

Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor *

Robert E. Hall
Hoover Institution and Department of Economics,
Stanford University
National Bureau of Economic Research
rehall@stanford.edu; website: Google “Bob Hall”

January 23, 2009

Abstract

Recessions appear to be times when the marginal rate of substitution between goods and workers’ time—the marginal value of time—falls below the marginal product of labor. If so, the allocation of workers’ time is inefficient. I develop a model of households and production that reconciles cyclical movements in the marginal value of time and the marginal product of labor. The household part of the model takes microeconomic research on household data seriously. The model embodies the findings of this research that the Frisch elasticity of labor supply is less than one, that the intertemporal elasticity of substitution in consumption is less than one, and that consumption and hours of work are complements. The most important element of the reconciliation is an explicit treatment of unemployment in a search-and-matching setup with non-Nash bargaining. In the model, recessions do not result in private, bilateral inefficiency in the allocation of labor, but the unemployment rate may be socially inefficiently high.

*This research is part of the program on Economic Fluctuations and Growth of the NBER. I thank the editor and referees, numerous participants in seminars, the Montreal Conference on Advances in Matching Models, the Rogerson-Shimer-Wright group at the 2006 and 2007 NBER Summer Institutes, and the 2006 Minnesota Workshop on Macroeconomic Theory, and to Susanto Basu, Max Flötotto, Felix Reichling, and Harald Uhlig for comments. An earlier version had the title, “Sources and Mechanisms of Cyclical Fluctuations in the Labor Market.” A file containing the calculations is available at my website.

1 Introduction

A fundamental efficiency condition holds that workers' marginal rate of substitution between goods and time spent working—the marginal value of time—should equal the marginal product of labor. Recessions appear to be times when the economy fails this condition—the marginal value of time falls dramatically, at least for the unemployed, while the marginal product of labor appears to remain close to its normal level. The central question about recessions, then, is why are private actions unable to restore the efficient use of workers' time?

The main goal of this paper is to show that standard economic principles embodied in the Frisch consumption-demand and hours-supply functions, with elasticities consonant with research based on household data, together with a model of unemployment in the extended Mortensen-Pissarides class, can account for diminished hours and employment in recessions without invoking private bilateral inefficiency. Workers in immediate contact with employers—those bargaining with prospective employers, setting hours of work in the course of employment, and facing potential layoff—achieve private efficiency. If unemployment is high, the unemployed, participating in a search process with externalities across millions of searchers and employers, may be inefficiently allocated, but this social inefficiency cannot be corrected by private bilateral action.

Figure 1 shows the data whose joint movements I seek to understand. The data are detrended to focus on the cyclical movements. The series are nondurables and services 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. Common movements associated with the business cycle are prominent in all four measures. Consumption and productivity are fairly well correlated with each other and so are hours and employment. The correlation of productivity with hours and employment is lower, especially in the last 15 years of the sample.

The paper builds a dynamic model of households interacting with employers. The model implies a factor structure in which the four variables shown in the figure depend on two factors: the marginal utility of consumption and the marginal product of labor. I compute values for the factors and corresponding values for the fitted values of the four variables. I interpret the differences between the actual and fitted values as disturbances; they turn out to be quite small. The disturbance for the marginal product of labor is the difference




Figure 1: Detrended 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.

between the marginal value of time and the marginal product of labor. Its small size leads to the conclusion that private bilateral inefficiencies are small.

The model in this paper considers a worker in a representative family that maximizes the expected discounted sum of future utility. The family’s marginal utility of goods consumption, set at the same level for all members, describes the long-run or permanent level of well-being in the economy. The marginal product of labor describes the immediate payoff to work. When the marginal product rises, hours will rise as workers take advantage of the added benefit of working. The family takes friction in the labor market as given—though the family would allocate all of its eligible members to working if it could, in fact only a fraction of them are working at any moment because the remainder are searching for work. A key element of the model is a function that relates the employment rate to the same variables that control the consumption and hours choices of the family. Even though the family takes the employment rate as a given feature of the labor market, the employment rate resulting from the interaction of all workers and firms depends on the marginal utility of goods con-

sumption and the marginal product of labor. Thus the model describes consumption, hours, and the employment rate in terms of just these two factors.

The class of models that generate the employment functions of the type considered here includes the Mortensen and Pissarides (1994) (MP) model, the canon of the theory of the employment rate widely in use today. Other members of the class of models yield much higher responses of the employment rate to the two variables than does the MP model, but employment remains a function of the two variables alone. The modification of the original MP model that allows the model to match the cyclical volatility of employment is the replacement of MP's Nash wage bargain by a generalization that allows the share of the employment surplus captured by employers to decline in recessions, resulting in less recruiting effort and higher unemployment.

I take the driving force of the movements in the economy to be (1) changes in total factor productivity, in the terms of trade, and in the prices of factors other than labor, and (2) changes in household welfare arising from wealth effects of government activity. Changes in these variables induce changes in households and the labor market that are captured in the two endogenous variables, marginal utility and the marginal product of labor. Consumption, hours of work, and the employment rate are governed by structural equations involving only these two variables, which I take as basic underlying factors in the statistical factor model.

In the extended MP model, the fundamental efficiency condition breaks into two parts. One is bilateral private efficiency between worker and employer. Once a worker matches with an employer, the decisions about starting a job and ending the job and the choice of the number of hours of work are bilaterally efficient. No worker suffers layoff because a wage is too high. The second is social efficiency in the employment rate. Unemployment is determined in the labor market as a whole by the anonymous interactions of all workers and employers. Coase (1960) pointed out the likelihood of efficiency in bilateral relationships and the failure of efficiency in relationships involving externalities among millions of people and firms. In the MP model, firms' recruitment activities tighten the labor market, with favorable external effects on job-seekers and unfavorable ones on other firms. Similarly, job-seekers have favorable effects on firms and unfavorable effects on other job-seekers. Hosios (1990) describes these effects and derives the condition for an efficient employment rate. The employment rate of 94 or 95 percent seen in normal times may well be efficient, but a decline to 91 or 92 percent in a severe recession is almost certainly inefficiently high.

The paper finds a fairly small and non-cyclical discrepancy between the marginal value of time, inferred from hours and other measures from the labor market, and the marginal product of labor, inferred from productivity and tax rates. The reconciliation does not rest on highly elastic labor supply—the Frisch elasticity is taken to be 0.7, consistent with the findings of research with household data. Most of the cyclical variation in work effort takes the form of variations in unemployment, governed by the extended MP model. The substantial wage elasticity of the employment rate—the new concept I derive in detail—is a key element of the reconciliation. Another feature of the model important for the reconciliation is the complementarity of consumption and work effort—unemployed workers, having more time to produce at home, choose lower levels of consumption of market goods than they would if employed, at the same level of well-being.

The small size of the discrepancy between the marginal value of time and the marginal product implies that bilateral private inefficiency is not a feature of recessions. Actions that are within the grasp of the matched worker and employer—determining if the worker should join the firm, setting hours, and determining if the worker should leave the firm—are resolved as Coase predicted. On the other hand, the small discrepancy does not imply that the level of unemployment is efficient.

In this paper, I do not consider the small procyclical movements of participation in the labor force—Hall (2008) documents these movements. For simplicity, I treat the labor force as exogenously determined.

I treat marginal utility and the common value of the marginal value of time and the marginal product of labor as unobserved latent state 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. I do not use macro data to estimate the model’s slope parameters. One reason is that the model falls short of identification. The main reason is that macro data are probably not the best way to estimate parameters—data at the household level are generally more powerful. I use information from extensive research on some of the coefficients.

2 Earlier Research

Macroeconomists have approached the question of possible inefficiency in recessions from various angles. The original real business cycle model of Kydland and Prescott (1982)

avored efficiency. It portrayed the decline in work effort in recessions as the result of a modest decline in the marginal product of labor and an equal decline in the marginal value of time. Highly elastic labor supply implied that a substantial decline in work effort accompanied the modest decline in the marginal value of time. This view contradicts a body of microeconomic research on labor supply showing Frisch wage elasticities below one.

The efficiency condition and the puzzle of inefficiency in recessions makes no reference to wages, but a long-standing view holds that a market economy uses wages to decentralize the allocation of labor and that recessions are times when wages are too high. The inefficiency arises in recessions because employers equate the marginal product of labor to the wage and do not expand hiring to absorb unemployed workers whose marginal value of time falls short of the wage. Barro (1977) pointed out the weakness in this view—it invokes a failure of what is generally seen as one of the most reliable predictions of economics, that two actors in direct contact with one another will arrange their relationship to achieve private bilateral efficiency.

A third view has evolved in the past 15 years. It starts from the observation that the conspicuous failure of the efficiency condition arises from unemployment rather than from the hours of work of employed workers. Neither the original real business cycle view nor the traditional sticky-wage view assigned an explicit role to unemployment. Mortensen and Pissarides (1994) contributed a coherent account of unemployment, the MP model. Shimer’s (2005) investigation showed that the original MP model failed to account for cyclical fluctuations in unemployment. His paper launched a wide variety of modified MP models with amplified unemployment fluctuations.

The findings of this paper are closely related to the literature on the labor-market “wedge,” the gap between the marginal product of labor and the marginal value of time. Shimer (2008) summarizes this literature. Some earlier work has concluded that the wedge includes substantial unexplained cyclical movements. This paper quantifies the difference between the marginal value of time and the tax-adjusted marginal product of labor, so it focuses on the unexplained element and not the inefficiency resulting from taxation, the topic of an important branch of the wedge literature not concerned specifically with cyclical issues.

3 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 system 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. Blundell, Pistaferri and Preston (2008) find evidence of substantial insurance of individual workers against transitory shocks such as unemployment.

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. I make no claim that workers are insured against idiosyncratic permanent changes in their earnings capacities, only that the transitory effects of unemployment can usefully be analyzed under the assumption of insurance.

4 Dynamic Labor-Market Equilibrium

This section develops a unified model of the labor market and production. The outcome is a set of four equations relating the four observed variables in Figure 1—consumption, weekly hours, the employment rate, and the tax-adjusted marginal product of labor—to a pair of latent variables, the marginal utility of consumption and the marginal value of time. For consumption and hours, I use the Frisch consumption-demand and hours-supply equations, which provide a direct connection to a large body of research on household behavior in the Frisch framework. For the employment rate, I show that an interesting class of MP-style search and matching models also result in an equation relating the employment rate to marginal utility and the marginal value of time.

I consider an economy with many identical families, each with a large number of mem-

bers. All workers face the same pay schedule and all members of all families have the same preferences. The family insures its members against personal (but not aggregate) risks and satisfies the Borch-Arrow condition for optimal insurance of equal marginal utility across individuals. In each family, a fraction n_t of workers are employed and the remaining $1 - n_t$ are searching. These fractions are outside the control of the family—they are features of the labor market. In my calibration, a family never allocates any of its members to pure leisure—it achieves higher family welfare by assigning all non-working members to job search and it never terminates the work of an employed member. Thus, as I noted earlier, I neglect the small variations in labor-force participation that occur in the actual U.S. economy. To generate realistically small movements of participation in the model I would need to introduce heterogeneity in preferences or earning powers.

This section develops a model that generalizes the canonical model of Mortensen and Pissarides (1994). I adopt the undirected search and matching functions of their model, but replace the Nash bargain with a more general characterization of the determination of a newly hired worker’s compensation. I also follow other authors in generalizing preferences and incorporating choice over hours of work. I will refer to the result as the *extended MP model*.

4.1 Concepts of the wage

In the exposition of the model in this section, I will refer to the variable w as the wage, in the sense of the common value of the marginal value of time and the marginal product of labor that would occur in equilibrium in an economy with a market wage with the property that hours of work are chosen on the supply side to equate the marginal value of time to w and hours of work are chosen on the demand side to equate the marginal product of labor to the same w . This terminology seems most natural for describing the model. On the other hand, I need to distinguish between the supply side and the demand side in the empirical section, because a major issue in the paper is the possible wedge separating the marginal value of time and the marginal product of labor. At that point, I will switch to calling the marginal value of time v and the marginal product of labor m .

A related point is that none of these three concepts of the wage is the amount workers receive per hour—they are all shadow concepts reflecting marginal rather than average wages. I use the term *compensation* for actual cash payments to workers. Except for a brief

discussion at the end, I do not consider data on compensation.

4.2 Search and matching

Employers post vacancies. Each period, the probability that a worker will become available to fill the vacancy is q . In tighter labor markets, vacancies are harder to fill and q is lower. The MP model characterizes the tightness of the labor market in terms of the vacancy/unemployment ratio θ . The job-finding rate is an increasing and concave function $\phi(\theta)$ and the vacancy-filling rate is the decreasing function $\phi(\theta)/\theta$. The model assumes a constant exogenous rate of job destruction, s . Employment follows a two-state Markov process with stochastic equilibrium

$$n = \frac{\phi(\theta)}{s + \phi(\theta)}. \quad (1)$$

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 paper. I invert equation (1) to find $\theta(n)$ and take the job-filling probability to be the decreasing function

$$q(n) = \phi(\theta(n))/\theta(n). \quad (2)$$

In a tighter labor market with higher employment rate n , the job-filling rate $q(n)$ is lower.

As in the MP model, employers incur a cost γ at the beginning of a period to maintain a vacancy for the period, with probability $q(n)$ of filling the job at the end of the period.

4.3 The employment contract

Prices are quoted in terms of output. Employers pay workers w_t units of output for each hour of work in period t . Employers collect an amount y_t from a new worker. Both workers and employers are price-takers with respect to w_t , so the employment contract embodies efficient two-part pricing. I discuss the determination of y_t shortly; it is a key feature of the model. For simplicity I develop the model as if y_t were collected at the beginning of the job, but the results would be identical if it were spread over the period of employment and y_t were the present value as of the beginning of the period of the amount deducted from $w_t h_t$ by the employer. This contract is in a much more general class studied in Brügemann and Moscarini (2008), who emphasize the complement to y , the amount of rent captured by

the job-seeker. They include bilateral idiosyncratic asymmetric information as well, a major complication I do not take up.

4.4 Production and the firm's decisions

The economy has a single kind of output, with production function

$$F(H_t, K_t, \eta_t). \quad (3)$$

Here $H_t = n_t h_t$ is total hours of work, K_t is the capital stock, and η is a vector of random disturbances.

Firms make three decisions: (1) the number of vacancies to try to fill each period, (2) the hours to demand from the existing work force, and (3) the demand for capital.

(1) Under the standard employment contract, firms break even from employing a new worker during the worker's tenure. They decide whether to recruit workers based upon the immediate payoff,

$$q(n_t)y_t - \gamma. \quad (4)$$

They invest γ in holding a vacancy open for the period and have a probability $q(n_t)$ of gaining the payoff y_t . Firms are large enough to absorb the fully diversifiable risk associated with the probability of successful recruiting. Firms would create infinitely many vacancies if the payoff were positive and zero if it were negative. Equilibrium requires that the payoff to recruiting be zero:

$$q(n)y = \gamma. \quad (5)$$

The employment rate that solves this zero-profit condition is a function $n(y)$, which I call the *employment function*.

(2) The number of employees at a firm is a (fast-moving) state variable. The first-order condition,

$$\frac{\partial F}{\partial H} = w_t, \quad (6)$$

describes the firm's demand for their hours. To embody labor demand in the model, I measure productivity as the average product of labor, m , the ratio of total output to total hours of labor input. I let α be the elasticity of the production function with respect to labor input. From

$$w = \frac{\partial F}{\partial H} = \alpha m, \quad (7)$$

I get the equation:

$$\log m = \log w - \log \alpha. \quad (8)$$

(3) A capital services market allocates the available capital efficiently among firms in proportion to their employment levels. The first-order condition,

$$\frac{\partial F}{\partial K} = r_t, \quad (9)$$

describes the firm's demand for capital.

4.5 The family's decisions

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.

The family orders levels of hours of employed members, h_t , consumption of employed members, $c_{e,t}$, and consumption of unemployed members, $c_{u,t}$, within a period by the utility function,

$$n_t U(c_{e,t}, h_t) + (1 - n_t) U(c_{u,t}, 0). \quad (10)$$

The family orders future uncertain paths by expected utility with discount factor δ . I view the family utility function as a reduced form for a more complicated model of family activities that includes home production.

The family solves the dynamic program,

$$V(W_t, \eta_t) = \max_{h_t, c_{e,t}, c_{u,t}} \{n_t U(c_{e,t}, h_t) + (1 - n_t) U(c_{u,t}, 0) + \mathbb{E}_{\eta_{t+1}} \delta V((1 + r_t)[W_t - d(\eta_t) - n_t c_{e,t} - (1 - n_t) c_{u,t}] - \phi(n_t)(1 - n_t)y_t + w_t n_t h_t, \eta_{t+1})\} \quad (11)$$

Here $V(W_t, \eta_t)$ is the family's expected utility as of the beginning of period t , W_t is wealth, and $d(\eta_t)$ is a deduction from wealth depending on the random driving forces η_t that could, for example, arise from the lump-sum component of taxation. The expectation is over the conditional distribution of η_{t+1} . The amount $\phi(n_t)(1 - n_t)$ is the flow of new hires of family members, each of which costs the family y_t .

4.6 Equilibrium

Let $\eta_{(t)}$ be the history of the random driving forces up to time t . An equilibrium in this economy is a wage function $w(\eta_{(t)})$, a return function $r(\eta_{(t)})$, and an employment rate function $n_t(\eta_{(t)})$ such that the supply of hours $h(\eta_{(t)})$ and the supply of savings, $W(\eta_{(t)})$, from the family's maximizing program in equation (11) equal the firm's demands from equations (6) and (9), and the recruiting profit in equation (4) is zero, for every $\eta_{(t)}$ in its support.

4.7 State variables

I let λ_t be the marginal utility of wealth (and also marginal utility of consumption):

$$\lambda_t = \frac{\partial V}{\partial W_t} = \delta(1 + r_t) \mathbb{E}_t \frac{\partial V}{\partial W_{t+1}}. \quad (12)$$

I take λ_t and the hourly wage w_t as the state variables of the economy relevant to labor-market equilibrium. Both state variables are complicated functions of the underlying exogenous driving forces, η . Marginal utility, λ_t , is an endogenous variable that embodies the entire forward-looking optimization of the household based on its perceptions of future earnings and deductions. The common value of the marginal product of labor and the marginal value of time, w , is an endogenous variable that depends on the amounts of capital and labor used in production, which depend in turn on all of the elements of the labor-market model and on features of the economy not included in that model.

The strategy pursued in the rest of the paper exploits the property that a vector of four key observable endogenous variables—consumption, hours of work, the employment rate, and the marginal product of labor—are all functions of the two endogenous state variables, λ and w . The four observable variables have a factor structure, with just two latent factors.

4.8 Hours, consumption, and employment

The family's first-order conditions for hours and the consumption levels of employed and unemployed members are:

$$U_h(c_{e,t}, h_t) = -\lambda_t w_t \quad (13)$$

$$U_c(c_{e,t}, h_t) = \lambda_t \quad (14)$$

$$U_c(c_{u,t}, 0) = \lambda_t. \quad (15)$$

These conditions define three Frisch functions, $c_e(\lambda_t, \lambda_t w_t)$, $h(\lambda_t, \lambda_t w_t)$, and $c_u(\lambda_t)$ giving the consumption and hours of the employed and the consumption of the unemployed. I write the functions in this form to connect with research on Frisch labor-supply and consumer-demand equations. With consumption-hours complementarity, the family assigns a lower level of consumption to the unemployed than to the employed: $c_u < c_e$.

4.9 The compensation bargain

Recall that the prospective worker and employer bargain over the part of compensation taking the form of a lump sum, $y(\lambda, w)$, that the worker pays the employer at the beginning of the job. During employment, the employer pays the worker the economy's common value of the marginal product of labor and the marginal value of time, w ; the parties do not bargain over this component of compensation. In terms of an Edgeworth box, the bargain over y determines a point on the contract curve while the adoption of a common marginal value of time and marginal product of labor places the parties on the contract curve. This setup is a convenient way to think about compensation because it separates the part, $y(\lambda, w)$, that controls the tightness of the labor market and thus the employment rate from the efficient determination of hours of work, based on w . The payment $y(\lambda, w)$ funds the employer's recruiting effort. Of course, I do not believe that workers pay up front for the right to hold a job—actual compensation arrangements annuitize $y(\lambda, w)$ over the duration of the job, so the cash compensation flow to the worker is less than wh .

I am agnostic about the principles underlying the bargain—the only restriction is that the bargained payment is a function $y(\lambda, w)$ of the two state variables. One could interpret this assumption as a Markov property, the exclusion of any other endogenous state variable arising from the bargaining game between worker and employer. This exclusion has substance, as it rules out a state variable that might capture the inertia of compensation. In the setup of this paper, compensation can be sticky in the sense of being unresponsive to the state of the labor market, but it cannot be sticky in the sense of being under the influence of a slow-moving state variable other than λ and w .

I note that the Nash wage bargain is a member of the class of models where y is a function of the two state variables alone. The reservation payment for the employer, having encountered a worker, is zero—the employer is indifferent to hiring at that point and comes out definitely ahead if the worker makes any positive payment. The family's upper limit on

the payment, Y , is the amount of the increase in its value function from shifting a member from unemployment to employment, stated in consumption units. I differentiate equation (11) with respect to the family's own n_t (keeping the market-wide n_t in $\phi(n_t)$ constant) and apply equation (12) to get

$$Y(\lambda, w) = \frac{U(c_e, h) - U(c_u, 0)}{\lambda} - c_e + c_u + w_t h_t. \quad (16)$$

This is the change in utility when a member moves from unemployment to employment, restated in consumption units by dividing by λ plus the budgetary effect of the increase in consumption spending (a negative consideration) plus the added earnings. All of the terms in this expression are functions of λ or w or both. Let the fixed Nash bargaining weight of the job-seeker be ν . The Nash-bargain upfront payment is $y(\lambda, w) = (1 - \nu)Y(\lambda, w)$.

The employment function $n(y(\lambda_t, w_t))$ can now be written $n(\lambda_t, w_t)$, so it joins consumption and hours as functions of the two state variables, a property I will exploit shortly in the empirical analysis.

4.10 Volatility

Volatility in the labor market occurs because of movements in the random disturbances η_t , which result in movements of marginal utility λ and the wage w . These movements include shifts in technology and in other factors that appear in the technology as a reduced form, such as changes in the terms of trade. The volatility of hours operates in the standard way—an increase in the wage raises $h(\lambda, w)$ through the direct effect of w but the resulting decline in λ , arising from the favorable effect of a higher wage on wealth, lowers hours. Most labor volatility in the U.S. economy takes the form of variations in the employment rate $n(\lambda, w)$. Here again a higher wage raises employment while the resulting higher wealth and lower value of λ lowers employment, but, with the parameter values used in this paper, employment is more sensitive to both variables than is the supply of hours.

Volatility also occurs because of movements in the deduction $d(\eta_t)$ —a higher deduction induces higher hours of work through the standard wealth effect in labor supply.

The response of the employment rate to changes in the driving forces depends directly on the payment $y(\lambda, w)$ that a newly hired worker makes to the employer—see equation (5). The higher this payment, the tighter is the labor market, because employers recruit new workers more aggressively when the payoff is higher. If the payment were fixed, the employment rate would also be fixed. In fact, when the driving forces raise the wage w , the

employment rate rises, according to the evidence later in this paper. So an increase in the wage induces an increase in the upfront payment, y . Because the payment is a deduction from the worker’s total compensation, the positive response of the payment to w means that compensation does not rise in proportion to the wage—it is sticky in that sense. If, as seems likely, the upfront payment is amortized over the duration of a job, then the elasticity of the compensation that workers receive with respect to the underlying wage w is less than one. A higher w delivers more value from the employment relation to the employer and induces greater recruiting effort and thus a tighter labor market with a higher employment rate n .

In this framework, I interpret Shimer (2005) as showing that the value of the upfront payment y resulting from a Nash bargain with roughly equal bargaining weights has low sensitivity to w and results in low volatility of the employment rate. At the other extreme, if compensation to the worker—the present value of wh over the job less the upfront payment y —were unresponsive to w , y would move in proportion to w . In this situation of completely sticky compensation, recruiting effort would rise sharply with w and the volatility of the employment rate would be high and procyclical. The calibration in this paper makes the employment rate quite sensitive to w and implies that newly hired workers let employers keep some important part of an increase in w because the worker makes a higher upfront payment y . In general, the high sensitivity of $n(\lambda, w)$ to w implies some stickiness of compensation.

5 Unemployment Theories

What theories of employment and unemployment fit the paradigm of the extended MP model, where the employment rate is a function of λ and w ? 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 compensation just offsets the disamenity of the loss of time at home. Labor supply is perfectly elastic at that level of compensation. The employed are those who wind up in jobs at the labor demand prevailing at that compensation.

Second, in *search-and-matching models*—surveyed by Rogerson, Shimer and Wright (2005)—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. The

Mortensen and Pissarides (1994) model is a leading example. 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 .

Hall and Milgrom (2008) develop an alternating-offer bargaining model and calibration in which compensation is sufficiently insensitive to labor-market conditions that productivity changes cause realistic changes in unemployment. Hagedorn and Manovskii (2008) generate similar responses with Nash bargaining by assuming low bargaining power for the worker and a flow value of non-work corresponding to a high elasticity of labor supply.

The efficiency-wage model of unemployment volatility, as developed by Alexopoulos (2004), also fits within the framework developed above. Her model omits explicit treatment of the search-and-matching process, but the substance is the same. Under the efficiency-wage principle, employers set compensation at the level needed to prevent short-run opportunism among workers—their share of the employment surplus needs to be large enough to keep them working effectively. When productivity rises, the benefits go mostly to employers, who respond by recruiting harder and tightening the labor market.

Third, *allocational sticky-wage models* invoke a state variable, the sticky wage, that controls the allocation of labor. See Blanchard and Gali (2007) for a representative allocational sticky-wage model, but without any treatment of unemployment, and Romer (2006), p. 467, for a brief discussion of unemployment in that context. 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, such as the capital stock. Thus, because they drop the key assumption of perfectly elastic labor demand, allocational sticky-wage models have rather different implications for the employment function. In particular, labor-market outcomes depend on more than the two variables λ and w .

The equilibrium model plainly belongs to the class of models where employment depends just on λ and w . 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. 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.

A quick summary of this discussion is that sticky-compensation models in the extended MP class are consistent with the model in this paper, while allocational sticky-wage models are not.

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 rate. The next step is to measure the response of the rate to the two determinants.

6 Parameter Values

6.1 Research on preferences

The empirical approach in this paper draws parameters relating to preferences from research on individual behavior. This section relates the three functions $h(\lambda, \lambda w)$, $c_e(\lambda, \lambda w)$, and $c_u(\lambda)$ to that research.

Consider the standard intertemporal consumption-hours problem without unemployment,

$$\max \mathbb{E}_t \sum_{\tau=0}^{\infty} \delta^\tau U(c_{t+\tau}, h_{t+\tau}) \quad (17)$$

subject to the budget constraint,

$$\sum_{\tau=0}^{\infty} R_{t,\tau} (w_{t+\tau} h_{t+\tau} - c_{t+\tau}) = 0. \quad (18)$$

Here $R_{t,\tau}$ is the price at time t of a unit of goods delivered at time $t + \tau$.

I let $c(\lambda, \lambda w)$ be the Frisch consumption demand and $h(\lambda, \lambda w)$ be the Frisch supply of hours per worker. See Browning, Deaton and Irish (1985) for a complete discussion of Frisch systems in general. The functions satisfy

$$U_c(c(\lambda, \lambda w), h(\lambda, \lambda w)) = \lambda \tag{19}$$

and

$$U_h(c(\lambda, \lambda w), h(\lambda, \lambda w)) = -\lambda w. \tag{20}$$

Here λ is the Lagrange multiplier for the budget constraint.

The Frisch functions have symmetric cross-price responses: $c_2 = -h_1$. They have three basic first-order or slope properties:

- *Intertemporal substitution in consumption*, $c_1(\lambda, \lambda w)$, the response of consumption to changes in its price
- *Frisch labor-supply response*, $h_2(\lambda, \lambda w)$, the response of hours to changes in the wage
- *Consumption-hours cross effect*, $c_2(\lambda, \lambda w)$, 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.

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 functions is negative semi-definite. Consequently, the derivatives satisfy the following constraint on the cross effect controlling the strength of the complementarity:

$$c_2^2 \leq -c_1 h_2. \tag{21}$$

Each of these responses has generated a body of literature. In addition, in the presence of uncertainty, the curvature of U controls risk aversion, the subject of another literature.

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.

The Appendix 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. I believe that a fair conclusion from the research is that Frisch elasticity of consumption demand is $\beta_{c,c} = -0.5$ and the Frisch elasticity of hours supply is $\beta_{h,h} = 0.7$.

The literature on measurement of the cross-elasticity is sparse, but a substantial amount of research has been done on the decline in consumption that occurs when a person moves from normal hours of work to zero because of unemployment or retirement. The ratio of unemployment consumption c_u to employment consumption c_e reflects the same properties of preferences as does the Frisch cross-elasticity. I use the parametric utility function in Hall and Milgrom (2008) to find the cross-elasticity that corresponds to the consumption ratio of 0.85. It is a Frisch cross-elasticity of $\beta_{c,h} = 0.3$.

6.2 Parameters of the employment function

Hall and Milgrom (2008), Table 3, report that the observed elasticity of the unemployment rate with respect to productivity is about 20. This calculation holds the flow value of non-employment constant, so it corresponds in the framework of this paper to holding λ constant. The corresponding elasticity of the employment rate with respect to w is 1.2, the value I use.

I have not found any outside benchmark for the elasticity of $n(\lambda, w)$ with respect to λ . The general view of wage bargaining developed earlier in the paper does not speak to the value of the elasticity. Accordingly, I choose a value, $\beta_{n,\lambda} = 0.6$, that yields approximately the best fit.

The parameters of the employment function are the only ones chosen on the basis of

fit to the aggregate data. Research on search-and-matching models with realistic nonlinear preferences, non-Nash wage bargaining, and other relevant features has flourished recently and may provide more guidance in the future.

7 Econometric Model

I approximate the consumption demands, hours supply, and employment functions 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 . The model is:

Consumption of the employed:

$$\log c_e = \beta_{c,c} \log \lambda + \beta_{c,h}(\log \lambda + \log w) \quad (22)$$

Consumption of the unemployed:

$$\log c_u = \beta_{c,c} \log \lambda \quad (23)$$

Hours:

$$\log h = -\beta_{c,h} \log \lambda + \beta_{h,h}(\log \lambda + \log w) \quad (24)$$

Employment rate:

$$\log n = \beta_{n,\lambda} \log \lambda + \beta_{n,w} \log w \quad (25)$$

Productivity:

$$\log m = \log w + \log \alpha \quad (26)$$

Table 1 summarizes the parameter values I use as the base case. The upper panel gives the elasticities described in the previous section and the lower panel restates them as the coefficients governing the relation (in logs) between the observed variables and the underlying factors, λ and w . The lower panel takes into account the double appearance of λ in the Frisch consumption-demand and hours-supply functions .

Notice that the employment equation resembles the hours equation, but with larger coefficients. The elasticities of annual hours, nh , with respect to λ and w , are the sums of the coefficients in the second and third lines of the lower panel of Table 1. The effect of

<i>Elasticities</i>		
Consumption with respect to λ		-0.5
Consumption with respect to λw		0.3
Hours with respect to λ		-0.3
Hours with respect to λw		0.7
Employment with respect to λ		0.6
Employment with respect to w		1.2
<i>Coefficients</i>		
	λ	w
Consumption	-0.2	0.3
Hours	0.4	0.7
Employment	0.6	1.2
Productivity	0	1

Table 1: Parameters and Corresponding Coefficients in the Equations of the Model

including a substantially elastic employment function is to make annual hours far more elastic than is labor supply in household studies. The introduction of an employment function is a way to rationalize the fact of elastic annual hours with the microeconomic finding that the weekly hours of individual workers are not nearly so elastic. The employment function is not a feature of individual choice, but of the interaction of all workers and all employers.

7.1 Long-run properties and detrending

Hours of work, h , were roughly constant over the past 60 years. Given constant hours, the family's budget constraint requires, roughly, that consumption grow at the same rate as the marginal product of labor, w . Putting these conditions into the equations above yields the standard conclusion that the own-price elasticity of consumption demand, $\beta_{c,c}$, is minus 1 (log preferences) and that the cross-effect, $\beta_{c,h}$, is zero. Neither of these conditions is consistent with evidence from household studies. Therefore, I interpret the model as describing responses at cyclical frequencies but not at low frequencies, where trends in household technology and preferences come into play. Thus I study detrended data, specifically, the residuals from regressions of the data, in log form, on a third-order polynomial in time. Figure 1 showed the detrended data. The detrending also removes the production elasticity, α , which I assume moves only at low frequencies.

The uncompensated hours-supply function is backward-bending for the parameter values I use. By uncompensated, I mean subject to a budget constraint where consumption equals the amount of earnings, wh . Solving equations (40) and (41) for the change in h and c for a doubling of w subject to constancy of $\log(wh/c)$, I find that hours would fall by 15 percent and consumption would rise by 70 percent. With preferences satisfying the restriction of zero uncompensated wage elasticity, hours would remain the same and consumption would double.

My approach here is the opposite of Shimer's (2008). He requires that preferences satisfy the long-run restrictions and therefore does not match the elasticities I use. Because preferences are a reduced form for a more elaborate specification including home production, where productivity trends might logically be included, it is a matter of judgment whether to impose the long-run restrictions. Of course, the best solution would be a full treatment of the household with explicit technology and measured productivity trends.

7.2 Consumption

The model disaggregates the population by the employed and unemployed, who consume c_e and c_u respectively. Only average consumption c is observed. It is the average of the two levels, weighted by the employment and unemployment fractions:

$$c = nc_e + (1 - n)c_u. \quad (27)$$

I solve this equation for c_e given the hypothesis that the consumption of the unemployed is a fraction ρ of the consumption of the employed:

$$c_e = \frac{c}{n + (1 - n)\rho}. \quad (28)$$

The evidence discussed in the Appendix suggests that $\rho = 0.85$. The effect of this calculation is to remove from c the mix effect that occurs when employment falls and more people are consuming the lower amount c_u . The adjustment is quite small. I drop c_u from the model because it is taken to be strictly proportional to c_e .

7.3 Disturbances and their variances

The data do not fit the model exactly. I hypothesize additive disturbances ϵ_c , ϵ_h , ϵ_n , and ϵ_m in the equations for the four observed variables. I assume that these are uncorrelated with λ and w . This assumption is easiest to rationalize if the ϵ s are measurement errors.

I write the model in matrix form as

$$x = zB + \epsilon. \quad (29)$$

Here x is the row vector of observed values of the logs of consumption, hours, employment, and productivity, z is the row vector of logs of λ and w , and B is the matrix of coefficients shown in the bottom panel of Table 1. The covariance matrix of x is

$$\Omega = V(x) = B'CB + D, \quad (30)$$

where C is the 2×2 matrix containing the variances of λ and w and their covariance and D is the diagonal matrix of variances of the ϵ s.

I use a minimum quadratic distance estimator of the three distinct elements of C . Given an estimate of C , the four variances in D can be chosen to match the diagonal elements of Ω exactly, by subtraction:

$$D_{i,i} = \Omega_{i,i} - (B'CB)_{i,i}. \quad (31)$$

The estimation problem for C is to

$$\min_C \sum_i \sum_{j>i} (\Omega_{i,j} - (B'CB)_{i,j})^2. \quad (32)$$

This is an ordinary-least-squares calculation. The 6 distinct off-diagonal elements of $\Omega_{i,j}$ form the left-hand variable and the derivatives of $(B'CB)_{i,j}$ with respect to the three distinct elements of C are the right-hand variables. There is no constant.

7.4 Inferring the latent variables

Because there are six latent variables— λ , w , and the four disturbances in ϵ —the model does not permit the recovery of unique values of the latent variables from the four observed variables. But the model is informative about the likely values of the latent variables. From the point of view of statistical theory, the problem of inferring the likely values is the same as making a forecast. In the forecasting situation, one knows the joint distribution of some forecasting variables and the variable to be forecast. The standard way to form the forecast is to calculate the coefficients for the regression of the forecasted variable, say z , on the forecasting variables, say x . According to the standard ordinary least squares formula, the coefficients are:

$$b = V(x)^{-1} \text{Cov}(x, z). \quad (33)$$

The forecast, say \hat{z} , is the fitted value,

$$\hat{z} = xb \tag{34}$$

The situation is no different if z is an unknown latent variable this period rather than a variable that is unknown because it will only be learned next period, provided that the covariance matrix of x and z is known. Here

$$\text{Cov}(x, z) = B'C. \tag{35}$$

Recall that B is the matrix of coefficients shown in the lower panel of Table 1 and C is the matrix containing the variances and covariance of $\log \lambda$ and $\log w$, estimated by the procedure just described. Also recall that Ω is the observed covariance matrix of x . Thus the regression is

$$b = \Omega^{-1}B'C. \tag{36}$$

Note that b is a 4×2 matrix of regression coefficients. The first column are those to be applied to the four observed variables to infer $\log \lambda$ and the second column to infer $\log w$:

$$\hat{z} = xb. \tag{37}$$

The companion estimated values of the disturbances ϵ are the residuals from the structural equations:

$$\hat{\epsilon} = x - \hat{z}B. \tag{38}$$

The variance of each inferred disturbance is somewhat less than the known variance of the disturbance. The shortfall arises for the same reason that the forecast of a random variable has less variance than the variable is known to have.

7.5 Marginal value of time and marginal product of labor

With its disturbance, the equation for productivity is

$$\log m = \log v + \epsilon_m. \tag{39}$$

I assume that detrending removes $\log \alpha$ from productivity, so I treat $\log m$ as the measured marginal product of labor. The addition of the disturbance introduces a distinction between the marginal value of time and the marginal product of labor. Here and in the remainder of

the paper, I use v for the marginal value of time—it is the shadow price of labor hours on the supply side of the labor market—and m for the marginal product of labor, the shadow price of labor hours on the demand side. Thus ϵ_m is the labor wedge, the difference between the shadow wage on the supply side and on the demand side.

7.6 Restatement of the model

The model as estimated is:

Consumption of the employed:

$$\log c_e = \beta_{c,c} \log \lambda + \beta_{c,h}(\log \lambda + \log v) + \log \epsilon_c \quad (40)$$

Hours:

$$\log h = -\beta_{c,h} \log \lambda + \beta_{h,h}(\log \lambda + \log v) + \log \epsilon_h \quad (41)$$

Employment rate:

$$\log n = \beta_{n,\lambda} \log \lambda + \beta_{n,v} \log v + \log \epsilon_n \quad (42)$$

Productivity:

$$\log m = \log v + \epsilon_m. \quad (43)$$

8 Data

To avoid complexities from durables purchases, I use nondurables and services 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 (2008).

I measure productivity as output per hour of all persons, private business, BLS series PRS84006093. To calculate the tax wedge, I include the marginal personal income tax rate using the NBER's TAXSIM software applied to data on the average income per personal return from the Internal Revenue Service, the marginal rate of FICA (social security) taxation, and the average direct tax rate on consumption from sales and other taxes, calculated as the ratio of revenue from these sources to total consumption. For further details of the tax calculations, see the spreadsheet on my website.

	<i>Consumption of employed</i>	<i>Hours per worker</i>	<i>Employment rate</i>	<i>Productivity</i>
Covariances $\times 10,000$				
Consumption of employed	1.27	-0.12	0.36	1.46
Hours per worker		1.11	0.96	-0.03
Employment rate			1.61	0.48
Productivity				2.80
Standard deviations (percents)	1.13	1.05	1.27	1.67
Correlations				
Consumption of employed	1.00	-0.10	0.25	0.78
Hours per worker		1.00	0.72	-0.01
Employment rate			1.00	0.23
Productivity				1.00

Table 2: Covariances, Standard Deviations, and Correlations of Logs of Consumption, Hours, Employment, and Productivity

Table 2 shows the covariance and correlation matrixes of the logs of the four detrended series. Consumption is correlated positively with employment—it is quite pro-cyclical. Consumption-hours complementarity can explain this fact. Not surprisingly, hours and employment are quite positively correlated. Consumption also has by far the highest correlation with productivity.

The standard deviation of the employment rate is about 25 percent higher than the standard deviation of hours—the more important source for the added total hours of work in an expansion is the reduction in unemployment. Hours are not very correlated with productivity. Note that productivity has the highest standard deviation of the four variables—amplification of productivity fluctuations need not be part of a model in which productivity is the driving force.

9 Results

The top two panels of Table 3 show the second moments estimated by the procedure described earlier. The top one describes the volatility of the disturbances in the four equations. The left column shows the standard deviations of the disturbances, calculated from the difference between the variance of the observed variable and the variance implied by the structural model. All the standard deviations are well under one percent. The right column compares

	<i>St. dev. of disturbance, percentage points</i>	R^2
Consumption of the employed	0.48	0.82
Hours of workers	0.78	0.45
Employment rate	0.59	0.78
Productivity	0.62	0.86
Standard deviation of inferred marginal utility, λ	3.02	
Standard deviation of inferred marginal value of time, v	1.55	
Correlation of marginal utility and marginal value of time	-0.81	
Coefficients, b , to infer unobserved factors		
	Marginal utility, λ	Marginal value of time, v
Consumption	-2.50	0.54
Hours	-0.97	0.32
Employment	1.80	-0.12
Productivity	-0.38	0.60

Table 3: Parameter Values and Estimated Variances and Covariance

the variance of the disturbance to the variance of the observed variable in the form of an R^2 , calculated as one minus the variance of the disturbance divided by the variance of the variable. The model is reasonably successful in accounting for the volatility of all of the variables except weekly hours, the variable with the lowest volatility and the highest standard deviation of its residual.

9.1 Implied values of marginal utility and marginal value of time

The middle panel of Table 3 describes the second moments of the latent variables, marginal utility, λ , and the marginal value of time, v . The standard deviation of v is close to the observed standard deviation of productivity. The standard deviation of marginal utility is close to double the standard deviation of productivity. If productivity were a random walk, so that consumers updated consumption by the full amount of any change in productivity, the standard deviation of marginal utility would be double that of productivity, given the elasticity, $\beta_{c,c} = -0.5$. The correlation of -0.81 suggests that productivity has a large persistent component that generates movements in the opposite direction in marginal utility. If productivity were the only driving force and consumers observed the current innovation in productivity but had no other information about future productivity, the correlation would

be essentially -1 . If persistent movements in productivity are the primary driving force, but fiscal changes affect λ without much affecting v , then a value of -0.81 would seem quite reasonable.

The bottom panel of Table 3 shows the regression coefficients, b , to calculate the most likely values of the latent factors from the observed data. As expected, the inference of marginal utility puts a big negative weight on consumption—increases in consumption signal improvements in well-being and thus lower values of marginal utility, λ . Because of noise in consumption, the equation puts negative weight on productivity, which also has a negative relation to λ . The inference puts a large positive weight on employment, as expected. The negative coefficient on hours is a result of the swamping of the direct, positive relation between hours and marginal utility by powerful indirect effects operating through consumption and the employment rate.

The other feature of Table 3 worth noting is that the weight on productivity in the inference of v is 0.60 , below the loading of productivity on v 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 2 shows the estimates of detrended log marginal utility, $\log \lambda$, and marginal value of time, $\log v$, resulting from the application of the regression coefficients in Table 3 to the data on the four observables. The figure shows the pronounced negative correlation (-0.81) between marginal utility and the marginal value of time.

9.2 Fitted values for observables

Figure 3 shows the fitted values for the four observables from the time series for λ and v , using the regression coefficients in the bottom panel of Table 1. 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 pairs of them, hours-employment and consumption-productivity, 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 based in part on preference parameters drawn from earlier research. The success of the model is not so much the good fit shown in the figure, but rather achieving the good fit with coefficients that satisfy economic reasonability.

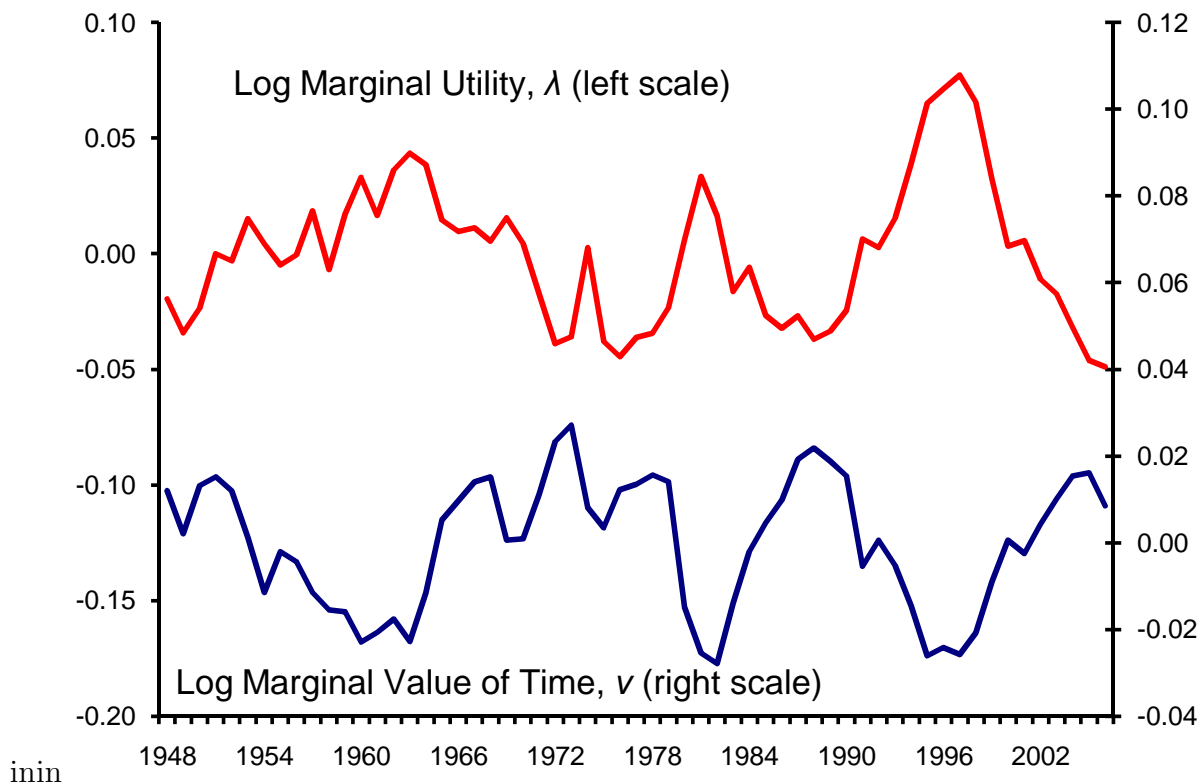


Figure 2: Inferred Values of Marginal Utility and Marginal Value of Time

9.3 Reconciliation of the marginal value of time and the marginal product of labor

The lower right plot in Figure 3 shows the extent that the model is able to generate estimates of the marginal value of time that track data on the tax-adjusted marginal product of labor. The marginal value of time follows measured productivity quite closely. The figure does not support any diagnosis of repeated or severe private bilateral inefficiency. Of course, one of the major factors accounting for the absence of bilateral inefficiency is to shift the efficiency issue from the bilateral situation of a worker and an employer to the economy-wide situation of job-seekers and employers interacting collectively and anonymously. Though the model portrays the movements in Figure 3 as privately bilaterally efficient, it does not portray socially efficient allocations.

The biggest departure from the normal view of the U.S. business cycle in the value of time-productivity plot occurs in the middle of the 1990s, usually viewed as a time of full employment and normal conditions, but portrayed here as an extended period of low productivity matched to low marginal value of time. Low productivity, trend- and tax-

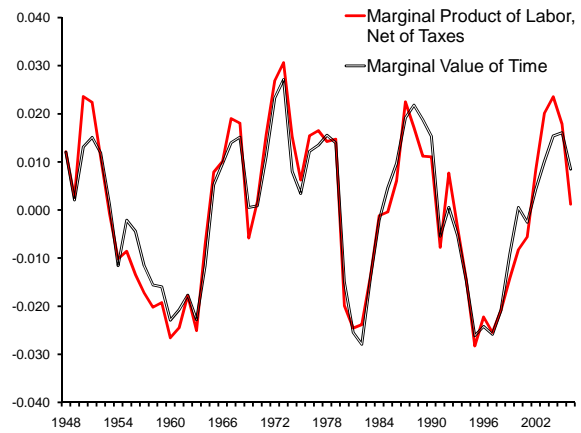





Figure 3: Actual and Fitted Values of the Four Observables

adjusted, comes straight from the data. How does the model infer that the marginal value of time was equally depressed? The upper left plot in Figure 3 shows that consumption was low, so marginal utility was high. Both hours and employment respond positively to λ —people work harder when they feel poorer, according to the standard theory of the household—and the employment function also responds positively—the labor market is tighter when λ is higher. Thus the slump in the mid-90s was a time when people worked hard because they did not feel well off. It was not a recession in the sense of a period of a slack labor market but it was a transitory period of depressed productivity and depressed value of time.

10 Comparison to Alternative Specifications

10.1 A traditional real-business-cycle specification

A major role of the employment function is to make annual hours of work, nh , more wage-elastic. From the dawn of formal modeling of aggregate fluctuations, researchers have invoked a fairly high wage elasticity to rationalize the high volatility of annual hours. To explore the traditional approach based on elastic labor supply where the employment rate is an aspect of household choice, I consider an alternative specification that drops the distinction between the employment function and the Frisch hours supply function by making them both respond with equally proportionate elasticities to λ and v . I alter the parameters to those shown in the second column of Table 4, labeled *RBC*. I use the standard assumption derived from long-run properties that the own-price elasticity of consumption is $\beta_{c,c} = -1$ corresponding to log consumption. I take preferences to be additively separable in consumption and hours, so $\beta_{c,h} = 0$ and consumption of employed and unemployed individuals are the same. I take the Frisch wage elasticity of hours supply to be 1.3 and the v and λ elasticities of the employment rate to be 2.7, so the overall elasticity of annual hours of work is 4.

Table 4 shows the results of estimating the variance parameters for the RBC specification. The first column repeats the results for the base case and the second column shows the results for the RBC specification. The estimation procedure fails, in that the estimated variance of the disturbance in the consumption equation is negative—the higher intertemporal elasticity of substitution makes the model predict a higher variance for the systematic part of consumption than is observed, so the variance of the disturbance, calculated as explained earlier as a residual, is negative. The standard deviation of the hours disturbance is a little higher and the standard deviation of the employment rate disturbance is lower. Most

	<i>Case</i>				
	<i>Base</i>	<i>RBC</i>	<i>RBC(0.6)</i>	<i>No comple- mentarity</i>	<i>Log c</i>
<i>Elasticities</i>					
Consumption with respect to own price, $\beta_{c,c}$	-0.5	-1	-0.6	-0.5	-1
Consumption with respect to wage, $\beta_{c,h}$	0.3	0	0	0	0.3
Hours with respect to consumption price, $-\beta_{c,c}$	-0.3	0	0	0	-0.3
Hours with respect to wage, $\beta_{h,h}$	0.7	1.3	1.3	0.7	0.7
Employment with respect to λ	0.6	2.7	2.7	0.7	0.6
Employment with respect to w	1.2	2.7	2.7	1.4	1.2
Ratio of unemployed consumption to employed	0.85	1	1	1	0.85
<i>Standard deviations of disturbances (percentage points)</i>					
Consumption of the employed	0.48	*	0.67	0.68	*
Hours of workers	0.78	0.88	0.86	0.74	0.89
Employment rate	0.59	0.38	0.21	0.72	0.85
Productivity	0.62	1.03	0.31	1.04	1.21
Standard deviation of inferred marginal utility	3.02	1.20	1.57	1.88	1.83
Standard deviation of inferred marginal product of labor	1.55	1.32	1.65	1.31	1.15
Correlation of marginal utility and marginal product	-0.81	-0.94	-0.96	-0.83	-0.74

*Negative variance

Table 4: Comparison of Base and Alternative Specifications

important, the standard deviation of the productivity disturbance is higher—the RBC specification is rather less successful than the base specification in matching the marginal value of time to the tax-adjusted marginal product of labor. Nonetheless, the RBC specification is fairly successful in matching the two. The conclusion that the wedge had large cyclical movements rested in part on the different econometric approach used in the earlier wedge literature. I discuss this further in the next section.

Figure 4 compares the actual and fitted values for productivity, interpreted as the marginal value of time and the marginal product of labor, for the RBC specification. Compare this figure to the lower right plot in Figure 3. The biggest problem with the RBC specification is its adherence to the long-run restriction of log utility of consumption, $\beta_{c,c} = -1$. With this elasticity set to -0.6 , the RBC specification matches the marginal value of time even better than the base specification does. The results for this alteration of the RBC specification are shown in Table 4 in the column labeled *RBC (0.6)*. This finding illustrates a key point of the paper: The traditional reliance on unrealistically elastic labor supply is well grounded in the facts about aggregate fluctuations. The accomplishment of this paper is not to rationalize the observed pattern of fluctuations without high elasticity, but to explain how the high elasticity is consistent with moderate elasticity arising from preferences. The base specification accomplishes this by assigning the bulk of the observed elasticity to the employment function, which is not an object determined by preferences alone.

10.2 Specifications with altered consumption elasticities

The column of Table 4 labeled *No complementarity* makes only one change to the base specification; it removes the complementarity between hours of work and consumption: $\beta_{c,h} = 0$. This change worsens the fit for consumption, the employment rate, and productivity and slightly improves it for hours. Thus complementarity is an important part of the success of the base specification in accounting for the joint movements of the four observed variables.

The last column of Table 4 makes a different single change to the base specification by setting $\beta_{c,c} = -1$ corresponding to log consumption in the utility function. As in the RBC specification, the systematic part of the model over-accounts for the variance of consumption, so the inferred variance of the disturbance is negative. The variances of the other disturbances are larger than in the base specification, especially productivity. Taming consumption volatility by setting $\beta_{c,c} = -0.5$ is another important element of the base specification.




Figure 4: Marginal Value of Time and Marginal Product of Labor, RBC Specification

11 Relation to the Labor-Wedge Literature

A number of authors have studied some of the issues in this paper in a framework that starts from the within-period first-order condition equating the marginal product of labor facing the worker to the worker's marginal rate of substitution,

$$(1 - \tau)m = \frac{-U_h(c_e, h)}{U_c(c_e, h)}. \quad (44)$$

Here τ is a wedge that plays the role of an extra tax on work effort, beyond the known taxes included in m . Shimer (2008) has an up-to-date discussion of the literature on the wedge and provides a new parametric specification for the marginal rate of substitution and new empirical estimates. In the setup of this paper, the latent variable v is the marginal rate of substitution. The disturbance ϵ_m is the log of the wedge.

Some authors view the wedge as inclusive of taxes, so that the wedge includes both observed and unobserved sources of inefficiency. Others view the wedge as only the unobserved element. I take the second view here by using data on productivity m measured net of observed taxes.

Shimer (2008) and all other contributions to the literature known to me, including Hall (1997), consider a system where the wedge is the only disturbance. In this paper, by contrast, each of the four observed variables has a disturbance—the productivity-MRS wedge is just one of four disturbances. The econometric method puts equal emphasis on the four disturbances rather than focusing on the wedge of earlier research. In this section, I show that the basic conclusion of the paper—that the productivity-MRS wedge is relatively small with a reasonable labor-supply elasticity, provided the employment rate is modeled explicitly and realistically—holds under an alternative procedure that makes the assumption implicit in earlier research that the disturbances for hours and consumption are zero.

Specifically, I solve equations (40) and (41) for the value of $\log v$ implied by the data on consumption of employed workers and their hours of work. By using the actual values of these two series, I am in effect setting their disturbances to zero. The result is a measure of the marginal rate of substitution on the same footing as in earlier studies of the wedge. I then subtract the estimate of $\log v$ from the data on $\log m$ to form the estimate of the log of the wedge.

Using the elasticities in the base case, I find that the wedge has a standard deviation of 1.32 percent. The wedge approach finds much more volatility in the difference between the




Figure 5: Marginal Value of Time and Marginal Product of Labor, Assuming Zero Disturbances in Consumption and Hours

marginal product of labor and the marginal value of time than does the econometric approach of this paper. The reason is that the—to me arbitrary—assumption of zero disturbances in the consumption and hours equations forces the wedge to account for any disturbances that actually belong in those equations. The wedge approach is somewhat rigged against the view that a structural model can account for the movements of key variables without invoking important departures from the efficiency condition that the marginal product of labor, net of taxes, should equal the marginal value of time. Nonetheless, as Figure 5 shows, the estimated wedge is not too big relative to the movements of the marginal product and the marginal value of time.

In the econometric framework of this paper, the disturbances in each of the four equations are uncorrelated with both of the latent variables. Because the latent variables capture cyclical variations quite successfully, the disturbances are not very correlated with the cyclical variables like the employment rate n . In the base specification, the correlation of the

productivity disturbance (the estimate of the wedge) and $\log n$ is -0.21 . Only a small fraction of declines in employment result in increases in the wedge, that is, the marginal value of time dropping below the marginal product of labor. But when the wedge is calculated assuming zero disturbances in consumption and hours, the correlation becomes quite a bit more negative, at -0.49 . Not only is the wedge quite a bit larger, but it is quite a bit more cyclical in the direction suggested by the hypothesis that the marginal value of time drops more than the marginal product of labor in recessions.

12 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 further 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.

13 Concluding Remarks

Data on cyclical variations in the labor market compel either (1) the acceptance that recessions involve a substantial inefficiency because the marginal value of workers’ time falls well below the marginal product of labor, or (2) belief in a high elasticity of employment. I support the second view by introducing a new way to think about the high elasticity. In

this model, much of the elasticity comes from the employment function, an object derived from search equilibrium in the labor market, not directly from workers' preferences. The fall in weekly hours of work recessions is consistent with microeconomic estimates of the Frisch elasticity of labor supply. I remove the larger fall in the employment rate from the domain of the labor supply function to the domain of an extended MP model.

The extension takes place in a framework that dictates the variables that matter for the employment function but is otherwise unrestrictive with respect to the elasticities of the function. Shimer (2005) demonstrated that Nash bargaining implied tiny elasticities. 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 compensation is sticky. The upfront payment y falls when productivity falls, so compensation is cushioned and does not fall as much. Because I take a purely empirical approach, there is nothing surprising or significant in itself in the model's ability to track variations in the employment rate or other measures of tightness.

The primary focus of this paper is the demonstration of the consistency of a model grounded in the theory of household behavior and in the MP class of unemployment models with the actual behavior of the key variables in the U.S. economy. The paper does not claim to reject other theories. What the model interprets as high complementarity of hours and consumption could arise from liquidity constraints that link current earnings to consumption more tightly than under the assumptions made here. Less-than-full insurance against the idiosyncratic risk of unemployment may contribute to the finding of high complementarity as well. With respect to unemployment, I noted earlier that the assumption that the determinants of the employment-payment bargain, y , are limited to those that are payoff-relevant, while often made in game-theoretic models, is not completely compelling. Until theory provides more guidance, it is hard to see how to characterize additional determinants of y and test for their exclusion.

References

- Aguiar, Mark and Erik Hurst, "Consumption versus Expenditure," *Journal of Political Economy*, October 2005, 113 (5), pp. 919–948.
- Alexopoulos, Michelle, "Unemployment and the Business Cycle," *Journal of Monetary Economics*, March 2004, 51 (2), 277–298.
- 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, 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.
- Barro, Robert J., "Long-Term Contracting, Sticky Prices, and Monetary Policy," *Journal of Monetary Economics*, July 1977, 3 (3), pp. 305–316.
- 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.
- Blanchard, Olivier and Jordi Gali, "Real Wage Rigidities and the New Keynesian Model," *Journal of Money, Credit and Banking*, 2007, 39 (1), 35–66.

- Blundell, Richard, Luigi Pistaferri, and Ian Preston, “Consumption Inequality and Partial Insurance,” *American Economic Review*, December 2008, 98 (5), 1887–1921.
- 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.
- Brügemann, Björn and Giuseppe Moscarini, “Rent Rigidity, Asymmetric Information, and Volatility Bounds in Labor Market,” December 2008. Department of Economics, Yale University.
- 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.
- Coase, Ronald H., “The Problem of Social Cost,” *Journal of Law and Economics*, October 1960, 3, 1–44.
- Cohen, Alma and Liran Einav, “Estimating Risk Preferences from Deductible Choice,” *American Economic Review*, June 2007, 97 (3), 745–788.
- 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,” *American Economic Review*, September 2008, *98* (4), 1692–1706.
- Hall, Robert E., “Macroeconomic Fluctuations and the Allocation of Time,” *Journal of Labor Economics*, January 1997, *15* (1), S223–S250.
- , “Cyclical Movements along the Labor Supply Function,” in Katherine Bradbury, Christopher L. Foote, and Robert K. Triest, eds., *Labor Supply in the New Century*, Federal Reserve Bank of Boston Boston 2008, pp. 241–264.
- and Paul R. Milgrom, “The Limited Influence of Unemployment on the Wage Bargain,” *American Economic Review*, September 2008, *98* (4), 1653–1674.
- Hansen, Lars Peter and Ravi Jagannathan, “Implications of Security Market Data for Models of Dynamic Economies,” *Journal of Political Economy*, April 1991, *99* (2), pp. 225–262.
- Harris, John R. and Michael P. Todaro, “Migration, Unemployment, and Development: A Two-Sector Analysis,” *American Economic Review*, 1970, *60* (1), pp. 126–142.
- Hosios, Arthur J., “On the Efficiency of Matching and Related Models of Search and Unemployment,” *Review of Economic Studies*, April 1990, *57* (2), 279–298.
- Hurst, Erik, “The Retirement of a Consumption Puzzle,” February 2008. NBER Working Paper No. 13789.
- Imbens, Guido W., Donald B. Rubin, and Bruce I. Sacerdote, “Estimating the Effect of Unearned Income on Labor Earnings, Savings, and Consumption: Evidence from a Survey of Lottery Players,” *American Economic Review*, September 2001, *91* (4), 778–794.

- 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.
- Kydland, Finn E. and Edward C. Prescott, “Time to Build and Aggregate Fluctuations,” *Econometrica*, November 1982, *50* (6), 1345–1370.
- Low, Hamish, Costas Meghir, and Luigi Pistaferri, “Wage Risk and Employment Risk over the Life Cycle,” September 2008. Institute for Fiscal Studies, London.
- Mehra, Rajnish and Edward C. Prescott, “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, March 1985, *15* (2), pp. 145–161.
- Merz, Monika, “Search in the Labor Market and the Real Business Cycle,” *Journal of Monetary Economics*, 1995, *36* (2), pp. 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.
- Pistaferri, Luigi, “Anticipated and Unanticipated Wage Changes, Wage Risk, and Intertemporal Labor Supply,” *Journal of Labor Economics*, 2003, *21* (3), pp. 729–754.
- Rogerson, Richard, “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics*, 1988, *21*, pp. 3–16.
- and Johanna Wallenius, “Micro and Macro Elasticities in a Life Cycle Model with Taxes,” *Journal of Economic Theory*, 2009. Forthcoming.
- , Robert Shimer, and Randal Wright, “Search-Theoretic Models of the Labor Market: A Survey,” *Journal of Economic Literature*, December 2005, *43*, pp. 959–988.
- Romer, David, *Advanced Macroeconomics*, New York: McGraw-Hill, 2006.

Shimer, Robert, “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,”
American Economic Review, 2005, 95 (1), pp. 24–49.

—, “Convergence in Macroeconomics: The Labor Wedge,” *American Economic Journal: Macroeconomics*, 2008, 1 (1), 280–297.

A Appendix: Research on Properties of Preferences

A.1 Research based on Marshallian and Hicksian labor supply functions

The Marshallian labor supply function gives hours of work as a function of the wage and the individual's wealth. The Hicksian labor supply function replaces wealth with utility. The elasticity of the Marshallian function with respect to the wage is the uncompensated wage elasticity of labor supply and the elasticity of the Hicksian function is the compensated wage elasticity. Both are paired with consumption demand functions with the same arguments.

For simplicity, I will discuss the relation of the Marshallian and Hicksian functions to the Frisch functions used in this paper with a normalization such that the elasticities are also derivatives. I consider the properties of the functions at a point normalized so that consumption, hours, the wage, and marginal utility λ are all one. In this calibration, non-wage wealth is taken to be zero; this is not a normalization. The research I consider treats wage changes as permanent, in which case one can examine a static Marshallian labor supply function where wealth is replaced by permanent income.

From the budget constraint,

$$c(\lambda, \lambda w) - wh(\lambda, \lambda w) = x, \quad (45)$$

where x is non-wage permanent income, I differentiate with respect to x , replace the derivatives of the Frisch functions with the β elasticities, and set $x = 0$, to find the Marshallian income effect:

$$-\frac{\beta_{h,h} - \beta_{c,h}}{\beta_{h,h} - \beta_{c,c} - 2\beta_{c,h}}. \quad (46)$$

By a similar calculation, the Marshallian uncompensated labor elasticity is

$$\beta_{h,h} - \frac{(\beta_{h,h} - \beta_{c,h})^2 + \beta_{h,h} - \beta_{c,h}}{\beta_{h,h} - \beta_{c,c} - 2\beta_{c,h}}. \quad (47)$$

The Hicksian compensated wage elasticity of labor supply is the difference between the Marshallian elasticity and the income effect:

$$\beta_{h,h} - \frac{(\beta_{h,h} - \beta_{c,h})^2}{\beta_{h,h} - \beta_{c,c} - 2\beta_{c,h}}. \quad (48)$$

The compensated elasticity is non-negative.

Chetty (2006) takes an approach similar to the one suggested by these relations, though without explicit reference to the Frisch functions. 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, $-1/\beta_{c,c}$) 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 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.

The following exercise gives results quite similar to Chetty's: From his Table 1, reasonable values for the income elasticity and the compensated wage elasticity from labor supply estimates in the Marshallian-Hicksian framework are -0.11 and 0.4 . For the income elasticity, the work of Imbens, Rubin and Sacerdote (2001) is particularly informative. The paper tracks the response of earnings of winners of significant prizes in lotteries. It finds an income elasticity of 0.10. The range of values of the Frisch parameters that are consistent with these responses is remarkably tight with respect to the wage elasticity $\beta_{h,h}$. If the elasticity is 0.45 , the complementarity parameter is $\beta_{c,h} = 0$, its minimum reasonable value, and the own-price elasticity of consumption is $\beta_{c,c} = -3.67$, an unreasonable magnitude. The minimum compensated wage elasticity is 0.4 , in which case the own-price elasticity of consumption is -0.4 and the complementarity parameter is 0.4 , at the outer limit of concavity. At $\beta_{h,h} = 0.402$, the other elasticities are $\beta_{c,c} = -0.53$ and $\beta_{c,h} = 0.38$, not too far from the values used in the paper of $\beta_{h,h} = 0.7$, $\beta_{c,c} = -0.5$ and $\beta_{c,h} = 0.3$. The static labor-supply literature is reasonably consistent with the other research considered in this appendix. It is completely inconsistent with compensated or Frisch elasticities of labor supply in the range of one or above.

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 subtracting an amount almost equal to consumption before measuring

the MRS. I am skeptical about applying this approach in 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 (2007) 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 take the most reasonable value of the Frisch own-price elasticity of consumption demand to be -0.5 . 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.

Pistaferri studies only men and most of the rest of the literature in the Frisch framework focuses on men. Studies of labor supply generally find higher wage elasticities for women.

Rogerson and Wallenius (2009) introduce a distinction between micro and macro estimates of the Frisch elasticity, with the conclusion that the two can be quite different. Their vocabulary is different from the one in this paper, so their conclusion does not stand in the way of the philosophy employed here, of building a macro model based on micro elasticities. By micro, they refer specifically to estimating the Frisch elasticity as ratio of the slope of the log of hours of work over the life cycle to the slope of log-wages over the life cycle. They build a life-cycle model where most of the effects of wage variation take the form of changes in the age when people enter the labor force and when they leave, so the slope while in the labor force seriously understates the true Frisch elasticity. A non-convex production technology is

key to the understatement. Although early attempts to measure the Frisch elasticity used the approach that Rogerson and Wellenius consider, the literature I have cited here uses more robust sources of variation.

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 that the elasticity of a family member's consumption with respect to family income is 56 percent, for declines in income related to unemployment of that member. The actual decline in consumption upon unemployment is 14 percent. Low, Meghir and Pistaferri (2008) confirm Browning and Crossley's finding in U.S. data from SIPP.

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. Hurst (2008) surveys this research.

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.

In the parametric preferences considered in Hall and Milgrom (2008), a difference in consumption between workers and non-workers of 15 percent corresponds to a Frisch cross-price elasticity of demand of 0.3, the value I adopt.