American Economic Review: Papers & Proceedings 2008, 98:2, 268–273
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.2.268

The Probability of Finding a Job

By ROBERT SHIMER*

Recent research has reaffirmed that the probability of an unemployed worker finding a job varies substantially over the business cycle. Robert E. Hall (2005, 101) concludes from his examination of a variety of data sources that "Unemployment is high in a recession because jobs are hard to find, not because more job-seekers have been dumped into the labor market by elevated separation rates." Shimer (2007) shows that movements in the job finding probability account for three-quarters of the fluctuations in the unemployment rate in the United States during the postwar period, while movements in the exit rate from employment to unemployment account for the other quarter. Michael W. Elsby, Ryan Michaels, and Gary Solon (2007) argue that movements in the job finding probability accounted for about 65 percent of unemployment fluctuations prior to the last two recessions, and more in 1990–1991 and 2001. Shigeru Fujita and Gary Ramey (2007) claim a more substantial role for the exit rate to unemployment but still find that the job finding probability accounts for at least half of the fluctuations in unemployment.

While it is theoretically convenient to discuss a single job finding probability for all workers, economists have long recognized that the job finding probability falls with unemployment duration (Hyman B. Kaitz 1970). This paper reexamines duration dependence and the cyclicality of duration dependence in the job finding probability, both empirically and theoretically. To start, I develop a simple model with a single parameter that determines both how the job finding probability varies with unemployment duration and how it varies with aggregate economic conditions. The model's main departure from most existing research is to think of unemployed workers as waiting for labor market conditions to improve, rather than searching for job opportunities (Boyan Jovanovic 1987; Fernando Alvarez and Shimer 2007). They continuously compare their lifetime utility in the best available job with their lifetime utility if they remain unemployed. Individual i works if this difference, $\delta_i(t)$, is positive at time t and not if it is negative. If δ_i is persistent, this leads to duration dependence in the hazard rate of exiting unemployment, since a newly unemployed worker is more likely to be near the threshold for taking a job than someone who is long term unemployed. The extent of duration dependence is governed by the stochastic process for δ_i, which also determines how the average job finding probability varies with economic conditions.

I then compare the predictions of this model with data, confirming that workers who have been unemployed for longer are less likely to find a job. Moreover, when the aggregate job finding probability is lower, the job finding probability falls uniformly at all unemployment durations. Both predictions are quantitatively consistent with the model. The last section of this paper compares this model of duration dependence with important alternatives, including those based on heterogeneity among the unemployed.

I. Model

Consider an individual i and let $\Delta_i(t, dt)$ denote the difference between the expected lifetime utility from taking the best available job for a short interval of time $(t, t + dt)$ and behaving optimally thereafter, and the expected lifetime utility from being unemployed during $(t, t + dt)$ and behaving optimally thereafter. Let $\delta_i(t) \equiv \lim_{dt \to 0} \Delta_i(t, dt)/dt$ denote the instantaneous value of working rather than remaining unemployed.

I make two key assumptions. First, $\delta_i(t)$ exists and satisfies $d\delta_i(t) = \mu dt + \sigma dz_i(t)$, where the

* University of Chicago, Department of Economics, 1126 East 59th Street, Chicago IL, 60637 (e-mail: robert.shimer@gmail.com). This paper builds on research with Fernando Alvarez. I am grateful for comments from Michael Elsby and participants at the American Economic Association Annual Meetings and for support by a grant from the National Science Foundation.

drift μ is nonnegative,[1] the standard deviation σ is positive, and $z_i(t)$ is a standard Brownian motion, with increments that are independent over time and across workers. A deeper model would derive the behavior of $\delta_i(t)$ from the equilibrium of a labor market hit continuously by idiosyncratic shocks and from time-variation in the worker's value of leisure. Here, I simply assume that the drift and standard deviation are common across individuals and constant over time. Second, building on Alvarez and Shimer's (2007) model of rest unemployment, I assume that individual i works if and only if $\delta_i(t)$ is nonnegative. I am interested in computing the probability that individual i is employed at $t + 1$, given that she is unemployed with duration τ at t.

As a preliminary step, and following the logic in Section 5.4 of Alvarez and Shimer (2007), I derive the hazard rate of finding a job for a worker with unemployment duration τ. Consider a worker who is currently unemployed, so the instantaneous value of working is $\delta < 0$. Standard results on first passage times (e.g., Samuel Karlin and Howard M. Taylor 1975, chap, 7, Theorem 5.3) imply that the density of the amount of time τ until the worker first reenters employment is

$$g(\tau,\delta) = \frac{-\delta}{\sigma\sqrt{2\pi\tau^3}}e^{-(\delta+\mu\tau)^2/2\sigma^2\tau}.$$

Let $G(\tau,\delta) \equiv 1 - \int_0^\tau g(s,\delta)\,ds$ denote the probability the worker is still unemployed after τ periods. Then, the hazard rate of finding a job after τ periods is $g(\tau,\delta)/G(\tau,\delta)$. Since a newly unemployed worker has $\delta = 0$, take the limit as $\delta \rightarrow 0$ using L'Hôpital's rule to get the hazard rate of finding a job for a worker whose current unemployment spell has lasted τ periods:

$$(1)\quad h(\tau) = \frac{\phi\left(\frac{\mu\sqrt{\tau}}{\sigma}\right)}{2\tau\left(\phi\left(\frac{\mu\sqrt{\tau}}{\sigma}\right) - \frac{\mu\sqrt{\tau}}{\sigma}\left(1 - \Phi\left(\frac{\mu\sqrt{\tau}}{\sigma}\right)\right)\right)},$$

where Φ is the cumulative distribution of the standard Normal and ϕ is the density. When $\mu = 0$, the hazard is simply $1/(2\tau)$. An increase in the drift μ or decrease in the standard deviation σ raises the hazard at all τ.

Now, suppose an individual has unemployment duration τ at some time t. The density of the time when she first finds a job, $t + s$, is a standard function of the hazard rate,

$$h(\tau + s)e^{-\int_0^s h(\tau+v)\,dv}.$$

But note that, conditional on becoming employed at $t + s \in (t, t + 1)$, the worker may no longer be employed at $t + 1$; the probability she is employed is equal to the probability that $\delta(t + 1) \geq 0$ conditional on $\delta(t + s) = 0$. Since $\delta(t + 1) - \delta(t + s)$ is normally distributed with mean $\mu(1 - s)$ and standard deviation $\sigma\sqrt{1 - s}$, the probability of this event is $\Phi(\mu\sqrt{1 - s}/\sigma)$. Combining these two pieces and integrating over $s \in [0, 1]$ gives the probability that a worker who has unemployment duration τ at t is reemployed at $t + 1$:

$$(2)\quad F(\tau) \equiv \int_0^1 \Phi\left(\frac{\mu\sqrt{1 - s}}{\sigma}\right)h(\tau + s)e^{-\int_0^s h(\tau+v)\,dv}\,ds.$$

This depends only on unemployment duration τ and on the ratio μ/σ, a single parameter. If $\mu = 0$, $F(\tau) = (1 - \sqrt{\tau/(1 + \tau)})/2$. If a time period is a month, this falls from 50 percent for a worker who has just entered unemployment to about 2 percent for a worker who has been unemployed for a year ($\tau = 12$). As with the hazard rate, one can prove that higher values of μ/σ raise the reemployment probability at all durations. In any case, $F(\tau)$ equals $\Phi(\mu/\sigma)$ at $\tau = 0$ and falls monotonically to $\Phi(\mu/\sigma) - \phi(\mu/\sigma) \times (\sqrt{\pi/2} + \mu/\sigma) \geq 0$ as τ gets large. For intermediate durations, closed-form solutions are not available, but the integrals are straightforward to compute numerically.

[1] One can analyze the model with $\mu < 0$, in which case a worker stays unemployed forever with positive probability. This case turns out to be quantitatively uninteresting and changes some expressions, and so I omit it for brevity.

II. Data

The theory provides a measure of the probability that a worker with unemployment duration τ at time t finds a job by time $t + 1$. This section constructs an empirical counterpart of that measure for the United States from January 1976 to October 2007.

I use the public-use microdata from the Current Population Survey (CPS), the main US household labor force survey. Each month, every unemployed individual reports her unemployment duration in weeks. I multiply this by 12/52 to convert to months. I then observe the individual's employment status the following month. I measure F_τ as the fraction of individuals who report duration τ one month and are employed the next month.

The actual implementation of this procedure is somewhat more cumbersome. First, I restrict the sample to individuals who are either job losers or job leavers, about 60 percent of unemployment, dropping new entrants and reentrants to the labor force. The model does not explain why individuals choose to enter or exit the labor force and hence does not predict an initial value for δ_i for new entrants and reentrants.[2] Second, to observe an individual's employment status the following month, I take advantage of the fact that the CPS is a rotating panel. Each address is in the survey for four consecutive months, out for eight months, and then in again for four months. Hence, there is a chance to observe the subsequent month employment status for about three-quarters of the CPS sample.[3] Unfortunately the attrition of households that move or otherwise leave the survey is not representative of the population. Third, since the redesign of the CPS instrument in 1994, unemployment duration data are often imputed for all workers except those in the first and fifth rotation groups, i.e., those who were not surveyed in the previous month. I use data only from these "incoming rotation groups" after 1994, which reduces the sample size by a factor of three. I was ultimately able to match 506,000 records during the 32-year period. Fourth, I weigh observations using the CPS final weights and seasonally adjust the number of people with duration τ with each subsequent employment status (employed, unemployed, or inactive) using a difference-from-moving-average.

Let $n^e_{\tau,t}$ denote the weighted number of people with unemployed duration τ in month t who are employed in $t + 1$, $n^u_{\tau,t}$ denote the number who are unemployed, and $n^i_{\tau,t}$ denote the number who are inactive (not in the labor force). I measure the job finding probability at each duration τ as[4]

$$F_\tau = \frac{\sum_t n^e_{\tau,t}}{\sum_t (n^e_{\tau,t} + n^u_{\tau,t} + n^i_{\tau,t})}. \tag{3}$$

The dots in Figure 1 show the results: 51 percent of workers reporting one week duration find a job before the next survey and the share declines sharply thereafter. For workers with duration less than six months, the job finding probability averages 31 percent. It falls to 19 percent during the next six months and just 14 percent for workers who have been unemployed for over a year.[5]

The line in Figure 1 shows the behavior of the model with $\mu = 0.6\sigma$. The probability that a worker who has just become unemployed ($\tau = 0$) is employed at the next survey date is $\Phi(\mu/\sigma) = 0.73$. It falls rapidly, reaching 50 percent after just one week of unemployment and then falling below 20 percent after 5.5 months. Although not an exact fit—the model overestimates the job finding probability during most of the first 3

[2] The job finding probability for reentrants as a function of duration is similar to the probability for job losers and job leavers, while the probability for new entrants is lower and falls less with duration.

[3] The procedures and pitfalls for matching observations in the CPS are well-known (John Abowd and Arnold Zellner 1985; James Poterba and Lawrence Summers 1986). I match individual records using rotation group, household identifiers, individual line numbers, race, sex, and age. There are a few gaps in the series due to changes in the household identifiers in the public-use files: it is impossible to match data for December 1977/January 1978, June 1985/July 1985, September 1985/October 1985, December 1993/January 1994, and May 1995/June 1995 to August 1995/September 1995.

[4] An alternative would be to use the relationship between the number of unemployed with duration τ at t and the number with duration $\tau + 1$ at $t + 1$ (Hal Sider 1985). This leads naturally to a measure of the unemployment continuation probability, the fraction of individuals who report duration τ one month and are unemployed the next month. Since the model does not draw a strong distinction between unemployment and inactivity, measuring the job finding probability is more useful.

[5] Week-to-week, the figure shows substantial fluctuations, for example increasing at 14 weeks, just over 3 months. This is a duration that is reported relatively infrequently due to "digit preferences" (Michael Baker 1992).

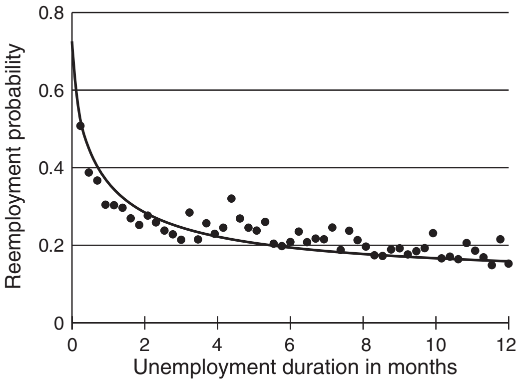

FIGURE 1

Notes: The dots show the empirical probability of being employed at the next survey as a function of unemployment duration, F_τ, for job losers and job leavers. The line shows the theoretical reemployment probability $F(\tau)$ as a function of duration for $\mu = 0.6\sigma$.

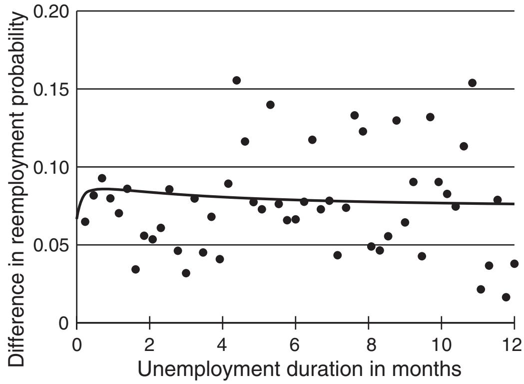

FIGURE 2

Notes: The dots show the difference between the empirical probability of being employed at the next survey as a function of unemployment duration in months when the reemployment hazard is at least 2 percent higher than the mean for job losers and job leavers, and the probability in months when the reemployment hazard is at least two percent lower than the mean, $F_\tau^H – F_\tau^L$. The line shows the difference between the theoretical reemployment probability $F(\tau)$ with $\mu = 0.7\sigma$ and $\mu = 0.5\sigma$.

months of unemployment and then subsequently underestimates the job finding probability—it is worth stressing that one parameter, μ/σ, determines the level, slope, and curvature in the model.[6]

To see how the duration-conditional job finding probability changes with the business cycle, I divide the available data into three bins based on the average job finding probability during the month. In the full sample, an average worker finds a job with probability

$$\frac{\sum_\tau \sum_t n^e_{\tau,t}}{\sum_\tau \sum_t (n^e_{\tau,t} + n^u_{\tau,t} + n^i_{\tau,t})} = 0.286.$$

I split the sample based on whether the average job finding probability in month t,

$$\frac{\sum_\tau n^e_{\tau,t}}{\sum_\tau (n^e_{\tau,t} + n^u_{\tau,t} + n^i_{\tau,t})},$$

lies within 2 percentage points of the mean, is above 0.306, or is below 0.266. This classification assigns 32 percent of months to a high job finding probability, $t \in H$, and 24 percent to a low job finding probability, $t \in L$. I then compute the job finding probability at each duration for these subsets of months,

$$F^H_\tau \equiv \frac{\sum_{t \in H} n^e_{\tau,t}}{\sum_{t \in H} (n^e_{\tau,t} + n^u_{\tau,t} + n^i_{\tau,t})},$$

and similarly for F^L_τ.

Figure 2 shows that F^H_τ lies on average about 8 percent above F^L_τ, with no systematic variation as a function of unemployment duration. In other words, F^H_τ is simply an upward-shifted version of F_τ, while F^L_τ is a downward shift. The curve in the same figure shows the difference between the model-generated value of $F(\tau)$ with $\mu = 0.7\sigma$ and the one with $\mu = 0.5\sigma$.[7] Again, this is

[6] One way to reduce the gap between model and data would be to introduce a small cost of taking a job, so an unemployed worker becomes employed only when $\delta_i(t) \geq \varepsilon \geq 0$, while an employed worker becomes unemployed when $\delta_i(t) < 0$. Positive values of ε reduce the job finding probability at short durations relative to that at long durations. Since, in the baseline model with $\varepsilon = 0$, almost all unemployment spells end with zero duration, this has the added benefit of eliminating many implausibly short spells.

[7] This simple comparative static ignores some potentially interesting dynamics. Suppose the drift suddenly falls from 0.7σ to 0.5σ. The distribution of δ_i changes only gradually, and hence the reemployment hazard takes some time to fully adjust to the shock. Still, t months after the change in the drift, the reemployment hazard for all

nearly independent of unemployment duration. To summarize, the model and data both imply that the magnitude of the decline in the job finding probability is similar for all workers, regardless of unemployment duration. Of course, this means that the percentage decline is largest for the long-term unemployed, so a reduction in the job finding probability disproportionately impacts the expected unemployment duration of those workers who are already long-term unemployed.

III. Discussion

I have developed a novel model of duration dependence in the job finding probability and shown that it is consistent with the data along some important dimensions. This observation may lend credence to the notion of rest unemployment developed in Jovanovic (1987) and Alvarez and Shimer (2007). The finding that aggregate shocks change the job finding probability by a similar amount for all workers implies that the percentage decline in the job finding probability is largest for the long-term unemployed. The difficult-to-insure risk of prolonged joblessness is concentrated in a small subset of the population, with potentially serious welfare consequences.

Neither the model nor the empirical analysis has mentioned heterogeneity among the unemployed. If workers are heterogeneous but each has a constant hazard of finding a job, dynamic sorting will lead to a declining aggregate hazard of finding a job as the good workers disappear from the stock of unemployed (see James Heckman and Burton Singer 1984; Stephen Machin and Alan Manning 1999 and the citations therein). Moreover, we know that observable characteristics help to predict which workers find jobs, so dynamic sorting must be empirically relevant. This suggests that ignoring heterogeneity may bias my analysis.

On the contrary, this paper does not ignore heterogeneity. Without heterogeneity, $z_i(t)$ and hence $\delta_i(t)$ would be the same for all workers, and so all unemployment spells would start and end at the same time. Heterogeneity drives the idiosyncratic Brownian motion, which in turn moves $\delta_i(t)$. To connect this with empirical evidence on heterogeneity, suppose $\sigma z_i(t) = \sigma_y y_i(t) + \sigma_\eta \eta_i(t)$, where $y_i(t)$ and $\eta_i(t)$ are independent Brownian motions. Although $y_i(t + \tau) - y_i(t)$ is Normally distributed with increments that are independent over time and across individuals for any t and τ, suppose an econometrician can perfectly forecast $y_i(t)$ using worker characteristics. For example, consider a worker for whom it happens to be the case that $y_i(t + \tau) = y_i(t) + \nu\tau$ for all τ, so $d\delta_i(t) = (\mu + \sigma_y \nu)\,dt + \sigma_\eta d\eta_i(t)$. Then, if the worker becomes unemployed at t, her reemployment probability conditional on this path of $y_i(t + \tau)$ is given by (2) but depends on the ratio of the conditional drift to the conditional standard deviation, $(\mu + \sigma_y \nu)/\sigma_\eta$, rather than the unconditional ratio μ/σ. More generally, conditioning on $y_i(t)$ will complicate the expression for the residual job finding probability, and may even eliminate "true" duration dependence, consistent with the evidence in Machin and Manning (1999, 3117).

The model in principle could also incorporate other explanations for duration dependence, such as loss of skill during unemployment (Christopher A. Pissarides 1992), which would affect the drift μ. Still, it seems unlikely that the sharp drop in the hazard rate during the first few months of an unemployment spell could reflect loss of skill alone, without the accompanying dynamic selection that I highlight in the present model. Still other explanations for duration dependence require a simple reinterpretation of the model. Melvyn G. Coles and Barbara Petrongolo (2003) emphasize the possibility that newly unemployed workers may quickly determine whether they are interested in any of the stock of available jobs. If not, they wait for a suitable job to enter. In reduced form, this looks like newly unemployed workers lie near the work-unemployment margin, which gradually recedes if there are no suitable jobs.

Finally, I comment on the model's strongest assumption, that $\delta_i(t)$ follows a Brownian motion. A reasonable question is whether this can be justified in a dynamic general equilibrium model. The tentative answer is "almost." Alvarez and Shimer (2007) show that a measure similar to $\delta_i(t)$ follows a *regulated* Brownian motion in a version of the Robert E. Lucas, Jr., and Edward C. Prescott (1974) search model. The drift and standard deviation of the

workers with current duration $\tau \leq$ t has reached its new steady-state value. A more thorough analysis of this issue would study an environment in which μ (or σ) follows a stochastic process.

Brownian motion are determined by idiosyncratic shocks to the productivity of intermediate goods producers, as well as the elasticity of substitution between different intermediate goods. The Brownian motion is regulated above by the possibility of new workers entering the market, and it is regulated below by the possibility of exiting for another market. Although this complicates the analysis, the qualitative and quantitative insights of this simple framework for the behavior of reemployment hazards carries over to that model.

REFERENCES

Abowd, John, and Arnold Zellner. 1985. "Estimating Gross Labor-Force Flows." *Journal of Business and Economic Statistics*, 3(3): 254–83.

Alvarez, Fernando, and Robert Shimer. 2007. "Search and Rest Unemployment." Unpublished.

Baker, Michael. 1992. "Digit Preference in CPS Unemployment Data." *Economics Letters*, 39(1): 117–21.

Coles, Melvyn G., and Barbara Petrongolo. 2003. "A Test Between Unemployment Theories Using Matching Data." Unpublished.

Elsby, Michael W., Ryan Michaels, and Gary Solon. 2007. "The Ins and Outs of Cyclical Unemployment." National Bureau of Economic Research Working Paper 12853.

Fujita, Shigeru, and Garey Ramey. 2007. "The Cyclicality of Separation and Job Finding Rates," Federal Reserve Bank of Philadelphia Working Paper 07–19.

Hall, Robert E. 2005. "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years." In *NBER Macroeconomics Annual 20*, ed. Mark Gertler and Kenneth Rogoff, 101–37. Cambridge, MA: MIT Press.

Heckman, James, and Burton Singer. 1984. "A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data." *Econometrica*, 52(2): 271–320.

Jovanovic, Boyan. 1987. "Work, Rest, and Search: Unemployment, Turnover, and the Cycle." *Journal of Labor Economics*, 5(2): 131–48.

Kaitz, Hyman B., 1970. "Analyzing the Length of Spells of Unemployment." *Monthly Labor Review*, 93(11): 10–20.

Karlin, Samuel, and Howard M. Taylor. 1975. *A First Course in Stochastic Processes, Second Edition*. San Diego: Academic Press.

Lucas, Robert E. Jr., and Edward C. Prescott. 1974. "Equilibrium Search and Unemployment." *Journal of Economic Theory*, 7(2): 188–209.

Machin, Stephen, and Alan Manning. 1999. "The Causes and Consequences of Longterm Unemployment in Europe." In *Handbook of Labor Economics*, Volume 3, ed. O. Ashenfelter and D. Card, 3085–3139. New York: Elsevier.

Pissarides, Christopher A., 1992. "Loss of Skill during Unemployment and the Persistence of Employment Shocks." *The Quarterly Journal of Economics*, 107(4): 1371–91.

Poterba, James, and Lawrence Summers. 1986. "Reporting Errors and Labor Market Dynamics." *Econometrica*, 54(6): 1319–38.

Shimer, Robert. 2007. "Reassessing the Ins and Outs of Unemployment." Unpublished.

Sider, Hal. 1985. "Unemployment Duration and Incidence: 1968–1982." *American Economic Review*, 75(3): 461–72.