The Labor Market is the Key to Understanding the Business Cycle *

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

September 23, 2004

1 Introduction

Modern economies experience booms and recessions with amplitudes that are puzzling within the framework of standard economic principles. Economists have trouble finding large enough external impulses to explain the observed fluctuations and equal trouble finding amplification mechanisms that might link the fluctuations to relatively small impulses.

J.M. Keynes opened a nonstandard line of inquiry that remains at the center of research on aggregate fluctuations. He suggested that labor and product markets do not work as smoothly as standard economics suggests. In all markets, prices should absorb some of the effect of a shock—and in markets such as the labor market, with inelastic supply, the price should absorb most of a shock. Instead, prices and wages seem to respond relatively little to the shocks that cause recessions, while employment, unemployment, and output respond sensitively.

Research immediately following Keynes showed that two non-standard assumptions would create a model that had some chance to replicate the cyclical

volatility of the economy. The first is wage-price rigidity. The second—whose importance is rarely adequately emphasized—is that sellers stand ready to serve the demand of buyers at the rigid price or wage. Barro (1977) stressed that the rigid-price or rigid-wage model would not explain output or employment volatility absent the second assumption. The essence of the traditional rigid-price model is that the resulting allocation is inefficient—the buyer equates its marginal benefit to the price, but the seller has a marginal cost below price under recession conditions, when output and employment are inefficiently low. As Barro noted, it is hard to understand why two parties, buyer and seller or employer and worker, cannot make a new efficient bargain on their contract curve. The issue is *not* to readjust the wage, but rather to move to the efficient level of output or employment.

Discussions of the rigidity of prices and wages have neglected Barro's observation and continued to focus primarily on rationalizing price-wage rigidity, not output and employment volatility. In the meantime, following Kydland and Prescott (1982), an important alternative explanation of volatility occupied center stage in macroeconomics—an equilibrium model with highly elastic labor supply. The evidence in favor of such high elasticity has not been persuasive, so the interest among macroeconomists has returned to price and wage rigidity.

My own reading of the evidence on price rigidity is skeptical. A few prices do fit the model, such as those for regulated utilities, where a price is fixed for a considerable period and customers choose the level of output. But markets for the great majority of final products resemble auctions, not regulated utilities. Prices for major purchases are dickered, an informal type of auction. And, in markets where customers choose among alternative posted prices, the resulting outcome is just the same as a classic second-price auction—see Klemperer (2003). Not only does an auction generate a flexible price, but it also results in an efficient allocation.

I will not develop the case against a model based on price rigidity in this paper. Rather, I want to make the case in favor of a new way that wage rigidity can account for employment volatility. The ideas in this new view might apply to product markets as well, so I have no interest in dismissing price rigidity.

I call the new view the *equilibrium sticky-wage* model. The model describes an equilibrium in the sense that workers and employers never pass up opportunities for bilateral joint improvement. Workers do not lose their jobs because their sticky wage exceeds their marginal products. There is no burst of job loss in a recession.

2

Instead, fluctuations in employment and unemployment occur because jobs are harder to find in recessions.

I start the paper with a review of the evidence on two empirical propositions that are central to the new view. First is that turnover dynamics in the labor market have essentially nothing to do with the time-series properties of unemployment or employment. One might think that a recession is kicked off by a burst of job loss and that the recession lasts as long as it takes for the victims to find new jobs. But the truth is quite different. Around 35 percent of the unemployed find jobs each month and another 25 percent leave the labor force, so the exit rate from unemployment is about 60 percent per month. A recession would last only two or three months if the reabsorption of job losers was the central mechanism for restoring normal unemployment. In reality, the adjustment rate of unemployment back to normal is two percent per month, not 60 percent. It turns out that turnover dynamics operate so rapidly that it is safe to ignore the transitory stochastic disequilibrium of the labor market when unemployment changes. This point is established in the first section of the paper.

The second point is that, in the modern U.S. economy, recessions are not times of unusual job loss. New data on separations show them to be remarkably constant from peak to trough. Bursts of job loss had some role in earlier recessions, but are still mostly a side issue for the reason just mentioned—a burst is quickly reabsorbed because of high job-finding rates.

To understand the basic facts of unemployment fluctuations, one must understand why the period of several years following a recession, when unemployment remains high, is also a period of subnormal job-finding rates. Unemployment is high because the typical searcher takes more time to find a new job. A theory of fluctuations must explain persistent changes in job-finding rates.

I examine the fundamentals of the determination of the job-finding rate. A searcher and an employer enjoy a joint surplus from the transition from unemployment to employment. The surplus arises because work yields a higher flow value than a worker receives when not working. The surplus provides the incentive for the two of them to try to find each other and make a job match. The costs of job creation include the shoe leather of the searcher, the help-wanted advertising and other costs of attracting a worker, and the costs of evaluating the worker for the job.

3

Absent frictions, responses to incentives would stabilize unemployment at an equilibrium point that would not change much over time. If unemployment rose a bit, the joint surplus from job creation would rise, recruiting effort would rise, and unemployment would return quickly to normal. The model would provide a deep explanation of the natural unemployment rate, but would not help understand fluctuations.

Frictions are a likely characteristic of the labor market, however. When the market is slack and unemployment is high, the joint surplus from job creation is high because workers' opportunity costs are low when jobs are hard to find. If most of the effort needed to form a match comes from the employer, then the greater value of job creation in a slack market needs to be communicated to the employer. The communication occurs before the pair agree on the terms of employment. If the employer believes that the terms will be favorable when the match is formed, then the employer will invest in extra match-forming effort in soft markets. Thus flexible wages would bring unemployment quickly back to normal. The flexible wages transfer the incentive from the worker, who bears the cost of a soft market, to the employer, who controls the means for tightening the market.

If wages are sticky, the incentive is not transferred—it remains with the worker. Under the assumption—critical to this line of thought—that workers cannot substitute their own job-creation efforts for those not made by employers—the labor market can remain soft for as long as the wage remains high. The model can explain fluctuations from outside shocks that would call for a lower wage in a flexible-wage model—such as a decline in productivity or adverse shift in the terms of trade—or from a spontaneous increase in the wage.

I am not offering an affirmative explanation for sticky wages. Notice that the wage that needs to be sticky in this setup is the anticipated present value of wages as of the time of hire, not the wage paid to incumbent workers. Rationalizations applicable to incumbents, such as resistance to opportunistic wage cuts, are not helpful here. We need to understand why employers cannot hire at lower wages even though the opportunity cost of the worker has declined. A standard view of bilateral bargaining—that the resulting wage responds to the threat points of both parties—does not hold if wages are sticky. In particular, the Nash bargain, where the wage is a convex combination of the two threat points, makes wages so flexible that the model cannot come close to explaining the observed volatility of

unemployment—see Shimer (2005).

An important point in favor of sticky wages, however, is that a fixed wage is an equilibrium of the fundamental bargaining problem between a worker and an employer when they first meet to agree on the terms of their relationship. Any wage in their bargaining set is an equilibrium in the basic sense that the pair cannot bargain to a jointly superior point. Consequently, a sticky wage—so long as it remains in the bargaining set—is one of many equilibrium selection rules that could operate in the labor market.

The essence of the equilibrium sticky-wage view, then, is that the labor market is tight and jobs are easy to find when the current wage happens to find itself toward the lower end of the bargaining set. Employers have strong incentives to recruit workers and create jobs for them in that case, because they receive a large share of the surplus resulting from the new employment. Unemployment is low. If the boundaries of the bargaining set move downward, employers receive less of the surplus and have a lower incentive to recruit. The labor market slows down because jobs are harder to find. Unemployment rises.

Although sticky wages provide a mechanism for explaining large fluctuations in recruiting effort, job-finding rates, and the unemployment rate, the general view of the paper extends to other mechanisms as well. I conclude the paper with a brief discussion of some alternatives.

This paper is deliberately sketchy as far as references to related literature, to descriptions of data, and formal models. For all of these, see my three related papers, Hall (2005b), Hall (2005a), and Hall (2004).

2 The irrelevance of turnover dynamics

The following discussion is adapted from Hall (2005a) which contains many further details. Let u be the unemployment rate, f be the exit rate (the fraction of unemployed workers in one month who are not unemployed in the next month), and s the entry rate (the number of newly unemployed as a fraction of employment). With the labor force normalized at 1, the unemployment rate follows the difference equation,

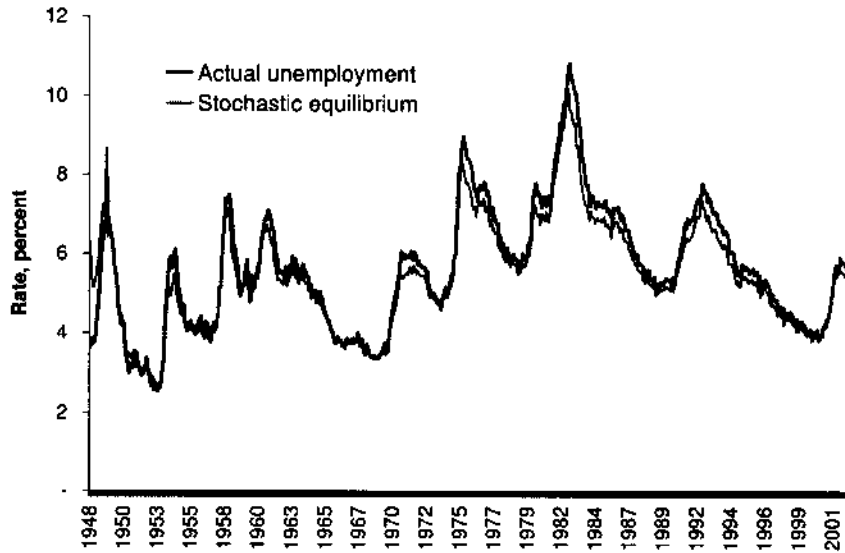$$u_t = (1 - f_{t-1})\, u_{t-1} + s_{t-1}\, (1 - u_{t-1}) \qquad (1)$$

Figure 1. Irrelevance of Turnover Dynamics

If the exit and entry rates are constant, then equation (1) describes a two-state Markov process with stationary unemployment,

$$u = \frac{s}{s+f}.$$

(2)

If turnover dynamics were an important part of the story of the movements of unemployment, then the stationary level of unemployment would lead the movements of actual unemployment. For example, during a period of higher flows into unemployment from job losses at the beginning of a recession, unemployment builds up to its new higher level, then recedes to its normal level after the inflow returns to normal. But the lead is tiny, because the exit rate is typically 60 percent per month. Figure 1 demonstrates the irrelevance of turnover dynamics. It compares the actual movements of unemployment to the movements of the stationary level, evaluated at the current estimates of the entry and exit rates. The actual unemployment rate tracks the stationary level almost exactly.

6

Hall (1995) and Cole and Rogerson (1999) noted earlier that unemployment movements have almost nothing to do with turnover dynamics. The implication of this finding is simple—to understand the movements of unemployment, we need only look at the movements of the entry rate s_t and the exit rate f_t. We do not need to worry about the fact that the unemployment rate is a distributed lag of past values of these variables, because the lag is so short. For the rest of this paper, I will refer interchangeably to the movements of the entry and exit rates and movements of the unemployment rate.

3 Separations and flows into unemployment

A second important simplification in thinking about the modern labor market is that the flow into unemployment is roughly constant. A recession is not a period of high flows into unemployment—it is a time of low rate of success of job-seekers per month. Historically, flows into unemployment did rise sharply at the outset of recessions, but the rise did not last long. The last two recessions saw no significant rise in flows into unemployment. A related finding is that the flow of workers out of jobs—the separation rate—did not rise in the recession of 2001.

The Current Population Survey is an economy-wide source of data on flows in the labor market. See Blanchard and Diamond (1990) for further discussion and cites to the earlier literature. Figure 2 shows the rate of entry to unemployment as derived from the data on the number of unemployed workers whose spells began within 5 weeks of the survey. NBER recessions are shown at the bottom. The flow has large low-frequency movements, rising to a peak in 1982 and then falling to its historical low in the last year reported, 2003. There is no sign of important increases in inflows to unemployment in the two most recent recessions, in 1990-91 and 2001. Earlier recessions, especially 1948-49 and 1981-82, did show bursts of entry to unemployment.

Figure 3 breaks down new unemployment by source starting in 1977. Job loss from both temporary layoff and other sources—permanent loss of jobs and the ending of temporary jobs—rose dramatically in the 1981-82 recession and rose a small amount in 1990-91 and 2001.

Beginning in December 2000, the Bureau of Labor Statistics has collected data on separations and hires for a large sample of employers. Fortuitously, the early
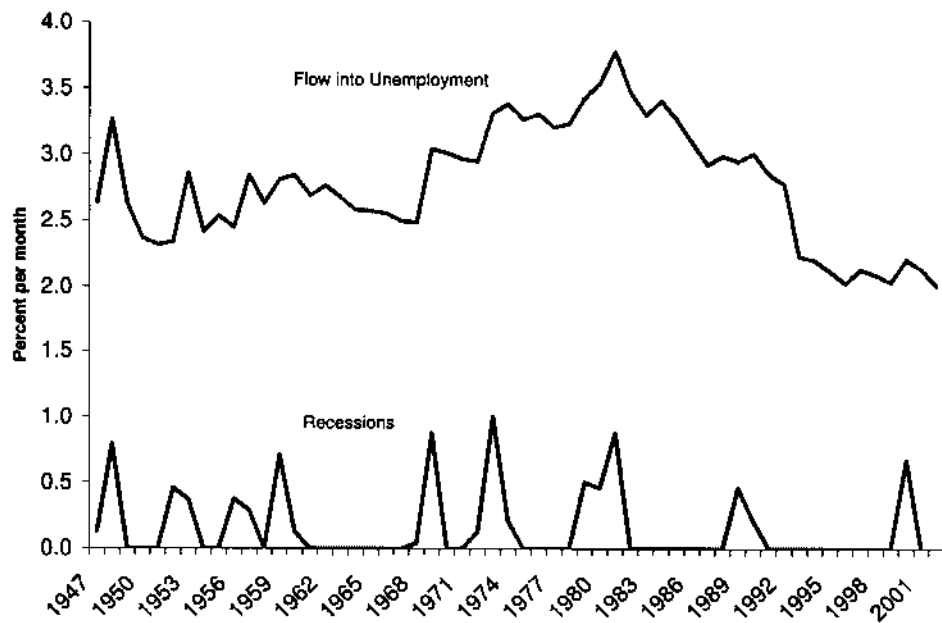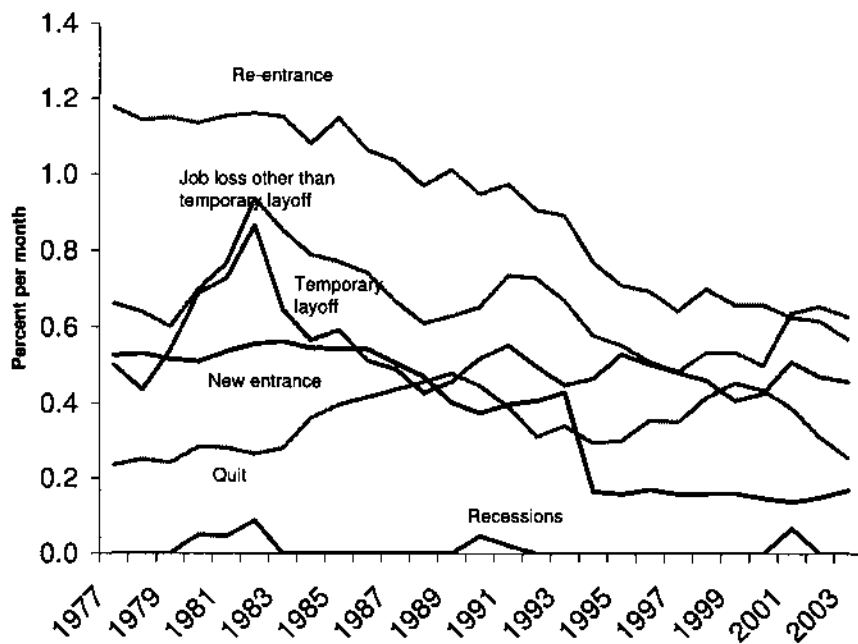
Figure 2. Flows into Unemployment

8

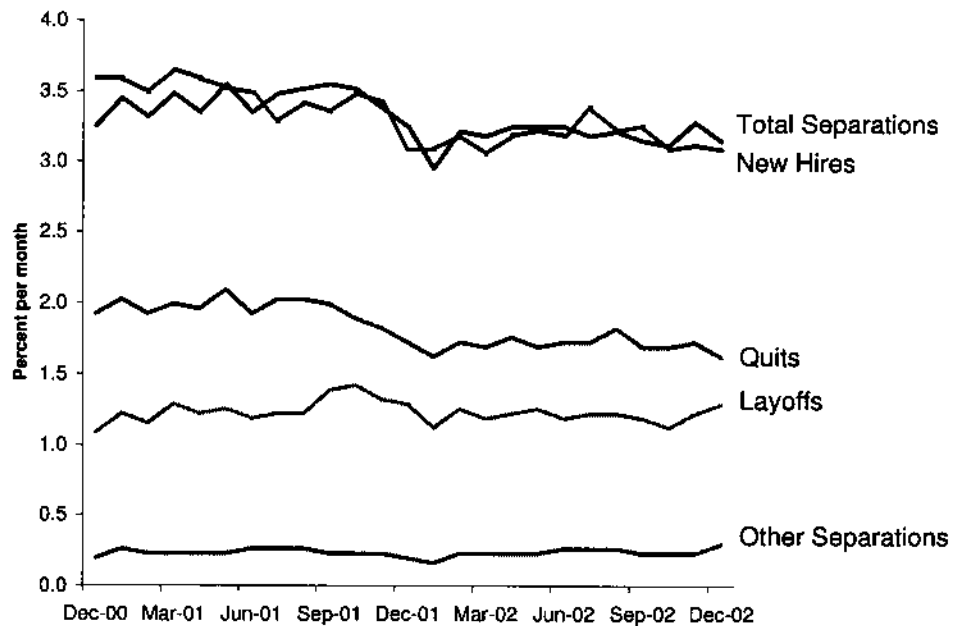Figure 3. Components of Flows into Unemployment

Figure 4. Separation Rate and Components

months of the new survey caught the labor market just before the peak of employ-
ment, so the period of the survey to date describes the differences between a strong
market and a weak market. The extended unemployment rate rose from 6.9 percent
in December 2000 to 9.1 percent in August 2003.

Figure 4 shows the turnover rates recorded in the survey. Most remarkable
is the behavior of the separation rate. Except for a bulge following September
11, 2001, layoffs remained almost exactly constant from the peak of the market
in December 2000 through the end of 2002, a period of continuing declines in
employment and rising unemployment. The recession did not begin with a burst of
job loss. Quits did decline later in the contraction, in accord with standard beliefs
about what happens in the labor market during recessions, so total separations fell
modestly.

Another source of information about flows in the labor market is the tabulation
of plant-level employment changes pioneered by Davis and Haltiwanger (1992).
These authors measure what they call *job destruction* as the sum of employment

10

declines across plants. They find that job destruction spikes during recessions. But job destruction does not measure separations—rather, it measures separations less new hires in plants where the difference is positive. Job destruction does not distinguish employment reductions that occur because of failure to replace normal attrition, from employment reductions that occur because of actual separations. In any period of declining employment, job destruction necessarily rises unless a surprising and unlikely change occurs in the shape of the distribution of employment changes across employers. In a model governed by the principle of efficient separations, and with heterogeneity across plants, job destruction would rise in response to a shock that caused a decline in employment, even though separations remained constant. Consequently, there is no contradiction between the finding of no spike in separations and entry to unemployment in recessions with Davis and Haltiwanger's finding of a spike in job destruction.

4 Unemployment exit rates and recruiting effort

Figure 5 shows the weekly exit rate from unemployment from the Current Population Survey. The exit rate plunges reliably in every recession, including especially the most recent in 2001.

Recruiting effort as measured by advertising is remarkably volatile. Figure 6 shows the Conference Board's index of help-wanted advertising since 1951. It is not uncommon for advertising to fall by 50 percent from peak to trough, as it did from 2000 to 2003.

A second measure of recruiting effort is vacant jobs. In many employers, the formal posting of a vacancy accompanies recruitment efforts of many kinds. The JOLTS survey collects data on vacancies across the economy. Table 1 shows data from JOLTS on vacancies by industry for the period of slackening of the labor market since late 2000. The figures confirm the high volatility of vacancies suggested by the data on help-wanted advertising. The data show that vacancies have declined in all industries. Although the forces that caused the downturn in the economy disproportionately affected a few industries far more than others—notably computers, software, and telecommunications equipment—the softening of the labor market was economy-wide.
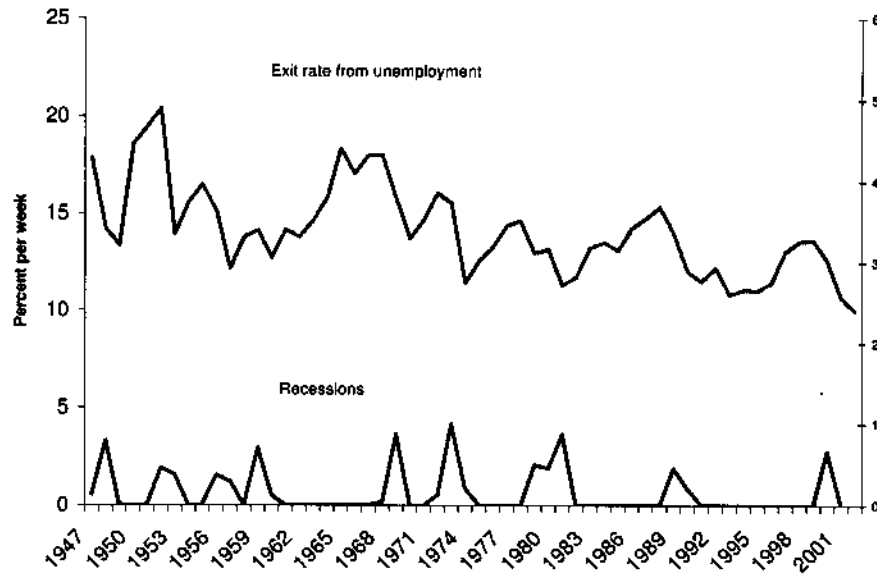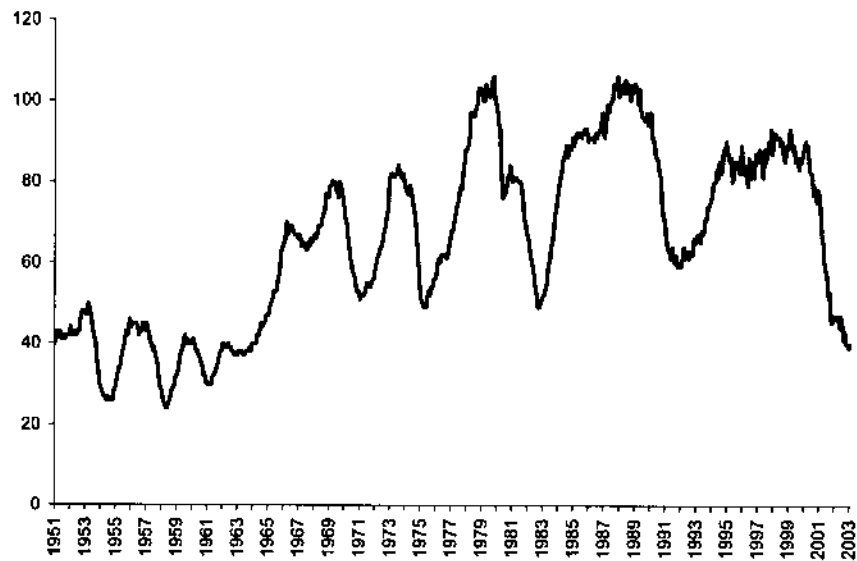
Figure 5. Weekly Exit rate from Unemployment

Figure 6. Index of Help-Wanted Advertising
Source: The Conference Board, http://www.globalindicators.org

12

Industry	Ratio of vacancy rates in 12/02 and 12/00
Mining	0.36
Construction	0.38
Durables	0.45
Nondurables	0.48
Transportation and utilities	0.80
Wholesale trade	0.52
Retail trade	0.60
Finance, insurance, and real estate	0.79
Services	0.68
Federal government	0.54
State and local government	0.70

Table 1. Change in Vacancy Rates by Industry in JOLTS, December 2000 to December 2002

5 Implications for the theory of fluctuations

The evidence about the modern economy points clearly in one direction about the labor market in recessions. In a recession, as unemployment rises, the market slows down. Jobs become harder to find and the unemployed take longer to find new jobs. Employers in all industries put less effort into recruiting. Employed workers, however, do not lose jobs more frequently in recessions than in other times. Unemployment rises because the exit rate from unemployment is lower, not because the entrance rate is higher.

The traditional view of wage rigidity emphasized job loss as much as lowered job-finding rates in recessions. When a negative shock hit, according to that view, workers lost their jobs in unusual numbers because their employment was no longer profitable to employers. These employers chose the employment level unilaterally to maximize profit given the wage. They did not take account of the economic loss inflicted on job losers in making their employment decisions. Given that the loss is higher when unemployment is high and jobs are hard to find, there is a presumption that the burst of layoffs in a recession is inefficient. Employers and workers failed to make a Pareto-improving deal to exploit the gap between the marginal product of labor, which remained unchanged, and the marginal value of workers' time, which fell in the recession.

The alternative view is that the influence of sticky wages on unemployment occurs only before the match is made and does not result in inefficient job loss post-match. In this view, the job-loss rate reflects idiosyncratic fundamentals. Aggregate shocks do not trigger inefficient job loss. But sticky wages do make the employer's benefit from job creation sensitive to driving forces. Employers make large adjustments in help-wanted advertising and a wide array of other recruiting and job-creation measures in response to changes in the payoff they perceive to adding new workers. In recessions, this payoff is low and the labor market operates sluggishly with high unemployment.

6 Theory

I will start an informal discussion of the theory with the case where unemployed workers and potential employers face no frictions in solving the joint problem of
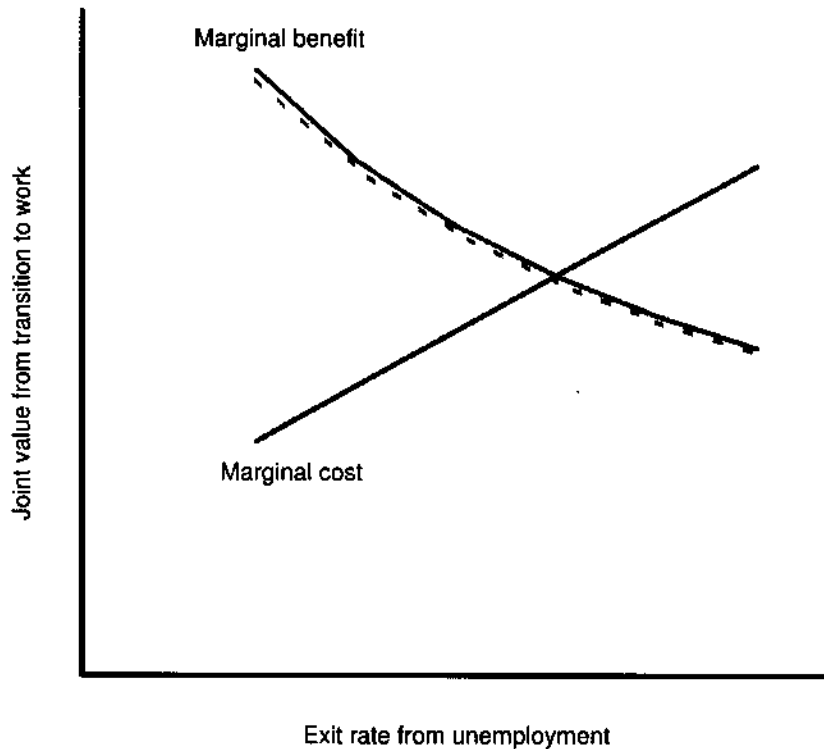
14

Figure 7. Equilibrium without frictions

forming matches and creating jobs. Figure 7 lays out the basics. On the vertical axis is the joint benefit of an unemployed worker starting a new job. On the horizontal axis is the job-finding rate.

The downward-sloping curve is the benefit from one-worker's transition from unemployment to employment. The benefit reflects the fact that the flow value created by the worker at the new job exceeds the flow value of being unemployed. The first of these flow values is the worker's marginal product at the new job and the second is the amenity value of not having to work plus the rate of unemployment compensation plus the flow probability of finding a job multiplied by the gain in value from being employed rather than unemployed. The benefit slopes downward because of the role of the job-finding rate in the searcher's opportunity cost. In a tight market with a high job-finding rate, the benefit of taking a particular job is

15

smaller because the probability is higher of finding another job soon.

The upward-sloping curve is the cost to the two parties of the incremental efforts needed to bring about the hire this period (in the formal version of this analysis, the marginal benefits and costs are measured in terms of infinitesimal changes in matching resources). Except at the end of the paper, I will follow the labor-market-matching literature in presuming that the marginal cost curve slopes upward, corresponding to diminishing returns in the production function that creates matches.

The labor market is in equilibrium at the intersection of the marginal benefit and marginal cost curves. As I noted earlier, the labor market is always very close to its equilibrium and approaches the equilibrium rapidly in a matter of a few months. Thus a recession in this model—a period with a lower than normal job-finding rate—would require a leftward shift of one of the curves.

A decline in the efficiency of the matching and job-creation process would shift the marginal cost curve to the left. In principle this would happen if matching became more difficult, help-wanted advertising more expensive, or employee evaluation more time-consuming, but none of these seem likely candidates. A shift in the marginal benefit curve is more plausible. The dashed version of the marginal benefit shows the shift from a moderate decline in productivity. Shimer (2005) explores a similar model in detail and asks how much change in unemployment results from a realistic decline in market-wide productivity. Note the the productivity of work is one of the determinants of the benefit of the transition from unemployment to employment. Shimer's answer is that the resulting change in unemployment is absolutely tiny. The small shift in the intersection shown in the figure illustrates Shimer's point.

The literature does not contain any other suggestions for sources of shifts in the marginal benefit curve. I conclude that the friction-free fundamental model offers little hope for an explanation of the observed volatility of unemployment. The model does provide a good starting point for understanding the general level of unemployment.

The coordination between searchers and employers implicit in the friction-free model is beyond the implausible. Recall that all of the relevant coordination occurs before the parties make their deal. Part of it occurs before they actually meet and the rest between meeting and making the deal.

The difference between the labor market in recession and in normal times is a lower job-finding rate and thus a lower opportunity cost for the searcher. In other words, the higher marginal benefit of a transition from search to employment in a recession is the result of a cost bearing on the searcher in a recessionary labor market. A primary coordination problem is to pass this benefit as an incentive to the actor who makes the important decision.

Empirical and theoretical analysis of the allocation of effort in matching and job creation is in its infancy. The standard model has attributed all of the variable effort to the employer. The employer incurs an elective cost to create a vacancy. Searchers make no choices about resources. In that model, the issue is simply what part of the total incentive for matching and job-creation faces the employer.

The wage is a natural tool for allocating incentives between the two parties. If the employer expected to pay a wage that reflected the opportunity cost of labor, then the employer's incentive would be strong to spend on additional matching and job-creation activities in recessions and the labor market would be governed by curves essentially the same as in Figure 7. (I should note by way of parenthesis that I will skip over issues of the difference between the social and private versions of the curves in Figure 7. Because search involves externalities, the two versions are different. See Hosios (1990) for a thorough discussion in the standard model.)

The strongest incentives for employer effort would be achieved in a setting where workers received only their opportunity cost and did not participate in the joint benefit, which would go entirely to the employer. But this arrangement does not make practical sense. First, searchers surely do make important decisions based on incentives. Matching would not occur if workers lacked any reason to play their parts. Second, the resulting equilibrium involves a huge gap between the wage and the marginal product of labor.

A thoroughly studied case has the wage set as a Nash bargain between the searcher and the employer. In the bargain, the parties share the surplus according to specified shares, often equal shares. The wage is the corresponding weighted average of the searcher's opportunity cost and the marginal product. In this case— the one studied by Shimer—employers face powerful incentives to correct incipient recessions because half of the decline in the searcher's opportunity cost accrues to the employer. A diagram for this case looks much like Figure 7, with the vertical axis labeled as the employer's value rather than the joint value.

17

If the wage anticipated by the recruiting employer is fixed, the situation is remarkably different. Figure 8 displays it as a diagram, from the employer's perspective. For the moment, I will retain the assumption that searchers do not make decisions about resources devoted to forming matches and creating jobs. In the new figure, the marginal benefit curve is flat rather than downward sloping—the wage no longer passes along the higher joint value of an unemployment-job transition associated with the lower opportunity cost of the searcher. Equally important, the vertical position of the employer's value is now extremely sensitive to its underlying determinants. A small decline, for example in productivity, lowers the marginal benefit curve substantially, as shown by the dashed line. The marginal benefit is the difference between the marginal product of labor and the wage. In a flexible-wage setup such as the symmetric Nash bargain, or in the joint analysis of Figure 7, the effect of a change in productivity is washed out almost entirely. In the Nash-bargain case, the wage falls by only a hair less than productivity falls. In the joint-value case, the opportunity cost of searchers—driven largely by the marginal product in other jobs—falls by almost exactly the amount that productivity falls. In the fixed-wage case, the marginal benefit is the marginal product less the wage, which has a high elasticity with respect to the marginal product.

The conclusion is that sticky wages are a powerful way to generate fluctuations in unemployment from small changes in driving forces, such as productivity. Shimer (2005) pointed this out initially. My other papers demonstrate it in a variety of settings, including full general equilibrium.

7 Sticky wages in a bargaining setting

My earlier papers make the following point: the bargaining problem that is at the heart of the labor-market analysis of this line of thought has a bargaining set for the wage (in the simplest version, the present value of the wage over the course of employment). The bargaining set runs from the opportunity cost of the searcher at the low end to the marginal product of the worker in the new job, at the high end. Fundamentally, any wage in the bargaining set is an equilibrium. This idea can be formalized by setting up the problem as a Nash demand game—see Hall (2005b). The natural instinct of the bargaining theorist is to put additional structure on the bargaining so as to get a determinate outcome—invariably the Nash bargain. But a
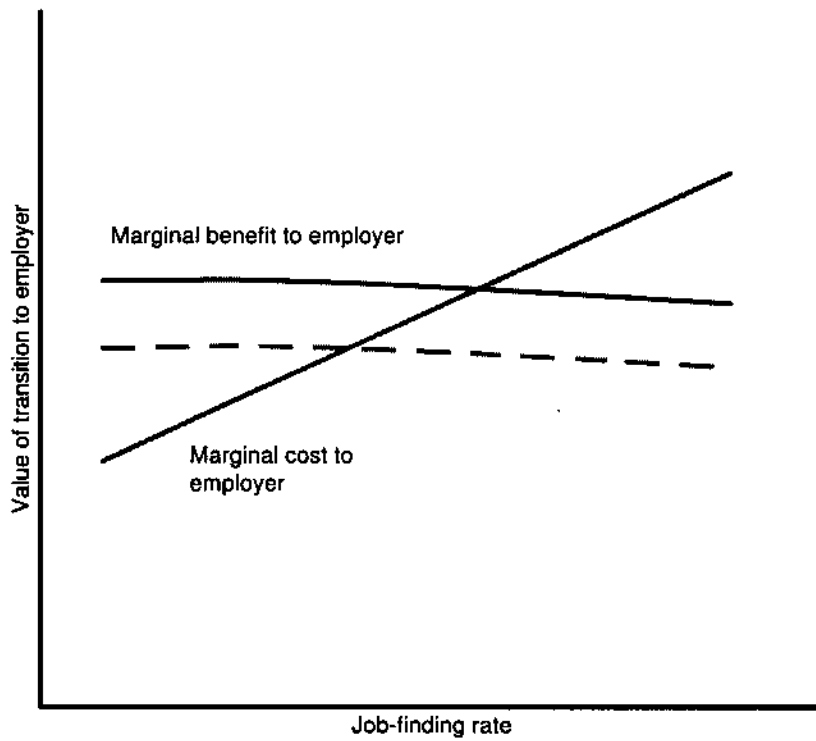
Figure 8. Equilibrium with a sticky wage and matching resources controlled by employers

wage that is fixed over time or over states of the world is also an equilibrium, provided the bargaining set always contains the wage. A sticky wage is an equilibrium selection rule. A possible rationalization for a sticky wage is as a social norm.

8 Worker and employer effort in matching and job creation

Sticky wages have the dramatic effect explained earlier because they isolate the employer from recession-correcting incentives and employers make all the decisions about matching and job-creation effort. The model would lose its power to explain recessions if searchers could respond to the strong incentive that a sticky wage gives to them for raising job-seeking effort when the labor market is slack. After all, the incentive that a sticky wage takes away from the employer is given to the searcher.

Thus a top priority for additional research in this area is better facts and theories about the roles of searchers as well as employers in matching and job creation. It may be useful to divide the process into two parts. One is the technology by which searchers and employers come into contact with each other. This includes help-wanted advertising, employment agencies, talking to friends, and making phone calls. All have taken on Internet versions over the past decade. The second part is what occurs at the employer's site between initial contact and making the employment bargain. Employers evaluate applicants, often using formal testing. They collect information from central sources, especially about criminal records. They check with references. Workers evaluate jobs, generally in a more informal way.

Help-wanted advertising is an interesting example of the asymmetry that matters in the model. In normal times, newspapers are full of help-wanted ads placed by employers. The volume of position-wanted ads is vastly lower. Employers make most of the resource-allocation decisions for printed advertisements. On the other hand, the allocation of effort and payment for employment agencies is more symmetrical.

I suspect that the greater share of the cost of matching and job-creation costs occur in the second phase, at the employer's site. Here the asymmetry is quite pronounced. The searcher contributes time, but the active resource allocation decision falls on the employer.

In some labor markets, searchers can buy credentials, which presents an opportunity for searchers to exert effort to make up for what is missing among employers when the labor market is soft. The rider of any subway is familiar with the amazing variety of vocational certificates available for a few months' study. The purchase of the certificate by the searcher replaces evaluation effort by the prospective employer, provided the employer respects the certificate. I believe that this part of the education industry is at least as counter-cyclical as other types of education that enroll adults.

9   Other ideas in this framework

Eva Nagypal has made the interesting observation that variations in the anticipated duration of employment have important effects on the surplus from the transition from unemployment to employment. If workers hired in a soft labor market are more likely to leave as the market improves, employers will put less effort into recruitment and the softening of the market will be amplified.

Another idea challenges the upward-sloping marginal cost schedule in Figures 7 and 8. Suppose that job-seekers create congestion costs for employers when the job-finding rate is low and unemployment is high. In particular, in that situation, searchers apply for jobs that are not likely to be good matches, but which offer something better than continuing unemployment. By contrast, in a strong labor market, searchers who apply for a particular job are self-selected as good candidates. Employers then have to spend more on screening in soft markets. If this effect is strong enough, it could result in a downward-sloping marginal benefit schedule. The external effect of congestion would swamp diminishing returns to recruiting effort. If the slope of the marginal cost curve were similar to the slope of the marginal cost curve, the model would make employment sensitive to small shifts in the marginal benefit.

10   Concluding remarks

Frictions in the labor market have long been seen as the most promising explanation for the substantial variations in aggregate resource utilization over the business cycle. Recent work is refining our understanding of the nature of the frictions. In

the modern labor market, frictions do not cause employed workers to lose their jobs. Unemployment does not rise in a recession because of waves of layoffs. Rather, the key friction is in re-employing the normal flow of people out of jobs and into unemployment. Recessions are times when unemployment rises because the rate of escape from unemployment is abnormally low.

Powerful incentives might bring unemployment back to normal rapidly. The joint gain to a worker and an employer when the worker moves from unemployment in a recessionary labor market to a new job is substantial—much higher than in normal times, because the opportunity cost of the unemployed is lower. The key friction appears to arise in the transmission of the incentive from unemployed workers to employers, who deploy the resources that result in re-employment. A likely form of the friction is a sticky wage. No principle of economic equilibrium is broken if the sticky wage lies within the bargaining set for a new worker and a potential employer. Sticky wages are a form of equilibrium selection mechanism.

New facts and ideas about the role of labor-market frictions in the business cycle have opened up important new topics for research. We need a better understanding of the process of recruiting, matching, and job creation. Existing theories relegate workers to a passive role in this process. We need to know if that is true, and, if so, why labor-market institutions are unsuccessful in passing incentives from workers to employers.

References

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

Blanchard, Olivier J. and Peter A. Diamond, "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity*, 1990, (2), pp. 85–143.

Cole, Harold and Richard Rogerson, "Can the Mortensen-Pissarides Matching Model Match the Business Cycle Facts?," *International Economic Review*, 1999, *40* (4), pp. 933–960.

Davis, Steven and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, August 1992, *107* (3), 819–863.

Hall, Robert E., "Lost Jobs," *Brookings Papers on Economic Activity*, 1995, *1995* (1), pp. 221–273.

_____ , "Macro Volatility," July 2004. Hoover Institution, Stanford University.

_____ , "Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market," *Review of Economics and Statistics*, 2005. Forthcoming.

_____ , "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, March 2005. forthcoming.

Hosios, Arthur J., "On the Efficiency of Matching and Related Models of Search and Unemployment," *Review of Economic Studies*, April 1990, *57* (2), 279–298.

Klemperer, Paul, "Why Every Economist Should Learn Some Auction Theory," in M. Dewatripont, L. Hansen, and S. Turnovsky, eds., *Advances in Economics and Econometrics: Invited Lectures to Eighth World Congress of the Econometric Society*, Cambridge, England: Cambridge University Press, 2003, p. 2555.

Kydland, Finn and Edward C. Prescott, "Time to Build and Aggregate Fluctuations," *Econometrica*, November 1982, *50* (00), 1345–1370.

Shimer, Robert, "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, March 2005. Forthcoming.