

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

June 12, 2007

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. About a third 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. Cyclical variations in the employment rate account for more than half the cyclical variation in total labor input, which declines substantially in recessions because of the rise in unemployment. I define and estimate an employment function, analogous to the supply function for hours per worker. My work differs from previous attempts to place cyclical movements of total hours on a labor supply curve 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 many seminars, 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 Max Flötotto and Felix Reichling for comments. A file containing the calculations is 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.

Figure 1 shows first differences of log nondurables consumption per person, weekly hours per worker, the employment rate (fraction of the labor force working in a given week, one minus the unemployment rate), and the average product of labor for the United States since 1949. Common movements associated with the business cycle are prominent in all four measures. Consumption, hours, and employment are fairly well correlated with each other, while their correlation with productivity is lower, especially in the last 15 years of the sample.

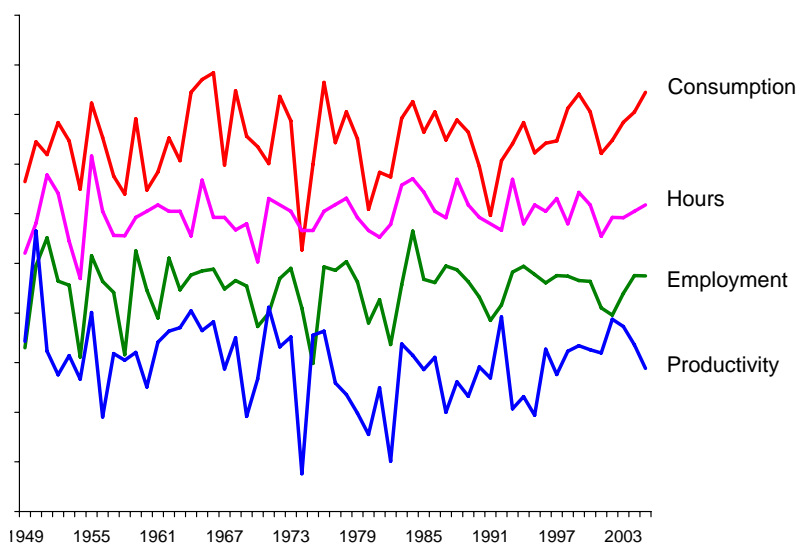


Figure 1. Growth in Consumption, Hours per Worker, Employment Rate, and Productivity

Note: The tick marks on the vertical axis are one percentage point apart. Constants are added to the series to separate them vertically.

I take the driving force of the movements shown in the figure to be changes in the marginal product of labor, arising from random changes in total factor productivity growth, in the terms of trade, and in the prices of factors other than labor. I portray the movements of hours per worker in terms of a standard labor supply schedule without extreme wage elasticity.

Understanding the cyclical movements of consumption in this framework is a challenge. With preferences additively separable in work and consumption, it is difficult to construct a model on standard principles that generates a strong hours response—as seen in Figure 1—and a strong pro-cyclical consumption response. The approach I take is to invoke fairly high complementarity between consumption and hours of work. High marginal productivity induces households to substitute purchased consumption goods and services to replace the diminished time at home resulting from longer hours of work.

The second big challenge is to understand the movements of the employment rate in this framework. I do so by making job search an integral part of the model and a distinct use of peoples' time. In this area, the model draws heavily on the Mortensen and Pissarides (1994) theory of equilibrium unemployment. I develop an employment function that is in some ways analogous to an hours supply function. But it does not depend solely on choices made by workers. That is, job search is not just a use of time determined by individual choice in response to a market wage. Rather, it is an equilibrium of choices made by jobseekers and by recruiting employers.

In broad summary, the model in this paper considers a worker who maximizes the expected 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, $c(\lambda, w)$ describes the corresponding choice of consumption. I view the increases in consumption that occur when w is unusually high (that is, relative to λ) as resulting from the positive response of consumption to w through the consumption-work complementarity.

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 an increasing function of both λ and w . In Appendix B, I define a broader class of models that differ by the principle governing the compensation paid to newly hired workers. Members of this class can yield much higher responses of unemployment to the two driving forces than is present in the MP model, but unemployment remains a function of the two driving forces alone.

In this paper, I do not consider the small procyclical movements of participation in the labor force—Hall (2007) documents these movements. The function

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

governs the total volatility of hours of work, apart from participation. 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.

I treat the two key driving forces λ and w as unobserved latent variables. I take each of the four indicators—consumption, hours, the employment rate, and productivity—as a function of the two latent variables plus an idiosyncratic residual. The model falls short of identification. I use information from extensive research on some of the coefficients to help identify the remaining coefficients. I also use inequalities derived from the model to limit the ranges of the coefficients. I embody the information in a prior distribution and compute the posterior distribution of the parameters from the prior and the sample evidence shown in Figure 1.

The posterior distribution shows that the empirical employment function is much more sensitive to λ and w than is its counterpart in the MP model. Although I treat this as an empirical finding not associated with any specific theory of compensation determination, it implies that compensation paid to newly hired workers is stickier than it would be with Nash bargaining—unemployment rises in recessions because the marginal product of labor falls relative to the compensation paid to newly hired workers, so employers cut their job-creation efforts.

The model provides an internally consistent account of cyclical movements in the labor market. It attains the goal of explaining the large observed cyclical volatility of labor input without invoking an unrealistically high elasticity of labor supply. The main way that it attains the goal is to explain the movements of unemployment as responses to the two driving forces. Because most of the decline in labor input that occurs in a recession takes the form of rising unemployment rather than reduced hours of those at work, the shift in emphasis from the elasticity of labor supply to the elasticity of unemployment is appropriate.

The Mortensen-Pissarides class of employment models imposes only loose constraints on the behavior of compensation paid to workers after employment. Because observed data on compensation are dominated by the behavior of the compensation of incumbents rather than the newly hired, I do not include compensation in the basic model of the paper. At the end, I examine the behavior of compensation in relation to the marginal product series derived from the model. Although the correlation of the rates of growth of real compensation and the marginal product is 0.63, there is no obvious pattern in the movements of real compensation. I also examine nominal compensation—its rate of growth appears somewhat smoothed in years of high and variable inflation.

2 Insurance

The analysis in this paper makes the assumption that workers are 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 sys-

tem of consumption demand and labor supply the ideal analytical framework. Absent the assumption, the model is an approximation based on aggregating employed and unemployed individuals, each with a personal state variable, wealth.

I do not believe that, in the U.S. economy, consumption during unemployment behaves literally according to the model with full insurance against unemployment risk. 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 for the response of workers to unemployment. I make no claim that workers are insured against permanent changes in their earnings capacities, only that the transitory effects of unemployment can usefully be analyzed under the assumption of insurance.

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’s (1985) 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.

Consider the standard intertemporal consumption-hours problem,

$$\max E_t \sum_{\tau=t}^T \delta^{\tau-t} U(c_\tau, h_\tau) \quad (2)$$

subject to the budget constraint,

$$\sum_{\tau=t}^T (w_\tau h_\tau - p_\tau c_\tau) = 0. \quad (3)$$

Here w_τ is the wage, taken for now as the slope of a personal budget constraint, and p_τ is the price of the consumption good. Both the wage and the price are quoted, for now, in units of abstract purchasing power, as of time t .

I let $c(\lambda p, \lambda w)$ be the Frisch consumption demand and $h(\lambda p, \lambda w)$ be the Frisch labor supply. See Browning, Deaton and Irish (1985) for a complete discussion of Frisch systems

in general. They satisfy, for consumption and hours at time zero,

$$U_c(c(\lambda p_0, \lambda w_0), h(\lambda p_0, \lambda w_0)) = \lambda p_0 \quad (4)$$

and

$$U_h(c(\lambda p_0, \lambda w_0), h(\lambda p_0, \lambda w_0)) = -\lambda w_0 \quad (5)$$

Here λ is the Lagrange multiplier for the budget constraint. Consumption in period t is $c(\lambda_t p_t, \lambda_t w_t)$ and similarly for hours. I will focus on time t and drop the time subscript in what follows.

The Frisch demands are symmetric: $c_2 = -h_1$. They have three basic first-order or slope properties:

- *Intertemporal substitution in consumption*, $c_1(\lambda p, \lambda w)$, the response of consumption to changes in its price
- *Frisch labor-supply response*, $h_2(\lambda p, \lambda w)$, the response of hours to changes in the wage
- *Consumption-hours cross effect*, $c_2(\lambda p, \lambda w)$, 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. In addition, in the presence of uncertainty, the curvature of U controls risk aversion, the subject of another literature.

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. People consume more when wages are high because they work more and consume less leisure. Browning et al. (1985) show that the Hessian matrix of the Frisch demand equations is negative semi-definite. Consequently, the derivatives satisfy the following constraint on the cross effect controlling the strength of the complementarity:

$$c_2^2 \leq -c_1 h_2. \quad (6)$$

To understand the three basic properties of consumer-worker behavior listed earlier, I 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 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.

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 -0.5 , a Frisch elasticity of labor supply of 0.7 , and a Frisch cross-elasticity of 0.3 . I use informative priors for these parameters. I use much less informative priors for parameters that have received less attention in past research—the elasticities of the employment function with respect to λ and w , the variances of the stochastic elements, and the correlation of λ and w .

3.1 Normalization of the price

In the rest of the paper, I normalize the price as $p_t = 1$. Thus in period t , values are stated in terms of units of period- t output. Further, λ_t becomes marginal utility in period t under this normalization and the Frisch functions are $c(\lambda, \lambda w)$ and $h(\lambda, \lambda w)$. Notice that the response of consumption to a change in marginal utility λ is:

$$\frac{\partial c}{\partial \lambda} = c_1 + w c_2 \tag{7}$$

and for hours:

$$\frac{\partial h}{\partial \lambda} = -c_2 + wh_2. \quad (8)$$

3.2 Supply of hours per worker

The supply function for hours, $h(\lambda, w)$, solves equations (4) and (5) 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.

3.3 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.

4 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 a feature of labor supply alone—it describes choices of firms as well as workers.

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

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.

Second, *search-and-matching models*—surveyed recently by Rogerson, Shimer and Wright (2005)—divide the labor market into many sub-markets, 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. 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 separability between hours and consumption and with consumption entering as the log.

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. Appendix B demonstrates 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 shifts 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.

To interpret the empirical employment function, I start with the MP 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. 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. I am agnostic about the principles that govern the compensation bargain. I allow the bargaining powers of the parties to depend on any endogenous variable in the MP model. Models in this class imply a wide range of volatilities of unemployment. Under Nash bargaining, volatility is low. With sticky wages, less responsive to current conditions, volatility is higher, possibly extremely high. I exclude some potentially interesting cases, notably bargaining theories (yet to be developed) that imply an endogenous state variable that imparts inertia to the wage.

In the extended MP model, 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.

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

over time.

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.

5 Latent Factor Model

Because consumption is conspicuously nonstationary and the other two variables are somewhat non-stationary, I work in first differences of logs, that is, rates of growth. I approximate the Frisch consumption demand, Frisch hours supply, and the employment function as log-linear, with $\beta_{c,c}$ denoting the elasticity of consumption with respect to its own price (the elasticity corresponding to the partial derivative c_1 in the earlier discussion), $\beta_{c,h}$ the cross-elasticity of consumption demand and hours supply, and $\beta_{h,h}$ the own-elasticity of hours supply. I further let $\beta_{n,\lambda}$ denote the elasticity of employment with respect to marginal utility λ and $\beta_{n,w}$ the elasticity with respect to the marginal product w .

5.1 Hours and employment

The factor equation for hours is:

$$\Delta \log h = (-\beta_{c,h} + \beta_{h,h})\Delta \log \lambda + \beta_{h,h}\Delta \log w + \epsilon_h \quad (9)$$

and for employment is:

$$\Delta \log n = \beta_{n,\lambda}\Delta \log \lambda + \beta_{n,w}\Delta \log w + \epsilon_n. \quad (10)$$

Here $\beta_{h,h}$ and $-\beta_{c,h}$ are the Frisch own- and cross-elasticities of hours supply for employed workers, $\beta_{n,\lambda}$ and $\beta_{n,w}$ are the elasticities of the employment function, and the ϵ s are idiosyncratic random components.

5.2 Consumption

I disaggregate the population by the employed and unemployed, who consume c_e and c_u respectively. The employed satisfy the two conditions

$$U_c(c_e, h) = \lambda \quad (11)$$

and

$$U_h(c_e, h) = -\lambda w. \quad (12)$$

The unemployed satisfy only the condition for consumption and are constrained to zero hours of work:

$$U_c(c_u, 0) = \lambda. \quad (13)$$

Under the assumption of full insurance against idiosyncratic personal shocks, the employed and unemployed have the same marginal utility λ .

Only average consumption c is observed. I hypothesize approximate constancy of the ratio c_u/c_e which I designate as $1 - \eta$. With separable preferences, the ratio is one and $\eta = 0$. With consumption-hours complementarity, η is positive. Observed consumption is the average of the two levels, weighted by the employment and unemployment fractions:

$$c = nc_e + (1 - n)c_u. \quad (14)$$

Taking first differences of the log-linearization in the variables, around the point $n = \bar{n}$, I find

$$\Delta \log c = \frac{c_e - c_u}{c} \bar{n} \Delta \log n + \bar{n} \frac{c_e}{c} \Delta \log c_e + (1 - \bar{n}) \frac{c_u}{c} \Delta \log c_u. \quad (15)$$

Further,

$$\frac{c_e}{c} = \frac{1}{\bar{n} + (1 - \bar{n})(1 - \eta)}, \quad (16)$$

and

$$\frac{c_u}{c} = \frac{1 - \eta}{\bar{n} + (1 - \bar{n})(1 - \eta)}. \quad (17)$$

The consumption changes relate to latent factors as

$$\Delta \log c_e = (\beta_{c,e} + \beta_{c,h}) \Delta \log \lambda + \beta_{c,h} \Delta \log w \quad (18)$$

and

$$\Delta \log c_u = \beta_{c,c} \Delta \log \lambda. \quad (19)$$

Substituting equations (18) and (19) into equation (15), I find, now including an idiosyncratic disturbance ϵ_c ,

$$\Delta \log c = \frac{c_e - c_u}{c} \bar{n} \Delta \log n + \beta_{c,c} \Delta \log \lambda + \beta_{c,h} \bar{n} \frac{c_e}{c} (\Delta \log w + \Delta \log \lambda) + \epsilon_c. \quad (20)$$

Finally, I substitute equation (10) for $\Delta \log n$ to get

$$\begin{aligned} \Delta \log c &= \left(\beta_{c,c} + \beta_{c,h} \frac{c_e}{c} \bar{n} + \beta_{n,\lambda} \frac{c_e - c_u}{c} \bar{n} \right) \Delta \log \lambda \\ &+ \left(\beta_{c,h} \frac{c_e}{c} \bar{n} + \beta_{n,w} \frac{c_e - c_u}{c} \bar{n} \right) \Delta \log w + \epsilon_c + \frac{c_e - c_u}{c} \bar{n} \epsilon_n. \end{aligned} \quad (21)$$

5.3 Productivity

I measure productivity as the average product of labor, $m = \frac{q}{h}$, where q is output per worker.

I let α be the elasticity of output with respect to labor input. From

$$w = \frac{\partial q}{\partial h} = \alpha \frac{q}{h}, \quad (22)$$

I get the equation for the log-change in m :

$$\Delta \log m = \Delta \log w - \Delta \log \alpha. \quad (23)$$

Notice that $\Delta \log \alpha = 0$ for a Cobb-Douglas technology. Finally, I define ϵ_m to include $-\Delta \log \alpha$ and any other disturbances, such as measurement error, so the equation for m in the model is

$$\Delta \log m = \Delta \log w + \epsilon_m. \quad (24)$$

5.4 Statistical model

I assume that the idiosyncratic components, ϵ , are uncorrelated with λ and w . This assumption is easiest to rationalize if the ϵ s are measurement errors. In the case of productivity, ϵ_m includes the change in the elasticity of output with respect to labor input. One would expect feedback from this disturbance to the endogenous variables, but the constancy of factor shares in the U.S. economy suggests that the volatility of ϵ_m from this source is quite small.

The model has 12 parameters: the 5 β slope coefficients, the variances and correlation of the latent factors, σ_λ^2 , σ_w^2 , and $\sigma_{\lambda,w}$, and the variances of the four idiosyncratic components, $\sigma_{\epsilon,c}^2$, $\sigma_{\epsilon,h}^2$, $\sigma_{\epsilon,n}^2$, and $\sigma_{\epsilon,m}^2$. The model implies 10 observed moments, the distinct elements of the covariance matrix of the observables, the employment-adjusted log-change in consumption

and the log-changes of hours, employment, and productivity. It is further restricted by non-negativity of the 6 variances and by the Cauchy inequality for the covariance:

$$\sigma_{\lambda,w}^2 \leq \sigma_{\lambda}^2 \sigma_w^2 \quad (25)$$

and by the concavity condition, equation (6).

Under the assumption that the random variables λ , w , ϵ_c , ϵ_h , ϵ_n , and ϵ_m are multivariate normal, any parameter set that matches the sample moments achieves the maximum of the likelihood function. The likelihood has a plateau of equal height for any set of parameters with this property. The posterior distribution is governed by the prior everywhere on the plateau. Stripped of an inessential constant, the log-likelihood function is

$$-\frac{T}{2} \left[\log \det \Omega + \text{tr} \left(\Omega^{-1} \hat{\Omega} \right) \right]. \quad (26)$$

Ω is the covariance matrix of the observables implied by the model and $\hat{\Omega}$ is the sample covariance matrix. On the plateau, $\Omega = \hat{\Omega}$ and the value of the log-likelihood is

$$-\frac{T}{2} \left(\log \det \hat{\Omega} + 4 \right). \quad (27)$$

The prior distribution is discrete. It takes the 12 parameters to be independent of one another. The marginal distribution of each parameter takes on equal values at four equally spaced points. Thus the posterior distribution is defined on a lattice of $4^{12} = 16.8$ million points. I calculate the marginals of the posterior distribution by enumeration of all of these points. Although this approach is usually inefficient in comparison to one based on the values of the posterior, it is so cheap to evaluate the likelihood here that brute force is the preferred approach. It is better suited to an underidentified model like the one at hand than to a model where the sample evidence dominates the prior.

5.5 Inferring the Values of λ and w

I write the model in matrix form as

$$x = \theta_{\lambda} \Delta \log \lambda + \theta_w \Delta \log w + \epsilon. \quad (28)$$

Here x is the vector of observed values of the log-changes of consumption, hours, employment, and productivity. I infer λ as a linear combination, $\hat{\lambda} = a'x$. I choose the weights a as the coefficients of the projection of λ on x , using the moments implied by the parameter values at the posterior mean. I calculate the inference of w , \hat{w} , similarly.

6 Prior Distributions

Table 1 shows the marginal prior distributions I use for the parameters. They are four-point distributions for all parameters. The priors are highly informative when drawn from the research summarized in Appendix A. They are less informative for parameters where earlier work is either sparse or nonexistent, for the variances of the random elements, and for the correlation of $\Delta \log \lambda$ and $\Delta \log w$. I constrain the cross-elasticity $\beta_{c,h}$ to satisfy concavity and the correlation of the latent factors to be greater than -1 .

The parameter η controlling the ratio of unemployment consumption c_u to employment consumption c_e reflects the same properties of preferences as does the Frisch cross-elasticity, $\beta_{c,w}$. Accordingly, I take the joint prior for the two parameters to have perfect correlation, with $\eta = 0.75\beta_{c,w}$. The proportionality factor 0.75 is derived from a parametric utility function that matches the means of the priors of the Frisch elasticities—when the cross-elasticity is 0.20, the consumption ratio is 0.85.

7 Data

To avoid complexities from durables purchases and measurement error in the consumption of services, I use nondurables consumption as an indicator of consumption. I take the quantity index for nondurables consumption from Table 1.1.3 of the U.S. National Income and Product Accounts and population from Table 2.1. I take weekly hours per worker from series LNU02033120, Bureau of Labor Statistics, Current Population Survey, and the unemployment rate from series LNS14000000. For further discussion of the labor-market data, see Hall (2007). I take output per hour of work during 1948 through 1986 as the ratio of real value-added output to labor input in the historical multifactor productivity data, bls.gov/mfp/historicalsic.htm. For 1987 through 2005, I use output per hour of all persons, private business, series MPU740021, bls.gov/mfp.

Table 2 shows the covariance and correlation matrixes of the log-differences of the four series. The most important fact in the table is the positive correlation of consumption with both hours and employment. Consumption is quite pro-cyclical. This fact is much easier to explain in a model with consumption-hours complementarity. Not surprisingly, hours and

<i>Parameter</i>	<i>Interpretation</i>	<i>Mean</i>	<i>Lowest value</i>	<i>Highest value</i>
$\beta_{c,c}$	Frisch own-price elasticity of consumption	-0.50	-0.6	-0.4
$\beta_{c,h}$	Frisch cross-price elasticity of consumption	0.30	0.0	0.6
$\beta_{h,h}$	Frisch wage elasticity of hours	0.75	0.6	0.9
$\beta_{n,\lambda}$	Elasticity of employment with respect to λ	0.35	0.0	0.7
$\beta_{n,w}$	Elasticity of employment with respect to w	0.95	0.4	1.5
σ_{λ}^2	Variance of latent λ	2.15	0.3	4.0
σ_w^2	Variance of latent w	2.15	0.3	4.0
ρ	Correlation of λ and w	-0.70	-0.9	-0.5
σ_c^2	Variance of consumption noise	1.00	0.5	1.5
σ_h^2	Variance of hours noise	0.30	0.2	0.4
σ_n^2	Variance of employment noise	0.25	0.1	0.4
σ_m^2	Variance of productivity noise	0.75	0.3	1.2

Table 1. Priors

	<i>Consumption</i>	<i>Hours</i>	<i>Employment</i>	<i>Productivity</i>
Covariances				
Consumption	2.08	0.54	1.03	0.96
Hours		0.76	0.63	0.13
Employment			1.26	0.28
Productivity				2.60
Correlations				
Consumption	1.000	0.511	0.702	0.414
Hours		1.000	0.645	0.093
Employment			1.000	0.157
Productivity				1.000

Table 2. Covariances and Correlations of Log-First Differences of Consumption, Hours, Employment, and Productivity

employment are quite positively correlated. Consumption has surprisingly high volatility, a property not explained in this paper. Consumption also has by far the highest correlation with productivity.

The variance of employment is about 70 percent higher than the variance of hours—the most important source for the added total hours of work in an expansion is the reduction in unemployment. Hours and employment are not very correlated with productivity.

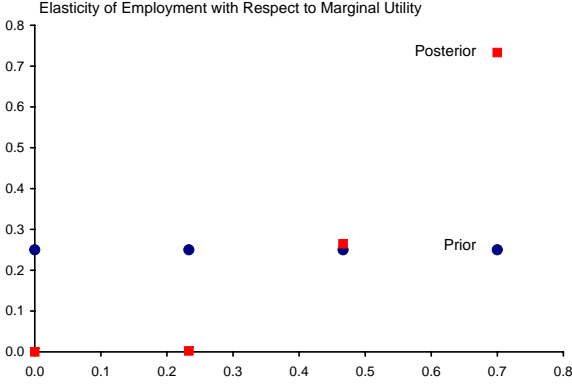
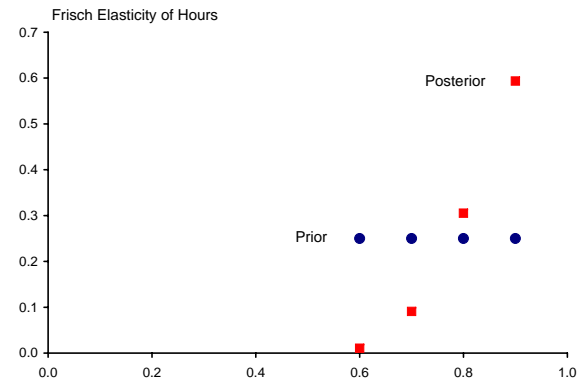
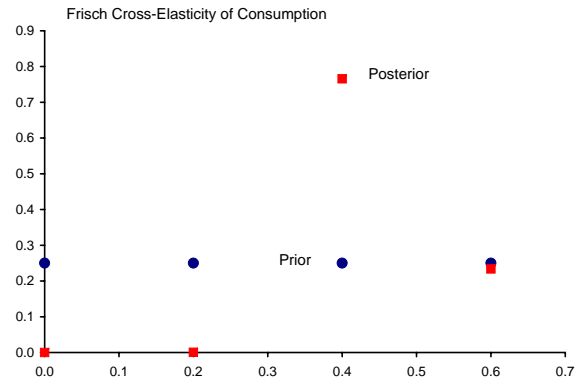
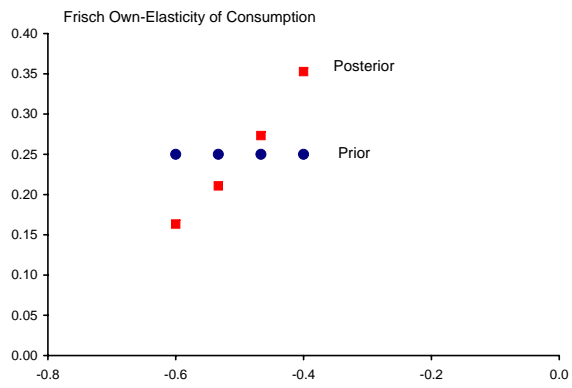
8 Results

Table 3 shows the means and standard deviations of the marginal posterior distributions of the 12 parameters of the model and Figure 2 plots the marginal distributions of the elasticities. For two key parameters, the two own Frisch elasticities of consumption and hours supply, the prior is highly informative, as it is based on a large body of existing research. For both of those parameters, the posterior mean is virtually the same as the prior mean and the posterior standard deviation is small. The sample evidence is not contributing much to knowledge of those parameters. Rather, priors derived from past research help identify other parameters in the model.

Figure 2 shows some hints about how the behavior of aggregate variables suggests pa-

<i>Parameter</i>	<i>Interpretation</i>	<i>Prior mean</i>	<i>Posterior mean</i>	<i>Posterior standard deviation</i>
$\beta_{c,c}$	Frisch own-price elasticity of consumption	-0.50	-0.48	0.07
$\beta_{c,h}$	Frisch cross-price elasticity of consumption	0.30	0.45	0.08
$\beta_{h,h}$	Frisch wage elasticity of hours	0.75	0.85	0.07
$\beta_{n,\lambda}$	Elasticity of employment with respect to λ	0.35	0.64	0.10
$\beta_{n,w}$	Elasticity of employment with respect to w	0.95	1.25	0.19
σ_{λ}^2	Variance of latent λ	2.15	3.75	0.53
σ_w^2	Variance of latent w	2.15	1.50	0.22
ρ	Correlation of λ and w	-0.70	-0.73	0.09
σ_c^2	Variance of consumption noise	1.00	1.14	0.23
σ_h^2	Variance of hours noise	0.30	0.34	0.06
σ_n^2	Variance of employment noise	0.25	0.29	0.10
σ_m^2	Variance of productivity noise	0.75	1.14	0.14

Table 3. Posterior Distribution



	λ	w
Consumption	0.16	0.83
Hours	0.40	0.85
Employment	0.64	1.25
Average product of labor	0.00	1.00

Table 4. Coefficients for Log-First Differences of Consumption, Hours, Employment, and Productivity on λ and w

parameter values that are different from those found in research on micro data. For the intertemporal elasticity of substitution/reciprocal of the coefficient of relative risk aversion in the upper left, the aggregate data seem to point to a lower value. For the cross-elasticity of consumption in the upper right, the aggregate data suggest a high value within the broad range in the prior. This finding may reflect feedback from the idiosyncratic component of consumption—that is, a failure of the assumption that the component is uncorrelated with the λ and w . For the Frisch elasticity of hours supply in the middle right of the figure, the aggregate data prefer a value at the high end of the range that is plausible from research in micro data. This result confirms the long-standing finding that aggregate behavior seems to involve higher elasticity of labor supply than is found at the household level. Here I constrain the elasticity far below the level implicit in most earlier aggregate work, however.

Table 4 shows the coefficients relating the observed variables to the latent variables λ and w at the posterior means of the parameters. The coefficients for employment and for the response of hours to w are the elasticities reported in Table 3 and those for productivity are zero on λ and one on w . The more complicated relations are for consumption and for the response of hours to λ , from equations (9) and (15).

The biggest surprise in Table 4 is the positive response of consumption to marginal utility λ . Although one might think that marginal utility is a declining function of consumption, theory does not require that property in a Frisch demand system. Recall from equation (15) that the coefficient on λ in the consumption equation is

$$\beta_{c,c} + \beta_{c,h} \frac{c_e}{c} \bar{n} + \beta_{n,\lambda} \frac{c_e - c_u}{c} \bar{n}. \quad (29)$$

The theoretical limit on the complementarity effect is, from equation (6),

$$\beta_{c,h} \leq \sqrt{-\beta_{c,c}\beta_{n,w}}. \quad (30)$$

From Table 3, the cross-elasticity is 0.45 while the square root is 0.64, comfortably larger. The key point is that the coefficient of consumption on λ is *not* the own-price effect, which is necessarily negative, but the own-price plus the cross-price effect, which can be positive if complementarity is strong enough. Because of the aggregation of consumption across workers and the unemployed, the complementarity effect has two components in equation (29). First, a higher λ (lower well-being) raises the consumption of workers through the direct effect of the complementarity, controlled by $\beta_{c,h}$. Second, a higher λ increases the employment rate. Because the employed consume more than the unemployed, average consumption rises on this account as well. The second effect is controlled by $\beta_{n,\lambda}$, whose posterior mean is 0.64.

Complementarity also explains the high response of consumption to the current marginal product of labor, w . Again from equation (15), this response is

$$\beta_{c,h} \frac{c_e}{c} \bar{n} + \beta_{n,w} \frac{c_e - c_u}{c} \bar{n}. \quad (31)$$

The second term describes the stimulus to employment (decline in unemployment) that accompanies an increase in w . The direct effect through $\beta_{c,h}$ is 0.45. The effect from employment change is $\beta_{n,w} = 1.25$ multiplied by the consumption-difference effect, which is 0.32.

The effect of λ on hours is correspondingly weak. The coefficient is $-\beta_{c,h} + \beta_{h,h}$. Complementarity enters negatively, offsetting the relatively strong own-elasticity effect. An increase in λ raises the price of consumption as it raises the reward to work. Because non-work time is a substitute for consumption, people shift toward non-work when the price of consumption rises.

The sample evidence is fairly influential in the posterior for the cross-elasticity $\beta_{c,h}$. Figure 2 shows that the distribution reaches quite a sharp peak at 0.4, despite the prior mean of only 0.30. The role of the cross-elasticity appears to be even more important than has been considered in the limited earlier discussion of models that do not assume separability with zero cross-elasticity. The high value results from two features of the data shown in Table 3—the generally high correlation of consumption with other cyclical variables and the

particularly high correlation, relative to the hours and employment, between consumption and productivity. Recall that productivity reveals the latent marginal product w except for its own noise.

The sample evidence is also influential about the elasticities of employment with respect to λ and w , a subject not previously investigated. The posterior reaches a sharp peak for the λ -elasticity at 0.7 and for the w -elasticity at 1.1. Despite the model’s lack of identification and the uninformative priors placed on these parameters (uniform from 0 to 0.7 for the first and from 0.4 to 1.5 for the second), the other priors combine with the sample evidence to provide useful information.

8.1 Implied values of marginal utility and marginal product

Table 5 shows the coefficients of the projection of the latent factors on the observed variables, at the posterior means of the parameter values. As expected, the inference of marginal utility puts negative weights on consumption and productivity—increases in them signal improvements in well-being and thus lower values of marginal utility, λ . The inference puts a positive weight on employment. The reason is shown in Table 4. An increase in λ raises employment by more than it raises consumption and hours, relative to the coefficients for w . Therefore, on the average, an increase in employment signals that an increase in λ has occurred. The other feature of Table 5 worth noting is that the weight on productivity in the inference of w is 0.41, well below the loading of productivity on w of 1. This finding reflects the noise in productivity. The inference puts weight on all of the variables positively correlated with productivity to filter out as much noise as it can.

Figure 3 shows the estimates of marginal utility, $\Delta \log \lambda$, and the marginal product, $\Delta \log w$, resulting from the application of the coefficients in Table 5 to the data on the four observables. The figure shows a pronounced negative correlation between the changes in marginal utility and in the marginal product of labor. News that raises the current marginal product of labor tends to raise lifetime well-being and thus to lower λ . If the economy were perturbed by a single shock and households had no advance information about the shock, the correlation would be -1 . With multiple shocks and advance information, the correlation would be less negative, in accord with the estimated correlation of -0.73 .

	<i>Inferred λ</i>	<i>Inferred w</i>
Consumption	-0.36	0.23
Hours	0.19	0.16
Employment	0.65	0.13
Average product of labor	-0.76	0.41

Table 5. Coefficients for inference of λ and w from Log-First Differences of Consumption, Hours, Employment, and Productivity

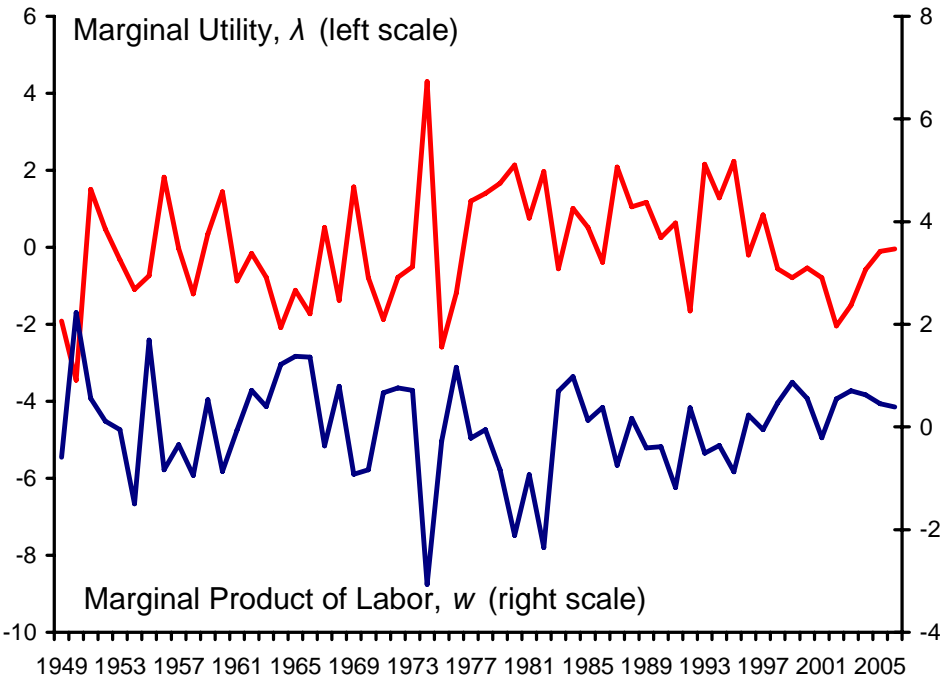


Figure 3. Inferred Growth of Marginal Utility, λ , and Marginal Product of Labor, w , Percentage Points

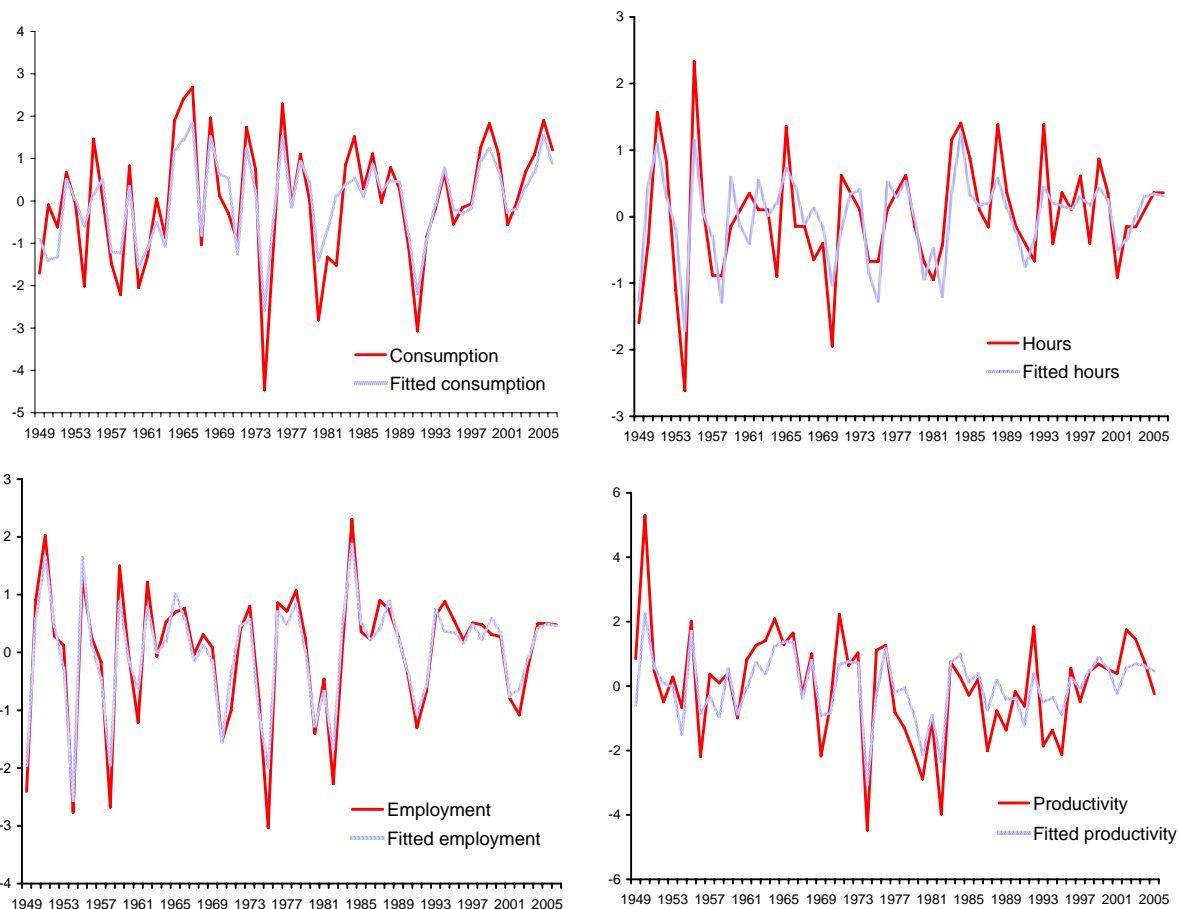


Figure 4. Actual and Fitted Values of the Four Observables

8.2 Fitted values for observables

Given the time series for λ and w , I can calculate the implied fitted values for the four observables. These are shown in Figure 4. The two-factor setup is highly successful in accounting for the observed movements of all four variables. Little is left to the idiosyncratic disturbances. Of course, two factors are likely to be able to account for most of the movement of four macro time series, especially when two of them, hours and employment, are fairly highly correlated. But the choices of the factors and the factor loadings are not made, as in principal components, to provide the best match. The loadings are influenced by the priors drawn from earlier research. The success of the model is not so much the good fit shown in Figure 4, but rather achieving the good fit with coefficients that satisfy economic reasonability.

9 Interpretation

The results in the previous section achieve the main goal of the paper—to show that standard economic principles embodied in the Frisch consumption demand and hours supply, together with a model of unemployment in the extended Mortensen-Pissarides class, can account for the higher-frequency movements of those variables. The accounting does not rest on implausible values of any parameters. Most importantly, it does not rest on exaggerated ideas about the elasticity of hours supply. Because the business cycle dominates the higher-frequency movements of the variables, the results give a coherent account of the business cycle.

The main way that the model escapes reliance on unrealistic elasticity of hour supply is to recognize that the primary dimension of fluctuations in labor input at higher frequencies is in the employment rate. Research on labor supply in household data does not reveal the elasticities of employment, which are not features of household choice alone, but reflect an equilibrium involving employer actions as well. This paper is the first to provide estimates of an employment function of the type implied by the MP model, though one could interpret Hagedorn and Manovskii (2006) in terms of its implications for the employment function. However, their approach rests on preferences that imply an extremely high wage elasticity of hours supply.

Thus the centerpiece of the account in this paper of movements in labor input over the cycle is the high elasticity of employment with respect to the marginal product of labor. Employment falls and unemployment rises in a contraction because w falls and the elasticity of employment with respect to w is something like 1.2. A rise in marginal utility offsets some of the decline in w in the typical recession, but its elasticity is only around 0.5.

Unlike most of the literature on the “Shimer puzzle,” I do not take a stand on the source of the high elasticity of the employment function with respect to w . In particular, I do not sponsor any particular bargaining principle in place of the Nash bargain. I take a purely empirical approach to the measurement of the elasticity. In a model that follows Mortensen and Pissarides in every respect except bargaining, my results imply that bargaining power shifts toward workers during recessions, or, to put it differently, that wages are sticky. But my finding is also consistent with other mechanisms, such as on the job search, that raise

the elasticity of employment with respect to w .

10 Compensation

The factor model does not consider the actual value of compensation paid to workers, despite the key role of compensation in the Mortensen-Pissarides class of employment models. In that class of models, 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 hiring, governs the employer’s vacancy-creation efforts. The class of models has no implications about the pattern of payment of compensation over the period of employment. 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. In particular, nothing rules out smoothing of compensation in relation to productivity. I am not aware of any way to introduce observed compensation, averaged over workers hired over the past 40 years, into the factor model without making special assumptions about the determination of compensation during the period of employment. Even if compensation is the result of period-by-period bargaining, one would have to take a stand on bargaining principles to pin down compensation.

To investigate this issue, I use data on compensation from the historical multifactor productivity measures of the BLS, bls.gov/pub/special.requests/opt/mp/prod3.mfptablehis.zip. This source reports nominal hourly compensation of private workers, with extensive adjustments for worker quality. I deflated that series with the personal consumption price index, Table 1.1.4 in the National Income and Product Accounts.

The correlation of the log-changes in real hourly compensation with the inferred value of w is 0.63. The variance of compensation growth is substantially higher than the variance of the inferred w and is equal to the variance of the measured growth rate of output per worker. This finding supports the view that compensation does not track w year by year. The typical duration of employment is sufficiently long to disconnect one year’s compensation from that year’s marginal product. About 82 percent of workers remain with the same employer from

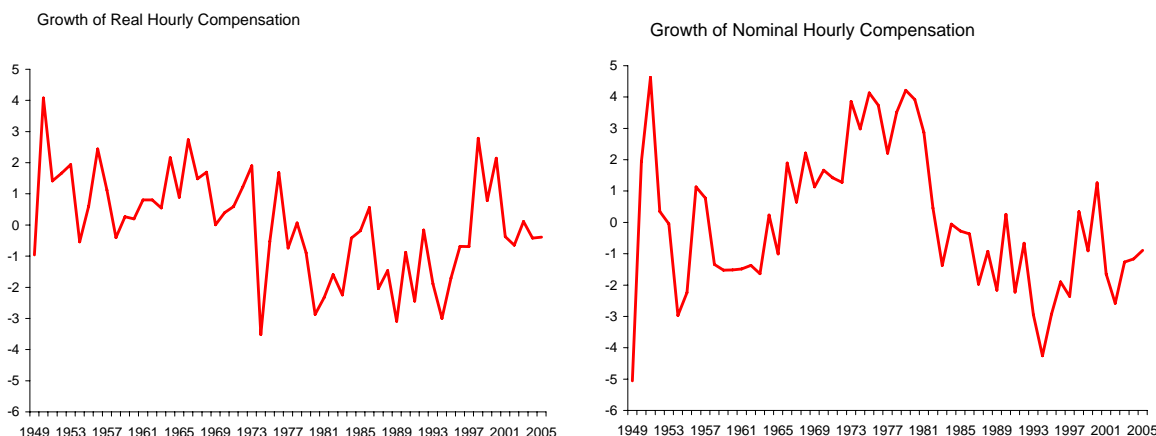


Figure 5. Real and Nominal Compensation Growth

one year to the next and most work occurs during jobs that last a decade or more—see Hall (1982). Observed compensation is not the anticipation for newly hired workers, but is the average paid to workers employed at the time, many hired years earlier.

Figure 5 shows, in the left graph, growth rates of real compensation per hour. No particular pattern, cyclical or otherwise, is apparent in the data. The right graph shows the growth of nominal compensation per hour. This graph gives at least a hint that compensation may be smoothed, during some periods, in nominal terms. From 1958 to 1980, nominal growth rose gradually, without the sawtooth movements observed in w or in real compensation. The period 1983 to 1990 also saw a period of relatively stable nominal consumption growth. Since 1990, nominal growth has been no more stable than real growth, because inflation has been held in a tight band, so the distinction between nominal and real is no longer important.

11 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 0.85, at the upper end of the range found in recent research using household data, does the job.

About a third of the volatility of cyclical fluctuations in hours per person takes the form of volatility of hours of job-holders. I argue that movements in the marginal product of labor

and in the marginal utility of consumption are plausible sources of the movements of hours. These are the arguments of the Frisch hours supply function.

The remaining larger part of cyclical fluctuations in labor input per person comes from unemployment. Labor input declines 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. An extended version of the Mortensen-Pissarides model makes unemployment depend on just these two variables. Further work on the employment function, either in the framework of the extended MP model or outside that framework, is clearly in order.

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.

- Browning, Martin and Thomas F. Crossley, “Unemployment Insurance Benefit Levels and Consumption Changes,” *Journal of Public Economics*, 2001, 80, pp. 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), pp. 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 New Method of Estimating Risk Aversion,” *American Economic Review*, December 2006, 96 (5), pp. 1821–1834.
- 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*, October 2006, 53 (7), pp. 1451–1472.
- 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., “The Importance of Lifetime Jobs in the U.S. Economy,” *American Economic Review*, September 1982, 72 (4), 716–724.
- , “Complete Markets as an Approximation to the Bewley Equilibrium with Unemployment Risk,” 2006. Hoover Institution, Stanford University, stanford.edu/~rehall.
- , “Cyclical Movements along the Labor Supply Function,” April 2007. Stanford.edu/~rehall.
- 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.
- 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 in Labor Market Equilibrium when Workers can Self-Insure,” November 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 coefficient of relative risk aversion (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's (1991) 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

The Mortensen-Pissarides model has four major components: (1) a search and matching technology, (2) valuations for the various states that workers and firms experience, (3) an equilibrium condition for employers' recruitment effort, and (4) a bargaining or other principle for the determination of workers' compensation. The extended MP model goes beyond the canonical model of Mortensen and Pissarides (1994) in two respects—evaluations of workers' states are based on concave rather than linear utility with a choice of hours of work and the principle for compensation determination is not necessarily the Nash bargain.

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:

$$U_c(c_0, 0) = \lambda. \tag{32}$$

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

$$U_c(c_1, h) = \lambda \tag{33}$$

and

$$-U_h(c_1, h) = \lambda w. \tag{34}$$

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

equilibrium

$$u = \frac{s}{s + \phi(\theta)}. \quad (35)$$

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

$$f_0 = \frac{U(c_0, 0)}{\lambda} - c_0 + by, \quad (36)$$

for the job-seeker and

$$f_1 = \frac{U(c_1, h)}{\lambda} - c_1 + y, \quad (37)$$

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

$$f_2 = -k, \quad (38)$$

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

$$f_3 = wh - y, \quad (39)$$

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

Equations (36) through (39) 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:

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

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

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

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

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

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

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

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

The matrix Bellman equation of the model is

$$V = f(y, \theta) + \delta\pi(\theta)V, \tag{40}$$

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

$$V(y, \theta) = (I - \delta\pi(\theta))^{-1} f(w, \theta). \tag{41}$$

B.4 Free entry

The MP model assumes 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:

$$V_2(y, \theta) = 0. \tag{42}$$

B.5 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 $V_3 - V_2$.

The parties choose compensation by bargaining or otherwise. I assume that the choice depends on the endogenous variables of the MP model and not on other variables. A fixed wage is one example of a compensation principle. The Nash bargain with bargaining weight β is another, described by the equation

$$\beta(V_3 - V_2) = (1 - \beta)(V_1 - V_0). \quad (43)$$

A convex combination of a fixed wage and the Nash wage is another member. In that case, the weight given to the fixed wage is a parameter of wage stickiness. With a stickier wage, the extended MP model yields higher unemployment volatility in response to shifts in w . Yet another admissible compensation principle links the Nash bargaining weight to labor-market tightness θ . If the bargaining weight moves in favor of the worker in a slacker market with lower θ , the model yields higher unemployment volatility.

The leading example of a compensation principle that lies outside the class of extended MP models as I have defined the class is one with a state variable capturing compensation inertia. No model to date has derived such a compensation principle from bargaining or other theory.

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, the zero-profit condition, equation (42), and the compensation principle.

B.7 Proposition

Unemployment and the other endogenous variables of the extended MP model depend only on λ and w and the parameters of the model.