Optimal Social Insurance with Heterogeneity

Isaiah Andrews and Conrad Miller[1]

June 3rd, 2013

Abstract

We analyze the effect of heterogeneity on the widely used analyses of Baily (1978) and Chetty (2006) for optimal social insurance. The basic Baily-Chetty formula is robust to heterogeneity along many dimensions but requires that risk aversion be homogeneous. We extend the Baily-Chetty framework to allow for arbitrary heterogeneity across agents, particularly in risk preferences. We find that heterogeneity in risk aversion affects welfare analysis through the covariance of risk aversion and consumption drops, which measures the extent to which larger risks are borne by more risk tolerant workers. Calibrations suggest that this covariance effect may be large.

Keywords: sufficient statistics; heterogeneity; unemployment insurance

1   Introduction

In this paper we investigate the implications of population heterogeneity for the sufficient statistic approach to welfare analysis developed in Baily (1978) and generalized in Chetty (2006). In a stylized model of unemployment, Baily obtains a simple formula for the optimal unemployment insurance (UI) benefit as a function of three parameters: (1) the elasticity of unemployment durations with respect to benefits; (2) the drop in consumption associated with unemployment as a function of UI benefits; and (3) the coefficient of relative risk aversion.[2] This framework has been applied extensively in both empirical and theoretical work on social insurance (e.g., Gruber 1997; Chetty and Looney 2006; Chetty and Saez 2010; Landais et al. 2010; Gross and Notowidigdo 2011; Kroft and Notowidigdo 2011; Schmieder et al. 2012). One potential shortcoming of the Baily (1978) and Chetty (2006) results is that they are derived using models where agents are homogeneous along some important dimensions, while in practice heterogeneity seems likely to be empirically relevant. In the UI context, for example, there may be heterogeneity across agents in search costs, ability to smooth consumption (e.g. via borrowing, savings, or spousal labor supply), and local risk aversion. This heterogeneity can affect how individuals value UI, and the need to aggregate heterogeneous individual preferences may significantly complicate welfare analysis.

[2]Chetty (2006) generalizes the intuition behind Baily's stylized model, demonstrating that with minor adjustments, the Baily formula holds in a more general setting that allows for a large class of realistic extensions, including arbitrary borrowing constraints, leisure benefits of unemployment, and endogenous asset accumulation or human capital investment.

As noted by Chetty (2006), the Baily-Chetty formulas are robust to a limited degree of heterogeneity provided one plugs in appropriate population averages.[3] This result, however, requires the assumption that agents share a common coefficient of relative risk aversion. This homogeneity assumption is used to relate differences in average utility across states to differences in average consumption. By considering the joint distribution of risk aversion and consumption drops, we extend the Baily-Chetty framework to allow for arbitrary heterogeneity in risk preferences, and hence unrestricted heterogeneity across agents.

We show that several different approaches to calculating aggregate welfare for heterogeneous agents yield equivalent welfare metrics. We find that heterogeneity in risk aversion affects welfare analysis through the covariance between risk aversion and consumption drops in the cross-section of the unemployed. This reflects the fact that unemployment insurance is more valuable if more risk averse agents are subject to larger risks. We refer to this as the *covariance effect*.[4] Our approach easily generalizes to accommodate a number of extensions including UI systems with taxes and benefits that are proportional to wages. Further, we show that our results extend to a heterogeneous version of the rich dynamic model studied by Chetty (2006), allowing for a range of additional behaviors and constraints including private insurance purchases and limits on borrowing.

To explore the potential importance of the covariance effect, we calibrate a stylized model of private consumption smoothing decisions using data on observed household consumption drops associated with unemployment. The results suggest that the covariance effect may be large: for plausible population distributions of risk preferences, we find that accounting for the covariance effect can change the approximate consumption smoothing benefit of UI by more than 50%.

Our results show that the value of social insurance depends on the extent to which risk exposure and risk tolerance are aligned in the unemployed population: for a given distribution of consumption drops, the lower the covariance of consumption drops faced by workers with individual risk aversion, the lower the value of additional social insurance. In contexts where risk aversion, ex-ante risk, and ability to self-insure are largely independent, we would generally expect this covariance to be negative because more risk averse agents will take private actions to reduce their risk. Moreover, we would expect the magnitude of this effect to be larger when workers are better able to self-insure. To take an extreme example, even if most agents are quite risk averse and the average consumption drop associated with unemployment is large, the marginal value of social insurance may be zero if all consumption risk is borne by a risk-neutral subpopulation, as could occur in the presence of actuarially fair private unemployment insurance. Without knowing the joint distribution of risk preferences, ex-ante risk, and ability to smooth consumption, however, the sign and magnitude of the covariance effect are *a priori* ambiguous and may depend on context; estimating this covariance is an important challenge for future research.

There is growing evidence of risk preference heterogeneity in insurance settings. In particular, a recent literature documents heterogeneity in risk preferences in insurance markets estimated using

---

[3]Chetty (2006) footnote 8, page 1894.

[4]As in Chetty (2006), although the model here refers to unemployment shocks the same model can be applied to other types of social insurance by relabeling the shock (e.g. injury or disability).

market outcomes. Cohen and Einav (2007) find evidence of substantial variation in absolute and relative risk aversion using observational data on deductible choice in Israeli auto insurance contracts. Barseghyan et al. (2011) and Einav et al. (2011) explore the stability of individual risk preferences across different insurance and investment domains and find evidence that person-specific risk preferences are significantly correlated across settings. Einav et al. (2011) also find that an individual's revealed risk preferences in every other insurance domain predict that individual's choice in a given insurance domain better than demographic characteristics. Together, these findings suggest that there exists substantial heterogeneity in risk preferences, even within demographic subgroups. We show that this heterogeneity has substantive implications for the Baily-Chetty approach to welfare analysis.

Our study is closely related to Chetty and Saez (2010), who apply a sufficient statistic approach to characterize the welfare gains from social insurance when the private sector provides partial insurance. In this setting, the validity of the Baily-Chetty formula depends crucially on whether private insurance generates moral hazard. If so, the standard formula must be modified to account for fiscal externalities; otherwise, the formula is unaffected. In contrast, while we do not account for fiscal externalities our results highlight the importance of accounting for private insurance markets because they may affect the covariance of risks faced by workers and risk preferences.

This paper also relates to a new literature on tests for efficient risk sharing across households when household preferences for risk are heterogeneous. Standard tests are based on the idea that, under full risk sharing, household consumption should not respond to idiosyncratic shocks after accounting for aggregate shocks. Schulhofer-Wohl (2011) and Mazzocco and Saini (2012) show that, under heterogeneity, such tests may reject efficiency even if households share risk efficiently. In particular, when some households are less risk averse than others, it is Pareto-efficient for those who are less risk averse to bear more aggregate consumption risk (Diamond 1967; Wilson 1968). Acknowledging this, both sets of authors derive tests of efficient risk sharing that allow for risk preference heterogeneity.[5] Similar to the present paper, this literature highlights that it is important to understand the relationship between risks faced by agents and risk preferences when preferences may be heterogeneous.

In the next section, we review the two-period model studied by Chetty (2008) and use this model to illustrate the basic sufficient statistics approach. In section 3, we propose a tractable approach to calculating aggregate welfare for heterogeneous agents. In section 4, using this approach we consider two separate cases: one where taxes are set to be actuarially fair within individual and the other where taxes are set uniformly, allowing for expected transfers between workers. We also discuss practical issues that arise in implementing our formula. In section 5, we show in a simple calibration exercise that the covariance effect is plausibly large. In section 6, we extend our results to cover benefits and taxes proportional to wages and higher order approximations to the utility function, and show that analogs of our baseline results hold in a heterogeneous version of the rich

---

[5]Moreover, the authors are able to reject the null that models constraining household risk preferences to be homogeneous are correctly specified using data from the U.S. and India, respectively.

dynamic model studied by Chetty (2006). Section 7 concludes.

2   The Baily-Chetty Formula

We begin by reviewing the basic sufficient statistics approach of Baily (1978) and Chetty (2006) as applied to unemployment insurance. The starting point of our analysis is the 2-period job search model described in Chetty (2008), which can be used to frame the analyses of Baily (1978) and Chetty (2006). Suppose an agent or worker becomes unemployed at $t = 0$ with assets A. The agent chooses search effort s, where s is normalized to be the probability that the agent finds a job. Let $\psi(s)$ denote the cost of search effort, where $\psi(\cdot)$ is strictly increasing and convex. If the agent remains unemployed at $t = 1$, he receives unemployment benefit b. If the agent successfully finds a job, he receives wage w at time $t = 1$ and pays a tax τ that is used to finance the UI system. Let $u(\cdot)$ denote the agent's utility over consumption at $t = 1$, where $u(\cdot)$ is strictly concave. Let $c_u = A + b$ and $c_e = A + w - \tau$ denote consumption if the agent remains unemployed or finds a job, respectively.

The social planner's problem is to choose the benefit level that maximizes the agent's expected utility subject to a budget balance constraint. In particular, the planner's problem is

$$\max_b \tilde{W}(b) = (1 - s(b))u(c_u) + s(b)u(c_e) - \psi(s(b))$$

such that

$$b(1 - s(b)) = s(b)\tau.$$

Taking the first order condition and applying the envelope theorem yields

$$\frac{d\tilde{W}}{db} = (1 - s)u'(c_u) - su'(c_e)\frac{d\tau}{db}.$$

Following Baily (1978) and Chetty (2006, 2008), consider the money metric marginal utility $\frac{dW}{db}$ obtained by scaling $\frac{d\tilde{W}}{db}$ by the welfare gain from a marginal increase in the wage, $su'(c_e)$. This yields

$$\frac{dW}{db} = \frac{d\tilde{W}/db}{su'(c_e)} = \frac{1 - s}{s}\left\{\frac{u'(c_u) - u'(c_e)}{u'(c_e)} - \frac{\epsilon_{1-s,b}}{s}\right\} \tag{1}$$

where $\epsilon_{1-s,b}$ is the elasticity of the probability of remaining unemployed with respect to the benefit level, or the elasticity of one minus search effort s with respect to benefits b.

In general, this expression does not allow us to analyze the welfare impact of UI benefit changes unless we are willing to assume a particular form for $u(\cdot)$. Rather than choosing a functional form, following Baily (1978) and Chetty (2006) we approximate $u(\cdot)$ by its second-order Taylor expansion around c_e. In particular,

$$u'(c_u) - u'(c_e) \approx u''(c_e)(c_u - c_e)$$

which yields

$$\frac{u'(c_u) - u'(c_e)}{u'(c_e)} \approx -\frac{u''(c_e)}{u'(c_e)} c_e \frac{c_e - c_u}{c_e}$$

$$= \gamma \frac{c_e - c_u}{c_e} = \gamma \Delta \tag{2}$$

where $\gamma = -\frac{u''(c_e)}{u'(c_e)} c_e$ is the coefficient of relative risk aversion evaluated at c_e and $\Delta = \frac{c_e - c_u}{c_e}$ is the proportional consumption drop from unemployment.[6] We are left with the familiar Baily-Chetty formula:

$$\frac{dW}{db} \approx \frac{1-s}{s} \left\{ \gamma \Delta - \frac{\epsilon_{1-s,b}}{s} \right\}. \tag{3}$$

This can be interpreted as the welfare change associated with raising the level of UI payments. Intuitively, the first bracketed term captures the value of transferring money from the agent's employed state to the agent's unemployed state, while the second term reflects the moral hazard cost of raising benefits: increasing benefits by a dollar increases the cost of the program by more than a dollar per unemployed agent because agents may respond to higher benefits by searching less intensely.

3   Welfare Analysis Under Heterogeneity

We are interested in extending the analysis above to contexts with heterogeneous agents. In particular, suppose that instead of a single unemployed agent there is a population of agents indexed by $i \in I$ for some index set I and that the distribution of agents is given by F. By the same logic as above, the marginal utility of agent i from an increase in the UI benefit level is $(1 - s^i)u_i'(c_u^i)$, while the marginal disutility of a tax increase is $s^i u_i'(c_e^i)$. Thus, the welfare change for agent i from a \$1 increase in the UI benefit together with a tax change $\frac{d\tau^i}{db}$ in the employed state is

$$\frac{d\tilde{W}^i}{db} = \left(1 - s^i\right) u_i'(c_u^i) - s^i u_i'(c_e^i) \frac{d\tau^i}{db}. \tag{4}$$

While the same approach as above can be used to approximate this marginal welfare for each agent, to measure the overall effect of a given change in the UI system we need some way to aggregate across individuals. Aggregation of individual welfare gains $\frac{d\tilde{W}^i}{db}$ is complicated by the well-known fact that Von-Neumann Morgenstern utilities are only defined up to an affine transformation: the behavior of an individual with utility function $u_i(\cdot)$ is in every way indistinguishable from that of an individual with utility function

$$\tilde{u}_i(\cdot) = a_1^i + a_2^i u_i(\cdot) \tag{5}$$

for a_1^i a real number and $a_2^i > 0$. Even if we take the utility function for individual i to have some

---

[6]Chetty (2006) notes that if the third order term of $u(c)$ is significant, the approximation should include an additional term that depends on the coefficient of relative prudence, $\rho = -\frac{u'''(c)}{u''(c)} c$, an adjustment for precautionary saving motives. We do not include this term in our baseline analysis, but discuss it in section 6.

cardinal meaning, that is meaning beyond merely representing preferences, we have no hope of recovering quantities that are sensitive to the constants a_1 and a_2 without additional assumptions. Hence, to obtain an empirically implementable expression for aggregate welfare we need either to ensure invariance to individual-specific affine transformations of utility (5) or to impose additional assumptions that rule out problematic transformations.

We see at least three ways to proceed. One is to normalize individual utilities to rule out problematic transformations of the form (5). Another is to choose welfare weights that, by construction, generate invariant measures of aggregate welfare. Finally rather than attempting to aggregate the marginal welfare gains (4) directly, we can instead consider quantities, for example money-metric utilities, that are more directly comparable across individuals and aggregate these. Below, we show that reasonable implementations of all three approaches lead to equivalent expressions for marginal aggregate welfare.

3.1 Normalization of Utilities

To rule out potentially problematic transformations of the form (5), we can normalize the utility function to eliminate a degree of freedom. In particular, note that (4) depends only on the marginal utility $u_i'(\cdot)$ and hence it suffices to rule out multiplication of the utility function by a constant a_2^i. One way to do this is to normalize marginal utility in the employed state to one, effectively fixing $a_2^i = u_i'(c_e^i)^{-1}$. Under this normalization, (4) becomes

$$\frac{d\tilde{W}^i}{db} = \left(1 - s^i\right) \frac{u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau^i}{db}.$$

If we consider a utilitarian welfare metric, which takes aggregate welfare to be $\hat{W} = E\left[\tilde{W}^i\right]$, we have that the welfare change from the reform considered is

$$\frac{d\hat{W}}{db} = E\left[\frac{d\tilde{W}^i}{db}\right] = E\left[\left(1 - s^i\right) \frac{u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau^i}{db}\right]$$

$$= (1 - \bar{s})\, E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right] + (1 - \bar{s}) - E\left[s^i \frac{d\tau^i}{db}\right] \tag{6}$$

where $E[x^i] = \bar{x} = \int_I x^i dF(i)$ denotes the mean of x^i in the population and $E^u[x^i] = \bar{x}^u = \int_I \frac{1-s^i}{1-\bar{s}} x^i dF(i)$ denotes the mean weighted by unemployment probability and is the average one obtains by considering a cross-section of unemployed agents at $t = 1$. Hence, under this normalization we see that marginal aggregate welfare depends on three terms: the unemployment rate at $t = 1$, $1 - \bar{s}$, the average proportional increase in marginal utility from unemployment among the unemployed population, $E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right]$, and the average expected tax increase, $E\left[s^i \frac{d\tau^i}{db}\right]$.

6

3.2 Stabilizing Welfare Weights

We could aggregate the marginal welfare gains (4) without imposing a normalization, but must then choose welfare weights carefully to arrive at an invariant measure of aggregate welfare. To this end, suppose we attach welfare weight α^i to individual i and are interested in measuring $\hat{W} = E\left[\alpha^i \tilde{W}^i\right]$. The change in this aggregate welfare measure from the UI reform considered is

$$\frac{d\hat{W}}{db} = E\left[\alpha^i \frac{d\tilde{W}^i}{db}\right] = E\left[\alpha^i \left(\left(1 - s^i\right) u_i'(c_u^i) - s^i u_i'(c_e^i)\frac{d\tau^i}{db}\right)\right]. \tag{7}$$

Note, however, that if $\alpha^i \left(\left(1 - s^i\right) u_i'(c_u^i) - s^i u_i'(c_e^i)\frac{d\tau^i}{db}\right)$ is non-invariant to the transformation (5) for a positive mass of agents then (7) is non-invariant as well. Hence, invariance of (7) implies that for almost every agent i we have $\alpha^i \frac{d\tilde{W}^i}{db} = \tilde{K}^i$ for a constant $\tilde{K}^i$ which is invariant to (5). Provided $\frac{d\tilde{W}^i}{db} \neq 0$ for almost every agent, we have that $\alpha^i = \tilde{K}^i / \frac{d\tilde{W}^i}{db}$ which can be seen to imply that $\alpha^i = \frac{K^i}{u_i'(c_e^i)}$, again for K^i invariant to (5). Hence, to construct an invariant marginal aggregate welfare (7) all that remains is to pick the constant K^i. The choice $K^i = 1$ yields

$$\alpha^i \frac{d\tilde{W}^i}{db} = \left(1 - s^i\right) \frac{u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau^i}{db}$$

and hence that as in the previous section,

$$\frac{d\hat{W}}{db} = (1 - \bar{s}) E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right] + (1 - \bar{s}) - E\left[s^i \frac{d\tau^i}{db}\right].$$

Note that this result depends on our choice of K^i; as might be expected, other choices of welfare weights yield different expressions for aggregate welfare.

3.3 Aggregating Money-Metric Utilities

A third approach to constructing invariant measures of aggregate benefit from a UI reform is to pick some non-utility measure of benefit that is more easily comparable across individuals and to aggregate that. In this section, we consider aggregating money-metric utilities.

The welfare gain for individual i from the considered UI and tax changes, relative to their welfare gain from a \$1 wage increase, is

$$\frac{dW^i}{db} = \frac{d\tilde{W}^i/db}{s^i u_i'(c_e^i)} = \frac{\left(1 - s^i\right) u_i'(c_u^i)}{s^i u_i'(c_e^i)} - \frac{d\tau^i}{db}.$$

Since agent i regains employment with probability i, the expected cost of delivering this increased wage to agent i in the employed state, holding search behavior fixed, is s^i. Hence, the expected cost, fixing search intensity s^i, of the combination of agent-specific wage increases that changes

each agent's expected utility by the same amount as the proposed UI reform is

$$E\left[s^i \frac{dW^i}{db}\right] = E\left[\frac{(1-s^i)\, u_i'(c_u^i)}{u_i'(c_e^i)} - s^i \frac{d\tau^i}{db}\right]$$

$$= (1-\bar{s})\, E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right] + (1-\bar{s}) - E\left[s^i \frac{d\tau^i}{db}\right]$$

which is exactly the same as the welfare expressions derived above. Note that this also has the interpretation as the cost, holding search behavior fixed, of the combination of agent-specific wage changes such that all agents would be indifferent between this wage change and the proposed reform. The cost per employed agent, which is more closely analogous to the money-metric utility considered in the single-agent case, is then

$$\frac{d\bar{W}}{db} = \frac{1-\bar{s}}{\bar{s}} E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right] + \frac{1-\bar{s}}{\bar{s}} - E^e\left[\frac{d\tau^i}{db}\right] \tag{8}$$

where $E^e[x^i] = \bar{x}^e = \int_I \frac{s^i}{\bar{s}} x^i dF(i)$ denotes an average weighted by agents' job-finding probability. This expression depends on the same terms as (6) above but has a more intuitive interpretation. In particular it depends on the average increase in marginal utility from unemployment among the unemployed population, $E^u\left[\frac{u_i'(c_u^i) - u_i'(c_e^i)}{u_i'(c_e^i)}\right]$, the number of unemployed agents per employed agent, $\frac{1-\bar{s}}{\bar{s}}$, and the average tax increase faced by employed workers, $E^e\left[\frac{d\tau^i}{db}\right]$. Hence, we can see that optimal UI policy based on this metric will balance the consumption smoothing benefit from UI for the unemployed against the marginal disutility of taxes for employed agents, taking into account the relative size of these two populations. This expression will be our focus for the remainder of the analysis, but by the results discussed above one could also multiply all our expressions by $\bar{s}$ and interpret our results as concerning a utilitarian welfare metric under a particular normalization or particular welfare weights.

3.4 Alternative Approaches to Evaluating Aggregate Welfare

Our analysis represents only one of many possible approaches to extending the Baily (1978) formula to models with heterogeneous agents. A natural alternative, implicit in Chetty (2006), is to instead consider the average marginal utility from a given UI and tax change relative to the average marginal utility of a $1 wage increase, i.e.

$$\frac{dW^*}{db} = \frac{E\left[d\tilde{W}^i/db\right]}{E\left[s^i u_i'(c_e^i)\right]} = \frac{E\left[(1-s^i)\, u_i'(c_u^i) - s^i u_i'(c_e^i) \frac{d\tau^i}{db}\right]}{E\left[s^i u_i'(c_e^i)\right]}.$$

The marginal welfare $\frac{dW^*}{db}$ values a UI reform in terms of the wage increase that generates the same average welfare change. We can interpret this quantity as the uniform wage change such that a rational agent would be indifferent between this change and the proposed reform if they knew they would subsequently be assigned an identity i at random according to the population distribution.

In contrast, $\frac{d\bar{W}}{db}$ calculates the wage increase equivalent to the reform for each agent i separately and then aggregates across the population. Analogous to the literature on efficiency in economies with incomplete information we can view $\frac{dW^*}{db}$ as an *ex-ante* money-metric measure of welfare from the proposed reform (see Holmstrom and Myerson, 1983), since in calculating both the utility gain from the proposed UI reform and the marginal utility from a wage increase we average across types i. The marginal welfare $\frac{d\bar{W}}{db}$ can instead be viewed as an *interim* money metric, since the money metric value of the reform for each agent i takes into account all relevant agent characteristics.

While both ex-ante and interim welfare criteria have their advantages, we argue that for the UI analysis considered here the interim money-metric $\frac{d\bar{W}}{db}$ has a number of important strengths. First, if one could survey a cross-section of agents and obtain truthful responses on the i-specific value (in wage terms) of the proposed reform $\frac{d\tilde{W}^i}{db}$ along with job-finding probability s^i this would suffice to calculate $\frac{d\bar{W}}{db}$ directly without any assumptions. Further, as we demonstrate in the next section even without such survey data obtaining approximations to $\frac{d\bar{W}}{db}$ for the heterogeneous case is straightforward, and the analysis readily accommodates extensions to more general contexts, for example benefits and taxes proportional to heterogeneous wages as discussed in section 6. While as noted in Chetty (2006) (footnote 8), if we take the tax change $\frac{d\tau^i}{db}$ to be uniform $\frac{d\tau^i}{db} = \frac{d\tau}{db}$ and assume that agents share a common coefficient of relative risk aversion $\gamma^i = \gamma$ we can use the analysis of that paper to approximate $\frac{dW^*}{db}$, given the empirical literature documenting preference heterogeneity in a wide range of settings the assumption of homogeneous risk aversion seems unappealing.[7] In contrast to the analysis for $\frac{d\bar{W}}{db}$, extending the approximations for $\frac{dW^*}{db}$ to accommodate unobserved preference heterogeneity is far from straightforward.

4 Baily-Chetty Under Heterogeneity

To extended the Baily-Chetty analysis to heterogeneous agents using marginal aggregate welfare (8), for a given collection of i-specific tax changes $\frac{d\tau^i}{db}$ we can apply the same second-order approximations to the utility function as in the single-agent case above. Since under our assumptions agents may exhibit different behavioral responses to changes in the level of the UI benefit, some changes to the UI system may generate net transfers across individuals in expectation and these transfers may matter for our welfare analysis. To separate these *transfer effects* from the consumption smoothing benefit of a given UI reform, we first calculate the marginal welfare gain from an increase in benefits in an actuarially-fair UI system that generates no transfers across agents in expectation and then turn to a UI reform with a uniform tax change $\frac{d\tau^i}{db} = \frac{d\tau}{db}$.

4.1 Actuarially Fair UI Under Heterogeneity

We begin by considering the case where taxes are actuarially fair within individual: each individual's taxes in the employed state reflect their personal search intensity in the unemployed state so there

---

[7]Note, moreover, that we still need to impose an appropriate normalization on the utility function to rule out transformations of the form (5).

is expected budget balance for each individual, that is

$$b(1 - s^i(b)) = s^i(b)\tau^i.$$

This implies that there are no expected transfers between individuals. Under this restriction one can show that a \$1 increase in the UI benefit must be accompanied by a tax increase

$$\frac{d\tau^i}{db} = \frac{1 - s^i}{s^i}\left(1 + \frac{\epsilon^i_{1-s,b}}{s^i}\right)$$

on individual i in the employed state. To approximate the change in aggregate welfare (8) from this UI reform, we consider the same second-order approximation to the utility function as in the single-agent case, $u'_i(c_u) - u'_i(c_e) \approx u''_i(c_e)(c_u - c_e)$. The marginal aggregate welfare (8) from a change in the UI benefit level can then be approximated by

$$
\begin{aligned}
\frac{d\bar{W}}{db} &\approx \frac{1 - \bar{s}}{\bar{s}}E^u\left[\gamma^i\frac{c^i_e - c^i_u}{c^i_e}\right] + \frac{1 - \bar{s}}{\bar{s}} - E^e\left[\frac{1 - s^i}{s^i}\left(1 + \frac{\epsilon^i_{1-s,b}}{s^i}\right)\right] \\
&= \frac{1 - \bar{s}}{\bar{s}}E^u\left[\gamma^i\Delta^i - \frac{\epsilon^i_{1-s,b}}{s^i}\right] \\
&= \frac{1 - \bar{s}}{\bar{s}}\left(\bar{\gamma}^u\bar{\Delta}^u + \text{cov}^u\left(\gamma^i, \Delta^i\right) - E^u\left[\frac{\epsilon^i_{1-s,b}}{s^i}\right]\right).
\end{aligned}
\tag{9}
$$

where $\text{cov}^u\left(x^i, y^i\right) = E^u[x^iy^i] - E^u[x^i]E^u[y^i]$ is defined analogously to $E^u[x^i]$ and corresponds to the covariance in a cross-section of the unemployed.

This expression for marginal welfare $\frac{d\bar{W}}{db}$ depends on three terms. The first term, $\bar{\gamma}^u\bar{\Delta}^u$, is the product of the average risk aversion in the cross-section of the unemployed and the average consumption drop. This term is analogous to the term $\gamma\Delta$ in the homogeneous case, but the weighting here is important. In particular, if more risk averse agents are more likely to be unemployed, this increases the value of raising the UI benefit all else equal. The second term, $\text{cov}^u\left(\gamma^i, \Delta^i\right)$, captures the covariance between risk aversion and consumption drops in the cross-section of the unemployed, and reflects the fact that unemployment insurance is more valuable if more risk averse agents are subject to larger risks. We refer to this new term as the *covariance effect*. The last term in the expression, $E^u\left[\frac{\epsilon^i_{1-s,b}}{s^i}\right]$, measures the behavioral response to higher benefits, and depends on the joint distribution of individual-specific elasticities $\epsilon^i_{1-s,b}$ and job-finding rates s^i. The form of this term reflects our assumption of actuarial fairness, and as we discuss in the next section relaxing this restriction allows us to obtain a more tractable expression. Nonetheless, we can see that if the population of agents is completely homogeneous (9) simplifies to (3), confirming that our analysis generalizes the Baily-Chetty formula to heterogeneous agents.

4.2 UI with Uniform Taxes Under Heterogeneity

In the previous section we considered actuarially fair UI systems, which set the tax τ^i on each agent based on that agent's probability of remaining unemployed. We may also be interested in UI systems that are not actuarially fair within individual. In particular, the actuarially fair tax depends not only on the individual-level elasticity of unemployment probability with respect to UI benefits, $\epsilon^i_{1-s,b}$, but on the individual-level probability of regaining employment s^i. In contexts with unrestricted heterogeneity identifying the joint distribution of these objects, to say nothing of their values of each individual, poses a daunting challenge. In this section we consider a UI change that, while maintaining budget balance in the aggregate, need not be actuarially fair on an individual level and hence may generate net transfers across individuals through the UI system. In particular, we consider UI systems such that the tax τ is uniform across individuals.

4.2.1 Transfers with Uniform Taxes

If we allow net transfers across individuals through the UI system, budget balance requires only that

$$b\left(1 - \bar{s}(b)\right) = \bar{s}(b)\tau$$

which implies that

$$\frac{d\tau}{db} = \frac{1 - \bar{s}}{\bar{s}}\left(1 + \frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right)$$

where $\epsilon_{1-\bar{s},b}$ is the elasticity of the unemployment rate in the second period (i.e. $1 - \bar{s}$) with respect to the level of UI benefits. Hence, we can see that the money-metric value of the expected transfer to individual i due to a \$1 increase in the UI benefit together with a uniform change in τ is

$$\frac{d\tau^i}{db} - \frac{d\tau}{db} = \frac{1 - s^i}{s^i}\left(1 + \frac{\epsilon^i_{1-s,b}}{s^i}\right) - \frac{1 - \bar{s}}{\bar{s}}\left(1 + \frac{\epsilon_{1-\bar{s},b}}{\bar{s}}\right).$$

Since the second term is constant across i, variation in this term across individuals is driven entirely by the behavior of the first term. In particular, individuals who are likely to remain unemployed (so $\frac{1-s^i}{s^i}$ is large) or who have a large behavioral response to increased benefits (so $\epsilon^i_{1-s,b}$ is large) will be subsidized by those whose probability of unemployment and behavioral response are smaller.

4.2.2 Welfare Analysis with Uniform Taxes

Again taking a second-order approximation to the utility function for each individual i, we can approximate the marginal welfare gain from an increase in the level of UI benefits together with a

uniform change in the tax rate as

$$\frac{d\bar{W}}{db} \approx \frac{1-\bar{s}}{\bar{s}} E^u \left[\gamma^i \frac{c_e^i - c_u^i}{c_e^i} \right] + \frac{1-\bar{s}}{\bar{s}} - E^e \left[\frac{1-\bar{s}}{\bar{s}} \left(1 + \frac{\epsilon_{1-\bar{s},b}}{\bar{s}} \right) \right]$$

$$= \frac{1-\bar{s}}{\bar{s}} E^u \left[\gamma^i \Delta^i - \frac{\epsilon_{1-\bar{s},b}}{\bar{s}} \right]$$

$$= \frac{1-\bar{s}}{\bar{s}} \left(\bar{\gamma}^u \bar{\Delta}^u + \text{cov}^u \left(\gamma^i, \Delta^i \right) - \frac{\epsilon_{1-\bar{s},b}}{\bar{s}} \right). \tag{10}$$

Hence, the welfare change from an increase in the UI benefit again depends on three terms: the product of the average risk aversion and average consumption drop in the unemployed population $\bar{\gamma}^u \bar{\Delta}^u$, the covariance between consumption drops and risk aversion $\text{cov}^u \left(\gamma^i, \Delta^i \right)$ and, distinct from the actuarially fair case, $\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}$, the elasticity of the unemployment rate with respect to UI benefits divided by the average job finding probability.

If we consider a homogeneous population (10) again simplifies to the Baily-Chetty sufficient statistic (3) for the single agent case. If agents are heterogeneous but risk aversion and consumption drops are uncorrelated, the covariance term disappears and we have that for uniform taxes

$$\frac{d\bar{W}}{db} \approx \frac{1-\bar{s}}{\bar{s}} \left(\bar{\gamma}^u \bar{\Delta}^u - \frac{\epsilon_{1-\bar{s},b}}{\bar{s}} \right),$$

which is a simple modification of the familiar Baily-Chetty formula where we have replaced each term by an analogous population quantity. To accommodate arbitrary heterogeneity we require only one additional term: the covariance between risk aversion and consumption drops in the unemployed population.

4.3 Implementing the Formulas

To apply these formulas we need estimates of four terms: the weighted average risk aversion, the weighted average consumption drop, the appropriate tax and transfer term, and the weighted covariance of risk aversion and consumption drops.[8] While the term $E^u \left[\frac{\epsilon_{1-s,b}^i}{s^i} \right]$ depends on individual-level elasticities and job-finding probabilities and may be quite difficult to calculate in practice, $\frac{\epsilon_{1-\bar{s},b}}{\bar{s}}$ depends only on aggregate quantities and estimates are available in the literature (e.g. Meyer 1990).

The final term, the covariance of individual risk aversion and consumption drops, is novel and does not enter in the homogeneous case. To the best of our knowledge, this covariance has not been investigated empirically. Note that the sign of $\text{cov}^u \left(\gamma^i, \Delta^i \right)$ is *a priori* ambiguous. If agents were identical except for their risk preferences, we might expect that $\text{cov}^u \left(\gamma^i, \Delta^i \right) \leq 0$ given that agents with higher relative risk aversion value proportional consumption smoothing more. However, this basic intuition may break down given other plausible types of heterogeneity. For example, suppose

---

[8]It is worth emphasizing that average relative risk aversion $\bar{\gamma}^u$ will not, in general, correspond to population estimates derived under a homogeneity assumption. For example, Cohen and Einav (2007) compute a 'back-of-the-envelope' absolute risk aversion estimate using population averages that is more than 6 times smaller than their estimate for mean absolute risk aversion derived from a structural model that allows for individual-level heterogeneity.

poorer workers are both more risk averse (locally, e.g. risk aversion is declining in wealth) and less able to smooth consumption across employment states (e.g. have less savings, poorer access to credit markets, or are less able to substitute labor supply across household members). This can reverse the above intuition and generate $\text{cov}^u\left(\gamma^i, \Delta^i\right) \geq 0$. In general, the sign and magnitude of $\text{cov}^u\left(\gamma^i, \Delta^i\right)$ may depend on the context.

5 Calibrating the Covariance Effect

Our analysis above demonstrates that when agents are heterogeneous, the optimal benefit depends on the covariance between relative risk aversion and consumption drops across agents. This suggests that further study of the covariance effect is necessary before we can apply the Baily-Chetty approach with confidence. However, the approach already relies on several approximations–perhaps we can ignore the covariance term without introducing substantive bias to the welfare analysis. To get a sense of how large this covariance may be, we use data on observed consumption drops associated with unemployment and plausible population distributions of risk preferences to calibrate a simple model of private consumption smoothing decisions.

The model is two periods. In the first period, a set of heterogeneous workers draw implicit 'unemployment insurance' prices randomly from a common distribution. We can think of these prices as reflecting the effective price of consumption smoothing, whether by borrowing, saving, private insurance, informal insurance (e.g. risk sharing across households), spousal labor supply, or another mechanism. Given this price, each agent decides how much insurance to purchase. In the second period, each agent faces an unemployment shock with some probability, and consumes his available resources. In the absence of insurance, the shock reduces the resources available to consume.

We assume each agent has CRRA preferences over consumption, $u(c) = \frac{c^{1-\gamma_i}}{1-\gamma_i}$ with relative risk aversion γ_i. Except for their preferences over consumption, agents are identical *ex-ante*. We also assume that each agent would face the same consumption profile in the absence of insurance, c_e and c_u, and has the same probability of becoming unemployed, $1-s$. We calibrate the distribution of prices to rationalize the observed distribution of consumption drops given the distribution of risk preferences that we assume. The object of interest is the covariance between observed consumption drops for those who become unemployed and their relative risk aversion.

Mathematically, each agent solves the optimization problem

$$\max_{\alpha} su(c_e - \alpha p_i) + (1-s)u(c_u + (1-p_i)\alpha)$$

where $p_i \in (0,1)$ is the per-unit price of insurance faced by agent i and α is the amount of insurance purchased.

Taking the first order condition yields

$$psu'(c_e - \alpha p) = (1-p)(1-s)u'(c_u + (1-p)\alpha)$$

$$\Rightarrow \Delta = 1 - \left(\frac{ps}{(1-p)(1-s)} \right)^{-\frac{1}{\gamma}}$$

where Δ is the consumption drop, or

$$\Delta_i = \frac{c_e - \alpha p - (c_u + (1-p)\alpha)}{c_e - \alpha p}$$

$$= \frac{c_e - c_u - \alpha}{c_e - \alpha p}.$$

Intuitively, higher γ means an agent values consumption smoothing more and will make insurance purchases such that Δ is closer to zero (i.e. $|\Delta|$ is decreasing in γ). Hence, with implicit prices that tend to be actuarially unfair, this simple model should produce a negative covariance between γ and consumption drops across agents. Note that these predictions do not depend on the particular form of preferences used here. For example, applying a second-order approximation for a general utility function $u(\cdot)$ as above yields

$$psu'(c_e - \alpha p) = (1-p)(1-s)u'(c_u + (1-p)\gamma)$$

$$\Rightarrow \Delta \approx \frac{1}{\gamma} \left(\frac{ps}{(1-p)(1-s)} - 1 \right).$$

Again, it can be shown that $|\Delta|$ is decreasing in γ.

To measure the distribution of consumption drops, we use data from the Panel Study of Income Dynamics (PSID) from 1968-1993. The PSID surveys a sample of families, and includes information on household demographics, labor market outcomes, and consumption. To measure consumption, we follow Gruber (1997) and use self-reported household food consumption expenditures deflated using the CPI. This includes the amount usually spent on food both at home and away from home, and the value of food stamps used. The sample consists of all heads of household who are employed at time $t-1$ and unemployed at time t. Following the literature (e.g., Gruber 1997; Chetty and Sziedl 2007; Kroft and Notowidigdo 2011), we approximate proportional consumption drops by the change in log consumption from $t-1$ to t. Following Gruber (1997), we (a) drop observations where food consumption is imputed; and (b) drop observations with more than a three-fold change in total food consumption. Following Chetty and Sziedl (2007), we also exclude households that change in size between years. This yields a mean log consumption change of 8.2 log points. The estimated dispersion is substantial: the standard deviation of these measured log consumption changes is 42 log points.

Consumption data is typically measured with considerable error, and so will tend to overstate the variability of consumption drops. Ahmed et al. (2010) analyze a Canadian survey similar to the Consumer Expenditures Survey (CEX) and suggest that about 75% of variation in consumption

14

data is due to measurement error. To correct for this measurement error, we take the true standard deviation of log consumption changes to be 25% of the measured variation, or about 10.5 log points.

Given that the typical UI replacement rate is around 50%, we normalize $c_e = 1$ and take $c_u = 0.50$. As in Chetty (2008), we take $s = 0.946$.

Given our uncertainty about the distribution of risk preferences in the population, we try a variety of distributions. For half of the calibrations, we assume that γ is uniformly distributed and for the other half, following Cohen and Einav (2007), we assume γ follows a lognormal distribution. We choose three ranges of values for each distribution type, corresponding to the 5th-95th percentile range when γ is lognormally distributed. The first range we try is 1 to 5, the range of values for γ typically seen in calibrations.[9] Motivated by the work of Chetty and Sziedl (2007), who argue that short-run consumption commitments can substantially amplify risk aversion over moderate stakes, we also try a larger range of 1 to 20.[10] In addition, we try an intermediate range of 1 to 10.

We assume that prices are drawn i.i.d. from a normal distribution with mean μ_p and standard deviation σ_p, and then censor these prices at 0.01 and 0.99.[11] We find the values for μ_p and σ_p that match the mean and standard deviation of log consumption changes found in the PSID to those implied by the simulations.

Calibration results are presented in Table 1, and are based on 10 million i.i.d. draws for each case. There are two main results to note. First, and perhaps unsurprisingly, the larger the spread of γ relative to the mean, or $\frac{\sigma_\gamma}{\bar{\gamma}}$, the larger in magnitude the covariance term is relative to to the mean term, $\bar{\gamma} \times \bar{\Delta}$. Second, the covariance term is substantial in magnitude, ranging from 15% to 54% of the mean term, $\bar{\gamma} \times \bar{\Delta}$. Hence, our calibrations suggest that the covariance effect may have empirically significant implications for the Baily-Chetty approach to welfare analysis.

While the correlation between γ and Δ is negative in this simple example, recall that the sign of the covariance term in a more general model is ambiguous. In particular, the probability of an unemployment shock and the implicit price of consumption smoothing may vary arbitrarily with local risk preferences.

6   Extensions

In this section we consider three extensions to the analysis developed above. In the first subsection we extend our sufficient statistics approach to accommodate taxes and benefits that are proportional to wages, as in many real-world UI systems. Next, we follow Chetty (2006) and consider the impact of third-order terms in the utility function, which Chetty argues may have a quantitatively

---

[9]For example, Gruber (1997) uses the range 1 to 4 and Kroft and Notowidigdo (2011) use the range 2 to 4. Chetty (2008) mentions that his results are consistent with $\gamma \approx 5$, while a more parametric approach in Chetty (2003) implies a γ of around 7. Both are estimated in the unemployment insurance context.

[10]While this dispersion is substantial, it is conservative compared to the distribution estimated in Cohen and Einav (2007). They find that, in the context of Israeli auto insurance, the standard deviation of absolute risk aversion is *an order of magnitude* larger than the mean. In the UI context, the empirical results in Chetty (2008) are consistent with substantial risk preference heterogeneity by (predicted) wealth quartile. Also see Barsky et al. (1997) who find evidence of substantial heterogeneity in implied risk aversion using survey responses.

[11]The estimates are generally insensitive to the censoring points chosen.

important effect in some plausible cases. Finally, we show that using results from Chetty (2006) we can apply our approach in a rich class of dynamic models, allowing heterogeneity-robust sufficient statistics to be used in a far wider range of contexts than the simple static model discussed above.

6.1 UI with Proportional Benefits and Taxes

In our analysis above we consider a UI system with a constant benefit b and constant tax τ, consistent with the canonical Baily-Chetty model. In practice, however, UI benefits and taxes are often set proportional to individual wages, at least up to a cap. While this makes little difference in the single worker case, it affects the welfare analysis when workers are heterogeneous. In this section, we consider a UI system where benefits and taxes are proportional to the wage rate.

As a first step, again consider the single worker case. Let b_ρ and τ_ρ denote the UI replacement rate and tax rate. We now have that $c_u = A + b_\rho w$ and $c_e = A + w(1 - \tau_\rho)$. Consequently, the marginal welfare gain from an increase in b_ρ is

$$\frac{d\tilde{W}_\rho}{db_\rho} = w(1-s)u'(c_u) - wsu'(c_e)\frac{d\tau_\rho}{db_\rho}$$

and we arrive at the original Baily-Chetty formula after applying a modified scaling to obtain a money metric. In particular, instead of scaling by $su'(c_e)$, the welfare gain from a marginal increase in the wage, we scale by $wsu'(c_e)$, the welfare gain from a marginal *proportional* increase in the wage.[12]

As in section 3, to conduct welfare analysis for heterogeneous agents we need some way to meaningfully aggregate welfare. We will focus on the generalization of the money-metric approach discussed in section 3.3, but suitable extensions of the approaches discussed in sections 3.1 and 3.2 will yield equivalent analyses. To aggregate money metric utilities, note first that the welfare gain for individual i from a given proportional UI and tax change, relative to their welfare gain from a proportional increase in the wage, is

$$\frac{dW_\rho^i}{db_\rho} = \frac{d\tilde{W}_\rho^i/db_\rho}{s^i w^i u_i'(c_e^i)} = \frac{\left(1 - s^i\right)u_i'(c_u^i)}{s^i u_i'(c_e^i)} - \frac{d\tau_\rho^i}{db_\rho}.$$

Since agent i regains employment with probability s^i and earns wage w^i while employed, the expected cost of delivering a proportional increase in agent i's wage, holding search behavior fixed, is $s^i w^i$. Hence, the expected cost of the combination of agent-specific proportional wage changes

---

[12]Note that $\epsilon_{1-s,b_\rho} = \epsilon_{1-s,b}$.

with the same effect on each agent's expected utility as the proposed UI change is

$$E\left[s^i w^i \frac{dW^i_\rho}{db_\rho}\right] = E\left[\frac{(w^i - s^i w^i)\, u'_i(c^i_u)}{u'_i(c^i_e)} - s^i w^i \frac{d\tau_\rho}{db_\rho}^i\right]$$

$$= (\bar{w} - \overline{sw})\, E^{u,w}\left[\frac{u'_i(c^i_u) - u'_i(c^i_e)}{u'_i(c^i_e)}\right] + (\bar{w} - \overline{sw}) - E\left[s^i w^i \frac{d\tau^i}{db}\right]$$

where $E^{u,w}[x^i] = \bar{x}^{u,w} = \int_I \frac{w^i - s^i w^i}{\bar{w} - \overline{sw}} x^i dF(i)$ denotes the average weighted by wage earnings lost due to unemployment. This is the average one obtains by considering a cross-section of unemployed agents at $t = 1$ weighted by individual wages. Hence the expected cost, relative to total wage income, for the combination of wage changes that would increase each individual's utility by the same amount as a one percentage point increase in the UI replacement rate together with individual-specific proportional tax changes $\frac{d\tau_\rho}{db_\rho}^i$ is

$$\frac{d\bar{W}_\rho}{db_\rho} = E\left[\frac{s^i w^i}{\overline{sw}} \frac{dW^i_\rho}{db_\rho}\right] = \frac{\bar{w} - \overline{sw}}{\overline{sw}} E^{u,w}\left[\frac{u'_i(c^i_u) - u'_i(c^i_e)}{u'_i(c^i_e)}\right]$$

$$+ \frac{\bar{w} - \overline{sw}}{\overline{sw}} - E^{e,w}\left[\frac{d\tau_\rho}{db_\rho}^i\right] \tag{11}$$

where $E^{e,w}[x^i] = \int_I \frac{s^i w^i}{\overline{sw}} x^i dF(i)$ is defined analogously to $E^{u,w}[x^i]$. This term is closely related to (8), the primary difference being that all quantities in (11) depend on wages. In particular, we can see that the marginal increase in aggregate welfare from a given proportional UI reform depends on the average increase in marginal utility from unemployment weighted by lost wage earnings $E^{u,w}\left[\frac{u'_i(c^i_u) - u'_i(c^i_e)}{u'_i(c^i_e)}\right]$, the ratio of wage earnings lost to unemployment to total wage earnings $\frac{\bar{w} - \overline{sw}}{\overline{sw}}$, and the increase in expected tax payments relative to total wage earnings $E^{e,w}\left[\frac{d\tau^i}{db}\right]$.

To obtain usable approximations to (11) we follow the same approach adopted in section 4. In particular, for actuarially fair proportional taxes the welfare change from a one percentage point increase in the replacement rate is

$$\frac{d\bar{W}_\rho}{db_\rho} \approx \frac{\bar{w} - \overline{sw}}{\overline{sw}}\left(\bar{\gamma}^{u,w}\bar{\Delta}^{u,w} + \text{cov}^{u,w}\left(\gamma^i, \Delta^i\right) - E^{u,w}\left[\frac{\epsilon^i_{1-s,b_\rho}}{s^i}\right]\right) \tag{12}$$

where $\text{cov}^{u,w}\left(x^i, y^i\right) = E^{u,w}[x^i y^i] - E^{u,w}[x^i]E^{u,w}[y^i]$ corresponds to the covariance in a cross-section of the unemployed weighted by individual wages. As we'd expect this expression is similar to (9), the corresponding formula for constant benefits and taxes. However, instead of weighting population averages by unemployment probability $1 - s^i$, we weight by expected wage earnings lost due to unemployment, $w^i - s^i w^i$.

As before, it is straightforward to extend our analysis to the uniform tax case. With uniform proportional taxes, budget balance requires that

$$b_\rho(\bar{w} - \overline{sw}(b_\rho)) = \overline{sw}(b_\rho)\tau_\rho$$

17

which implies that

$$\frac{d\tau_\rho}{db_\rho} = \frac{\bar{w} - \overline{sw}}{\overline{sw}} \left(1 + \epsilon_{\bar{w} - \overline{sw}, b_\rho} \frac{\bar{w}}{\overline{sw}} \right)$$

where $\epsilon_{\bar{w} - \overline{sw}, b}$ is the elasticity of lost wage earnings with respect to the replacement rate. Hence, for uniform proportional taxes the welfare gain associated with a one percentage point increase in the replacement rate is

$$\frac{d\bar{W}_\rho}{db_\rho} \approx \frac{\bar{w} - \overline{sw}}{\overline{sw}} \left(\bar{\gamma}^{u,w} \bar{\Delta}^{u,w} + \mathrm{cov}^{u,w} \left(\gamma^i, \Delta^i \right) - \epsilon_{\bar{w} - \overline{sw}, b_\rho} \frac{\bar{w}}{\overline{sw}} \right). \tag{13}$$

This expression is similar to (10), the corresponding formula for a constant benefit and tax system. Again, to accommodate heterogeneity in wages we weight all averages by lost wage earnings rather than unemployment duration. In this case, the relevant elasticity is also that of lost wage earnings (with respect to the replacement rate) rather than unemployment probability, accounting for the impact of heterogeneous wages on the budget balance constraint.

In practice UI systems often have both constant and proportional components, for example setting benefits and taxes proportional to wages up to some cap, w^c, so $b^i = b_\rho \min\{w^i, w^c\}$ and $\tau^i = \tau_\rho \min\{w^i, w^c\}$. The results of this section can also be applied to this mixed case, provided we replace the wage w^i by the *taxable* wage $\tilde{w}^i = \min\{w^i, w^c\}$ in all expressions. The resulting welfare expression $\frac{d\bar{W}_\rho}{db_\rho}$ can again be viewed as a money-metric for welfare, though the interpretation is more involved.

6.2 Coefficient of Relative Prudence

In our baseline analysis we assume that the utility function of agent i is well-approximated by a two-term Taylor expansion, i.e. that

$$u'_i(c_u) - u'_i(c_e) \approx u''_i(c_e)(c_u - c_e).$$

As noted in Chetty (2006) this approximation may problematic if third-order terms in the utility function are large. In particular Chetty discusses a calibration based Gruber (1997) with CRRA utility and risk aversion ranging from 1 to 5, and shows that the formula (3) for the single agent case sometimes underestimates the optimal level of UI benefits by more than 30%. To correct this issue Chetty suggests including third-order terms in the approximation to $u(\cdot)$, considering

$$u'(c_u) - u'(c_e) \approx u''(c_e)(c_u - c_e) + \frac{1}{2} u'''(c_e)(c_u - c_e)^2 \tag{14}$$

for the single-agent case, and shows that this approximation yields a new approximation to dW/db that depends on the coefficient of relative prudence $\rho = \frac{u'''(c_e)}{u''(c_e)} c_e$, which accounts for precautionary savings motives. In Chetty's calibration the inclusion of this term reduces the maximal error in estimating the optimal benefit level to less than 4%, though as Chetty notes there are certainly examples where this distortion would be larger.

Our heterogeneity-robust approach can easily accommodate such third order terms. Taking the exact expression for marginal welfare (8) and applying the approximation (14) for each individual we obtain

$$\frac{d\bar{W}}{db} \approx \frac{1-\bar{s}}{\bar{s}} E^u \left[\frac{u_i''(c_e^i)(c_u^i - c_e^i) + \frac{1}{2}u_i'''(c_e^i)(c_u^i - c_e^i)^2}{u_i'(c_e^i)} \right] + \frac{1-\bar{s}}{\bar{s}} - E^e \left[\frac{d\tau^i}{db} \right]$$

$$= \frac{1-\bar{s}}{\bar{s}} E^u \left[\gamma^i \Delta^i + \frac{1}{2}\gamma^i \rho^i (\Delta^i)^2 \right] + \frac{1-\bar{s}}{\bar{s}} - E^e \left[\frac{d\tau^i}{db} \right].$$

Note that the addition of a third-order term in our approximation to u_i has no effect on the tax term $E^e \left[\frac{d\tau^i}{db} \right]$. Hence, for actuarially-fair individual taxes $\frac{d\tau^i}{db} = \frac{1-s^i}{s^i} \left(1 + \frac{\epsilon_{1-s^i,b}}{s^i} \right)$ we have

$$\frac{d\bar{W}}{db} \approx \frac{1-\bar{s}}{\bar{s}} \left\{ E^u \left[\gamma^i \Delta^i + \frac{1}{2}\gamma^i \rho^i (\Delta^i)^2 \right] - E^u \left[\frac{\epsilon_{1-s^i,b}}{s^i} \right] \right\}$$

$$= \frac{1-\bar{s}}{\bar{s}} \left\{ \bar{\gamma}^u \bar{\Delta}^u + \frac{1}{2} E^u \left[\gamma^i \rho^i \right] E^u \left[(\Delta^i)^2 \right] + \text{cov}^u \left(\gamma^i, \Delta^i \right) \right.$$

$$\left. + \frac{1}{2}\text{cov}^u \left(\gamma^i \rho^i, (\Delta^i)^2 \right) - E^u \left[\frac{\epsilon_{1-s^i,b}}{s^i} \right] \right\}.$$

This differs from the approximation (9) which neglects third-order terms in including the term $\frac{1}{2} \left(E^u \left[\gamma^i \rho^i \right] E^u \left[(\Delta^i)^2 \right] + \text{cov}^u \left(\gamma^i \rho^i, (\Delta^i)^2 \right) \right)$. This term can be re-written as

$$\frac{1}{2} \left(\bar{\gamma}^u \bar{\rho}^u + \text{cov}^u \left(\gamma^i, \rho^i \right) \right) \left((\bar{\Delta}^u)^2 + \text{var}^u \left(\Delta^i \right) \right) + \frac{1}{2}\text{cov}^u \left(\gamma^i \rho^i, (\Delta^i)^2 \right)$$

which makes it clear that the third-order approximation to $\frac{d\bar{W}}{db}$ depends on four terms that did not appear in the second-order approximation (9): the mean coefficient of relative prudence $\bar{\rho}^u$, the covariance between risk aversion and relative prudence $\text{cov}^u \left(\gamma^i, \rho^i \right)$, the variance of consumption drops $\text{var}^u \left(\Delta^i \right)$, and the covariance of the product $\gamma^i \rho^i$ with the squared consumption drop, $\text{cov}^u \left(\gamma^i \rho^i, (\Delta^i)^2 \right)$, all in the population of unemployed agents. Many of these additional terms are fairly intuitive: $\bar{\rho}^u$ can be viewed as analogous to the parameter ρ in the single agent case while $\text{cov}^u \left(\gamma^i, \rho^i \right)$ captures any co-movement between risk aversion and relative prudence. Our use of a third-order approximation to the utility implies a second-order approximation to the marginal utility $u_i'(\cdot)$, making the variance of consumption drops $\text{var}^u \left(\Delta^i \right)$ potentially important for calculating expected marginal utility.

Analysis of the uniform-tax case with $\frac{d\tau^i}{db} = \frac{d\tau}{db} = \frac{1-\bar{s}}{\bar{s}} \left(1 + \frac{\epsilon_{1-\bar{s},b}}{\bar{s}} \right)$ proceeds along the same lines, yielding the same transfer terms as before and marginal welfare equal to

$$\frac{d\bar{W}}{db} \approx \frac{1-\bar{s}}{\bar{s}} \left\{ \bar{\gamma}^u \bar{\Delta}^u + \frac{1}{2} E^u \left[\gamma^i \rho^i \right] E^u \left[(\Delta^i)^2 \right] + \text{cov}^u \left(\gamma^i, \Delta^i \right) \right.$$

$$\left. + \frac{1}{2}\text{cov}^u \left(\gamma^i \rho^i, (\Delta^i)^2 \right) - \frac{\epsilon_{1-\bar{s},b}}{\bar{s}} \right\}$$

where the term $\frac{1}{2}\left(E^u\left[\gamma^i\rho^i\right]E^u\left[\left(\Delta^i\right)^2\right] + \operatorname{cov}^u\left(\gamma^i\rho^i, \left(\Delta^i\right)^2\right)\right)$ can be re-written as above.

6.3 Sufficient Statistics for Dynamic Models

The central result of Chetty (2006) is that a simple sufficient-statistic formula analogous to those in the previous section continues to hold under reasonable conditions in a rich dynamic model with a single agent. As discussed in section 3.4 above, Chetty notes that this result can be extended to the case with heterogeneous agents provided risk aversion is homogeneous, though his approach corresponds to a different welfare metric than the one considered in this paper. In this section we argue that provided Chetty's assumptions hold for each agent i a natural generalization of our approach for static models allows us to extend our analysis to cover dynamic models with arbitrarily heterogeneous agents. Essentially, we show that under welfare metrics analogous to those considered in section 3 we can use Chetty's results to obtain approximate marginal welfare expressions for each agent and then aggregate across agents as before.

We consider a potentially heterogeneous version of the model studied by Chetty (2006). This model is considerably more elaborate than the simple static model described in section 2: we briefly introduce the key terms for our analysis but refer the interested reader to Chetty (2006) for a full exposition. As in Chetty's model, assume that time is continuous and that all agents i live from $t = 0$ to $t = 1$. The state of the world at time t is indexed by a state variable ω_t that follows some arbitrary stochastic process and whose unconditional distribution at time t is $F_t(\omega_t)$. Agents choose behavior at time t conditional on ω_t including consumption $c^i(t, \omega_t)$ and a vector of M other behaviors $x^i(t, \omega_t)$, for example, search effort and private insurance purchases. We assume that utility is time-separable and that the flow utility at time t is $u_i\left(c^i\left(t, \omega_t\right), x^i\left(t, \omega_t\right)\right)$ where we will typically suppress the dependence on $x^i(t, \omega_t)$. Agent i's employment status at time t is tracked by $\theta^i_t(t, \omega_t)$, which is equal to one if the agent is employed and zero if the agent is unemployed. Assume that the agents are subject to a budget constraint as in Chetty (2006), earn income w^i and pay tax τ^i in the employed state, and receive lump sum benefit b in the unemployed state. Chetty's model also allows the agent's activities x^i to generate income $f^i\left(x^i\left(t, \omega_t\right)\right)$, and accommodates the imposition of N other constraints on agent i's choices.

To show that the welfare analysis discussed above can be extended to this case, let us denote by $V^i\left(b, \tau\right)$ agent i's maximized expected utility, i.e.

$$V^i\left(b, \tau^i\right) = \max_{c^i, x^i} \int \int u_i\left(c^i\left(t, \omega_t\right), x^i\left(t, \omega_t\right)\right) dF_t(\omega_t)dt$$

subject to the constraints on the agent's choice set (see page 1887 in Chetty (2006) for a formal description of this optimization problem, where our case differs only allowing all terms to depend on the identity of agent i). Define D^i to be agent i's expected fraction of lifetime unemployed,

$$D^i = \int \int \left(1 - \theta^i\left(t, \omega_t\right)\right) dF_t(\omega_t)dt.$$

Provided Assumptions 1-3 and 5 of Chetty (2006) hold for each agent i (which we'll assume to be the case), Lemma 1 in Chetty (2006) establishes that

$$\frac{dV^i}{db} = D \cdot Eu'_i(c_u) - \frac{d\tau^i}{db}(1 - D) \cdot Eu'_i(c_e)$$

where

$$Eu'_i(c_e) = \frac{\int \int \theta^i(t, \omega_t) u'_i(c^i(t, \omega_t), x^i(t, \omega_t)) dF_t(\omega_t) dt}{\int \int \theta^i(t, \omega_t) dF_t(\omega_t) dt}$$

$$Eu'_i(c_u) = \frac{\int \int (1 - \theta^i(t, \omega_t)) u'_i(c^i(t, \omega_t), x^i(t, \omega_t)) dF_t(\omega_t) dt}{\int \int (1 - \theta^i(t, \omega_t)) dF_t(\omega_t) dt}$$

are agent i's marginal utility in the employed and unemployed states, respectively, averaging over both time and states of the world.

To extend our welfare analysis for heterogeneous agents to this case, note that we face the same difficulties aggregating welfare across heterogeneous agents as discussed in section 3. To overcome this challenge we adapt our approach for the static case, normalizing agent i's utility function by $\frac{1}{Eu'_i(c_e)}$ or, equivalently, attaching welfare weight $\alpha^i = \frac{1}{Eu'_i(c_e)}$ to agent i. The expected normalized marginal welfare of agent i is then

$$\frac{dW^i}{db} = D^i \cdot \frac{Eu'_i(c_u)}{Eu'_i(c_e)} - \frac{d\tau^i}{db}(1 - D^i)$$

and marginal aggregate welfare (dividing through by the constant $1 - \bar{D}$) is

$$\frac{d\bar{W}}{db} = \frac{1}{(1 - \bar{D})} \int \frac{dW^i}{db} dF_i = \frac{\bar{D}}{(1 - \bar{D})} E^u \left[\frac{Eu'_i(c_u) - Eu'_i(c_e)}{Eu'_i(c_e)} \right]$$
$$+ \frac{\bar{D}}{(1 - \bar{D})} - E^e \left[\frac{d\tau^i}{db} \right] \tag{15}$$

where $\bar{D} = \int D^i dF_i$ while

$$E^u[X^i] = \frac{1}{\bar{D}} \int D^i X^i dF_i$$

$$E^e[X^i] = \left(\frac{1}{1 - \bar{D}} \right) \int (1 - D^i) X^i dF_i$$

are averages over agents weighted by expected fraction of life spent unemployed and employed, respectively.

To give a money metric interpretation for (15) note that by equation (14) in Chetty (2006) the marginal utility of agent i with respect to a permanent increase in consumption in the employed state is $(1 - D^i) Eu'_i(c_e)$. Note, further, that the cost of providing a permanent \$1 consumption increase to agent i in the employed state, holding the expected fraction of lifetime unemployed D^i fixed, is $1 - D^i$. Hence, the aggregate welfare (15) has the interpretation as the total expected

cost (holding job-finding rates fixed) of the bundle of agent-specific consumption changes in the employed state that increase the expected utility of each agent by the same amount as the proposed UI reform, directly generalizing the money metric for the static model to this much richer dynamic case.[13]

Chetty's Lemma 2 shows that for each individual i, provided third and higher order terms in the utility function are small we can approximate average marginal utility in each state by the marginal utility of the average consumption in that state, i.e. $Eu'_i(c_e) \approx u'_i(\bar{c}^i_e)$ and $Eu'_i(c_u) \approx u'_i(\bar{c}^i_u)$ for

$$\bar{c}^i_e = \frac{\int \int \theta^i(t, \omega_t) c^i(t, \omega_t) dF_t(\omega_t) dt}{\int \int \theta^i(t, \omega_t) dF_t(\omega_t) dt}$$

$$\bar{c}^i_u = \frac{\int \int (1 - \theta^i(t, \omega_t)) c^i(t, \omega_t) dF_t(\omega_t) dt}{\int \int (1 - \theta^i(t, \omega_t)) dF_t(\omega_t) dt}$$

and a third-order expansion of the utility function $u_i(\cdot)$ then yields that

$$\frac{Eu'_i(c_u) - Eu'_i(c_e)}{Eu'_i(c_e)} \approx \gamma^i \Delta^i + \frac{1}{2} \gamma^i \rho^i (\Delta^i)^2$$

for $\Delta^i = \frac{\bar{c}^i_e - \bar{c}^i_u}{\bar{c}^i_e}$. Note that in computing Δ^i we use agent i's consumption in each employment state averaged over time and states of the world, so Δ^i can be interpreted as a particular measure of the drop in consumption for agent i due to unemployment. Substituting this approximation into (15) for actuarially fair tax change $\frac{d\tau^i}{db} = \frac{D^i}{1-D^i}\left(1 + \frac{\epsilon_{D^i,b}}{1-D^i}\right)$ yields marginal aggregate welfare

$$\frac{d\bar{W}}{db} \approx \frac{\bar{D}}{(1-\bar{D})} \left(\bar{\gamma}^u \bar{\Delta}^u + \frac{1}{2} E^u[\gamma^i \rho^i] E^u[(\Delta^i)^2] + \text{cov}^u(\gamma^i, \Delta^i) \right.$$
$$\left. + \frac{1}{2}\text{cov}^u(\gamma^i \rho^i, (\Delta^i)^2) - E^u\left[\frac{\epsilon_{D^i,b}}{1-D^i}\right] \right) \tag{16}$$

for $\epsilon_{D^i,b}$ the elasticity of agent i's expected unemployment duration with respect to the benefit level. The expected money-metric value of the transfer to agent i under a uniform tax change $\frac{d\tau}{db} = \frac{\bar{D}}{1-\bar{D}}\left(1 + \frac{\epsilon_{\bar{D},b}}{1-\bar{D}}\right)$ is $\frac{D^i}{1-D^i}\left(1 + \frac{\epsilon_{D^i,b}}{1-D^i}\right) - \frac{\bar{D}}{1-\bar{D}}\left(1 + \frac{\epsilon_{\bar{D},b}}{1-\bar{D}}\right)$, and the marginal aggregate welfare for a UI reform with uniform taxes $\frac{d\tau^i}{db} = \frac{d\tau}{db}$ is

$$\frac{d\bar{W}}{db} \approx \frac{\bar{D}}{(1-\bar{D})} \left(\bar{\gamma}^u \bar{\Delta}^u + \frac{1}{2} E^u[\gamma^i \rho^i] E^u[(\Delta^i)^2] \right.$$
$$\left. + \text{cov}^u(\gamma^i, \Delta^i) + \frac{1}{2}\text{cov}^u(\gamma^i \rho^i, (\Delta^i)^2) - \frac{\epsilon_{\bar{D},b}}{1-\bar{D}} \right) \tag{17}$$

where $\epsilon_{\bar{D},b}$ is the elasticity of the average unemployment duration with respect to the UI benefit level. Hence, we see that the third-order approximations obtained in section 6.2 above can be

---

[13]Note that in our baseline static model, a \$1 increase in the wage is equivalent to a \$1 consumption increase while employed.

generalized to a heterogeneous version of the dynamic model considered by Chetty (2006).

Note that we have followed the primary exposition in Chetty (2006) in using a second-order approximation to $u_i(\cdot)$ to approximate $Eu_i'(c_e) \approx u_i'(\bar{c}_e^i)$ but then using a third order expansion of $u_i(\cdot)$ evaluated at the mean consumption level. As in Chetty (2006) we could instead use a third order approximation to $u_i(\cdot)$ at both steps at the cost of introducing additional terms reflecting the relative variability of consumption in the employed and unemployed states. Alternatively, using a second-order approximation to $u_i(\cdot)$ at both steps yields (16) and (17) with $\rho^i \equiv 0$, and can be viewed as the generalization of our baseline approximation developed in section 4 to the dynamic case.

To understand (16) and (17) it is important to think carefully about precisely what averages are being taken. The expectation

$$E^u\left[X^i\right] = \frac{\int \int \int \left(1 - \theta^i\left(t, \omega_t\right)\right) X^i dF_t(\omega_t) dt dF_i}{\int \int \int \left(1 - \theta^i\left(t, \omega_t\right)\right) dF_t(\omega_t) dt dF_i}$$

takes the mean of X^i where agents are weighted by expected unemployment duration and corresponds to the mean of X^i that we would calculate by averaging over a random sample of agent-year observations in the unemployed state. Likewise, $\text{cov}^u\left(X^i, Y^i\right)$ is the covariance between X^i and Y^i that we would obtain by pooling repeated cross-sections of the unemployed population.

One aspect of the formulas (16) and (17) that could make them challenging to implement is the presence of the term Δ^i, which depends on the average consumption level for agent i across both time and ω_t conditional on each employment state. If we're willing to assume that agents face little lifetime consumption risk in each state, so that for each individual $\bar{c}_e^i$ and $\bar{c}_u^i$ are well-proxied by the average realized consumption in each state, then we can use any sample which contains the realized distribution of lifetime consumption paths and employment states to calculate $\bar{\Delta}^u$, $\text{var}^u\left(\Delta^i\right)$, and so on. Without such an assumption, however, the problem is more challenging and we will in general need additional assumptions (for example on the distribution of consumption risks faced by individuals) to recover the distribution of Δ^i from data on realized consumption profiles.

7   Conclusion

The Baily-Chetty formula is robust to a degree of heterogeneity but requires the assumption that agents share a common coefficient of relative risk aversion. In this paper, we extend the Baily-Chetty framework to allow for arbitrary heterogeneity across agents. We find that heterogeneity affects welfare analysis through the covariance effect: welfare gains depend on the covariance between risk aversion and consumption drops in the cross-section of the unemployed. This reflects the fact that unemployment insurance is more valuable if more risk averse agents are subject to larger risks. Calibration results suggest that the covariance effect may be large: for plausible population distributions of risk preferences, we find that accounting for the covariance effect can change the approximate consumption smoothing benefit of UI by more than 50%.

Our results may have important implications for existing applications of the Baily-Chetty approach. For example, a recent literature extends the approach to investigate how UI benefits should vary over the business cycle (Landais et al. 2010; Kroft and Notowidigdo 2011; Schmieder et al. 2012). These papers emphasize that how optimal benefits vary over the business cycle depends on the cyclicality of the duration elasticity and consumption drops as a function of UI benefits. Our analysis makes clear that optimal benefits will also depend on the cyclicality of the covariance effect, which could arise if consumption smoothing mechanisms like consumer credit and spousal labor supply become less available during recessions.

Our results demonstrate that the value of social insurance depends on the covariance of risk exposure and risk aversion in the population: for a given distribution of consumption drops, the lower the covariance of risks faced by workers with individual risk aversion, the lower the value of additional social insurance. The sign and magnitude of the covariance effect are *a priori* ambiguous and will depend on the joint distribution of risk preferences, ex-ante risk, and ability to smooth consumption in a given context. Estimating this covariance is an important area for future research.

References

Ahmed, Naeem, Matthew Brzozowski, and Thomas F. Crossley, "Measurement Errors in Recall Food Consumption Data," April 2010. Unpublished manuscript.

Baily, Martin Neil, "Some Aspects of Optimal Unemployment Insurance," *Journal of Public Economics*, December 1978, *10*, 379–402.

Barseghyan, Levon, Jeffrey Prince, and Joshua C. Teitelbaum, "Are Risk Preferences Stable Across Contexts? Evidence from Insurance Data," *American Economic Review*, April 2011, *101* (2), 591–631.

Barsky, Roberet B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics*, May 1997, *112* (2), 537–579.

Chetty, Raj, "Consumption Commitments, Unemployment Durations, and Local Risk Aversion," December 2003. Unpublished manuscript.

——— , "A general formula for the optimal level of social insurance," *Journal of Public Economics*, 2006, *90*, 1879–1901.

——— , "Moral Hazard versus Liquidity and Optimal Unemployment Insurance," *Journal of Political Economy*, 2008, *116* (2), 173–234.

——— **and Adam Looney**, "Consumption Smoothing and the Welfare Consequences of Social Insurance in Developing Economies," *Journal of Public Economics*, December 2006, *90* (12), 2351–2356.

___ **and Adam Szeidl**, "Consumption Commitments and Risk Preferences," *Quarterly Journal of Economics*, May 2007, *122* (2), 831–877.

___ **and Emmanuel Saez**, "Optimal Taxation and Social Insurance with Endogenous Private Insurance," *American Economic Journal: Economic Policy*, 2010, *2* (2), 85–114.

Cohen, Alma and Liran Einav, "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, June 2007, *97* (3), 745–788.

Diamond, Peter A., "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty," *American Economic Review*, 1967, *57* (4), 759–776.

Einav, Liran, Amy Finkelstein, Iuliana Pascu, and Mark Cullen, "How general are risk preferences? Choices under uncertainty in different domains," *American Economic Review*, October 2012, *102* (6), 2606–2638.

Gross, Tal and Matt Notowidigdo, "Health Insurance and the Consumer Bankruptcy Decision: Evidence from Expansions of Medicaid," *Journal of Public Economics*, 2011, *95* (7-8), 767–788.

Gruber, Jonathan, "The Consumption Smoothing Benefits of Unemployment Insurance," *American Economic Review*, March 1997, *87* (1), 192–205.

Holmstrom, Bengt and Roger B. Myerson, "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, November 1983, *51* (6), 1799–1891.

Kroft, Kory and Matthew J. Notowidigdo, "Should Unemployment Insurance Vary With the Unemployment Rate? Theory and Evidence," December 2011. Working paper.

Landais, Camille, Pascal Michaillat, and Emmanuel Saez, "Optimal Unemployment Insurance over the Business Cycle," November 2010. Working paper.

Mazzocco, Maurizio and Shiv Saini, "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," *American Economic Review*, 2012, *102* (1), 428–468.

Meyer, Bruce, "Unemployment Insurance and Unemployment Spells," *Econometrica*, July 1990, *58*, 757–782.

Schmieder, Johannes, Till von Wachter, and Stefan Bender, "The Effects of Extended Unemployment Insurance over the Business Cycle: Evidence from Regression Discontinuity Estimates over 20 Years," *Quarterly Journal of Economics*, 2012, *127* (2), 701–752.

Schulhofer-Wohl, Sam, "Heterogeneity and Tests of Risk Sharing," *Journal of Political Economy*, October 2011, *119* (5), 925–958.

Wilson, Robert, "The Theory of Syndicates," *Econometrica*, 1968, *36* (1), 119–132.

Table 1: Covariance Term Calibration

Distribution:	Uniform			Lognormal		
Range:	1-5	1-10	1-20	1-5	1-10	1-20
$\mathrm{Cov}(\gamma, \Delta)$	-0.034	-0.107	-0.284	-0.036	-0.109	-0.272
$\bar{\gamma} \times \bar{\Delta}$	0.221	0.406	0.778	0.186	0.299	0.503
$\frac{|Cov(\gamma,\Delta)|}{\bar{\gamma}\times\bar{\Delta}}$	0.154	0.263	0.365	0.195	0.363	0.540
$\bar{\gamma}$	3.000	5.500	10.501	2.521	4.038	6.772
σ_γ	1.155	2.598	5.484	1.311	3.207	7.707
$\frac{\sigma_\gamma}{\bar{\gamma}}$	0.385	0.472	0.522	0.520	0.794	1.138
$\bar{p}$	0.067	0.075	0.090	0.064	0.066	0.068
σ_p	0.013	0.018	0.022	0.010	0.010	0.006

Notes: For lognormal distribution, the range is the 5th-95th percentile range. Results are based on 10 million i.i.d. draws for each case. The mean proportional consumption drop, $\bar{\Delta}$, is stable across cases at 7.4%.