Why do Unemployment Benefits Raise Unemployment Durations? Moral Hazard vs. Liquidity

Raj Chetty*

UC-Berkeley and NBER

March 2006

Abstract

It is well known that unemployment benefits raise unemployment durations. This result has traditionally been interpreted as a substitution effect caused by a distortion in the price of leisure relative to consumption, leading to moral hazard. This paper questions this interpretation by showing that unemployment benefits can also affect durations through an income effect for agents with limited liquidity. The empirical relevance of liquidity constraints and income effects is evaluated in two ways. First, I divide households into groups that are likely to be constrained and unconstrained based on proxies such as asset holdings. I find that increases in unemployment benefits have small effects on durations in the unconstrained groups but large effects in the constrained groups. Second, I find that lump-sum severance payments granted at the time of job loss significantly increase durations among constrained households. These results suggest that unemployment benefits raise durations primarily because of an income effect induced by liquidity constraints rather than moral hazard from distorted incentives.

Keywords: borrowing constraints, income effects, insurance

*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, Emmanuel Saez, Robert Shimer, Adam Szeidl, Ivan Werning, 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 classic 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.[1] 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 incentive to work. In their handbook chapter on social insurance, 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 the substitution effect and moral hazard interpretation of the link between unemployment benefits and durations. The analysis is motivated by recent studies which give strong evidence that many unemployed individuals face liquidity constraints. Gruber (1997) finds that increases in UI benefits reduce the consumption drop during unemployment, indicating that agents are unable to smooth consumption relative to permanent income. Browning and Crossley (2001) and Bloemen and Stancanelli (2005) provide more direct evidence for the liquidity constraint mechanism by showing that the UI-consumption link identified by Gruber holds only for the subset of individuals who report holding few assets at the time of job loss. Nearly half of job losers in the United States report zero liquid wealth at the time of job loss, suggesting that liquidity is a concern for many of the unemployed.

I show how liquidity constraints affect the UI-duration link using a stylized model in which unemployed agents make search decisions and face a constraint on the amount of debt they can hold. I first establish that UI has negligible wealth effects for agents who are able to smooth consumption during unemployment spells, since UI benefits do not change permanent income very much.[2] Hence, for unconstrained individuals, UI has essentially a pure substitution effect, consis-

[1] See Atkinson and Micklewright (1990) and Krueger and Meyer (2002) for comprehensive reviews of this literature.

[2] For frequent job losers, the wealth effects of UI may be non-trivial even without liquidity constraints. This issue is addressed in greater detail in section 2.

tent with the traditional interpretation of the empirical evidence. However, for agents who face liquidity constraints, behavior while unemployed is determined by cash on hand rather than lifetime resources. UI benefits provide liquidity and raise consumption while unemployed. As a result, agents face less pressure to find a new job quickly, and may choose to lengthen their unemployment duration by raising their reservation wage or lowering search effort. Hence, for agents with limited liquidity, UI benefits can raise durations for a second reason: an *income* effect that occurs because agents have more cash on hand while unemployed.[3] This income effect occurs independently of and in addition to the conventional substitution effect.

The distinction between the income and substitution effect is of interest because the two effects lead to divergent views about behavioral responses to insurance. The substitution effect is a socially suboptimal response to the creation of a wedge between private and social marginal costs. The lengthening of unemployment durations by the substitution effect – sometimes labelled "gaming the system" – creates a moral hazard cost of providing insurance. In contrast, the income effect is a socially optimal response to the correction of a market failure (insufficient liquidity). An increase in durations is thus less desirable from a normative perspective if it arises from a substitution effect (see section 5 for details). Therefore, estimating the magnitude of the substitution effect is a key step in assessing the relevance of moral hazard in UI as well as other insurance programs.

I use two complementary empirical strategies to investigate the importance of income vs. substitution effects in UI. The first method provides estimates of the effect of UI on durations separately for liquidity-constrained and unconstrained households. This heterogeneity analysis is helpful in determining the plausibility of income effects. For example, if UI benefits were to affect durations only among unconstrained households, income effects could not be very large. But if the UI-duration link were driven primarily by constrained households, income effects generated by liquidity constraints could be relevant.

An obvious difficulty in implementing the heterogeneity analysis is that one cannot directly observe which households face liquidity constraints in the data. To overcome this latent variable

[3]The term "income effect" is used somewhat differently in this paper than in typical models where there are no liquidity constraints. To clarify, I use the term "income effect" to describe the effect of having additional income (i.e., additional liquidity) while unemployed. This differs from the effect of raising lifetime wealth, which I refer to as a "wealth effect."

problem, I use several proxies that have been shown to predict constraints in studies of consumption (e.g. Zeldes 1989, Browning and Crossley 2001). The first proxy is the household's liquid asset holdings net of unsecured debt (or assets relative to earnings) prior to job loss. Households with higher levels of assets are less likely to become constrained than those who have a smaller buffer stock. The second proxy is whether the individual had a working spouse prior to job loss. Dual-earner households are more likely to have the resources and credit access necessary to smooth consumption when one of them loses a job. The third proxy is whether the individual was making a mortgage payment before losing his job, which is a rigid commitment that effectively reduces liquid wealth.

I examine the effect of UI benefits on unemployment exit hazards in each of the constrained and unconstrained groups using nonparametric graphical methods and Cox hazard models. Visual analysis, non-parametric tests, and hazard model estimates uniformly indicate that a 10% increase in UI benefits raises unemployment durations by 6-8% in all the constrained groups, but has little or no effect on durations in the unconstrained groups. These results are robust to the inclusion of rich controls and other specification checks such as the permission of unobserved heterogeneity in baseline hazards. In addition, there is no association between UI benefits and durations in a "control group" of UI-ineligible and non-claiming individuals, supporting the exogeneity of the UI benefit rates.

These results show that the link between unemployment benefits and durations documented in prior studies is driven by a subset of the population that is likely to be liquidity constrained. This point requires careful interpretation. Barring additional assumptions, the evidence does not establish that liquidity constraints *cause* larger responses to UI benefits. It simply shows that UI benefits have different effects in constrained and unconstrained groups. Whether these differences arise because of the constraints themselves or because of correlation between preferences and asset holdings, which are endogenous to preferences, is unclear. What is clear – and most relevant for the main point of this paper – is that the substitution effect for the unconstrained group is small, while the total benefit elasticity of durations in the constrained group is large. These findings are consistent with the hypothesis that an income effect is involved in the UI-duration link, but do not establish the existence of an income effect by themselves (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 explicitly decompose the benefit elasticity of unemployment durations into an income and substitution effect using variation in lump-sum severance payments. To do so, I use a new dataset that matches survey data collected by Mathematica with administrative records on unemployment durations. Non-parametric analysis shows 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, suggesting that income effects are indeed large. An obvious concern is that this finding may reflect correlation rather than causality because severance pay is not randomly allocated. 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 model. Second, severance payments have a large positive effect on durations among constrained (low asset) households, but have no effect on durations among unconstrained households, consistent with the model's predictions. Since there is no a priori reason to expect a differential effect of severance pay by asset holdings under the most plausible omitted-variable hypotheses, this evidence supports the claim that lump-sum income grants lengthen unemployment durations.

Combining the point estimates from the two empirical approaches, a simple calculation indicates that roughly 70% of the UI-duration link is caused by an income effect.[4] Hence, the liquidity effect identified here appears to be more important than the marginal incentive distortions typically emphasized as an explanation of why unemployment benefits lengthen unemployment durations. This finding for UI raises the possibility that moral hazard is less prevalent in other types of private and social insurance than suggested by the existing literature as well.

I would like to emphasize two limitations of this paper before proceeding. First, a full welfare analysis of UI would require a complete model of job separations and finding with endogenous determination of saving behavior based on unemployment benefits. This analysis is outside the

[4]The non-zero estimate of the substitution effect is consistent with the spike in the hazard rate around benefit exhaustion (Katz and Meyer 1990). This spike could partially be generated by an income effect as agents anticipate losing benefits, but its magnitude does suggest some intertemporal substitution.

scope of this paper. My goal here is simply to identify the key empirical patterns that should inform such a welfare analysis, using a stylized model that makes the main intuition transparent. A second limitation is that the empirical analysis in this paper is not based on randomization, and thus one may have concerns about omitted variable biases in interpreting some aspects of the evidence. In defense of the results, most alternative explanations would not simultaneously explain all the patterns observed in the two datasets. For example, though constrained households might be more responsive to UI benefits because of unobserved heterogeneity in preferences, the finding that both severance pay and UI benefits affect the behavior of only these households still implies an income effect. Nonetheless, a study that uses randomized variation in lump sum grants is needed to obtain the most compelling and precise estimates of income effects. In view of these limitations, this paper should be viewed as a first step that calls for more research on disentangling income and substitution effects to better understand the relevance of moral hazard in insurance provision.

The remainder of the paper proceeds as follows. The next section formalizes the two potential behavioral responses to UI in a lifecycle model with borrowing constraints. Section 3 describes the estimation strategy, data, and results for the borrowing constraint and heterogeneity tests. Section 4 examines the effect of severance payments on durations. Section 5 discusses some normative implications of the results, and section 6 concludes.

2 Theory

I analyze the effects of borrowing constraints and UI benefits in a stylized model similar to that used by Zeldes (1989) to analyze the effect of borrowing constraints on consumption dynamics. The only differences are that the model below ignores portfolio choice but introduces endogenous labor supply to study unemployment durations.

Let c_s denote consumption at time s and $\widetilde{w}$ denote the agent's wage, which is constant over time. Normalize the interest and discount rates at zero. Assume that the agent lives for T years (in continuous time) and loses his job at time $t < T$.

I make several assumptions to simplify the analysis, and discuss later why relaxing them does not change the main results. First, I model the borrowing constraint by assuming that the agent must

always maintain positive wealth. Second, I assume that agents can control their unemployment duration (d) deterministically by varying search effort, as in Moffitt and Nicholson (1982). Search costs, the leisure value of unemployment, and the benefits of additional search via improved job matches are all captured in a reduced-form manner by a concave, increasing function $\varphi(d)$. Third, I assume that the agent never loses his job again after he finds a new job, and supplies one unit of labor permanently after that point.

The agent receives an unemployment benefit of b while he is not working. The government finances the benefits by taxing the worker at a rate τ while employed, so his net-of-tax wage is $w = \widetilde{w}(1-\tau)$. To focus on the duration margin, assume that the probability of job loss does not vary with b.

Assuming the Inada condition $u_c(c=0) = \infty$, the technological constraints $c_s \geq 0$ will never bind and can be ignored in the maximization. Therefore, the agent chooses the path of c_s and d to

$$\begin{aligned} &\max \int_t^T u(c_s)ds + \psi(d) \\ &\text{s.t. } A_T = A_t + bd + w(T-d) - \int_t^T c_s ds = 0 \\ &\qquad A_s \geq 0 \forall s \in [t,T) \end{aligned}$$

where A_s denotes asset holdings at time s.

Since there is no uncertainty or discounting, and no income growth both when unemployed and employed, the optimal consumption path is flat in both states. Let c_u denote consumption while unemployed and c_e consumption while employed. Note that the only time the borrowing constraint could possibly bind is at the end of the unemployment spell. The agent's problem can therefore be rewritten as

$$\begin{aligned} & & \max du(c_u) + (T-d)u(c_e) + \psi(d) & \\ \text{s.t. } [\lambda]\ A_T &= A_t + bd + w(T-d) - dc_u - (T-d)c_e = 0 & (1) \\ [\mu]\ A_d &= A_t + bd - dc_u \geq 0 & (2) \end{aligned}$$

Let λ denote the multiplier associated with the intertemporal budget constraint (1) and μ the multiplier associated with the borrowing constraint (2). These multipliers represent the marginal value of relaxing each of the constraints at the optimum. Let $\Delta u = u(c_e) - u(c_u)$ denote the change in the flow utility of consumption from the unemployed to the employed period. The first order conditions for the agent's maximization problem are:

$$u'(c_u) = \lambda + \mu \tag{3}$$

$$u'(c_e) = \lambda \tag{4}$$

$$\varphi'(d) = (\lambda + \mu)(w - b) + (\lambda + \mu)(c_u - c_e) + \Delta u \tag{5}$$

The intuition for these optimality conditions can be seen with standard perturbation arguments. First consider the case where (2) does not bind. If the borrowing constraint is slack at the optimum, there cannot be any marginal value in loosening it further. Hence, $\mu = 0$ and $u'(c_u) = \lambda = u'(c_e)$. In this case, the optimality condition for the duration choice simplifies to $\varphi'(d) = \lambda(w - b)$. Intuitively, the marginal benefit of remaining unemployed one week longer should offset the marginal consumption utility loss of losing $w - b$ in income.

Now consider the case where the borrowing constraint (2) binds. Now the provision of an extra dollar of wealth at time t relaxes both the borrowing constraint and the intertemporal budget constraint, raising utility by $\lambda + \mu$. Since it is strictly optimal to consume that dollar immediately if the borrowing constraint is binding, the marginal utility of consumption while unemployed must equal the sum of these two multipliers. But additional wealth when employed does not relax the borrowing constraint, so $u'(c_e) = \lambda$. When the agent is constrained, the optimality condition for duration has additional terms because the agent exhausts his assets before finding a new job and thus consumption is not smooth over time: $c_u < c_e = w$.

Let $g(\cdot)$ denote the inverse of the $\psi'(d)$ function and define $Z = (\lambda + \mu)(w - b) + (\lambda + \mu)(c_u - c_e) + \Delta u$. Then the agent's unemployment duration can be written as

$$d = g(Z) \tag{6}$$

This equation is very similar to the Frisch labor supply expression obtained from intertemporal labor supply models (MaCurdy 1981; Blundell and MaCurdy 1999).[5] It differs from the standard Frisch expression only because of the borrowing constraint. In the unconstrained case, where $Z = \lambda(w-b)$, the agent's unemployment duration (or, equivalently, labor supply) is fully determined by the marginal utility of wealth, λ, and the net wage, $w - b$. As shown by MaCurdy, this representation for the optimal labor supply decision permits a transparent separation of wealth and substitution effects, because wealth effects affect behavior only by changing λ. I now use this observation to compute the wealth and income effects of UI benefits for unconstrained and constrained individuals.

2.1 Income and Wealth Effects of UI Benefits

Unconstrained Case. Consider an individual for whom (2) does not bind at the optimum ($\mu = 0$). An increase in b raises d for this individual through two channels. The first is the reduction in $w - b$, which creates an incentive to substitute toward leisure and raise d. This response to distorted incentives is the moral hazard effect of insurance provision. The second channel is the wealth effect: the wealth provided by an increase in b reduces λ. For unconstrained individuals, the wealth effect is negligible in practice. To see this, observe that the wealth effect of the UI benefit on duration is given by

$$\varepsilon^{\mathrm{W}}_{d,b}(\mu = 0) = \frac{\partial \log g}{\partial \log Z} \frac{\partial \log \lambda}{\partial \log W} \frac{\partial \log W}{\partial \log b}$$

Define $\delta = -\frac{\partial \log g}{\partial \log Z}$. Let $\gamma_W = -\frac{\partial \log \lambda}{\partial \log W}$ denote the elasticity of the marginal utility of wealth with respect to wealth, i.e. the coefficient of relative risk aversion of the value function over wealth. With this notation,

$$\varepsilon^{\mathrm{W}}_{d,b}(\mu = 0) = \delta \gamma_W \varepsilon_{W,b} \tag{7}$$

[5] Despite the similarity in the equations, the Frisch elasticity of labor supply and the elasticity of duration w.r.t. UI benefits may differ in magnitude. Job search and hours worked could respond differently to wage and benefit changes (Krueger and Meyer 2002).

In the aggregate, UI is a balanced-budget transfer program, and induces no change in lifetime wealth if behavior is unchanged. Higher benefits are fully offset by higher taxes. Hence, in a benchmark case with identical agents and no behavioral responses, $\varepsilon^{W}_{d,b} = \varepsilon_{W,b} = 0$. In an environment with heterogeneity, higher UI benefits can generate increases in net wealth for some individuals. To bound the magnitude of $\varepsilon^{W}_{d,b}$ in this case, consider the effect of an increase in UI benefits without any change in the UI tax. In this case, $\frac{\partial W}{\partial b} = 1$ and $\varepsilon_{W,b}$ equals the fraction of lifetime wealth accounted for by UI benefits. To obtain a rough estimate of this fraction, I use data on weeks of unemployment from the PSID for household heads followed from 1968 to 1998. Among individuals who report being unemployed at least once, the median number of weeks unemployed between 1968 and 1998 is 32.5 and the mean is 50. Since the wage replacement rate for UI is typically 50% and males work for roughly 40 years in the U.S., UI benefits account for approximately $\frac{0.5\times 50}{40\times 52} = 1.2\%$ of lifetime wealth. This calculation likely overstates $\varepsilon_{W,b}$ because (1) it ignores non-labor income, (2) not all weeks of unemployment are covered by UI, and (3) it ignores behavioral responses that lengthen unemployment durations and further reduce lifetime wealth. Hence $\varepsilon_{W,b} < 0.012$, i.e. doubling UI benefits permanently raises lifetime wealth by at most 1.2% for the mean job loser.

The small impact of UI benefits of lifetime wealth implies that UI has small wealth effects on unemployment durations for unconstrained agents. This can be established formally by bounding the other parameters in equation (7). To bound δ, let $\varepsilon_{d,b}$ denote the total elasticity of durations with respect to benefits, and recall that empirical studies of UI have found $\varepsilon_{d,b} \in (0.4, 0.8)$. Differentiating (6) w.r.t. b yields $\varepsilon_{d,b} > -\frac{\partial \log g}{\partial \log Z}\frac{b}{w-b}$, which implies $\delta = -\frac{\partial \log g}{\partial \log Z} < 1$ given that $\frac{b}{w} \simeq \frac{1}{2}$ in practice. Given a plausible value for the coefficient of relative risk aversion (e.g. $\gamma_W < 5$), it follows that $\varepsilon^{W}_{d,b} < 0.06$. Hence, a 10% increase in UI benefits raises duration by at most 0.6% via the wealth effect. The lifetime wealth effect thus accounts for a minor fraction of $\varepsilon_{d,b}$ for the typical unconstrained UI claimant, even in the extreme case where higher benefits are not offset at all by higher taxes.[6] It follows that moral hazard induced by distorted incentives accounts for

[6] Of course, individuals who are laid off very frequently, such as seasonal workers, might experience a significantly larger wealth effect from UI benefit changes. Although these responses do not arise from the liquidity constraint mechanism emphasized in this paper, they reinforce the general point that much of the UI-duration link could be due to income or wealth effects rather than substitution effects.

virtually all of the UI-duration link among unconstrained households.

Constrained Case. Now consider an individual for whom (2) binds, perhaps because he experienced shocks before period t that depleted his wealth or because he has a high discount rate and did not build up a large buffer stock.[7] An increase in b affects d through two channels for this agent as well. The first is the substitution effect, which again arises from the reduction in $w-b$. The second channel is the income effect: the liquidity provided by an increase in b reduces $\lambda+\mu$ and Δu. Using the Taylor approximation $\Delta u = (\lambda+\mu)(c_e - c_u)$ and recalling that $c_e = w$ when $\mu > 0$, the income effect component of $\varepsilon_{d,b}$ can be expressed as

$$\varepsilon^{\text{INC}}_{d,b}(\mu > 0) \simeq \frac{\partial \log g}{\partial \log Z}\frac{\partial \log \lambda+\mu}{\partial \log c_u}\frac{\partial \log c_u}{\partial \log b} \tag{8}$$

$$= \delta\gamma_c\varepsilon_{c_u,b} \tag{9}$$

where $\gamma_c = -\frac{\partial \log u'(c_u)}{\partial \log c_u}$ denotes the coefficient of relative risk aversion over consumption and $\varepsilon_{c_u,b}$ is the elasticity of consumption while unemployed with respect to benefits. To gauge the magnitude of $\varepsilon^{\text{INC}}_{d,b}$, first observe that $\gamma_c \geq \gamma_W$: since individuals can adjust labor supply and other margins over their lifetime, the curvature of indirect utility over wealth must be lower than the curvature of utility over consumption (Bodie, Merton, and Samuelson 1992). Therefore, $\frac{\varepsilon^{\text{INC}}_{d,b}}{\varepsilon^{\text{W}}_{d,b}} \geq \frac{\varepsilon_{c_u,b}}{\varepsilon_{W,b}}$. Empirical studies of consumption-smoothing have found that $\varepsilon_{c_u,b} \geq 0.2$ among constrained groups. It follows that the income elasticity of UI for constrained agents can be 10-20 times larger than the corresponding wealth effect for unconstrained agents. Thus, the income effect could potentially be as or more important than moral hazard in driving the UI-duration link among liquidity constrained individuals.

The intuition for this result is straightforward. When an agent has insufficient liquidity to smooth consumption relative to permanent income, raising the unemployment benefit level can have a large effect on consumption while unemployed. This reduces pressure to find a job to maintain consumption, creating the potential for a large income effect. In contrast, when agents are unconstrained, this channel is virtually shut down because UI benefits are a trivial fraction of

[7] The important question of why many job losers have virtually no assets is left to future research. In this study, I focus only on ex-post search behavior conditional on asset holdings, ignoring the reasons for initial asset choices.

lifetime wealth and have little impact on consumption while unemployed.

The simplicity of this intuition indicates why this result holds in more general environments. UI has an income effect on durations for similar reasons in a search model with liquidity constraints.[8] The result also holds with weaker assumptions about the nature of liquidity constraints. For example, suppose agents are unconstrained when they lose their jobs, but are uncertain about their unemployment durations and have a non-zero probability of hitting a borrowing constraint later in the spell. Carroll (1997) shows that consumption while unemployed can be highly sensitive to transitory income in such an environment because agents must be prepared for the possibility of becoming constrained. Hence, UI can have a substantial income effect for a broad group of individuals with limited liquidity, and not just those facing a binding constraint.[9]

Irrespective of the modelling details, the key point is that the fraction of the UI-duration link due to moral hazard is an empirical question when individuals have liquidity concerns. I provide evidence on this issue in the next section.

3 Empirical Analysis I: The Role of Constraints

3.1 Estimation Strategy

The model suggests a natural first step in evaluating whether liquidity constraints and income effects are empirically relevant in the UI-duration link: Estimate the effect of UI benefits on durations for constrained individuals ($\mu > 0$) and unconstrained individuals ($\mu = 0$) separately. If the UI-duration link is driven primarily by the $\mu = 0$ group, it would be implausible that income effects are important; but if the link comes from the $\mu > 0$ group, income effects might matter.

I implement this heterogeneity analysis by replicating the identification strategy of Moffitt (1985) and Meyer (1990) on various subsets of the data. In particular, I divide individuals into unconstrained and constrained groups and estimate equations of the following form using cross-state

[8] In a search model, reservation wages depend on the marginal utility of wealth and the level of unemployment benefits. The marginal utility of wealth becomes sensitive to the level of UI benefits for constrained individuals, creating an income effect on durations.

[9] In view of this point, I refer to any household that has a positive probability of becoming liquidity constrained while unemployed as "constrained" below.

and time variation in UI benefit levels:

$$\log d_{it} = \beta_0 + \beta_1 \log b + \beta_2 X_{i,t} + \theta_{i,t} \tag{10}$$

where $X_{i,t}$ denotes a set of covariates and $\theta_{i,t}$ denotes an idiosyncratic error. The key identifying assumption is the same as that underlying the Moffitt and Meyer estimates – the UI benefit rate must be orthogonal to the error in the estimating equation:

$$Eb \times \theta_{i,t} = 0 \tag{11}$$

Some evidence supporting this assumption is described in the next section.

For simplicity, I ignore the small lifetime wealth effect of a change in UI benefits for unconstrained individuals by assuming $\varepsilon^{\mathrm{W}}_{d,b}(\mu = 0) = 0$ below. This assumption leads us to slightly overstate the true magnitude of substitution effects and understate the magnitude of income effects. With this simplification, when (10) is estimated for the unconstrained ($\mu = 0$) group, the coefficient

$$\beta_1^{\mu=0} = \varepsilon_{d,b}^{\mu=0} = \varepsilon_{d,b}^{s,\mu=0}$$

gives the pure substitution effect of UI benefits on durations for unconstrained individuals.

When (10) is estimated for a group of constrained individuals ($\mu > 0$), we obtain

$$\beta_1^{\mu>0} = \varepsilon_{d,b}^{\mu>0} = \varepsilon_{d,b}^{s+I,\mu>0}$$

which is an estimate of the total effect of UI on durations for this group, including *both* substitution and income effects. Note that the composition of this elasticity in terms of the income and substitution components cannot be directly identified with the empirical strategy implemented in this section.

One might wonder why I focus on UI benefits to test whether liquidity affects unemployment durations, rather than using other forms of variation in wealth holdings at the time of job loss. The main reason is that the variation in UI benefits is credibly exogenous, insofar as it comes

from differences across states and time in laws. In contrast, wealth holdings at the time of unemployment are endogenous and correlated with other factors that could influence durations such as skills. Indeed, Gruber (2001) finds that agents with low levels of wealth also tend to have short job tenures and limited labor force experience, inducing a negative correlation between wealth and duration.[10]

Defining the constrained group. The main difficulty in implementing (10) is that μ is a latent variable, making it impossible to classify households into groups directly based on whether they are liquidity constrained. As shown by equation (9), an ideal definition of the constrained group would be the set of households whose consumption while unemployed is sensitive to transitory UI benefits. Unfortunately, the SIPP data used below lack information on consumption. I therefore use several proxies for being constrained that have been identified as predictors of liquidity constraints by existing studies of consumption behavior.

The primary proxy is liquid wealth net of unsecured debt, which I term "net wealth." Browning and Crossley (2001) and Sullivan (2005) find that households with no financial assets prior to job loss suffer consumption drops during unemployment that are mitigated by provision of UI benefits. In contrast, households with positive balances are able to smooth consumption and do not show excess sensitivity of consumption to UI benefits. This evidence suggests that households that report positive net wealth prior to job loss are likely to have $\mu = 0$, while the remainder of the households are more likely to face a liquidity constraint while unemployed.[11]

The second proxy is whether the individual has a spouse who is also working prior to job loss. Households that rely on a single income are more likely to be constrained when that individual loses his job; those with a second income source may have additional sources of liquidity, including better access to credit because at least one person is employed. The validity of this proxy is also substantiated by the analysis of Browning and Crossley, who find larger consumption drops and higher sensitivity to UI among single-earner households. The third proxy for μ is an indicator

[10] This endogeneity problem could explain why Lentz (2003) and others generally find little association between wealth holdings and unemployment durations in the cross-section.

[11] As a robustness check, I also proxy for constraints using net liquid wealth divided by pre-unemployment wage. This measure captures how much of the lost income each household can replace using its assets (Gruber, 2001). Results with this alternative definition (not reported) are very similar.

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 liquid wealth, and is more likely to be liquidity constrained than a renter.[12]

Though existing studies support the validity of these proxies for liquidity constraints, these markers are obviously imperfect predictors of who is constrained. Some households with $\mu = 0$ will be misallocated to the $\mu > 0$ group and vice-versa. Such classification error will pull the estimated elasticities for the two groups closer together, thereby causing us to *underestimate* the importance of liquidity constraints in the UI-duration link.

A concern in implementing (10) is that households may become constrained as an unemployment spell elapses. This issue does not arise in the model in section 2 because it assumes that households can anticipate their unemployment durations perfectly at the time of job loss. In practice, however, households update their expectations over time while depleting their buffer stocks. As a result, the probability that a household faces a binding constraint could rise as the spell elapses. Since asset data are available only once for each household, the only feasible way to account for this possibility is to estimate a model that permits a time-varying effect of UI benefits on unemployment exit rates within the ex-ante constrained and unconstrained groups. This issue, and more importantly the fact that many unemployment spells are censored in the data, motivates estimation of a hazard model with time-varying covariates rather than estimation of (10) using OLS. Letting $h_{i,s}$ denote the unemployment exit hazard rate for household i in week s of an unemployment spell and $X_{i,s}$ denote a set of controls, the primary estimating equation for the constraint tests is thus

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

[12] 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, 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.

3.2 Data

The data used to estimate (12) are from the 1985-1987, 1990-1993, and 1996 panels of the Survey of Income and Program Participation (SIPP). The SIPP collects information from a sample of approximately 30,000 households every four months for a period of two to three years. The interviews I use span the period from the beginning of 1985 to the middle of 2000. 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.

I make five exclusions on the original sample of job leavers to arrive at my core sample. First, following previous studies of UI, I restrict attention to prime-age males (over 18 and under 65). Second, I include only the set of individuals who report 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.[13] Fourth, I exclude individuals who have less than three months work history within the survey because there is insufficient information to estimate pre-unemployment wages for this group. Finally, I focus primarily 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. The potential sample selection bias related to UI takeup that arises from this exclusion is addressed below.

These exclusions leave 4,560 individuals in the core sample. Note that asset data is 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. The most

[13] Katz and Meyer (1990) show that whether an individual considers himself to be on temporary layoff is endogenous to the duration of the spell; recall may be expected early in a spell but not after some time has elapsed since a layoff. Excluding temporary layoffs can therefore potentially bias the estimates. To check that this is not the case, I include temporary layoffs in some specifications of the model.

germane statistic for the present analysis is pre-unemployment wealth: median liquid wealth net of unsecured debt is only $128.

The raw data on UI laws were obtained from the Employment and Training Administration (various years), and supplemented with information directly from individual states.[14] The computation of weekly benefit amounts deserves special mention. 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 approaches to proxy for each claimant's (unobserved) actual UI benefits, all of which yield similar results. First, I use published average benefits for each state/year pair 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. 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.[15] In the second stage, I use the predicted wage as a proxy for the true wage, and assign the claimant unemployment benefits using the simulation program.

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. I then show the robustness of these results to controls, sample selection, and other potential specification concerns. First consider the asset proxy for

[14] 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.

[15] 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.

constraints. I divide households into four quartiles based on their net liquid wealth. The median level of net liquid wealth is close to zero, suggesting (based on the Browning and Crossley evidence) that households in the lowest two quartiles are likely to be constrained, while those in the third and fourth quartiles are not. 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 may be 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 qualitative 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 ($\mu > 0$). 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.[16]

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

[16] 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.

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

An important assumption in this analysis is that the variation in UI benefits is orthogonal to unobservable determinants of durations, i.e. that (11) holds. To evaluate this identification 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.[17]

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.[18] 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

[17] Results are similar for the set of job losers who are ineligible for UI, who may be a better "control" because takeup of UI is endogenous. However, the UI-ineligible group consists of workers who have very low levels of earned income before unemployment and may differ from the average UI claimant.

[18] 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.

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: 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, pre-unemployment spousal work status, and being on the seam between interviews to adjust for the seam effect. Standard errors in this and all subsequent specifications are clustered by state.

The estimate in column 1 of Table 2a indicates that a 10% increase in the UI benefit rate reduces the hazard rate by 4% in the pooled sample. Reassuringly, this unconditional estimate is in the range found by 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 as in the graphical analysis. 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.[19] Among households

[19] The β_2^j coefficients on the time interactions are generally insignificant and do not exhibit any strong patterns in

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). The coefficients of interest are virtually unchanged when this rich set of covariates is introduced. The fact that controlling for observed heterogeneity does not affect the results suggests that the estimates are unlikely to be very sensitive to unobservable heterogeneity as well.

The preceding specifications maximize sample size by using data on post-unemployment assets for households where pre-unemployment asset data are unavailable. Specification (4) addresses the concern that post-unemployment asset levels are endogenous to durations by estimating (3) on the subsample of households who have pre-unemployment asset data. Since the sample size is reduced by more than 50%, the standard errors in this specification are larger. However, the pattern of the coefficients remains very similar to that in specification (3). The hypothesis that the effect of UI on exit rates of below-median wealth and above-median wealth households is the same can be rejected with $p = 0.05$.

Specification (5) introduces state fixed effects in addition to the full set of controls. In this model, the variation in the UI benefit level comes purely from within-state law changes. Results remain similar, with monotonically increasing β_1^j coefficients as wealth rises.

Table 2b reports a series of additional robustness checks. All of these specifications include the full control set. Column (1) shows that the estimates are similar when individuals with pre-unemployment annual wages above \$24,720, the 75th percentile of the wage distribution, are excluded from the sample. This specification addresses the concern that UI may have a smaller

this and subsequent specifications. In the interest of space, these coefficients are not presented in the tables.

effect in the high asset quartiles because it replaces a smaller fraction of income for high-income households.

Columns (2) and (3) examine robustness to changes in the definition of b_i. Column (2) uses the maximum UI benefit level in individual i's state/year and column (3) uses the simulated benefit for each individual i using the two-stage procedure described above. Both specifications give similar results to the baseline case. Column (4) shows that including individuals who report being on temporary layoff does not affect the results.

Finally, column (5) replicates the baseline specification 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 constraints play a role in the UI-duration link, 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 5.

As a robustness check, I predict assets based on age, income, education, and marital status. I then define net liquid wealth quartiles on the basis of predicted wealth and replicate the preceding analysis. Results from this procedure (not reported) are quite similar to those above. Individual predictors of assets also support the results. For example, the effect of UI benefits on durations is much stronger among individuals younger than the median age (31 years) – who have fewer assets on average – than among those above the median age.

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 income effects. If one were to assume that substitution effects are similar across unconstrained and constrained groups, this evidence would be sufficient to infer that income 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.[20] For this reason, I avoid identification of the income effect based on such cross-group comparisons, and turn to a second empirical strategy to investigate this issue.

[20] In Hurst's (2004) terminology, households with low liquidity may be myopic "grasshoppers," while unconstrained households may be "ants" who plan ahead. Comparing grasshoppers with ants could be problematic.

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 an income and substitution elasticity. The empirical strategy is based on 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 pure income effects and do not distort marginal incentives for unemployed agents. I estimate the income elasticity of unemployment durations 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:[21]

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

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, tests of the key orthogonality condition are discussed.

[21] Unfortunately, the data on severance amounts is too sparse to be useful.

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.[22] 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

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

(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.

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 liquidity-constrained households. In contrast, there is no clear reason to expect a differential effect of severance pay across constrained and unconstrained households under the alternative explanation described above. Hence, studying the heterogeneity of the severance pay effect provides a means of distinguishing between the causal and most natural omitted-variable hypotheses.

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 liquidity constraints 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. Recall that results based on predicted assets (using exactly the same prediction equation) and actual reported assets were very similar in the SIPP data: the total elasticity of duration w.r.t. UI benefits is large among households with predicted assets below the median, but near zero among households with predicted assets 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. As in the UI benefit analysis, 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 income effects drive much of the

UI-duration link.[23]

As a second approach to examining the causality of severance pay, I assess the sensitivity of the severance pay effect to a rich set of controls. 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. Since controlling for observed heterogeneity does not affect the results, the estimates are not likely to be very sensitive to unobservable heterogeneity either.

4.4 How Big is the Income Effect?

The severance pay estimates can be combined with the total elasticity estimates from the SIPP data to get a rough estimate how much of the UI-duration link is due to moral hazard vs. liquidity. Note that the median severance payment in the sample is approximately equal to 5 weeks of wages, which equals 10 weeks of UI benefits. The median unemployment duration is also around 10 weeks. Hence, receipt of severance pay is roughly equivalent to doubling UI benefits for the typical individual.

The estimate from specification 2 in Table 2 implies that a one log unit increase in the UI benefit reduces the hazard rate by 55% in the constrained groups (households below the median in net liquid wealth). Analogously, the estimate from specification 3 of Table 5 indicates that receipt of severance pay reduces hazard rates by 38% for constrained households. Therefore,

[23] 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.

roughly $38/55 = 70\%$ of the UI-duration link is attributable to an income effect for the constrained groups. In the unconstrained groups, both income and substitution effects appear negligible; the null hypothesis that both effects are zero cannot be rejected. If both elasticities are roughly zero for the unconstrained groups, the income effect still accounts for 70% of the total UI-duration elasticity in the pooled sample of all households. While this calculation suggests that UI benefits raise durations primarily by providing liquidity rather than by distorting marginal incentives, the 70% figure should be interpreted cautiously because it involves a comparison of point estimates in two different samples. Further work is necessary to pin down the magnitude of the income effect precisely.

5 Normative Implications

I now discuss some normative implications of the results. Given the stylized nature of the model analyzed here, these points should not be viewed as policy recommendations, but as a description of some of the issues that should be considered in subsequent normative work.

1. Moral hazard cost of insurance. 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.

To quantify the moral hazard cost, consider a policy that reduces the proportional benefit rate b in order to reduce d. Suppose that the saved UI benefits are returned as a lump-sum grant while the agent is unemployed, so that c_u does not fall. In Appendix A, I show that the marginal welfare gain of this policy change for an agent who is liquidity constrained is

$$MW(b) = \varepsilon^c_{d,b} du'(c_e) \tag{16}$$

where $\varepsilon^c_{d,b}$ is the substitution elasticity of the proportional benefit on duration. Equation (16)

shows that the efficiency gain from reducing the incentive distortion is proportional to $\varepsilon^c_{d,b}$, not the total elasticity $\varepsilon_{d,b}$. Similar expressions can be obtained for the marginal welfare gain from raising search costs or providing a job finding bonus when the accompanying taxes/subsidies are lump-sum. Intuitively, if the proportional UI benefit affects search behavior only through an income effect, changing marginal incentives would leave search behavior and social surplus unchanged. But if UI affects search behavior through a substitution effect, correcting marginal incentives without reducing consumption while unemployed would bring d closer to the first-best, raising welfare. Hence, only the substitution effect generates a moral hazard cost.

Using (16), the empirical evidence suggests that efforts to shorten durations would yield efficiency gains 70% smaller than suggested by studies that attribute the entire duration response to a substitution effect. However, there are some important caveats to this point. First, this result applies only to 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. UI may generate substantial deadweight burdens because of substitution effects on such margins. Second, the results apply only locally at the observed level of benefits (approximately 40-50% of pre-unemployment wages) in the U.S. If benefits were closer to full wage replacement, substitution effects could become much more important. Hence, these results cannot be generalized beyond a moderate benefit level.

2. *Optimal level and provision of UI.* As one might expect, the smaller moral hazard cost implies that the welfare-maximizing unemployment benefit level is higher. Recent studies have shown that the optimal benefit rate is higher when the consumption-smoothing benefit of UI is large (Chetty 2006) and when the sensitivity of reservation wages to the benefit level is high (Shimer and Werning 2006). In Chetty's formula, a large income effect implies a higher optimal benefit level because $\varepsilon^{\text{INC}}_{d,b}$ is proportional to the consumption-smoothing benefit of UI ($\gamma\varepsilon_{c_u,b}$), as shown in equation (9). In Shimer and Werning's formula, generating a large income effect would require more "hand-to-mouth" consumers, raising the sensitivity of reservation wages to benefits and again driving up the optimal benefit rate. Intuitively, the income elasticity estimate gives an additional moment to be matched when calibrating a model as in Acemoglu and Shimer (2000). Matching a large income effect would likely require tight liquidity constraints, which in turn yields a high

optimal benefit rate.[24]

While insuring unemployment may be beneficial, government intervention is not necessarily the best means of providing such insurance. Indeed, the empirical evidence in this paper suggests that liquidity is a greater concern for unemployed agents than "insurance" in the classic sense of pooling risk across households. Feldstein and Altman (1998) and Shimer and Werning (2005) point out that grants or low-interest loans (e.g. via privately held UI accounts) are a better means of providing liquidity than government transfers. In this sense, the results of this paper may ultimately favor the use of private insurance or liquidity mechanisms.

3. *Optimal path of UI benefits.* The provision of benefits late in an unemployment spell is valuable from an insurance perspective but has the drawback of magnifying incentive distortions (see Karni 1999 for a survey). Insofar as substitution effects are relatively small, an upward-sloping benefit path may be more desirable than previously thought. Calibrating models to match the income and substitution elasticities estimated here could be a fruitful direction for further work in this area.

4. *Means-Testing.* There is a controversial 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 because they perceive a lower private return to work (moral

[24] The income elasticity itself does not fully pin down the optimal benefit rate because it only reveals the marginal sensitivity of behavior to UI benefits. To calculate the optimal benefit rate, one must also determine how far current behavior is from the first-best, which introduces additional parameters into the welfare calculation.

hazard). The evidence documented here suggests a more benign view of the UI-duration link: unemployment durations rise mainly because households have more liquidity while unemployed, reducing the pressure to find work quickly.

It is important to reiterate that while the analysis of this paper suggests that the moral hazard cost of UI is smaller than previously thought, it does not make clear recommendations regarding government policy. In particular, we have not considered why the private sector cannot solve the unemployment insurance problem. Indeed, the finding that duration responses to insurance are primarily the consequence of liquidity rather than distorted incentives would apply equally well to private unemployment insurance contracts, and may ultimately favor private-market approaches to liquidity provision.

In future research, it would be interesting to identify the benefits of extending unemployment durations through the provision of liquidity. Does UI allow constrained individuals to find better job matches or does it facilitate additional consumption of leisure? Existing empirical studies of this question have generally failed to detect an effect of UI benefits on wages, but confidence intervals are large because of the high degree of noise in wage data. There is some evidence, however, that raising UI benefits lowers subsequent job turnover rates (Centeno 2004). Future work might be able to shed light on this issue by examining the effect of unemployment benefits on wage and turnover rates for individuals with low liquidity in a large dataset.

Finally, while this paper has focused on unemployment, the main conceptual point applies to a broad range of social and private insurance markets. For example, empirical studies have documented large responses to health insurance, disability insurance, workers compensation, and social security (see e.g. Krueger and Meyer 2002; Autor and Duggan 2003; Finkelstein 2005). Identifying the fraction of these responses that is due to a substitution effect could shed light on the importance of moral hazard in these insurance markets.[25]

[25] Nyman's (2003) recent theory of health insurance essentially proposes that income effects account for the large behavioral responses to health insurance documented in the empirical literature.

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, "The Rise In The Disability Rolls And The Decline In Unemployment," Quarterly Journal of Economics 118(1): 157-205, 2003.

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.

Carroll, Christopher (1997) "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," Quarterly Journal of Economics 112, 1-56.

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).

Finkelstein, Amy, "The Aggregate Effects of Health Insurance: Evidence from the Introduction of Medicare," NBER Working Paper 11619, 2005.

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: Efficiency cost of UI benefits

Consider a UI system that consists of two parts: a lump-sum grant (severance payment) g and a proportional benefit b. Both benefits are financed through a tax τ when the agent returns to work. The UI provider maintains a balanced-budget, so

$$\tau = bd + g$$

For an agent with $\mu > 0$, $c_u = \frac{A_t+g}{d} + b$ and $c_e = w - \frac{\tau}{T-d}$. Indirect utility as a function of b and g is

$$V(b,g) = \max_d du(\frac{A_t+g}{d} + b) + (T-d)u(w - \frac{\tau}{T-d}) + \psi(d)$$

To compute the efficiency cost of lowering the proportional UI benefit and raising the grant by an equivalent amount, I calculate $\frac{dV}{db}$ and $\frac{dV}{dg}$:

$$\begin{aligned} \frac{dV}{db} &= du'(c_u) - u'(c_e)[d + b\frac{\partial d}{\partial b}] \\ \frac{dV}{dg} &= u'(c_u) - u'(c_e)[1 + b\frac{\partial d}{\partial g}] \end{aligned}$$

Lowering b by \$1 reduces income while unemployed by \$$d$. The change in welfare from a \$1 reduction in b compensated by an increase in the lump-sum grant of \$$d$ is:

$$\begin{aligned} MW(b) &= d\frac{dV}{dg} - \frac{dV}{db} \\ &= bu'(c_e)[\frac{\partial d}{\partial b} - d\frac{\partial d}{\partial g}] \end{aligned}$$

Using a Slutsky decomposition, the substitution effect $\frac{\partial d^c}{\partial b}$ for duration (which equals one minus labor supply here) is given by

$$\frac{\partial d^c}{\partial b} = \frac{\partial d}{\partial b} - d\frac{\partial d}{\partial g} \tag{17}$$

It follows that

$$MW(b) = \frac{\partial d^c}{\partial b} bu'(c_e) = \varepsilon^c_{d,b} du'(c_e)$$

Appendix B: Measurement of unemployment durations in SIPP data

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. This appendix describes the method used to compute durations, which follows Cullen and Gruber (2000).

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. The duration of unemployment is computed 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 included as "searching" if they report ES = 4 at any point during their spell.

This method of computing durations results in a slightly different mean duration than that found in the CPS data. The mean spell in the SIPP lasts for 20.95 weeks before ending or being censored, whereas the US Department of Labor reports a mean duration of approximately 15 weeks. The official figure is computed from the length of ongoing spells for the cross-section of unemployed individuals who report they are looking for work in the CPS. The official definition therefore excludes the spells of individuals who become discouraged and stop searching for work. Unfortunately, these individuals cannot be identified in the SIPP because of the lack of reliable information on search behavior. At a weekly frequency, reports of job search are frequently interspersed with reports that the individual is not looking for a job; moreover, individuals often find jobs after reporting that they were not looking for one. Therefore, the only feasible measure of the length of an unemployment spell is to count the weeks from job separation to either job finding or censoring.

TABLE 1a

Summary Statistics by Wealth Quartile in SIPP Sample

		Net Liquid Wealth Quartile			
		1	2	3	4
	Pooled	(< -$1,115)	(-$1,115-$128)	($128-$13,430)	(>$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	Yes	No	Yes
	(0.63)	(0.37)	(0.55)	(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 2a

Hazard Model Estimates by Quartile of Net Liquid Wealth

	(1)	(2)	(3)	(4)	(5)
	Pooled	By Quartile of Net Liquid Wealth			
	Full cntrls	No cntrls	Full cntrls	Pre-wave Full cntrls	State FE's
log UI ben	-0.399				
	(0.197)				
Q1 x log UI ben		-0.794	-0.765	-0.756	-1.023
		(0.300)	(0.331)	(0.626)	(0.357)
Q2 x log UI ben		-0.710	-0.547	-0.803	-0.829
		(0.458)	(0.406)	(0.546)	(0.412)
Q3 x log UI ben		-0.165	-0.292	-0.287	-0.520
		(0.301)	(0.272)	(0.386)	(0.343)
Q4 x log UI ben		0.111	0.113	0.193	-0.081
		(0.344)	(0.288)	(0.392)	(0.402)
Q1=Q4 p-val		0.043	0.045	0.245	0.024
Q1+Q2=Q3+Q4 p-val		0.012	0.052	0.050	0.028
Observations	83834	81307	75739	35291	75739

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 two rows of table report p-values from F-test for equality of 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: 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, pre-unemployment spousal work status, and being on the seam between interviews to adjust for the seam effect. Spec 5 also includes state fixed effects.

Spec 4 includes only households for whom asset data is available prior to job loss. Spec 1 includes log UI ben, log UI ben interacted with weeks unemployed, and full control set reported above.

In all specs, log UI ben is defined as average UI benefit in claimant's state/year pair.

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

Number of observations equals total risk set (i.e., total number of unemployed weeks observed in the dataset).

TABLE 2b

Additional Hazard Model Estimates by Quartile of Net Liquid Wealth

	(1) Low-wage Full cntrls	(2) Maximum Benefits	(3) Actual Benefits	(4) Temp Layoffs	(5) Home Equity
Q1 x log UI ben	-1.133	-0.765	-0.479	-0.754	-0.325
	(0.342)	(0.331)	(0.177)	(0.321)	(0.628)
Q2 x log UI ben	-0.763	-0.547	-0.542	-0.539	-0.143
	(0.415)	(0.406)	(0.169)	(0.382)	(0.527)
Q3 x log UI ben	-0.119	-0.292	-0.522	-0.238	-0.344
	(0.255)	(0.272)	(0.145)	(0.264)	(0.616)
Q4 x log UI ben	0.211	0.113	0.048	0.145	-0.126
	(0.483)	(0.288)	(0.282)	(0.289)	(0.411)
Q1=Q4 p-val	0.028	0.045	0.127	0.032	0.784
Q1+Q2=Q3+Q4 p-val	0.003	0.052	0.174	0.041	0.990
Observations	52805	75739	75739	80574	29549

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 two rows of table report p-values from F-test for equality of coefficients across quartiles. All specifications 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). All specs also include 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, pre-unemployment spousal work status, and being on the seam between interviews to adjust for the seam effect. Spec 1. excludes households in the upper quartile of the wage distribution. Spec 2. proxies for UI ben using state/year maximums rather than averages. Spec 3. uses individual-level simulated benefits based on wage histories. Spec 4 includes temporary layoffs. Spec 5. defines asset quartiles by home equity holdings instead of net liquid wealth and includes only homeowners. In all specs, baseline hazards are stratified by net liquid wealth quartile. Number of observations equals total risk set (i.e., total number of unemployed weeks observed in the dataset).

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

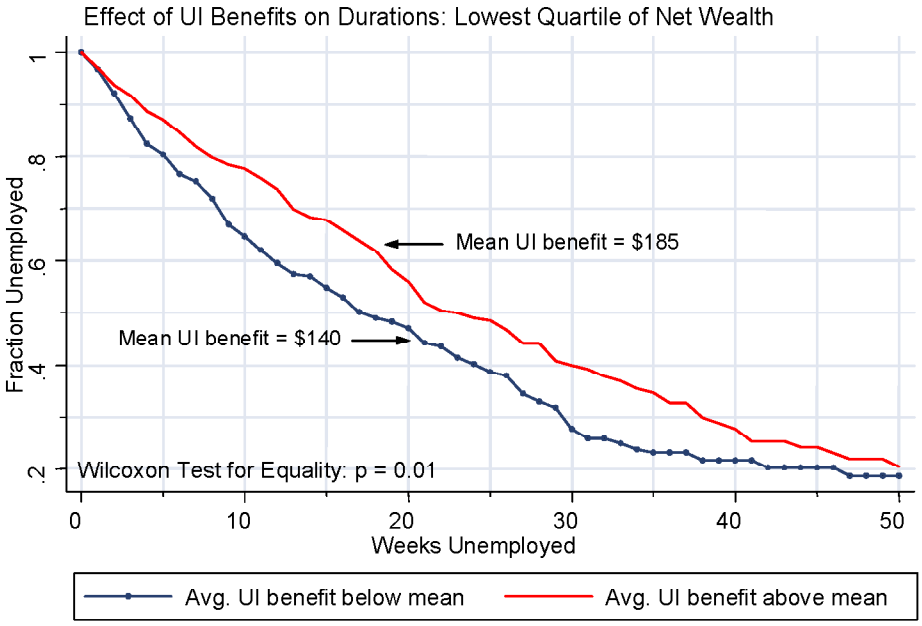
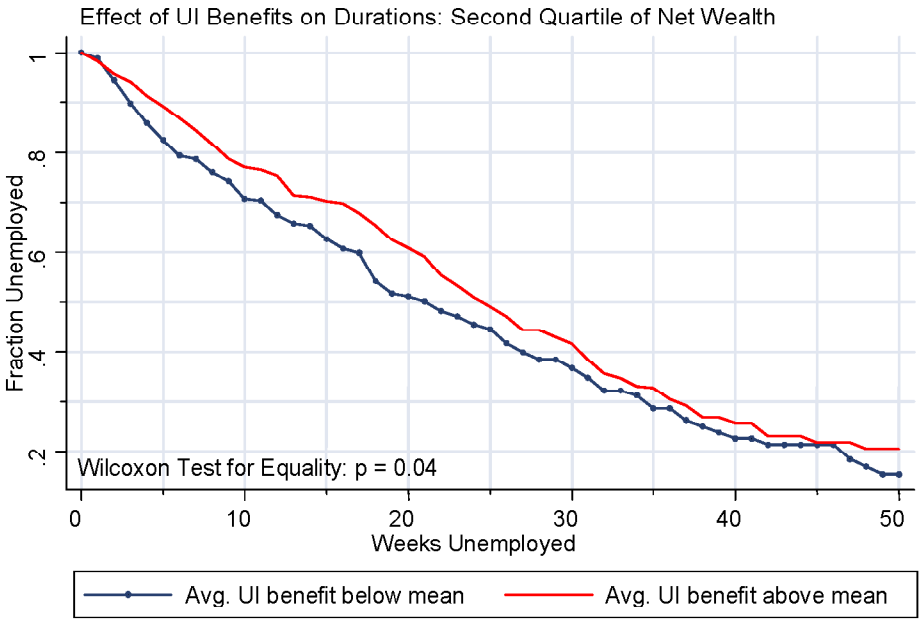


Figure 1b



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

Figure 1c

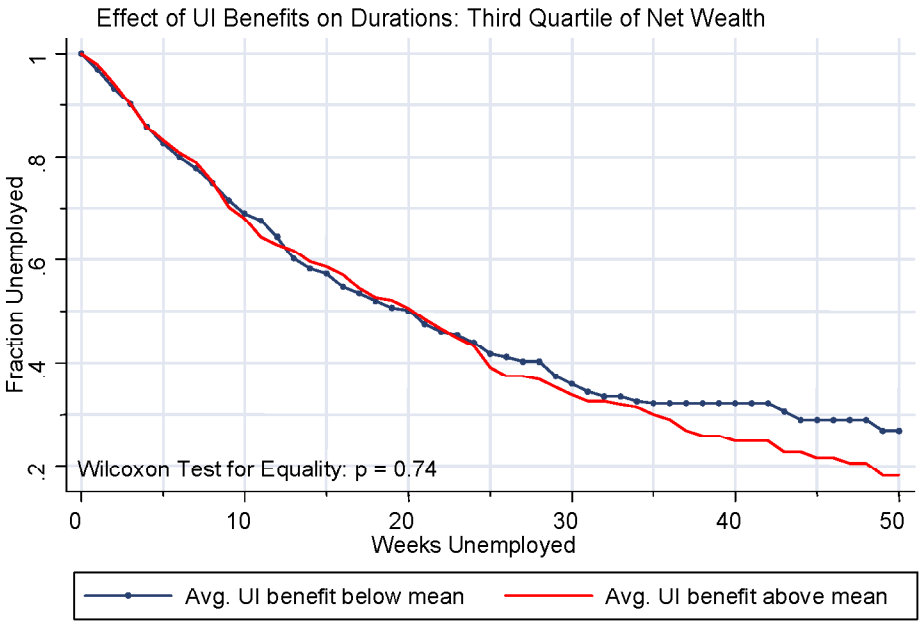
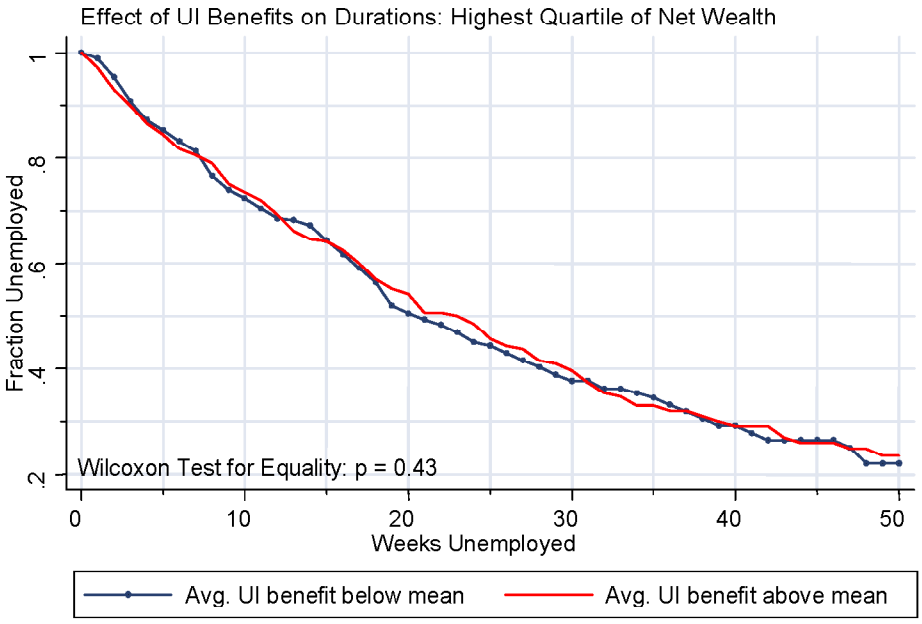


Figure 1d



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

Figure 2a

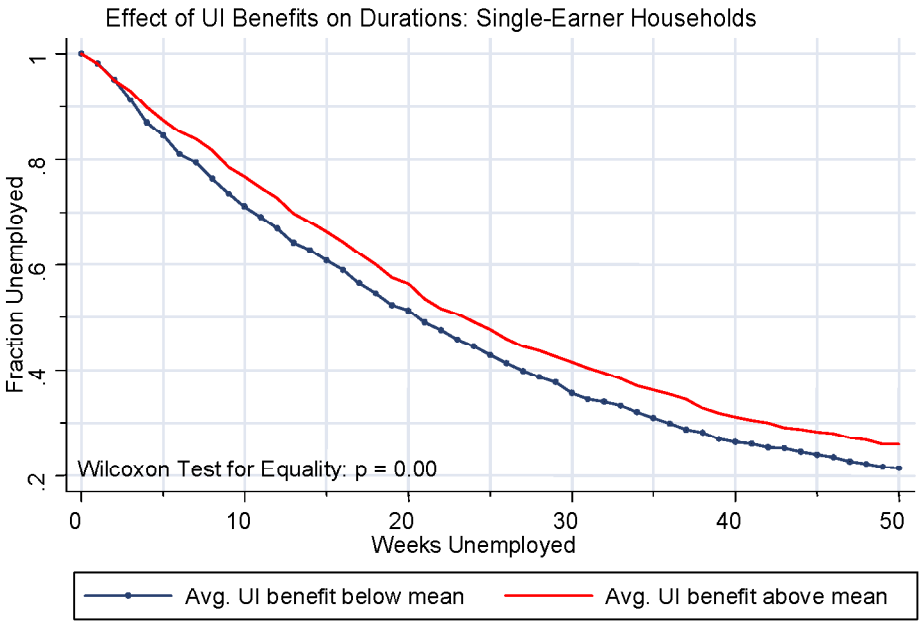
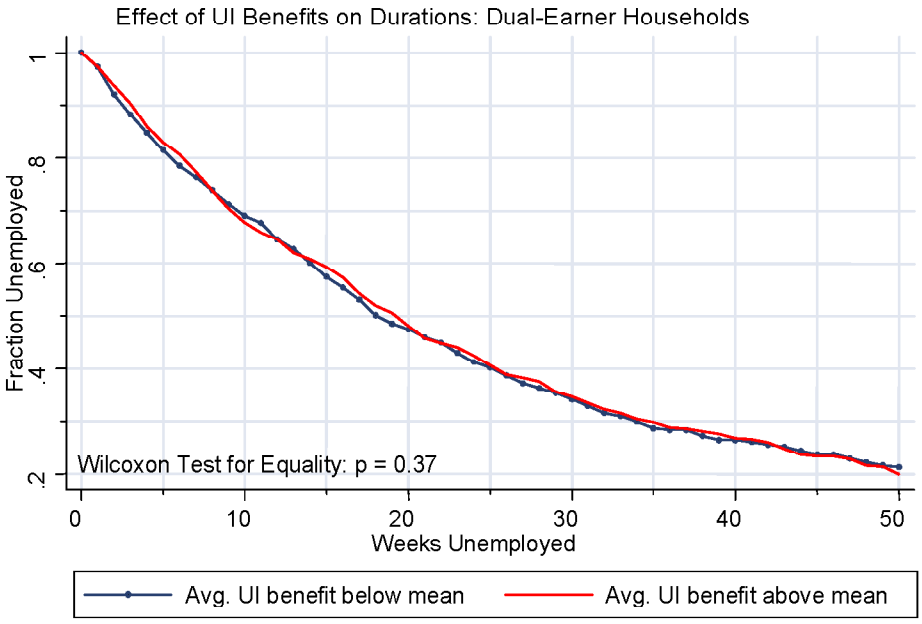


Figure 2b



NOTE–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

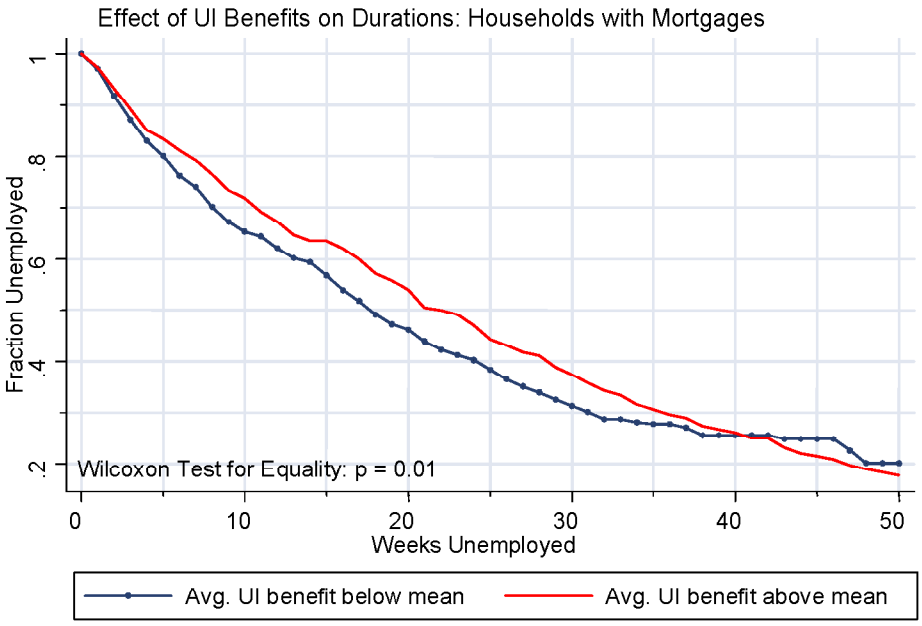
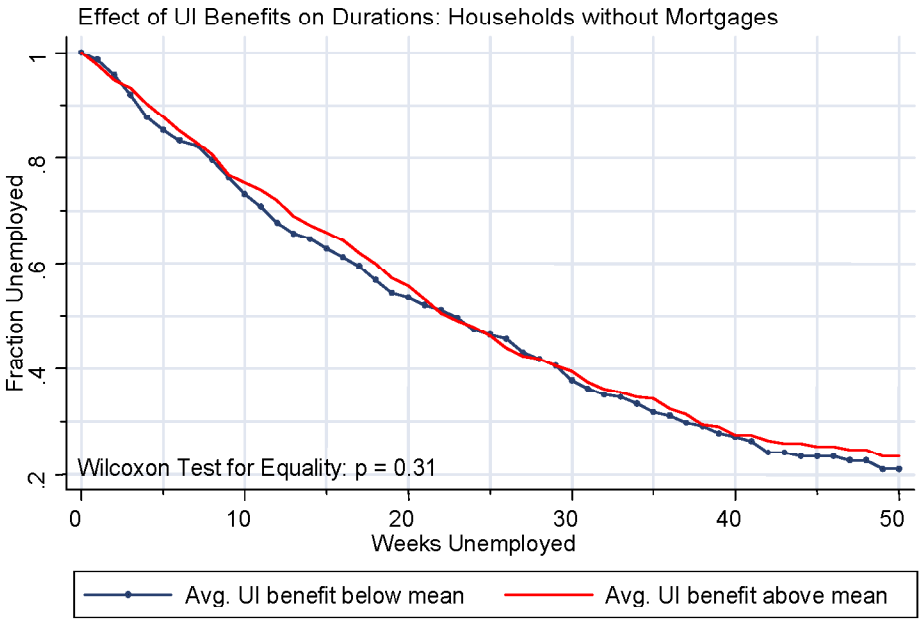
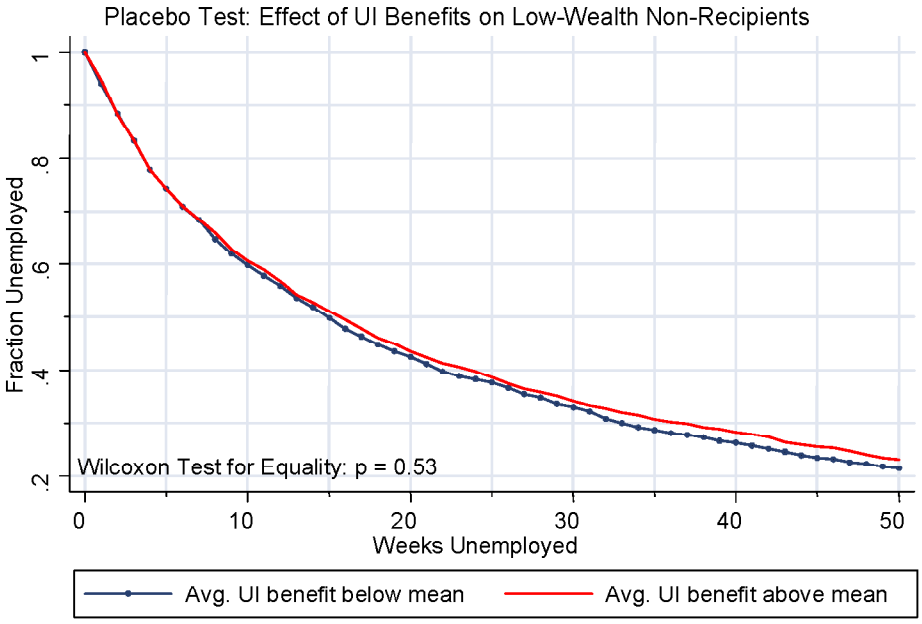


Figure 3b



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.

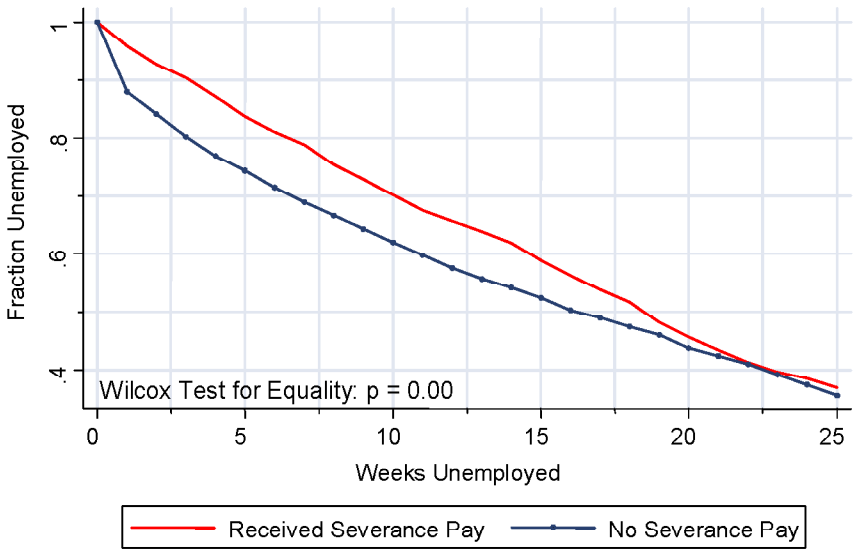
Figure 4



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

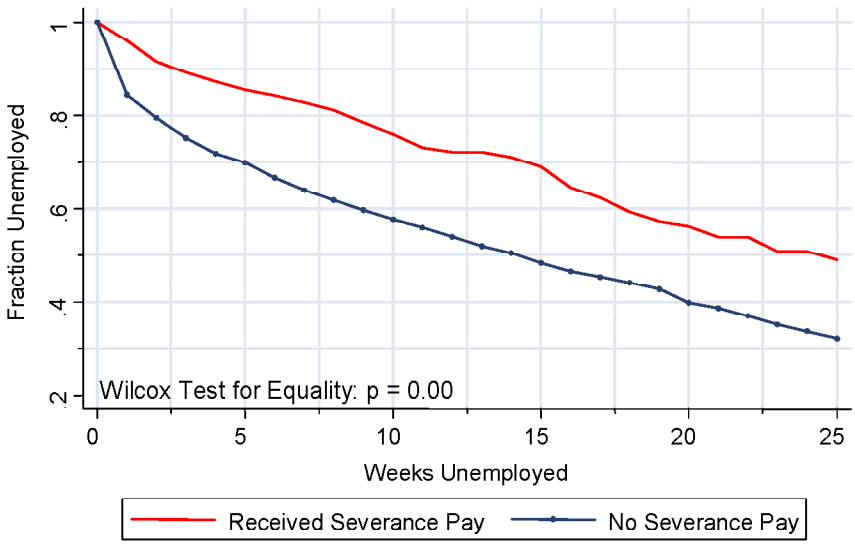


Figure 6b

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.