Cyclical Movements along the Labor Supply Function *

Robert E. Hall
Hoover Institution and Department of Economics
Stanford University
National Bureau of Economic Research
rehall@stanford.edu stanford.edu/∼rehall

June 12, 2007

Abstract

I decompose cyclical movements of labor input to the aggregate U.S. economy into three components: (1) labor-force participation, (2) the employment rate of participants, and (3) weekly hours of work of the employed. I proceed under the identifying hypothesis—implicit in the work of most schools of modern macroeconomics—that higher-frequency movements of labor input and its components reflect shifts in the demand for labor much more than shifts in supply. Under this hypothesis, I derive an index of higher-frequency shifts in labor demand. I estimate the elasticity of each of the components with respect to the index of demand. The largest elasticity, accounting for about 56 percent of the cyclical movements of labor input, is the employment rate of participants, or, equivalently, the unemployment rate. Participation accounts for 12 percent of cyclical variation and hours for the remaining 32 percent. All of the elasticities tend to be higher for people under the age of 25 and over the age of 55. I conclude that research on the high cyclical elasticity of unemployment—an active area in recent years—is central to the understanding of the cyclical volatility of labor input.

*Prepared for the 52nd Economic Conference of the Federal Reserve Bank of Boston, *Labor Supply in the New Century*, Cape Cod, June 18 to 20, 2007. This research is part of the Economic Fluctuations and Growth program of the National Bureau of Economic Research. Data and programs will be available from stanford.edu/∼rehall

1 Introduction

A consensus in macroeconomics holds that the observed higher-frequency movements in employment and hours of work are movements along a labor-supply function caused by shifts of the labor demand function. Recent thinking has extended this view to include fluctuations in unemployment, so that macroeconomics can speak coherently of movements along an unemployment function caused by shifts in labor demand.

I develop a measurement framework for the movements along the labor supply function and for shifts of labor supply. I review data sources for the U.S. economy and conclude that the household survey is the only source of data that supports a clean set of measures of hours and employment. I note the discrepancy between short-run movements of employment from the household survey and the employer payroll survey, but am unable to make any further contribution to reconciling the puzzle of the higher amplitude of employment fluctuations in the employer survey.

The measurement framework rests on the inference of an underlying single unobserved variable that determines labor supply. In other research, Hall (2007), I have derived this variable from the Frisch framework for labor supply and consumption demand. Here I use an econometric model with a latent variable to infer the variable. It turns out to move closely with unemployment but to have a high correlation with weekly hours as well, though there is much more noise in the measurement of hours from the household survey.

2 Labor Supply and Labor Demand

Figures 1 through 3 describe four different views of changes that occur in the aggregate labor market when labor demand shifts outward. The shift may be the result of improved aggregate productivity, declines in the prices of inputs other than labor, or a favorable shift in the terms of trade. The horizontal axis is total labor input measured in hours per year. The vertical axis is the hourly real wage.

Figure 1 shows the standard neoclassical view. Labor supply is fairly inelastic. The labor market clears at all times at the intersection of supply and demand. A large outward shift in demand raises labor input by a small amount and the wage by a substantial amount. As a theory of fluctuations, the neoclassical view fails in both dimensions, as cyclical fluctuations in hours are large and in wages are small.

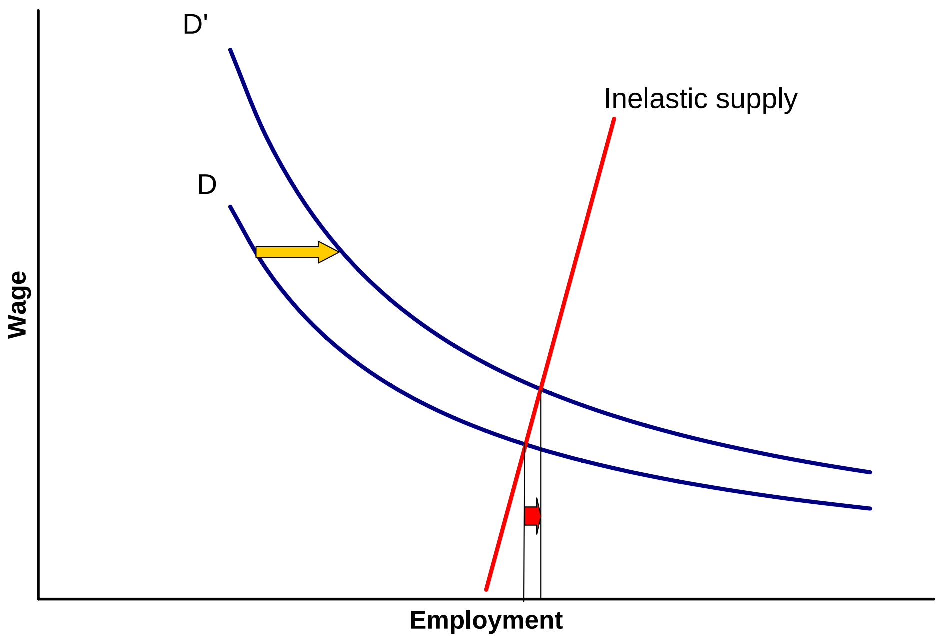

Figure 1. Demand Shift with Inelastic Supply

Figure 2 shows two views with the same properties but very different rationalizations. In the real business cycle model, labor supply is highly elastic—the aggregate labor-supply schedule is essentially flat. Real business cycle theorists, notably Rogerson (1988), have provided analytical foundations for elastic supply and have addressed the important question of why studies of labor supply estimated by investigators at the level of individuals find relatively low elasticities—see Rogerson and Wallenius (2006). Whether micro and macro labor-supply elasticities can be reconciled is under lively debate today.

The other interpretation of Figure 2 takes the horizontal line to express rigidity of the real wage. It is not a standard supply schedule derived from the choices of workers about participation and hours, but the operation of a system of employment governance in which employers choose labor input given a fixed real wage. Theoretical rationalizations of this system of governance have not fared well in recent years, after early enthusiasm about the possibility that contracts made under asymmetric information might take this form. The rigid real wage model carries with it an explanation of unemployment—it is the horizontal distance between the actual level of employment and the labor-supply curve of Figure 1. This is a gap theory of disequilibrium unemployment. Little theoretical work has been done in this framework in recent years, especially in the American context.

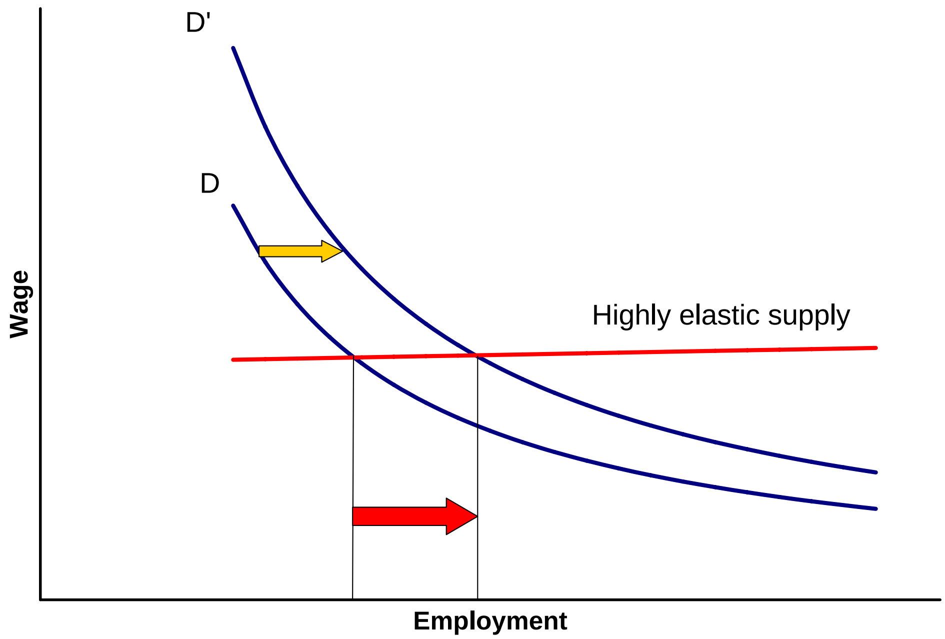

Figure 2. Demand Shift with Elastic Supply or Rigid Wage

Figure 3 illustrates the theory of labor-market fluctuations underlying the measurement work in this paper. Labor supply has its inelastic neoclassical form. Absent frictions in the labor market, shifts in labor demand would cause small changes in hours and large changes in hourly real wages. But the model embodies an economic, equilibrium view of unemployment derived from an explicit consideration of frictions. Unemployment is not a gap but is the result of the interaction of search and matching frictions and compensation determination. The search and matching elements are from Mortensen and Pissarides (1994). As Shimer (2005) demonstrated, search and matching frictions are not enough to explain cyclical fluctuations in unemployment. Shimer's paper set off an enthusiastic investigation of many different modifications of the Mortensen-Pissarides model. It is too early to say which will emerge as the leading explanation.

The Mortensen-Pissarides model describes physical frictions in the labor market but not wage frictions. Wages clear the market in a sense that Shimer's investigation explains. The simplest way to alter the model in a way that makes it consistent with Figure 3 is to introduce what I call equilibrium wage stickiness (Hall (2005a)). With this form of wage rigidity, the extended Mortensen-Pissarides model implies that an outward shift of labor demand, as it tries to push the wage up, will also reduce unemployment substantially. The result, as Figure 3 shows, is an increase in labor input that is much larger than the movement along the labor supply schedule because of the added effect of drawing people out of unemployment and putting them to work.

The line of thought expressed in Figure 3 embodies a full economic treatment of three activities—out of the labor force (specializing in non-market activity), looking for work, and working. In that sense, it is a natural extension of modeling from two activities, as in the first two figures, to three activities. But it is important to understand that the unemployment curve shown in Figure 3 is not just an expression of individual choice about how much time to spend looking for work. Rather, it is the equilibrium of the search and matching process together with wage determination. Individual choices about search and job acceptance are only one component of that equilibrium. For further discussion, see Hall (2007).

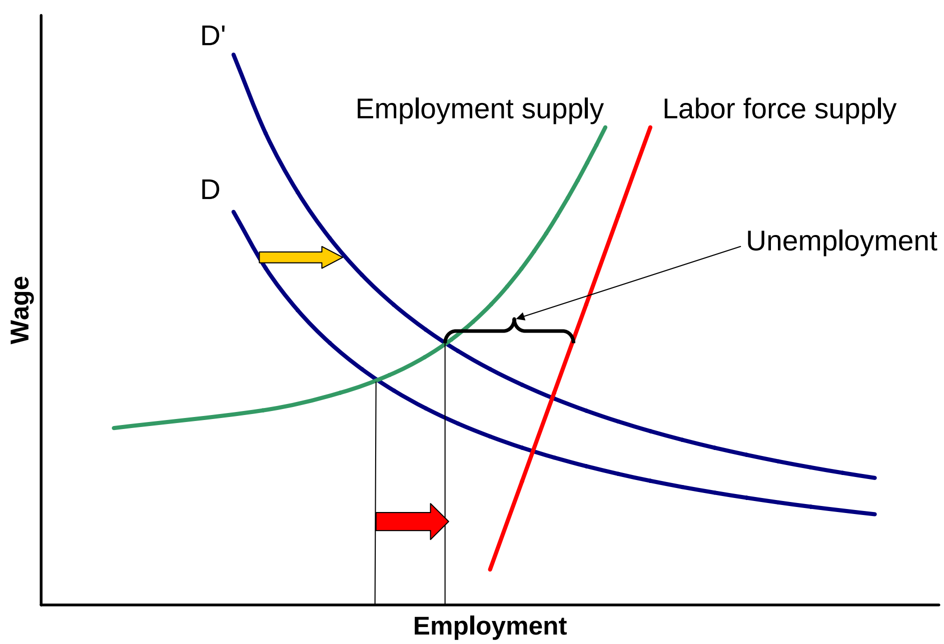

Figure 3. Demand Shift with Equilibrium Unemployment

3 Earlier Work on Cyclical Fluctuations in the Labor Market

All of the earlier research that I have located so far takes unemployment, employment, or output as the measure of the business cycle in the labor market. I am not aware of work that infers an unobserved index.

3.1 Participation

Tella (1964) was an early and influential investigation of higher-frequency movements in aggregate labor-force participation. He considered the relation between the participation rate and the employment-to-population ratio. He focused on higher frequencies by using first differences and found coefficients of 0.40 for men and 0.62 for women. These figures are substantially higher than those found in later work and in this paper, probably because participation is one of the components of the right-hand variable and because he used data from 1948 through 1962.

Wachter (1977) found that participation increased for all age groups for men in tight labor markets with low unemployment, though the effects are small except for teenagers and those over 65. For women, he found similar results for all but the older groups, where participation declines in tight markets.

3.2 Hours

Raisian (1983) studied the cyclical variation of hours per week and weeks per year as a function of experience, using data from the Panel Study of Income Dynamics. He found that the elasticity of hours per week with respect to the employment rate (1 minus the unemployment rate) was 0.30 and that the elasticity of weeks per year was 1.14. The latter figure implies an elasticity of participation of 0.14.

Cho and Cooley (1994) took as a stylized fact of the U.S. business cycle that one quarter of the variation in total hours of work is in hours per worker and three-quarters is in workers per member of the population. These are approximately the relative standard deviations of Hodrick-Prescott filtered hours per worker and employment, given in their Table 2.

4 Framework and Data

The objective of this work is to develop a conceptual framework and corresponding data in which the three dimensions of labor supply—participation, employment rate, and hours—play roles derived from the theory of labor supply and unemployment.

The modern theory that provides the logical starting point for the framework runs as follows: Individuals have preferences defined over hours spent at home, hours of search, and hours of work. Each period (month in this paper) they choose an allocation of hours out of a set of available choices. Hours spent looking for work and hours spent at home are not restricted, but hours at work depend on the jobs available that period—workers do not have unilateral choice over jobs or the hours of jobs.

Individuals' choices map into observed activities. The CPS uses certain important conventions in assigning individuals to activities. Although the CPS is a monthly survey, it uses a combination of time periods in the assignment process. The first convention is that work trumps any other activity, in the sense that a person who worked even one hour in the week before the survey is counted as employed, notwithstanding any other time spent at home or in job search. The second convention is that a person not recorded as employed is recorded as unemployed if the person was not employed in the previous week but made any of a variety of designated types of efforts to find a job in the preceding four weeks. Those who fail to meet the criteria for employment or unemployment are counted as out of the labor force.

Recently the launching of the American Time Use Survey will provide a far more complete view of the allocation of household time. The new survey focuses on measuring all uses of time rather than assigning individuals to categories based on partial measures. However, the size of the ATUS sample is not large enough to support good national estimates of labor-market status monthly.

Flinn and Heckman (1983) make the reasonable proposal that the unemployed should be taken to be non-working individuals who have a probability of finding work in the coming period above a designated threshold. The CPS definition of unemployment appears to implement a rough approximation to the Flinn-Heckman definition. Along with Flinn-Heckman, the CPS definition does not classify people as unemployed if they have decided that no job realistically likely to become available would be superior to non-work activities. The CPS has a separate category for these people, often called discouraged workers.

The home activities that occupy all individuals, working or not, include home production as well as leisure. As the ATUS shows, these activities include shopping, cooking, and caring for others, together with sleep and pure leisure, such as watching TV or socializing.

5 Measuring Employment

The Bureau of Labor Statistics runs two independent surveys aimed at determining a seemingly simple concept, the number of people at work at a given moment. In addition to the CPS count of employment, the BLS surveys employers about the number of workers on their payrolls. Almost from the beginning of the household survey, economists noted cyclical discrepancies between the two surveys—the payroll measure of employment rises faster in booms and falls faster in recessions than does the CPS measure. Economists affiliated with the party in power find reasons to praise the CPS measure during recessions—especially the most recent one—while others cite the payroll survey as the more accurate description of the ravages of the recession.

Figure 4 compares employment counts from the two sources. It shows the raw ratio of the payroll count to the household count together with its higher-frequency component. The latter comprises the residuals from a regression of the ratio on a fourth-order polynomial in time. The payroll count rose irregularly from 82 percent of the CPS level in 1959 to 97 percent at its maximum at the end of the 1990s and then fell to its current level of 94 percent.

The higher-frequency component of the ratio is conspicuously correlated with the business cycle. In each recession, the payroll count falls by one to three percent of the CPS count. The decline was particularly large in the most recent recession. It was large in the worst postwar recession, in 1981-82, but not as large in the other deep recession, 1973-75.

The cyclical discrepancy remains almost entirely unexplained. Table 1 shows a dissection of the conceptual differences between the two employment measures based on Bowler and Morisi (2006). The top line is the percentage shortfall of the payroll count from the CPS count. During the expansion years, 1994 through 2000, the shortfall shrank and then

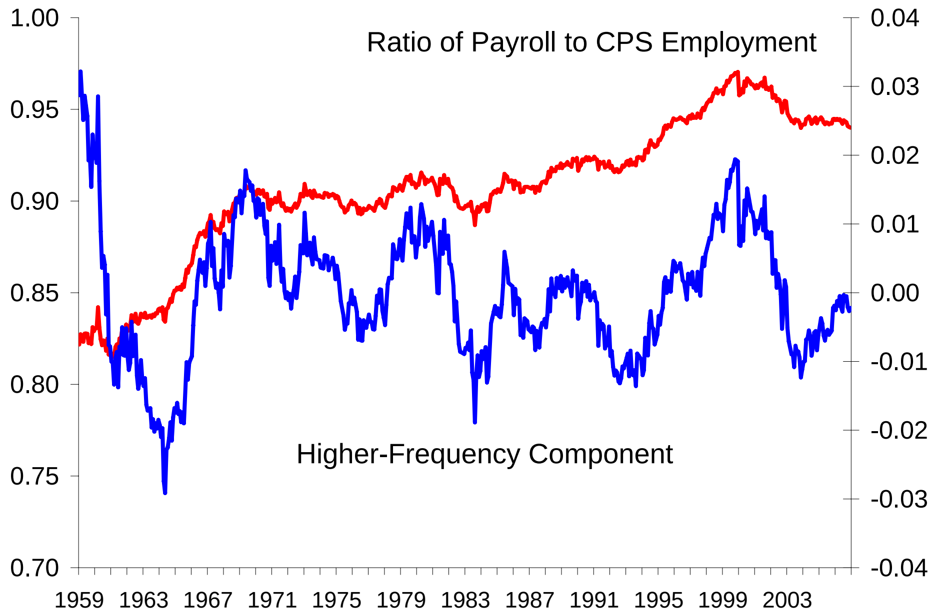

Figure 4. Ratio of Payroll Employment to CPS Employment

expanded during the recession and following years, 2000 through 2004. The column headed *Cycle* is the percentage growth from 2000 through 2004 plus two-thirds of the shrinkage from 1994 though 2004. This figure is zero if the figures to the left grow linearly with time and is positive if the figures to the left fall during the expansion and rise during the contraction, as the payroll shortfall plainly does.

The lines labeled "Components from CPS" report components of CPS employment that are conceptually different from the payroll data, stated as percentages of the total CPS employment count. The cycle measure is given in the right column for each adjustment. A positive cycle measure means that the component helps explain the pro-cyclical discrepancy between the payroll and CPS counts.

The first of the conceptual differences between the two employment measures is that the CPS includes the self-employed and wage-earners in agriculture, whereas the payroll data exclude agricultural employment. The cycle measure is negative for this component—the strong labor market of 2000 resulted in an upward deflection in agricultural employment. This phenomenon only deepens the mystery of the cyclical discrepancy, as it would make the CPS more cyclical than the payroll data by itself.

The second adjustment shows an important source of the cyclical discrepancy—self

	1994	2000	2004	*Cycle*
Shortfall of payroll jobs	7.13	3.73	5.58	4.12
Components from CPS				
Agriculture	2.77	2.47	1.60	-0.67
Non-ag self employed	7.32	6.40	6.80	1.00
Non-ag unpaid family workers	0.11	0.08	0.06	0.01
Private household workers	0.78	0.66	0.56	-0.01
Unpaid absence	1.62	1.47	1.38	0.01
Multiple jobholders	-5.51	-5.20	-5.07	-0.07
Total components	7.09	5.88	5.33	0.27

Table 1. Components of CPS Employment Related to Conceptual Differences from Payroll Data

employment—which declined sharply as a fraction of CPS employment during the expansion and rose a bit during the recession and aftermath. The payroll data exclude the self-employed.

The other four components shown in Table 1 account for trivial percentages of the cyclical movements. Unpaid family workers and private household workers, included in the CPS but excluded from the payroll data, are tiny fractions of total employment and have no cyclical component. People who have jobs but are not currently being paid—counted in the CPS but not the payroll data—make no contribution to the cycle. And second jobs—counted twice in the payroll count of jobs but only once in the CPS count of employed people—make a small contribution in the wrong direction to explain the discrepancy

Notice that the components almost perfectly match the CPS and payroll counts in the years 1994 and 2004, but result in an excess of payroll employment in the peak year, 2000.

According to Table 1, the cyclical discrepancy in employment counts between the two surveys is almost completely a mystery. The table covers all but one of the important conceptual differences between the surveys, the length of the reference period. In the CPS, a person who worked one hour or more in the week before the survey counts as employed. The payroll survey counts the number of people on the payroll at any time during the pay period that includes the 12th of the month. My impression is that pay periods are generally two weeks or half a month (I believe that the law prohibits longer pay periods for most workers, but this is a subject for further research).

The relation between the length of the pay period and the overstatement of snapshot unemployment by the payroll data is simple: The overstatement is the weekly rate of new hires times the number of weeks in the pay period. Hall (2005b) discusses evidence on cyclical variation in the new hire rate. The JOLTS survey measures the rate directly and shows little variation in the only cycle that has occurred since it was launched in 2000. The separation rate is an excellent proxy for the new hire rate—the two differ only by the rate of change of employment, which is insignificant at all times in comparison to the levels of new hires and separations. The CPS has measured total separations since 1994, so it too includes only the most recent cycle. Figure 2.4 in Hall (2005b) shows that the monthly separation rate fell by about half a percent from the strong labor market of 2000 to the weak market of 2003. The weekly rate thus fell by a little over a tenth of a percent. Even if the pay period is 4 weeks or a month, cyclical variations in the overstatement caused by longer pay periods is tiny in relation to the observed discrepancy in the cyclical behaviors of CPS and payroll employment.

Absent an understanding of the source of the extra cyclical movements of the payroll employment data, it is not possible to use the data in the three-activity framework normally used in research on labor-market dynamics. The difficulty is that the fractions of the population in the three activities—out of the labor force, unemployed, and working—must sum to one. The payroll survey provides no measure of the first two. One would have to adjust the fractions from the CPS for out of the labor force and unemployed to satisfy adding up. There is no basis for making the fraction out of the labor force and the fraction unemployed more countercyclical than is reported in the CPS, but these adjustments would be needed to incorporate the payroll employment data.

6 Data on Hours

The CPS asks the respondent (often not the actual worker), a question like the following, "So, for last week, how many hours did he actually work at his job?" (the computer tailors the question to the individual worker). The process gathers hours separately by job for multiple job-holders. The respondent decides what constitutes an hour of work—whether it includes breaks, setup time, and the like.

The CPS measure of hours drops dramatically at random, when a holiday falls in the reference week. The choice of the reference week as the one including the 12th of the month dodges Thanksgiving, Christmas, and several other holidays, but cannot exclude every holiday. Monthly plots of hours show these drops.

The BLS also provides a comprehensive measure of hours based primarily on the payroll data, extended to agriculture and self-employment with CPS data (Major Sector Productivity and Costs Index or MSPC, bls.gov/lpc/home.htm). The payroll survey determines hours paid per job from employers. The MSPC restates the results on the basis of hours worked rather than hours paid, using another survey that collects both. The MSPC also uses CPS hours for workers not covered by the payroll survey. The result is presented as an index. Although the MSPC measure of hours is mainly hours per job rather than hours per worker, there is so little cyclical variation in jobs per worker that the distinction is unimportant for the study of cyclical phenomena.

Figure 5 compares the two sources of data on hours. Hours as measured by the MSPC fell by about 10 percent relative to CPS hours from 1959 to 2005. I am not aware of any discussion or explanation of this behavior. As in Figure 4, I also show the higher-frequency component. It is relatively small and not conspicuously cyclical. Apart from the differing trend, there seems no important discrepancy between the measures.

7 The Single Driving Force of Movements along the Labor Supply Function

The consensus of modern macroeconomics is that shifts of labor supply are not a significant driving force of the business cycle. Rather, productivity shocks, oil shocks and other shifts in the terms of trade, and changes in other factor prices move workers along their labor supply schedules. Hours of work reflect variations over time in the current payoff to work relative to the value of other activities. Choices about whether to participate in the labor market also reflect a similar choice. Both of these derive from perfectly standard models of labor supply.

A more recent extension, deriving from the work of Mortensen and Pissarides (1994), has

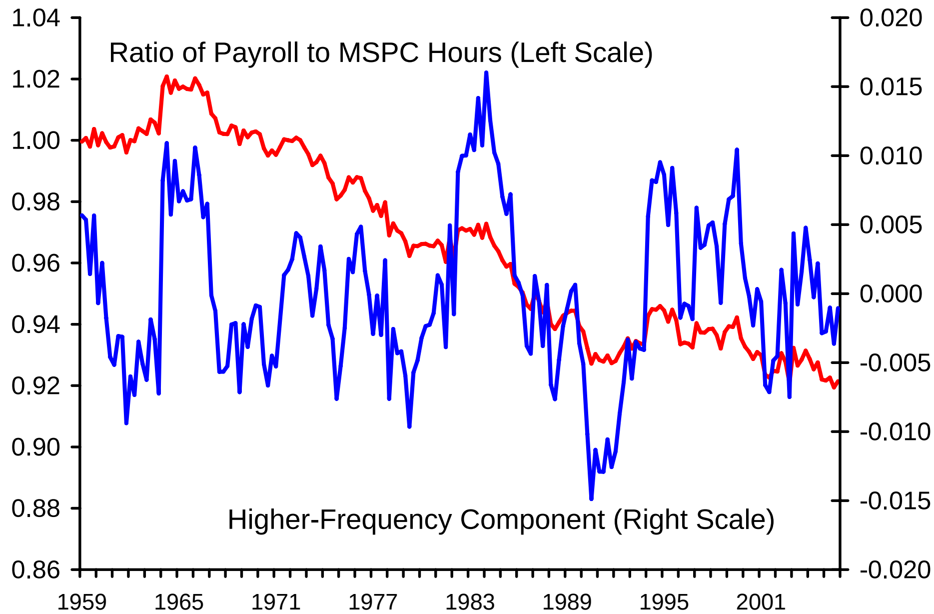

Figure 5. Ratio of MSPC Hours to CPS Hours

developed a model of the third use of time, job search, that responds to the same factors. Hall (2007) shows how unemployment fits into a model of labor-market fluctuations. That paper derives two indexes that jointly capture the driving forces of labor-market fluctuations. One index describes the overall well-being of households, based on expectations of future earnings. The other describes the current state of the labor market. The two are highly correlated, so it is a reasonable approximation to treat the labor market as having a single driving force, the approach taken here.

The point to be taken from this line of thought is that the single force drives all three key measures—participation, unemployment, and hours of work. The single force is the current position of the labor demand function in relation to its typical level. When demand is unusually strong, participation rises, unemployment falls, and hours of work rise. The rest of this paper will derive a measure of the single driving force from the multiple indicators and measure the relative cyclical sensitivities of participation, unemployment, and hours.

The model underlying this work—and the conclusion about a single driving force—does not necessarily rest on any ideas of the kind usually labeled Keynesian. In fact, all of the conclusions except the magnitude of the fluctuations will hold in a neoclassical, real-business-cycle model, extended only in the direction of Mortensen-Pissarides. Although the easiest way to explain the observed amplitude of the responses of labor-market variables to the driving force is with sticky wages, it is an open and very interesting question whether other mechanisms may be involved as well or if any wage or price stickiness is needed.

To derive a measure of the single driving force, I use three monthly measures that track the business cycle. Two are from the labor market: unemployment and hours. To put unemployment in a form that makes it interchangeable (except for sign) with employment per participant in a log-additive framework, I measure unemployment as the negative of the log of the employment rate. The third measure in the cyclical system is real disposable personal income per capita (NIPA Table 2.6).

The following econometric setup enables the measurement of the common driving force, z_t:

$$y_t = \gamma_y z_t + \tau_y(t) + \epsilon_{y,t} \tag{1}$$

$$u_t = \gamma_u z_t + \tau_u(t) + \epsilon_{u,t} \tag{2}$$

$$h_t = \gamma_h z_t + \tau_h(t) + \epsilon_{h,t} \tag{3}$$

Here y_t is log real income, u_t is the unemployment rate, and h_t is the log of weekly hours of work. The γs are the loading factors of the observed variables on the unobserved driving force z_t. They are interpreted as elasticities of the component with respect to the cyclical driving force. The $\tau(t)$ functions capture slower-moving non-cyclical determinants of the observed variables and the ϵs are the idiosyncratic higher-frequency movements not associated with the cyclical driving force z_t—they are assumed to be uncorrelated with z_t. I assume that z_t, whose units are arbitrary, has a variance of one. I also assume that γ_u is negative, so z_t is procyclical.

I specify the $\tau(t)$ functions as fourth-order polynomials in time. I also include seasonal dummies for hours because the data are not seasonally adjusted. The model has two sets of moment conditions. The first are standard regression conditions—orthogonality of the time variables in the τ functions with the disturbances. The regression part—like all regressions—has the same number of moment conditions and unknown parameters and is exactly identified.

The second set of moment conditions describes the latent-variable structure of the disturbances. This part of the model has 6 observed moments: 3 variances of the ϵ disturbances, V_y, V_u, and V_h, and 3 covariances, $C_{u,y}$, $C_{u,h}$, and $C_{h,y}$. It has 6 unknown parameters, γ_y, γ_u, γ_h, σ_y, σ_u, and σ_h, where the last three are the standard deviations of the idiosyncratic components. The latent-variable model is exactly identified and has the following moment conditions:

$$C_{u,y} = \gamma_u \gamma_y, \tag{4}$$

$$C_{u,h} = \gamma_u \gamma_h, \tag{5}$$

$$C_{h,y} = \gamma_h \gamma_y, \tag{6}$$

$$\sigma_y^2 = V_y - \gamma_y^2, \tag{7}$$

$$\sigma_u^2 = V_u - \gamma_u^2, \tag{8}$$

and

$$\sigma_h^2 = V_h - \gamma_h^2. \tag{9}$$

The overall model is exactly identified. Its moment conditions are block-triangular—I can solve for the regression parameters first and then derive the latent-variable parameters. The first step is to estimate regressions of the three variables on the components making up the τ functions (powers of t and seasonal dummies). I denote the residuals from these regressions as $\hat{y}_t$ and similarly for u and h. The variances and covariances in the moment conditions for the latent-variable model then refer to the hatted residuals.

From the moment conditions, I derive

$$\gamma_y = \sqrt{\frac{C_{u,y} C_{h,y}}{C_{u,h}}}, \tag{10}$$

with the square root taken as positive. The remaining parameters come directly from the moment conditions. Notice that the model imposes a condition on the signs of the covariances—the expression under the square root is non-negative. In addition, the implied values of the squared values of three σ parameters must be non-negative.

To infer the values of the single driving force z_t, I use the projection of z on the observed variables; that is, the fitted values of the regression of z on those variables. The regression coefficients are the inverse of the covariance matrix of the variables (observed) multiplying the vector of covariances of z and the variables. The covariances are just the estimated parameters γ, because the variance of z is one.

Table 2 shows the results of these calculations. The top panel shows the variances and covariances of the residuals from the preliminary regressions. The unemployment rate is in percent and real income and hours in 100 times their natural logs. Hours and unemployment have about the same variances but the variance of real income, around its lower-frequency trend, is quite a bit higher. The covariances of the three variables are as expected—unemployment is countercyclical and income and hours are procyclical.

	Unemployment	*Real income*	*Hours*
Moments			
Unemployment	1.14	-1.33	-0.54
Real income		4.20	0.78
Hours			1.64
Parameters			
Loading on *z*, γ	-0.96	1.39	0.56
Variance, σ^2	0.22	2.27	1.32
Coefficients for *z, a*	-0.696	0.097	0.068

Table 2. Inference of Cyclical Driving Force from Data on Unemployment, Real Income, and Hours

The first line in the lower panel of Table 2 shows the loading coefficients, γ, for the three variables. Unemployment has a loading coefficient on the cyclical driving force of just under one. The next line shows that unemployment has a fairly low idiosyncratic movement—the variance of its non-cyclical higher-frequency movements is only 0.22. Real income loads on the cyclical component with an elasticity of 1.39 and has an idiosyncratic variance of 2.27, about half its total variance of 4.20. Hours load on the cyclical driving force with an elasticity of 0.56, leaving a large idiosyncratic variance of 1.32 out of its total variance of 1.64.

The a coefficients for extracting the implied time series for the driving force z show that the optimal inference places a large negative coefficient on unemployment and smaller positive coefficients on real income and hours. Figure 6 shows the index $\tilde{z}_t$.

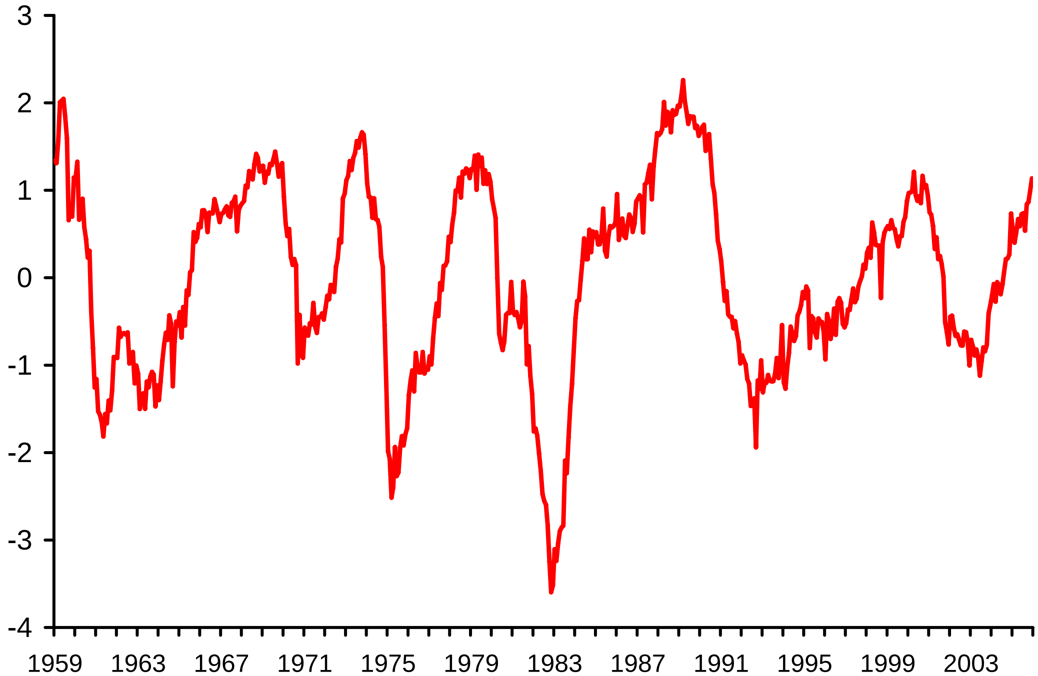

Figure 6. Index of the Driving Force

8 Cyclical Sensitivity of Participation, Unemployment, and Hours

I am now equipped to answer the basic question of the cyclical sensitivities of participation, unemployment, and hours. Table 3 shows the loading factors for the three dimensions of work effort on the driving force, z. For employment, the coefficient is the positive value of the one shown in Table 2 and for hours, it is the value shown there. For participation, not included in the earlier model, I show the coefficient of the regression of log of the CPS participation rate on the inferred measure, $\tilde{z}$; the regression also includes the fourth-order polynomial in t as in the earlier regressions. For all three, I measure the standard error from that type of regression. The total loading shown at the bottom is just the sum of the loadings of the three components.

The first line of Table 3 shows the small but statistically unambiguous cycle in participation. Recall that the units of the cyclical driving forces are standard deviations of cyclical

	Loading	*Standard error*	*Percent of total*
Participation	0.197	(0.008)	11.6
Employment	0.957	(0.008)	56.5
Hours	0.534	(0.099)	31.5
Total	1.696	(0.075)	

Table 3. Loading Coefficients for the Three Dimensions of Work on the Cyclical Driving Force

movements in the labor market. A one standard deviation tightening of the market raises participation by 0.2 percent. Because the level of participation is around 60 percent, this is about 0.12 percentage points. The response of participation is 11.6 percent of the total response of labor input.

Employment, shown in the second line, is a bit more than half of the total cyclical variation. A tightening of the market by one standard deviation raises employment and lowers unemployment by just under one percentage point.

Weekly hours, shown in the third line, account for a third of total cyclical variation in labor input.

Tables 2 and 3 deal with labor measures per person. Table 4 considers the employment count, the product of population, participation, and the employment rate. I will not consider the employment count in the rest of the paper, but it does permit a further consideration of the difference between the CPS and payroll data, as the latter take the form of employment counts only, without the breakdown into population, participation, and the employment rate.

Table 4 needs to put population on the same footing as the other measures, as the higher-frequency component obtained as residuals from the regression of the log of population on a fourth-order polynomial in time. Then, to reconcile the CPS measures including population with the payroll measure in the framework of the paper, it needs to measure the loading of population on the cyclical driving force. The loading is -0.147 with a standard error of 0.018. How can population be countercyclical? Obviously population does not respond to the forces that cause the cycle, but population swings could be a contributor to the cycle. This hypothesis seems to be part of the explanation—unusually low population growth led

	Loading	*Standard error*
Population	-0.147	(0.018)
Participation	0.197	(0.008)
Employment rate	0.958	(0.008)
CPS employment	1.006	(0.029)
Payroll employment	1.512	(0.048)

Table 4. Cyclical Loadings for Number of Employed Workers

to a tighter labor market in the late 1960s and unusually high growth to a slacker market in the period containing the weakest labor market, 1973 through 1983. Another part of the explanation is discontinuous increases in the population estimates used in the CPS at the beginning of 1990 and 2000, both near cyclical peaks.

The loading of the CPS employment count on the cyclical driving force is very close to one. The loading is only slightly higher than the loading for employment per participant in Table 3, because the negative effect of population offsets the positive role of participation.

By contrast, the loading of the log of payroll employment on the cyclical driving force is much higher, at 1.512. The stronger cycle in payroll employment shows through prominently in the framework of the cyclical driving force, even though the driving force is derived completely independently of the payroll data.

Although higher-frequency changes in the working-age population are shifts of labor supply rather than movements along a labor-supply function, the movements in participation, unemployment, and hours considered in this paper are movements along their respective functions. The discovery that population movements are part of the driving force of those movements is quite consistent with the overall framework of this paper.

9 Cyclical Responses by Demographic Groups

Table 5 breaks down the responses shown in Table 3 by age and sex, to the extent that the data are available from the BLS. Long historical tabulations of the data are incomplete, though the important features of the differences among demographic groups are visible and in accord with prior beliefs. The hours data in Table 5 begin in June 1976.

	Sex	*Age*	*Loading*	*Standard error*
Participation	Male	16 to 24	0.285	(0.060)
		25 to 54	0.123	(0.013)
		55+	0.294	(0.055)
	Female	16 to 24	0.571	(0.080)
		25 to 54	0.123	(0.039)
Unemployment	Male	16 to 24	-1.920	(0.024)
		25 to 54	-0.910	(0.015)
	Female	16 to 24	-1.292	(0.023)
		25 to 54	-0.729	(0.011)
Hours	Male	16 to 19	1.911	(0.089)
		20 to 24	0.992	(0.048)
		25+	0.588	(0.048)
	Female	16 to 19	1.787	(0.090)
		20 to 24	0.954	(0.055)
		25+	0.482	(0.059)

Table 5. Loading Coefficients for Participation, Unemployment, and Hours by Age and Sex

Table 5 confirms that the participation elasticity is higher for younger (under 25) and older (over 54) workers and higher for women than for men among younger workers. The more elastic groups contain a larger fraction of people who are close to the margin between work and specializing in non-work activities, primarily activities at home and attending school. Unemployment is more sensitive to the driving force among men and among younger workers.

The elasticities of hours with respect to the cyclical driving force are slightly lower for women than for men. For both sexes, the response of hours is much higher for the youngest workers.

10 Interpretation

I have not tested the dogma of modern macro analysis of the labor market—that shifts in labor demand account for most of the cyclical variation in labor input. But it holds up well in the analysis of this paper. First, all three components, participation, the employment rate for participants, and hours per week of workers, respond positively to my measure of cyclical shifts in labor demand. Because these shifts are transitory, they involve mostly substitution effects. Basic labor-supply theory shows that the substitution effect in participation and in hours per worker should be positive. The extended Mortensen-Pissarides theory requires the substitution effect for the employment rate to be positive as well.

More than half of the extra labor input in a cyclical upswing is drawn from unemployment. No model of the cycle in the labor market can claim any realism unless it takes this finding seriously. It is inappropriate to lump those assigned by the CPS to unemployment together with those found to be out of the labor force, because the unemployed are much more likely to be employed a month later.

Research trying to explain the high cyclical elasticity of unemployment has made exciting advances in the past few years, but a great deal remains to be done.

References

Bowler, Mary and Teresa L. Morisi, "Understanding the Employment Measures from the CPS and CES Surveys," *Monthly Labor Review*, February 2006, pp. 23–38.

Cho, Jank-Ok and Thomas F. Cooley, "Employment and Hours over the Business Cycle," *Journal of Economic Dynamics and Control*, 1994, *18*, 411–432.

Flinn, Christopher J and James J. Heckman, "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?," *Journal of Labor Economics*, January 1983, *1* (1), 28–42.

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

____ , "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years," *NBER Macroeconomics Annual*, 2005, *20* (1), 101–137.

____ , "Sources and Mechanisms of Cyclical Fluctuations in the Labor Market," June 2007. stanford.edu/∼rehall.

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.

Raisian, John, "Contracts, Job Experience, and Cyclical Labor Market Adjustments," *Journal of Labor Economics*, April 1983, *1* (2), 152–170.

Rogerson, Richard, "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 1988, *21*, 3–16.

____ and Johanna Wallenius, "Micro and Macro Elasticities in a Life Cycle Model With Taxes," December 2006. Arizona State University.

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

Tella, Alfred, "The Relation of Labor Force to Employment," *Industrial and Labor Relations Review*, April 1964, *17* (3), 454–469.

Wachter, Michael L., "Intermediate Swings in Labor-Force Participation," *Brookings Papers on Economic Activity*, 1977, (2), 545–576.