reduce parametric restrictions, the baseline hazards are allowed to vary arbitrarily across the constrained and unconstrained groups. Columns 2-5 of Table 2 report estimates of $\{\beta_1^j\}_{j=1,2,3,4}$ from the following stratified Cox model:

$$\log h_{isj} = \alpha_{s,j} + \beta_1^j Q_{i,j} \log b_i + \beta_2^j Q_{i,j}(s \times \log b_i) + \beta_3 X_{isj} \tag{14}$$

In this equation, β_1^j corresponds to the elasticity of the hazard rate w.r.t. UI benefits in quartile j of the net wealth distribution. Specification (2) of Table 2a reports estimates of (14) with no controls (no X). The estimates indicate that β_1^j is rising in j, i.e. the effect of UI benefits monotonically declines as one moves up in the net liquid wealth distribution.[12] Among households in the lowest quartile of net wealth, a 10% increase in UI benefits reduces the hazard rate by 7.9%, an estimate that is statistically significant at the 5% level. In contrast, there is a small, statistically insignificant 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. These findings support the conclusion that UI benefits have much stronger effects on durations for households constrained by low net liquid wealth.

Specification (3) replicates (2) with the full set of controls used in column (1), including state

[12] The β_2^j coefficients on the time interactions are generally insignificant and do not exhibit any strong patterns in this and subsequent specifications. In the interest of space, these coefficients are not presented in the tables.

and year fixed effects so that the coefficients are identified from changes in UI laws within states rather than cross-state comparisons. 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. The fact that controlling for observed heterogeneity has little impact on the results suggests that the estimates are unlikely to be very sensitive to unobservable heterogeneity as well.

In specifications (4) and (5), I explore robustness to changes in the definition of b_i. Both of these specifications include the full control set used in (1). 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 uppermost wealth quartile, consistent with the earlier results. In addition, the coefficient estimates are declining from quartiles 2 to 4. However, the coefficient 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.[13]

Finally, specification (6) replicates (3) but defines the quartiles of wealth in terms of home equity rather than net liquid wealth and restricts the sample to homeowners. Home equity is much less accessible than liquid wealth during an unemployment spell, since borrowing even against secured assets is difficult when one is unemployed (Hurst and Stafford 2004). If liquidity matters, the differences in the effect of UI benefits across quartiles of home equity should be weak compared to quartiles of net liquid wealth. Consistent with this prediction, there is no systematic pattern in the coefficients on the UI benefit variable across the quartiles in column 6.

I have fit a wide 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 (with pre-unemployment annual wages above the 75th percentile of the wage distribution) are excluded, addressing the concern that UI may have a smaller effect in the high asset quartiles because it

[13] One potential explanation for the difference in results is that there may be greater measurement error in the individual-level benefits in the lowest quartiles, leading to attenuation of those coefficients. In future work, it would be very useful to obtain administrative measures of individual-level benefits matched with asset data obtain more precise estimates.

replaces a smaller fraction of income for high-income households. Including individuals who report being on temporary layoff also does not affect the results. Defining the quartiles in terms of net wealth divided by wages rather than the absolute amount of net wealth also yields similar results.

Spousal Work Status. Table 3a reports estimates of specifications analogous to (14) for the spousal work proxy. Instead of quartiles of liquid wealth, the UI benefit coefficient is interacted with a dummy for whether the agent lived in a single-earner or dual-earner household prior to job loss. The baseline hazards are also stratified by this dummy. The first specification includes all observations in the core sample without any controls. In this group, there is a moderate but statistically insignificant difference in the UI benefit coefficient for the single-earner and dual-earner groups.

To explore this result in greater detail, observe that households with very low net wealth (who typically have substantial debt) are likely to be constrained irrespective of whether they have two earners or not, and households with very high net wealth are likely to be unconstrained regardless of spousal work status. Specification (2) therefore focuses on households in the middle two quartiles of the net wealth distribution, who are most likely to be on the margin of being liquidity constrained. In this subgroup of households, the effect of spousal work status emerges much more clearly. A 10% increase in the UI benefit reduces the mean unemployment exit hazard by 5.5% for single-earners but has a small, statistically insignificant effect for dual earners. The null hypothesis that the effect is identical in the two groups is rejected with $p = 0.06$. The third column shows that this result is robust to including the full set of controls described above. The fourth column adds state fixed effects, and shows that the general pattern is preserved although standard errors rise in this specification. Column 5 restricts attention to the households in the lowest quartile of net wealth. Consistent with the hypothesis that these households are constrained regardless of spousal work status, UI benefits have a strong effect on durations in both single-earner and dual-earner families in this category.

Mortgage Status. Table 3b shows results for the mortgage proxy using the observations for which pre-unemployment mortgage data is available. The first specification supports the graphical evidence in Figure 3, indicating that UI benefits have a much larger effect on durations among households that have mortgages. Equality of coefficients on the UI benefit variable among mortgage-holders and non-holders is rejected with $p < 0.01$. The second and third specifications confirm that this result is robust to the full set of controls and state fixed effects. The fourth specification includes only households with net liquid wealth below the sample median. The estimates indicate that low-wealth households who have to pay a mortgage – who are perhaps especially likely to face a liquidity constraint – are extremely sensitive to unemployment benefits in their search behavior.

Sample Selection Concerns. One might worry that endogeneity of takeup with respect to the level of benefits biases the estimate of the UI benefit elasticity. In my sample, a 10% increase in the benefit rate is associated with a 1% increase in the probability of UI takeup in the first month of unemployment. If the marginal individuals who decide to take up UI when benefits rise tend to have shorter unemployment spells on average, estimates of the UI benefit elasticity will be biased toward zero.

This issue is unlikely to affect the results above for two reasons. First, the takeup elasticity is similar across all the constrained and unconstrained subgroups. Hence, there is no reason that it should artificially bias down the estimate only in the unconstrained group. Second, even if there were differential biases across groups, the effects on the estimated UI benefit elasticity would be quite small. The magnitude of the bias can be gauged by assuming that the individuals who are added to the sample through this selection effect are drawn randomly from the group who do not takeup UI. The empirical hazards for the non-UI group are on average 1.1 times as large as those of the UI recipients. In practice, the marginal individual who takes up UI is likely to anticipate a longer UI spell than the average agent who does not take up UI, so the 1.1 ratio provides an upper bound for the size of the selection bias. Starting from an initial takeup rate of 50%, a 10% increase in benefits will cause the average hazard rate to rise through this selection effect by approximately $\frac{1\%}{50\%} * (1.1 - 1) = 0.2\%$. But the difference in the hazard rates across constrained and unconstrained groups induced by a 10% benefit increase was an order of magnitude larger (approximately 5%), suggesting that this selection effect is not critical.

Summary and Interpretation. The SIPP data exhibit two patterns: (1) UI benefits induce small substitution effects among households likely to be unconstrained; and (2) UI benefits cause large duration increases among the constrained. This pattern is consistent with the existence of substantial liquidity effects. 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 substantial. However, this assumption may be untenable: households with low liquidity might have different preferences that generate larger substitution effects than unconstrained households. For this reason, I avoid identification of the liquidity effect based on such cross-group comparisons, and turn to a second empirical strategy to investigate this issue.

4 Empirical Analysis II: Severance Pay and Durations

4.1 Estimation Strategy

The goal of this section is to decompose the effect of benefits on durations within the constrained group into a liquidity and substitution effect. The empirical strategy exploits the fact that many firms in the United States make severance payments to employees they lay off. According to a recent survey of Fortune 1000 firms (Lee Hecht Harrisson 2001), the most common policy for regular (non-executive) full-time workers is a severance payment of one week of pay for each year of service at the firm. However, some companies have flatter or steeper severance pay profiles with respect to job tenure. Many companies have minimum job tenure thresholds to be eligible for severance pay, ranging from 3 to 5 years. For regular salaried employees, there is very little variation in severance packages within a given firm and tenure bracket (presumably because individuals are reluctant to negotiate with firms about severance pay). Hence, conditional on tenure, the primary source of variation in severance pay comes from cross-firm differences in policies.

The key characteristic of severance payments for the present analysis is that they are lump-sum, i.e. they are not proportional to the length of unemployment spells. Receipt of standard tenure-based severance pay does not delay eligibility for UI benefits. Severance payments therefore have a pure liquidity effect, and do not distort marginal incentives for unemployed agents. I estimate the liquidity effect using models similar to those above, changing the key independent variable from the UI benefit to sev_i, a dummy for receipt of severance pay:

$$\log h_{i,s} = \alpha_s \theta_1 \text{sev}_i + \theta_2 X_{i,s} \tag{15}$$

The coefficient θ_1 reveals the causal effect of lump sum grants on unemployment exit hazards 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, matched with administrative data from state UI records. The first 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. In addition to administrative data on prior wages and weeks of UI paid, there are a large set of survey variables that give information on demographic characteristics, household income, job characteristics (tenure, occupation, industry), and most importantly for this study, receipt of severance pay.

The second dataset is the "Pennsylvania Reemployment Bonus Demonstration." This data was collected as part of an experiment to evaluate the effect of job reemployment bonuses on search behavior. It contains information on 5,678 durations for a representative sample of job losers in Pennsylvania in 1991. The information in the dataset is similar to that in the exhaustees study.

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.[14] First, I include only prime-age males. Second, I exclude temporary layoffs by discarding all individuals who expected a recall at the time of layoff, but check to make sure that including these observations do not change the results. These exclusions leave 2,730 individuals in the sample, of whom 521 report receiving a severance payment at job loss. Throughout the analysis, the data are reweighted using the sampling weights to obtain estimates for a representative sample of job losers.

Two measures of "unemployment duration" are available in this data. The first is the number of weeks for which UI benefits were paid in the base year. This definition has the advantage of accuracy since it comes from administrative records. It also has two disadvantages: it is censored at the time of benefit exhaustion, and it captures total weeks unemployed in a given year rather than the length of a particular spell (which could be different for individuals with multiple short spells). The second measure is the survey measure, constructed from individual's recollection (typically one-two years after the job loss event) of when they lost their initial job and when they found a new one. I focus on the administrative measure here given its accuracy. However, results are quite similar (with larger standard errors) for the survey measure.

Table 4 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.

[14] Similar results are obtained within each of the two datasets when examined separately.

4.3 Results

I begin again with graphical evidence. Figure 5 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. 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 unemployment exit rates. As a result, 66% of individuals who received severance pay claimed more than 10 weeks of UI benefits, compared with 59% among those who received no severance payment. Equality of the two survival curves is rejected by a nonparametric test with $p < 0.01$.

An obvious concern with this result is that it may reflect correlation (via omitted variables) rather than causality because severance pay recipients differ from non-recipients in many respects. As noted above, conditional on tenure, severance pay is determined primarily by firms and is therefore unlikely to be correlated with individual-specific characteristics. Hence, any omitted-variables explanation of the results must arise primarily from differences between firms that pay severance and those that do not. A plausible alternative explanation of the result is that firms that offer severance packages require very specific skills, making it difficult for job losers to find new jobs, leading to long durations.

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, many alternative omitted-varable explanations, such as the one proposed above, would not predict a differential effect of severance pay across constrained and unconstrained households. Hence, studying the heterogeneity of the severance pay effect provides 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 exactly the same prediction equation) and actual reported assets are very 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 6a-b replicate Figure 5 for the two groups. Figure 6a shows that receipt of severance pay is associated with a large and statistically significant increase in survival probabilities for constrained (low asset) households. Figure 6b shows that severance pay has little 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, and 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.[15]

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 Cox hazard models first with only a linear tenure control and then with the following control set: ten piece linear splines for wages, household income, job tenure; dummies for prior industry, occupation, race, state, and year; and linear controls for age, marital status, education, and household size. The first two columns of Table 5 show that receipt of severance pay lowers the job-finding hazard rate by about 12% in both the tenure-control and full-control specifications. Specifications (3) and (4) estimate separate severance pay coefficients for constrained (below-median predicted assets) and unconstrained (above-median) households. Consistent with Figure 6, the estimates indicate that severance pay has a significant effect on hazard rates only in the constrained group.

[15] Severance pay recipients look better on observables than non-recipients within both the high predicted-asset and low predicted-asset groups. This further supports the causal interpretation of the link between severance pay and durations in the constrained group, given that similar differences in observables do not generate a correlation between severance pay and durations among unconstrained households.

5 Welfare Implications

The severance pay estimates can be combined with the benefit elasticity estimates from the SIPP data to estimate how much of the UI-duration link is due to moral hazard vs. liquidity. This requires rescaling the coefficient estimates so that a \$1 increase in total UI benefits is compared with a \$1 increase in assets. To compute the scaling factor, note that the mean severance payment in the sample is roughly equal to 4 weeks of wages, which equals 8 weeks of UI benefits. The mean weeks of compensated unemployment in the SIPP is 14.9. Hence, receipt of severance pay is roughly equivalent to increasing UI benefits by $8/14.9 = 54\%$ at the mean.

The estimate from specification 1 in Table 2 implies that doubling the UI benefit will reduce the unemployment exit hazard by 44%. Analogously, the estimate from specification 2 of Table 5 indicates that receipt of severance pay reduces the hazard rate by 12%. It follows that

$$\frac{\partial s/\partial a}{\partial s/\partial b} = \frac{.12/.54}{.44} = 0.5$$

The evidence thus suggests that roughly half of the UI-duration link is due to a liquidity effect for UI recipients in the U.S. It is straightforward to use this estimate to implement (10) and calculate the welfare gain from raising the UI benefit level. Assume that the agent is unemployed for 5% of his life on average ($\sigma = 0.95$) and that the elasticity of duration $\varepsilon_{d,b} = 0.58$, which follows from the estimate in specification 1 using a constant-hazard approximation. Then it follows that

$$\frac{\partial W}{\partial b} = \frac{1-0.95}{0.95}\{1 - \frac{0.58}{0.95}\} = 0.02$$

This calculation shows that starting from the benefit levels observed in the U.S., a \$1 increase in the weekly benefit level would raise welfare by the equivalent of a 2 cent increase in weekly wage income. It should be emphasized that this welfare gain calculation applies only *locally* at the observed level of benefits in the data. Welfare gains of raising the benefit level could be much larger if the initial benefit level were near zero (because of larger liquidity effects), and much smaller if the initial benefit level were near full wage replacement (because substitution effects might be much larger). This is why the formula should be viewed as a "test" for optimal benefits rather than a means of calculating the welfare gain from UI as a function of the benefit level. The fact that the welfare gain is estimated to be small but positive suggests that the current UI benefit levels in the U.S. nearly satisfy the optimality condition in (10). Hence, the "revealed preference" approach proposed here suggests an optimal wage replacement rate for UI around 50%, considerably higher than the rates implied by existing consumption-based studies such as Baily (1978) and Gruber (1997).

The calculation above should be viewed as illustrative because it abstracts from many factors. First, the analysis focuses only on the duration margin. As emphasized by Feldstein (1978), Topel (1983), and others, UI benefits can distort other margins of behavior such as the incidence of layoffs, which may make a lower benefit level desirable. 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 level desirable. Finally, the estimate itself should be interpreted cautiously because it involves a comparison of point estimates from two different samples. Further work is needed to obtain a more precise estimate of the key parameter $\frac{\partial s/\partial a}{\partial s/\partial b}$.

It is worth noting that the moral hazard vs. liquidity analysis also has some policy implications beyond the level of benefits that could be explored more formally in future work. First, following the substitution effect interpretation of the benefit-duration link, a commonly held view is that policies which shorten unemployment durations at the margin can raise social surplus significantly by reducing the moral hazard problem. Examples of policies that aim to correct marginal incentives include more stringent search requirements that impose a cost on extending durations, provision of a job-finding bonus, or marginal reductions in benefits to induce shorter spells. This paper's analysis suggests that efforts to shorten durations would yield efficiency gains 50% smaller than suggested by studies that attribute the entire duration response to a substitution effect. Second, there is a debate on whether temporary income assistance programs should be means-tested, as in the United Kingdom. Browning and Crossley (2001) and Bloemen and Stancanelli (2003) find that UI does not smooth consumption for those who have high levels of pre-unemployment assets, which points in favor of asset-testing. However, UI does not appear to significantly affect unemployment durations for this group either. Since means-testing can generate an additional efficiency cost by creating an incentive to save less, a universal benefit may maximize welfare.

6 Conclusion

This paper proposes a new interpretation of the well known empirical relationship between unemployment benefits and durations. The existing 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 (moral hazard). The evidence documented here suggests a different, more favorable view of the UI-duration link: benefits increase durations largely because households have more liquidity while unemployed, reducing the pressure to find work quickly. I estimate that half of the UI-duration link is due to the liquidity effect and half is due to moral hazard. Based on a formula derived from a search theoretic model, this implies that increasing the UI benefit level from the current rate (roughly 50% of wages) would yield a modest welfare gain, suggesting that a wage replacement rate of 50% is near optimal.

While the analysis indicates that providing households' with substantial resources to smooth income shocks is valuable, one must be careful in drawing implications for government policy from this finding. Even if insuring unemployment is beneficial, it is not clear that government transfers are the best means of providing such insurance. For instance, Feldstein and Altman (1998) and Shimer and Werning (2005) argue that grants or low-interest loans (e.g. via privately held UI accounts) are a better means of providing liquidity than government transfers.

Finally, although this paper has focused on unemployment, the revealed preference approach to valuing insurance proposed here can be applied to a broad range of social and private insurance markets. For example, empirical 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.

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.

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 XX, May 2007.

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

Blundell, Richard and Thomas MaCurdy, "Labor Supply: A Review of Alternative Approaches," in Ashenfelter, Orley and David Card, eds., Handbook of Labor Economics 3 (Amsterdam: North-Holland, 1999).

Bodie, Zvi, Robert Merton, and William Samuelson, "Labor Supply Flexibility and Portfolio Choice in a Life Cycle Model", Journal of Economic Dynamics and Control 16 (1992), 427-450.

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

Centeno, Mário. 2004. "The Match Quality Gains from Unemployment Insurance." Journal of Human Resources 39(3): 839-863.

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

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).

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. (2001), "The wealth of the unemployed," Industrial and Labor Relations Review, 55(1), October, pp. 79-94.

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

Hurst, Erik, "Grasshoppers, Ants and Pre-Retirement Wealth: A Test of Permanent Income Consumers," Univ. of Chicago mimeo, 2004.

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.

Karni, E. (1999), "Optimal Unemployment Insurance: A Survey." Southern Economic Journal 66, 442–465.

Katz, Lawrence and Bruce Meyer, The Impact of the Potential Duration of Unemployment Benefits on the Duration of Unemployment," Journal of Public Economics 41 (1990) 45-72.

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, “Optimal Unemployment Insurance in an Estimated Job Search Model with Savings,” Boston University mimeo, 2003.

MaCurdy, Thomas (1981). ”An Empirical Model of Labor Supply in a Life Cycle Setting.” Journal of Political Economy. 89(6): 1059-1085.

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.

Moffitt, R., and W. Nicholson (1982), “The effect of unemployment insurance on unemployment: the case of federal supplemental benefits”, The Review of Economics and Statistics 64:1-11.

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

Shimer, Robert and Ivan Werning, “Reservation Wages and Unemployment Insurance,” Univ. of Chicago mimeo, 2006.

Sullivan, James. “Borrowing During Unemployment: Unsecured Debt as a Safety Net,” Notre Dame mimeo 2005.

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: Test for Optimal Benefits in General Case

Let b denote the (constant) UI benefit level and w the (constant) wage. Suppose the agent receives an annuity payment a in all periods in all states. 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. Recall that

$$\begin{aligned} J_0 &= (1-s_0)U_0(b,\tau) + sV_0(b,\tau) \\ \implies \frac{dJ_0}{db} &= (1-s_0)\{\frac{\partial U_0}{\partial b} - \frac{\partial U_0}{\partial w}\frac{d\tau}{db}\} - s_0\frac{\partial V_0}{\partial w}\frac{d\tau}{db} \end{aligned}$$

Since $\tau = \frac{bD}{T-D}$, it follows that $\frac{\partial \tau}{\partial b} = \frac{D}{T-D} + \frac{bT}{(T-D)^2}\frac{\partial D}{\partial b}$. Observe that

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

Let the average marginal utility of consumption while employed be denoted by $Eu'(c^e) = \frac{1}{T-D}[s_0\sum_{t=0}^{T-1}u'(c^e_{0,t}) + \sum_{t=1}^{T-1}[\prod_{i=1}^{t}(1-s_{i-1})]s_t\sum_{j=t}^{T-1}u'(c^e_{t,j})]$. Then

$$\begin{aligned} \frac{dJ_0}{db} &= (1-s_0)\frac{\partial U_0}{\partial b} - \frac{d\tau}{db}(T-D)Eu'(c^e) \\ &= (1-s_0)\frac{\partial U_0}{\partial b} - [DEu'(c^e) + \frac{bT}{T-D}\frac{\partial D}{\partial b}Eu'(c^e)] \end{aligned}$$

Next, consider the effect of raising a on U_0 and V_0:

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

Note that for any variable $x \in \{a,b,w\}$,

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

Using this expression and the preceding formulas, it is easy to show that

$$\frac{\partial s_0}{\partial a} = \frac{\partial s_0}{\partial b} + \frac{\partial s_0}{\partial w}$$

Next, observe that $\frac{\partial s_0}{\partial a} = \frac{1}{\psi''}[\Delta - \frac{\partial U_0}{\partial b}]$ where

$$\begin{aligned}
\Delta &= \frac{\partial s_0}{\partial w}\psi'' = \sum_{t=0}^{T-1} u'(c^e_{0,t}) - (s_1 \sum_{j=1}^{T-1} u'(c^e_{1,j}) + \sum_{t=2}^{T}[\prod_{i=2}^{t}(1-s_{i-1})]s_t \sum_{j=t}^{T} u'(c^e_{t,j})) \\
&= \sum_{t=0}^{T-1} u'(c^e_{0,t}) + \frac{s_0}{1-s_0}\sum_{t=0}^{T-1} u'(c^e_{0,t}) - \frac{1}{1-s_0}\{s_0 \sum_{t=0}^{T-1} u'(c^e_{0,t}) + \sum_{t=1}^{T-1}[\prod_{i=1}^{t}(1-s_{i-1})]s_t \sum_{j=t}^{T-1} u'(c^e_{t,j}))\\
&= \frac{1}{1-s_0}\{\sum_{t=0}^{T-1} u'(c^e_{0,t}) - (T-D)Eu'(c^e)\}
\end{aligned}$$

Let $Eu'(c^e_0) = \frac{1}{T}\sum\limits_{t=0}^{T-1} u'(c^e_{0,t})$ denote the average marginal utility of consumption while employed if the agent finds a job in period 0. Then

$$\Delta = \frac{1}{1-s_0}\{TEu'(c^e_0) - (T-D)Eu'(c^e)\}$$

Under the approximation $Eu'(c^e_0) = Eu'(c^e)$, it follows that

$$\begin{aligned}
\Delta &= \frac{1}{1-s_0}DEu'(c^e)\\
\frac{\partial s_0}{\partial a} &= \frac{1}{\psi''}[\frac{D}{1-s_0}Eu'(c^e) - \frac{\partial U_0}{\partial b}]
\end{aligned}$$

and hence

$$\begin{aligned}
\frac{dJ_0}{db} &= (1-s_0)\frac{\partial U_0}{\partial b} - [DEu'(c^e) + \frac{bT}{T-D}\frac{\partial D}{\partial b}Eu'(c^e)]\\
&= -(1-s_0)\frac{\partial s_0}{\partial a}\psi'' - \frac{bT}{T-D}\frac{\partial D}{\partial b}Eu'(c^e)
\end{aligned}$$

Note that

$$\frac{\partial s_0}{\partial w} = \frac{1}{\psi''}\Delta = \frac{1}{\psi''}\frac{1}{1-s_0}DEu'(c^e)$$

Normalizing the expression for $\frac{dJ_0}{db}$ by the expected welfare gain from increasing the wage by \$1, $\frac{dJ_0}{dw} = (T-D)Eu'(c^e)$ yields

$$\begin{aligned}
\frac{\partial W}{\partial b} &= \frac{1}{T-D}\{-\frac{1}{Eu'(c^e)}(1-s^*_0)\frac{\partial s^*_0}{\partial a}\psi'' - \frac{bT}{T-D}\frac{\partial D}{\partial b}\}\\
&= \frac{1}{T-D}\{-D\frac{\partial s^*_0}{\partial a}/\frac{\partial s^*_0}{\partial w} - \frac{bT}{T-D}\frac{\partial D}{\partial b}\}\\
&= \frac{D}{T-D}\{-\frac{\partial s^*_0}{\partial a}/\frac{\partial s^*_0}{\partial w} - \frac{T}{T-D}\varepsilon_{D,b}\}
\end{aligned}$$

Finally, recalling that $-\frac{\partial s_0}{\partial w} = \frac{\partial s_0}{\partial b} - \frac{\partial s_0}{\partial a}$ and defining $\sigma = \frac{T-D}{T}$, it follows that

$$\frac{\partial W}{\partial b} = \frac{1-\sigma}{\sigma}\{\frac{\partial s^*_0/\partial a}{\partial s^*_0/\partial b - \partial s^*_0/\partial a} - \frac{\varepsilon_{D,b}}{\sigma}\}.$$

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 and 90,616 unemployment spells.

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 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 in Appendix C. This leaves me with a sample of 16,784 individuals (3.6% of original sample) who experienced a total of 21,796 unemployment spells. Finally, to arrive at the core analysis sample, I restrict attention to individuals who were not temporarily laid off (leaving 21,107 spells) and those individuals who take up benefits within the first month of unemployment. 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.

TABLE 1a

Summary Statistics by Wealth Quartile in SIPP Sample

		Net Liquid Wealth Quartile			
	Pooled	1 (< -$1,115)	2 (-$1,115-$128)	3 ($128-$13,430)	4 (>$13,430)
Median Liq. Wealth	$1,763	$466	$0	$4,273	$53,009
Median Unsecured Debt	$1,000	$5,659	$0	$353	$835
Median Home Equity	$8,143	$2,510	$0	$11,584	$48,900
Median Annual Wage	$17,780	$17,188	$14,374	$18,573	$23,866
Mean Years of Education	12.07	12.21	11.23	12.17	13.12
Mean Age	36.99	35.48	35.18	36.64	41.74
Fraction Renters	0.39	0.43	0.61	0.35	0.16
Fraction Married	0.61	0.64	0.59	0.60	0.63

TABLE 1b

Summary Statistics by Spousal Work and Mortgage Status in SIPP Sample

	Dual Earner?		Has Mortgage?	
	No (0.63)	Yes (0.37)	No (0.55)	Yes (0.45)
Median Liq. Wealth	$1,193	$3,001	$630	$4,855
Median Unsecured Debt	$778	$1,357	$523	$1,725
Median Home Equity	$3,838	$15,801	$0	$30,421
Median Annual Wage	$16,472	$20,331	$15,946	$20,792
Mean Years of Education	11.84	12.46	11.88	12.53
Mean Age	35.33	39.79	35.96	38.66
Fraction Renters	0.44	0.30	0.71	0.00
Fraction Married	0.38	1.00	0.55	0.70

NOTE--Data source is 1985-87, 1990-93, and 1996 SIPP panels. All monetary values are in real 1990 dollars. 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. Sample size is 4,560 observations. Liquid wealth is defined as total wealth minus all home equity, business equity, and vehicle equity. Net liquid wealth is liquid wealth minus unsecured debt. Dual earner families are those where spouse is working in month immediately preceding layoff.

TABLE 2

Hazard Model Estimates by Quartile of Net Liquid Wealth

	(1)	(2)	(3)	(4)	(5)	(6)
	Pooled	By Quartile of Net Liquid Wealth				
	Full cntrls	No cntrls	Full cntrls	Max WBA	Simulated Indiv WBA	Home Equity
log UI ben	-0.580					
	(0.269)					
Q1 x log UI ben		-0.794	-1.022	-0.695	-0.589	-0.471
		(0.300)	(0.357)	(0.287)	(0.206)	(0.949)
Q2 x log UI ben		-0.710	-0.829	-0.289	-0.668	-0.249
		(0.458)	(0.412)	(0.341)	(0.177)	(0.662)
Q3 x log UI ben		-0.165	-0.520	-0.049	-0.508	-0.669
		(0.301)	(0.343)	(0.363)	(0.143)	-0.848
Q4 x log UI ben		0.111	0.081	0.237	-0.149	-0.24
		(0.344)	(0.402)	(0.344)	(0.243)	-0.743
state, year, industry, and occup. fixed effects	x		x	x	x	x
Q1=Q4 p-val		0.043	0.024	0.001	0.229	0.77
Q1+Q2=Q3+Q4 p-val		0.012	0.028	0.004	0.179	0.846
Number of Spells	4529	4337	4054	4054	4054	1583

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. Bottom rows of table report p-values from F-test for equality of reported coefficients across quartiles. Specifications 2-5 include log UI ben interacted with asset quartile dummies as well as log UI ben interacted with weeks unemployed interacted with asset quartile dummies to capture time-varying effects of UI (see text for details). Specs 3-5 include in addition the following controls: 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, pre-unemployment spousal work status, and being on the seam between interviews to adjust for the seam effect.

Spec 1 includes log UI ben, log UI ben interacted with weeks unemployed, and full control set reported above.

In specs 2-5, baseline hazards are stratified by net liquid wealth quartile.

In specs 1-3 and 6, UI ben is defined as average UI benefit in claimant's state/year pair.

Spec 2. proxies for UI ben using state/year maximums rather than averages.

Spec 5. uses individual-level simulated benefits based on wage histories (see text for details).

TABLE 3a

Hazard Model Estimates by Spousal Work Status

	(1)	(2)	(3)	(4)	(5)
	Full sample No cntrls	Middle netliq Qs No cntrls	Middle netliq Qs Full cntrls	Middle netliq Qs State FE's	netliq Q=1 Full cntrls
Single earner x log UI ben	-0.443	-0.550	-0.607	-0.969	-0.814
	(0.267)	(0.284)	(0.259)	(0.330)	(0.446)
Dual earner x log UI ben	-0.308	0.109	0.138	-0.451	-0.730
	(0.286)	(0.457)	(0.423)	(0.363)	(0.453)
Single = Dual p-val	0.590	0.057	0.070	0.217	0.901
Observations	84363	40905	36828	36828	19130

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. Bottom row of table reports p-values from F-test for equality of coefficients across single and dual earners. All specs include log UI ben interacted with dummies for spousal work as well as log UI ben interacted with weeks unemployed interacted with spousal work dummies to capture time-varying effects of UI (see text for details). Specs 3-5 include in addition the following controls: Industry, occupation, and year dummies; 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. Spec 4 also includes state fixed effects.
Spec 1 includes all observations; specs 2-4 only observations that lie between 25th and 75th percentile of net liquid wealth distribution; spec 5 includes only observations below 25th percentile of net liquid wealth distribution.
In all specs, baseline hazards are stratified by spousal work status.
Number of observations equals total risk set (i.e., total number of unemployed weeks observed in the dataset).

TABLE 3b

Hazard Model Estimates by Mortgage Holding Status

	(1)	(2)	(3)	(4)
	Full sample No cntrls	Full sample Full cntrls	Full sample State FE's	netliq Q <=2 Full cntrls
No mortgage x log UI ben	0.269	0.322	0.263	-0.135
	(0.292)	(0.218)	(0.455)	(0.390)
Mortgage x log UI ben	-0.976	-0.938	-0.957	-1.552
	(0.424)	(0.419)	(0.484)	(0.667)
No mortg. = Mortg. p-val	0.002	0.005	0.008	0.047
Observations	37087	35291	35291	16656

NOTE-Coefficients reported are elasticities of hazard rate w.r.t. UI bens. Standard errors clustered by state in parentheses. Sample consists of households in core sample for whom pre-unemp mortgage data is available. See note to Table 1 for definition of core sample. Bottom row of table reports p-values from F-test for equality of coefficients across non-mortgage and mortgage holders. All specs include log UI ben interacted with mortgage dummies as well as log UI ben interacted with weeks unemployed interacted with mortgage dummies to capture time-varying effects of UI (see text for details). Specs 2-4 include in addition the following controls: Industry, occupation, and year dummies; 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, spousal work, and being on the seam between interviews to adjust for the seam effect. Spec 3 also includes state fixed effects. Specs 1-3 includes all observations; spec 4 only observations that lie below 50th percentile of net liquid wealth distribution. In all specs, baseline hazards are stratified by mortgage dummy. Number of observations equals total risk set (i.e., total number of unemployed weeks observed in the dataset).

TABLE 4

Summary Statistics for Mathematica Sample

	Pooled	No Severance	Severance
		(0.82)	(0.18)
Median pre-unemp job tenure (years)	1.9	1.6	4.8
Median pre-unemp annual wage	\$20,828	\$19,183	\$29,874
Percent dropouts	14%	16%	6%
Percent college grads	17%	13%	34%
Percent married	58%	56%	68%
Mean age	36.4	35.5	40.6
Mean number of persons in hhold	2.21	2.24	2.04

NOTE--Data source is Study of Unemployment Insurance Exhaustees and Pennsylvania Reemployment Bonus Demonstration (Mathematica surveys matched to administrative UI records). These datasets are publicly available through the Upjohn Institute. All monetary values are in real 1990 dollars.
Sample includes prime-age male UI claimants who are not on temporary layoff.
Sample size is 2,730 observations. Data is reweighted using sampling probabilities to yield estimates for a representative sample of job losers. Pre-unemp job tenure is defined as number of years spent working at firm from which worker was laid off.

TABLE 5

Effect of Severance Pay on Hazard Rates

	Pooled		By Net Liquid Wealth	
	(1)	(2)	(3)	(4)
	Tenure Control	Full Controls	Tenure Control	Full Controls
Severance Pay Dummy	-0.115	-0.127		
	(0.030)	(0.035)		
(Netliq < Median) x Sev Pay			-0.476	-0.445
			(0.084)	(0.093)
(Netliq > Median) x Sev Pay			0.068	0.058
			(0.058)	(0.060)
Equality of coeffs p-val			<0.001	<0.001
Observations	2730	2426	2561	2426

NOTE-Coefficients reported are elasticities of hazard rate w.r.t. UI bens. Standard errors clustered by state in parentheses. See note to Table 4 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. Columns 1 and 3 include only a linear control for tenure at pre-job loss employer in addition to reported coefficients. Columns 2 and 4 include the following controls: ten piece linear splines for wages, household income, job tenure; dummies for prior industry, occupation, race, state, and year; linear controls for age, marital status, education, and household size. Baseline hazards in specs 3-4 are stratified by Netliq < Median. Netliq < Median is an indicator variable for whether the household's predicted assets (using an equation estimated from the SIPP data; see text) are below the sample median. Netliq > Median is defined analogously.

Figure 1a

Effect of UI Benefits on Durations: Lowest Quartile of Net Wealth

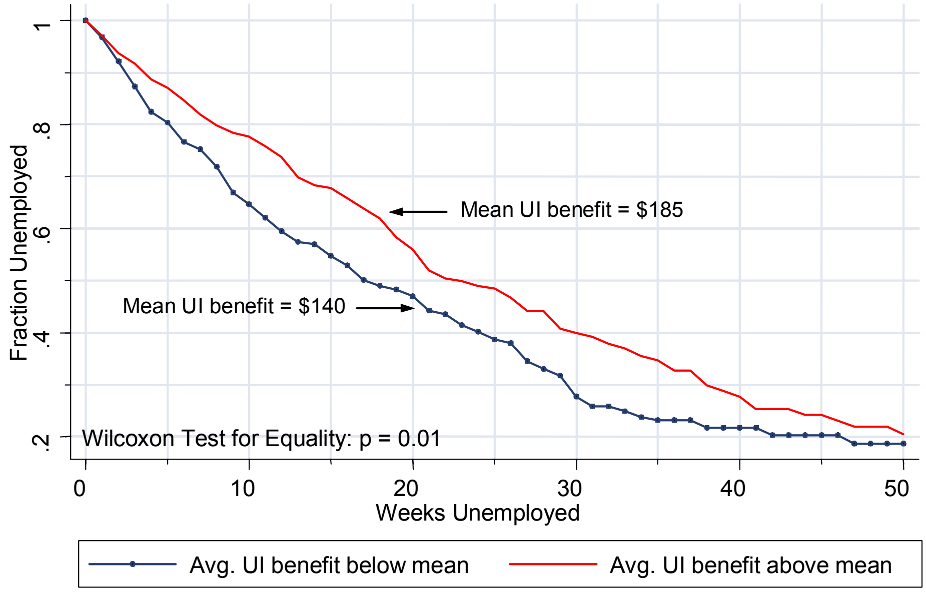

Figure 1b

Effect of UI Benefits on Durations: Second Quartile of Net Wealth

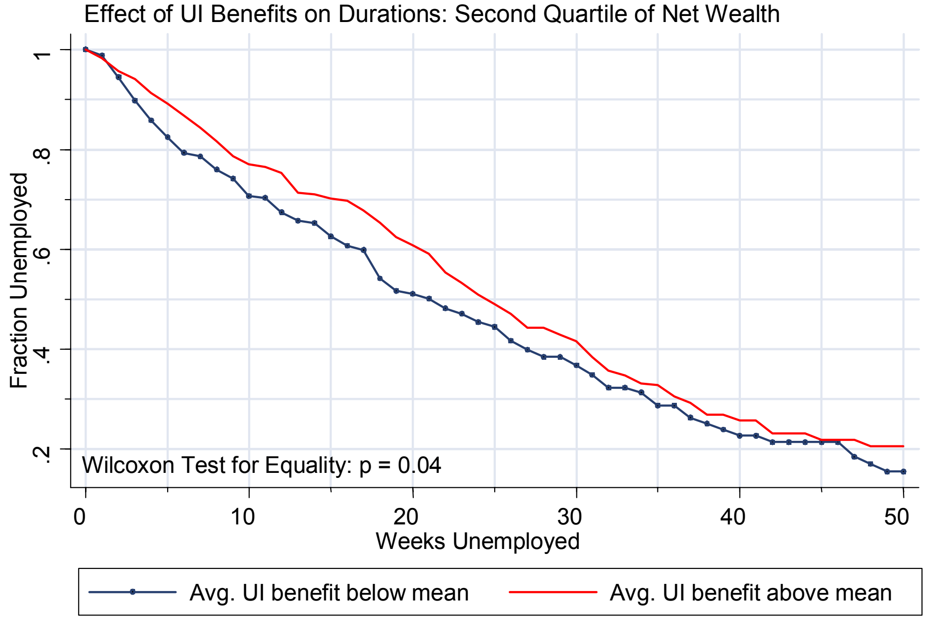

NOTE–Sample for both figures consists of observations in the core SIPP sample for which pre-unemployment wealth data are available. See Table 1 for definition of core sample and definition of net liquid wealth. Figure 1a includes households in lowest quartile of real net liquid wealth. Figure 1b includes 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. Survival curves are adjusted for seam effect by fitting a Cox model and recovering baseline hazards as described in text.

Figure 1c

Effect of UI Benefits on Durations: Third Quartile of Net Wealth

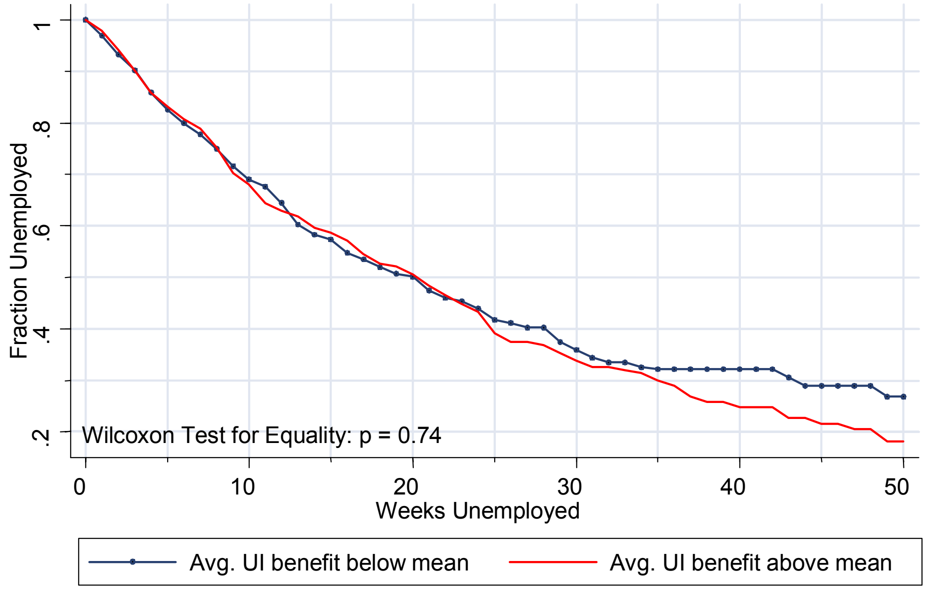

Figure 1d

Effect of UI Benefits on Durations: Highest Quartile of Net Wealth

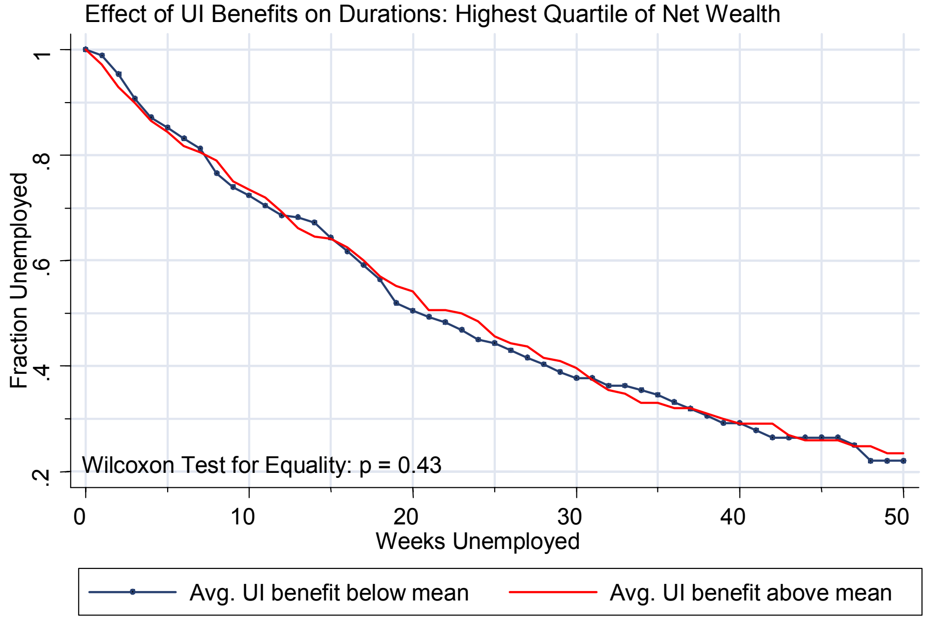

NOTE–These figures are constructed in the same way as Figures 1a-b using observations in the third and fourth quartiles of net wealth. See notes to Figures 1a-b for details.

Figure 2a

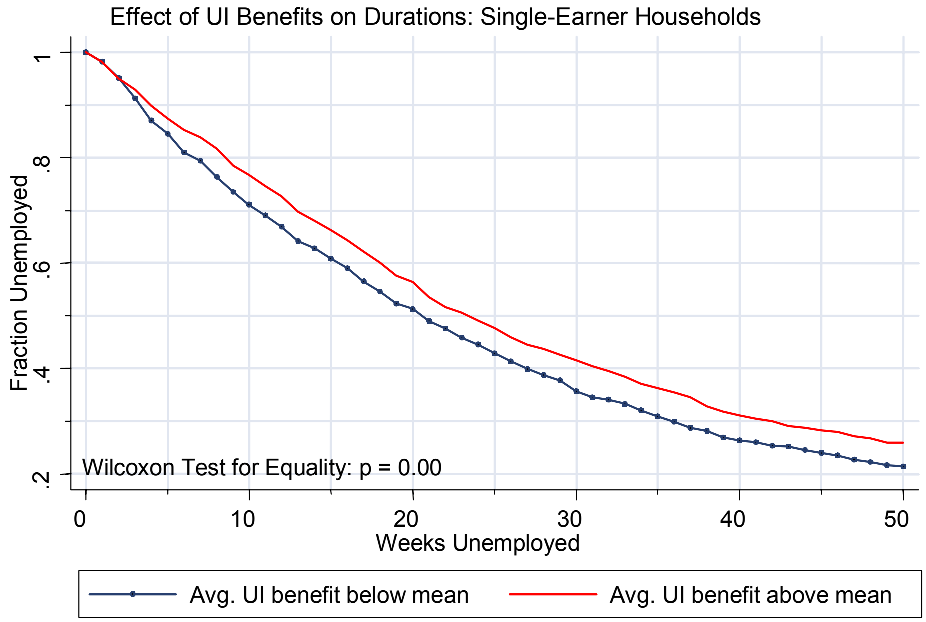

Figure 2b

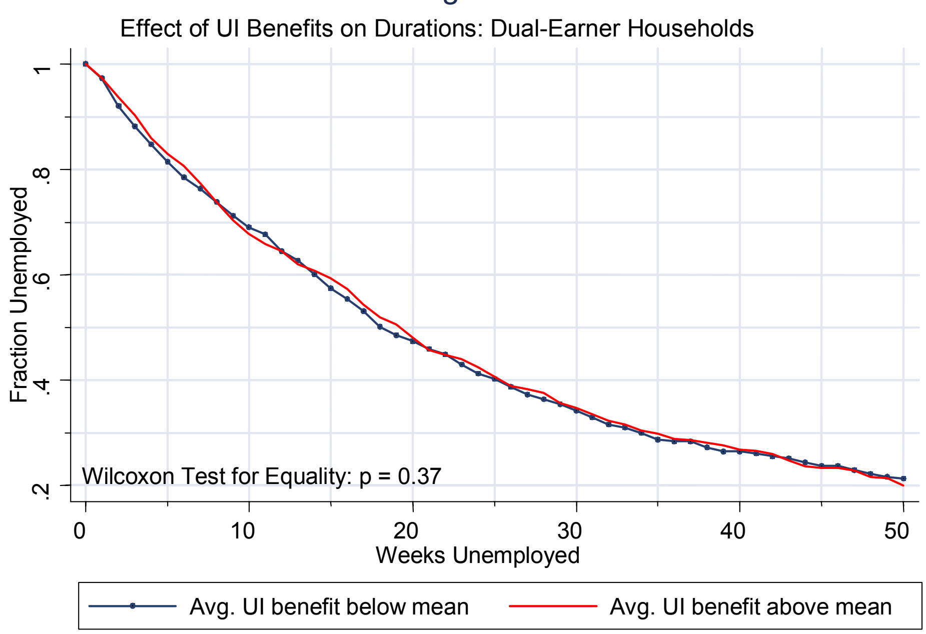

NOTE–Figure 2a includes households in the core SIPP sample with one earner in the month prior to job loss. Figure 2b includes households with two earners. See Table 1 for definition of core sample. Each figure plots Kaplan-Meier survival curves for two groups of individuals: Those in state/year pairs with average weekly benefit amounts (WBA) below the sample mean and those in state/year pairs with WBAs above the mean. Survival curves are adjusted for seam effect by fitting a Cox model and recovering baseline hazards as described in text.

Figure 3a

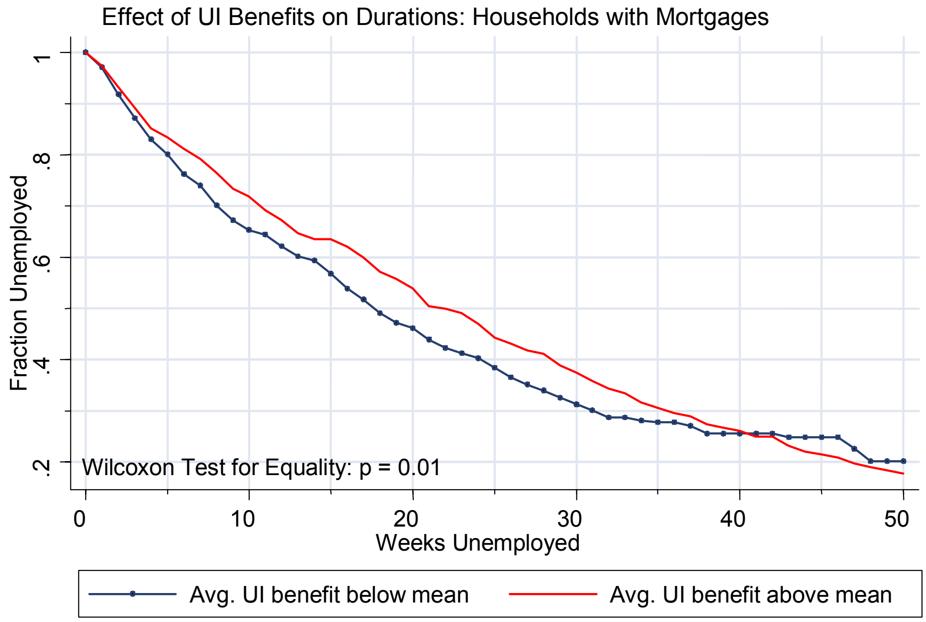

Figure 3b

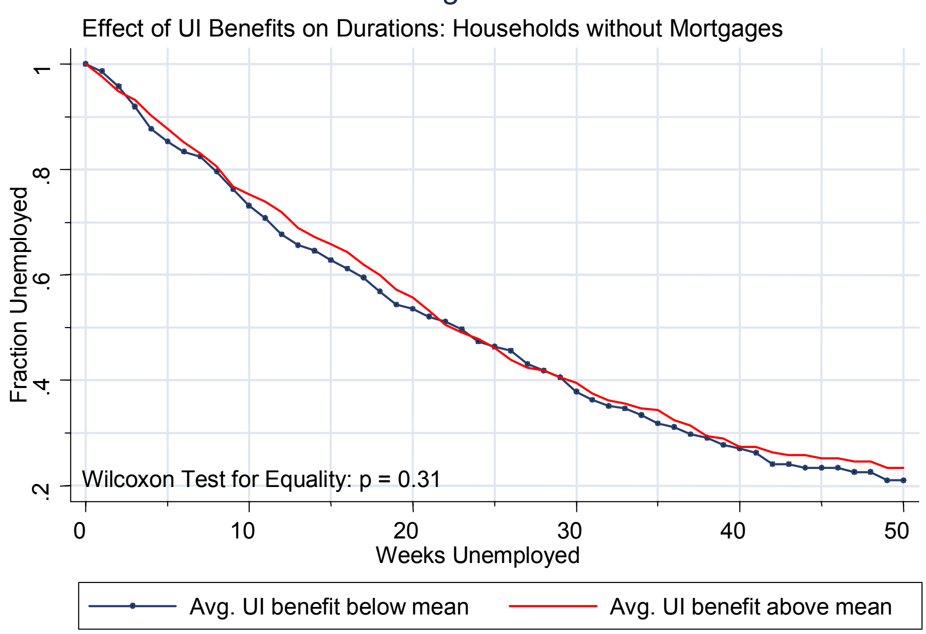

NOTE–These figures are constructed in the same way as Figures 2a-b. Figure 3a includes households who make mortgage payments; 3b includes all others. Only observations with mortgage data prior to job loss are included. See notes to Figures 2a-b for additional details on construction of these figures.

Figure 4

Placebo Test: Effect of UI Benefits on Low-Wealth Non-Recipients

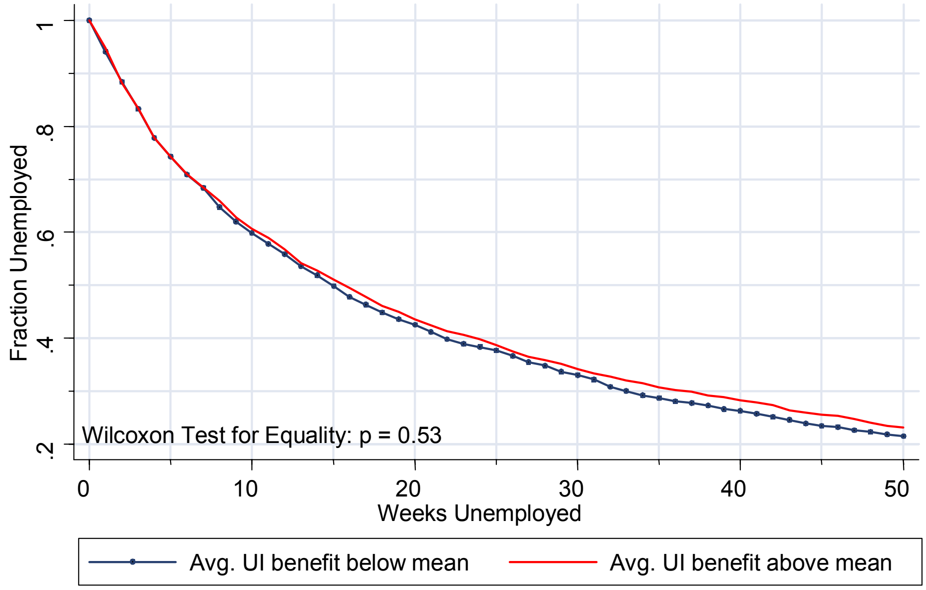

NOTE–The sample for this figure consists of prime-age male job losers in the SIPP data who (a) do not report receiving UI benefits while unemployed, (b) report searching for a job, (c) are not on temporary layoff, (d) have at least 3 months of wage history prior to job loss, and (e) have below-median net liquid wealth prior to job loss. Kaplan-Meier survival curves are plotted 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. Survival curves are adjusted for seam effect by fitting a Cox model and recovering baseline hazards as described in text.

Figure 5

Effect of Severance Pay on Durations

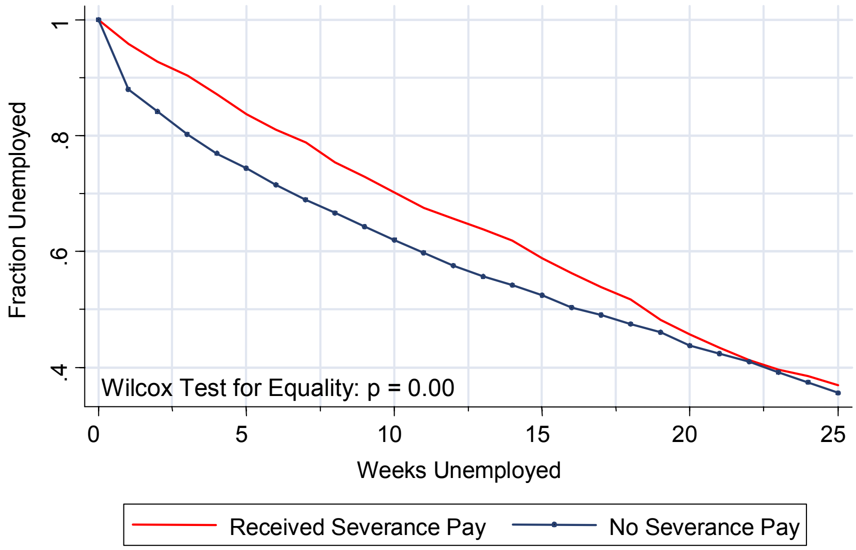

NOTE–Data are from Mathematica surveys matched to administrative UI records. See note to Table 4 for additional details on data and sample definition. Data is reweighted using sampling probabilities to yield estimates for a representative sample of job losers. Kaplan-Meier survival curves are plotted 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 6a

Effect of Severance Pay on Durations: Below Median Net Wealth

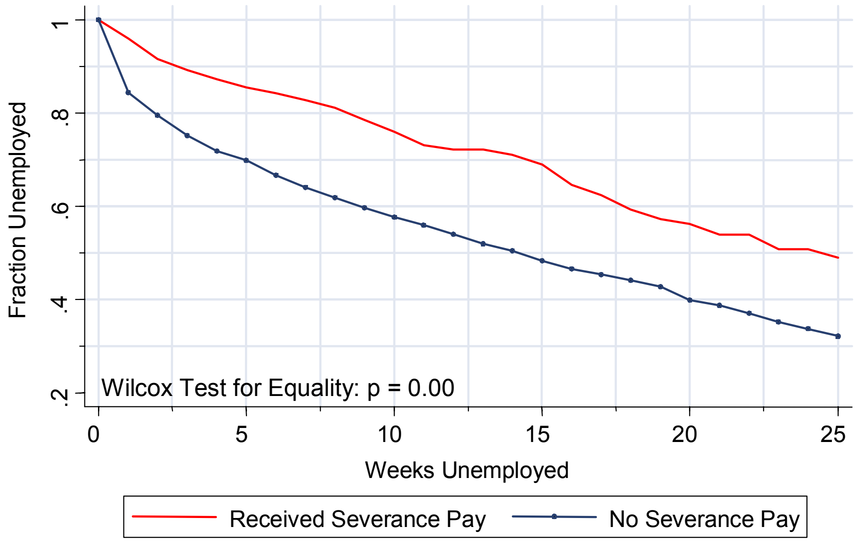

Figure 6b

Effect of Severance Pay on Durations: Above Median Net Wealth

NOTE–See Figure 5 for sample definition. Each of these figures is constructed in exactly the same way as Figure 5. Figure 6a includes observations where predicted net wealth is below the sample median; Figure 6b 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.