THE CYCLICALITY OF THE OPPORTUNITY COST OF EMPLOYMENT*

Gabriel Chodorow-Reich
Harvard University and NBER

Loukas Karabarbounis
University of Chicago and NBER

June 2014

Abstract

The flow opportunity cost of moving from unemployment to employment consists of foregone public benefits and the foregone value of non-working time in units of consumption. We construct a time series of the opportunity cost of employment using detailed microdata and administrative or national accounts data to estimate benefits, consumption by employment status, and preference parameters. Our estimated opportunity cost is procyclical and volatile over the business cycle. The estimated cyclicality implies far less unemployment volatility in search and matching models of the labor market than that observed in the data. This result holds irrespective of the level of the opportunity cost or whether wages are set by Nash bargaining or by an alternating-offer bargaining process. We conclude that appealing to aspects of labor supply does not help search and matching models explain aggregate employment fluctuations.

JEL-Codes: E24, E32, J64.

Keywords: Opportunity Cost of Employment, Unemployment Fluctuations.

---

*Chodorow-Reich: Harvard University Department of Economics, Littauer Center, Cambridge, MA 02138 (e-mail: chodorowreich@fas.harvard.edu); Karabarbounis: University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637 (email: loukas.karabarbounis@chicagobooth.edu). We are especially grateful to Bob Hall for many insightful discussions and for his generous comments at various stages of this project. This paper also benefited from conversations with Steve Davis, Peter Ganong, Erik Hurst, Greg Kaplan, Larry Katz, Pat Kehoe, Guido Lorenzoni, Giuseppe Moscarini, Casey Mulligan, Richard Rogerson, Rob Shimer, and Gianluca Violante. We also thank seminar participants at Brown, Chicago Booth, the Empirical Macroeconomics Workshop in New Orleans, the Federal Reserve Bank of Chicago, the Federal Reserve Bank of New York, the Federal Reserve Bank of Minneapolis, Harvard, LSE, the NBER EFG meeting, Northwestern, NYU, Princeton, the Tsinghua Workshop in Macroeconomics, UT Austin, Wharton, and Yale. Much of this paper was written while Gabriel Chodorow-Reich was visiting the Julis-Rabinowitz Center at Princeton University. Loukas Karabarbounis thanks Chicago Booth for summer financial support. The Appendix to this paper is available at the authors' webpages.

1 Introduction

Understanding the causes and the consequences of labor market fluctuations ranks among the most important and difficult issues in economics. In recent years, the theory of unemployment with search and matching frictions described in Mortensen and Pissarides (1994) (hereafter MP model) has emerged as the workhorse building block of the labor market in macroeconomic models. As emphasized in influential work by Shimer (2005) and Hall (2005), the standard MP model with wages set according to Nash bargaining fails to account quantitatively for the observed volatility of unemployment. This has led to a significant amount of research effort devoted to reconciling the search and matching model with the data.

The flow value of the opportunity cost of employment (which we denote by z) plays a crucial role in the MP model and in some of the leading proposed solutions to the unemployment volatility puzzle that we revisit below. The importance of this variable has generated debate about its level, but the literature has almost uniformly adopted the assumption that z is constant over the business cycle. Our contribution starts from the observation that not only the level, but also the cyclicality of z matters for unemployment fluctuations. Movements in z correspond loosely to shifts in labor supply, making it unsurprising that they would affect unemployment. While this insight goes back as far as Pissarides (1985), to date the cyclical properties of the opportunity cost in the data remain unknown.

We find that the opportunity cost z is procyclical and quite volatile over the business cycle. Our estimated cyclicality poses a strong challenge to models that rely on a fixed z to solve the unemployment volatility puzzle. This is because the procyclical opportunity cost undoes the wage rigidity generated by these models.

We construct a time series of z using detailed microdata, administrative data, and the structure of the search and matching model with concave preferences and an explicit value of non-working time.[1] We call this model the MP/RBC model, as it combines elements from

[1] Our approach complements recent research that uses surveys to ask respondents directly about their reservation wage (Hall and Mueller, 2013; Krueger and Mueller, 2013). Relative to survey estimates, our measurement allows to construct a much longer time series for z, which is crucial for studying cyclical patterns.

both the MP model and the real business cycle (RBC) model. In its basic form, the MP/RBC model has been studied extensively in the literature (Merz, 1995; Andolfatto, 1996; Shimer, 2010). We use an extended version of the model to derive an expression for the opportunity cost which we can take to the data. We write z as the sum of two components, $z = b + \xi$.

The b component is the value of public benefits that an unemployed person forgoes upon employment. Our approach to measuring b departs from the literature in three significant ways. First, we focus on effective rather than statutory benefit rates. Second, we include both unemployment insurance (UI) benefits, which are directly related to unemployment status, and non-UI benefits such as supplemental nutritional assistance (SNAP), welfare assistance (AFDC/TANF), and health care (Medicaid). The latter belong in the opportunity cost to the extent that receipt of these benefits changes with unemployment status. Third, we take into account UI benefits expiration, and we model and measure the utility costs (e.g. job search costs and other filing and time costs) associated with taking up UI benefits. These utility costs allow the model to match the fact that roughly one-third of eligible unemployed do not actually take up benefits.

We combine micro survey data with program administrative data to measure b. Using household and individual-level data from the Current Population Survey (CPS) and the Survey of Income and Program Participation (SIPP), we estimate the shares of UI, SNAP, AFDC/TANF, and Medicaid spending that belong in b. To circumvent the noise and the undercounting of benefits in the microdata, we then apply these shares to benefit totals from administrative data sources. Our estimated b is countercyclical, rising around every recession since 1961. However, because our estimates reflect effective rather than statutory rates, and because they account for costs associated with take-up and for expiration, we find a level of b much smaller than what the literature has traditionally calibrated. The sample average of b is 4.1 percent of the sample average of the marginal product of employment (which we denote by p^e).

The ξ component of z results from consumption and work differences between the employed and unemployed. This component resembles the marginal rate of substitution between non-

working time and consumption in the RBC model, with the difference being that the extra value of non-working time is calculated along the extensive margin. In the RBC model, an intraperiod first-order condition equates the marginal rate of substitution between non-working time and consumption to the marginal product of labor. While the search and matching literature has appealed to this equality to motivate setting the *level* of z close to that of the marginal product p^e, the same logic suggests that the ξ component of z would move *cyclically* with p^e just as in the RBC model.

To construct the time series of ξ, we require time series of hours per worker, the consumption of the employed, the consumption of the unemployed, and preference parameters. We discipline the preference parameters by requiring that the consumption decline upon unemployment in the steady state of the model match its empirical analog from microdata. Both our estimates of the drop in consumption from the Consumption Expenditure Survey (CE) and the Panel Study of Income Dynamics (PSID) and the resulting preference parameters are broadly consistent with those found in the literature, including Aguiar and Hurst (2005), Hall and Milgrom (2008), and Hall (2009). We generate time series of consumption of the employed and unemployed using the model's structure jointly with our estimates of relative consumption by labor force status, population shares by labor force status, and NIPA consumption per capita. We find that ξ is highly procyclical. Intuitively, the household values most the contribution of the employed (through higher wage income) relative to that of the unemployed (through higher non-working time) when consumption is low and non-working time is high.

Combining the opportunity cost associated with benefits b with the opportunity cost associated with consumption and work differences ξ, we show that $z = b + \xi$ is procyclical and quite volatile over the business cycle. The procyclicality of z occurs despite the countercyclical and volatile b. Because the level of b is small, the ξ component of the opportunity cost accounts for the majority of the fluctuations in z.

The cyclicality of z dampens unemployment fluctuations in models where z affects the wage determination process. The logic of this result is quite general, and does not depend on the

set of primitive shocks driving the business cycle. Relative to the fixed z case, a procyclical z increases the worker's surplus from accepting a job *at a given wage* during a recession, which puts downward pressure on equilibrium wages and ameliorates the increase in unemployment. The extent to which actual wages vary cyclically remains an open and important question. We interpret our result as showing that any such wage rigidity cannot be justified by mechanisms that appeal to aspects of the opportunity cost.

We illustrate the consequences of a procyclical z in the context of two leading proposed solutions to the unemployment volatility puzzle in the MP class of models, both of which take the marginal product of employment p^e as the driving force. As a first step, we estimate the comovement of the cyclical components of z and p^e, and find an elasticity of z with respect to p^e of close to one.

Hagedorn and Manovskii (2008) show that increasing the level of z close to that of p^e and making z constant over the business cycle allows the MP model to generate realistic unemployment fluctuations. Intuitively, a high level of z makes the total surplus from an employment relationship small on average. Then even modest increases in p^e can generate large percent increases in the surplus, incentivizing firms to significantly increase their job creation.[2] However, if changes in p^e are accompanied by equal percent changes in z, the surplus from a new hire remains relatively stable over the business cycle. As a result, the fluctuations in unemployment generated by the model are essentially neutral with respect to the level of z.

Hall and Milgrom (2008) generate volatile unemployment fluctuations by replacing the assumption of Nash bargaining over match surplus with an alternating-offer wage setting mechanism. With Nash bargaining, the threat point of an unemployed depends on the wage other jobs would offer in case of bargaining termination. In the alternating-offer bargaining game, the threat point depends instead mostly on the worker's flow value z if bargaining continues.

[2] A number of papers have followed this reasoning to set a relatively high level of z. In Hagedorn and Manovskii (2008), $z = 0.955$ and $p^e = 1$. Examples of papers before Hagedorn and Manovskii (2008) include Mortensen and Pissarides (1999), Mortensen and Pissarides (2001), Hall (2005), and Shimer (2005), which set z at 0.42, 0.51, 0.40, and 0.40. Examples of papers after Hagedorn and Manovskii (2008) include Mortensen and Nagypal (2007), Costain and Reiter (2008), Hall and Milgrom (2008), and Bils, Chang, and Kim (2012), which set z at 0.73, 0.745, 0.71, and 0.82. See Hornstein, Krusell, and Violante (2005) for a useful summary of this literature.

With constant z, wages respond weakly to increases in p^e, which incentivizes firms to significantly increase their job creation. Allowing instead z to comove with p^e as in the data, the unemployed's threat point becomes again sensitive to aggregate conditions. This increases the flexibility of wages and reduces the volatility of unemployment.

Finally, we extend our results to allow for heterogeneity across workers. First, we construct separate z's for workers with different educational attainment. While this exercise reveals interesting variation in the level and composition of z across workers, each of the skill-specific z's is roughly as cyclical as the aggregate z. Second, we relax the assumption of full risk sharing between employed and unemployed in the MP/RBC framework and examine instead a setting with incomplete markets and uninsured idiosyncratic unemployment shocks. Even with incomplete markets, in general the presence of aggregate shocks generates cyclical fluctuations in the (different) marginal utilities of workers, leading to a procyclical z.

Section 2 presents the MP/RBC model and derives the opportunity cost z. In Section 3, we use microdata, administrative data, and labor market data to estimate key parts of b and derive empirical moments necessary for estimating ξ. Section 4 discusses the remainder of the calibration. Section 5 reports the cyclicality of z. Section 6 presents implications for unemployment under Nash bargaining and under alternating-offer bargaining. We construct z for different educational categories in Section 7 and discuss the cyclicality of z under alternative risk sharing arrangements in Section 8. Section 9 concludes.

2 Baseline Model

We develop our measure of the opportunity cost of employment within the context of the labor market search and matching framework of Mortensen and Pissarides (1994) and Pissarides (2000) as embedded in a real business cycle model by Merz (1995) and Andolfatto (1996). Following this literature, we start our analysis by assuming that wages are set according to the generalized Nash bargaining solution. We discuss the alternating-offer wage setting mechanism used by Hall and Milgrom (2008) in Section 6.2.

Time is discrete and the horizon is infinite, $t = 0, 1, 2, ...$ We denote the vector of exogenous shocks by $\mathbf{Z}_t$. Consumption is the numeraire good. There is a representative firm producing output with capital and labor. There is a representative household that owns the firm and rents its capital stock K_t in a perfect capital market at a rate R_t. The household consists of a continuum of ex-ante identical workers of measure one. At the beginning of each period t, there are e_t employed who produce output and $u_t = 1 - e_t$ unemployed who search for jobs.

The labor market is subject to search and matching frictions. The firm posts vacancies v_t to increase employment in the next period. Each vacancy costs κ_t in terms of the numeraire good. Trade in the labor market is facilitated by a constant returns to scale matching technology that converts searching by the unemployed and vacancies by the firm into new matches, $m_t = m_t(v_t, u_t)$. We denote market tightness by $\theta_t = v_t/u_t$. An unemployed worker matches with a firm with probability $f_t(\theta_t) = m_t/u_t$ and the firm fills a vacancy with probability $q_t(\theta_t) = m_t/v_t = f_t(\theta_t)/\theta_t$. At the end of each period fraction s_t of the employed exogenously separate and become unemployed. Hence, employment evolves according to $e_{t+1} = (1 - s_t)e_t + m_t$.

2.1 Household

The representative household maximizes the expected discounted utility flows of its members by choosing consumption for the employed and the unemployed, C_t^e and C_t^u, purchases of investment goods X_t, and the share ζ_t of eligible unemployed to take up UI benefits. The household takes as given the path of prices and the outcome of the bargaining game described below. There is perfect risk sharing among the members of the household, so the household allocates consumption between employed and unemployed to equalize their marginal utilities. The assumption of perfect risk sharing simplifies the analysis, facilitates comparison to existing literature, and allows us to estimate the opportunity cost in the data in a transparent way. In Section 8 we show that relaxing this assumption does not qualitatively change our results.

The maximization problem is:

$$W^h\left(e_0, \omega_0, K_0, \mathbf{Z}_0\right) = \max E_0 \sum_{t=0}^{\infty} \beta^t \left[e_t U^e(C_t^e, N_t) + (1 - e_t)U^u(C_t^u, 0) - (1 - e_t)\omega_t \psi(\zeta_t)\right], \quad (1)$$

where $U^e(C_t^e, N_t)$ is the flow utility of the employed, $U^u(C_t^u, 0)$ is the flow utility of the unemployed excluding costs associated with taking up benefits, ω_t is the share of eligible unemployed for UI benefits, ψ_t denotes the household's costs per eligible unemployed from taking up UI benefits, and N_t is hours per employed worker.

The new element in the household's objective function is the utility costs of UI take-up. In a seminal study, Blank and Card (1991) found that roughly one-third of unemployed workers eligible for UI do not claim the benefit. To explain this fact, we introduce utility costs associated with claiming UI.[3] These costs capture foregone time and effort associated with providing information to the UI agency, any stigma from claiming benefits, and time spent searching for a job. We assume that the cost per eligible unemployed depends on the fraction of those eligible that take up UI benefits, $\psi_t = \psi(\zeta_t)$. Appendix B microfounds this assumption in an environment with heterogeneous take-up costs across potential recipients and an optimal take up decision by each eligible member.

The budget constraint of the household is given by:

$$e_t C_t^e + (1 - e_t) C_t^u + C_t^o + X_t + T_t = w_t e_t N_t + (1 - e_t) B_t + R_t K_t + \Pi_t, \tag{2}$$

where T_t are lump sum taxes, C_t^o is total expenditure by persons out of the labor force and the government, w_t is the wage per hour worked, B_t is benefits received per unemployed, and Π_t is dividends from ownership of the firm. Capital K_t accumulates as $K_{t+1} = (1 - \delta)K_t + X_t$, where δ denotes the depreciation rate.

Benefits B_t received from the government include UI benefits as well as other transfers such as supplemental nutritional assistance, welfare assistance, and health care. B_t includes only the part of the benefit that a worker loses upon moving from unemployment to employment.[4] We

[3] Blank and Card (1991) also provide state-level evidence that take-up responds to benefit levels, a finding confirmed by Anderson and Meyer (1997) using administrative microdata and by our own findings in Section 4 using aggregate time series data. The fact that some of those eligible forgo their UI entitlement indicates either an informational friction or a cost associated with take-up. The comovement of take-up with benefit levels suggests that informational frictions cannot fully explain the low take-up rate (unless these frictions are correlated with benefits).

[4] Benefits are financed by lump-sum taxes, $T_t = (1 - e_t)B_t$. The assumption that B_t includes only the part of the benefit that a worker loses upon employment is without loss of generality. Benefits not depending on employment status do not affect the value of moving a worker from unemployment to employment. If these benefits are also financed with lump sum taxes, then the budget constraint of the household does not change.

split B_t into non-UI benefits per unemployed $B_{n,t}$, for which we do not model take up costs, and UI benefits per unemployed $B_{u,t}$, so $B_t = B_{n,t} + B_{u,t}$.[5] Benefits per unemployed from UI are the product of the fraction of unemployed who are eligible for benefits ω_t, the fraction of eligible unemployed who receive benefits ζ_t, and benefits per recipient unemployed $\tilde{B}_t$, so $B_{u,t} = \omega_t \zeta_t \tilde{B}_t = \phi_t \tilde{B}_t$, where ϕ_t denotes the fraction of unemployed receiving UI.

The fraction of eligible unemployed ω_t is a state variable that depends on past eligibility, expiration policies, and the composition of the newly unemployed. In the U.S., UI eligibility depends on sufficient earnings during previous employment (monetary eligibility), the reason for employment separation (non-monetary eligibility), and the number of weeks of UI already claimed (expiration eligibility). We model expiration eligibility with a simple process under which eligible unemployed who do not find a job in period t maintain their eligibility in period $t+1$ with probability ω^u_{t+1}. We combine monetary and non-monetary eligibility into a single term ω^e_{t+1} which gives the probability that a newly unemployed worker in period t is eligible for UI in the next period. The stock of eligible unemployed in period $t+1$ is $u^E_{t+1} = \omega^u_{t+1}(1 - f_t)u^E_t + \omega^e_{t+1} s_t e_t$. Therefore, the fraction of eligible unemployed $\omega_{t+1} = u^E_{t+1}/u_{t+1}$ follows the law of motion:

$$\omega_{t+1} = \left(\omega^u_{t+1}(1 - f_t)\frac{u_t}{u_{t+1}}\right)\omega_t + \omega^e_{t+1} s_t \frac{e_t}{u_{t+1}}. \tag{3}$$

Denoting by λ_t the multiplier on the budget constraint, the first-order conditions are:

$$\lambda_t = \frac{\partial U^e_t}{\partial C^e_t} = \frac{\partial U^u_t}{\partial C^u_t}, \tag{4}$$

$$\lambda_t = E_t \beta \lambda_{t+1} \left(R_{t+1} + 1 - \delta\right), \tag{5}$$

$$\psi'(\zeta_t) = \lambda_t \tilde{B}_t. \tag{6}$$

Equation (4) says that the household allocates consumption to different members to equate their marginal utilities. Equation (5) is the Euler equation for capital. Equation (6) is the first-order condition for the optimal take-up rate ζ_t. The household directs eligible unemployed to claim benefits up to the point where the marginal cost $\psi'(\zeta_t)$ equals the utility value of

[5] While non-UI programs have take-up rates below unity, we do not include take-up costs for those programs because the decision and timing of take-up does not generally coincide with the timing of an unemployment spell.

benefits $\lambda_t \tilde{B}_t$. In Appendix B, we show that with heterogeneous utility costs across potential recipients, the marginal cost $\psi'(\zeta_t)$ also equals the utility cost of the marginal recipient.

A key object for unemployment fluctuations is the household's marginal value of an additional employed worker, $J_t^h = \partial W^h(e_t, \omega_t, K_t, \mathbf{Z}_t)/\partial e_t$. We express the marginal value in consumption units by dividing it by the marginal utility of consumption λ_t. Appendix C shows that this value is given by:

$$\frac{J_t^h}{\lambda_t} = w_t N_t - \left[b_t + (C_t^e - C_t^u) - \frac{U_t^e - U_t^u}{\lambda_t} \right] + (1 - s_t - f_t) E_t \left(\frac{\beta \lambda_{t+1}}{\lambda_t} \right) \frac{J_{t+1}^h}{\lambda_{t+1}}. \tag{7}$$

The marginal value of an employed worker in terms of consumption consists of a flow value plus the expected discounted marginal value in the next period. The flow value consists of a flow gain from increased wage income, $w_t N_t$, and a flow loss associated with moving a worker from unemployment to employment.

Following Hall and Milgrom (2008), we define the (flow) opportunity cost of employment as the bracketed term in equation (7):

$$z_t = b_t + (C_t^e - C_t^u) - \frac{U_t^e - U_t^u}{\lambda_t} = b_t + \xi_t, \tag{8}$$

where b_t denotes the component of the opportunity cost related to benefits and ξ_t denotes the component of the opportunity cost related to consumption and work differences between the employed and the unemployed. We now discuss each of these components in further detail.

2.1.1 Opportunity Cost of Employment: Benefits

Appendix C derives an expression for the opportunity cost of employment related to benefits:

$$b_t = B_{n,t} + B_{u,t} \left(1 - \frac{1}{\alpha} \right) \left[1 - E_t \left(\frac{\beta \lambda_{t+1} \tilde{B}_{t+1} \zeta_{t+1}}{\lambda_t \tilde{B}_t \zeta_t} \right) \left(\frac{\omega_{t+1}^e}{\omega_t} - \omega_{t+1}^u \right) \left(\frac{s_t (1 - f_t)}{1 - e_{t+1}} \right) \Gamma_{t+1} \right], \tag{9}$$

where $\Gamma_{t+1} = \left(1 - \frac{\beta \lambda_{t+1}}{\lambda_t} \omega_{t+1}^u (1 - f_t) \frac{u_t}{u_{t+1}} \right)^{-1} > 1$ and $\alpha = \psi'(\zeta_t)\zeta_t / \psi(\zeta_t) > 1$.

The first term in equation (9) for b_t is simply non-UI benefits per unemployed, $B_{n,t}$. The second term consists of UI benefits per unemployed $B_{u,t}$, multiplied by an adjustment for the disutility of take-up and an adjustment for benefits expiration. This term is smaller than UI benefits per unemployed $B_{u,t}$.

The term $1 - 1/\alpha$ captures the surplus from receiving UI benefits. The average surplus per recipient equals the benefit per recipient less the average utility cost per recipient, $\lambda_t \tilde{B}_t - \psi(\zeta_t)/\zeta_t$. Using the first order condition (6), the average surplus is equivalently given by the difference between the marginal and the average cost, $\psi'(\zeta_t) - \psi(\zeta_t)/\zeta_t$. This difference depends on the elasticity of the cost function $\alpha = \psi'(\zeta_t)\zeta_t/\psi(\zeta_t) > 1$. If this elasticity is close to one, average cost per recipient is roughly constant, and there is a small surplus from receiving benefits as the household always incurs a cost per recipient that approximately equals the benefit per recipient. If this elasticity is much greater than one, average cost per recipient is below the marginal cost, and the household enjoys a larger surplus from receiving benefits.

The term in brackets captures the adjustment for benefits expiration. This term is lower than one when the probability that newly separated workers receive benefits, ω^e_{t+1}, exceeds the probability that previously eligible workers continue to receive benefits, $\omega^u_{t+1}\omega_t$. Intuitively, increasing employment in the current period entitles workers to future benefits which lowers the opportunity cost. The term Γ_{t+1} captures the dynamics of this effect over time, since increasing employment in the current period affects the whole path of future eligibility.

2.1.2 Opportunity Cost of Employment: Consumption and Work Differences

The second component of the opportunity cost of employment, ξ_t, results from consumption and work differences between employed and unemployed. It is useful to write it as:

$$\xi_t = \frac{[U^u(C^u_t, 0) - \lambda_t C^u_t] - [U^e(C^e_t, N_t) - \lambda_t C^e_t]}{\lambda_t}. \tag{10}$$

The first term in the numerator, $U^u_t - \lambda_t C^u_t$, is the total utility of the unemployed less the utility of the unemployed from consumption. It has the interpretation of the utility the unemployed derive solely from non-working time. Similarly, the term $U^e_t - \lambda_t C^e_t$ represents the utility of the employed from non-working time. The difference between the two terms represents the additional utility the household obtains from non-working time when moving a worker from employment to unemployment. The denominator of ξ_t is the common marginal utility of consumption. Therefore, ξ_t represents the value of non-working time in units of

consumption. This is similar to the marginal rate of substitution between non-working time and consumption in the RBC model, with the difference that the additional value of non-working time is calculated along the extensive margin.[6]

To understand the cyclical properties of ξ_t qualitatively, we linearize it around its trend. Letting x_t^* denote the approximation point of a variable x_t and $\hat{x}_t = x_t/x_t^* - 1$ the percent deviation from the approximation point, we obtain:

$$\xi_t = (\xi_t)^* - \left[\frac{(U_t^u)^* - (U_t^e)^*}{(\lambda_t)^*}\right] \hat{\lambda}_t + (p_t^e)^* \hat{N}_t, \tag{11}$$

where

$$\hat{\lambda}_t = -\rho_t^* \hat{C}_t^e + \sigma_t^* \hat{N}_t = -\rho_t^* \hat{C}_t^u. \tag{12}$$

The parameter $\rho_t > 0$ denotes the absolute value of the elasticity of the marginal utility of consumption with respect to consumption, $\sigma_t > 0$ denotes the elasticity of the marginal utility of consumption with respect to hours per employed worker, and p_t^e denotes the marginal product of an employed worker. In deriving this equation we have used the first-order condition (16) for hours in the bargaining problem described below to substitute in the term $(p_t^e)^*$.

Equation (11) states that cyclical variation in ξ_t comes from two sources. First, movements in the marginal utility of consumption affect ξ_t. When λ_t rises, the value of earning income that can be used for market consumption rises relative to the value of non-working time. Second, variation in hours per employed N_t affect ξ_t. Because N_t gives the difference in non-working time between the unemployed and the employed, when N_t falls the contribution of the unemployed relative to the employed to household utility declines. In sum, the household values most the contribution of the employed (who generate higher wage income) relative to that of the unemployed (who have higher non-working time) during recessions, when consumption is lower and the difference in non-working time between employed and unemployed is smaller.

[6] When flow utilities are $U_t^u = U_t^e$, we take $\xi_t = C_t^e - C_t^u$. In this case, our estimates of consumption differences imply that z_t is roughly 15 percent on average. To justify a z_t higher than that we need $(U_t^u - U_t^e)/\lambda_t > 0$. The interpretation of $(U_t^u - U_t^e)/\lambda_t > 0$ is that non-working time is valued at a sufficiently high level relative to consumption. This is a standard assumption in the literature (see Rogerson and Wright, 1988). In Section 8 with incomplete asset markets, we show that the unemployed's expected present value of discounted utility flows V_t^u is lower than the employed's expected present value of discounted utility flows V_t^e, even when flow utilities are $U_t^u > U_t^e$.

2.1.3 Opportunity Cost of Employment: Comparison to the Literature

The MP literature has motivated z_t as either benefits or the foregone value of non-working time, and typically assumes a constant $z_t = z$. If the value of benefits does not fluctuate, $b_t = b$, then z_t is constant if ξ_t is constant. We nest a constant ξ under two sets of restrictions on utility:

1. No disutility from hours worked and utility functions that do not depend on employment status (e.g. Shimer, 2005):

$$U_t^s = U\left(C_t^s\right),\ s \in \{e, u\} \Longrightarrow C_t^e = C_t^u, U_t^e = U_t^u \Longrightarrow \xi_t = 0 \Longrightarrow z_t = b.$$

2. Linearity in consumption, separability, and constant values of hours per worker N and the net value of home production Q (e.g. Hagedorn and Manovskii, 2008):

$$U_t^e = C_t^e - v\left(N\right), U_t^u = C_t^u + Q \Longrightarrow \xi_t = v\left(N\right) + Q \Longrightarrow z_t = b + v\left(N\right) + Q.$$

In general, the component ξ_t will vary over time if N_t or Q_t enters as an argument into either utility function, and either (i) N_t or Q_t varies over time, or (ii) utility is not linear in consumption.

2.2 Firm

The firm chooses vacancies and capital to maximize the discounted present value of dividends. It produces output using a constant returns to scale technology $Y_t = F_t(K_t, e_t N_t)$, with marginal products given by $p_t^k = \partial F_t / \partial K_t$, $p_t^n = \partial F_t / \partial (e_t N_t)$, and $p_t^e = \partial F_t / \partial e_t = p_t^n N_t$. In solving its problem the firm takes as given the path of prices and the outcome of the bargaining game. The firm maximizes its value:

$$W^f\left(e_t, \mathbf{Z_t}\right) = \max_{K_t, v_t} \left\{ Y_t - R_t K_t - w_t e_t N_t - \kappa_t v_t + E_t \tilde{\beta}_{t+1} W^f\left(e_{t+1}, \mathbf{Z_{t+1}}\right) \right\}, \tag{13}$$

subject to the law of motion for employment $e_{t+1} = (1 - s_t)e_t + m_t = (1 - s_t)e_t + q_t v_t$. In the maximization problem, the firm takes as given the stochastic discount factor of the household $\tilde{\beta}_{t+1} = \beta \lambda_{t+1} / \lambda_t$, market tightness θ_t, and the vacancy-filling probability $q_t(\theta_t)$.

The firm sets the marginal product of capital equal to the rental rate of capital, $p_t^k = R_t$. The first-order condition for vacancies is:

$$\frac{\kappa_t}{q_t(\theta_t)} = E_t \tilde{\beta}_{t+1} \left(\left(p_{t+1}^n - w_{t+1} \right) N_{t+1} + \frac{\kappa_{t+1}(1 - s_{t+1})}{q_{t+1}(\theta_{t+1})} \right). \tag{14}$$

The marginal value of an additional employed worker for the firm J_t^f consists of the increase in flow profits plus the expected discounted future marginal value:

$$J_t^f = \frac{\partial W^f(e_t, \mathbf{Z}_t)}{\partial e_t} = \left(p_t^n - w_t \right) N_t + (1 - s_t) E_t \tilde{\beta}_{t+1} J_{t+1}^f. \tag{15}$$

2.3 Labor Market Matching and Bargaining

The household and the firm split the surplus from an additional match according to the generalized Nash bargaining solution. Matching is random and the firm cannot discriminate between unemployed of different durations. Let μ denote the bargaining power of the household. Bargaining takes place over the wage w_t and hours worked N_t. The total surplus associated with the formation of an additional match, in terms of the numeraire good, is $S_t = J_t^h/\lambda_t + J_t^f$, where J_t^h/λ_t is given by equation (7) and J_t^f is given by equation (15).

Hours are determined implicitly from the first-order condition:

$$\frac{\partial S_t}{\partial N_t} = 0 \Longrightarrow -\frac{\partial U_t^e}{\partial N_t} = \lambda_t p_t^n, \tag{16}$$

which equates the marginal product of labor to the employed's marginal utility of non-working time relative to the marginal utility of consumption. With efficient bargaining, hours are chosen to maximize the joint surplus. The wage allocates the surplus between the household and the firm using the surplus-splitting rule, $(1 - \mu) J_t^h/\lambda_t = \mu J_t^f$. In Appendix C we show that this results in a standard wage equation:

$$w_t = \left(\frac{1}{N_t} \right) \left(\mu p_t^e + (1 - \mu) z_t + \mu \kappa_t \theta_t \right). \tag{17}$$

3 Data and Measurement

We construct a dataset of U.S. time series at quarterly frequency between 1961(1) and 2012(4), drawing on microdata from the CPS, SIPP, CE, and PSID, published series from the NIPA,

BLS, and various other government agencies, and historical data collected from print issues of the Economic Report of the President. Appendix A provides greater detail on the source data.

We begin by discussing a few general principles of our measurement exercise. The first is an aggregation result. Following Mortensen and Nagypal (2007), we assume that employers cannot discriminate ex-ante in choosing a potential worker with whom to bargain. Then, even if individuals have heterogeneous opportunity costs, the vacancy creation decision of the firm depends on the average opportunity cost over the set of unemployed persons. Accordingly, we estimate foregone government benefits and the expenditure decline for the average unemployed. Section 7 relaxes this assumption and allows z to differ across skill groups.

Our second general principle concerns the definition of the unemployed. Our model follows much of the literature in abstracting from the labor force participation margin. We recognize that this abstraction omits potentially important flows into and out of participation, and that it affects our measurement insofar as people move directly from non-participation to employment.[7] Nonetheless, lacking good data on search intensity, we conform whenever possible to the official Bureau of Labor Statistics U-3 definition of unemployment.

3.1 Benefits $B_{n,t}$ and $B_{u,t}$

We depart from the literature in measuring b_t in three significant ways. First, following the logic of our aggregation result, we measure the average benefit across all unemployed, rather than statutory benefit rates. This matters because, for example, only about one-third of unemployed persons receive UI on average in our sample. Second, the safety net includes a number of other programs such as supplemental nutritional assistance payments (SNAP, formerly known as food stamps), welfare assistance (TANF, formerly AFDC), and health care (Medicaid). Income from all of these programs belongs in $B_{n,t}$ to the extent that unemployment status correlates with receipt of these benefits. Finally, for UI benefits we differentiate between monetary benefits

[7] Allowing for endogenous labor force participation would not, however, affect our expression for z. For example, allowing non-employed workers to choose between unemployment and non-participation would add a first order condition to the model requiring indifference between the two states. The marginal value of adding an employed worker would remain unchanged and given by equation (7), and equation (8) would still describe the flow opportunity cost of employment.

per unemployed $B_{u,t}$ and the part of these benefits associated with the opportunity cost of employment. As equation (9) shows, the latter is lower than $B_{u,t}$ both because there exist utility costs associated with taking up benefits and because benefits expire.

Our empirical approach to measuring benefits combines micro survey data with program administrative data. Let $B_{k,t}$ denote each of the four components of total benefits, with $B_t = \sum_k B_{k,t}$ for programs $k \in \{\text{UI, SNAP, AFDC/TANF, Medicaid}\}$.[8] We measure $B_{k,t}$ as:

$$B_{k,t} = \left(\frac{(\text{survey dollars tied to unemployment status})_{k,t}}{(\text{total survey dollars})_{k,t}} \right) \left(\frac{(\text{total administrative dollars})_{k,t}}{(\text{number of unemployed})_t} \right). \tag{18}$$

We use the micro data to estimate the term in the first parentheses in equation (18), the fraction of total program spending in the survey that depends on unemployment status, and call this ratio $B_{k,t}^{\text{share}}$. We then apply $B_{k,t}^{\text{share}}$ to the ratio of dollars from program administrative data to the number of unemployed (the term in the second parentheses). Equation (18) adjusts the survey estimate of dollars tied to unemployment status by the ratio of administrative to survey dollars to correct for the well known undercounting of program benefits in surveys (Meyer, Mok, and Sullivan, 2009).

We now explain and implement our procedure to estimate $B_{k,t}^{\text{share}}$. Define $y_{k,i,t}$ as income from category k received by household or person i. We use the microdata to estimate the change in $y_{k,i,t}$ following an employment status change. To solve the time aggregation problem that arises because an individual may spend part of the reporting period employed and part unemployed, we model directly the instantaneous income of type k for an individual with labor force status $s \in \{e, u\}$. This is given by:

$$y_{k,i,t}^{s} = \phi_k \mathbf{X}_i + y_{k,t}^{e} + \beta_{k,t} \mathbb{I}\{s_{i,t} = u\} + \epsilon_{k,i,t}, \tag{19}$$

where $\mathbf{X}_i$ denotes a vector of individual characteristics, $y_{k,t}^{e}$ the income of a hypothetical employed, and $\mathbb{I}\{s_{i,t} = u\}$ is an indicator function taking the value of one if the individual is unemployed at time t. According to this process, income from program k increases discretely

[8] We also investigated the importance of housing subsidies. We found their importance quantitatively trivial, so we omit them from the analysis.

by $\beta_{k,t}$ during an unemployment spell. Integrating over the reporting period and taking first differences to eliminate the person fixed effect yields:

$$\Delta y_{k,i,t} = \beta^0_{k,t} + \beta_{k,t}\Delta D^u_{i,t} + \Delta\beta_{k,t}D^u_{i,t-1} + \Delta\epsilon_{k,i,t}, \tag{20}$$

where $\beta^0_{k,t} = \Delta y^e_{k,t}$, and the variable $D^u_{i,t}$ measures the fraction of the reporting period that an individual spends as unemployed.

By definition, $B^{\text{share}}_{k,t}$ is:

$$B^{\text{share}}_{k,t} = \frac{(\text{survey dollars tied to unemployment status})_{k,t}}{(\text{total survey dollars})_{k,t}} = \beta_{k,t}\frac{\sum_i \omega_{i,t}D^u_{i,t}}{\sum_i \omega_{i,t}y_{k,i,t}}, \tag{21}$$

where $\omega_{i,t}$ is the survey sampling weight for individual i in period t. Substituting equation (21) into equation (20) gives a direct estimate of $B^{\text{share}}_{k,t}$ from the regression:

$$\Delta y_{k,i,t} = \beta^0_{k,t} + B^{\text{share}}_{k,t}\Delta\tilde{D}_{i,t} + \Delta\beta_{k,t}D^u_{i,t-1} + \Delta\epsilon_{k,i,t}, \tag{22}$$

where $\Delta\tilde{D}_{i,t} = \Delta D^u_{i,t}\sum_i \omega_{i,t}y_{k,i,t}/\sum_i \omega_{i,t}D^u_{i,t}$.

We implement equation (22) using both the March CPS with households matched across consecutive years, starting in 1989, and the SIPP starting in 1996. Appendix A describes the surveys and our sample construction. In each survey, we construct a measure of unemployment at the individual level that mimics the BLS U-3 definition. The U-3 definition of unemployment counts an individual as working if he had a job during the week containing the 12th of the month (the survey reference week), and as in the labor force if he worked during the reference week, spent the week on temporary layoff, or had any search in the previous four weeks.[9]

We aggregate unemployment and income up to the level at which the benefits program is administered. In particular, in the regressions with UI income as the dependent variable, the

[9] In the March Supplement, we count an individual as in the labor force only for those weeks where he reports being on temporary layoff or actually searching during the previous year. In the SIPP, we count an individual as employed if he worked in any week of the month, rather than only if he worked during the BLS survey reference week. Accordingly, we define the fraction of time an individual is unemployed as:

$$D^{u,\text{CPS}}_{i,t} = \left[\frac{\text{weeks searching or on temporary layoff in year } t}{\text{weeks in the labor force in year } t}\right]_i,$$

$$D^{u,\text{SIPP}}_{i,t} = \frac{1}{4}\sum_{m=1}^{4}\mathbb{I}\left\{[\text{non-employed, at least 1 week of search or layoff}]_{i,t-m}\right\}.$$

Table 1: Share of Government Program Benefits Belonging to B

	UI	SNAP	TANF	Medicaid
CPS (1989-2012)				
B^{share}	0.880	0.072	0.063	0.026
Standard error	(0.018)	(0.005)	(0.011)	(0.003)
Observations	455,216	255,310	296,340	255,310
SIPP (1996-2012)				
B^{share}	0.632	0.037	0.036	
Standard error	(0.009)	(0.002)	(0.004)	
Observations	1,480,993	968,718	968,779	
Mean of B^{share} (CPS and SIPP)	0.756	0.054	0.049	0.026

The table reports summary statistics based on OLS regressions of equation (22), where B^{share} is defined by equation (21). The regressions are weighted using sampling weights in each year, with the weights normalized such that all years receive equal weight. Standard errors are based on heteroskedastic robust (CPS, non-UI), heteroskedastic robust and clustered by family (CPS, UI), or heteroskedastic robust and clustered by household (SIPP) variance matrix.

unit of observation is the individual and we cluster standard errors at the household level. In regressions for SNAP, TANF, and Medicaid, the unit of observation is the family average of unemployment and the family total of income.

Table 1 reports results based on OLS regressions of equation (22) and constraining $B^{\text{share}}_{k,t}$ to be constant over time.[10] For UI, the average B^{share} is 0.76. If only unemployed persons received UI, then this share would equal one. In fact, roughly one-quarter of UI income reported in a year goes to recipients who report having had no unemployment spells. These individuals may have had part-time employment in states that have positive labor income caps for receipt of UI, or may have claimed UI without actually exerting search effort.[11]

[10] We have correlated the cyclical component of an estimated time-varying $B^{\text{share}}_{k,t}$ with the cyclical component of the unemployment rate, and in almost all cases we cannot reject the hypothesis that $B^{\text{share}}_{k,t}$ is acyclical. The largest absolute correlation is 0.33, and the mean correlation is 0.07. Only in the case of SNAP in the SIPP can we reject a zero correlation at a ten percent confidence level.

[11] The fraction of UI income reported by non-unemployed has also risen since the early 1990s, such that part of the difference in the B^{share} found in the CPS and the SIPP stems from the longer CPS sample. Using a separate audit methodology, the Department of Labor estimates that roughly 10 percent of UI payments go to ineligible recipients (http://ows.doleta.gov/unemploy/improp_pay.asp).

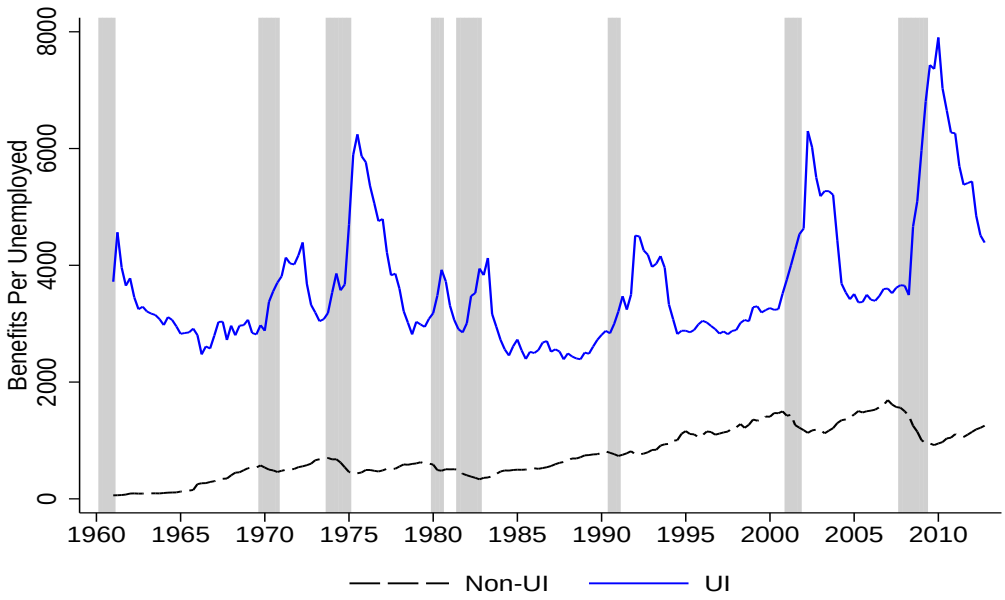

Figure 1: Time Series of Benefits Per Unemployed

Only five percent of SNAP and TANF and three percent of Medicaid spending appear in $B_{n,t}$. We find these estimates reasonable. Roughly two-thirds of Medicaid payments accrue to persons who are over 65, blind, or disabled (Centers for Medicare and Medicaid Services, 2011, Table II.4). Moreover, even prior to implementation of the Affordable Care Act, all states had income limits for coverage of children of at least 100 percent of the poverty line, and half of states provided at least partial coverage to working adults with incomes at the poverty line (Kaiser Family Foundation, 2013). Similarly for SNAP, tabulations from the monthly quality control files provided by Mathematica indicate that no more than one-quarter of SNAP benefits go to households with at least one member unemployed. Given observed statutory phase-out rates and deductions, 5 percent appears as a reasonable estimate.

To summarize, to measure $B_{n,t}$ and $B_{u,t}$ we first use micro survey data to estimate the share of each program's total spending associated with unemployment, B_k^{share}. We then apply this share to the total spending observed in administrative data. As a result, B_t inherits directly the cyclical properties of the program administrative data. Although the B_k^{share}'s for the non-UI programs are small, the standard errors strongly indicate that they are not zero. We plot the resulting time series of $B_{n,t}$ and $B_{u,t}$ in constant 2009 dollars in Figure 1.

3.2 Consumption

To measure the ξ_t component of the opportunity cost, we require time series of consumptions of the employed C_t^e and the unemployed C_t^u, and estimates of preference parameters. Let $s \in \{e, u, n, r\}$ denote persons 16 years or older who are employed, unemployed, out of the labor force but of working age (16-64), and older than 65 years old, respectively. Let π_t^s be the fraction of the population belonging in each group. Time series of π_t^s come directly from published tabulations by the BLS.

Denote by C_t^s consumption expenditures on non-durables and services per member of group s. We have the adding-up identity:

$$\pi_t^e C_t^e + \pi_t^u C_t^u + \pi_t^n C_t^n + \pi_t^r C_t^r = C_t^{\text{NIPA}}, \tag{23}$$

where C_t^{NIPA} is NIPA consumption per person 16 years or older. Defining $\gamma_t^s = C_t^s / C_t^e$ as the ratio of consumption in status s to consumption when employed, we solve equation (23) for the consumption of an employed:

$$C_t^e = \frac{C_t^{\text{NIPA}}}{\sum_s \pi_t^s \gamma_t^s}. \tag{24}$$

Equation (24) together with estimates of the consumption ratios γ_t^s provide the basis for deriving the time series of consumptions for the employed and unemployed and for calibrating the model.

We now turn to our estimates of γ_t^s. Let $C_{i,k,t}^s$ denote the instantaneous expenditure on consumption category k of an individual i in group s at time t. We assume that while employed, individual i has expenditure $C_{i,k,t}^e = \exp\{\phi_{k,t}\mathbf{X}_{i,t} + \epsilon_{k,i,t}\}\tilde{C}_{k,t}$, where $\mathbf{X}_{i,t}$ denotes a vector of demographic characteristics, $\phi_{k,t}$ a vector of parameters, $\epsilon_{k,i,t}$ a mean zero idiosyncratic component uncorrelated with employment status, and $\tilde{C}_{k,t}$ is a base level of consumption. Using the definition of $\gamma_{k,t}^s$:

$$C_{k,i,t}^s = \gamma_{k,t}^s \exp\{\phi_{k,t}\mathbf{X}_{i,t} + \epsilon_{k,i,t}\}\tilde{C}_{k,t}. \tag{25}$$

For a working age individual with potential status e, u or n, we integrate over the reporting

period and take logs to obtain:[12]

$$\ln C_{k,i,t} = \gamma^0_{k,t} + \phi_{k,t}\mathbf{X}_{i,t} + \left(\tilde{\gamma}^u_{k,t} - 1\right) D^u_{i,t} + \left(\tilde{\gamma}^n_{k,t} - 1\right) D^n_{i,t} + \epsilon_{k,i,t}, \tag{26}$$

where $\gamma^0_{k,t} = \ln \tilde{C}_{k,t}$, and the variables $D^u_{i,t}$ and $D^n_{i,t}$ measure the fraction of time an individual spends as unemployed and out of the labor force, respectively.[13] In equation (26), $\tilde{\gamma}^s_{k,t} - 1$ denotes the difference between the log of consumption of an individual in group s and the log of consumption of an employed. Therefore, to recover the actual consumption ratios $\gamma^s_{k,t}$ from the log point differences we use the formula $\gamma^s_{k,t} = \exp\left(\tilde{\gamma}^s_{k,t} - 1\right)$.

We estimate equation (26) using the Consumer Expenditure Survey (CE). The CE asks respondents for the number of weeks worked over the previous year, but does not ask questions about search activity while not working. We set $D^n_{i,t} = 1$ if the respondent reports zero weeks worked over the previous year, and does not give "unable to find job" as the reason for not working. For the rest of the respondents, we define $D^u_{i,t} = 1 - (\text{weeks worked})_{i,t}/52$. We average $D^u_{i,t}$ and $D^n_{i,t}$ at the household level.

We focus our discussion of results on the unemployment margin because γ^u will directly inform our calibration of preferences. Figure 2 reports γ^u_t by year, for the aggregate category of nondurable goods and services, less housing, health, and education. The mean of γ^u_t implies a 21 percent decline in expenditure on nondurable goods and services during unemployment. The series does not exhibit any apparent cyclicality, with a correlation between the cyclical components of γ^u_t and the unemployment rate of -0.03. We also test for cyclicality parametrically by interacting $D^u_{i,t}$ in equation (26) with both the state and local unemployment rates, and again cannot reject the hypothesis that the consumption ratio is acyclical (see Table A.1).

In Appendix D we present further analysis of the consumption ratios. We first discuss the rich set of controls that we include in $\mathbf{X}_{i,t}$ to capture taste shocks and proxies for permanent

[12] In deriving our estimating equation we replace the term $\ln\left[1 - \sum_s \left(1 - \gamma^s_{k,t}\right) D^s_{i,t}\right]$ with $\sum_s D^s_{i,t}\left(\tilde{\gamma}^s_{k,t} - 1\right)$, where the coefficients $\tilde{\gamma}^s_{k,t}$ are related to the coefficients $\gamma^s_{k,t}$ and to terms of order higher than one in the linear approximation of the left-hand side around $\gamma^s_{k,t} = 1, \forall s$. Derivations for the estimating equation for γ^r_t proceed analogously.

[13] The derivation of equation (26) assumes that $\tilde{\gamma}^u_{k,t}$ does not vary with unemployment duration $D^u_{i,t}$. In unreported regressions, we have estimated $\tilde{\gamma}^u_{k,t}$ non-parametrically by grouping households into bins of weeks unemployed. Our estimated $\tilde{\gamma}^u_{k,t}$ for each bin indicates a duration-independent $\tilde{\gamma}^u_{k,t}$. This finding supports the assumption in the model that the instantaneous consumption of the unemployed does not depend on duration.

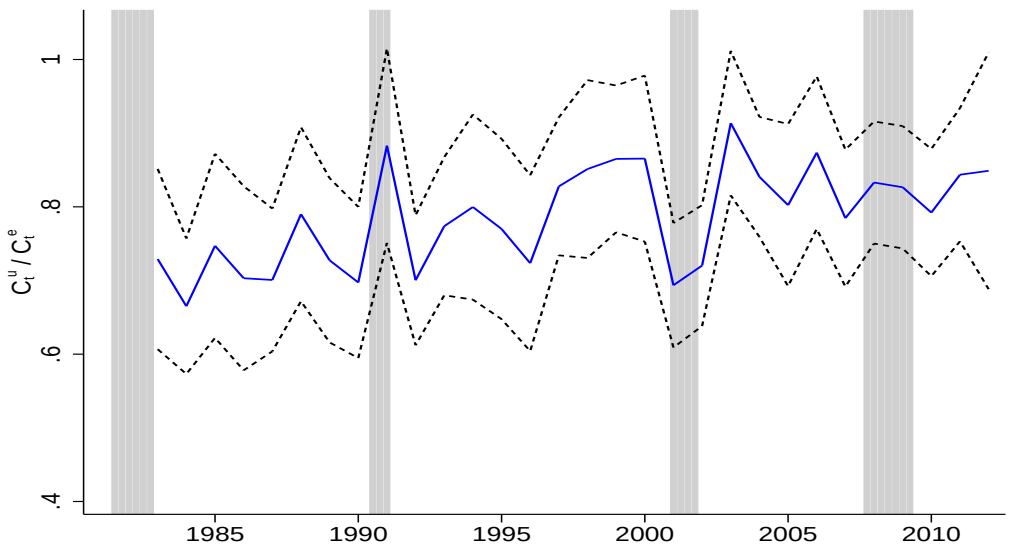

Figure 2: Decline in Nondurables and Services Upon Unemployment

Notes: The solid line reports the estimates of $\gamma_t^u = \exp(\tilde{\gamma}_{u,t} - 1)$, where $\tilde{\gamma}_{u,t}$ is estimated from equation (26) using data from the CE. The dotted lines give 95 percent confidence interval bands based on robust standard errors. Regressions are weighted using survey sampling weights. See Appendix D for included covariates.

income in our regression in the CE. Second, we use the panel dimension of the PSID and a first-differenced version of equation (26) to validate the sufficiency of the control variables to absorb differences in permanent income in the cross-section of households. Our estimates of γ^u from the PSID line up very well with our estimates from the CE for overlapping categories of consumption. Third, we show that our estimates of the consumption decline upon unemployment are broadly consistent with estimates from previous studies.

We apply the consumption ratios to equation (24) in two steps. First, the calibration of the preference parameters in Section 4 requires data on the mean level of C_t^e (denoted by C^e) and the mean level of C_t^u (denoted by C^u). For these, we impose constancy of the consumption ratios $\gamma_t^s = \gamma^s$ in equation (24) and obtain a time series for C_t^e and $C_t^u = \gamma^u C_t^e$.[14] We then define $C^e = (1/T)\sum_t C_t^e$ and $C^u = (1/T)\sum_t C_t^u$. Second, the time series of ξ_t requires time series of C_t^e and C_t^u. We use the first-order condition (4) to define a time-varying consumption ratio $\bar{\gamma}_t^u$ that makes the first-order condition in the model hold exactly in the data. We apply

[14] We set $\gamma^u = 0.793$, the value from estimating equation (26) for a constant γ^u. For the other categories, we estimate $\gamma^n = 0.743$ and $\gamma^r = 0.940$. Similarly to our estimates of γ^u, we cannot reject acyclicality of these consumption ratios.

this time-varying ratio $\gamma_t^u = \bar{\gamma}_t^u$ in equation (24) to obtain a time series for C_t^e and then define $C_t^u = \bar{\gamma}_t^u C_t^e$.[15]

3.3 Other Variables

The number of employed comes from the monthly CPS for consistency with our unemployment variable (BLS series LNS12000000). With a constant labor force, the number of newly unemployed workers equals the product of the previous period's separation rate s_{t-1} and stock of employed workers e_{t-1}. We therefore define the separation rate s_t at quarterly frequency as the ratio of the number of workers unemployed for fewer than 15 weeks in quarter $t+1$ (using the sum of BLS series LNS13008397 and LNS13025701) to the number of employed workers in t. The separation rate and the unemployment rate allow us to calculate the job-finding rate f_t from the law of motion for unemployment $u_{t+1} = u_t(1-f_t) + s_t(1-u_t)$.[16]

We construct estimates of UI benefits per recipient $\tilde{B}_t$, the fraction of unemployed receiving UI benefits ϕ_t, the fraction eligible ω_t, and the fraction of eligible who take up benefits ζ_t, as follows. The Department of Labor provides data on the number of UI recipients in all tiers (state regular benefits, extended benefits, and federal emergency benefits) beginning in 1986. We extend this series back to 1961 using data from Statistical Appendix B of the Economic Report of the President. Dividing the NIPA total of UI benefits paid by the number of UI recipients gives a time series of UI benefits per recipient $\tilde{B}_t$. The fraction of unemployed receiving benefits is $\phi_t = B_{u,t}/\tilde{B}_t$, where $B_{u,t}$ is our estimated UI benefit per unemployed in Section 3.1.

We estimate ω_t using its law of motion in equation (3) and data on u_t, s_t, f_t, ω_t^e, and ω_t^u. We measure the probability ω_t^e that a newly unemployed worker is eligible using the fact that

[15] This approach ensures the internal consistency of our estimated parameters with the model's analog of the first-order condition for risk sharing in the data. The time-varying $\bar{\gamma}_t^u$ is extremely smooth, falling comfortably in the confidence interval of the estimated γ_t^u. See Appendix F for more details.

[16] We recognize the point of Shimer (2012) that this procedure understates the amount of gross flows between unemployment and employment because some workers will separate and find a new job within the period. However, a discrete time calibration must accept this shortcoming if both the law of motion for unemployment holds and the share of newly unemployed matches the share in the data. For our purposes, matching the share of newly unemployed matters more than matching the level of gross flows. Estimating s_t and f_t at a monthly frequency, which should substantially mitigate the bias from within-period flows, and then averaging at the quarterly level makes little difference for our results.

workers who quit their jobs and new labor force entrants are ineligible for UI.[17] From the CPS basic monthly microdata, we construct the number of unemployed for less than five weeks who report "job loser" as their reason for unemployment. We add to this total the product of the number of re-entrants who have worked in the past 12 months, and the 6 month lag of the fraction of job losers among those moving from employment to unemployment. Dividing by the number of total unemployed for less than five weeks then gives an estimate of the fraction of the newly unemployed that satisfy non-monetary eligibility. We tie cyclical movements of ω_t^e to cyclical movements in this fraction.[18] We center ω_t^e around 0.61 to target a mean take-up rate ζ_t of roughly 0.65.

We set ω_t^u, the probability that an unemployed remains eligible, such that the expected potential duration of eligibility equals the national maximum. We adjust this probability for the fact that not every unemployed individual has the maximal potential duration (see Appendix A for further details). Evaluating equation (3) using the time series of u_t, s_t, f_t, ω_t^e, and ω_t^u gives our time series of eligibility ω_t. The take-up rate equals $\zeta_t = \phi_t/\omega_t$.

We measure hours per worker N_t from CPS basic monthly microdata starting from 1968. We extend the series back to 1961 and fill in some missing months between 1968 and 1975 using data from Cociuba, Prescott, and Ueberfeldt (2012). In Appendix A we describe how we seasonally adjust the resulting series. We normalize the mean of N_t in the sample to be one.

The marginal product of employment p_t^e is defined as $1-\nu$ multiplied by real GDP and then divided by the number of employed, where $\nu = 0.333$ is the elasticity of output with respect to capital in the production function. The marginal product of labor is defined as $p_t^n = p_t^e/N_t$. We divide in the data variables such as real GDP, consumption, benefits, and the opportunity cost of employment by the mean of the marginal product of employment in the sample. Therefore, all variables (both in the data and in the model) are expressed relative to the mean level of $p^e = 1$.

[17] We do not have information on monetary eligibility at cyclical frequencies. It seems reasonable that monetary eligibility would be procyclical, as newly unemployed transition from weaker labor markets during recessions. In that case, ignoring monetary eligibility leads us to understate the volatility of ζ and ultimately of z.

[18] Prior to 1968, we impute this share using the fitted values from a regression of the share on leads and lags of the unemployment rate and of the fraction of job losers among all durations of unemployed.

4 Parameterization

We parameterize the model so that in the deterministic steady state it matches key features of the United States economy between 1961(1) and 2012(4). Appendix E defines formally the equilibrium of the model. Variables without time subscripts denote steady state values. We start by discussing the functional forms. The production function and matching technology are Cobb-Douglas:

$$Y_t = A_t K_t^{\nu} \left(e_t N_t\right)^{1-\nu}, \tag{27}$$

$$m_t = M_t v_t^{\eta} u_t^{1-\eta}, \tag{28}$$

where A_t is the exogenous technology level, ν denotes the elasticity of output with respect to capital, M_t is the exogenous matching efficiency, and η is the elasticity of matches with respect to vacancies. The cost function for taking up benefits is given by:

$$\psi_t = \frac{\Psi}{\alpha} \zeta_t^{\alpha}, \tag{29}$$

where Ψ denotes an exogenous shifter of the cost function and α denotes the elasticity of the cost function with respect to the take-up rate.

The utility functions of the employed and the unemployed are given by:

$$U_t^e = \left(\frac{1}{1-\rho}\right)\left((C_t^e)^{1-\rho}\left(1 - \frac{(1-\rho)\chi\epsilon}{1+\epsilon} N_t^{1+1/\epsilon}\right)^{\rho} - 1\right), \tag{30}$$

$$U_t^u = \left(\frac{1}{1-\rho}\right)\left((C_t^u)^{1-\rho} - 1\right) + Q. \tag{31}$$

The preferences given by equation (30) are consistent with balanced growth and feature a constant Frisch elasticity of labor supply ϵ (Shimer, 2010; Trabandt and Uhlig, 2011). The parameter χ determines the disutility of hours worked. The parameter ρ governs both the intertemporal elasticity of substitution $(1/\rho)$ and the degree of complementarity between consumption and hours worked. When $\rho \to 1$ utility becomes separable between consumption and hours worked. The employed consume more (market) goods than the unemployed in our model when $\rho > 1$, that is, when consumption and hours worked are complements.

Table 2: Parameter Values

Description	Parameter	Value	Description	Parameter	Value
		Externally Calibrated			
Discount Factor	β	0.990	Depreciation Rate	δ	0.025
Capital Elasticity	ν	0.333	Frisch Elasticity	ϵ	0.700
Matching Elasticity	η	0.400	Bargaining Power	μ	0.600
UI Initial Eligibility	ω^e	0.609	UI Expiration	ω^u	0.494
UI per Recipient	$\tilde{B}$	0.215	Non-UI Benefits	B_n	0.015
Separation Rate	s	0.045	Utility Shifter	Q	varies with z
		Internally Calibrated			
Matching Efficiency	M	0.802	Technology Level	A	0.619
Vacancy Cost	κ	varies with z	Preference Parameter	χ	1.689
Preference Parameter	ρ	1.376	Exogenous Spending	C^o	varies with z
UI Cost Elasticity	α	1.759	UI Cost Level	Ψ	0.954

The parameter Q in the utility of the unemployed is a separable shifter that we will use to target different levels of z. This parameter has the interpretation of the additional utility from consuming non-market goods produced at home by the unemployed. In our baseline results Q does not vary over time. Equivalently, the value of home production in terms of market consumption, Q/λ_t, changes over time because of variations in the marginal utility of market consumption and not because of home production shocks. Below we argue that this is a conservative assumption for our results.

Table 2 summarizes the calibrated parameters. Recall that a model period equals a quarter. We set the discount factor to $\beta = 0.99$, the depreciation rate to $\delta = 0.025$, and the elasticity of output with respect to capital to $\nu = 0.333$. Following Pistaferri (2003) and Hall (2009), we set the Frisch elasticity of labor supply to $\epsilon = 0.70$. Following Mortensen and Nagypal (2007), we set the elasticity in the matching function to $\eta = 0.40$. We set the worker's bargaining power to $\mu = 0.60$ to satisfy the Hosios condition.

To calibrate the model we take averages over 1961(1) and 2012(4) of the variables con-

structed in Section 3. We estimate an average separation rate of $s = 0.045$. Together with an average job-finding probability of $f = 0.704$, this implies a steady state unemployment rate of $u = 0.06$. We find in-sample averages of the UI eligibility parameters $\omega^e = 0.609$ and $\omega^u = 0.494$. Finally, we estimate $\tilde{B} = 0.215$ and $B_n = 0.015$, both expressed relative to the mean marginal product of employment.

We calibrate eight parameters $\{M, A, \kappa, \chi, \rho, C^o, \alpha, \Psi\}$ to match eight targets. Inverting the matching function and imposing that in steady state the job-finding rate is $f = 0.704$ yields a value for the matching efficiency parameter M. The technology level A is chosen such that the marginal product of labor equals $p^n = 1$. To calibrate the vacancy creation costs κ we use the fact that the steady state equilibrium tightness is given by:

$$\theta = \left(\frac{1}{\kappa}\right)\left(\frac{f\beta(1-\mu)}{1-\beta(1-s-\mu f)}\right)(p^e - z)\,. \tag{32}$$

den Haan, Ramey, and Watson (2000) estimate a monthly job-filling probability of 71 percent, which translates to $q = 0.975$ at a quarterly frequency. Given the value of $f = 0.704$, this produces a market tightness of $\theta = f/q = 0.722$. Below we will set parameters such that $N = 1$ and hence $p^e = p^n N = 1$. Given a value of z, this leaves κ as the only free parameter in equation (32), so for each z we calibrate κ to hit the same level of $\theta = 0.722$.

The parameters χ and ρ in the utility function are calibrated from the first-order condition for risk sharing (4) and for hours per worker (16). In these conditions we target levels of consumption $C^e = 0.543$ and $C^u = 0.430$ that we obtain from our estimates in Section 3.2. The calibrated value of $\rho = 1.376$ implies an intertemporal elasticity of substitution for consumption of around 0.73 and an elasticity of hours with respect to the marginal utility of wealth equal to $\epsilon_{N\lambda} = \epsilon/\rho = 0.51$.[19] Finally, the parameter C^o is calibrated such that the resource constraint $Y = eC^e + (1-e)C^u + C^o + X + \kappa v$ holds at the values $N = 1$, $C^e = 0.543$ and $C^u = 0.430$. We present details of this calibration in Appendix C.

To parameterize the costs of taking up benefits, we note that the first-order condition for

[19] These values are close to the values of 0.5 and 0.4 used in Hall (2009), with the difference explained by the fact that our estimated ρ is lower than the ρ implicit in Hall's formulation.

the take-up rate (6) implies:

$$\hat{\zeta}_t = \left(\frac{1}{\alpha - 1}\right)\left(\hat{\lambda}_t + \hat{\tilde{B}}_t\right). \tag{33}$$

In our sample, a regression of the percent deviation of ζ_t from its trend on the percent deviation of $\lambda_t \tilde{B}_t$ from its trend yields an estimated value of $\alpha = 1.759$ (standard error 0.151). Given this value of α, we pick the parameter Ψ to target a steady state take-up rate of $\zeta = 0.652$.

We consider three baseline values for the level of the opportunity cost of employment z. The value $z = 0.452$ comes from setting $Q = 0$ in the unemployed's utility in equation (31). This value turns out to be close to the value of 0.4 used in Shimer (2005). However, it results from a calibration strategy different from that of Shimer (2005), and one more similar in spirit to that used in Hall and Milgrom (2008). Hall and Milgrom also compute the part of the opportunity cost associated with consumption and work differences ξ using a utility function with curvature. Our calibration differs from theirs, however, in that our estimated opportunity cost associated with benefits b is much lower. Specifically, because we find $b = 0.041$, when we set $Q = 0$ we obtain $\xi = 0.411$ and $z = b + \xi = 0.452$. Hall and Milgrom (2008) have $z = 0.71$ with $b = 0.25$, and so their implied $\xi = 0.46$ is close to our $\xi = 0.411$ under $Q = 0$.

The second level of z that we consider is the Hall and Milgrom (2008) value of $z = 0.71$. To achieve this value of z, we set Q in the utility function of the unemployed at a level such that Q/λ equals 0.259 (relative to a marginal product of one). The final value that we consider is taken from Hagedorn and Manovskii (2008). Specifically, we calibrate Q to achieve a level of $z = 0.955$. This requires setting the utility value of home production goods in terms of consumption goods to $Q/\lambda = 0.504$.

5 The Opportunity Cost of Employment in the Data

Figure 3 plots the percent deviation of the opportunity cost of employment from its trend, $\hat{z}_t$, along with the percent deviation of output from its trend, $\hat{Y}_t$, between 1961(1) and 2012(4). We plot $\hat{z}_t$ for the middle level of $z = 0.71$ discussed above. Table 3 presents business cycle statistics of the opportunity cost z_t for each of the three cases $z = 0.452$, $z = 0.71$, and

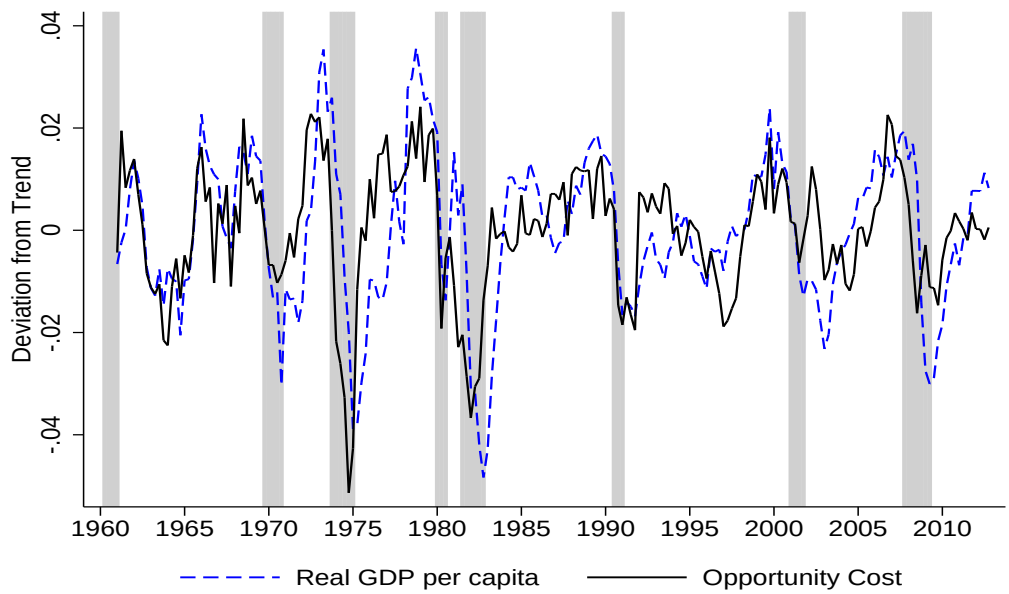

Figure 3: Cyclicality of the Opportunity Cost of Employment

Notes: The solid line is the percent deviation of the cyclical component of Y_t (real GDP per person aged 16 and older) from its trend. The sold line is the percent deviations of the cyclical components of the opportunity cost of employment z_t from its trend, generated under the value of the parameter Q in equation (31) that yields a steady state $z = 0.71$. Variables are logged and HP-filtered with a smoothing parameter of 1,600.

$z = 0.955$. The opportunity cost is more volatile than the marginal product over the business cycle and comoves strongly with output.

Table 3 also reports the elasticities $\epsilon(\hat{x}_1, \hat{x}_2)$ of two variables x_1 and x_2 as a metric that takes into account both the correlation between the two variables and the relative volatilities. For consistency with prior literature that focuses on A_t shocks as the driving force in the model, and to correct for measurement error in $\hat{p}^e_t$, in constructing the elasticities we instrument for $\hat{p}^e_t$ using the Fernald (2012) unadjusted TFP series.[20] The resulting elasticities of $\hat{z}$ with respect to $\hat{p}^e$ range between 0.9 and 1.2, leading to our preferred value of roughly one.

To better understand the forces driving the procyclicality of the opportunity cost, Figure 4 plots the cyclical component of the opportunity cost (for the case of $z = 0.71$, the solid line) together with four counterfactual series (the dashed lines). The upper left panel shows the opportunity cost under the counterfactual that the numerator of ξ_t in equation (10) does

[20] Without instrumenting all of the elasticities decline by roughly 0.1-0.2, consistent with the existence of measurement error in productivity causing attenuation bias.

Table 3: Cyclicality of Wages and Unemployment Fluctuations

	$z = 0.452$		$z = 0.710$		$z = 0.955$	
Statistic	HP-filter	BK-filter	HP-filter	BK-filter	HP-filter	BK-filter
$\text{sd}\left(\hat{Y}\right)$	1.51	1.48	1.51	1.48	1.51	1.48
$\text{sd}\left(\hat{u}\right)$	11.82	11.58	11.82	11.58	11.82	11.58
$\text{sd}\left(\hat{z}\right)$	1.47	1.38	1.24	1.20	1.17	1.15
$\text{sd}\left(\hat{\xi}\right)$	1.61	1.51	1.31	1.28	1.23	1.22
$\text{sd}\left(\hat{b}\right)$	8.08	7.36	8.08	7.36	8.08	7.36
$\text{sd}\left(\hat{p}^e\right)$	0.89	0.85	0.89	0.85	0.89	0.85
$\text{corr}\left(\hat{z}, \hat{Y}\right)$	0.56	0.59	0.58	0.60	0.57	0.60
$\text{corr}\left(\hat{\xi}, \hat{Y}\right)$	0.82	0.87	0.77	0.80	0.72	0.75
$\text{corr}\left(\hat{b}, \hat{Y}\right)$	−0.41	−0.46	−0.41	−0.46	−0.41	−0.46
$\text{corr}\left(\hat{p}^e, \hat{Y}\right)$	0.76	0.79	0.76	0.79	0.76	0.79
$\epsilon\left(\hat{u}, \hat{p}^e\right)$	−7.13	−7.49	−7.13	−7.49	−7.13	−7.49
	(0.99)	(1.00)	(0.99)	(1.00)	(0.99)	(1.00)
$\epsilon\left(\hat{z}, \hat{p}^e\right)$	1.13	1.18	1.00	1.06	0.95	1.02
	(0.12)	(0.10)	(0.11)	(0.09)	(0.11)	(0.09)

Notes: We denote the percent deviation of some variable x_t from its trend by $\hat{x}_t$. We compute trends of variables either with the HP-filter or with the BK-filter. For the HP-filter we use a smoothing parameter of 1,600. For the BK-filter, we isolate fluctuations with a period of length ranging between 6 and 32periods. The elasticity $\epsilon\left(\hat{x}_1, \hat{x}_2\right)$ is the regression coefficient of $\hat{x}_1$ on $\hat{x}_2$, instrumented using the Fernald (2012) TFP series. Newey-West standard errors with four lags are in parentheses.

not vary cyclically. Recall that the numerator of ξ_t has the interpretation of the utility value of non-working time that a person forgoes when moving from unemployment to employment, and fluctuates procyclically because, *inter alia*, hours per worker move procyclically. Shutting down the cyclicality of this component makes the opportunity cost slightly less volatile.

The upper right panel shows the opportunity cost when the denominator of ξ_t always equals its trend. The denominator of ξ_t is the marginal utility of consumption λ_t, and converts the utility flow into units of consumption. Removing cyclical fluctuations from the denominator of ξ_t leads to a visibly less cyclical and volatile opportunity cost.

The bottom two panels show the opportunity cost when the entire ξ_t term always equals

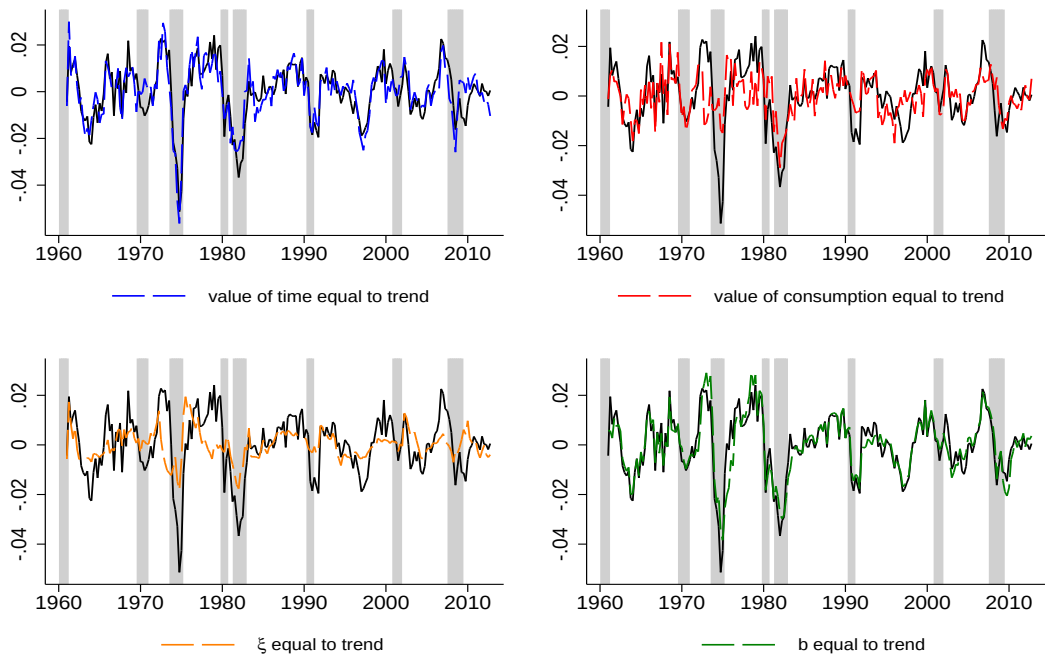

Figure 4: Cyclicality of the Opportunity Cost of Employment

Notes: The solid line is the percent deviation of the cyclical component of z_t from its trend for the case of $z = 0.71$. The dashed lines show the counterfactual cyclical components of z_t. Variables are logged and HP-filtered with a smoothing parameter of 1,600.

its trend (left panel) and when b_t always equals its trend (right panel). Comparing the two panels, the ξ_t component clearly dominates the movements of the sum $z_t = b_t + \xi_t$. Setting ξ_t always equal to its trend makes the opportunity cost roughly acyclical. In contrast, setting b_t always equal to its trend barely affects the cyclicality of the opportunity cost.

The low average level of b_t explains why its cyclicality barely affects the cyclicality of z_t. Why do we estimate so low a level of b, despite the fact that we add non-UI benefits? The sample-average of benefits per recipient, $\tilde{B}$, is 21.5 percent of the marginal product, close to the statutory replacement rates used by Mortensen and Nagypal (2007) and Hall and Milgrom (2008) and the rate suggested by Hornstein, Krusell, and Violante (2005). However, only about one-third of unemployed actually receive benefits, $\phi = \omega\zeta = 0.334$. Therefore, the "effective" rate is much smaller than benefits per recipient, $B_u = \phi\tilde{B} = 0.072$. Adjusting for benefits expiration and take up costs reduces further the UI component such that, even after adding

the sample mean of $B_{n,t}$, we obtain a sample average of $b = 0.041$.[21]

To assess the sensitivity of the cyclicality of ξ_t to our estimation procedure, recall that the mapping from data to ξ_t depends on our estimate of $\rho = 1.376$. Under the preferences given by equation (30), ρ affects both the curvature of the utility function and the degree of non-separability between consumption and hours. A natural question is how our results would change under the log-separable preferences common to the RBC literature. To answer this question, we re-estimate our model for the case of $z = 0.71$ assuming a counterfactual zero decline in consumption upon unemployment (instead of a 21 percent decline). With $C_t^e = C_t^u = C_t$ our calibration yields $\rho = 1$, in which case the utility function of the employed collapses to the standard log-separable utility $U^e = \log C_t^e - \frac{\chi\epsilon}{1+\epsilon} N_t^{1+1/\epsilon}$, and the utility of the unemployed is $U^u = \log C_t^u + Q$. This gives $\xi_t = \left(\frac{\chi\epsilon}{1+\epsilon} N_t^{1+1/\epsilon} + Q\right)/\lambda_t$, with $\lambda_t = 1/C_t$.[22]

Both ξ_t and z_t become slightly more procyclical and volatile when $\rho = 1$. The confluence of three channels leads to this result. First, the smaller curvature in the utility function causes the volatility of the cyclical component of the marginal utility λ_t to decline from 1.15 percent when $\rho = 1.376$ to 0.96 percent when $\rho = 1$. By itself, lower volatility of λ_t makes ξ_t less volatile. Second, the non-separable preferences in our baseline specification attenuate the correlation of λ_t with output from -0.70 ($\rho = 1.376$) to -0.47 ($\rho = 1$), as the decline in hours during recessions partly counteracts the decline in consumption. Third, the correlation between the numerator and the denominator of ξ_t decreases from 0.11 under our baseline estimate ($\rho = 1.376$) to -0.46 ($\rho = 1$), again because of the positive comovement of consumption and hours per worker. The absence of non-separability therefore has the effect of increasing the correlation of ξ_t with the business cycle and raising its volatility. Putting these forces together, with $\rho = 1$ the correlation of $\hat{z}_t$ and $\hat{y}_t$ rises to 0.68, and the standard deviation of $\hat{z}_t$ rises to 1.34 percent.

Two last aspects of our construction of z_t merit discussion. First, in constructing our time

[21] The sensitivity of reported reservation wages to UI benefits suggest that, if anything, our b_t may be too large. In our model, the increase of the reservation wage for individuals already receiving UI is given by the bracketed term in equation (9), which has a sample average value of 0.83. Estimates of the increase in reservation wages when UI benefits increase range from zero (Krueger and Mueller, 2013) to as large as 0.42 (Feldstein and Poterba, 1984).

[22] One could reconcile this calibration with our estimated consumption ratio by assuming that preferences are separable but asset markets are incomplete. We return to this possibility in Section 8.

series of consumptions we have assumed that the relative consumption of the unemployed is smooth over the business cycle. In Appendix F we show that our results are qualitatively robust to alternative assumptions about the cyclicality of the consumption ratio. Second, we have assumed that Q_t is acyclical. Recall that Q_t has the interpretation of the value (net of utility costs) an unemployed derives from producing goods in the home sector relative to an employed. A countercyclical Q_t would ameliorate the procyclicality of z_t. However, in Appendix G we use data from the American Time Use Survey to show that hours per unemployed on home production actually fall when unemployment rises. Therefore, a constant Q is conservative for our results.

6 Implications for Unemployment Fluctuations

We now discuss the implications of the cyclicality of z for unemployment fluctuations. We begin with the case in which wages are set according to Nash bargaining. Section 6.2 considers the alternating-offer bargaining model of Hall and Milgrom (2008).

6.1 Implications Under Nash Bargaining

We develop our analysis by following much of the literature in treating steady state movements in the marginal product of employment p^e and the opportunity cost of employment z as exogenous.[23] Differentiating equation (32) with respect to p^e, recognizing that f is a function of θ, and holding constant κ, μ, β, s, and M gives an expression for the elasticity of labor market tightness θ with respect to p^e shocks:

$$\epsilon(\hat{\theta}, \hat{p}^e) = \left[\frac{\mu f + \frac{1-\beta(1-s)}{\beta}}{\mu f + (1-\eta)\frac{1-\beta(1-s)}{\beta}} \right] \left(\frac{p^e - z\epsilon(\hat{z}, \hat{p}^e)}{p^e - z} \right). \tag{34}$$

The response of unemployment is then given by $\epsilon(\hat{u}, \hat{p}^e) = -\eta(1-u)\epsilon(\hat{\theta}, \hat{p}^e)$. Equation (34) generalizes the expressions given in Shimer (2005), Mortensen and Nagypal (2007), and Hagedorn and Manovskii (2008) to allow z to change in response to changes in p^e. The magnitude of this response is given by $\epsilon(\hat{z}, \hat{p}^e)$.

[23] In Appendix H we simulate the model and confirm our results when p^e and z endogenously change over time.

Table 4: Steady State Elasticity of Unemployment With Respect to the Marginal Product

	$z = 0.452$	$z = 0.710$	$z = 0.955$
$\epsilon(\hat{z}, \hat{p}^e) = 0.00$	-0.72	-1.36	-8.76
$\epsilon(\hat{z}, \hat{p}^e) = 0.25$	-0.64	-1.12	-6.67
$\epsilon(\hat{z}, \hat{p}^e) = 0.50$	-0.56	-0.88	-4.58
$\epsilon(\hat{z}, \hat{p}^e) = 0.75$	-0.48	-0.64	-2.49
$\epsilon(\hat{z}, \hat{p}^e) = 1.00$	-0.39	-0.39	-0.39
$\epsilon(\hat{z}, \hat{p}^e) = 1.25$	-0.31	-0.15	1.70

Table 4 presents the elasticity of unemployment with respect to the marginal product of employment $\epsilon(\hat{u}, \hat{p}^e)$ as a function of the level of the opportunity cost z and the cyclicality of the opportunity cost $\epsilon(\hat{z}, \hat{p}^e)$. Recall from Table 3 that $\epsilon(\hat{u}, \hat{p}^e)$ is between -7 and -7.5 in the data. In the first row of the table z is constant. The response of unemployment to shocks in the marginal product is small when the calibrated value of z is small, consistent with the result in Shimer (2005). Moving across columns to higher levels of z, the response of unemployment increases. As pointed out by Hagedorn and Manovskii (2008), a higher z reduces firm's steady state profits. An increase in the productivity of a match then causes a larger percent increase in profits, increasing the incentive to create vacancies and making unemployment more volatile.

A key result of our analysis can be seen by moving down the rows of Table 4, as we allow z to vary cyclically. A positive value of $\epsilon(\hat{z}, \hat{p}^e)$ means that in response to p^e shocks, z increases. The higher is the responsiveness of z, the smaller is the increase in the net flow surplus of the match, $p^e - z$, and the weaker is the firm's incentive to create vacancies. As a result, holding constant the level of z, the response of unemployment becomes smaller when $\epsilon(\hat{z}, \hat{p}^e)$ is higher.

Equation (34) shows analytically that under our preferred estimate of $\epsilon(\hat{z}, \hat{p}^e) = 1$, so that both z and p^e change by the same percent, the elasticity of θ and u with respect to the marginal product is independent of the level of z.[24] Table 4 shows that when $\epsilon(\hat{z}, \hat{p}^e) = 0.75$, the elasticity

[24] In a model with labor market frictions and concave utility, Blanchard and Gali (2010) show that unemployment is neutral with respect to fluctuations in productivity. In their model consumption moves one-to-one with productivity due to the lack of capital. Given that they do not consider benefits and fluctuations in hours per

of unemployment with respect to the marginal product is 66 percent of the elasticity obtained under a constant $z = 0.452$, 47 percent of the elasticity obtained under a constant $z = 0.710$, or 28 percent of the elasticity obtained under a constant $z = 0.955$. When $\epsilon(\hat{z}, \hat{p}^e) > 1$ and z is relatively high, it is even possible that the sign of the response of unemployment changes.

The role of cyclical movements of z in unemployment fluctuations with Nash bargaining is quite general. In particular, the numerator of the second term in equation (34) changes from $p^e - z\epsilon(\hat{z}, \hat{p}^e)$ to $p^e\epsilon(\hat{p}^e, \hat{x}) - z\epsilon(\hat{z}, \hat{x})$ for any shock x other than shocks to κ, μ, β, s, and M. The crucial determinant of unemployment volatility is the responsiveness of z relative to the responsiveness of p^e when some shock x hits the economy.

6.2 Implications Under Alternating Offers Bargaining

Hall and Milgrom (2008) replace Nash bargaining with an alternative wage setting mechanism. In their alternating-offer bargaining game, when a firm with a vacancy meets an unemployed worker, the firm offers a compensation package $\tilde{w}$. The worker can accept the offer and commence work, or prolong the bargaining and make a counteroffer $\tilde{w}'$. Crucially, z parameterizes the flow opportunity cost to the worker of prolonging the bargaining, and hence the threat point if the worker deems the employer's initial offer too low.[25] With a constant z, wages therefore respond weakly to increases in p^e. The rigidity of wages incentivizes firms to significantly increase their job creation.[26] Allowing instead z to comove with p^e in the alternating-offer bargaining model makes the unemployed's threat point again sensitive to aggregate conditions. This increases the flexibility of wages and reduces the volatility of unemployment.

We illustrate this point using the linear search and matching model presented in Hall and Milgrom (2008). We first replicate their results for three linear models, the Nash bar-

worker, their z also moves one-to-one with productivity. In our case unemployment fluctuations are small but not zero when $\epsilon(\hat{z}, \hat{p}^e) = 1$. The difference is explained by the fact that in equation (34) we have assumed that vacancy costs κ are constant. Blanchard and Gali instead assume that hiring costs scale one-to-one with productivity.

[25] Another important parameter is the probability (denoted by δ in Hall and Milgrom, 2008) that the bargaining exogenously falls apart and the worker returns to the general search pool. This probability governs the extent to which the wage depends on z rather than on wage offers at other firms.

[26] Christiano, Eichenbaum, and Trabandt (2013) embed the Hall and Milgrom (2008) model of wage bargaining into a New-Keynesian dynamic general equilibrium model and show that the estimated model outperforms the standard MP model in several dimensions including volatility in the labor market.

Table 5: Cyclicality of Wages and Unemployment Fluctuations

Statistic	Standard MP		Hagedorn-Manovskii		Hall-Milgrom	
	Fixed z	Cyclical z	Fixed z	Cyclical z	Fixed z	Cyclical z
Slope $d\tilde{U}^u/d\tilde{p}^e$	1.14	1.30	0.87	1.32	1.19	1.31
Slope $d\tilde{w}/d\tilde{p}^e$	0.93	0.98	0.71	0.98	0.69	0.91
Elasticity $\epsilon(\hat{u}, \hat{p}^e)$	-1.51	-0.44	-5.87	-0.40	-6.02	-1.75

gaining model with $z = 0.71$ ("Standard MP"), the Nash bargaining model with $z = 0.93$ ("Hagedorn-Manovskii"), and the alternating-offer bargaining model with $z = 0.71$ ("Hall-Milgrom"). Then, we introduce in these models a cyclical z with $\epsilon(\hat{z}, \hat{p}^e) = 1$. Appendix I presents the equations and parameters of Hall and Milgrom (2008), which we adopt here for our analysis.

Table 5 summarizes our results. We first discuss results under Nash bargaining, building on the intuition of the previous section. The first row shows the slope of the expected present value of utility flows for the unemployed $\tilde{U}^u$ with respect to the expected present value of a newly hired worker's product $\tilde{p}^e$. With Nash bargaining, $\tilde{U}^u$ is the outside option of the unemployed while bargaining. It helps to separate $\tilde{U}^u$ into the sum of two components, the expected present value from receiving z discounted by the probability the worker remains unemployed, and the value of obtaining a job in a future period discounted by the probability of exiting unemployment in that period. In the Standard MP model with constant z, $\tilde{U}^u$ responds substantially when $\tilde{p}^e$ increases. Intuitively, low z means that future job prospects contribute relatively more to $\tilde{U}^u$, and higher $\tilde{p}^e$ increases the probability of an unemployed finding a high-wage job. In the Hagedorn-Manovskii calibration, a high fixed z makes the expected discounted value of future z's a more important component of $\tilde{U}^u$. As a result, total $\tilde{U}^u$ responds less to the better job prospects created by higher $\tilde{p}^e$.

The second row shows the slope of the expected present value of wage payments $\tilde{w}$ with

respect to $\tilde{p}^e$. With constant z, the increase in the worker's outside option in the standard MP model makes wages respond flexibly to productivity as well. In the Hagedorn-Manovskii model, the insensitivity of the outside option to movements in productivity makes the wage more rigid. This difference explains the success of the Hagedorn-Manovskii calibration in generating volatile unemployment fluctuations, shown in the third row of Table 5.

Turning to the Hall and Milgrom model with constant z, here too the change in job prospects of an unemployed makes $\tilde{U}^u$ sensitive to variations in $\tilde{p}^e$. However, with alternating-offer bargaining, returning to the general search pool with value $\tilde{U}^u$ no longer constitutes the worker's outside option. Instead, the unemployed's threat point is to continue to bargain, in which case he receives a flow value z. Therefore, wages do not respond significantly to productivity variations, and the volatility of unemployment increases.

Summing up, both the Hagedorn and Manovskii (2008) calibration and the Hall and Milgrom (2008) alternating-offers model achieve volatile unemployment in part by generating endogenous wage rigidity. In both cases, the wage rigidity comes from increasing the importance of z to the worker's outside option, in Hagedorn and Manovskii (2008) by calibrating a higher z, and in Hall and Milgrom (2008) by changing the bargaining game to increase the weight of z in the outside option. This logic makes clear why both models no longer generate volatile unemployment if z moves cyclically. In that event, the outside option in both models again becomes sensitive to productivity, wages become volatile, and the firm's incentive to increase employment following a positive shock to $\tilde{p}^e$ becomes weaker. The columns labeled Cyclical z in Table 5 illustrate this point quantitatively.[27]

7 Heterogeneity Across Skills

Our measurement of z for the average unemployed followed from the assumption that all unemployed search for the same jobs and employers cannot discriminate ex-ante in choosing a

[27] With cyclicality in z, the Hall-Milgrom model performs better than the Hagedorn-Manovskii model. This is because in the Hall-Milgrom model wages partly depend on a firm-specific cost of continuing bargaining (denoted by γ) which is assumed to be constant over time. Making γ comove with the aggregate state ameliorates even more the unemployment fluctuations generated by the model.

potential worker with whom to bargain. We now relax this assumption and allow workers to differ along observable characteristics that may be correlated with their opportunity costs.

The economy consists of J heterogeneous households. In the data we separate workers into four educational attainment categories. Each household j contains fraction l_j of the population. Within each group j, a fraction e_{jt} is employed and a fraction u_{jt} is unemployed. There are J separate labor markets. In each labor market, firms post vacancies v_{jt} at a cost of κ_{jt} per vacancy, and a matching function converts search by the unemployed and vacancies into matches $m_{jt} = m_j(v_{jt}, l_{jt}u_{jt})$. We denote by $f_{jt} = m_{jt}/\left(l_{jt}u_{jt}\right)$ the job finding rate in market j, by $q_{jt} = m_{jt}/v_{jt}$ the vacancy filling rate, and by $\theta_{jt} = v_{jt}/\left(l_{jt}u_{jt}\right)$ the market tightness. Given an exogenous separation rate s_{jt}, employment for group j evolves as $e_{jt+1} = f_{jt}+(1-s_{jt}-f_{jt})e_{jt}$.

The problem of each household j is:

$$W_j^h = \max E_0 \sum_{t=0}^{\infty} \beta^t \left[l_j e_{jt} U_j^e(C_{jt}^e, N_{jt}) + l_j(1-e_{jt}) U_j^u(C_{jt}^u, 0) - l_j(1-e_{jt})\omega_{jt}\psi(\zeta_{jt}) \right], \quad (35)$$

subject to the budget constraint:

$$l_j e_{jt} C_{jt}^e + l_j(1-e_{jt}) C_{jt}^u + C_{jt}^o + K_{jt+1} + T_{jt} = w_{jt} l_j e_{jt} N_{jt} + l_j(1-e_{jt}) B_{jt} + (1+R_t-\delta) K_{jt} + \Pi_{jt}, \quad (36)$$

and the law of motion for eligibility:

$$\omega_{jt+1} = \left(\omega_{jt+1}^u (1-f_{jt}) \frac{u_{jt}}{u_{jt+1}} \right) \omega_{jt} + \omega_{jt+1}^e s_{jt} \frac{e_{jt}}{u_{jt+1}}. \quad (37)$$

We note that flow utilities U_j^e and U_j^u are allowed to vary by j. In the budget constraint, benefits per unemployed of type j are given by $B_{jt} = B_{n,jt} + B_{u,jt}$, where $B_{n,jt}$ denote non-UI benefits per unemployed in group j and $B_{u,jt}$ denote UI benefits per unemployed in group j.

To derive the opportunity cost of employment by group j, we proceed analogously to the aggregate case analyzed in Section 2. We first derive the marginal value of employment for household j, $J_{jt}^h = \partial W_{jt}^h / \partial \left(l_{jt} e_{jt}\right)$, as the sum of a flow payoff (wages minus opportunity cost) and a continuation value. Then, we define the opportunity cost of employment similarly to the aggregate case:

$$z_{jt} = b_{jt} + (C_{jt}^e - C_{jt}^u) - \frac{U_{jt}^e - U_{jt}^u}{\lambda_{jt}} = b_{jt} + \xi_{jt}, \quad (38)$$

Table 6: Heterogeneity Sample Statistics: 1969(1)–2012(4)

	Less Than High School	High School	Some College	College Or More
$B_{n,j}$	0.020	0.019	0.013	0.006
$B_{u,j}$	0.040	0.081	0.087	0.102
γ_j^u	0.767	0.804	0.786	0.798
C_j^e	0.436	0.502	0.567	0.679
N_j	0.902	1.005	0.989	1.071
s_j	0.095	0.046	0.037	0.018
f_j	0.702	0.661	0.685	0.634
p_j^e	0.661	0.909	1.034	1.569
ρ_j	1.599	1.358	1.399	1.270
χ_j	1.587	1.637	1.703	1.929
$-(1-\rho_j)\chi_j$	0.950	0.585	0.680	0.521
z_j when $Q_j = 0$	0.303	0.421	0.469	0.684
corr $\left(\hat{z}_j, \hat{Y}\right)$ when $Q_j = 0$	0.732	0.585	0.514	0.294
corr $\left(\hat{z}_j, \hat{Y}\right)$ when $z_j/p_j^e = 0.710$	0.714	0.593	0.531	0.422
corr $\left(\hat{z}_j, \hat{Y}\right)$ when $z_j/p_j^e = 0.955$	0.662	0.581	0.523	0.474

Notes: $B_{n,j}, B_{u,j}, C_j^e, p_j^e$, and z_j are expressed as a fraction of the mean aggregate marginal product of employment in the sample. N_j is expressed as a fraction of the mean aggregate hours per worker in the sample.

where b_{jt} is given by equation (9) taking into account j-specific values of variables. The thrust of our procedure for constructing z_{jt} follows that for the aggregate described in Section 3. Here we sketch briefly our estimation and refer the reader to Appendix J for more details.

For our estimates of the benefits per unemployed $B_{n,jt}$ and $B_{u,jt}$, we use the March CPS to measure the fraction of survey dollars in each program accruing to the unemployed of category j and the CPS basic monthly files to measure the fraction of unemployed belonging to group j. The first two rows of Table 6 report the sample averages of $B_{n,j}$ and $B_{u,j}$, expressed relative to the mean aggregate marginal product of employment in the sample. The opportunity cost of low skilled workers contains higher non-UI benefits than that of high skilled workers. This difference reflects the asset and income tests for non-UI benefits, which disqualify many high

skill workers. By contrast, UI benefits per unemployed increase monotonically with skill level. The average benefit doubles for workers with a high school diploma relative to those without, and increases another 25 percent for those with a college degree. The statutory linking of UI benefit receipt to previous wages explains the positive relationship between skill level and UI benefits per unemployed.

For our estimates of consumptions of employed C^e_{jt} and unemployed C^u_{jt}, we use the CE to measure the consumption ratios upon unemployment by group γ^u_j and the relative consumptions of employed of different skills $\gamma^e_j = C^e_j/C^e_i$. Applying these consumption ratios to a modified version of the adding-up equation (23), we obtain times series for consumptions. Table 6 shows the consumption ratios upon unemployment γ^u_j to be quite stable across different skill groups. It also shows large differences across groups in the consumption of the employed, with C^e_{jt} increasing monotonically with skill.

We measure hours per worker N_j, the separation rate s_j, and the job finding rate f_j in an analogous way to the aggregate. N_j and f_j appear relatively stable across groups, whereas the mean separation rate s_j declines sharply with skill level. Table 6 also reports the mean marginal product of each group p^e_j. We construct p^e_j using a CES aggregator of the J different labor inputs and calibrate parameters such that, in the long run, the ratio of marginal products across groups equals the ratio of labor earnings. Given estimates of C^e_j, C^u_j, N_j, and p^e_j, we calibrate the group-specific parameters ρ_j and χ_j in the utility functions (30) and (31) using the risk sharing condition (4) and the first-order condition for hours (16) for each group j.

Figure 5 plots the cyclical components of z_{jt}, for the case in which the utility shifter Q_j is set such that $z_j/p^e_j = 0.71$ for all j. The z_j's are highly synchronized across groups. Table 6 reports the correlation of the cyclical component of z_j with the cyclical component of output for the three cases of $Q_j = 0$, Q_j such that $z_j/p^e_j = 0.71$, and Q_j such that $z_j/p^e_j = 0.955$. The z's of lower skill groups correlate more strongly with output. This result partly reflects the fact that countercyclical UI benefits are a smaller fraction and procyclical non-UI benefits are a larger fraction of their b_{jt}. The ξ_{jt} component of the opportunity cost behaves quite similarly

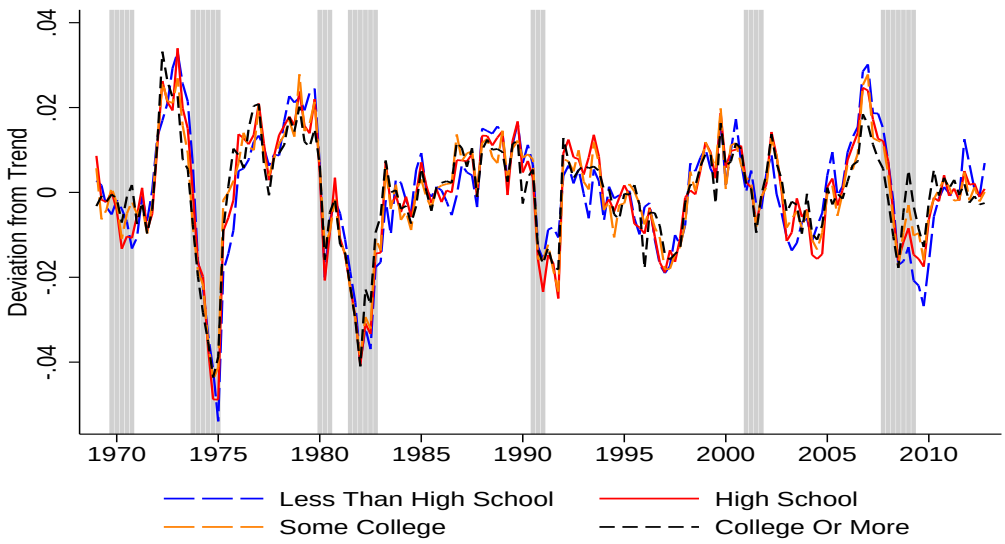

Figure 5: Opportunity Costs Across Skills

across groups, with the highest skilled group having a somewhat less cyclical ξ_{jt} due to a lower cyclicality of their N_{jt}.

To summarize, the main message of our analysis is that the procyclicality of the opportunity cost is also found across heterogeneous educational groups. While there are interesting compositional differences across groups, the same economic forces that drive the aggregate z_t to fluctuate over the business cycle also influence the skill-specific z_{jt}'s.

8 Risk Sharing versus Self Insurance

The construction of our aggregate z assumed that idiosyncratic employment risk is perfectly shared across agents. This follows from the assumption of a single budget constraint with a household-level asset (the capital stock), from which the marginal utilities of consumption of the employed and unemployed must equalize in all states of the world. We now show that the economic intuition for a cyclical z is quite robust to allowing for imperfections in risk sharing. With full risk sharing, the opportunity cost of employment changes in response to uninsurable aggregate risk that causes changes in the (common) marginal utility of consumption. When employed and unemployed cannot share their idiosyncratic risks, the existence of aggregate risk still leads to fluctuations of the (different) marginal utilities that generate a procyclical

opportunity cost.

To illustrate this result, we focus on a stripped down version of our model. The most prominent simplification is that we look at the problem of a household from a partial equilibrium perspective. We denote by x_t the exogenous aggregate state of the economy and by $\pi(x_{t+1}|x_t)$ the probability that the exogenous state transits from some x_t to some x_{t+1}. In the model with perfect risk sharing, the household's value function is:

$$V(x_t, a_t) = e_t U^e(C_t^e, N_t) + (1 - e_t)U^u(C_t^u, 0) + \beta \sum_{x_{t+1}} \pi(x_{t+1}|x_t)V(x_{t+1}, a_{t+1}), \tag{39}$$

subject to the single budget constraint $e_t C_t^e + (1-e_t)C_t^u + a_{t+1} = e_t w_t N_t + (1-e_t)b_t + (1+r_t)a_t + Y_t^o$ and the borrowing constraint $a_{t+1} \geq \bar{a}$. As before, the opportunity cost of employment is $z_t^R = b_t + (U_t^u/\lambda_t - C_t^u) - (U_t^e/\lambda_t - C_t^e)$, where λ_t is the common marginal utility of consumption, and the R superscript stands for "risk sharing."

Now consider the problem of a worker who cannot share risks perfectly with other members of the household, but instead accumulates assets to self insure against idiosyncratic employment shocks. The value function of an employed worker starting with assets a_t is:

$$V^e(x_t, a_t) = U^e(C_t^e, N_t) + \beta \sum_{x_{t+1}} \pi(x_{t+1}|x_t) \left((1 - s_t)V^e(x_{t+1}, a_{t+1}^e) + s_t V^u(x_{t+1}, a_{t+1}^e)\right), \tag{40}$$

subject to the budget constraint $C_t^e + a_{t+1}^e = w_t N_t + (1+r_t)a_t + Y_t^o$ and the borrowing constraint $a_{t+1}^e \geq \bar{a}$. The value function of an unemployed worker starting with assets a_t is:

$$V^u(x_t, a_t) = U^u(C_t^u, 0) + \beta \sum_{x_{t+1}} \pi(x_{t+1}|x_t) \left((1 - f_t)V^u(x_{t+1}, a_{t+1}^u) + f_t V^e(x_{t+1}, a_{t+1}^u)\right), \tag{41}$$

subject to the budget constraint $C_t^u + a_{t+1}^u = b_t + (1 + r_t)a_t + Y_t^o$ and the borrowing constraint $a_{t+1}^u \geq \bar{a}$. We denote by Y_t^o all "other income," which includes resources such as spousal income, other intrafamily transfers, and transfers or taxes from the government that do not depend on employment status.

We derive the flow opportunity cost of employment as the wage payment that leaves the worker indifferent between accepting a job in period t and not, holding the effect on employment status in $t+1$ constant. A worker entering period t with assets a_t receives surplus from moving

from unemployment to employment of $J_t^S = V^e(x_t, a_t) - V^u(x_t, a_t)$, where S stands for "self insurance." We define $J_{t+1}^S = V^e(x_{t+1}, a_{t+1}^e) - V^u(x_{t+1}, a_{t+1}^e)$. Evaluating both terms of J_{t+1}^S at a_{t+1}^e restricts the $t+1$ surplus to only that part associated with entering $t+1$ in the employed state. Substituting (40) and (41) into J_t^S, we obtain:

$$\frac{J_t^S}{\lambda_t^e} = w_t N_t - z_t^S + (1 - s_t - f_t) E_t \frac{\beta \lambda_{t+1}^e}{\lambda_t^e} \frac{J_{t+1}^S}{\lambda_{t+1}^e}, \tag{42}$$

where

$$z_t^S = b_t + \left(\frac{U_t^u}{\lambda_t^e} - C_t^u \right) - \left(\frac{U_t^e}{\lambda_t^e} - C_t^e \right) + z_t^A, \tag{43}$$

and where z_t^A denotes a component of the opportunity cost related to the differential asset accumulation between the employed and the unemployed.[28] We divide by the marginal utility of the employed λ_t^e because the wage negotiated during bargaining is paid in the state of the world in which the worker accepts the offer.

In the aggregate analysis we attributed the entirety of the 21 percent decline in consumption upon unemployment to non-separabilities between consumption and hours. Here we make the opposite extreme assumption that the consumption decline results only from market incompleteness, and so we use separable preferences ($\rho = 1$). We pick the level of "other income" Y_t^o such that, in the space of assets considered below, consumption drops on average by 21 percent upon unemployment. The exogenous state vector x_t includes the employment rate e_t, the separation rate s_t, the wage w_t, hours per employed worker N_t, and benefits per unemployed b_t. We define an aggregate "good state" in which the shocks e_t, w_t, and N_t are one standard deviation above their trend and s_t and b_t are one standard deviation below their trend, and we define an aggregate "bad state" symmetrically.[29]

[28] This term is $z_t^A = -\frac{\beta}{\lambda_t^e} E_t \left[f_t \left(V^e(x_{t+1}, a_{t+1}^e) - V^e(x_{t+1}, a_{t+1}^u) \right) + (1 - f_t) \left(V^u(x_{t+1}, a_{t+1}^e) - V^u(x_{t+1}, a_{t+1}^u) \right) \right] + a_{t+1}^e - a_{t+1}^u$. Moving from unemployment to employment (holding constant initial assets at a_t) causes a "budgetary loss" equal to $a_{t+1}^e - a_{t+1}^u$ due to the fact that employed accumulate more assets. There is an offsetting gain as the worker starts $t+1$ with higher assets. Because all of the surplus associated with a higher probability of having a job in $t+1$ is included into J_{t+1}^S, the value function gains from entering $t+1$ with assets a_{t+1}^e instead of a_{t+1}^u are evaluated as if the worker obtains employment in period $t+1$ with probability f_t.

[29] This procedure yields $Y^o = 1.15$. We calibrate $\chi = 0.71$ to target an opportunity cost $z^R = 0.710$ in the space of assets shown in Figure 6. We set the borrowing constraint to $\bar{a} = -0.5$, which corresponds to a fraction 23% of the total non-capital income of the employed $wN + Y^o$. The rest of the parameters are $w = 1$, $N = 1$, $b = 0.041$, $\beta = 0.98$, $r = 0.01$, $\epsilon = 0.7$, $Q = 0$, $s = 0.045$, and $f = 0.704$. Finally, we discretize the state in three values and assume that the transition matrix is given by $\pi(x_{t+1} = x_j | x_t = x_i) = 0.98$ for $j = i$ and $\pi(x_{t+1} = x_j | x_t = x_i) = 0.01$ for $j \neq i$.

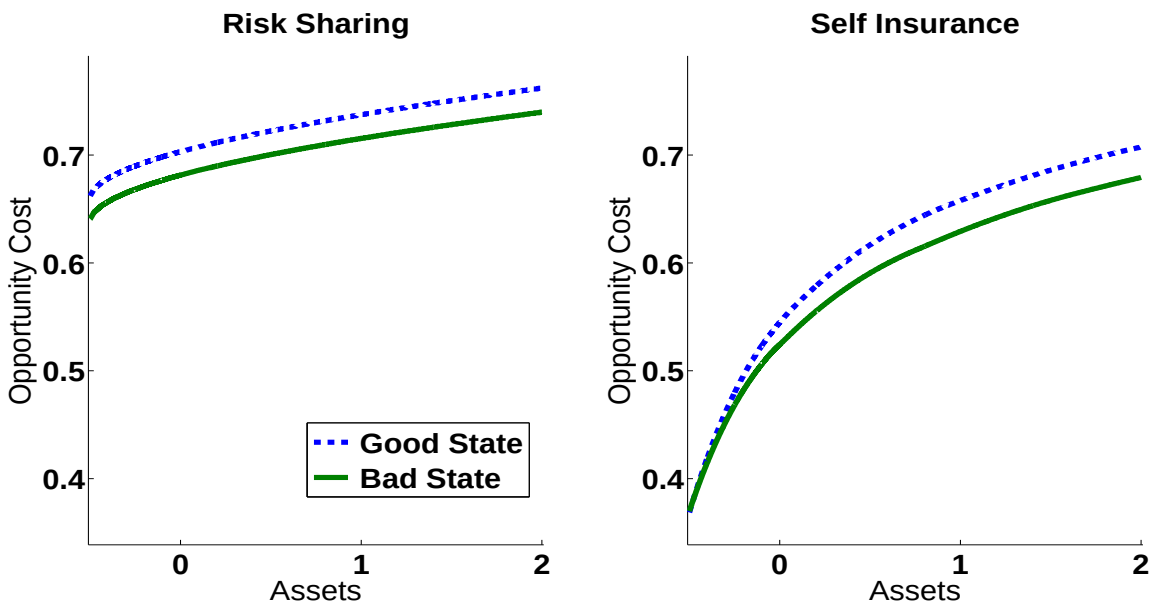

Figure 6: Opportunity Cost Under Alternative Risk Sharing Arrangements

Notes: The left panel plots the opportunity cost of employment z_t^R for the model with perfect risk sharing in a good state (dashed line) and in a bad state (solid line). The right panel plots the opportunity cost of employment z_t^S for the model with self insurance.

Figure 6 plots the opportunity costs of employment, z_t^R and z_t^S, in the models with perfect risk sharing and self insurance for different starting assets and aggregate states. Several results are worth highlighting. First, the opportunity cost is in general lower in the model with self insurance. With imperfect risk sharing, workers save more to insure against idiosyncratic shocks. Lower consumption for a given level of assets means that the marginal utility of consumption for both the unemployed and the employed is higher relative to the model with perfect risk sharing. As initial assets increase, the probability of hitting the borrowing constraint becomes smaller, and the opportunity cost in the model with self insurance approaches the level of the opportunity cost in the model with perfect risk sharing.

Second, in both models the opportunity cost increases in the level of assets. The increase is much sharper for workers close to the borrowing constraint in the model with self insurance. These workers have a very high marginal utility of consumption, making them more desperate to obtain employment. The positive relationship between assets and the opportunity cost implies procyclical movements of the opportunity cost in both models if in recessions average wealth of the unemployed declines.

Third, in both models the opportunity cost is in general procyclical for a given level of assets, as evidenced by the upward shift of the dashed line relative to the solid line. Just as with complete markets, with incomplete markets the marginal utility of consumption falls relative to the value of non-working time in the good state. For very low levels of initial assets, however, the opportunity cost becomes less cyclical. Intuitively, consumption is relatively insensitive to aggregate shocks when workers are close to their borrowing constraint.[30]

9 Conclusion

This paper has shown that the flow value of the opportunity cost of employment falls during recessions. The key mechanism is that the household values most the contribution of the employed (through higher wage income) relative to that of the unemployed (through higher non-working time) when market consumption is low and non-working time is high. This more than offsets the effect of the increase in government benefits.

A procyclical opportunity cost reduces unemployment volatility in models where z affects the wage bargain. Our preferred estimate of the elasticity of the opportunity cost with respect to the marginal product of employment is unity. With this value and Nash bargaining, fluctuations in unemployment generated by the model are essentially neutral with respect to the level of z, and remain far smaller than unemployment fluctuations in the data. We reach a similar outcome in a model in which wages are determined by alternating offers.

An interpretation of our results is that endogenous forms of wage rigidity, such as accomplished by Hagedorn and Manovskii (2008) and by Hall and Milgrom (2008), do not survive the introduction of a cyclical flow opportunity cost. Without rigid wages, these models cannot generate volatile unemployment. This pessimistic conclusion does not apply to models where wages are exogenously sticky or selected according to some process that does not depend on worker's opportunity cost of employment. Alternatively, using the Brugemann and Moscarini

[30] Nakajima (2012) develops a model with incomplete markets, leisure, and a benefit replacement rate of 64% that generates high volatility in unemployment. He argues that changes in borrowing constraints do not matter much for the performance of search and matching models as workers save and self-insure sufficiently to overcome these constraints. We hypothesize that the opportunity cost generated by his model is less cyclical than what we estimate because benefits constitute two-thirds of the opportunity cost.

(2010) decomposition of wages into payments covering opportunity costs and rents due to frictions, the procyclicality of z implies that wage rigidity requires substantial countercyclicality in rents. The extent to which actual wages vary cyclically remains an open and important question (see Pissarides (2009) and Hall and Milgrom (2008) for contrasting views).

Our results also bear on recent work emphasizing the role of social safety net expansions in propagating the increase in unemployment during the Great Recession (Hagedorn, Karahan, Manovskii, and Mitman, 2013). We find, contrary to this hypothesis, that fluctuations in the value of benefits have only a small effect on the opportunity cost. However, we have not modeled the complicated set of benefit phase-out schedules that give rise to high implicit marginal tax rates at the low end of the income distribution as in Mulligan (2012). The same economic reasoning that makes the employment margin choice sensitive to the marginal utility of consumption relative to the value of non-working time would also affect the choice of hours along the intensive margin and the movements between in and out the labor force.

References

Aguiar, M., and E. Hurst (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5), 919–48.

Anderson, P., and B. Meyer (1997): "Unemployment Insurance Takeup Rates and the After-Tax Value of Benefits," *Quarterly Journal of Economics*, 112(3), 913–37.

Andolfatto, D. (1996): "Business Cycles and Labor-Market Search," *American Economic Review*, 86(1), 112–32.

Bils, M., Y. Chang, and S.-B. Kim (2012): "Comparative Advantage and Unemployment," *Journal of Monetary Economics*, 59, 150–65.

Blanchard, O., and J. Gali (2010): "Labor Markets and Monetary Policy: A New Keynesian Model with Unemployment," *American Economic Journal: Macroeconomics*, 2(2), 1–30.

Blank, R., and D. Card (1991): "Recent Trends in Insured and Uninsured Unemployment: Is There an Explanation?," *Quarterly Journal of Economics*, 106(4), 1157–89.

Brugemann, B., and G. Moscarini (2010): "Rent Rigidity, Asymmetric Information, and Volatility Bounds in Labor Markets," *Review of Economic Dynamics*, 13(3), 575–96.

Centers for Medicare and Medicaid Services (2011): "Data Compendium 2011 Edition," Online Data.

Christiano, L., M. Eichenbaum, and M. Trabandt (2013): "Unemployment and Business Cycles," NBER Working Paper No. 19265.

Cociuba, S., E. Prescott, and A. Ueberfeldt (2012): "U.S. Hours and Productivity Behavior Using CPS Hours," Arizona State University Working Paper.

Costain, J., and M. Reiter (2008): "Business Cycles, Unemployment Insurance, and the Calibration of Matching Models," *Journal of Economic Dynamics and Control*, 32, 1120–55.

den Haan, W., G. Ramey, and J. Watson (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90(3), 482–98.

Feldstein, M., and J. Poterba (1984): "Unemployment Insurance and Reservation Wages," *Journal of Public Economics*, 23(12), 141–67.

Fernald, J. (2012): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Federal Reserve Bank of San Francisco Working Paper No. 2012-19.

Hagedorn, M., F. Karahan, I. Manovskii, and K. Mitman (2013): "Unemployment Benefits and Unemployment in the Great Recession: The Role of Macro Effects," NBER Working Paper No. 19499.

Hagedorn, M., and I. Manovskii (2008): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 98(4), 1692–706.

Hall, R. (2005): "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95(1), 50–65.

——— (2009): "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor," *Journal of Political Economy*, 117(2), 281–323.

Hall, R., and P. Milgrom (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98(4), 1653–74.

Hall, R., and A. Mueller (2013): "Wage Dispersion and Search Behavior," Stanford University Working Paper.

Hornstein, A., P. Krusell, and G. Violante (2005): "Unemployment and Vacancy Fluctuations in the Matching Model: Inspecting the Mechanism," *Federal Reserve Bank of Richmond Economic Quarterly*, 91(3), 19–51.

Kaiser Family Foundation (2013): "Where Are States Today?," Online Data.

Krueger, A., and A. Mueller (2013): "A Contribution to the Empirics of Reservation Wages," Columbia University Working Paper.

Merz, M. (1995): "Search in the Labor Market and the Real Business Cycle," *Journal of Monetary Economics*, 36(2), 269–300.

Meyer, B., W. Mok, and J. Sullivan (2009): "The Under-Reporting of Transfers in Household Surveys: It's Nature and Consequences," NBER Working Paper No. 15181.

Mortensen, D., and C. Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61(3), 397–415.

——— (1999): "Unemployment Responses to 'Skill-Biased' Technology Shocks: The Role of Labour Market Policy," *Economic Journal*, 109(455), 242–65.

——— (2001): "Taxes, Subsidies and Equilibrium Labor Market Outcomes," CEP Discussion Papers No. 519.

Mortensen, D. T., and E. Nagypal (2007): "More on Unemployment and Vacancy Fluctuations," *Review of Economic Dynamics*, 10(3), 327–47.

Mulligan, C. (2012): *The Redistribution Recession*. Oxford University Press.

Nakajima, M. (2012): "Business Cycles in the Equilibrium Model of Labor Market Search and Self-Insurance," *International Economic Review*, 53(2), 399–432.

Pissarides, C. (1985): "Short-Run Equilibrium Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review*, 75(4), 676–90.

——— (2000): *Equilibrium Unemployment Theory*. MIT Press.

Pissarides, C. A. (2009): "The Unemployment Volatility Puzzle: Is Wage Stickiness the Answer?," *Econometrica*, 77(5), 1339–69.

Pistaferri, L. (2003): "Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply," *Journal of Labor Economics*, 21(3), 729–54.

Rogerson, R., and R. Wright (1988): "Involuntary Unemployment in Economies with Efficient Risk Sharing," *Journal of Monetary Economics*, 22, 501–15.

Shimer, R. (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1), pp. 25–49.

——— (2010): *Labor Markets and Business Cycles*. Princeton University Press.

——— (2012): "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15, 127–48.

Trabandt, M., and H. Uhlig (2011): "The Laffer Curve Revisited," *Journal of Monetary Economics*, 58(4), 305–27.