Cyclical Movements along the Labor Supply Function

Robert E. Hall

I. 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 theoretical thinking has extended this view to include fluctuations in unemployment, so that macroeconomists can speak coherently of movements along an unemployment function caused by shifts in labor demand.

I develop an empirical framework for measuring the movements along the labor supply function and for measuring 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. While recognizing the discrepancy between short-run movements of employment from the household survey and the employer payroll survey, at this point in my study I 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 presented in this paper rests on the inference of an underlying single unobserved variable that determines labor supply. At the paper's end I discuss how this single driving force is related to modern macroeconomic labor-market theory. Here I use an econometric model with a latent variable to infer the unobserved variable, which turns out to move closely with unemployment. This variable has a high correlation with weekly hours as well, though there is much more noise in the measurement of hours from the household survey than in the employer payroll survey.

II. Labor Supply and Labor Demand

Figures 5.1 through 5.3 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 5.1 shows the standard neoclassical view of the labor market, which holds that 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 describing both dimensions of demand and supply, as cyclical fluctuations in hours are large but small for wages.

Figure 5.2 shows two views with the same properties but very different rationalizations. In the real business cycle model, labor supply is highly

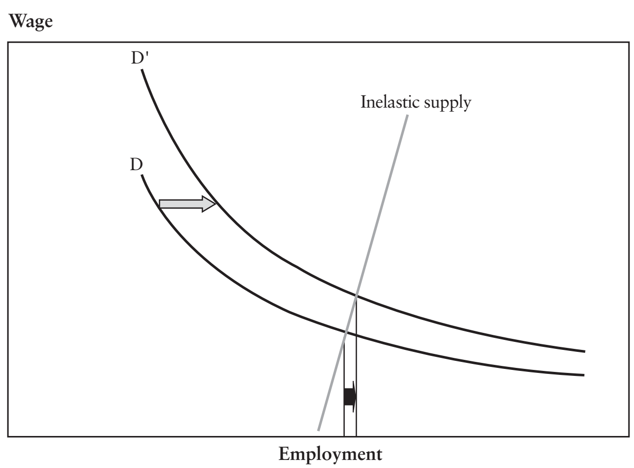

Figure 5.1
Demand Shift with Inelastic Supply

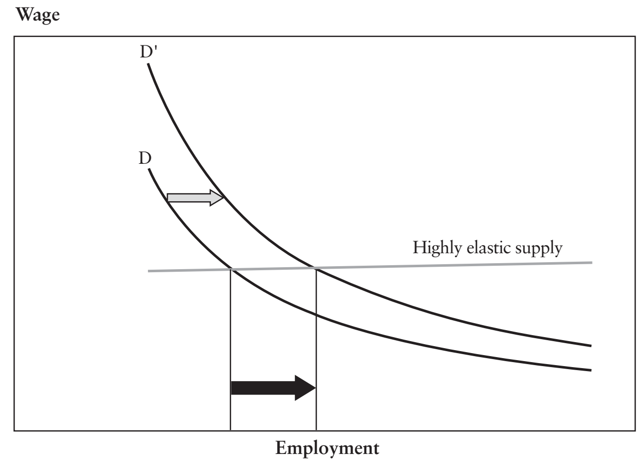

Figure 5.2
Demand Shift with Elastic Supply or Rigid Wage

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 at the level of individuals find relatively low elasticities—see Rogerson and Wallenius (2007). Whether microeconomic and macroeconomic estimates of labor-supply elasticities can be reconciled is a lively topic of debate today.

The other interpretation of Figure 5.2 takes the middle 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 a level of labor input given a fixed real wage rate. 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 5.1. This rigid wage model is a gap theory of disequilibrium unemployment. Yet little theoretical work has been done in this framework in recent years, especially in the American context.

Figure 5.3 illustrates the theory of labor-market fluctuations underlying the measurement work discussed in this paper. As mentioned earlier, 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 used in this paper 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 (such as the search and hiring process on the part of firms and workers) in the labor market but not wage frictions. Wages clear the market in a sense that they are the result of an unrestricted voluntary bargain between employers and workers. The simplest way to alter the Mortensen-Pissarides model in a way that makes it consistent with Figure 5.3 is to introduce what

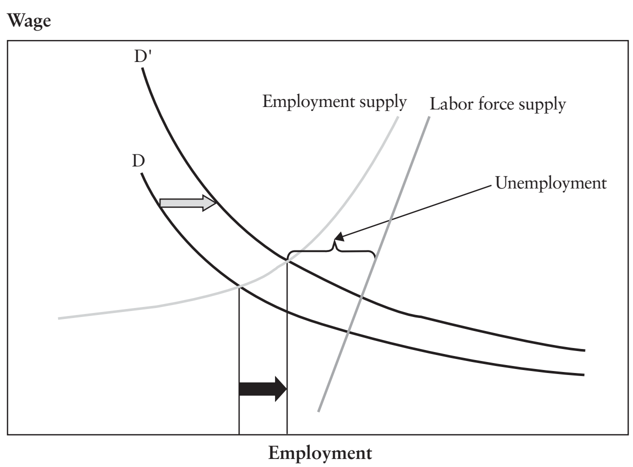

Figure 5.3
Demand Shift with Equilibrium Unemployment

I call equilibrium wage stickiness (see 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 5.3 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 5.3 embodies a full economic treatment of an individual's three possible activities related to the job market—remaining completely out of the labor force (specializing in nonmarket activity), looking for work, and being employed. In that sense, it is a natural extension of modeling from two activities, as illustrated in the first two figures, to three activities. But it is important to understand that the unemployment curve shown in Figure 5.3 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 (2008).

III. Earlier Work on Cyclical Fluctuations in the Labor Market

All of the earlier macroeconomic 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.

Participation

Tella (1964) was an early and influential investigator of higher-frequency movements in aggregate labor-force participation. He considered the relation between the participation rate and the employment-to-population ratio, focusing on higher frequencies by using first differences and finding 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, when women's labor force participation was lower than it has been in subsequent decades.

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

Hours

Raisian (1983) studied the cyclical variation of hours per week and weeks per year as a function of experience in the labor market, 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 the remaining three-quarters is in workers per member of the working-age population. These are approximately the relative standard deviations of Hodrick-Prescott filtered hours per worker and employment.

IV. Framework and Data

The objective of this paper, which is part of a larger research agenda, 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 macroeconomic theory of labor supply and unemployment.

The modern theory that provides the logical starting point for this paper's measurement framework runs as follows: Individuals have preferences defined over hours spent at home, hours of search, and hours of work. Each period (months are used 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 spent at work depend on the jobs available that period—workers do not have unilateral choice over jobs or the hours of jobs.

The macroeconomic theory of unemployment that has emerged from Mortensen and Pissarides's pioneering work focuses on the interacting behavior of job-seekers and employers. Hall (2007) gives an extended discussion of a generalization of their model. The job-creation efforts of employers control the unemployment rate. Employers respond to the job-creation incentive defined by the gap between the marginal product of labor (that is, labor demand) and the wage they expect to pay a newly hired worker. Wage flexibility is a key issue. If an increase in labor demand results in an equal increase in the wage, job creation and thus unemployment remain unchanged. The stickier the wage, the stronger is the decline in unemployment in response to an increase in labor demand. Some bargaining models imply sticky wages—see Hall and Milgrom (2008). Another source of wage stickiness is efficiency wages—see Alexopolous (2004).

Individuals' and firms' choices map into the three observable labor market activities for individuals. The Current Population Survey (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 a job search. The second convention is that a person not recorded as employed is recorded as unemployed if the person was not participating in the workforce 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.

The recent launch of the American Time Use Survey (ATUS) 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 monthly labor-market status.

Flinn and Heckman (1983) make the reasonable proposal that the unemployed should be taken to be non-working individuals who have a probability, above a designated threshold, of finding work in the coming period. 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 these individuals have decided that no job realistically likely to become available would be superior to engaging in non-work activities. The CPS has a separate category for these people, often called discouraged workers.

The home activities that occupy all individuals, employed 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 television or socializing.

V. Measuring Employment

The Bureau of Labor Statistics (BLS) runs two independent surveys aimed at determining a seemingly simple concept: the number of people at work in the United States 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 expansions and falls faster in recessions than does the CPS measure. The report of the President's Committee to Appraise Employment and Unemployment Statistics commented extensively on the issue in 1962. Economists affiliated with the party in power find reasons to praise the CPS measurement during recessions—especially the most recent one—while others cite the payroll survey as the more accurate description of the ravages of the downturn.

Figure 5.4 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 comparing the household survey with the employer survey is conspicuously correlated with

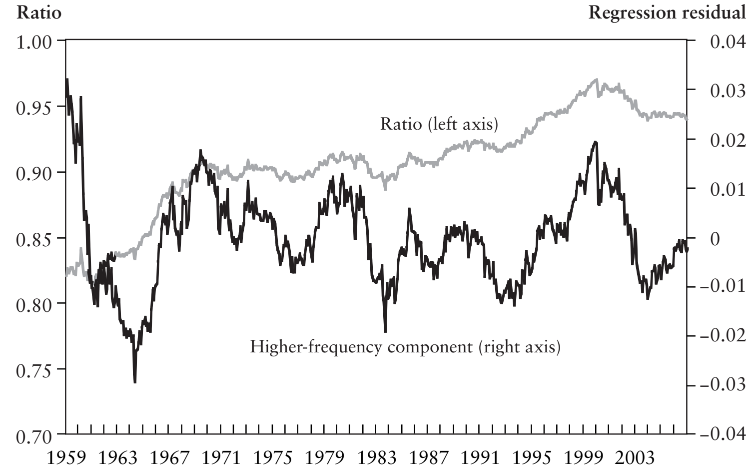

Figure 5.4
Ratio of Payroll Employment to Current Population Survey Employment
Source: U.S. Bureau of Labor Statistics, U.S. Census Bureau, and author's calculations.
Note: The higher-frequency component comprises the residuals from a regression of the ratio on a fourth-order polynomial in time.

the business cycle. In each recession, the payroll count falls by 1 to 3 percent of the CPS count. The decline was particularly large in the most recent recession. It was large in the worst postwar recession that took place in 1981–1982, but not as large in the other deep postwar recession that occurred between 1973 and 1975.

The cyclical discrepancy between these two measures remains almost entirely unexplained. Table 5.1 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 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.

Table 5.1
Components of CPS Employment Related to Conceptual Differences from Payroll Data

	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

The entries collectively labeled "Components from CPS" report components of CPS employment measures that are conceptually different from the payroll data, stated as percentages of the total CPS employment count. The business-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 employer payroll measurement and CPS household counts.

The first of the conceptual differences between the two employment measures is that the CPS one 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 measurement by itself more cyclical than the payroll data.

The second adjustment shows an important source of the cyclical discrepancy—self-employment—which declined sharply as a fraction of CPS employment during the 1994–2000 expansion, and rose a bit during the 2001 recession and its aftermath. The payroll data exclude the self-employed.

The four remaining CPS components shown in Table 5.1 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, comprise tiny fractions of total employment and have no cyclical component. People who have jobs but are not currently being paid—individuals counted in the CPS household measurement but not in 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 between the two employment estimates.

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

According to Table 5.1, 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 measurement, 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 an employer's payroll at any time during the pay period that includes the twelfth of the month. My impression is that pay periods are generally two weeks, half a month, or a month.

Explaining the relationship between the length of the pay period and the overstatement of monthly snapshot of 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 Job Openings and Labor Turnover Survey (JOLTS) measures the new hire rate directly and shows little variation in the only business cycle that has occurred since it was launched in 2000. The separation rate is an excellent proxy for the new hire rate—the two measures differ only by the rate of change of employment, which at all times is insignificant 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 one-tenth of one percent. Even if the pay period is monthly, or four weeks, cyclical variations in the overstatement of employment caused by longer pay periods are tiny in relation to the observed discrepancy in the cyclical behaviors of CPS and payroll employment measurements.

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 individual fractions of the population engaging in the three possible activities—out of the labor force, unemployed, and working—must sum to one. The payroll survey provides no measure of the first two activities. One would have to adjust the fractions from the CPS household survey for calculating the fractions of the population that are out of the labor force and are unemployed to satisfy adding up to one. There is no basis for making the fraction of those out of the labor force and the fraction of those unemployed more countercyclical than is reported in the CPS, but these adjustments would be needed to incorporate the payroll employment data.

VI. Data on Hours

The CPS asks the household respondent (often not the actual worker whose activity is reported), a question like, "So, for last week, how many hours did the individual actually work at her or his job?" (the computer tailors the question to the individual worker). The survey procedure 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, such as when a holiday falls in the reference week. The choice of the reference week as the one including the twelfth of the month dodges Thanksgiving, Christmas, and several other major holidays, but cannot exclude every holiday celebrated in the United States. 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; see the Major Sector Productivity and Costs (MSPC) index

at http://bls.gov/lpc/home.htm. The payroll survey determines hours paid per job from employers. The MSPC index restates the results on the basis of hours worked rather than hours paid, using another survey that collects both. The MSPC index also uses CPS hours for workers not covered by the payroll survey. 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 5.5 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 5.4, 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.

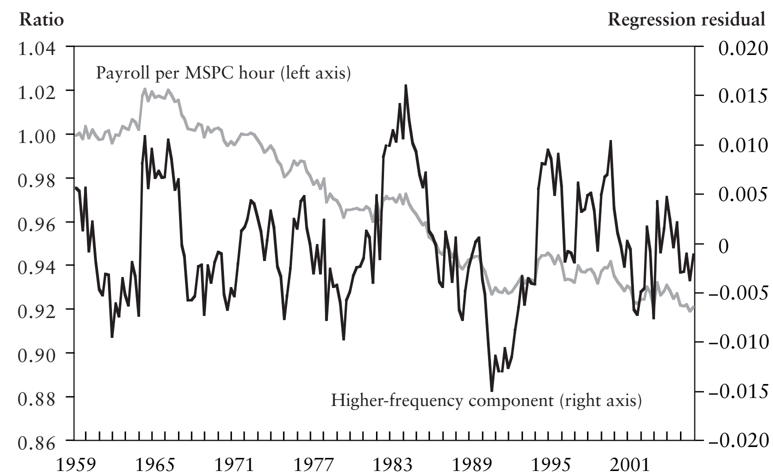

Figure 5.5
Ratio of Major Sector Productivity and Costs Index Hours to Current Population Survey Hours
Source: U.S. Bureau of Labor Statistics, U.S. Census Bureau, and author's calculations.
Note: The higher-frequency component comprises the residuals from a regression of the ratio on a fourth-order polynomial in time.

VII. 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), of this consensus view has developed a model incorporating job search, the third use of an individual's time, that responds to the same factors affecting labor supply. 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 indexes are highly correlated, so it is a reasonable approximation to treat the labor market as having a single driving force, the approach taken in this paper.

The important point that derives from this line of thought is that a single force drives all three key measures—participation, unemployment, and hours of work. This single force is the current position of the labor demand function in relation to its typical level trend. When labor demand is unusually strong, labor force 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 paper's 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 required to explain the labor supply response.

To derive a measure of the single driving force, I use three monthly measures that track the business cycle. Two measures 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 model is real disposable personal income per capita (see National Income and Product Accounts, Table 2.6).

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

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

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

$$(3) \quad h_t = \gamma_h z_t + \tau_h(t) + \varepsilon_{h,t}.$$

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. These factors 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—the εs 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 six observed moments: three variances of the ε disturbances, V_y, V_u, and V_h, and three covariances, $C_{u,y}$, $C_{u,h}$, and $C_{h,y}$. It has six 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:

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

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

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

(7) $\sigma_y^2 = V_y - \gamma_y^2$,

(8) $\sigma_u^2 = V_u - \gamma_u^2$,

(9) $\sigma_h^2 = V_h - \gamma_h^2$.

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

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

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 5.2 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.

The first line in the lower panel of Table 5.2 shows the loading coefficients, γ, for the three variables. Unemployment has a loading coefficient on the cyclical driving force of just under 1. 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 5.6 shows the index $\hat{z}_t$.

Table 5.2
Inference of Cyclical Driving Force from Data on Unemployment, Real Income, and Hours

	Unemployment	Real income	Hours
Moments			−0.54
Unemployment	1.14	−1.33	0.78
Real income		4.20	1.64
Hours			
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

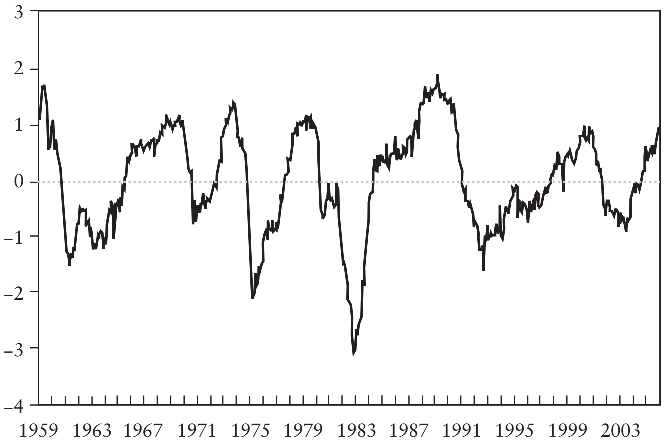

Figure 5.6
Index of the Single Driving Force
Source: U.S. Bureau of Labor Statistics, U.S. Census Bureau, and author's calculations.

VIII. 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 5.3 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 5.2 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 components, 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 5.3 shows the small but statistically unambiguous cycle in participation. Recall that the units of the cyclical driving forces are standard deviations of cyclical movements in the labor market. A one standard deviation tightening of the labor market raises participation by 0.2 percent. Because the aggregate level of participation is around 60 percent, this increase in the driving force increases labor force participation by 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 5.2 and 5.3 deal with labor measures per person. Table 5.4 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 5.4 needs to put population on the same footing as the other measures, as the higher-frequency component obtained as residuals from

Table 5.3
Loading Coefficients for the Three Dimensions of Work on the Cyclical Driving Force

	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 5.4
Cyclical Loading Coefficients for Number of Employed Workers

	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)

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 measurement framework used in the paper, it needs to measure the loading of population on the cyclical driving force. The population loading component 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 business cycle, but population swings could be a contributor to the cycle. This hypothesis seems to be part of the explanation—unusually low population growth during World War II led to a tighter labor market in the late 1960s, before the baby boomers entered, and unusually high population growth led to a slacker market in the period containing the weakest labor market, 1973 through 1983, as the boomers started work. 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 1. The loading is only slightly higher than the loading for employment per participant in Table 5.3, 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.

IX. Cyclical Responses by Demographic Groups

Table 5.5 breaks down the responses shown in Table 5.3 by age and sex, to the extent that the data are available from the BLS. Long historical

Table 5.5
Loading Coefficients for Participation, Unemployment, and Hours by Age and Sex

	Sex	Age	Loading	Standard error
Participation	Male	16 to 24	0.339	(0.071)
		25 to 54	0.146	(0.016)
		55+	0.349	(0.065)
	Female	16 to 24	0.679	(0.095)
		25 to 54	0.147	(0.047)
Unemployment	Male	16 to 24	–2.283	(0.029)
		25 to 54	–1.081	(0.018)
	Female	16 to 24	–1.537	(0.028)
		25 to 54	–0.867	(0.013)
Hours	Male	16 to 19	2.273	(0.105)
		20 to 24	1.179	(0.057)
		25+	0.699	(0.057)
	Female	16 to 19	2.125	(0.107)
		20 to 24	1.135	(0.066)
		25+	0.573	(0.070)

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 5.5 begin in June 1976.

Table 5.5 confirms that the participation elasticity is higher for younger workers (those under 25 years of age) and for older workers (those over 54 years of age), and among younger workers is higher for women than for men. The more elastic groups contain a larger fraction of people who are close to the margin between choosing to participate in the labor force and choosing to specialize in non-work activities, primarily activities at home and attending school. Unemployment among men and among younger workers is more sensitive to the driving force.

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 younger workers.

X. A More Complete View of Driving Forces

Hall (2008) develops a model of labor supply and unemployment derived from the underlying principles of labor supply and a generalization of the Mortensen-Pissarides search-and-matching view of unemployment. In the Hall model, hours of work and unemployment are linked to the rest of the economy through two variables. One is the marginal product of labor, which conveys the demand for labor. The other is the marginal utility of consumption, which conveys the long-run well-being of workers. Because the marginal product of labor is the primary determinant of well-being and because the marginal product of labor tends to evolve as a random walk, the marginal product of labor and the marginal utility of consumption are quite highly correlated. In this paper, I make the approximation that the correlation is so high that these can be treated as a single variable, which I call the driving force, z.

The analysis of some issues of the response of labor-market variables requires the two-variable framework. In particular, all macroeconomic models agree that an expansion in government purchases of goods and services tightens the labor market, raises hours of work, and decreases unemployment. The increased government spending does not change the demand for labor as measured by the marginal product of labor. Instead, it depresses long-term well-being because any increase in government purchases must be paid for, sooner or later, by reduced consumption. Hours of work increase because of the negative wealth effect in labor supply—higher government purchases decrease wealth and raises hours. Hall (2007) shows that the same effect operates on unemployment—lower wealth results in lower unemployment.

In the framework with two variables, the consensus I noted in the introduction might be phrased more precisely as "a consensus in macroeconomics holds that the observed higher-frequency movements in employment and hours of work are movements along labor-supply and employment functions caused by changes in fundamentals acting through the marginal product of labor and the marginal utility of consumption." The consensus viewpoint rules out shifts of the labor-supply and employment functions as important sources of fluctuations.

XI. Interpretation

I have not tested the consensus viewpoint that shifts in labor demand account for most of the cyclical variation in labor input. But it holds up well provisionally in the analysis of this paper. First, all three components of the labor demand function—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, these 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 the ranks of the unemployed. 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 workers found to be out of the labor force, because the unemployed are much more likely to be employed a month later. The unemployed are truly different from other people who are not working because they generally wind up working within a few months.

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.

■ *Prepared for the 52nd Economic Conference of the Federal Reserve Bank of Boston, 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 are available from the author—Google "Robert Hall".*

References

Alexopoulos, Michelle. 2004. "Unemployment and the Business Cycle." *Journal of Monetary Economics* 51(2): 277–298.

Bowler, Mary, and Teresa L. Morisi. 2006. "Understanding the Employment Measures from the CPS and CES Surveys." *Monthly Labor Review* 129(2): 23–38.

Cho, Jank-Ok, and Thomas F. Cooley. 1994. "Employment and Hours Over the Business Cycle." *Journal of Economic Dynamics and Control* 18(2): 411–432.

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

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

Hall, Robert E. 2005. "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years." In *NBER Macroeconomics Annual,* Volume 20, ed. Mark Gertler and Kenneth Rogoff, 101–137. Cambridge, MA: The MIT Press.

Hall, Robert E. 2008. "Sources and Mechanisms of Cyclical Fluctuations in the Labor Market." Available at http://www.stanford.edu/~rehall/SM022208.pdf.

Hall, Robert E., and Paul R. Milgrom. 2008. "The Limited Influence of Unemployment on the Wage Bargain." *American Economic Review*, forthcoming.

Mortensen, Dale T., and Christopher Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61(208): 397–415.

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

Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21(1): 3–16.

Rogerson, Richard, and Johanna Wallenius. 2007. "Micro and Macro Elasticities in a Life Cycle Model With Taxes." NBER Working Paper 13017.

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

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

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

Comments on "Cyclical Movements along the Labor Supply Function" by Robert E. Hall

Katharine G. Abraham

I very much enjoyed reading Hall's paper. From an empirical point of view, the paper does two things. First, it develops an innovative measure of high frequency shifts in labor demand. Second, given this monthly index of labor demand, it decomposes the total response of labor input to shifts in labor demand into the pieces represented by changes in labor force participation, changes in the share of people in the labor force who are employed, and changes in weekly hours. I will discuss these topics in turn.

As conceptualized in the paper, the index of labor demand is a latent variable that drives observable outcomes in the labor market. The behavior of this driving force is inferred from de-trended data series on the log of real personal disposable income per capita, the unemployment rate (measured as the negative of the log of the employment rate), and the log of weekly hours of work. Three regression equations are specified, each of which relates one of the observed variables just mentioned to the unobserved cyclical driving force (z_t), a fourth-order polynomial in time and, in the hours equation, a set of seasonal dummies. The values of the z_t variable are inferred using a set of moment conditions. These implied values are somewhat surprising, in the sense that they do not conform especially well to what most people think they know about the relative severity of recessions over the last few decades. In particular, the behavior of this index of labor demand is virtually the same in the recession of the early 1990s as in the recession of the mid-1970s, whereas conventional measures show the mid-1970s recession to have been much more severe than the recession of the early 1990s. This leads me to wonder what it is about the way in which the labor demand index is constructed that produces this result. One observation I would make is that the index of labor demand in essence reflects the behavior of the residuals of the three observable variables—personal income per capita, unemployment, and weekly hours—from a quadratic trend. The quadratic trend may pick up not only the longer-term influences on the observable variables it is intended to capture, but also movements that more properly could be considered cyclical.

A second observation relates to the use of personal disposable income in the system of equations from which the labor demand index is derived. Given the purpose for which the system is estimated, it would seem that the dependent variables in the individual equations should be outcomes that depend rather directly on the demand for labor. Much of personal disposable income can be categorized as labor income, but personal disposable income also includes the part of proprietor's income that represents returns to capital as well as a substantial amount of government transfer payments. I wonder whether it might not have been better to construct the system using a measure of employee compensation as a dependent variable in place of the measure of personal disposable income.

The most serious question I have about the system of equations Hall used to estimate the labor demand index, however, is the presence of separate equations for employment and hours. Thinking about how adjustments to changes in the demand for labor occur, employment and hours can be viewed as substitutes for one another: if the demand for labor falls, firms may reduce employment, cut workers's hours, or a combination of both. The point here is that employment and hours need not respond in the same way to changes in the demand for labor. With the estimating equations as specified, however, it is the co-movements in these variables' responses that will be reflected in the index of labor demand. This then leads me to wonder whether it might not have been better to use a measure of total hours, rather than separate measures of employment and weekly hours, in the estimating equations.

Let me turn now to the second part of the paper, which looks at the responsiveness of participation, employment, and weekly hours of work to changes in the index of labor demand. The equations used to estimate these responses are specified in a rather parsimonious fashion. For starters, the model makes no allowance for lags in the response of any of the three labor input variables to changes in labor demand. It also makes no allowance for changes over time in the way in which the different labor input variables respond to changes in labor demand. As an aside, I might also mention that one needs to be careful in characterizing the estimates obtained from the model. The paper is very careful in its choice of words—talking about labor input rather than labor supply—but this is a conference on labor supply and, to the extent that decisions about hiring and hours are made by the employer rather than by the worker, it should be emphasized that the outcomes that are observed cannot be interpreted as labor supply responses, but rather can be traced to labor demand.

The main findings in this part of the paper decompose the response of labor input to changes in labor demand into shares due to changes in labor force participation, changes in the (un)employment rate, and changes in hours of work. Changes in labor force participation are relatively unimportant. A suggestion to be made here is that, in carrying out this sort of decomposition exercise, there could be value in thinking about where the boundary between being in and being out of the labor force is drawn. In the official data series that underlie the estimates reported in the paper, people are considered to be in the labor force if they are either counted as working or as being unemployed, with the unemployed including people who want to work, are available for work, and have actively looked for work in the last four weeks. But this boundary could be drawn differently, with either more or fewer people counted as belonging to the labor force and, at least in principle, this could affect the results obtained. The Bureau of Labor Statistics produces several alternate measures that draw the boundary somewhat differently – for example, changing the position of the boundary to include people who searched for work in the past year but not in the past four weeks as belonging to the labor force. In practice, the cyclical behavior of these alternate measures appears to be very similar to the cyclical behavior of the official unemployment rate, which makes me think this is probably not a major issue for the analysis. Still, this seems like something that might merit a bit more careful exploration.

Most of the response of labor input to changes in the index of labor demand is attributable to changes in the (un)employment rate (the share of the labor force that is (un)employed), and changes in weekly hours. At least to my eyes, in Hall's paper, changes in weekly hours account for a large share of this response. My reading of the message to be taken from the literature on the dynamics of labor demand is that, in the United States, adjustment to changing labor demand conditions tends to occur through changes in employment rather than through changes in weekly hours. In the results reported here, however, fluctuations in weekly hours account for about a third of the overall response. The framework here is rather different than the framework used in the labor economics literature with which I am more familiar, but I nonetheless find this paper's results a bit surprising. In the remainder of my comments, I would like to consider why hours adjustment appears to be so important here.

In part, I suspect, these findings may reflect the way that the labor demand index is constructed. As I have already noted, by construction, the labor demand index reflects only the common movements in income, employment, and hours. I worry that the approach adopted in this paper builds in an association between weekly hours worked and the labor demand index that would not be there if, for example, a measure of total hours had been used in place of the separate employment and hours measures.

A second issue, discussed at some length in the paper, is that the measure of employment derived from the Current Population Survey (CPS) on which the paper rests is much less cyclical than the measure derived from the monthly employer payroll survey. One might suspect that the different behavior of the two measures could be explained by their different conceptual underpinnings—for example, the inclusion of the self-employed in the CPS measure but not in the payroll survey measure—but accounting for these conceptual differences in fact does little to make the two series more comparable. If the payroll survey does a better job of measuring employment than the household survey, then, properly measured, the responsiveness of hours is relatively less important than implied by the estimates reported here.

Which measure of employment—the CPS measure or the payroll survey measure —is more believable? To answer this question, we need first to understand why the cyclical behavior of the two employment measures has been so different. A variety of explanations have been proposed, but

at this point there is no obvious "smoking gun." Potential explanations that seem to me to merit further investigation include possible problems with the classification of people as self-employed versus wage-and-salary workers in the CPS, cyclical movements in the amount of "off-the-books" employment recorded in the CPS that do not show up in the payroll survey, and problems with the CPS population controls related to difficulties in accounting for immigration—but I cannot tell you that any of these potential explanations will end up offering an answer. In the meantime, the uncertainty about which employment series should be believed makes it difficult to know how much confidence to place in this paper's estimates.

In summary, Hall's paper offers a creative and interesting approach to the measurement of fluctuations in labor demand and their effects on the labor market. At this point, I have a number of questions about the results obtained. My hope would be that Hall's further work on this issue will provide a better basis for assessing the robustness of the conclusions reported here.

Comments on "Cyclical Movements along the Labor Supply Function" by Robert E. Hall

Susanto Basu

Bob Hall is famous for working in different areas of macroeconomics, making important contributions in each one, and then forsaking his current area of research to work on other issues. But there is one topic to which he has returned again and again since his earliest days as an academic, understanding the behavior of unemployment and hours worked over the business cycle. The last few years have seen Hall produce a burst of papers that collectively deepen our understanding of this central issue in macroeconomics and labor economics. The present paper is another step in this important research program.

Let us review the facts that Hall takes as his starting point. Over the business cycle, we see large changes in employment and total hours worked, with relatively small changes in real wages. If one adopts the perspective that Hall takes in this paper, namely that business cycles are due to shifts in labor demand along a stable labor supply curve, these facts suggest that the labor supply curve must be quite flat. However, microeconomic estimates of labor supply elasticities suggest that the elasticity of hours worked with respect to wages for continuously employed workers is small—that is, the labor supply curve for employed workers is steep, not flat.

But of course, changes in hours worked by the employed account for only a small fraction (perhaps 15 percent) of the cyclical variation in total hours worked. The vast majority of the variation in labor supply over the business cycle comes from changes in the number of people employed. Together, these facts suggest that the underpinnings of "macro labor supply"—the supply of total hours worked—are quite different from those assumptions typically used to explain micro labor supply, which involve utility-maximizing choices made by individuals in a frictionless, neoclassical framework.

Hall thus moves to a framework where the determinants of micro and macro labor supply are different. In the paper that derives formally the framework explained heuristically in this paper, Hall (2008) shows how a variant of the popular Mortensen-Pissarides (MP) search-matching model yields a labor supply function of the form

$$(1) \quad L = H(\lambda, w)N(\lambda, w)$$

where L is total hours worked, H is hours per worker, N is the number of workers, w is the real wage, and λ is the marginal utility of wealth (assumed equal for everyone due to perfect consumption insurance). The supply of hours for employed workers takes the standard Frisch form: it is a function of the real wage and (expected) lifetime wealth. Hall's striking achievement is to show that the number of workers available for employment is also a function of the same two variables. In the same paper, which I recommend highly to all who are interested in this topic, Hall shows that one can use a mix of calibration and estimation to infer the shapes of the H and N functions from household studies, plus aggregate data on consumption, employment, and hours worked. The MP model fits the facts well once one modifies the model as in Hall (2005), by assuming that the real wage is constant as long as it is within the bargaining range that is efficient for workers and firms.

While Hall's interpretation of the facts is insightful and consistent with labor market paradigms that are currently in vogue, it is useful to ask whether there are alternative models of the labor market that can also explain the basic stylized facts. And if there are indeed alternative theoretical explanations, how might one use data to discriminate empirically among these competing models?

Models based on fixed costs of going to work that are incurred by workers, such as Hansen (1985) and Rogerson (1988), try to explain the difference between micro and aggregate labor supply in a neoclassical framework. But Mulligan (2001) shows that the strong implications of these models come from the assumption in both papers that the fixed cost

of going to work is identical across all workers. Once Mulligan allows for a distribution of this parameter, he finds that there are few implications of fixed costs per se, and none that are significant for labor supply over the business cycle.

The early Keynesian interpretation of the cyclical facts that motivate Hall also differentiates between micro and macro labor supply. But where Hall modifies the neoclassical model by introducing the information friction central to the MP model, the Keynesian story dismisses the neoclassical paradigm altogether, at least for the purpose of understanding short-run aggregate fluctuations in the labor market. In the Keynesian framework, workers agree to work as much as employers demand at a pre-set nominal wage. As in the model that Hall develops, the effective labor supply curve in the (Old) Keynesian model is indeed flat—the pre-set wage is independent of employment—and fluctuations in employment and hours are determined by labor demand. If prices are approximately as sticky as wages, then real wages change relatively little over the business cycle, matching what the data show.

One might object that the Keynesian framework does not determine the split of total hours worked between overall employment and hours per worker. But a small change in the set-up that introduces fixed per-worker costs incurred by the firm—for example, benefits like health insurance, the value of which often is independent of the number of hours worked—would remedy this problem.

The more fundamental problem that Hall sees in the Keynesian framework is that its predictions are not derived from optimal decision-making by workers and firms. In a paper that has been central to Hall's thinking on these issues, Barro (1977) pointed out that the contracts assumed in sticky-wage models are not Pareto-efficient—that is, both firms and workers would gain by renegotiating to reduce wages in a downturn instead of cutting employment. Hall has adopted Barro's view that even if observed nominal (or real) wages are sticky, the decisions on employment and hours are made in accordance with a long-term implicit contract between workers and firms in which wages are fully flexible, with the observed spot wages being "installment payments" of the agreed-upon total lifetime wages due workers.

A different set of models, based on efficiency-wage considerations, also has implications for the difference between micro and macro labor supply. As an example, consider the efficiency-wage model of Shapiro and Stiglitz (1984). Unlike the Keynesian model, the Shapiro-Stiglitz model assumes optimizing behavior by firms and workers. In fact, the famous labor-market diagram summarizing the predictions of the model bears a striking resemblance to Figure 5.3 of Hall's paper. Both models predict equilibrium unemployment. Both display steep supply functions for total hours in the absence of frictions, but a relatively flat effective labor supply curve in the environment with frictions. In Hall's model the friction is an informational one—people do not automatically know where to find a good match for their skills, and must invest in job search. In the Shapiro-Stiglitz model, the friction is imperfect monitoring—workers and firms contract over the number of hours spent on the job, but the firm cannot contract over how hard the worker works each hour, since effort is not observed perfectly. Is one friction clearly more important than the other? The answer is not obvious to me. But macro labor research in recent years has focused almost exclusively on the MP model and its variants, and ignored the Shapiro-Stiglitz model.

Which of these four classes of models best explains the data? One would hope that such questions would be settled by confronting these competing models with data in a systematic fashion. Unfortunately, few such efforts have been made, and most are tests of a single model against an unspecified alternative. The failure to test the models against one another is due partly to the fact that these models explain different features of the data. A major strength of the Mortensen-Pissarides search-matching framework, for example, is its ability to explain data on worker flows and job vacancies, but the other models have little or nothing to say about such issues. On the other hand, all the models of the labor market discussed above are able to match the key stylized facts of the data—in most cases, they were created to do so! Unfortunately, beyond matching these facts, the models make surprisingly few empirical predictions, and the predictions they make are often ambiguous and subject to varying interpretations. A good example is the literature on efficiency-wage models and inter-industry wage differentials. Efficiency-wage models do

predict that workers in some industries might be paid higher wages than identical workers in other industries, but it is impossible to rule out the possibility that workers in different industries who appear identical to the econometrician actually differ in their labor market characteristics.

Thus, in practice the choice between the matching-based framework that Hall advocates and the other models is guided as much by aesthetic considerations (what makes a good model?) and intuitive ones (what are the basic institutions and frictions in the labor market?) as by any formal empirical testing.

These are my comments on Hall's current research agenda in labor economics, of which the current paper is a part. Now let me turn to the novel contribution of the conference paper at hand. Relative to Hall (2008), the contribution in this paper is to derive an index of labor demand, and study how each component of labor input responds to labor demand. Hall is very clear in noting that his exercise is possible only under the assumption that the labor supply curve is stable over the business cycle. In terms of equation, the assumption is that changes in λ happen at low frequencies, but are not relevant for business cycles.

Unfortunately, this assumption does not hold for an important category of shocks, namely shocks to government expenditure. The empirical literature on these shocks, summarized in Perotti (2007), shows that these increase output and hours worked at high frequencies. Even in the post-World War II period, there have been large, exogenous changes in U.S. government expenditure, usually associated with national security crises, notably the Korean and Vietnam Wars. But in the neoclassical labor supply framework in which Hall operates, spending shocks increase output and hours worked by raising λ, that is, by making consumers feel poorer which results in them working (and producing) more.

Another way of making this point is to note that in a neoclassical framework with distortionary taxes, the labor demand curve can be expressed as:

$$(2) \quad w = (1 - \tau)AF_L(K, AL),$$

where τ is the (labor income) tax rate, F is a production function, K is capital, A is the state of technology, and the L subscript denotes the marginal product of labor. Since technology does not change in response to a fiscal shock and capital, a state variable, is essentially unchanged as well, the labor demand curve in (w, L) space shifts only if the tax rate changes. But one would think that current tax rates would rise in response to an increase in government spending, since the government requires more revenue, and indeed the empirical work confirms this conjecture. In this situation, then the labor demand curve must shift inward. So if output and hours must rise to match the empirical evidence, then an outward shift in labor supply must be responsible for more than 100 percent of the increase in hours worked after an increase in government purchases.

In some non-neoclassical equilibrium models with imperfect competition (for example, Rotemberg and Woodford 1992), labor demand can shift out due to a fall in the mark-up of price over marginal cost. Even so, the mark-up changes in response to a change in output, and output would not change unless the labor supply curve shifts outward.

To confirm that government purchase shocks do indeed have their expected effects on unemployment and hours per worker, I took the preferred series of government purchase shocks from Perotti (2007). (Roberto Perotti kindly supplied the data.) Figure 5.7 shows the effects of government purchase shocks on employment and hours per worker. As one would surmise, an unexpected increase in government purchases increases weekly hours and lowers unemployment. Some of the fluctuations are quite large, on the order of a half-percentage point change in the unemployment rate. Thus, under the maintained hypothesis that Hall's framework provides the correct interpretation, the evidence shows that it is not safe to assume that changes in λ are negligible at high frequencies.

In sum, Bob Hall is revolutionizing our understanding of one of the most perplexing issues in macroeconomics, the behavior of labor supply over the course of the business cycle. I remain an interested observer. Since I am unconvinced that a lack of information is the most important friction preventing the labor market from functioning smoothly over the business cycle, I am somewhat skeptical that this research program will attain all its promised objectives, but I remain hopeful. However, I am fairly sure that this particular paper will not contribute in any important way to the attainment of the desired outcome.

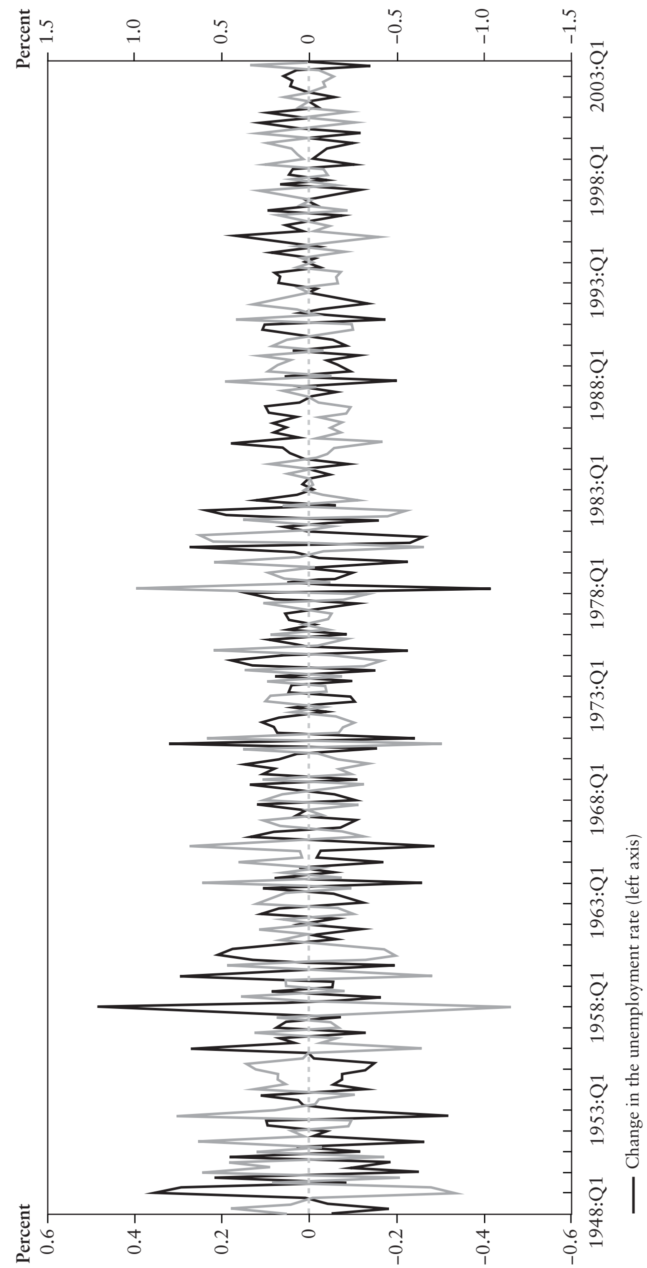

Figure 5.7
Changes in Unemployment and Weekly Hours Due to Shocks in Government Spending
Source: Blanchard and Perotti (2002).

References

Barro, Robert J. 1977. "Long-Term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics* 3(3): 305–316.

Blanchard, Olivier, and Robert Perotti. 2002. "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output." *The Quarterly Journal of Economics* 117(4): 1329–1368.

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

Hall, Robert E. 2008. "Sources and Mechanisms of Cyclical Fluctuations in the Labor Market." Available at http://www.stanford.edu/~rehall/SM022208.pdf.

Hansen, Gary D. 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16(3): 309–327.

Mulligan, Casey B. 2001. "Aggregate Implications of Indivisible Labor." Advances in Macroeconomics 1(1). Available at: http://www.bepress.com/bejm/advances/vol1/iss1/art4.

Perotti, Roberto. 2007. "In Search of the Transmission Mechanism of Fiscal Policy." NBER Working Paper 13143.

Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics* 21(1): 3–16.

Rotemberg, Julio J., and Michael Woodford. 1992. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy* 100(6): 1153–1207.

Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74(3): 433–444.

6

Labor Supply and Labor Demand in the Long Run