

Sources and Mechanisms of Cyclical Fluctuations in the Labor Market *

Robert E. Hall

Hoover Institution and Department of Economics

Stanford University

National Bureau of Economic Research

rehall@stanford.edu stanford.edu/~rehall

September 20, 2006

Abstract

I develop a model that accounts for the cyclical movements of hours and employment in the U.S. over the past 60 years. The model pays close attention to evidence about preferences for work and consumption. More than half of cyclical variations in total hours of work are in hours per worker. I show that reasonable volatility in the driving force and a reasonable elasticity of labor supply provide a believable account of the observed cyclical movements in hours per worker. The other part of cyclical variations in total hours comes from variations in the employment rate—hours decline in recessions in part because of the rise in unemployment. I define an employment function, analogous to the supply function for hours per worker, and estimate its parameters. My work differs from previous attempts to place cyclical movements of total hours on a labor supply curve by using a measure of the marginal product of labor that has much more transitory cyclical variation than previous measures based solely on productivity and by using a correspondingly lower elasticity of labor supply. It also differs by its explicit treatment of unemployment in a framework parallel to the supply of hours of work by the employed.

*This research is part of the program on Economic Fluctuations and Growth of the NBER. I am grateful to numerous participants in the Montreal Conference on Advances in Matching Models, the Rogerson-Shimer-Wright group at the 2006 NBER Summer Institute, and the 2006 Minnesota Workshop on Macroeconomic Theory, and to Felix Reichling for comments. A file containing the calculations will be available at Stanford.edu/~rehall

1 Introduction

I take up the challenge of accounting for volatility in the labor market, in hours per worker and in the employment rate, without contradicting the evidence about the elasticity of labor supply. Many contributions to the literature on aggregate labor-market volatility rest on explicit or implicit assumptions of unreasonably high elasticities of labor supply. The model of the paper describes labor supply in a broad sense, including unemployment. The model integrates labor supply and consumption demand. It takes labor demand as given, measured by the marginal product of labor. An important conclusion of the paper is that the marginal product of labor depends on a number of variables in addition to productivity and that the marginal product has correspondingly higher transitory cyclical variation.

Figure 1 shows detrended hours and employment for the United States since 1948. Somewhat more than half of the volatility of total hours of work comes from hours per worker and the remainder from workers per member of the labor force (the employment rate, one minus the unemployment rate). Virtually no cyclical fluctuations occur in the labor-force participation rate.

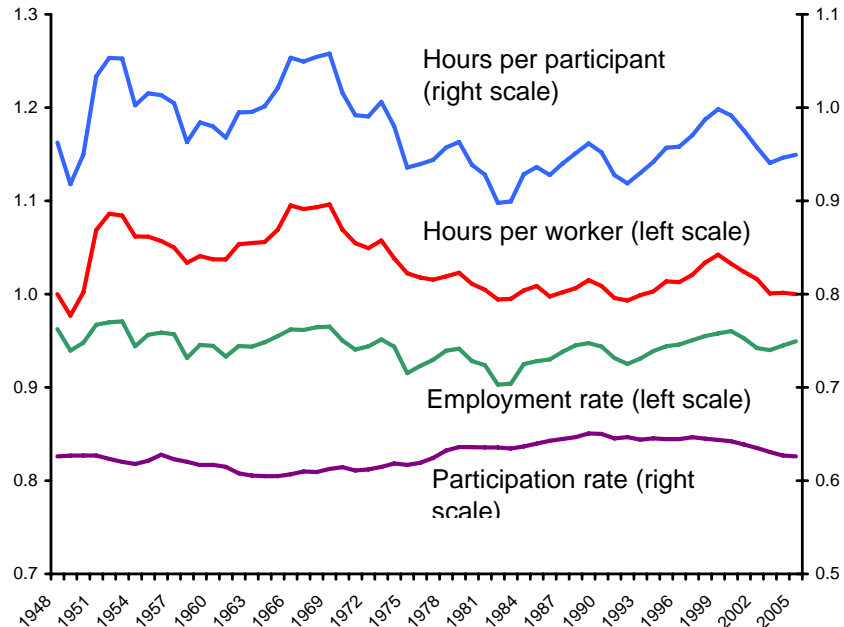


Figure 1. Hours per Participant, Hours per Worker, Employment Rate, and Participation Rate, 1948-2005

In broad summary, the model in this paper considers a worker who maximizes the discounted sum of future utility, which depends positively on the level of consumption and

negatively on the number of hours worked. The worker has an hourly marginal product w . I denote it w because it functions as the wage in the determination of the worker's hours. The worker has an intertemporal budget constraint—the economy offers an opportunity to save and earn a return. The Lagrange multiplier for the budget constraint, λ , is the marginal utility of consumption. λ describes the long-run or permanent level of well-being in the economy. The marginal product w captures the deviation of current conditions from normal. When w is higher than the level corresponding to current consumption, hours will be higher than normal as workers take advantage of the temporarily exceptional benefit of working.

Given λ and w , hours of work per worker, h , is a function $h(\lambda, w)$, the Frisch hours-supply function, expressing the level that equates the marginal disutility of work to λw . Hours supply is an increasing function of both λ and w . A companion function, the Frisch consumption demand function, describes the corresponding choice of consumption.

I argue that it is reasonable to portray the employment rate as a function $n(\lambda, w)$ of the same two variables. I show that this is true for the Mortensen and Pissarides (1994) (MP) model, the basic statement of the theory of unemployment widely in use today. The employment rate is also an increasing function of both λ and w .

The function

$$h(\lambda, w)n(\lambda, w). \tag{1}$$

governs the total volatility of hours of work. When the marginal product w rises temporarily above the level corresponding to λ , employment and hours rise, creating a cyclical bulge in total hours per person. Recessions are times when the opposite occurs.

For the supply function for hours per worker, $h(\lambda, w)$, I use explicit, calibrated preferences, with parameters drawn from recent research on household behavior. For the employment function, $n(\lambda, w)$, my approach is empirical. I infer λ and w from data on hours and consumption—I reverse engineer the Frisch hours-supply, consumption-demand system. Then I estimate slope parameters from the regression of the employment rate on the inferred times series for λ and w . The employment rate is sensitive to both variables in the expected positive direction. This finding is not surprising—it reflects the fairly tight correlation between hours per worker and unemployment.

Because the fit of the empirical employment function is quite good and because the hours supply function $h(\lambda, w)$ fits perfectly by construction, the supply function for total hours, $h(\lambda, w)n(\lambda, w)$, fits the actual data for hours well—the model reproduces the movements

shown in Figure 1.

The empirical employment function is much more sensitive to λ and w than is its counterpart in the MP model. I introduce a modification of the MP model with an additional parameter that shifts bargaining power along with the unemployment rate in a way that matches the slope coefficients of the empirical employment function.

A second phase of the empirical work asks if the marginal product w inferred from hours of work matches a measure constructed from the production side of the economy. I show that, with constant returns to scale and full adjustment of other factors of production, the marginal product of labor is value added less the rental cost of capital, divided by hours of work. Though data problems prevent a close comparison, the two time series share the same basic movements over the period since 1948.

I make a sharp distinction between compensation per worker and the marginal product of labor, w . Compensation is stable in relation to the marginal product. Expected compensation is the result of a bargain between worker and employer struck at the time the worker is hired. The model is unrestrictive about the evolution of compensation post-hire and is consistent with the large amount of compensation-smoothing visible in the data.

2 Insurance

The analysis in this paper makes best sense under the assumption that workers are fully insured against the personal risk of the labor market and that the insurance is actuarially fair. The insurance makes payments based on outcomes outside the control of the worker that keep all workers' marginal utility of consumption the same. This assumption—dating at least back to Merz (1995)—results in enormous analytical simplification. In particular, it makes the Frisch system of consumption demand and labor supply the ideal analytical framework.

I do not believe that, in the U.S. economy, consumption during unemployment behaves literally according to the model with full insurance. But families and friends may provide partial insurance. I view the fully insured case as a good and convenient approximation to the more complicated reality, where workers use savings and partial insurance to keep consumption close to the levels that would maintain roughly constant marginal utility. See Hall (2006) for evidence supporting the view that the fully insured case is a good approximation.

The approximation is good for the transient loss of income that occurs when a worker

goes through a spell of unemployment. It cannot possibly apply when a life event results in a permanent change in earnings. Such an event resets the worker’s marginal utility once and for all. Although I do not spell out the argument here, a richer model with a personal random-walk component of earnings would share many of the properties of the model in this paper.

3 Preferences

As in most research on choices over time, I assume that preferences are time-separable, though I am mindful of Browning, Deaton and Irish (1985)’s admonition that “the fact that additivity is an almost universal assumption in work on intertemporal choice does not suggest that it is innocuous.” In particular, additivity fails in the case of habit.

In this section, I will refer to w as the wage, because I am dealing with research based on the idea that w is visible to the worker as the slope of a personal budget constraint.

To discuss preferences in the Frisch framework, it is convenient to include a consumption price p that I will later normalize to one and omit. Thus I let $c(\lambda, p, w)$ be the Frisch consumption demand and $h(\lambda, p, w)$ be the Frisch labor supply. See Browning et al. (1985) for a complete discussion of Frisch systems in general. Given a utility kernel $U(c, h)$, the consumption demand and labor supply satisfy

$$U_c(c(\lambda, p, w), h(\lambda, p, w)) = \lambda p \tag{2}$$

and

$$U_h(c(\lambda, p, w), h(\lambda, p, w)) = -\lambda w \tag{3}$$

Frisch demands are symmetric: $\partial c/\partial w = -\partial h/\partial p$. They have three basic first-order or slope properties:

- *Intertemporal substitution in consumption*, $\partial c/\partial p$, the response of consumption to changes in its price
- *Frisch labor-supply response*, $\partial h/\partial w$, the response of hours to changes in the wage
- *Consumption-hours cross effect*, $\partial c/\partial w$, the response of consumption to changes in the wage (and the negative of the response of hours to the consumption price). The expected property is that the cross effect is positive, implying substitutability between

consumption and hours of non-work or complementarity between consumption and hours of work.

Each of these responses has generated a body of literature, which I will draw upon to calibrate the utility kernel, $U(c, h)$. In addition, in the presence of uncertainty, the curvature of U also controls risk aversion, the subject of another literature.

3.1 Functional form

I posit that individuals order consumption-hours pairs according to the period utility of Malin (2006),

$$U(c, h) = \frac{1}{1 - \eta} \left[\frac{c^{-(1/\sigma-1)} - c^{-(1/\sigma-1)}}{1/\sigma - 1} - \frac{\gamma}{1/\psi + 1} h^{1/\psi+1} \right]^{1-\eta}. \quad (4)$$

The kernel inside the brackets governs the marginal rate of substitution between consumption and work within a month and state of the world. A key feature of the specification is to take a concave transformation of the kernel before adding across months or states. The parameter η controls the concavity—positive values of η correspond to complementarity of hours and work, with $U_{c,h} > 0$.

Consumption and hours are Frisch complements if consumption rises when the wage rises (work rises and non-work falls)—see Browning et al. (1985) for a discussion of the relation between Frisch substitution and Slutsky-Hicks substitution. In the specification with $\eta = 0$, consumption and work are at the boundary. If η is positive, consumption and work are unambiguously Frisch complements. People consume more when wages are high because they work more and consume less leisure. Nothing can be deduced about limited insurance or consumption-smoothing from the reaction of consumption to changes in work opportunities, without studying the related movements of wages or, alternatively, of hours.

3.2 Calibration

The Frisch consumption demand and labor supply are:

$$c^{-1/\sigma} \left[\frac{c^{-(1/\sigma-1)} - c^{-(1/\sigma-1)}}{1/\sigma - 1} - \frac{\gamma}{1/\psi + 1} h^{1/\psi+1} \right]^{-\eta} = \lambda p \quad (5)$$

$$\gamma h^{1/\psi} \left[\frac{c^{-(1/\sigma-1)} - c^{-(1/\sigma-1)}}{1/\sigma - 1} - \frac{\gamma}{1/\psi + 1} h^{1/\psi+1} \right]^{-\eta} = \lambda w. \quad (6)$$

I calibrate at the point where consumption c , hours h , the consumption price p , and the wage w are all 1. By taking the ratio of equations (5) and (6), I infer that $\gamma = 1$ under this normalization, so I will presume that value in what follows.

The four remaining parameters of preferences are the intercept in the kernel, \underline{c} , the complementarity parameter, η , the curvature parameter for consumption, σ and the curvature parameter for work, ψ .

The intercept \underline{c} controls the point of disaster for consumers with $\eta > 0$ (it is irrelevant for those with $\eta = 0$). Disaster occurs when the term in brackets reaches zero and marginal utility rises to infinity. If $\sigma < 1$, the range of values I find appropriate, disaster occurs at consumption levels above the level of zero that drives $c^{-1/\sigma}$ to infinity. Consumers with Malin preferences who are not working ($h = 0$) reach the disaster point when $c = \underline{c}$, provided $\sigma < 1$. Disasters are only likely for people who are not working, because earnings are generally above the disaster level. I take the disaster level of consumption to be $\underline{c} = 0.2$.

I calibrate the three curvature parameters to the three basic properties of consumer-worker behavior listed earlier. In all cases, I will draw primarily upon research at the household rather than the aggregate level. The first property is risk aversion and intertemporal substitution in consumption. With additively separable preferences across states and time periods, the coefficient of relative risk aversion (CRRA) and the intertemporal elasticity of substitution are reciprocals of one another. But there is no widely accepted definition of measure of substitution between pairs of commodities when there are more than two of them. Chetty (2006) discusses two natural measures of risk aversion when hours of work are also included in preferences. In one, hours are held constant, while in the other, hours adjust when the random state becomes known. He notes that risk aversion is always greater by the first measure than the second. The measures are the same when consumption and hours are neither complements nor substitutes. In this case, the CRRA is $-\frac{cU_{c,c}}{u_c}$.

For the purposes of this calibration, I assume that research on intertemporal substitution/risk aversion measures the Frisch elasticity. I believe this is a reasonable approximation for two reasons: First, the evidence suggests that complementarity is not strong. When it is absent, all measures of the elasticity agree. Second, my results later in the paper on the relevance of the fully insured model even when, in truth, insurance is only partial, imply that the research actually reveals the Frisch elasticity. When confronted with a change in relative prices from one period to the next, the fully insured consumer responds along the

Frisch demand function—the Borch-Arrow insurance condition has the same content as the Frisch demand.

Appendix A summarizes the findings of recent research on the three key properties of the Frisch consumption demand and labor supply system. The own-elasticities have been studied extensively. The literature on measurement of the cross-elasticity is sparse, but a substantial amount of research has been done on an equivalent issue, the decline in consumption that occurs when a person moves from normal hours of work to zero because of unemployment or retirement. I believe that a fair conclusion from the research is that a person in the middle of the joint distribution of the three properties has a Frisch elasticity of consumption demand of -0.4 , a Frisch elasticity of labor supply of 0.7 , and a level of consumption when not working 15 percent below the level when working.

I incorporate the third property in the calibration by requiring

$$0.85^{-1/\sigma} \left[\frac{\underline{c}^{-(1/\sigma-1)} - .85^{-(1/\sigma-1)}}{1/\sigma - 1} \right]^{-\eta} = \lambda \quad (7)$$

Notice that this calibration does not require me to take a stand on whether people who are not working have chosen that condition voluntarily, against other available choices. Equation (7) holds whether the choice is voluntary or involuntary. Some of the research on the effects of unemployment and retirement on consumption has interpreted the decline as the result of frictions in capital and insurance markets. I make the hypothesis in this part of the paper that the decline arises mainly from the Frisch substitutability of work and consumption, not from failures of insurance and capital markets. A number of the studies cited in the appendix support this hypothesis.

3.3 Parameter values

The values of the parameters of preferences corresponding to the three values of the Frisch properties are $\sigma = 0.48$, $\psi = 1.27$, and $\eta = 3.48$, obtained by numerical solution of equations (5), (6), and (7). The uncompensated elasticity of static labor supply is -0.38 , so supply is backward-bending, consistent with many studies of static labor supply—see Blundell and MaCurdy (1999). The compensated elasticity is 0.35 , also in line with these studies.

4 Supply of Hours per Worker

The supply function for hours, $h(\lambda, w)$, solves equations (5) and (6) under the normalization $p = 1$. It is the supply function for the hours of employed people, the first factor of the overall supply function for hours, which is $h(\lambda, w)n(\lambda, w)$, hours per worker times the employment rate.

4.1 Transitory and permanent changes in the marginal product w

An increase in w with λ held constant implies that the change is temporary—it has no effect on the worker’s overall well being, measured by marginal utility. A permanent increase in w lowers λ by raising the lifetime value of earnings. A simple way to measure the effect of a permanent increase is to set λ to the marginal utility of consumption at the level that would prevail if a person chose a level of consumption equal to earnings given w . This occurs at the solution to the static labor-supply conditions, $\lambda = U_c(wh, h)$ and $w\lambda = -U_h(wh, h)$. The change in h resulting from this calculation is the uncompensated response of labor supply to a change in the wage.

The elasticity of hours with respect to w is the Frisch elasticity, 0.7, for transitory changes (λ held constant) and the uncompensated elasticity, -0.39 , for permanent changes. Figure 2 shows the hours supply function $h(\lambda, w)$ for transitory and permanent changes in w .

5 The Employment Function

The employment function maps conditions in the labor market into the employment rate, n . Although the employment rate is one of the two factors in the overall supply of labor, the employment function is not strictly a feature of labor supply alone—it describes choices of firms as well as workers.

5.1 General considerations

What aspects of the labor-market environment determine the employment rate? The answer depends on beliefs about the class of theories of employment and unemployment that best describe the actual labor market. I distinguish three broad classes of theories.

First, the *pure equilibrium model* of employment launched by Rogerson (1988) places workers at their points of indifference between work and non-work, so the wage just offsets




Figure 2. Hours Supply for Transitory and Permanent Changes in the Marginal Product w

the disamenity of the loss of time at home. Labor supply is perfectly elastic at that wage. The employed are those who wind up in jobs at the labor demand prevailing at that wage. An important feature of this model is the fixity of hours—workers choose between not working and working a prescribed number of hours.

Second, *search-and-matching models*—surveyed recently by Rogerson, Shimer and Wright (2005)—divide the labor market into many submarkets, each in equilibrium. Unemployment arises because some workers are in markets where their marginal products do not cover the disamenity of work. The canonical Mortensen and Pissarides (1994) model is a leading example: Workers are either in autarky, unmatched with any employer, in which case they have zero marginal product by assumption, or they are matched and are employed at a marginal product above their indifference point. Job-seekers enjoy a capital gain upon finding a job. Although most search-and-matching models assume fixity of hours, that assumption is not essential and is straightforward to relax—Andolfatto (1996) was a pioneer on this point. A key assumption of the MP model is that the firm’s demand for labor is perfectly elastic. This assumption only makes sense if the labor market is at the point where the total supply of hours equals the total demand for hours at the marginal product w .

Third, *allocational sticky-wage models* invoke a state variable, the sticky wage, that controls the allocation of labor. Employers choose total labor input to set the marginal product

of labor to the sticky wage. In that case, the sticky wage is the marginal product, w , as well. As far as I know, the literature lacks a detailed, rigorous account of the resulting equilibrium in the labor market comparable to the MP model. Blanchard and Gali (2006)'s recent contribution in this area is not what I have in mind. Their model has sticky compensation, as in Hall (2005a), not an allocational sticky wage. One simple view is that employed workers work $h(\lambda, w)$ hours and that the number employed, n , is the total number of hours demanded divided by $h(\lambda, w)$. Unemployment of the rent-seeking type in Harris and Todaro (1970) results whenever n falls short of the labor force. In that case, the unemployed are those queued up for scarce jobs. The arguments of the employment function $n(\cdot)$ include λ , w , and the other determinants of labor demand. But n depends *negatively* on λ because a higher value results in more hours of work by the employed and thus fewer jobs. And n depends negatively on w for a similar reason and because labor demand falls with w . Finally, n depends on the other determinants of labor demand. Thus allocational sticky-wage models have rather different implications for the employment function.

My approach here is to posit that the two key variables for the employment function are the same as for hours supply— λ to capture the perceived long-term well being of workers, and w to capture the immediate payoff to work. In the class of models where employment depends just on these aspects of the environment, a value of w that is high in relation to λ tightens the labor market and results in high employment. An important implication of this property is that the response of unemployment to changes in w is stronger when λ remains constant—a transitory change in w —than when the change is permanent and λ changes as well. Pissarides (1987) made this point early in the development of the MP literature, though without a full development of the underlying preferences. Blanchard and Gali (2006) make the same point for the special case of the preferences in this paper when $\eta = 0$ (separability between hours and consumption) and $\sigma = 1$ (log of consumption).

The equilibrium model plainly belongs to this class. In that model, labor supply is perfectly elastic at a value of w dictated by λ . The employment function $n(\lambda, w)$ is a correspondence mapping the two variables into 1.0 if w is above the critical value, into the unit interval at that value, and into zero below the value. Shortly I will demonstrate that the extended MP model is also a member of the class of models with employment function $n(\lambda, w)$. On the other hand, allocational sticky-wage models are not in the class because they require that employment shift along with the non-wage determinants of labor demand.

I will proceed on the assumption that a function $n(\lambda, w)$ that gives the employment rate n in an environment where marginal utility is λ and the marginal product is w is a reasonable way to think about the employment factor in the total labor supply function. The next step is to measure the response of employment to the two determinants.

5.2 The empirical employment function

According to the model, fluctuations in overall well being, measured by λ , and in the immediate benefit of work, w , result in movements of hours and unemployment along a fixed path. When w is high in relation to λ , hours are unusually high and unemployment is unusually low. In this section, I show that the behavior of the U.S. labor market over the past 5 decades confirms this aspect of the model.

The model with intertemporal substitution parameter σ less than one is inconsistent with steady-state growth. The substantial decline in hours per worker over the past 50 years supports this proposition. Other influences, notably the vacancy cost, probably have trends. Changes in the technology of home production create trends in the parameters of preferences, which are a reduced form for a more complicated setup in the household. To sidestep these issues, I work with detrended data. I remove a constant exponential trend from each time series.

In a model with unemployment, the employed, facing a continuous choice of hours of work, choose a level of consumption c_1 and hours of work h to satisfy the Frisch system of equations (5) and (6). The unemployed are constrained at zero hours of work and choose a different level of consumption, c_0 , from equation (5) by itself.

Recall that I calibrate to the observation that the consumption of the unemployed is 85 percent of the consumption of the employed. I start by extracting employed consumption c_1 from data on consumption per capita by solving the system,

$$c = uc_0 + (1 - u)c_1 \tag{8}$$

and

$$c_0 = 0.85c_1. \tag{9}$$

I calculate λ as marginal utility from equation (5) and then w from the marginal disutility of hours of work from equation (6), divided by my estimate of λ .




Figure 3. Inferred Detrended Marginal Utility, λ , and Marginal Product of Labor, w

Figure 3 shows the resulting estimate of marginal utility, λ , and the marginal product, w , using data from the U.S. National Income and Product Accounts, Table 1.1.6, for consumption, divided by total U.S. population, and hours per worker from Tables 6.4 and 6.9. I converted the results to indexes and detrended them, with equal starting and ending values.

Figure 3 shows a pronounced negative correlation between marginal utility and the marginal product. Times of high marginal utility, such as the recession of 1981-82, tend to be times when the marginal product is low. Consumption was low at the same time as the marginal product, presumably because the public believed that times were bad in a persistent way.

To gain a rough idea of the empirical elasticities, I regressed the log of the employment rate (one minus the standard BLS unemployment rate) on the logs of λ and w . The implicit identifying condition is that the disturbance in the regression is orthogonal to the two variables. Table 1 shows the estimated elasticities.

The regression in log-levels shows evidence of slow-moving omitted influences—the Durbin-Watson statistic is 0.47. The regression in first differences of logs puts less weight on lower

<i>Variable</i>	<i>Interpretation</i>	<i>Levels</i>		<i>First differences</i>	
		<i>Elasticity</i>	<i>Standard error</i>	<i>Elasticity</i>	<i>Standard error</i>
λ	Marginal utility	0.151	(0.032)	0.176	(0.038)
w	Shadow wage	0.246	(0.054)	0.389	(0.045)
	R^2		0.45		0.64
	Durbin-Watson statistic		0.47		2.20
	Standard error of the regression		0.0120		0.0069

Table 1. Estimated Elasticities of Employment with Respect to Marginal Utility λ and the Marginal Product w

frequencies and probably yields the preferred estimates. As expected, employment responds positively to both determinants. Neither elasticity is large. The volatility of the two driving forces, shown in Figure 2, is sufficient to account for a good part of the movements of the employment rate without invoking large elasticities.

The results in Table 2 are moderately favorable for a model of the employment rate that emphasizes positive responses to λ and w —such as the extended MP model—and moderately unfavorable for the allocational sticky-wage type of model. I say “moderately” because I lack a strong rationale for the identification of the parameters and because the right-hand variables are constructed from hours, not measured independently.

5.3 The extended Mortensen-Pissarides employment function

To interpret the empirical employment function, I turn to Mortensen and Pissarides (1994)’s search and matching model, which has proven remarkably helpful in understanding unemployment. In its simplest version, workers and jobs are homogeneous. I will retain this assumption—so in effect I am studying the labor market for a particular type of worker, not the market in general—but I extend the model’s treatment of labor demand and labor supply to relate them to w and λ respectively, and I generalize the compensation bargain, to match the observed relation between these driving forces and unemployment. I will define and estimate a key parameter of the compensation bargain to match the employment function I found in the data. I also make hours of work part of the bargain.

Appendix B lays out the extension of the MP model. In the model, matching frictions delay the process of finding a new job after an earlier job has ended. The job-finding rate is a key variable because the separation rate governing the flow of new job-seekers into the

pool of the unemployed is assumed, realistically, to be constant. The job-finding rate, in turn, depends on the recruiting efforts of employers in relation to the number of job-seekers. Employers put resources into recruiting sufficient to drive the economic benefit from a new hire to zero.

The firm and worker take the marginal product w as given. They bargain over hours and compensation. Their bargaining problem fits the Edgeworth-box paradigm—the choice of hours at the point where the marginal rate of substitution equals w places them on their contract curve and they bargain over compensation, the location along the contract curve. See Reichling (2006) for a much more extensive discussion of the MP model with choice of hours.

The marginal product w influences employment through the recruiting decision. A higher payoff to employment resulting from a higher w causes employers to recruit harder, raises the job-finding rate and the employment rate. Part of a higher w goes to workers as a rent because they can bargain for higher compensation. As Shimer (2005) showed, if the compensation bargain follows the principle of Nash—say equal splitting of the surplus from the bargain—essentially all of the increase goes as rent, employers gain almost no benefit, raise recruiting by little, and employment hardly rises at all. If the compensation bargain results in less of an increase, the recruiting and employment response is stronger.

I introduce a parameter κ that controls the stabilization of the compensation bargain. The less responsive is compensation to the marginal product, the higher are the elasticities of the MP model’s employment function. The Nash bargain has $\kappa = 0$ and elasticities of essentially zero. At a substantially positive value of κ , the employment function has the same elasticities I found in the data. Compensation does not shift nearly as much in the worker’s favor with higher w as it would with Nash bargaining.

The marginal utility of consumption, λ , enters the extended MP model by determining the value of time at home in relation to the value of work. When λ is high, job-seekers are more interested in finding work because they value time away from work less. Workers have lower reservation wages as a result, and the wage bargain is more favorable to the employer. Thus employment is an increasing function of λ .

A review of the details of the extended MP model in Appendix B shows that w and λ are the only two outside determinants of employment. Changes in some of the parameters of the model do cause changes in employment, but none seem likely candidates to cause




Figure 4. The Elasticities of the Employment Rate in the MP Model as Functions of the Compensation Stabilization Parameter, κ

changes over the business cycle—they appear more likely to generate small, smooth changes over time.

Figure 4 shows the key feature of the MP model’s employment function, its elasticities with respect to w and λ , as functions of the compensation stabilization parameter κ . The elasticities are small unless κ is reasonably close to one. The elasticity of employment with respect to λ is just less than half the wage elasticity for all values of κ . The two variables w and λ enter the employment function in a separable combination, so the ratio of the derivatives and elasticities of n with respect to them do not depend on κ .

From Figure 4, one can see that the estimated elasticities of 0.176 for λ and 0.389 for w correspond to a value of κ of 0.88. This value is an estimate based on an informal application of indirect inference. Notice that the two coefficients have close to the theoretical ratio of two to one.

5.3.1 Implications of the Frisch elasticity of labor supply

The Frisch elasticity of labor supply of 0.7 refers to men. Estimates for women are generally higher than those for men. In addition, some authors have suggested reasons for a

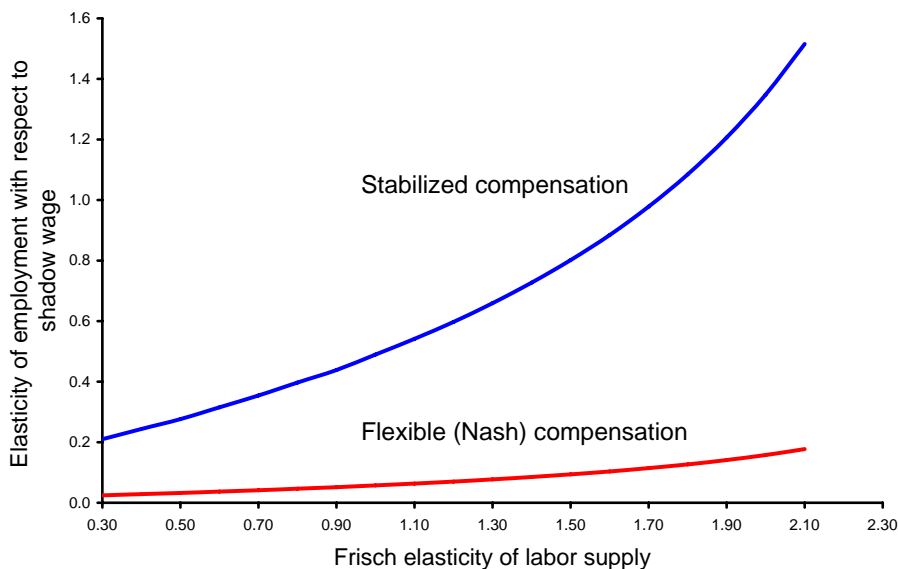


Figure 5. Elasticity of Employment with Respect to w , as a Function of the Frisch Elasticity of Labor Supply, for Nash and Stabilized Compensation

downward bias in estimated Frisch elasticities. On the other hand, many estimates of the Frisch elasticity are lower than 0.7. Here I explore the implications of a range of values of the elasticity.

I keep the Frisch elasticity of consumption demand at -0.4 and non-worker/worker consumption ratio at 0.85. I recompute the values of σ , ψ , and η for a range of values of the Frisch elasticity of labor supply. The values of σ and η are not sensitive to the elasticity, while ψ varies in approximate proportion to it. I keep the worker's share of the surplus at one-half, I solve for new values of the parameter k_0 governing the flow cost of a vacancy, so that I always match the unemployment rate of 5.5 percent.

Figure 5 illustrates the importance of the Frisch elasticity of labor supply, shown on the horizontal axis. Evidence from household data finds a value of about 0.7, as I discussed earlier. The upper curve shows the elasticity of employment with respect to the marginal product w , in the case of a stabilized wage with $\kappa = 0.88$. The lower curve is the same in the case of a Nash wage bargain. These responses refer to transitory changes in w , in the sense that λ is held constant.

Figure 5 shows that a somewhat lower value of the compensation stabilization parameter κ would match the observed elasticity of employment with respect to w if labor supply is more elastic. It also confirms Shimer (2005)'s finding of low amplification with Nash bargaining,

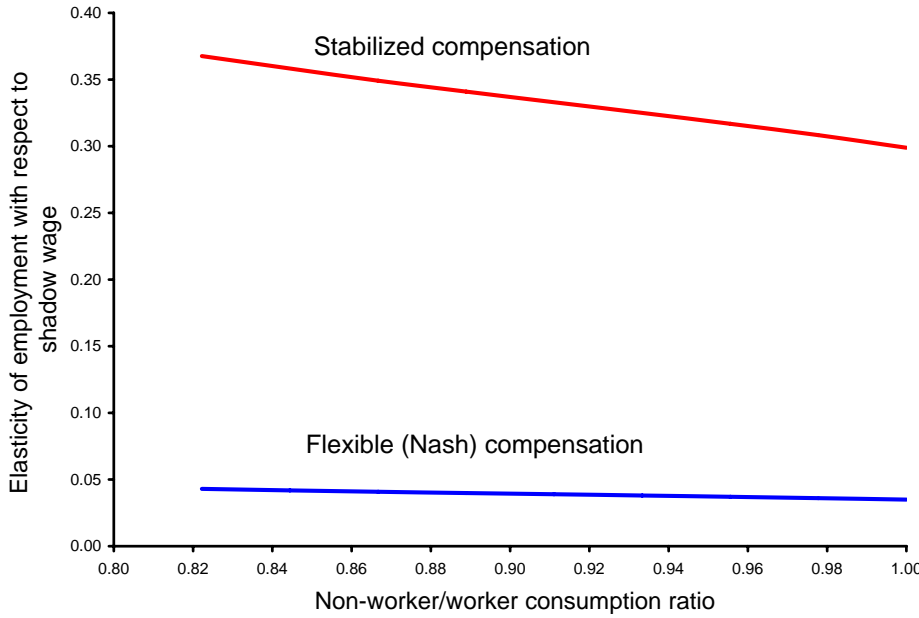


Figure 6. Elasticity of Employment with Respect to w , as a Function of the Non-Worker/Worker Consumption Ratio

even for values of the Frisch elasticity of labor supply well above Shimer’s implicit value of around 0.4.

5.3.2 Implications of the complementarity of hours and consumption

Figure 6 shows how the response of employment to changes in w varies with the complementarity of hours and consumption. The horizontal axis is the ratio of consumption while unemployed to consumption while employed. Stronger complementarity, at the left in the figure, corresponds to higher sensitivity of employment to the marginal product w .

5.3.3 The employment function

Figure 7 shows the employment function, $n(\lambda, w)$, with constant λ , corresponding to transitory changes in w .

5.4 Fitted unemployment

Given the time series for λ and w and the estimate $\kappa = 0.88$, I evaluate the MP unemployment function $u(\lambda, w)$. Figure 8 compares the fitted and actual time series for U.S. unemployment. Most of the important movements of unemployment are captured by the




Figure 7. The Employment Function, $n(\lambda, w)$, with Constant λ

fitted series, both at cyclical and lower frequencies. The facts lying behind this favorable finding are straightforward. Hours per worker is a good measure of the state of the labor market. When labor is scarce, hours are high. At the same time, unemployment is low.

Previous research in this area has lacked a well-founded measure of scarcity. Productivity, the primary earlier measure, has two major disadvantages. It is difficult to measure, and many of its movements are permanent, in which case theory suggests little effect on unemployment.

6 The Total Labor Supply Function

The total labor supply function is the product, $h(\lambda, w)n(\lambda, w)$, of the function in Figure 2 for hours and Figure 7 for employment. The elasticity of the first with respect to w , holding λ constant, is the Frisch elasticity, 0.7. The elasticity of the employment function is the estimated figure, about 0.4, so the total elasticity is about 1.1, again with respect to transitory fluctuations in w .

7 Measuring the Marginal Product w from Value Added

In Figure 3, I inferred the value of the marginal product w from the marginal disutility of labor. In this section, I measure the marginal product from the firm's perspective. I assume




Figure 8. Actual and Fitted Unemployment Rates, 1948-2005

the firm has a constant-returns technology, with

$$Q = AF(N, K, M), \quad (10)$$

showing how hours of labor N , hours of capital services, K , and materials M combine with efficiency A to make output Q . From the first-degree homogeneity of F ,

$$Q = A \frac{\partial F}{\partial N} N + A \frac{\partial F}{\partial K} K + A \frac{\partial F}{\partial M} M. \quad (11)$$

The firm sells its output in a competitive market at a price normalized at one, rents capital in a competitive market with rental price ρ , and purchases materials in a competitive market at price p_M . From the first-order condition for profit maximization,

$$A \frac{\partial F}{\partial K} = \rho \quad (12)$$

$$A \frac{\partial F}{\partial M} = p_M. \quad (13)$$

I substitute these prices for the marginal products and solve for $w = \partial F / \partial N$:

$$w = \frac{Q - \rho K - p_M M}{N}. \quad (14)$$

The difference between revenue and materials purchases is value added, $V = Q - p_M M$. The calculation of w from data takes the simple form,

$$w = \frac{V - \rho K}{N}. \quad (15)$$

7.1 Cobb-Douglas example

Let the production function be

$$Q = AN^\alpha K^\gamma M^{1-\alpha-\gamma}, \quad (16)$$

With Cobb-Douglas technology, the factor shares are in proportion to the elasticities, so

$$\rho K = \frac{\gamma}{\alpha} w N \quad (17)$$

$$p_M M = \frac{1 - \alpha - \gamma}{\alpha} w N \quad (18)$$

Inserting these values and solving for w yields

$$w = A' \rho^{-\gamma/\alpha} p_M^{-(1-\alpha-\gamma)/\alpha}. \quad (19)$$

Here A' is a constant proportional to $A^{1/\alpha}$. If consumers pay a sales tax at rate τ on their purchases, the value of w to be equated to the marginal rate of substitution is $1/(1 + \tau)$ times the value above.

Notice that the efficiency level, A , the labor elasticity, α , the tax rate, τ , and the rental price of capital, ρ , all enter through the marginal product, w . Another potential driving force that operates through w is market power. Without a fuller specification of the product demand side of the economy, one cannot embody market power explicitly in w . But the tax, τ , can be considered a proxy for market power. Collusion among sellers can be viewed as a tax imposed on sales that is rebated as a lump sum to the colluding firms. Although one normally thinks of collusion as raising prices, here, where output is taken as numeraire and its price normalized at one, the effect of collusion or tax is to lower the wage in units of output. But the effects of market power will escape measurement using equation (15) unless the lump-sum rebate is deducted from value added.

7.2 Data

Measuring w at the level of the total economy from value added is a challenge for several reasons. First, value added in government is not based on the market value of output, but is

built up from factor payments, so w cannot be inferred from a residual—there would be no residual if the construction of the NIPAs could be reproduced fully. Second, the same is true of value added arising from the services of owner-occupied housing, another large component of GDP. Third, almost all studies of non-corporate business, especially small proprietorships, has found that revenue falls short of a reasonable imputed wage and rental cost of capital. For these reasons, I measure the marginal product w for corporate business.

Data on the components of value added appear in the NIPAs. I take the sum of profits before tax, interest paid, and compensation of employees (Table 1.15), together with capital consumption allowances (Table 6.22). I calculate the rental price of capital using the approach of Hall and Jorgenson (1967). I obtain the quantity of corporate capital from the Fixed Assets Tables of the NIPAs, Table 4.1, and estimate corporate inventories as the corporate share of total inventories, NIPA Table 5.12. I calculate the corporate marginal product w_{Corp} from equation 15. For comparability to the earlier results, I then approximate the marginal product for the total economy as

$$w = y \frac{w_{Corp}}{y_{Corp}}, \quad (20)$$

where y , as before, is compensation and y_{Corp} is corporate compensation, both per worker.

7.3 Comparison of measures of the marginal product

Figure 9 compares the marginal product w , computed as the residual from value added on the firm’s side according to equation (15), with the earlier measure from the worker’s side. Although the two differ to some extent at medium frequencies, there is substantial agreement in cyclical movements. Earlier work, including Hall (1997), found a substantial and variable gap between the two. This paper makes progress in closing the gap, by taking a more empirical approach to measuring the marginal product.

The volatilities of the two series in Figure 9 are quite similar. The idea of this paper—that the marginal product is a reasonably volatile driving force, able to account for the observed volatilities of hours and employment with reasonable elasticities—receives some support from Figure 9. In first differences, the two measures of the marginal product have standard deviations almost double that of productivity, as measured by the BLS’s Multifactor Productivity Index.

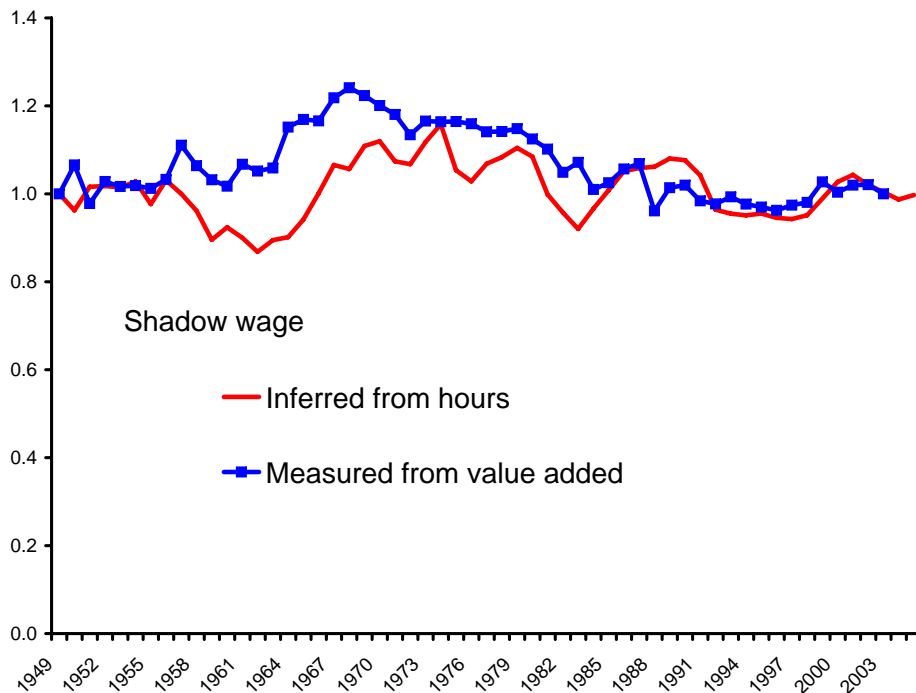


Figure 9. Marginal Product w , Measured from Value Added and Inferred from Hours, 1948-2005

7.4 Compensation-smoothing

I measure actual compensation per hour as the ratio of total compensation, NIPA Table 6.2, to hours of work, Table 6.9. Figure 10 compares this measure of compensation to the stabilized hourly compensation predicted by the extended MP model, y/h , and to the marginal product w inferred from hours. Fitted hourly compensation from the model is much more volatile than is actual compensation. Further, fitted hourly compensation is almost indistinguishable from the marginal product w —I raise the fitted line wage above w so that both show in the figure. Actual compensation involves much more smoothing than does fitted compensation, even though the bargain underlying fitted compensation involves stabilization and therefore some smoothing.

The bargained level of compensation has no allocational role once a job-seeker and an employer find each other—it only divides the surplus from the match. Compensation gains its influence over unemployment through the non-contractible, pre-match effort of employers in attracting workers. These efforts—which take the form of the creation of vacancies in the model—govern the tightness of the labor market and thus the unemployment rate. The difference between the marginal product and compensation, anticipated at the time of




Figure 10. Actual and Fitted Hourly Compensation and Inferred Marginal Product w , 1948-2005

hiring, governs the employer’s vacancy-creation efforts. This proposition has two important implications for Figure 10.

First, the actual wage shown in the figure is not the anticipation for newly hired workers, but is the average paid to workers employed at the time, many hired years earlier. The smoothness of the actual wage is the result of even greater stability of wages paid to workers not recently hired. Consequently, average compensation has little direct role in evaluating the idea of compensation stabilization. A full analysis would need data on compensation by seniority and would need to evaluate the difference between the forward-looking compensation cost of a newly hired worker and the forecast of the worker’s marginal product.

Second, the gap between the marginal product and hourly compensation in my calibration is small—in the middle year, 1976, they are 1.040 and 1.016. The Nash level of hourly compensation in the same year is 1.014. Thus the effects of compensation stabilization relative to the marginal product or to Nash bargaining are hard to detect.

Although the central issue in unemployment volatility appears to be stabilization of employers’ expectations of compensation for new hires, Figure 10 makes it clear that a study of compensation behavior cannot add to understanding of volatility within the framework of this paper.

8 Concluding Remarks

Contrary to earlier impressions, one can make sense out of the fairly large cyclical fluctuations in hours of work per person without invoking either unreasonably high elasticity of labor supply—as in real business cycle models—or allocational sticky wages. A Frisch elasticity of labor supply of 0.7, easily in the range found in recent research using household data, does the job.

More than half of cyclical fluctuations in hours per person is in movements of hours of job-holders. I argue that transitory movements in the marginal product of labor (not measured compensation per hour, which appears to be smoothed) are a plausible source of the movements of hours. I reverse-engineer the Frisch supply function for hours and demand for consumption to show the required amount of cyclical volatility. I show that it is similar to a measure of the marginal product derived as a residual from value added.

The remaining part of cyclical fluctuations in hours per person comes from unemployment. Hours per person decline in recessions because fewer people work and more are looking for work. I show that the U.S. labor market appears to have a well defined employment function with reasonable positive elasticities for both the marginal product of labor and the marginal utility of consumption.

The only analytical benchmarks for movements of unemployment in response to changes in incentives come from search-and-matching models, notably the Mortensen-Pissarides model. The MP model falls far short of accounting for the observed cyclical volatility of employment, even when the relatively volatile marginal product and marginal utility of consumption are taken as the driving forces. Shimer's finding of inertness of unemployment in the MP model is hard to shake. It takes a rather large departure from the Nash bargain of the MP model to rationalize the observed elasticities of the employment function, though the rationalization does work.

Further work on the employment function, either in the framework of the marginal product and marginal utility of consumption or outside that framework, is clearly in order.

Another area for further work is the improvement of the direct measurement of the marginal product from value added.

References

- Aguiar, Mark and Erik Hurst, “Consumption versus Expenditure,” *Journal of Political Economy*, October 2005, *113* (5), pp. 919–948.
- Andolfatto, David, “Business Cycles and Labor-Market Search,” *American Economic Review*, 1996, *86* (1), 112–132.
- Attanasio, Orazio P. and Guglielmo Weber, “Consumption Growth, the Interest Rate, and Aggregation,” *Review of Economic Studies*, July 1993, *60* (3), pp. 631–649.
- and — , “Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey,” *Journal of Political Economy*, December 1995, *103* (6), pp. 1121–1157.
- and Hamish Low, “Estimating Euler Equations,” *Review of Economic Dynamics*, 2004, *7* (0), pp. 406–435.
- , James Banks, Costas Meghir, and Guglielmo Weber, “Humps and Bumps in Lifetime Consumption,” *Journal of Business & Economic Statistics*, January 1999, *17* (1), pp. 22–35.
- Banks, James, Richard Blundell, and Sarah Tanner, “Is There a Retirement-Savings Puzzle?,” *American Economic Review*, September 1998, *88* (4), pp. 769–788.
- Barsky, Robert B., F. Thomas Juster, Miles S. Kimball, and Matthew D. Shapiro, “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study,” *Quarterly Journal of Economics*, May 1997, *112* (2), pp. 537–579.
- Basu, Susanto and Miles S. Kimball, “Long-Run Labor Supply and the Elasticity of Intertemporal Substitution for Consumption,” December 2000. University of Michigan.
- Blanchard, Olivier and Jordi Gali, “A New Keynesian Model with Unemployment,” July 2006. MIT Department of Economics.
- Blundell, Richard and Thomas MaCurdy, “Labor Supply: A Review of Alternative Approaches,” in Orley C. Ashenfelter and David Card, eds., *Handbook of Labor Economics, Volume 3A*, Amsterdam: North-Holland, 1999, chapter 27, pp. 1559–1695.

- Browning, Martin and Thomas F. Crossley, “Unemployment Insurance Benefit Levels and Consumption Changes,” *Journal of Public Economics*, 2001, 80, 1–23.
- , Angus Deaton, and Margaret Irish, “A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle,” *Econometrica*, 1985, 53 (3), pp. 503–544.
- Campbell, John Y and John H. Cochrane, “By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior,” *Journal of Political Economy*, April 1999, 107 (2), 205–251.
- Carroll, Christopher D., “Death to the Log-Linearized Consumption Euler Equation! (And Very Poor Health to the Second-Order Approximation),” *Advances in Macroeconomics*, 2001, 1. Issue 1, Article 6.
- Chetty, Raj, “A Bound on Risk Aversion Using Labor Supply Elasticities,” *American Economic Review*, 2006. forthcoming.
- Cohen, Alma and Liran Einav, “Estimating Risk Preferences from Deductible Choice,” June 2005. NBER Working Paper 11461.
- Domeij, David and Martin Floden, “The Labor-Supply Elasticity and Borrowing Constraints: Why Estimates Are Biased,” *Review of Economic Dynamics*, 2006, 9 (0), pp. 242–262.
- Fisher, Jonathan, David S. Johnson, Joseph Marchand, Timothy M Smeeding, and Barbara Boyle Terrey, “The Retirement Consumption Conundrum: Evidence from a Consumption Survey,” December 2005. Center for Retirement Research at Boston College, Working Paper 1005-14.
- Guvenen, Fatih, “Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective,” *Journal of Monetary Economics*, 2006. forthcoming.
- Hagedorn, Marcus and Iourii Manovskii, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited,” April 2006. University of Frankfurt and University of Pennsylvania.

- Hall, Robert E, “Macroeconomic Fluctuations and the Allocation of Time,” *jle*, January 1997, *15* (1), S223–S250.
- Hall, Robert E., “Employment Fluctuations with Equilibrium Wage Stickiness,” *American Economic Review*, March 2005, *95* (1), 50–65.
- , “Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years,” *NBER Macroeconomics Annual*, 2005, *20* (1), 101–137.
- , “Complete Markets as an Approximation to the Bewley Equilibrium with Unemployment Risk,” 2006. Hoover Institution, Stanford University, stanford.edu/ rehall.
- and Dale W. Jorgenson, “Tax Policy and Investment Behavior,” *American Economic Review*, 1967, *57* (3), 391–414.
- Hansen, Lars Peter and Ravi Jagannathan, “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy*, April 1991, *99* (2), 225–262.
- Harris, John R. and Michael P. Todaro, “Migration, Unemployment, and Development: A Two-Sector Analysis,” *American Economic Review*, 1970, *60* (1), 126–142.
- Kimball, Miles S and Matthew D. Shapiro, “Labor Supply: Are the Income and Substitution Effects Both Large or Both Small?,” May 2003. Department of Economics, University of Michigan.
- King, Robert G., Charles I. Plosser, and Sergio T. Rebelo, “Production, Growth, and Business Cycles: I. The Basic Neoclassical Model,” *Journal of Monetary Economics*, March 1988, *21* (0), 195–232.
- Malin, Benjamin, “Lower-Frequency Macroeconomic Fluctuations: Living Standards and Leisure,” February 2006. Department of Economics, Stanford University.
- Mehra, Rajnish and Edward C. Prescott, “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, March 1985, *15* (2), 145–161.
- Merz, Monika, “Search in the Labor Market and the Real Business Cycle,” *Journal of Monetary Economics*, 1995, *36* (2), 269–300.

- Miniaci, Raffaele, Chiara Monfardini, and Guglielmo Weber, “Is There a Retirement Consumption Puzzle in Italy?,” July 2003. Institute for Fiscal Studies, Working Paper 03/14.
- Mortensen, Dale T. and Christopher Pissarides, “Job Creation and Job Destruction in the Theory of Unemployment,” *Review of Economic Studies*, 1994, 61 (0), pp. 397–415.
- Mulligan, Casey B., “Substitution over Time: Another Look at Life-Cycle Labor Supply,” *NBER Macroeconomics Annual*, 1998, pp. pp. 75–134.
- Pissarides, Christopher A., “Search, wage Bargains and Cycles,” *Review of Economic Studies*, 1987, 54, 473–483.
- Pistaferri, Luigi, “Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply,” *Journal of Labor Economics*, 2003, 21 (3), pp. 729–754.
- Reichling, Felix, “Optimal Unemployment Insurance with Worker Self-Insurance and Realistic Preferences,” October 2006. Department of Economics, Stanford University.
- Rogerson, Richard, “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics*, 1988, 21, 3–16.
- , Robert Shimer, and Randal Wright, “Search-Theoretic Models of the Labor Market: A Survey,” *Journal of Economic Literature*, December 2005, 43, 959–988.
- Shimer, Robert, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 2005, 95 (1), 24–49.

Appendixes

A Research on Properties of Preferences

A.1 Approaches

Chetty (2006) considers the issues surrounding the calibration of household preferences. He shows that the value of the coefficient of relative risk aversion (or, though he does not pursue the point, the inverse of the intertemporal elasticity of substitution in consumption) is implied by a set of other measures. He solves for the consumption curvature parameter by drawing estimates of responses from the literature on labor supply. One is the third item on the list above, consumption-hours complementarity. The others are the compensated wage elasticity of static labor supply and the elasticity of static labor supply with respect to unearned income. These are functions of the derivatives listed above, so information about static labor supply does not add anything that those derivatives miss. In principle, as long as the mapping has adequate rank, one could take any set of measures of behavior and solve for the slopes of the Frisch functions or any other representation of preferences. My procedure links the empirical measures more directly to the underlying basic properties of preferences. I do, however, study the implications of my calibration for static labor supply. My calibration lies within the space of values that Chetty extracts from a wide variety of studies of static labor supply.

Basu and Kimball (2000) pursue an idea related to Chetty's. They calibrate preferences to an outside estimate of the intertemporal elasticity of substitution in consumption and to zero uncompensated elasticity of static labor supply with respect to the wage. They constrain the complementarity of consumption and hours to have the multiplicative form of King, Plosser and Rebelo (1988).

A.2 Risk aversion

Research on the value of the CRRA falls into several broad categories. In finance, a consistent finding within the framework of the consumption capital-asset pricing model is that the CRRA has high values, in the range from 10 to 100 or more. Mehra and Prescott (1985) began this line of research. A key step in its development was Hansen and Jagannathan (1991)'s demonstration that the marginal rate of substitution—the universal stochastic discounter in the consumption CAPM—must have extreme volatility to rationalize the equity premium.

Models such as Campbell and Cochrane (1999) generate a highly volatile marginal rate of substitution from the observed low volatility of consumption by the trick of subtracting an amount almost equal to consumption before measuring the MRS. This trick does not seem plausible as a model of household consumption.

A second body of research considers experimental and actual behavior in the face of small risks and generally finds high values of risk aversion. For example, Cohen and Einav (2005) find that the majority of car insurance purchasers behave as if they were essentially risk neutral in choosing the size of their deductible, but a minority are highly risk-averse, so the average coefficient of relative risk aversion is about 80. But any research that examines small risks, such as having to pay the amount of the deductible or choosing among the gambles that an experimenter can offer in the laboratory, faces a basic obstacle: Because the stakes are small, almost any departure from risk-neutrality, when inflated to its implication for the CRRA, implies a gigantic CRRA. The CRRA is the ratio of the percentage price discount off the actuarial value of a lottery to the percentage effect of the lottery on consumption. For example, consider a lottery with a \$20 effect on wealth. At a marginal propensity to consume out of wealth of 0.05 per year and a consumption level of \$20,000 per year, winning the lottery results in consumption that is 0.005 percent higher than losing. So if an experimental subject reports that the value of the lottery is one percent—say 10 cents—lower than its actuarial value, the experiment concludes that the subject's CRRA is 200!

Remarkably little research has investigated the CRRA implied by choices over large risky outcomes. One important contribution is Barsky, Juster, Kimball and Shapiro (1997). This paper finds that almost two-thirds of respondents would reject a new job with a 50 percent chance of doubling income and a 50 percent chance of cutting income by 20 percent. The cutoff level of the CRRA corresponding to rejecting the hypothetical new job is 3.8. Only a quarter of respondents would accept other jobs corresponding to CRRAs of 2 or less. The authors conclude that most people are highly risk averse. The reliability of this kind of survey research based on hypothetical choices is an open question, though hypothetical choices have been shown to give reliable results when tied to more specific and less global choices, say, among different new products.

A.3 Intertemporal substitution

Attanasio, Banks, Meghir and Weber (1999), Attanasio and Weber (1993), and Attanasio and Weber (1995) are leading contributions to the literature on intertemporal substitution in consumption at the household level. These papers examine data on total consumption (not food consumption, as in some other work). They all estimate the relation between consumption growth and expected real returns from saving, using measures of returns available to ordinary households. All of these studies find that the elasticity of intertemporal substitution is around 0.7.

Barsky et al. (1997) asked a subset of their respondents about choices of the slope of consumption under different interest rates. They found evidence of quite low elasticities, around 0.2.

Guvenen (2006) tackles the conflict between the behavior of securities markets and evidence from households on intertemporal substitution. With low substitution, interest rates would be much higher than are observed. The interest rate is bounded from below by the rate of consumption growth divided by the intertemporal elasticity of substitution. Guvenen's resolution is in heterogeneity of the elasticity and highly unequal distribution of wealth. Most wealth is in the hands of those with elasticity around one, whereas most consumption occurs among those with lower elasticity.

Finally, Carroll (2001) and Attanasio and Low (2004) have examined estimation issues in Euler equations using similar approaches. Both create data from the exact solution to the consumer's problem and then calculate the estimated intertemporal elasticity from the standard procedure, instrumental-variables estimation of the slope of the consumption growth-interest rate relation. Carroll's consumers face permanent differences in interest rates. When the interest rate is high relative to the rate of impatience, households accumulate more savings and are relieved of the tendency that occurs when the interest rate is lower to defer consumption for precautionary reasons. Permanent differences in interest rates result in small differences in permanent consumption growth and thus estimation of the intertemporal elasticity in Carroll's setup has a downward bias. Attanasio and Low solve a different problem, where the interest rate is a mean-reverting stochastic time series. The standard approach works reasonably well in that setting. They conclude that studies based on fairly long time-series data for the interest rate are not seriously biased. My conclusion favors studies with that character, accordingly.

I calibrate to a Frisch elasticity of consumption demand of -0.4 . Again, I associate the evidence described here about the intertemporal elasticity of substitution as revealing the Frisch elasticity, even though many of the studies do not consider complementarity of consumption and hours explicitly.

A.4 Frisch elasticity of labor supply

The second property is the Frisch elasticity of labor supply. Pistaferri (2003) is a leading recent contribution to estimation of this parameter. This paper makes use of data on workers' personal expectations of wage change, rather than relying on econometric inferences, as has been standard in other research on intertemporal substitution. Pistaferri finds the elasticity to be 0.70 with a standard error of 0.09. This figure is somewhat higher than most earlier work in the Frisch framework or other approaches to measuring the intertemporal elasticity of substitution from the ratio of future to present wages. Here, too, I proceed on the assumption that these approaches measure the same property of preferences as a practical matter. Kimball and Shapiro (2003) survey the earlier work.

Mulligan (1998) challenges the general consensus among labor economists about the Frisch elasticity of labor supply with results showing elasticities well above one. My discussion of the paper, published in the same volume, gives reasons to be skeptical of the finding, as it appears to flow from an implausible identifying assumption.

Kimball and Shapiro (2003) estimate the Frisch elasticity from the decline in hours of work among lottery winners, based on the assumption that the uncompensated elasticity of labor supply is zero. They find the elasticity to be about one. But this finding is only as strong as the identifying condition.

Domeij and Floden (2006) present simulation results for standard labor supply estimation specifications suggesting that the true value of the elasticity may be double the estimated value as a result of omitting consideration of borrowing constraints.

I calibrate to a Frisch elasticity of labor supply of 0.7. I also discuss results for a range of higher values.

A.5 Consumption-hours complementarity

The third property is the relation between hours of work and consumption. A substantial body of work has examined what happens to consumption when a person stops working,

either because of unemployment following job loss or because of retirement, which may be the result of job loss.

Browning and Crossley (2001) appears to be the most useful study of consumption declines during periods of unemployment. Unlike most earlier research in this area, it measures total consumption, not just food consumption. They find a 14 percent decline on the average from levels just before unemployment began.

A larger body of research deals with the “retirement consumption puzzle”—the decline in consumption thought to occur upon retirement. Most of this research considers food consumption. Aguiar and Hurst (2005) show that, upon retirement, people spend more time preparing food at home. The change in food consumption is thus not a reasonable guide to the change in total consumption.

Banks, Blundell and Tanner (1998) use a large British survey of annual cross sections to study the relation between retirement and nondurables consumption. They compare annual consumption changes in 4-year wide cohorts, finding a coefficient of -0.26 on a dummy for households where the head left the labor market between the two surveys. They use earlier data as instruments, so they interpret the finding as measuring the planned reduction in consumption upon retirement.

Miniaci, Monfardini and Weber (2003) fit a detailed model to Italian cohort data on non-durable consumption, in a specification of the level of consumption that distinguishes age effects from retirement effects. The latter are broken down by age of the household head. The pure retirement reductions range from 4 to 20 percent. This study also finds pure unemployment reductions in the range discussed above.

Fisher, Johnson, Marchand, Smeeding and Terrey (2005) study total consumption changes in the Consumer Expenditure Survey, using cohort analysis. They find small declines in total consumption associated with rising retirement among the members of a cohort. Because retirement in a cohort is a gradual process and because retirement effects are combined with time effects on a cohort analysis, it is difficult to pin down the effect.

B The Extended Mortensen-Pissarides Model

Individuals face a transition probability $\pi_{i,i'}$ from state i (0 if unemployed; 1 if employed) to state i' . The job-finding rate is $\pi_{0,1}$ and the job-separation rate is $\pi_{1,0}$. Employed workers work h_1 hours. Naturally $h_0 = 0$. Unemployed individuals receive benefits y_0 . Workers'

earnings are y_1 . I will also use the name w for y_1 and b for y_0/w , as in the MP literature, where b is the earnings replacement rate in the unemployment insurance system.

B.1 Consumption and hours

An unemployed individual chooses consumption to satisfy:

$$U_c(c_0, 0) = \lambda. \quad (21)$$

Under efficient employment governance, the employed individual will choose c_1 and work the number of hours h that satisfy:

$$U_c(c_1, h) = \lambda \quad (22)$$

and

$$-U_h(c_1, h) = \lambda w. \quad (23)$$

B.2 Labor turnover and recruiting cost

The MP model characterizes the tightness of the labor market in terms of the vacancy/unemployment ratio θ . The job-finding rate is the increasing and concave function $\phi(\theta)$ and the vacancy-filling rate is the decreasing function $\phi(\theta)/\theta$.

Employers incur a cost k to maintain a vacancy. I take the technology of recruiting to be Cobb-Douglas in labor and goods with labor elasticity χ . The opportunity cost of labor devoted to recruiting is m per hour. Thus

$$k = k_0 m^\chi. \quad (24)$$

The model assumes a constant exogenous rate of job destruction, s . Unemployment follows a two-state Markoff process with stochastic equilibrium

$$u = \frac{s}{s + \phi(\theta)}. \quad (25)$$

The corresponding level of vacancies is $v = \theta u$. Because the job-finding rate $\phi(\theta)$ is high—more than 25 percent per month—the dynamics of unemployment are rapid. Essentially nothing is lost by thinking about unemployment as if it were at its stochastic equilibrium and treating it as a jump variable. I will adopt this convention in the rest of the discussion.

B.3 Bellman values

In a compact matrix representation of the model, I assign subscripts as follows: 0 to the job-seeker, 1 to the employed worker, 2 to the unfilled vacancy, and 3 to a filled job.

Workers receive compensation y for the h hours they work. While unemployed, they receive unemployment benefits by . The given value of the marginal utility of consumption, λ , translates values stated in consumption units into values stated in utility units. I let

$$f_0 = \frac{U(c_0, 0)}{\lambda} - c_0 + by, \quad (26)$$

for the job-seeker and

$$f_1 = \frac{U(c_1, h)}{\lambda} - c_1 + y, \quad (27)$$

for a worker. These flow values are in consumption units. Notice that the value when working depends on compensation, y , not on the marginal value of hours worked, wh .

The key object in the analysis is the loss in flow value, $f_1 - f_0$, that would occur if the worker turned down a job opportunity instead of bargaining to a successful conclusion with an employer. This loss is measured in consumption units, using the marginal utility, λ , to translate the utility gain from not working into its consumption equivalent.

The flow values for the employer are

$$f_2 = -k, \quad (28)$$

for the cost to a firm of maintaining a vacancy, and

$$f_3 = wh - y, \quad (29)$$

for the net revenue of the firm while a job is filled.

Equations (26) through (29) define a function $f(y, \theta)$ giving the four flow values in the vector, f .

The model implies a vector of Bellman values: the job-seeker's value, V_0 , the worker's value, V_1 , the employer's value of a vacancy, V_2 , and the employer's value of a filled job, V_3 . The model has a matrix of transition probabilities, $\pi(\theta)$, with non-zero values:

Unemployment persistence, $\pi_{1,1} = 1 - \phi(\theta)$

Job-finding rate, $\pi_{1,2} = \phi(\theta)$

Separation rate, $\pi_{2,1} = s$

Job retention rate, $\pi_{2,2} = 1 - s$

Vacancy persistence rate, $\pi_{3,3} = 1 - \phi(\theta)/\theta$

Vacancy filling rate, $\pi_{3,4} = \phi(\theta)/\theta$

Employee retention rate, $\pi_{4,4} = 1 - s$

Note that π is *not* a Markoff transition matrix.

The matrix Bellman equation of the model is

$$V = f(y, \theta) + \delta\pi(\theta)V, \quad (30)$$

where δ is the discount factor, so the Bellman values for given y and θ are

$$V(y, \theta) = (I - \delta\pi(\theta))^{-1} f(y, \theta). \quad (31)$$

B.4 Compensation determination

Recall that y is total compensation for a bargained number of hours. Hours are set at the efficient level. In terms of an Edgeworth box, the parties make a bargain on the contract curve (described by the efficiency condition for hours) and y is their bargained position along that curve, within the bargaining set. Workers gain $V_1 - V_0$ from forming a match and employers gain $V_3 - V_2$. I assume a generalization of the Nash bargain in which conditions in the labor market, measured by tightness θ , affect bargaining. The worker and the employer split the surplus of a job with bargaining share $\beta(\theta)$ for the worker, so

$$\beta(\theta)(V_3 - V_2) = (1 - \beta(\theta))(V_1 - V_0). \quad (32)$$

I take the bargaining share to be

$$\beta(\theta) = \min \left(\frac{\beta_0}{\beta_0 + (1 - \beta_0) (\theta/\theta^*)^\kappa}, \bar{\beta} \right), \quad (33)$$

with $\bar{\beta} > \beta_0$. The parameter θ^* is the normal value of the vacancy/unemployment ratio. If $\kappa = 0$, the bargain is Nash with the worker share constant at β_0 . For positive values of κ , the worker's share is higher in a slacker market with lower θ , resulting in a smaller decline in the wage.

I cap $\beta(\theta)$ at $\bar{\beta}$ to prevent spurious equilibria at extremely low values of θ .

B.5 Free entry

In company with most earlier versions of the MP model, I assume free entry to the product and labor markets. In consequence, the value of a vacancy, V_2 , is zero—with free entry, the stream of expected future payments of the vacancy cost k just equals the benefit from hiring a worker:

$$V_2(y, \theta) = 0. \tag{34}$$

B.6 Equilibrium

An equilibrium of the unemployment model given λ and w and thus c_0 , c_1 , and h , comprises values of θ and y satisfying the wage bargain, equation (32), and the zero-profit condition, equation (34).

B.7 Calibration

Because the model does not consider movements into and out of the labor force and because it omits consideration of the large flows of workers from job to job without intervening unemployment, it cannot be more than a rough approximation to the actual labor market of the United States. I take the job-finding rate to have the following form:

$$\phi(\theta) = \phi_0 \theta^{0.5} \tag{35}$$

I take the exogenous separation rate, s , to be 3 percent per month. I calibrate to 5.5 percent unemployment and a vacancy/unemployment ratio of $\theta = \theta^* = 0.5$ by choosing the matching efficiency parameter ϕ_0 to be 0.729.

I normalize the marginal product of labor w at one. I take the discount factor δ to be $0.95^{1/12}$, and the bargaining parameters to be $\beta_0 = 0.5$, $\kappa = 0.88$, and $\bar{\beta} = 0.7$. The choice of the compensation stabilization parameter κ matches the observed volatility of unemployment. I take the elasticity of recruiting with respect to labor to be $\chi = 0.5$. Finally, I take the unemployment benefit replacement rate to be $b = 0.111$, the average of the values computed since 1960 in Hall (2005b). This figure probably understates the actual replacement rate for the unemployed, because it does not adjust for the different skill composition of the unemployed in relation to the employed. None of the results of the paper would be significantly different at rates of 0.20 or 0.30.

B.8 The flow value of unemployment, z , in the MP model

In Mortensen and Pissarides (1994), preferences appear in an abbreviated form: Workers attribute a flow value, z , to time spent not working. This flow value includes unemployment benefits and the benefit of not having to work. While working, workers enjoy a flow value equal to their compensation, w . In my framework, I include more than compensation in the flow value, f_1 , for working—see equation (27)—but it is easy to renormalize the flow values to extract the value of z implicit in my version of the MP model. Of course, an important feature of my version is that z is endogenous, rather than being taken as a fixed parameter, as in earlier versions of the MP model. In the case of transitory shocks, where I hold the marginal utility of wealth, λ , constant, z changes because of changes in consumption, hours, compensation, and, possibly, the replacement rate b . These effects are not large. But for permanent shocks, the added change from λ is tremendously important.

To accommodate the normalization in the MP model that equates the flow value of working to earnings, I will add a constant ω to the flow utility. Equations (26) and (27) imply, for employed workers consuming c_1 , working h hours, and earning w ,

$$f_1 = \frac{\omega + U(c_1, h)}{\lambda} - c_1 + w = w. \quad (36)$$

The w on the right-hand side corresponds to the MP normalization that the flow value of employment is earnings. Similarly, the flow value of not working and receiving benefits of bw is

$$f_0 = \frac{\omega + U(c_0, 0)}{\lambda} - c_0 + bw = zmh. \quad (37)$$

I write the flow value as zmh because the value of z is almost always stated as a fraction of flow productivity, mh , often normalized at one. The solution is

$$z = \frac{bw + c_1 - c_0 + \frac{U(c_0, 0) - U(c_1, h)}{\lambda}}{mh}. \quad (38)$$

In my baseline calibration, I find that consumption levels are $c_0 = 0.82$ and $c_1 = 0.98$ and hours $h = 1.05$, corresponding to utility levels of $U(c_0, 0)/\lambda = -1.03$ and $U(c_1, h)/\lambda = -1.34$. The preference component is

$$\frac{1}{mh} \left(c_1 - c_0 + \frac{U(c_0, 0) - U(c_1, h)}{\lambda} \right) = 0.44 \quad (39)$$

so the total flow value, from benefits and the preference for not working, is $z = 0.55$, somewhat higher than the value of 0.4 in Shimer (2005). Shimer's value corresponds to

a Frisch elasticity slightly below 0.4. I will show shortly that there is little substantive difference between the two values.

These findings shed some light on Hagedorn and Manovskii (2006)'s estimate of z obtained from a radically different calibration strategy. They require their version of the MP model to match the derivative of the wage with respect to productivity and they calibrate to an outside estimate of the cost of posting a vacancy. They show that the criteria imply that the flow value of non-work has the high value of $z = 0.955$ (in units of productivity) and that the bargaining power of labor is weak, with only 5.2 percent of the surplus accruing to the worker. The corresponding Frisch elasticity of labor supply is above 4, far out of the range of any empirical finding.