Review of Economic Studies (2016) **83**, 1092–1124 doi:10.1093/restud/rdw009

Advance access publication 18 February 2016

Should Unemployment Insurance Vary with the Unemployment Rate? Theory and Evidence

KORY KROFT
University of Toronto and NBER

and

MATTHEW J. NOTOWIDIGDO
Northwestern University and NBER

First version received January 2012*; final version accepted October* 2015 (*Eds.*)

We study how the marginal welfare gain from increasing the unemployment insurance (UI) benefit level varies over the business cycle. We do this by estimating how the moral hazard cost and the consumption smoothing benefit of UI vary with labour market conditions, which we identify using variation in the interaction of UI benefit levels with the unemployment rate within U.S. states over time. We find that the moral hazard cost is procyclical, greater when the unemployment rate is relatively low. By contrast, we do not find evidence that the consumption smoothing benefit varies with the unemployment rate. We use these empirical results to estimate the marginal welfare gain, and we find that it is modest on average, but varies positively with the unemployment rate.

Key words: Moral Hazard, Consumption Smoothing, Business Cycle, Optimal Unemployment Insurance

JEL Codes: H50, J64, J65

1. INTRODUCTION

During the Great Recession, expenditures on unemployment insurance (UI) benefits increased substantially, from \$33 billion in 2007 to \$94 billion in 2012.[1] This has triggered a debate among economists about the social costs and benefits of UI during recessions. Some argue that the social costs of UI are lower during recessions. As Alan Krueger and Bruce Meyer (2002, p. 64-65) remark:

> [F]or some programs, such as UI, it is quite likely that the adverse incentive effects vary over the business cycle. For example, there is probably less of an efficiency loss from reduced search effort by the unemployed during a recession than during a boom. As a consequence, it may be optimal to expand the generosity of UI during economic downturns ... Unfortunately, this is an area in which little empirical research is currently available to guide policymakers.

1. See https://www.cbo.gov/publication/43734.

Similarly, the Congressional Budget Office writes that the availability of long-term unemployment benefits “could dampen people’s efforts to look for work, [but that concern] is less of a factor when employment opportunities are expected to be limited for some time” (Elmendorf, 2010, p. 12). On the benefit side, Piketty and Saez (2012, p. 65) write:

> In recessions, the ability to smooth consumption might be reduced, as the long-term unemployed might exhaust their buffer stock savings and might face credit constraints. This implies that the gap in social marginal utility of consumption between workers and non-workers might grow during recessions, further increasing the value of redistributing from workers to the unemployed.

This article makes two main contributions to the UI literature. First, this article is the first to empirically investigate how the social marginal benefit and social marginal cost associated with the UI benefit level vary over the business cycle. Secondly, this article uses the job search model introduced by Lentz and Tranaes (2005) and further developed by Chetty (2008) to theoretically examine how the benefits and costs of UI vary over the business cycle.

We derive a standard formula for the marginal welfare gain from increasing the UI benefit level that illustrates the classic trade-off between consumption smoothing and moral hazard. The consumption smoothing term jointly depends on the consumption drop upon unemployment and the coefficient of relative risk aversion, and the moral hazard term is given by the elasticity of unemployment duration with respect to the benefit level. We depart from the prior literature by explicitly allowing these reduced-form parameters to depend on the unemployment rate. Identifying the relationship between these reduced-form parameters and the unemployment rate is therefore sufficient to characterize the marginal welfare gain over the business cycle.[2]

The welfare gain formula structures our two-part empirical strategy. In both parts, we use publicly available survey microdata from the U.S., combined with rich variation within states over time in both UI benefit levels and labour market conditions.[3] The first part examines how the duration elasticity varies with the state unemployment rate, our primary proxy for labour market conditions. We estimate Cox proportional hazard models allowing the effect of the UI benefit level on the exit rate from unemployment to vary with the state unemployment rate. Our findings indicate that the duration elasticity is approximately 0.6 at the average state unemployment rate, similar to Chetty (2008). Our new empirical result is that the duration elasticity varies with local labour market conditions; specifically, we find that the duration elasticity is statistically significantly lower when the state unemployment rate is relatively high. Furthermore, the magnitude of this interaction effect is economically large: in our preferred specification, a one standard deviation increase in the unemployment rate (an increase of 2.3 percentage points from a base of 6.2%) reduces the magnitude of the duration elasticity to roughly 0.3 (a decline in magnitude of roughly 50%). Although our precision is somewhat limited, we show that our results are fairly robust and qualitatively similar across a range of alternative samples, specifications, and measures of our key independent variables.

The second part of our empirical strategy builds on Gruber (1997) by estimating how the consumption smoothing benefit of UI varies with the unemployment rate. We regress the

2. Kiley (2003), Sanchez (2008), and Andersen and Svarer (2011) theoretically explore optimal UI over the business cycle in partial equilibrium job search models without savings.

3. Identification comes from exploiting variation in UI benefits within states over time interacted with within- and between-state variation in the unemployment rate. We pursue this time-series, cross-sectional research design using MSA and state unemployment rates rather than a purely time-series design using the national unemployment rate in order to have sufficient variation in UI benefit levels across a wide range of labour market conditions. This motivation is very similar to the motivation in Autor *et al*. (2013), who estimate the impact of trade on labour markets, Aguiar *et al*. (2013), who estimate the impact of unemployment on time use patterns during recessions, and Mian and Sufi (2012), who examine household debt and the impact of the 2007–2009 recession.

consumption drop upon unemployment on the UI replacement rate, allowing for an interaction between the replacement rate and the state unemployment rate. Overall, we find that a 10 percentage point increase in the UI replacement rate reduces the average consumption drop upon unemployment by 2.7%, which is similar to the preferred estimates in Gruber (1997). In contrast to our duration elasticity results, we do not find evidence that the consumption smoothing benefit of UI varies with the unemployment rate. Our estimate of the consumption smoothing interaction effect is both economically and statistically insignificant, and though our statistical power is limited, we can rule out a large interaction effect.

While these results extend the influential work of Gruber (1997), according to our marginal welfare gain formula, it is only necessary to estimate the average consumption drop upon unemployment at current benefit levels. To see intuition for this, consider a dollar transfer from the employed to the unemployed, ignoring behavioural responses. A government with a utilitarian objective values this transfer as simply the difference between the marginal utility of consumption for the unemployed and the employed, respectively. Under the specific assumptions we discuss in Section 2, this difference can be approximated by the consumption drop upon unemployment, scaled by the coefficient of relative risk aversion.[4] In particular, it is not necessary to identify the causal effect of UI benefits on the consumption drop. We believe this is an important (and under-appreciated) conceptual point because the data requirements for estimating the consumption drop upon unemployment are much less demanding. Our empirical results indicate that the mean consumption drop is 6.9% at the mean state unemployment rate and that this is slightly larger in magnitude when the unemployment rate is high, although our estimates are imprecise.

In the last part of our article, we use our estimates for the duration elasticity and consumption drop to calibrate the marginal welfare gain from increasing the UI benefit level. We find that it is negative on average and modest in magnitude, implying that reducing UI benefits from current levels would be welfare-improving. This conclusion holds across a wide range of assumptions regarding the level of risk aversion. Our new result is that the marginal welfare gain varies with the unemployment rate. In particular, at high levels of unemployment, we estimate a positive marginal welfare gain.

Since our welfare gain formula is stated in terms of reduced-form parameters, our analysis is in the spirit of the "sufficient statistics" approach discussed in detail in Chetty (2009). The primary advantages of this approach are that it is simple and transparent to empirically implement and it does not place restrictions on the model primitives. For example, our welfare analysis is valid for a wide range of underlying mechanisms that cause the duration elasticity and the consumption drop upon unemployment to vary with the unemployment rate. The primary disadvantages of this approach are that it is not well suited to estimating welfare effects for large (non-marginal) policy changes, nor is it well suited for conducting counterfactual policy simulations. As a result of this, we focus throughout the article on estimating the *marginal* welfare gain (as opposed to the globally optimal UI benefit level). This corresponds to the welfare effects of small changes in the UI benefit level, relative to the current UI benefit level. This analysis implicitly assumes that the reduced-form parameters we estimate are approximately constant over the range of relatively small policy changes considered in our calibrations.

4. It is important to emphasize that the consumption drop alone is not sufficient to characterize the welfare gain of UI. For instance, if workers engage in costly ways to smooth consumption, the consumption drop may be minimal, but the welfare effect of UI benefits may be substantial. Rather, one needs to jointly identify both the consumption change *and* the coefficient of relative risk aversion. This argument is made formally in Chetty and Looney (2006). We do not estimate risk aversion directly but instead calibrate the marginal welfare gain formula for alternative values of the coefficient of relative risk aversion.

An important conceptual point for our welfare analysis is whether the duration elasticity corresponds to a microeconomic or macroeconomic elasticity. The "micro-elasticity" measures the effect of an increase in benefits for a small randomly selected subset of individuals within a labour market on unemployment durations. As such, it accounts only for changes in worker search effort and job acceptance decisions in response to changes in benefits. By contrast, the "macro-elasticity" measures the effect of an increase in benefits, for all individuals, on unemployment durations. It also captures both wage and vacancy responses to changes in benefits. Theoretically, as we discuss below, our marginal welfare gain formula is valid under the assumption that the micro- and macro-elasticities are equal. This is very much in the spirit of research on optimal income taxation in economies with search frictions that abstract from inefficiencies by imposing the Hosios condition (Hungerbühler *et al.*, 2006). On the other hand, if the elasticities are not equal, then the formula depends additionally on the ratio of the macro- to micro-elasticity. In this case, a welfare analysis requires estimates of *both* micro- and macro-elasticities, as shown in Landais *et al*. (2014).

Our empirical strategy, which relies on variation in unemployment benefits across local labour markets, defined by state-year cells, comes closest to measuring the macro-elasticity. This is because the state-year variation in unemployment benefits affects all unemployed individuals in a given state and so plausibly captures general equilibrium responses.[5] We interpret our evidence as suggesting that the macro-elasticity declines with state unemployment rates and is consistent with crowding during recessions. Additionally, as long as the micro- and macro-elasticities are approximately equal, then one may use our empirical estimates of the duration elasticity to calibrate our marginal welfare gain formula, and one does not need to take a stand on whether our estimates correspond to micro- or macro-elasticities.

Our article builds on and contributes to several recent studies that have empirically examined the moral hazard cost of UI extensions during recessions. Schmieder *et al*. (2012) find that in Germany, the impact of UI extensions on unemployment durations is moderately lower during recessions. Consistent with this evidence, Rothstein (2011) and Farber and Valletta (2013) examine the impact of UI extensions on unemployment durations in the U.S. during the Great Recession and find that the responses are relatively small.[6] While these studies examine partial equilibrium (microeconomic) behavioural responses on the worker side since they hold labour demand and wage responses constant, several recent empirical papers focus on the general equilibrium (macroeconomic) response to changes in UI benefits. The findings are mixed and depend on the setting and research design. Hagedorn *et al*. (2013) estimate a large impact of UI extensions on job creation during the Great Recession using a border design to isolate variation in UI policy between U.S. states. Their results suggest large macroeconomic effects of UI on employment and unemployment. By contrast, Lalive *et al*. (2015) exploit a large policy reform in Austria and find substantial negative job search spillovers onto workers who are ineligible for UI, especially when labour market conditions are poor. This is consistent with the results in Crépon *et al*. (2013) who find similar evidence of "crowding" in areas of France with high unemployment. Both of these papers imply that the macroeconomic effects of UI are smaller than the microeconomic effects, especially in recessions. Lastly, Johnson and Mas (2015) analyse a sharp change in UI policy in the U.S. using both administrative data and survey data, and they find results that suggest similar microeconomic and macroeconomic responses to UI. The assumption

5. However, one must exercise caution in interpreting our results as capturing all of the relevant macroeconomic effects of UI since our empirical analysis excludes several subgroups from the labour market who may be directly or indirectly affected by UI policy, such as women, individuals ineligible for UI, and labour market non-participants.

6. Moffitt (1985) and Jurajda and Tannery (2003) also examine the business cycle effects of changes in the potential duration of unemployment benefits.

of equal microeconomic and macroeconomic effects of UI that underlies both our theoretical model and welfare analysis is most consistent with the results in Johnson and Mas (2015).[7]

The remainder of the article proceeds as follows. The next section develops the search model and describes our sufficient statistics approach. Section 3 presents our empirical analysis that estimates how the duration elasticity and consumption smoothing benefit of UI vary with the unemployment rate. Section 4 considers the welfare implications of our empirical findings. Section 5 concludes.

2. THEORY

In this section, we present a standard discrete-time, finite-time horizon, job search model following Lentz and Tranaes (2005) and Chetty (2008). There are two main purposes of the model. First, we use the model to analytically characterize the marginal welfare gain from UI in terms of reduced-form parameters. Secondly, we numerically simulate the model to explore how the marginal welfare gain varies with labour market conditions. We limit the focus here to the setup of the model and a discussion of the intuition underlying the main theoretical results, and we present detailed proofs in the Online Appendix.

2.1. *Agent's Problem*

Consider a single worker who lives for T periods $\{0,1,...,T-1\}$. The interest rate and the agent's time discount rate are assumed to be zero. The individual is unemployed at $t=0$ with exogenous, predetermined assets A_0. When unemployed, the individual receives unemployment benefits b in each period that are available for a maximum of B periods. If the individual is employed in period t, he works for $T-t$ periods and earns a net wage $w-\tau$, where w is the gross wage, which we assume to be fixed, and τ is a lump-sum tax, which finances unemployment benefits. Each period, we assume that the individual receives non-labour income z irrespective of his employment status.[8] Let c_t^u denote consumption in period t if unemployed, and let c_t^e denote consumption of an employed worker who finds a job in period t.[9] We model liquidity constraints by assuming that the individual cannot deplete assets below $L(<0)$ at any time.

In each period t, there is some probability that the unemployed individual finds a job. The individual exerts costly effort, e_t, to search for jobs, and we assume a separable cost of search, denoted by $\psi(e_t)$. The job-finding rate depends on search effort according to the function $\lambda(e_t,\alpha)$, where $\frac{\partial\lambda}{\partial e}\geq 0$, $\frac{\partial^2\lambda}{\partial e^2}\leq 0$, $\frac{\partial^2\lambda}{\partial e\partial\alpha}\geq 0$ and $\frac{\partial\lambda}{\partial\alpha}\geq 0$. This formulation generalizes Chetty (2008), who imposes the normalization that $\lambda(e_t,\alpha)=e_t$, which implicitly assumes that the return to search does not vary with labour market conditions. We generate variation in labour market conditions through the parameter α, which we assume is constant within an unemployment spell. For notational simplicity, we sometimes suppress the arguments of $\lambda(e_t,\alpha)$ and refer to it as λ_t.

Worker preferences over consumption are represented by a standard utility function $u(c)$, which we assume to be strictly increasing and concave. The value function of finding a job at

7. Also related is Marinescu (2015), who does not find evidence that UI affects vacancy creation, using data from a large online job board. This is also consistent with limited macroeconomic effects of UI.

8. We include exogenous non-labour income in the model to produce consumption drops upon unemployment that are realistic in the simulations.

9. Note that since the interest rate and the discount rate are equal, consumption is flat during an employment spell because this model does not allow for separations (and therefore there is no role for precautionary savings). If the agent finds a job in period t, consumption will be equal to c_t^e for each period $\{t,t+1,...,T-1\}$.

the beginning of period t, $V_t(A_t)$, and the value function of not finding a job at the beginning of period t, $U_t(A_t)$, are given, respectively, by:

$$V_t(A_t) = \max_{A_{t+1} \geq L} u(A_t - A_{t+1} + w + z - \tau) + V_{t+1}(A_{t+1})$$
$$U_t(A_t) = \max_{A_{t+1} \geq L} u(A_t - A_{t+1} + b + z) + J_{t+1}(A_{t+1})$$

where

$$J_t(A_t) = \max_{e_t} \lambda(e_t, \alpha) V_t(A_t) + (1 - \lambda(e_t, \alpha)) U_t(A_t) - \psi(e_t)$$

is the value of entering period t without a job. It is straightforward to show that the optimal effort choice in this model solves the following first-order condition:

$$\psi'(e_t) = \frac{\partial \lambda_t}{\partial e_t} (V_t(A_t) - U_t(A_t)) \tag{1}$$

which simply equates the marginal cost (left-hand side) and the marginal benefit (right-hand side) of additional search effort.

We conclude with several definitions that are useful for the results that follow. Let $S_t \equiv \prod_{i=0}^{t} (1 - \lambda_i)$ be the survival function at time t (with $S_0 = 1$) and $f_t \equiv \prod_{i=0}^{t} (1 - \lambda_i) \lambda_t = S_{t-1} \lambda_t$ be the probability that the unemployment spell lasts exactly t periods or (equivalently) the probability the employment spell is $T - t$ periods. Let $D = \sum_{t=0}^{T-1} S_t$ be the worker's expected unemployment duration and $D_B = \sum_{t=0}^{B-1} S_t$ be the worker's expected benefit duration (i.e. the length of time a worker collects unemployment benefits). We define the elasticity of unemployment duration and benefit duration with respect to the UI benefit level, respectively, as $\varepsilon_{D,b} \equiv \frac{d \log D}{d \log b}$ and $\varepsilon_{D_B,b} \equiv \frac{d \log D_B}{d \log b}$. Let $\mu_t^u \equiv \frac{S_t}{D_B}$ and $\mu_t^e \equiv \frac{f_t(T-t)}{T-D}$, and let $\overline{c_u} \equiv \sum_{t=0}^{B-1} \mu_t^u c_t^u$ be the weighted-average consumption over B periods of unemployment starting in period 0 and $\overline{c_e} \equiv \sum_{t=0}^{T-1} \mu_t^e c_t^e$ be the weighted-average consumption over T periods during employment starting in period 0. Finally, let $u \equiv \frac{D}{T}$ be the unemployment rate.

2.2. *Social Planner's Problem*

In this section, we derive the marginal welfare gain from a change in the benefit level b in a given labour market state (i.e. at a given value of α), taking the maximum duration of UI benefits, B, as given. We then focus on how the welfare gain varies with α. The social planner's problem is to maximize the worker's expected utility at time 0 subject to a balanced-budget condition and agent optimization. Budget balance requires that the planner set taxes so that expected UI benefits paid equals expected lump-sum taxes collected (i.e. $D_B b = (T - D)\tau$). Agent optimization requires that the first-order condition for optimal search effort in equation (1) is satisfied. For given values of b and τ, the agent's indirect utility at the beginning of an unemployment spell is defined as J_0.

The following proposition and corollary characterize the (money-metric) marginal welfare gain of increasing benefits by \$1 subject to budget balance, which we label as dW/db. We compute this as the ratio of the welfare effect of a \$1 increase in b (dJ_0/db) to the effect of a \$1 increase in w (dJ_0/dw), or $dW/db = \frac{dJ_0}{db} / \frac{dJ_0}{dw}$.

Proposition 1. *If the borrowing constraint is not binding at time B, the money-metric welfare gain of raising b by* \$1 *is given by*

$$\frac{dW}{db}=\frac{D_B}{T-D}\left\{\frac{\sum_{t=0}^{B-1}\mu_t^u u'(c_t^u)-\sum_{t=0}^{T-1}\mu_t^e u'(c_t^e)}{\sum_{t=0}^{T-1}\mu_t^e u'(c_t^e)}-\left(\varepsilon_{D_B,b}+\varepsilon_{D,b}\frac{D}{T-D}\right)\right\} \tag{2}$$

Corollary 1. *If the coefficient of relative prudence is zero ($\rho=-\frac{u'''(c)}{u''(c)}c=0$) and the duration elasticities are equal ($\varepsilon_{D,b}=\varepsilon_{D_B,b}$), then dW/db can be approximated by*

$$\frac{dW}{db}\approx\frac{D_B}{D}\frac{u}{1-u}\left\{\gamma\frac{\Delta c}{c}-\frac{\varepsilon_{D,b}}{1-u}\right\} \tag{3}$$

where $\gamma=-\frac{u''(c)}{u'(c)}c$ is the coefficient of relative risk aversion, $\frac{\Delta c}{c}=\frac{\overline{c_e}-\overline{c_u}}{\overline{c_e}}$ is the consumption drop upon unemployment, and $\overline{c_e}$ and $\overline{c_u}$ are the weighted-average consumption of the employed and unemployed, respectively.

Proof See Online Appendix. ||

Equations (2) and (3) present the exact and approximate formulas, respectively, for the welfare gain from UI.[10] In equation (3), the product of the consumption drop, $\frac{\Delta c}{c}$, and the coefficient of relative risk aversion, γ, capture the social marginal benefit of UI.[11] The duration elasticity, $\varepsilon_{D,b}$, captures the social marginal cost or moral hazard of UI. Moral hazard arises because agents do not fully internalize the social planner's budget. Setting $dW/db=0$ in equation (2) delivers the "Baily–Chetty formula" for the optimal level of UI (Baily, 1978; Chetty, 2006).[12]

The marginal welfare gain formula in equation (2) is derived under the assumption that w and λ_t are constant with respect to a change in unemployment benefits. In this case, the duration elasticity, $\varepsilon_{D,b}$, measures the micro-elasticity. From this standpoint, the model in this section is a micro-elasticity model. An alternative approach is to allow changes in unemployment benefits to affect wages and job creation and hence the job-finding rate.[13] In this case, the wedge between

10. We also derive a more precise approximate formula in the Online Appendix, and we use this formula in the simulations and in alternative calibrations. This alternative formula allows for the coefficient of relative prudence to be non-zero, but assumes that the coefficient of variation in consumption when unemployed is zero and the duration elasticities are approximately equal ($\varepsilon_{D,b}=\varepsilon_{D_B,b}$). In this case, dW/db can be approximated by

$$\frac{dW}{db}\approx\frac{D_B}{D}\frac{u}{1-u}\left\{\gamma\frac{\Delta c}{c}\left(1+\frac{1}{2}\rho\frac{\Delta c}{c}\right)-\frac{\varepsilon_{D,b}}{1-u}\right\}.$$

11. In a richer utility function with preferences over consumption and leisure, the formulas in equations (2) and (3) remain valid as long as the marginal utility of consumption does not depend on leisure choice. If there is no such separability, then the formulas need to be modified to account for the strength of complementarity between consumption and leisure.

12. Although our model assumes fixed wages, our formula for the marginal welfare gain remains valid in a richer model with stochastic wage offers and endogenous reservation wage choices. The intuition for this is that since the reservation wage is an optimized variable that the agent controls, the marginal welfare gain still depends on the gap in expected marginal utilities, as a result of the envelope theorem.

13. See Landais *et al*. (2014) for such an approach.

the micro- and macro-elasticities enters the welfare formula, as an additional term, along with the macro-elasticity (which replaces the micro-elasticity). When the micro- and macro-elasticities are equal, the formula in the general case collapses to equation (2). It is important to note therefore that our formula is valid even when wages and vacancies are endogenous with respect to the benefit level; the key issue is whether the micro- and macro-elasticities are equal. To see this more clearly, consider a Diamond (1982), Mortensen and Pissarides (1999), and Pissarides (2000) matching economy with proportional bargaining. An increase in benefits leads to higher unemployment due to reduced search effort (micro-effect). This acts to increase the bargained wage and thus the wage, offsetting some of the unemployment (wage spillover effect). A higher wage, however, leads to lower vacancy posting (labour demand spillover effect). The macro-effect is the sum of all three effects. If the wage effect and labour demand effect just offset, then the macro-elasticity equals the micro-elasticity; in this case, the Hosios condition is satisfied. From this perspective, one can interpret our welfare formula as holding in an efficient labour market. We will calibrate our welfare reform using our empirical estimates below that correspond more closely to the macro-elasticity. In the case where the micro- and macro-elasticities are similar, our welfare analysis is valid. On the other hand, if the micro- and macro-elasticities are far apart, the formula in equation (3) is no longer valid and one must incorporate the elasticity wedge.

Previous research on consumption smoothing assumes that $\frac{\Delta c}{c}(b) = \alpha_0 + \delta_1 b$ and estimates α_0 and δ_1 (Gruber, 1997; Browning and Crossley, 2001). The parameter α_0 is the implied consumption drop in the absence of UI, and the parameter δ_1 is the effect of UI on the consumption drop and is interpreted as the consumption smoothing effect of UI. Equation (3) shows that computing dW/db requires identifying the mean consumption drop at existing benefit levels, $\frac{\Delta c}{c}(\bar{b})$. This is given by $\alpha_0 + \delta_1 \bar{b}$ or equivalently by the sample mean of $\frac{\Delta c}{c}(\bar{b})$.[14] This shows that, in principle, to calculate the marginal welfare gain, it is not necessary to identify δ_1. The intuition for this is simple. Consider transferring one dollar from the employed to the unemployed. According to the envelope theorem, behavioural responses have no first-order effect on utility, other than through the fiscal externality on the government's budget, which is captured by $\varepsilon_{D,b}$ in equation (3). Therefore, we can assume that a dollar increase in unemployment benefits will be spent solely on consumption while unemployed when computing welfare changes and express the welfare gain purely in terms of marginal utilities which in turn may be approximated by the consumption drop scaled by the coefficient of relative risk aversion. We believe this is an important (and under-appreciated) conceptual point because the data requirements for estimating the mean consumption drop are much less stringent than the data requirements for estimating δ_1. In particular, one does not need data on (or exogenous variation in) unemployment benefits in order to estimate the consumption smoothing benefits of UI.

Of course, to solve for the globally optimal level of UI benefits, one needs to know the mapping between $\frac{\Delta c}{c}(b)$ and $\varepsilon_{D,b}$ and b. However, we do not attempt to solve for the globally optimal benefit level, b^*. Doing so would require estimating how the sufficient statistics in equation (3) vary with the benefit level, and we do not have enough statistical power to carry out this exercise in our empirical analysis. We think this insight is general and applicable to other forms of social insurance, such as disability insurance and health insurance. In the empirical section below, we therefore focus on estimating the average consumption drop upon unemployment, and use this estimate when calibrating the marginal welfare gain.

Our expression for the marginal welfare gain also emphasizes that the appropriate consumption measures correspond to the *weighted-average* consumption, where the weights μ_t^u and μ_t^e depend on the unemployment duration distribution (and therefore implicitly on the unemployment

14. We use the sample mean, $\bar{b}$, in our calibrations, which is similar to the current benefit level.

survival function). To the extent that consumption falls over the unemployment spell, this distinction is important. Moreover, our formula emphasizes that even if consumption does not vary over the business cycle, the consumption drop will vary if the weights vary over the business cycle. Intuitively, in a recession, unemployment durations increase and so there will be relatively more weight put on (the lower) consumption of the long-term unemployed. This will act to increase the consumption drop. This relates to results in Schmieder *et al*. (2012) who derive a formula for the marginal welfare gain from a small benefit extension using the same model. They show that the welfare benefit from an extension of UI benefits depends on the rate of benefit exhaustion and an appropriate measure of the gap in consumption between the unemployed at exhaustion and the employed. This is intuitive since unemployed workers only benefit from additional weeks of UI if they are going to be unemployed during these weeks. Since the exhaustion rate increases in recessions, their welfare formula implies that, all else equal, it is more beneficial to extend unemployment benefits in a recession.

Lastly, the welfare gain is defined at a given value of α. The next section considers simulations of the model to show how this varies with labour market conditions.

2.3. *Numerical Simulations of Model*

We simulate the model numerically and report results in Figure 1 through Figure 3. The goal is to evaluate how the marginal welfare gain varies with labour market conditions by studying how the duration elasticity and consumption drop vary with labour market conditions. We also assess the quality of the approximation formula in equation (3).

The details of the simulations are summarized in the notes to the figures. In all of the simulations that follow, we use the unemployment rate, u, as a proxy for labour market conditions, and we use variation in α (which affects the return to search effort through $\lambda(e,\alpha)$) in order to generate variation in it. All of the simulations assume that $\lambda(e,\alpha)=e\cdot\alpha$.[15] Finally, the simulations implicitly capture the "budget effect" in the sense that changes in benefits are constrained to be revenue-neutral, so that the planner adjusts taxes to balance the budget in expectation.[16]

Figure 1 plots the unemployment duration elasticity against the unemployment rate, which ranges from roughly 3%–8%. The figure shows that the duration elasticity is strongly decreasing in the unemployment rate, consistent with the speculation of Krueger and Meyer (2002) discussed in the Introduction. The model generates this strong relationship due to the strong complementarity between search effort and labour market conditions: the return to search effort is higher when labour market conditions are strong.[17]

Figure 2 plots the relationship between the consumption drop upon unemployment and the unemployment rate. The consumption drop is fairly low on average (around 5–10%, similar to

15. In the Online Appendix, we consider a standard matching model with unemployed workers, vacancies and search effort, following Pissarides (2000). We show in this standard model that tougher labour market conditions (i.e. a higher unemployment-vacancy ratio) make it harder to raise the job-finding rate through search effort. This is consistent with the cross-restrictions imposed on the job-finding rate in our model. We thank the editor for pointing out this connection to us.

16. It can be shown that if the planner can run deficits in bad times and surpluses in good times and balance the budget across labour market states, the budget effect disappears.

17. In Online Appendix Figure OA.2, we plot the simulated duration elasticity against the log of the unemployment rate. The figure shows that the relationship is close to linear. This is part of our motivation for our preferred empirical specification in Section 3 which assumes that the duration elasticity varies linearly with the log of the unemployment rate. Additionally, in Online Appendix Figure OA.3, we show the duration elasticity for high and low (initial) assets, which mirrors the results in Chetty (2008) that highlight the importance of liquidity constraints in determining the duration elasticity. Lastly, in Online Appendix Figure OA.1, we show that the unemployment duration elasticity and benefit duration elasticity are similar in magnitude and vary similarly with the unemployment rate.

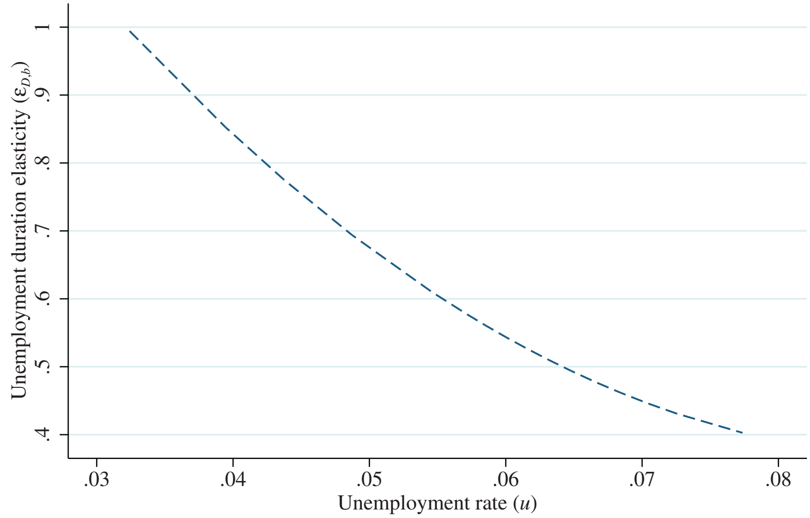

FIGURE 1

Unemployment duration elasticity and labour market conditions

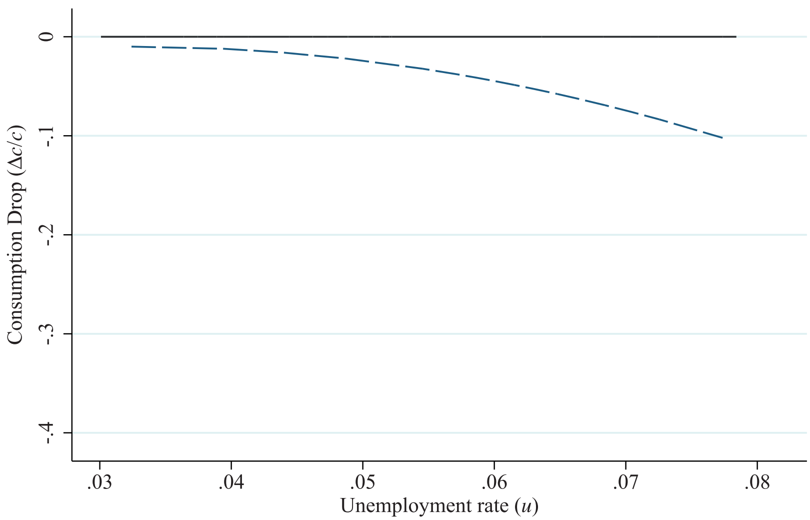

FIGURE 2

Consumption drop and labour market conditions

Notes: The figures above are generated by calibrating the job search model in the main text. The model is a finite-horizon, discrete-time dynamic program, with the time period in weeks. The total time is $T = 180$ weeks, and UI benefits are available for the first $B = 26$ weeks of unemployment. The benefit level is set to $b = \$170$ (per week). The fixed wage of employed workers is set to $w = \$340$ (per week). There is no discounting and the interest rate is set to 0. The job offer arrival rate is $\lambda(e, \alpha) = \alpha \cdot e$, where α is a parameter that affects the return to search, and e is endogenous search effort. The cost of search is given by $\phi e^{1+\kappa}/(1+\kappa)$, with $\phi = 2$ and $\kappa = 0.02$. The coeffcient of relative risk aversion (in the individual's CRRA utility function) is set to $\gamma = 2$, and the assets at $t = 0$ at the start of the unemployment spell are $A_0 = \$5000$. The individual can never have assets below $L = -\$1000$, which is an ad hoc liquidity constraint. The individual receives exogenous unearned income of $0.11w$ each period. The lump sum tax τ is set to balance budget in expectation and is only paid when the individual is employed. Choosing alternative values of α in the simulation generates variation in the unemployment rate ($u = D/T$), which is shown on the x-axis. In Figure 1, the y-axis shows the elasticity of the expected unemployment duration with respect to the UI benefit level (i.e. $\varepsilon_{D,b}$). In Figure 2, the y-axis shows the (weighted-average) consumption change between employment and unemployment (i.e. $\Delta c/c$).

Gruber, 1997). This comes from the initial assets of the agent ($A_0 = \$5000$) and the ability of the agent to borrow to finance consumption. The other notable feature of this figure is that the consumption drop is greater in magnitude when the unemployment rate is high. Intuitively, in a recession, the unemployed expect to be unemployed longer. This drives consumption down during the unemployment spell since the unemployed need to finance consumption out of a combination of savings, borrowing, and UI benefits. Additionally, in the simulation the unemployed are more likely to exhaust their UI benefits and face liquidity constraints when the unemployment rate is high, consistent with the speculation of Piketty and Saez (2012) discussed in the Introduction. As a result of these forces, the gap between consumption when employed and unemployed widens as the unemployment rate increases. When initial assets are very low, these constraints are even more likely to bind, so the consumption drop varies even more with the unemployment rate (Online Appendix Figure OA.4). Similarly when there are negligible borrowing constraints, the relationship between consumption drop and the unemployment rate is weaker (Online Appendix Figure OA.15).[18]

Figure 3 computes the marginal welfare gain using the approximation formula in equation (3) and the exact numerical solution, defined as $\frac{dJ_0}{db} / \frac{dJ_0}{dw}$. Note that the latter is not the expression given in equation (2). In particular, it is valid even when the borrowing constraint binds at time $t = B$. The figure shows that the marginal welfare gain is increasing in the unemployment rate, which is to be expected given that Figures 1 and 2 show that the duration elasticity is declining in the unemployment rate and the consumption drop upon unemployment is increasing with the unemployment rate. Overall, the exact and approximate curves are very close, although a small gap emerges at high unemployment rates. Part of the gap is approximation error, since the formula in equation (3) ignores third-order and higher-order terms, which understates the true consumption smoothing benefit. Online Appendix Figure OA.5 shows that including higher-order terms closes some of this gap. However, part of the gap is also due to the fact that as the unemployment rate increases, the probability that the individual hits the binding liquidity constraint increases. This is another source of approximation error, since equation (3) is derived from equation (2) which assumes that the liquidity constraint is not binding at benefit exhaustion.[19] Despite the existence of these two sources of error, we interpret the results in Figure 3 as suggesting that the approximation formula is reasonably accurate over a wide range of labour market conditions. Having demonstrated that the approximation formula for the marginal welfare gain is valid, we use it to structure our two-part empirical analysis in the next section.

3. EMPIRICAL ANALYSIS

The theoretical model highlights that the unemployment duration elasticity ($\varepsilon_{D,b}$) and the consumption drop at unemployment ($\Delta c/c$) may vary with labour market conditions. This structures our two-part empirical strategy: in part 1 we estimate how the duration elasticity varies with the unemployment rate, and in part 2 we estimate how the consumption drop varies

18. The simulations in Figure 2 assume that initial assets do not vary over the business cycle; however, it is possible that the unemployed have lower assets during a recession, perhaps because asset values tend to be low when the unemployment rate is high. We explored this possibility by assuming that initial assets at unemployment decline with the unemployment rate. For low levels of unemployment, the consumption drop is small, and as the unemployment rate increases, assets decline and the consumption drop increases. Compared with the case where initial assets are assumed to be constant over the business cycle, the consumption drop falls much more sharply with the unemployment rate in this case.

19. The assumption that liquidity constraints are not binding at the time benefits are exhausted appears elsewhere in the literature (see, e.g. Landais, 2014), but we have not seen a quantitative analysis of the importance of this assumption for estimating the marginal welfare gain, which is what we provide in Figure 3.

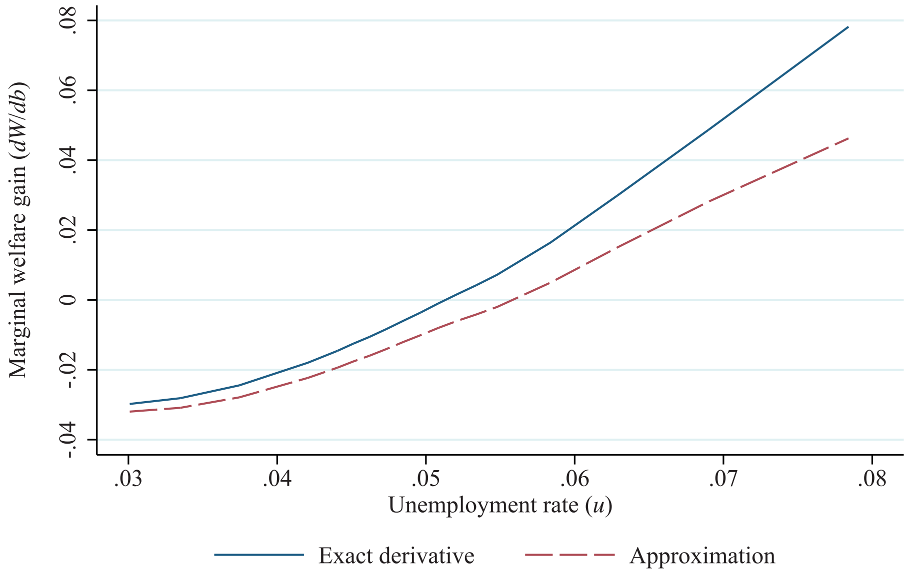

FIGURE 3

Marginal welfare gain and labour market conditions

Notes: This figure is generated by calibrating the job search model in the main text with the parameters described in the notes to Figures 1 and 2, except that the assets at the start of the unemployment spell are $A_0 = \$500$. The y-axis shows the marginal welfare gain (dW/db) of increasing the UI benefit level by \$1 and the x-axis shows the unemployment rate. The solid line shows the exact numerical derivative (i.e. $(dJ_0/db)/(dJ_0/dw)$), while the dashed line shows the second-order approximation based on the equation in footnote 10, which extends equation (3) in the main text.

with the unemployment rate. In both parts, we make three assumptions. First, we assume that the *predetermined* unemployment rate in the month at the start of an unemployment spell is a valid proxy for α. Using the predetermined unemployment rate—as opposed to the actual unemployment rate at a given time during an unemployment spell—partially addresses the concern that the unemployment rate is endogenous to the UI benefit level. Second, we assume that the unemployment rate is constant within an unemployment spell. This assumption is motivated by the fact that almost all of the variation in unemployment rates is across-spell variation, with negligible within-spell variation. Last, we rely on variation in unemployment rates between and within states, which implicitly assumes that the relevant local labour market conditions can be proxied by the state-level unemployment rate. We pursue this time-series, cross-sectional research design in order to have sufficient variation in UI benefit levels across a wide range of labour market conditions.

3.1. *Data*

The first part of the empirical analysis estimates how the duration elasticity varies with the unemployment rate. We present two pieces of evidence: (1) graphical evidence and non-parametric tests of survival curves and (2) semi-parametric estimates of proportional hazard models (Cox models). The empirical strategy closely follows Chetty (2008), extended to exploit cross-state variation in labour market conditions. We use unemployment spell data from the Survey of Income and Program Participation (SIPP) spanning 1985–2000. We impose the following sample restrictions: we focus on prime-age males who (1) report searching for a job, (2) are not on temporary layoff, (3) have at least three months of work history, and (4) claimed UI benefits. We

TABLE 1

Descriptive statistics for the unemployment duration and consumption samples

	Full sample		State unemp. rate < Median		State unemp. rate > Median		p-value of difference in means
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	
Panel A: Unemployment duration sample (SIPP)							
Unemployment duration (weeks)	18.51	14.35	16.95	13.61	19.37	14.68	0.000
Average UI weekly benefit amount ($)	163.33	26.80	163.08	26.07	163.46	27.21	0.660
Maximum UI weekly benefit amount ($)	226.93	45.74	219.57	45.63	231.00	45.30	0.000
Age	37.17	11.07	36.59	11.11	37.48	11.03	0.011
Years of education	12.17	2.88	12.12	2.87	12.20	2.88	0.372
1{Married}	0.62	0.49	0.61	0.49	0.62	0.49	0.501
Pre-unemployment wage income ($000s)	20.92	13.57	20.93	13.55	20.92	13.58	0.979
Number of unemployment spells	4307		2774		1533		
Panel B: Consumption sample (PSID)							
Change in the log of food consumption at unemployment	−0.073	0.504	−0.081	0.485	−0.065	0.521	0.492
After-tax UI replacement rate	0.57	0.16	0.56	0.15	0.57	0.16	0.276
Age	34.80	10.40	34.65	10.23	34.93	10.56	0.546
Years of education	11.49	2.67	11.51	2.72	11.46	2.62	0.685
1{Married}	0.72	0.45	0.70	0.46	0.73	0.44	0.133
1{White}	0.56	0.50	0.55	0.50	0.57	0.49	0.405
1{Black}	0.37	0.48	0.40	0.49	0.35	0.48	0.022
Number of children in household	1.18	1.35	1.16	1.35	1.20	1.35	0.470
Change in the log of food needs	−0.01	0.30	−0.02	0.29	−0.01	0.31	0.819
Number of observations	2003		962		1041		

Notes: In Panel A, the data are individual-level unemployment spells from the 1985–2000 SIPP data set. Average UI weekly benefit amount and maximum UI weekly benefit amount and all other dollar values are reported as 2000 CPI-U-adjusted dollars. In Panel B, the data are individual-level observations from the 1968–1997 PSID data set. The after-tax UI replacement rate is constructed using the UI benefit calculator in Gruber (1997). The sample for both data sets is restricted to men only. See the main text and the Online Appendix for more details.

also censor unemployment spells at 50 weeks. Due to the difficulty of constructing a precise measure of each individual's actual benefit level, we use the average benefit level for each state-year and the (statutory) maximum weekly benefit amount in the state-year in our baseline specifications. The maximum weekly benefit amount is the primary source of policy variation in average UI benefit levels across states. All dollar values in the data are adjusted to real dollars using the 2000 CPI-U series. The descriptive statistics for the SIPP sample are presented in Panel A of Table 1. The table presents summary statistics for the overall sample and subsamples divided into high and low unemployment rates. The two subsamples are broadly similar, although unemployed individuals are slightly older in states with high unemployment rates.

The second part of our empirical analysis replicates and extends previous work on the consumption smoothing benefit of UI (Gruber, 1997). Specifically, we estimate how the effect of UI on the consumption drop upon unemployment varies with the state unemployment rate in the previous year. The empirical strategy uses the after-tax UI replacement rate and the change in total food consumption as a proxy for the change in total consumption. On the surface, using food consumption rather than a broad-based consumption measure may seem restrictive. However, from a normative perspective, it is without loss of generality to use a single category of consumption (such as food). If agents are making optimal consumption choices when both employed and unemployed, then as long as the curvature of utility *over food* is used, along with the consumption smoothing elasticity for food consumption, one can conduct a valid welfare analysis (Chetty, 2006). For this part of the article, we use data from the Panel Study of Income Dynamics (PSID) between 1968 and 1997.[20] We impose the same sample restrictions as in Gruber (1997), focusing on all heads of household who are employed at interview date $t-1$ and unemployed at date t, and we define individuals as unemployed if they are looking for a new job and are not on temporary layoff. We make one additional restriction, which is to focus on prime-age men as in the SIPP sample above, so that the duration elasticity estimates and consumption smoothing estimates are estimated for a similar sample. We exclude observations if any element of food consumption is imputed or there is more than a 3-fold change in total food consumption. We present descriptive statistics for our PSID sample in Panel B of Table 1.

3.2. *Background on UI Benefits in the U.S.*

The UI program in the U.S. is a federal program, but both benefit levels and durations are set by each state, and states can freely adjust these parameters over time. In the Online Appendix, we provide a detailed overview of the UI program and how benefits are determined. In this section, we present evidence that the source of variation in unemployment benefits that we use in both parts of the empirical analysis is plausibly exogenous.

Figure 4 shows that changes in the average statutory maximum UI benefit level (averaged across states each year) are highly correlated with the national unemployment rate. When the national unemployment rate is high, states are (on average) more likely to increase UI benefit levels. This creates a potential for endogeneity bias, because national unemployment rates are also likely correlated with latent factors that affect both consumption changes and expected unemployment durations. Because of this concern, we interact UI benefit levels with year fixed effects and use these interactions as controls in all specifications below. This means that any variation in UI benefits that is the result of unobserved common shocks affecting all states in a given year is not used to identify the primary interaction term of interest, between UI benefits and the state unemployment rate.

Even with these controls, however, there is still the possibility that UI benefits at the state level are endogenous to state labour market conditions. We discuss this concern in detail below, but we do not believe it is a source of substantial bias. The reason for this is summarized in Figure 5, which shows that changes in state UI benefit levels are not significantly correlated with changes in state unemployment rates (conditional on year fixed effects).

3.3. *Part 1: Duration Elasticity*

To investigate how the duration elasticity varies with the unemployment rate, we estimate a set of Cox proportional hazard models. All results report standard errors clustered by state, and in our main results we use the log of the state unemployment rate. The decision to use log unemployment rates is based on the simulation results presented above, which showed that the duration elasticity is approximately linear in the log of the unemployment rate; this also follows Bertrand (2004).[21]

20. We extend the end of the sample from 1987 (which is the endpoint in Gruber, 1997) to 1997 in order to be closer to the 2000 endpoint in the SIPP sample. In the Online Appendix we extend to 2007 in both the SIPP and the PSID. For the PSID this introduces substantial complexity because the survey is only conducted every other year instead of every year after 1997. We therefore restrict to 1968–1997 to focus on the time period where we always have annual data to analyze.

21. To be clear, our objective is not to identify a supermodular hazard specification; rather, we want to consider a specification that is consistent with it.

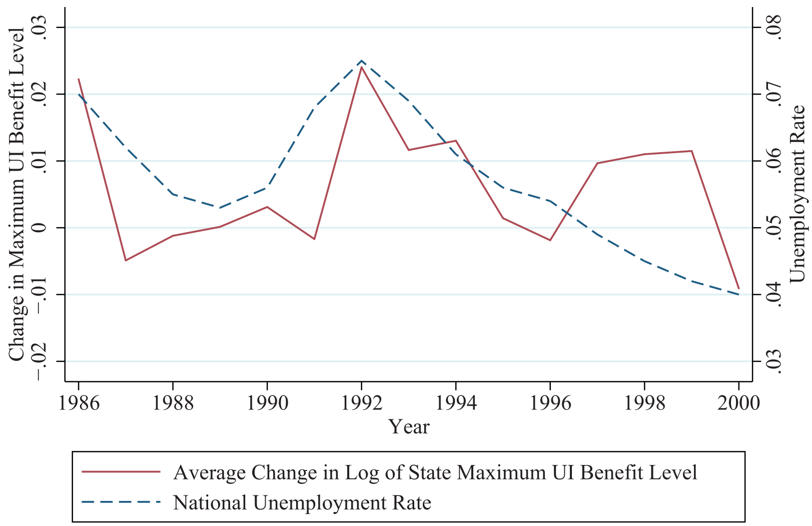

FIGURE 4

UI benefits and the national unemployment rate

Notes: Figure 4 shows the national unemployment rate and the annual average change in the log of the statutory maximum UI weekly benefit for each year between 1986 and 2000 (taking a simple average across states each year). The correlation between these two variables is 0.45.

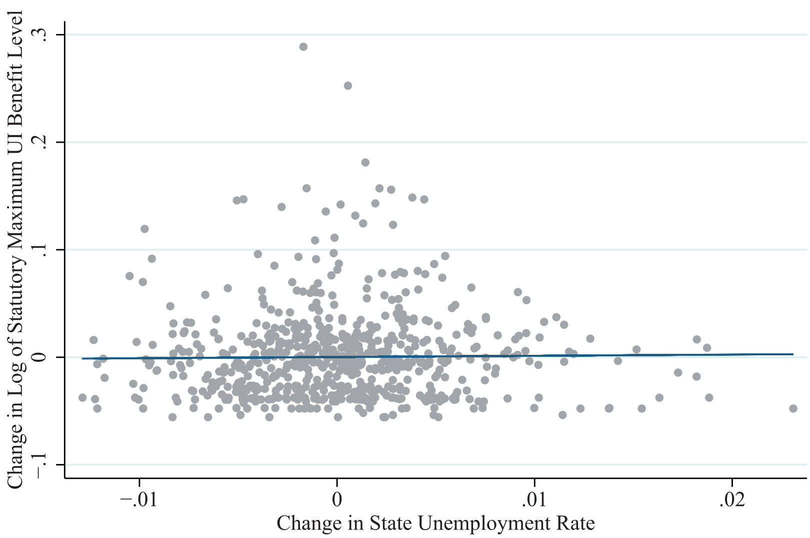

FIGURE 5

Within–state variation in UI benefits and the unemployment rate

Notes: Figure 5 shows the within-state change in the log of the statutory maximum UI weekly benefit amount and the within-state change in the state unemployment rate. Each point reports an annual within-state change in the log of the maximum UI weekly benefit amount and the within-state change in unemployment rate, where both measures are residuals from a regression of the raw measure on year fixed effects. The sample is all within- state changes for all years between 1986 and 2000 for all states in the analysis sample. The solid line in Figure 5 shows the fitted values from an unweighted OLS regression. The correlation between these variables is 0.01. The unemployment rate comes from the Bureau of Labour Statistics, and the statutory maximum UI benefit levels come from the Department of Labour.

We show below that our results are similar when we use the unemployment rate in levels. The baseline estimating equation is the following:[22]

$$\log h_{i,s,t} = \alpha_t + \alpha_s + \beta_1 \log(b_{s,t}) + \beta_2(\log(b_{s,t}) \times \log(u_{s,t_0})) + \beta_3 \log(u_{s,t_0}) + \mathbf{X}_{i,s,t}\Gamma + e_{i,s,t} \quad (4)$$

where $h_{i,s,t}$ is the hazard rate of exit out of unemployment for individual i in state s at time t, α_t and α_s represent year and state fixed effects, respectively, $b_{s,t}$ is the unemployment benefit for individual i at the start of the spell based on the state the individual resided in at the start of the spell, and $\mathbf{X}_{i,s,t}$ is a set of control variables. Our primary proxy for local labour market conditions, u_{s,t_0}, is the state unemployment rate at the start of the spell. We assign the monthly state unemployment rate based on the month at the start of the spell and the individual's state of residence. For example, if an individual in New York became unemployed in July 2000 and his spell lasted until October 2000, we use the New York unemployment rate in July 2000.

All control variables are de-meaned so that $-\beta_1$ represents the elasticity of unemployment durations with respect to the UI benefit level at the average state unemployment rate (for an average individual).[23] The coefficient on the interaction term ($-\beta_2$) is the incremental change in the duration elasticity for a one log point change in the state unemployment rate.

The identifying assumption that allows us to interpret β_2 as a test of whether the duration elasticity varies with the unemployment rate is the following: conditional on the UI weekly benefit amount, state unemployment rate, state fixed effects, year fixed effects, and control variables, there are no omitted determinants of the duration of an unemployment spell that vary with the *interaction* of the UI weekly benefit amount and the state unemployment rate.

An immediate concern with this identification assumption is that UI benefits may be correlated with unobserved labour market conditions. If so, then the direct effect of unemployment benefits, β_1, will be biased. In the Online Appendix, we formally show that we can consistently estimate β_2 so long as the correlation between unobserved labour market conditions and benefits does not depend on the state of the local labour market. If it does, then the estimate of β_2 will be biased. To see the logic, consider the case where benefits are chosen at random in good times, but are endogenous to local labour demand conditions in bad times. Intuitively, in this case we will consistently estimate the duration elasticity in good times; however, if variation in benefits is correlated with unobserved labour market conditions during bad times, then this will cause upward bias in the magnitude of the duration elasticity during bad times. Notice that to the extent that the magnitude of the duration elasticity is significantly *smaller* during bad times, this policy endogeneity likely causes us to understate the magnitude of the interaction term, causing our estimates to be a lower bound on the true magnitude.

In our robustness analysis below, we empirically address the concern about policy endogeneity in two ways. First, we control for the unemployment rate using increasingly flexible functional forms. Second, we use the MSA unemployment rate instead of the state unemployment rate. By doing so, we can exploit within-state, across-MSA variation in local labour market conditions, holding UI benefit levels fixed.

22. The notation of the estimating equation is a simplified presentation of the true model. The (latent) hazard rate is not actually observed in the data, and there is a flexible (non-parametric) baseline hazard rate which is also estimated when fitting the Cox proportional hazard model. Also, following Chetty (2008), we fit a separate baseline hazard rate for each quartile of net liquid wealth, although our results are very similar when a single non-parametric baseline hazard rate is estimated instead (see Online Appendix Table OA.4).

23. We will use the approximation $\log(D) \approx \log(1/h) = -\log(h)$ throughout for the expected unemployment duration, so that the duration elasticity and other marginal effects of interest are given by the negative of the coefficient in the hazard model.

An additional potential violation of our identifying assumption concerns composition bias. As the local unemployment rate fluctuates, there may be compositional changes in the pool of unemployed workers receiving UI benefits. For example, if there is heterogeneity in moral hazard across demographic groups, and the distribution of demographics of the unemployed varies with the level of unemployment, then this compositional change could generate an observed change in the average duration elasticity. We first note that the appropriate measure for the welfare calibrations below is how the average duration elasticity varies with the unemployment rate, and that this is true whether or not the change in the average duration elasticity is primarily due to compositional changes or individual-level changes in moral hazard. Nevertheless, we investigate the extent to which compositional changes can account for our findings, as understanding this may be important for other economic problems.

3.3.1. Empirical Results. We begin by providing graphical evidence on the effect of unemployment benefits on durations. We split the sample into two subsamples according to whether individuals began their unemployment spell in states with above-median unemployment rates or in states with below-median unemployment rates, where the median unemployment rate is defined across states in a given year. We then assign monthly state unemployment rates to unemployment spells based on the unemployment rate in the state that the individual resided in when his spell began. Lastly, we categorize unemployment spells based on whether the prevailing UI benefit level at the start of the spell in a given state and year is above or below the median UI benefit level across the sample.

Figures 6 and 7 show the effect of UI benefits on the probability of unemployment for individuals in above-median and below-median unemployment state-years, respectively.[24] In each figure, we plot Kaplan–Meier survival curves for individuals in low-benefit and high-benefit states. The results in Figure 6 show that the curves are fairly similar in both low-benefit and high-benefit states when the unemployment rate in a state-year is above the median unemployment rate. The curve in high-benefit states is slightly higher, indicating that UI benefits may marginally increase durations, but a non-parametric test that the curves are identical does not reject at conventional levels ($p=0.599$).[25] By contrast, the results in Figure 7 show the curves are noticeably different when the unemployment rate in a state-year is below the median unemployment rate; in particular, durations are significantly longer in high-benefit states, and the difference between the survival curves is statistically significant ($p=0.004$).[26]

These figures suggest that the moral hazard cost of UI benefits depends crucially on whether unemployment is high or low. In particular, our findings suggest that the effect of UI benefits on durations is not statistically significant when the unemployment rate is high but is strongly statistically significant when the unemployment rate is low.[27] These effects are based on simple comparisons across spells. It is possible, however, that the characteristics of individuals vary

24. The survival curves for the full sample are contained in Online Appendix Figure OA.6.

25. Across all the figures, we report p-values of log-rank tests of equality across the two survival curves. This is the appropriate test to use when data are censored (as is the case in our data). Results using Wilcoxon rank sum test, as reported in Chetty (2008), are similar.

26. While the survival curves are statistically significantly different in Figure 7 but not in Figure 6, one might ask whether the difference-in-difference (DD) across the two figures is statistically significant. To answer this question, we construct a semi-parametric test by estimating a Cox proportional hazard model with separate non-parametric baseline hazard estimates for above-median and below-median unemployment state-years. We include two covariates in this Cox model, an indicator for above-median benefits and a DD term which is 1 for state-years with above-median benefits and above-median unemployment state-years and 0 otherwise. The p-value on the estimated DD coefficient is 0.050.

27. We have also looked at the subsample of workers with above-median liquid wealth, and we find broadly similar results (see Online Appendix Figures OA.7 and OA.8). These results suggest that liquidity effects are not primarily

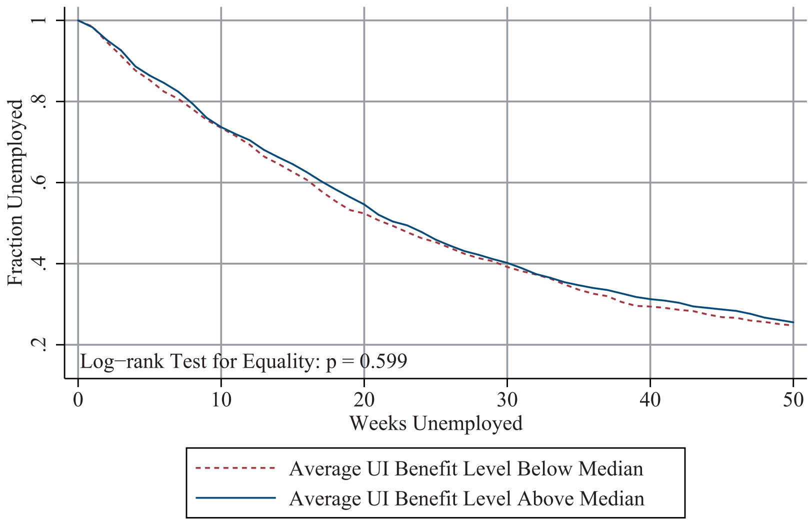

FIGURE 6

Survival curves under high unemployment rate

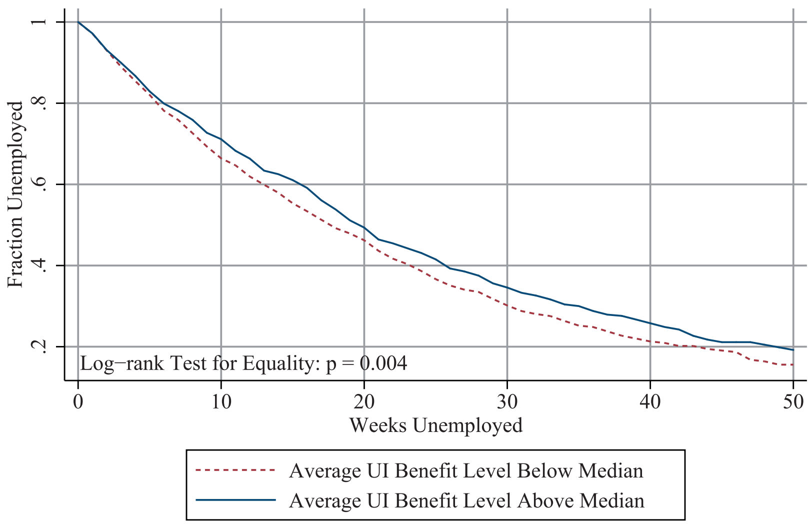

FIGURE 7

Survival curves under low unemployment rate

Notes: Data are individual-level unemployment spells from the 1985–2000 SIPP. In Figure 6, the sample includes spells in states with unemployment rates above the median across states in the year of spell; Figure 7 includes below-median unemployment rates defined similarly. Each figure plots (Kaplan–Meier) survival curves for two groups of individuals based on whether or not the Average UI weekly benefit amount in an individual's state is above or below the overall sample median. The survival curves are adjusted following Chetty (2008), which parametrically adjusts for a "seam effect" by fitting a Cox proportional hazard model with a seam dummy and then recovering the baseline hazard.

TABLE 2
How does the effect of UI on unemployment duration vary with the unemployment rate?

	Specification: Cox proportional hazard regression model	
	Average UI benefit amount	Maximum UI benefit amount
	(1)	(2)
(A) log(Average UI benefit amount)	−0.632	−0.384
	(0.332)	(0.291)
	[0.057]	[0.186]
(B) log(Average UI benefit amount) × log(State unemployment rate)	**1.346**	**1.009**
	(0.457)	**(0.544)**
	[0.003]	**[0.064]**
log(State unemployment rate)	0.035	0.020
	(0.124)	(0.135)
	[0.779]	[0.882]
Age	−0.017	−0.017
	(0.002)	(0.002)
	[0.000]	[0.000]
1{Married}	0.211	0.213
	(0.040)	(0.040)
	[0.000]	[0.000]
Years of education	0.004	0.004
	(0.006)	(0.006)
	[0.498]	[0.491]
Number of Unemployment Spells	4307	4307
Post-estimation:		
High unemployment elasticity ($u = 8.5\%$) (A)$+\sigma\times$ (B)	−0.277	−0.118
	(0.364)	(0.318)
	[0.446]	[0.712]
Low unemployment elasticity ($u = 4.9\%$) (A)$-\sigma\times$ (B)	−0.987	−0.650
	(0.343)	(0.330)
	[0.004]	[0.049]

Notes: All columns report semi-parametric Cox proportional hazard model results from estimating equation (4). Data are individual-level unemployment spells from 1985–2000 SIPP. The average UI benefit amount is the average weekly benefit paid to individuals claiming UI in a given state-year. The maximum UI benefit amount is the statutory weekly benefit amount paid to high wage earners in a given state-year. All specifications include state fixed effects, year fixed effects, industry and occupation fixed effects, a 10-knot linear spline in the log of the annual (pre-unemployment) wage income, and an indicator for being on the seam between interviews. All specifications also include year fixed effects interacted with the log of the UI benefit amount. All columns estimate non-parametric baseline hazards stratified by quartile of net liquid wealth. The final two rows report linear combinations of parameter estimates to produce the duration elasticity when the state unemployment rate is one standard deviation above/below the mean. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

with the unemployment rate in a way that would bias these effects. To investigate this issue and other potential biases, as well as to quantify the magnitude of this interaction effect, we report results from the estimation of semi-parametric proportional hazard models that include a rich set of individual-level controls. Overall, we find that the results from the hazard models are broadly consistent with the patterns of results in these figures.

The main results are reported in Table 2. Following Chetty (2008), the baseline specification controls for age, marital status, years of education, a full set of state, year, industry, and occupation fixed effects, and a 10-knot linear spline in log annual wage income.[28] Column (1) reports

accounting for the differential duration elasticity between high and low unemployment, which is broadly consistent with the results in Online Appendix Table OA.13 described below.

28. The only change to the baseline empirical specification in Chetty (2008) that we make is that we do not include the interaction of log(Average UI Benefit Amount) with unemployment duration (i.e. number of weeks elapsed in current

estimates of equation (4). The key coefficient of interest is the interaction term between the UI benefit level and the log state unemployment rate. The results indicate that the elasticity of unemployment durations with respect to the UI benefit level ($-\beta_1$) is 0.63 (SE 0.33) at the average unemployment rate. The (average) duration elasticity estimate is broadly similar to the previous literature (Moffitt, 1985; Meyer, 1990; Chetty, 2008). The results in column (1) show an estimate of $-\beta_2$ of -1.35 (SE 0.46). The bottom two rows of Table 2 report the duration elasticity when the state unemployment rate is one standard deviation (1.3 percentage points) above and below the mean unemployment rate (6.2%). At one standard deviation above the mean, the duration elasticity is 0.28 (SE 0.36), while at one standard deviation below the mean, the duration elasticity is 0.99 (SE 0.34). In column (2), the average UI benefit level is replaced by the statutory maximum UI benefit level in the state-year, and the results are very similar. In the robustness tests that follow, we will present results which use both the average and the maximum UI benefit level.

These results imply that the magnitude of the duration elasticity decreases with the unemployment rate and suggest that the moral hazard cost of UI is lower when the unemployment rate is relatively high. This empirical finding is consistent with a parameterization of our model where search effort (e) and labour demand conditions (α) are strongly complementary, as in the simulation reported in Figure 1. Moreover, the range of elasticities that we estimate empirically are within the range of elasticities that we find in Figure 1.

3.3.2. Robustness Analysis. We next implement a series of robustness tests. The collection of evidence in this section suggests that our baseline results are fairly robust to alternative specifications, samples, and measures of key variables.

In Table 3, we report results which control flexibly for the local unemployment rate and for unobserved trends.[29] Column (1) reports our baseline specification for comparison. Columns (2) and (3) include quadratic and cubic polynomial functions of the state unemployment rate, respectively, to addresses concerns that UI benefits respond non-linearly to labour market conditions. Additionally, to the extent that the flexible polynomial in the unemployment rate more thoroughly controls for unobserved local labour market conditions, this specification can be used to gauge the extent of the bias due to policy endogeneity. Although the results are somewhat less precise, the estimated interaction term is similar when these more flexible controls are included.[30] Columns (4) and (5) report results from modifications of our baseline specification which focus on alternative assumptions regarding contemporaneous trends across states within a region and within states over time. We find that the estimated interaction term is similar with controls for these unobserved trends.

The remainder of Table 3 reports results using metropolitan area (MSA) unemployment rates, rather than state unemployment rates.[31] Columns (6) through (8) report results using the average UI benefit level; analogous results using the maximum UI benefit level are reported

spell). This control is intended to capture duration dependence in the response to UI benefits, but because it is difficult to interpret this coefficient and it is always statistically and economically insignificant, we do not include it in any specifications. All results with this interaction term included are extremely similar.

29. All of the results in Table 3 are replicated in Online Appendix Tables OA.1 using the maximum UI benefit level instead of the average UI benefit level, and the results are similar.

30. We have also investigated robustness to including various polynomial functions of the UI benefit level as well as interactions between the state unemployment rate and state fixed effects and interactions between the state unemployment rate and year fixed effects. These specifications capture the possibility that in certain states and/or certain years, UI benefits may be unusually responsive to changes in local labour market conditions. These results are in Online Appendix Table OA.3, and the results are similar to our baseline specification.

31. To preserve the sample size, we assign the state unemployment rate to all unemployed individuals who do not have an MSA code in the data, which is roughly 50% of the sample.

TABLE 3

Robustness to controlling for observed and unobserved trends and to using variation across metropolitan areas within states

	Specification: Cox proportional hazard regression model							
Local labor market conditions proxy:	State Unemployment Rate					Metropolitan Area Unemployment Rate		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(A) log(Average UI benefit amount)	−0.632	−0.625	−0.688	−0.830	−0.788	−0.442	−0.877	
	(0.332)	(0.391)	(0.409)	(0.400)	(0.516)	(0.315)	(0.508)	N/A
	[0.057]	[0.110]	[0.092]	[0.038]	[0.126]	[0.160]	[0.084]	
(B) log(Average UI benefit amount) × log(State unemployment rate)	**1.346**	**1.340**	**1.369**	**1.432**	**1.361**			
	(0.457)	**(0.517)**	**(0.482)**	**(0.472)**	**(0.482)**			
	[0.003]	**[0.010]**	**[0.004]**	**[0.002]**	**[0.005]**			
(B) log(Average UI benefit amount) × log(Metropolitan area unemp. rate)						**1.019**	**1.364**	**2.211**
						(0.470)	**(0.538)**	**(1.027)**
						[0.030]	**[0.011]**	**[0.031]**
State FEs and year FEs	✓	✓	✓	✓	✓			
Metropolitan area FEs and year FEs						✓	✓	✓
Quadratic in state unemployment rate		✓						
Cubic in state unemployment rate			✓					
Region-specific linear time trends				✓				
State-specific linear time trends					✓		✓	
State × year FEs								✓
Post-estimation:								
High unemployment elasticity	−0.277	−0.271	−0.327	−0.453	−0.429	−0.111	−0.435	
($u=8.5\%$) (A)$+\sigma\times$(B)	(0.364)	(0.387)	(0.406)	(0.426)	(0.556)	(0.351)	(0.564)	
	[0.446]	[0.484]	[0.421]	[0.288]	[0.440]	[0.751]	[0.441]	*N/A*
Low unemployment elasticity	−0.987	−0.979	−1.050	−1.208	−1.147	−0.773	−1.321	
($u=4.9\%$) (A)$-\sigma\times$(B)	(0.343)	(0.439)	(0.449)	(0.412)	(0.505)	(0.349)	(0.508)	
	[0.004]	[0.026]	[0.019]	[0.003]	[0.023]	[0.027]	[0.009]	

Notes: All columns report Cox proportional hazard model results from estimating equation (4). Data are individual-level unemployment spells from 1985–2000 SIPP. Number of unemployment spells = 4307. See Table 2 for more details on the baseline specification and Online Appendix Tables OA.1 and OA.2 for analogous results using the Maximum UI Benefit Amount. To preserve sample size, observations without MSA codes are grouped together within a state and assigned the state unemployment rate in columns (6) through (8). The final two rows reports linear combinations of parameter estimates to produce the unemployment duration elasticity when the state unemployment rate is one standard deviation above or below the mean. Standard errors are shown in parentheses and are clustered by state, and *p*-values are shown in brackets.

in the Online Appendix Table OA.2. Overall, the MSA-level results are fairly similar to the baseline specification. In column (8), we report results which include a full set of state-by-year fixed effects, so that the only variation used to estimate the interaction term is within-state-year, across-MSA variation in the unemployment rate, holding the state-year UI benefit level constant. The interaction term is now somewhat larger than the baseline specification, though the statistical precision is substantially reduced with the inclusion of state-by-year fixed effects. These results suggest that policy endogeneity is unlikely to fully account for our results, since UI policy is the same across MSAs within a state. The results in this table also alleviate the concern that our estimates confound state UI potential duration effects with state UI benefit level effects, since state-year fixed effects capture all of the variation in the maximum potential duration of UI benefits in our data. Overall, the results in Tables 3 suggest that policy endogeneity is not primarily responsible for our findings. The next set of robustness tests explore additional threats to validity and alternative explanations for our findings.

In Table 4, we explore several alternative measures of the interaction term in each row. The first row reproduces our baseline estimates for comparison. The second row replaces the

TABLE 4
Robustness to alternative measures of the key interaction term

	Cox proportional hazard model estimates		Post-estimation	
	Direct effect, **(A)**	Interaction effect, **(B)**	**(A)** $+\sigma\times$ **(B)**	**(A)** $-\sigma\times$ **(B)**
(1) log(Average UI benefit amount) × log(State unemployment rate)	−0.632 (0.332) [0.057]	**1.346** **(0.457)** **[0.003]**	−0.277 (0.364) [0.446]	−0.987 (0.343) [0.004]
(2) log(Average UI benefit amount) × **1**{state unemp. rate ≥ median unemp. rate}	−1.443 (0.405) [0.000]	**1.068** **(0.211)** **[0.000]**	−0.374 (0.345) [0.278]	*N/A*
(3) log(Average UI benefit amount) × state unemployment rate	−0.651 (0.353) [0.065]	**0.164** **(0.073)** **[0.024]**	−0.374 (0.362) [0.302]	−0.928 (0.385) [0.016]
(4) log(Average UI replacement rate) × log(State unemployment rate)	−0.391 (0.301) [0.194]	**1.524** **(0.558)** **[0.006]**	0.012 (0.380) [0.976]	−0.793 (0.283) [0.005]
(5) log(Average UI benefit amount) × log(predicted vacancy/unemployment ratio)	−0.498 0.309 [0.107]	**0.293** **(0.134)** **[0.029]**	−0.197 (0.335) [0.557]	−0.799 (0.341) [0.019]
(6) log(Average UI WBA) × − log(predicted employment-to-pop ratio)	−0.683 (0.402) [0.089]	**0.672** **(0.808)** **[0.406]**	−0.506 (0.318) [0.112]	−0.860 (0.559) [0.124]
(7) log(Average UI WBA) × Average of log(state unemp. rate), 1985–2000	−0.698 (0.406) [0.086]	**2.813** **(2.179)** **[0.197]**	−0.261 (0.389) [0.502]	−1.136 (0.639) [0.076]
log(Average UI WBA) × (log(state unemp. rate) − average,1985–2000)		**1.167** **(0.512)** **[0.023]**	−0.420 (0.451) [0.327]	−0.955 (0.390) [0.014]
			p-value of test of equality of two interaction terms in (7) = 0.489	

Notes: All rows report semi-parametric Cox proportional hazard model results from estimating equation (4); each column reports separate parameter estimates. Data are individual-level unemployment spells from 1985–2000 SIPP. Number of unemployment spells = 4307 in all rows except for row (6) where number of spells = 4296. See Table 2 for more details on the baseline specification. In row (2), the median unemployment rate is calculated separately each year. In row (3), the Average UI WBA is the average weekly benefit amount paid to individuals claiming UI. In row (4), the average UI replacement rate is the average UI WBA divided by average weekly wages in a given state-year for prime-age males (computed from the CPS). In row (5), the Predicted Vacancy/Unemployed Ratio is constructed by imputing missing vacancy data using the log(Vacancy Index) from the 26 metropolitan areas covered by the monthly HWI data set produced by the Conference Board for 1985–2000. We estimate a regression of log(Vacancy Index) on log(metropolitan area unemployment rate), metropolitan area fixed effects and year-month fixed effects and impute log(Vacancy Index) for the other metropolitan areas not covered by the HWI data. In row (6), the Predicted Employment to Population Ratio is computed following the "shift share" procedure of Bartik (1991); see text for details. In row (7), the interaction term is split into two separate interaction terms to decompose the variation in the relative unemployment rate into within-state and between-state variation. Across all rows, the final two columns report linear combinations of parameter estimates to produce the unemployment duration elasticity when the labor market conditions variable is one standard deviation above or below the mean. In row (2), σ is set to 1.0 because the interaction term includes a dummy variable rather than a continuous measure of unemployment. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

state unemployment rate with a dummy variable for whether or not the unemployment rate is greater than the median state unemployment rate in a given year. This specification corresponds more closely to the non-parametric results presented in the figures above. The third row reports results using the state unemployment rate in levels (rather than logs). In both cases, the results are similar to the baseline specification. The fourth row shows that results are similar using the average replacement rate rather than the average weekly UI benefit amount.[32] The fifth row reports results using the vacancy/unemployment ratio as the proxy for labour market conditions. We use the Conference Board's Help-Wanted Index (HWI) to measure vacancies in metropolitan areas.[33] Since some areas are not covered by the HWI, we impute the HWI in these areas based on the observed unemployment rate which is used to generate a predicted vacancy index.[34] The results in this row are qualitatively similar to the baseline results. The sixth row uses an alternative proxy for local labour demand instead of the local unemployment rate, which addresses the concern that the unemployment rate may reflect both labour demand and labour supply shocks. We construct variation in the employment-population ratio that is driven by plausibly exogenous shifts in local labour demand using the procedure in Bartik (1991).[35] The results in the sixth row indicate that the magnitude of the estimated interaction term is somewhat similar to our baseline specification, but is very imprecisely estimated.

In our baseline specification, identification of the interaction term of interest comes from both across-state and within-state variation in unemployment rates. Online Appendix Figures OA.11 and OA.12 report survival curves analogous to Figures 6 and 7 using only within-state variation in unemployment rates, while Online Appendix Figures OA.13 and OA.14 show similar results using only cross-state variation in the state unemployment rate.[36] We quantify these within-state and cross-state patterns in the last two rows of Table 4, where we report results from a specification where we decompose the variation in the state unemployment rate into across-state and within-state variation. This specification allows us to see separately how the two sources of variation affect the duration elasticity. We find that both interaction terms are the same sign as the

32. We also report reduced form results using a "simulated instrumental variable" following Currie and Gruber (1996) in Online Appendix Table OA.4. By construction, the variation in this UI benefit variable is only due to changes in program parameters, holding sample composition constant. This variable is constructed by using a fixed 20% 1993 (national) sample and computing the average weekly UI benefit in this fixed sample for every state-year combination in the data set. In Online Appendix Table OA.4, we also report similar results when this simulated instrumental variable is used as an instrument for the average UI WBA. This instrumental variables estimation is implemented using a two-step control function approach.

33. The Conference Board is a non-profit company that creates the HWI, which is a monthly index of classified ads for jobs that are found in print newspapers of 51 major metropolitan areas. The monthly counts from each of the newspapers are adjusted to account for seasonality as well as differences in the number of weekdays and Sundays across months. The adjusted figures are then normalized to a 1967 = 100 base and aggregated using non-agricultural payroll employment weights to form the national HWI. According to Abraham (1987), in 1974, the index represented cities that accounted for roughly 50% of the total non-agricultural employment in the U.S. For further details on the construction of the index, see Abraham (1987).

34. To construct the predicted vacancy index, we regress log(Vacancy Index) on year-month fixed effects, metropolitan area fixed effects, and the log(metropolitan area unemployment rate). We then use the estimated model to calculate fitted values for the metropolitan areas that are not covered by HWI data. We then use the log(Vacancy Index/unemployment rate) as a proxy for labour market conditions, where the Vacancy Index is the true Vacancy Index for metropolitan areas that are covered by HWI data, and predicted otherwise. Similar to our MSA-level analysis in Table 3, to preserve sample size we group all non-metro areas in a state together and predict the vacancy index for these areas, as well (and use the state unemployment rate for these areas).

35. We closely follow the implementation of the Bartik procedure in Autor and Duggan (2003). We predict the state employment-to-population ratio by interacting initial cross-sectional distribution of state-level employment shares with national industry employment trends. Online Appendix Figures OA.9 and OA.10 plot survival curves comparing the effect of UI benefits across high and low predicted employment-to-population ratios. Consistent with Figures 6 and 7, this non-parametric evidence indicates that the effect of UI benefits is largest during periods of high predicted employment.

36. To construct these figures, we compute the average state unemployment rate over the sample period and divide the states based on whether they are above or below the median for the cross-state figures and we subtract off this average unemployment rate for the within-state figures. These figures show that the same pattern in Figures 6 and 7 emerges when using only within-state variation or only cross-state variation in the state unemployment rate. In particular, there is always a statistically significant difference between the high- and low-benefit survival curves when the unemployment rate is relatively low, but not when it is relatively high.

interaction term in the baseline specification. Although our statistical power is limited, we cannot reject the hypothesis that the correlation between the duration elasticity and the unemployment rate is the same across both sources of variation.

We present a number of additional results to assess the robustness of our results in the Online Appendix. We briefly summarize these results in the remainder of this section. In Online Appendix Table OA.5, we report estimates of an augmented version of our baseline specification where we add interactions between UI benefits and the following demographic variables: age, marital status, years of education, pre-unemployment wage, pre-unemployment occupation indicator variables, and pre-unemployment industry indicators.[37] If the estimated interaction term in our baseline specification is primarily due to compositional changes among demographic groups with different duration elasticities, then we would expect to see a reduction in the magnitude of the estimated interaction term as we include additional interactions between UI benefits and demographic controls. The results in Online Appendix Table OA.5 show that our main result is robust to including such controls; including interactions between demographics and UI benefits has a negligible effect on our main coefficient of interest.[38] An important caveat to this analysis is that composition bias may be limited through sample selection; since our analysis sample is restricted to prime-age men, this may reduce scope for composition effects. The final column of Online Appendix Table OA.5 investigates a related source of compositional bias, which is selection bias due to endogenous take-up. We find that the effect of UI benefits on take-up varies with the unemployment rate, with those induced to take up benefits when the unemployment rate is high being less responsive to UI benefit levels.[39] We also investigate the role of measurement error in labour market conditions. In Online Appendix Table OA.8, we report results limiting the sample to the largest 20 states by population in 2000. This sample restriction is intended to reduce the influence of measurement error in the unemployment rate, and the results are similar to our baseline specification.[40] Lastly, while our main results focus on the 1985–2000 SIPP sample, in Online Appendix Table OA.9, we extend the SIPP sample to 2007, so as to cover the entire 1985–2007 period. We find that the estimated interaction term is similar, but the average duration elasticity is somewhat smaller in this later period.[41]

To summarize, across all the specifications in this section, we find no evidence that our baseline results are primarily due to compositional changes, sample selection, mismeasurement, or other spurious factors. We therefore conclude that the most likely explanation for our findings is that the moral hazard cost of UI decreases with the unemployment rate. The next section describes the second part of the empirical analysis, which investigates consumption smoothing.

37. These tests are motivated by recent work which finds evidence that the composition of unemployed workers varies over the business cycle (Mueller 2012). In Online Appendix Table OA.6, we directly investigate how the composition of workers is associated with the unemployment rate. We do not find significant evidence that the composition of unemployed workers varies with the unemployment rate.

38. In Online Appendix Table OA.7, we show that results are similar to Online Appendix Table OA.5 when using the (statutory) maximum UI benefit amount instead of average UI benefit amount.

39. These results raise concerns about possible selection bias, though the results in the rest of the columns in Online Appendix Table OA.5, suggest negligible effects of selection on observables. To address the selection bias concern, we evaluated the following two-step estimator. In the first step, we estimate a probit model of UI receipt on interaction term using the same set of controls used in the baseline proportional hazard model using the expanded sample which includes eligibles who do not receive UI benefits. In the second step, we estimate the baseline hazard specification including as an additional control the inverse Mills ratio evaluated at the fitted values. The two-step estimation results (not shown) were similar to our main results.

40. Additionally, this table also reports alternative standard errors that account for uncertainty in unemployment rate estimates, which are roughly 20% larger than the baseline specification.

41. Online Appendix Table OA.10 reports results analogous to Table 3 in this extended sample.

3.4. *Part 2: Consumption Smoothing*

To investigate how the consumption smoothing benefit of UI varies with the unemployment rate, we study how UI affects the consumption drop upon unemployment and how this effect varies with the unemployment rate. Specifically, we use the PSID data described above and OLS to estimate the following linear model:

$$\Delta \log C_{i,t} = \alpha_0 + \alpha_t + \alpha_s + \delta_1 b_{i,s,t} + \delta_2 (b_{i,s,t} \times \log(u_{s,t-1})) + \delta_3 \log(u_{s,t-1}) + \mathbf{X}_{i,s,t}\Gamma + e_{i,s,t} \quad (5)$$

where $\Delta \log C_{i,t}$ is the difference in log total food consumption for individual i between year $t-1$ and year t, $b_{i,s,t}$ is the after-tax UI replacement rate, $u_{s,t-1}$ is the state unemployment rate in year $t-1$, α_t and α_s are year and state fixed effects, and $\mathbf{X}_{i,s,t}$ is the same set of control variables used in Gruber (1997).[42] We de-mean the state unemployment rate so that the interpretation of δ_1 is the effect of UI at the average state unemployment rate, and we de-mean the other variables so that α_0 represents the average consumption drop upon unemployment across all individuals in the sample. The coefficient δ_3 represents how the average consumption drop varies with the unemployment rate, and δ_2 represents how the consumption smoothing effect of UI varies with the state unemployment rate.

3.4.1. Empirical Results. Table 5 reports results of estimating equation (5) using the 1968–1987 PSID sample, following Gruber (1997). Column (1) reproduces column (4) in Table 1 of Gruber (1997), and column (2) reports our replication effort, which shows the average consumption smoothing benefit of UI using our replication sample. We find that a 10 percentage point increase in the UI replacement rate reduces the consumption drop upon unemployment by 2.7% (SE 0.9%), which is very similar to the estimate of δ_1 in column (1). Column (3) reports our preferred specification that includes the interaction between the replacement rate and the unemployment rate. The estimate on the coefficient of our interaction term is economically and statistically insignificant ($\delta_2 = 0.015$, SE 0.236), and the remaining columns of Table 5 show similar results for a variety of alternative specifications which control for various unobserved trends. As with the duration elasticity analysis above, the final two rows report estimates at one standard deviation above and below the mean unemployment rate. Unlike the duration elasticity results (which showed that the duration elasticity was statistically significantly lower when the unemployment rate was relatively high), the final two rows in Table 5 consistently show that the effect of UI on the consumption drop at unemployment does not meaningfully vary with the unemployment rate.

While the results in Table 5 indicate that unemployment benefits have a causal effect on the consumption drop upon unemployment, the marginal welfare gain formula in equation (3) emphasizes that the consumption smoothing benefits can be calculated using simpler empirical objects—specifically, the average consumption drop upon unemployment, and how this varies with the labour market conditions. With this insight, Table 6 re-interprets the regression results from equation (5) above. The key coefficient of interest is now the average consumption drop upon unemployment (α_0) and how this varies with the state unemployment rate (δ_3), rather than δ_1 and δ_2 as in Table 5.[43] The results of estimating equation (5) are reported in Table 6.[44] Columns

42. We use the previous year's unemployment rate because we do not observe individuals at the start of their spell, and we want to ensure that the unemployment rate is predetermined, for the reasons discussed above.

43. For consistency, we estimate the same regression model in both Table 5 and Table 6, although the results in Table 6 are similar if we drop the UI replacement rate (and its interaction with the state unemployment rate) from equation (5) and re-estimate the model using this simpler specification.

44. To increase comparability with the duration elasticity estimates for these results, the PSID sample is extended from 1968–1987 to 1968–1997. Results for 1968–1987 are similar and are reported in Online Appendix Table OA.12.

TABLE 5
How does the effect of UI on the average consumption change upon unemployment vary with the unemployment rate?

	Gruber (1997), Table 1, Column (4)	Replication sample and men only subsample				
	(1)	(2)	(3)	(4)	(5)	(6)
(A) UI replacement rate	0.280	0.278	0.292	0.323	0.312	0.276
[1.00 = 100% replacement]	(0.105)	(0.087)	(0.094)	(0.139)	(0.146)	(0.154)
		[0.003]	[0.003]	[0.025]	[0.037]	[0.079]
(B) Implied consumption drop at replacement rate of 0	−0.231	−0.229	−0.237	−0.269	−0.263	−0.241
UI replacement rate ×			**0.015**	**0.060**	**0.106**	**0.119**
log(State Unemployment Rate)			**(0.236)**	**(0.304)**	**(0.323)**	**(0.350)**
			[0.949]	**[0.843]**	**[0.745]**	**[0.734]**
N	1604	1605	1605	1249	1249	1249
R^2		0.122	0.122	0.133	0.138	0.162
State and year FEs	✓	✓	✓	✓	✓	✓
Men only subsample				✓	✓	✓
Region-specific linear trends					✓	
State-specific linear trends						✓
Post-estimation:						
Marginal effect of UI for high unemp.			0.296	0.340	0.343	0.310
(u=8.5%) (A)+$\sigma\times$(B)			(0.122)	(0.137)	(0.137)	(0.148)
			[0.015]	[0.013]	[0.013]	[0.037]
Marginal effect of UI for low unemp.			0.288	0.305	0.282	0.241
(u=4.9%) (A)−$\sigma\times$(B)			(0.110)	(0.188)	(0.202)	(0.213)
			[0.009]	[0.104]	[0.163]	[0.258]

Notes: Column (1) reproduces the results from column (4) in Gruber (1997), Table 1. The remainder of the columns report OLS results from estimating equation (5) on a replication sample. Data are individual-level observations from 1968–1987 PSID. See text for more details on the baseline specification. The implied consumption drop is computed as the average fitted value across the sample when the replacement rate is set to 0 for all observations. Standard errors are shown in parentheses and are clustered by state, and p-values are shown in brackets.

(1) through (3) show unweighted results for different sets of controls (adding region-specific trends and state-specific trends in (2) and (3), respectively). The first row of the table reports the average consumption drop upon unemployment (α_0) and the second row reports how this varies with the state unemployment rate (δ_3). The results in column (1) of Table 6 indicate that the mean consumption drop upon unemployment is roughly 6.9% at the mean state unemployment rate and at existing levels of UI benefits. The estimated magnitude of α_0 and δ_3 is similar across columns (1) through (3). Column (3) reports an estimate of $\delta_3=-0.041$ (SE 0.066), which suggests that the consumption drop is slightly larger when the state unemployment rate is high, although this estimate is small in magnitude and not statistically significant at conventional levels.

The marginal welfare gain formula in equation (3) indicates that the survival function should be used to construct weights to calculate the appropriate weighted-average consumption drop upon unemployment. The PSID data are not ideal for this exercise because they do not contain precise information on unemployment duration. Despite this limitation, we estimate weeks unemployed using the date of the interview, assuming all individuals began unemployment at the beginning of the year, and we use the survival functions from the SIPP data to construct appropriate weights. These weights are used to constructed weighted OLS estimates in column (4) through (6) analogous to columns (1) through (3). Overall, the results are very similar to the unweighted results.

3.4.2. Robustness Analysis. In the Online Appendix, we report estimates of an augmented version of equation (5) which includes alternative controls (specifically,

TABLE 6
How does the average consumption drop upon unemployment vary with the unemployment rate?

	Sample period: 1968–1997					
Weights:	Unweighted sample			Survival curve weights		
	(1)	(2)	(3)	(4)	(5)	(6)
Average consumption drop upon unemployment	−0.069	−0.070	−0.067	−0.069	−0.070	−0.064
[−0.10 = −10%]	(0.003)	(0.003)	(0.003)	(0.007)	(0.007)	(0.006)
	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]	[0.000]
The effect of log(State unemployment rate) on the	−0.004	0.013	−0.041	0.006	0.028	−0.047
average consumption drop upon unemployment	(0.062)	(0.065)	(0.066)	(0.070)	(0.068)	(0.066)
	[0.945]	[0.839]	[0.538]	[0.931]	[0.683]	[0.486]
N	2003	2003	2003	2003	2003	2003
R^2	0.114	0.118	0.132	0.132	0.137	0.160
State and year FEs	✓	✓	✓	✓	✓	✓
Region-specific linear trends		✓			✓	
State-specific linear trends			✓			✓
Post-estimation:						
Average consumption drop for high unemp.	−0.070	−0.066	−0.079	−0.067	−0.061	−0.078
($u = 8.5\%$)	(0.016)	(0.017)	(0.017)	(0.019)	(0.018)	(0.018)
	[0.000]	[0.000]	[0.000]	[0.001]	[0.001]	[0.000]
Average consumption drop for low unemp.	−0.068	−0.074	−0.055	−0.071	−0.078	−0.050
($u = 4.9\%$)	(0.020)	(0.022)	(0.022)	(0.024)	(0.024)	(0.023)
	[0.001]	[0.001]	[0.013]	[0.003]	[0.001]	[0.026]

Notes: All columns report results from estimating equation (5). Data are individual-level observations from 1968–1997 PSID. Columns (1) through (3) are unweighted, while columns (4) through (6) use weights based on estimated survival curves in Figures 6 and 7. See text for more details on the specification. The final two rows report linear combinations of parameter estimates to produce the consumption change when the state unemployment rate is one standard deviation above/below the mean. Standard errors are shown in parentheses and are clustered by state, and *p*-values are shown in brackets.

pre-unemployment wage, age, years of education, number of kids, marital status, and race). If the estimate of the coefficient on the state unemployment rate is primarily driven by compositional changes among demographic groups with different average consumption changes, then we would expect to see a change in the magnitude of the estimated coefficient as we change the set of demographic controls. The results in Online Appendix Table OA.11 show that our main result is robust to including/excluding such controls, since these modifications have a negligible effect on our main coefficient of interest. We also report results in Online Appendix OA.12 which extend the sample in several ways: including both men and women in the sample, as in Gruber (1997); extending the sample period to 2007; and restricting the sample to the same years as in Gruber (1997). Across all of these alternative specifications and samples, we find no evidence that average consumption drop upon unemployment varies with the unemployment rate, although our precision is often limited.

In addition to evaluating the consumption smoothing benefits of UI based on consumption data, we can also modify the duration elasticity specifications from above to study the consumption smoothing effect of UI by looking across individuals with different levels of wealth. Chetty (2008) presents evidence that a component of the observed duration elasticity represents a "liquidity effect" that proxies for the consumption smoothing benefit of UI. Following this insight, we can construct alternative tests of consumption smoothing benefits by estimating how the liquidity effect of UI varies with the state unemployment rate. These tests are reported in Online Appendix Table OA.13. Column (1) reports our baseline specification for comparison, while column (2) reports results for the subsample of workers in the third and fourth quartiles of net liquid wealth, where liquidity effects are likely to be less important. The coefficient on the interaction term is slightly larger than in the baseline specification. This is also true when using the subsample of unemployed workers without a mortgage (another proxy for individuals who are not liquidity constrained). In both cases, we would expect to estimate a smaller coefficient if the liquidity benefits of UI was larger when the unemployment rate is high, and therefore, these results provide no evidence that liquidity effects are larger when the unemployment rate is high. Columns (3) and (4) report results that include a full set of liquid wealth quartile dummy variables interacted with a combination of occupation fixed effects, industry fixed effects, unemployment duration, and the UI benefit level, and the results are also similar to the main results in column (1). Lastly, column (5) directly estimates how the interaction term of interest varies with liquid wealth, and the estimate is not statistically significant. We therefore conclude that the results support the interpretation that the liquidity effect of UI does not vary with the unemployment rate. These results are broadly consistent with the results in Tables 5 and 6 in that both results imply that the consumption smoothing benefit of UI does not vary strongly with the unemployment rate. Although both the consumption results and the liquidity effect results are limited by sample size and statistical power, taken together they show that the consumption smoothing benefit of UI does not vary strongly with the unemployment rate.

3.4.3. Discussion. We conclude this section with a brief discussion of the (perhaps somewhat surprising) empirical results as well as the important limitations of this analysis. Given the simulation results from the theoretical model in Section 2, we expected to see larger declines in consumption upon unemployment during times of high unemployment. The point estimates in Tables 5 and 6 generally line up with this prediction, but they are not large in magnitude and never statistically significant (although the lack of statistical significance is likely due to a lack of statistical power given the small sample size). The numerical simulations in Section 2.3 are useful because they make clear what one would expect to see in the data given the parameters chosen for the model calibration. In Figure 2, the numerical simulations indicated that the consumption drop increases in magnitude from roughly -5% to -10% as unemployment increases from 4% to 8%. The point estimates in the last column of Table 6 are fairly similar to these results: the empirical results show an average consumption drop when the unemployment rate is low of roughly -5%, while the consumption drop when unemployment rate is high is roughly -8%. The data therefore do not have the power to rule out magnitudes of variation in $\Delta c/c$ that one would predict from the model calibrations above, which is one important limitation of the analysis.[45] This points to the need for better estimates of the mean consumption change upon unemployment (and how this varies with labour market conditions) as an important area of future work. Another important limitation is that our analysis is restricted to food consumption, which may not be representative of broader changes in consumption upon unemployment. This highlights the need for future work that considers a broader measure of consumption.

Despite these limitations, our results suggest important consumption smoothing benefits of UI that can be compared against the duration elasticity estimates from the first part of the empirical

45. Another possibility is that the consumption drop does not vary with unemployment rate for reasons outside of the model. For example, unemployed individuals could be severely liquidity constrained or highly impatient. In this case, consumption would always fall to the benefit level regardless of labour market conditions. This is hard to reconcile with results suggesting small consumption changes on average, however. Alternatively, the types of individuals who lose jobs during recessions may have more access to liquidity. Although we explored the role of cyclical changes in the composition of the unemployed, this is a kind of compositional change that we are not able to fully investigate given the limitations of our data. This explanation is also consistent with the results in Mueller (2012), who finds that the unemployed are positively selected during downturns. Last, it is possible that the unemployed are overly optimistic about the probability that they find a job and may not realize it will take longer to find a job in recessions; see Spinnewijn (2015) for suggestive empirical evidence on this. In this case, unemployed individuals may consume too much initially.

TABLE 7
Calibrations of the marginal welfare gain of UI ($dW/db \times \$10$) by unemployment rate

Coefficient of relative risk aversion, γ		Unemployment rate and implied duration elasticity and consumption change upon unemployment				
	$u=$	4.1%	5.4%	6.7%	8.0%	9.3%
	$\varepsilon_{D,b}=$	1.293	0.922	0.632	0.393	0.191
	$\Delta c/c=$	0.067	0.068	0.069	0.070	0.070
$\gamma=1$		−\$0.49	−\$0.47	−\$0.39	−\$0.28	−\$0.13
$\gamma=2$		−\$0.47	−\$0.43	−\$0.35	−\$0.23	−\$0.06
$\gamma=3$		−\$0.44	−\$0.40	−\$0.30	−\$0.17	+\$0.00
$\gamma=4$		−\$0.42	−\$0.36	−\$0.26	−\$0.12	+\$0.07
$\gamma=5$		−\$0.39	−\$0.33	−\$0.22	−\$0.06	+\$0.13

Notes: Each cell reports the marginal welfare gain calculation according to the approximation formula in equation (3), which is based on the following statistics: (1) an assumed coefficient of relative risk aversion as indicated in the first column, (2) estimates of the elasticity of unemployment duration with respect to the UI benefit level (and how this elasticity varies with the unemployment rate) from Table 2, and (3) estimates of the average consumption change upon unemployment (and how this varies with the unemployment rate) from Table 6. See Section 4 for more details on the computations and Online Appendix Table OA.14 for additional calibration results.

analysis. In the next section, we conduct a welfare analysis using our empirical estimates as inputs in the marginal welfare gain formula in equation (3).

4. CALIBRATING THE MARGINAL WELFARE GAIN

Our empirical results suggest that the duration elasticity decreases with the unemployment rate, while the consumption drop upon unemployment is approximately constant. We use these estimates to calibrate the marginal welfare gain formula in equation (3), which we reproduce below:

$$\frac{dW}{db} \approx \frac{D_B}{D}\frac{u}{1-u}\left\{\gamma\frac{\Delta c}{c} - \frac{\varepsilon_{D,b}}{1-u}\right\}.$$

The purpose of this section is to empirically calibrate this formula for several values of the unemployment rate (u). Given the empirical specifications in the two-part empirical analysis above, we calibrate this formula assuming that $\varepsilon_{D,b}$ and $\Delta c/c$ are linear functions of $\log(u)$:

$$\varepsilon(u) = -\beta_1 - \beta_2 \times (\log(u) - \log(\overline{u}))$$
$$\frac{\Delta c}{c}(u) = -\alpha_0 - \delta_3 \times (\log(u) - \log(\overline{u})).$$

Our preferred empirical estimates are that $-\widehat{\beta}_1 = 0.63$ and $-\widehat{\beta}_2 = -1.35$ and that $-\widehat{\alpha}_0 = 0.067$ and $-\widehat{\delta}_3 = 0.041$. To calibrate the expected benefit duration and unemployment duration (D_B and D), we use the estimated baseline survival curve from Table 2 (and how this survival curve varies with the unemployment rate) to estimate D and D_B for each value of u. We find that D_B/D is roughly 0.75 on average and slightly lower when the unemployment rate is high. Lastly, because of the considerable uncertainty in the appropriate value of the coefficient of relative risk aversion, we report results for a range of the coefficients (from $\gamma = 1$ to $\gamma = 5$).

The results of this calibration exercise are reported in Table 7. The columns report the results for alternative values of the unemployment rate, while the rows report alternative values of the coefficient of relative risk aversion. The numbers are negative on average, which is consistent with an increase in UI benefits (from current levels) reducing welfare. However, when the unemployment rate is high ($u > 9\%$), the marginal welfare gain is positive, though it is small in magnitude across all rows and columns.[46]

At $\bar{u} = 6.7\%$ and $\gamma = 4$, the marginal welfare gain is $-\$0.26$ for a \$10 increase in UI benefits from existing levels. This hypothetical change would reduce utility by the equivalent of a 26 cent decline in the weekly wage.[47] At an unemployment rate of 8.0% (roughly one standard deviation above the mean unemployment rate), the formula implies a marginal welfare gain of $-\$0.12$. Thus, we see that variation in the unemployment rate can meaningfully affect the marginal welfare gain. To give a sense of the quantitative importance, the magnitude is roughly equivalent to a three unit change in the coefficient of relative risk aversion in the model (e.g. from $\gamma = 1$ to $\gamma = 4$), holding the unemployment rate constant. While the previous literature has emphasized the sensitivity of the optimal UI benefit level to the level of risk aversion, our results suggest that the marginal welfare gain is also approximately equally sensitive to labour market conditions. This sensitivity highlights the value of future work that produces more precise estimates of how the duration elasticity and consumption drop at unemployment vary with the unemployment rate.

5. CONCLUSION

In this article, we showed that the relationship between both the moral hazard cost and the consumption smoothing benefit of UI may vary with the unemployment rate, causing the marginal welfare gain from increasing UI benefits to vary with the unemployment rate. This insight structured our two-part empirical strategy which: (1) estimated how the elasticity of unemployment duration with respect to the UI benefit level varies with the unemployment rate and (2) estimated how the consumption drop upon unemployment varies with the unemployment rate. Overall, our empirical findings indicate that the moral hazard cost of UI is lower when unemployment is high, consistent with the speculation of Krueger and Meyer (2002), who claimed that there is likely less of an efficiency loss from reduced search effort by the unemployed when local labour market conditions are poor. On the other hand, we do not find evidence that the consumption smoothing benefit of UI varies with the unemployment rate, although our statistical precision limits strong conclusions. We use our empirical results to calibrate a formula for the marginal welfare gain and find that the welfare gain is modest in magnitude on average, but varies significantly with the unemployment rate.

There are several limitations to our analysis that should be addressed in future work. First, our job search model assumes that job seekers face a constant offer arrival rate that is independent of the length of the unemployment spell. In related work, we conducted an audit study and found evidence that interview rates are negatively correlated with the length of the spell and this correlation is attenuated when the unemployment rate is relatively high (Kroft *et al*., 2013). Studying how such "negative duration dependence" affects optimal UI policy represents an interesting area of future work.

Second, our results are based on variation in *local* labour market conditions. Local recessions and national recessions may have very different underlying mechanisms, and therefore, we believe

46. Online Appendix Table OA.14 shows results from an alternative calibration which is based on an alternative formula for the marginal welfare gain that is robust to allowing for a non-zero coefficient of relative prudence, as described in footnote 9. We assume that the coefficient of relative prudence is equal to $\gamma + 1$, as implied by a CRRA utility function. The results in this table suggest slightly larger welfare gains, although they are fairly similar to the results in Table 7.

47. In most of the scenarios, the marginal welfare gain is negative and small in magnitude, implying that reductions in benefits at current levels would raise welfare by a small amount. The approximation formula used for these calculations is only valid if liquidity constraints are not binding, however; the formula will underestimate the change in welfare from raising UI benefits if liquidity constraints are binding. See discussion of equation (2) and (3) in Section 2.3 for more details.

that caution should be exercised in extrapolating our results to national recessions, as would be the case in any local labour markets analysis.

Third, while we focus on how the marginal welfare gain with respect to the UI benefit level varies over the business cycle, there is little work that jointly characterizes the benefit level and potential duration over the business cycle.[48] In particular, *how much* should benefits be extended during a recession? Is it optimal to extend benefits to 99 weeks, leaving the benefit level unchanged, as occurred during the Great Recession? The answer to this key policy question requires information on the concavity of gains from extending benefits. Existing formulas for marginal welfare gains provide little information on this. Additionally, useful progress on this question requires estimating the consumption smoothing benefits of extended UI benefits (and how this varies with labour market conditions), with model-based estimates that permit counterfactual policy analysis based on non-local changes.

Fourth, our work is complementary to recent theoretical and empirical work that study optimal UI over the business cycle (Schmieder *et al*., 2012; Landais *et al*., 2014). While our article and Schmieder *et al*. (2012) contain empirical estimates of the fiscal costs of UI, across all three papers, our article is unique in presenting direct evidence on the consumption smoothing benefits over the cycle. While we believe that this represents a useful first step and the best available evidence to date, we hope that future research will be able to explore the consumption smoothing benefits of UI using high-frequency administrative data, such as retail scanner data, consumer credit report data, or credit card transaction data.

We view the concept that the consumption smoothing benefit and moral hazard cost of social policies may vary with local labour market conditions as quite general, with applications extending beyond the specific case of UI considered in this article. We find it plausible that the social marginal benefit and social marginal cost of other government policies may also vary with the unemployment rate. For example, if the labour supply response to tax changes is lower during recessions, it may be more efficient to redistribute during recessions. In the case of disability insurance and workers compensation, the benefits and costs of such programs may also be influenced by the business cycle. It would therefore be worthwhile to study how the behavioural responses of these programs vary over the business cycle. In contrast to these examples, there may also be circumstances when the behavioural responses to taxation and transfer programs are *smaller* in boom times; e.g. recent evidence by Edgerton (2010) suggests that financial constraints (or other factors) make firms more responsive to investment incentives during downturns in the business cycle. More generally, to the extent that the marginal benefits and costs of social insurance programs vary with labour market conditions, one should draw caution in comparing estimates across studies to the extent that there are different labour market conditions underlying the estimates.

Acknowledgments. We thank three anonymous referees for very helpful feedback and guidance. We also thank Joe Altonji, Hoyt Bleakley, Judy Chevalier, Jonathan Guryan, Erzo Luttmer, Andrew Metrick, Giuseppe Moscarini, Fiona Scott Morton, and seminar participants at Chicago Booth, Federal Reserve Bank of Philadelphia, Harvard, LSE, MIT, NBER, University of Texas at Austin, University of Toronto, Wisconsin, and Yale for providing excellent comments. We thank Jesse Burkhardt, Robert French, and Christian Goldammer for providing outstanding research assistance, Raj Chetty and Jon Gruber for sharing their data and code with us, and Gad Levanon and Jeanne Shu at the Conference Board for assistance with the Help-Wanted Index. M.J.N. gratefully acknowledges the National Institute of Aging (NIA grant number T32-AG000186) for financial support.

48. Landais (2014) jointly estimates the effect of UI benefit levels and durations on unemployment durations using a regression kink design, but does not study in detail how the effects vary over the business cycle.

Supplementary Data

Supplementary Data are available at *Review of Economic Studies* online.

REFERENCES

ABRAHAM, K. (1987), "Help-Wanted Advertising, Job Vacancies, and Unemployment", *Brookings Papers on Economic Activity*, **1987**, 207–248.

AGUIAR, M., HURST, E. and KARABARBOUNIS, L. (2013), "Time Use During Recessions", *American Economic Review*, **103**, 1664–1696.

ANDERSEN, T. M. and SVARER, M. (2011), "State Dependent Unemployment Benefits", *Journal of Risk and Insurance*, **78**, 325–344.

AUTOR, D. and DUGGAN, M. (2003), "The Rise in the Disability Rolls and the Decline in Unemployment", *Quarterly Journal of Economics*, **118**, 157–206.

AUTOR, D., DORN, D. and HANSON, G. (2013), "The China Syndrome: Local Labor Market Effects of Import Competition in the United States", *American Economic Review*, **103**, 2121–2168.

BAILY, M. N. (1978), "Some Aspects of Optimal Unemployment Insurance", *Journal of Public Economics*, **10**, 379–402.

BARTIK, T. J. (1991), "Who Benefits From State and Local Economic Development Policies" (Kalamazoo, Mich: W.E. Upjohn Institute for Employment Research).

BERTRAND, M. (2004), "From the Invisible Handshake to the Invisible Hand? How Import Competition Changes the Employment Relationship", *Journal of Labour Economics*, **22**, 723–766.

BROWNING, M. and CROSSLEY, T. F. (2001), "Unemployment insurance benefit levels and consumption changes", *Journal of Public Economics*, **80**, 1–23.

CHETTY, R. (2006), "A General Formula for the Optimal Level of Social Insurance", *Journal of Public Economics*, **90**, 1879–1901.

CHETTY, R. (2008), "Moral Hazard versus Liquidity and Optimal Unemployment Insurance", *Journal of Political Economy*, **116**, 173–234.

CHETTY, R. (2009), "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced Form Methods", *Annual Review of Economics*, **1**, 451–488.

CHETTY, R. and LOONEY, A. (2006), "Consumption Smoothing and the Welfare Consequences of Social Insurance in Developing Economies", *Journal of Public Economics*, **90**, 2351–2356.

CRÉPON, B., DUFLO, E., GURGAND, M., RATHELOT, R. and ZAMORA, P. (2013), "Do Labour Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment", *Quarterly Journal of Economics*, **128**.

CURRIE, J. and GRUBER, J. (1996), "Saving babies: The efficacy and cost of recent changes in the Medicaid eligibility of pregnant women", *The Journal of Political Economy*, **104**, 1263–1296.

DIAMOND, P. A. (1982), "Wage Determination and Efficiency in Search Equilibrium", *The Review of Economic Studies*, **49**, 217–227.

EDGERTON, J. (2010), "Investment Incentives and Corporate Tax Asymmetries", *The Journal of Public Economics*, **94**, 936–952.

ELMENDORF, D. (2010), "Policies for Increasing Economic Growth and Employment in the Short Term", Congressional Budget Office testimony.

FARBER, H. and VALLETTA, R. (2013), "Do Extended Unemployment Benefits Lengthen Unemployment Spells? Evidence from Recent Cycles in the U.S. Labor Market" (Working Paper No. 19048, NBER).

GRUBER, J. (1997), "The Consumption Smoothing Benefits of Unemployment Insurance", *American Economic Review* **87**, 192–205.

HAGEDORN, M., KARAHAN, F., MANOVSKII, I. and MITMAN, K. (2013), "Unemployment benefits and unemployment in the great recession: The role of macro effects" (Working Paper No. 19499, NBER).

JOHNSON, A. C. and MAS, A. (2015), "Potential Unemployment Insurance Duration and Labour Supply: The Individual and Market-Level Response to a Benefit Cut" (Working Paper, Princeton University).

JURAJDA, S. and TANNERY, F. J. (2003), "Unemployment Durations and Extended Unemployment Benefits in Local Labour Markets", *Industrial and Labour Relations Review*, **56**, 324–348.

HUNGERBÜHLER, M., LEHMANN, E., PARMENTIER, A., LINDEN, B. (2006), "Optimal Redistributive Taxation in a Search Equilibrium Model", *The Review of Economic Studies*, **73**, 743–767.

KILEY, M. T. (2003), "How Should Unemployment Benefits Respond to the Business Cycle?", *The B.E. Journal of Economic Analysis and Policy*, **3**, 1–30.

KROFT, K., LANGE, F. and NOTOWIDIGDO, M. (2013), "Duration Dependence and Labour Market Conditions: Evidence from a Field Experiment", *Quarterly Journal of Economics*, **128**, 1123–1167.

KRUEGER, A. and MEYER, B. (2002), "Labour Supply Effects of Social Insurance" in Auerbach, A. and Feldstein, M., (eds) *Handbook of Public Economics* Vol 4. (Amsterdam: North-Holland).

LALIVE, R., LANDAIS, C. and ZWEIMULLER, J. (2015), "Market Externalities of Large Unemployment Insurance Extension Programs", *American Economic Review*, (forthcoming).

LANDAIS, C. (2014), "Assessing the Welfare Effects of Unemployment Benefits Using the Regression Kink Design", *American Economic Journal: Economic Policy*, forthcoming.

LANDAIS, C., MICHAILLAT, P. and SAEZ, E. (2014), "Optimal Unemployment Insurance over the Business Cycle" (Working Paper No. 16526, NBER).
LENTZ, R. and TRANAES, T. (2005), "Job Search and Savings: Wealth Effects and Duration Dependence", *Journal of Labour Economics*, **23**, 467–489.
MARINESCU, I. (2015), "The General Equilibrium Impacts of Unemployment Insurance: Evidence from a Large Online Job Board", (Working Paper).
MEYER, B. (1990), "Unemployment Insurance and Unemployment Spells", *Econometrica*, **58**, 757–782.
MIAN, A. and SUFI, A. (2012), "What Explains High Unemployment? The Aggregate Demand Channel", (Working Paper No. 17830, NBER).
MOFFITT, R. (1985), "Unemployment Insurance and the Distribution of Unemployment Spells", *Journal of Econometrics*, **28**, 85–101.
MORTENSEN, D. T. and PISSARIDES, C. A. (1999), "New Developments in Models of Search in the Labour Market", *Handbook of Labour Economics*, **3**, 2567–2627.
MUELLER, A. (2012), "Separations, Sorting and Cyclical Unemployment" (Working Paper, Columbia University).
PIKETTY, T. and SAEZ, E. (2012), "Optimal Labour Income Taxation" (Working Paper No. 18521, NBER).
PISSARIDES, C. (2000), *Equilibrium Unemployment Theory* 2nd edn (Cambridge, Massachusetts and London, England: MIT Press).
ROTHSTEIN, J. (2011), "Unemployment insurance and job search in the great recession", *Brookings Papers on Economic Activity*, 143–196, Fall 2011.
SANCHEZ, J. (2008), "Optimal State-Contingent Unemployment Insurance", *Economics Letters*, **98**, 348–357.
SCHMIEDER, J. F., VON WACHTER, T. and BENDER, S. (2012), "The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates over Twenty Years", *The Quarterly Journal of Economics*, **127**, 701–752.
SPINNEWIJN, J. (2015), "Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs", *Journal of the European Economic Association*, **13**, 130–167.