Hiring chains and the dynamic behavior of job and worker flows

by Christopher Phillip Reicher

Hiring chains and the dynamic behavior of job and worker flows *

Christopher Phillip Reicher

Abstract:

In this paper, I discuss three sets of links which I uncover in the data on aggregate US job and worker flows. Job flows are strongly related to aggregate employment growth, while worker flows are strongly related to employment growth and the unemployment rate. I show that a simple frictionless business cycle model with heterogeneity and a simple form of on-the-job search can explain these links. Job flows respond simply to the cross-section of firm growth, which responds to aggregate employment growth. Worker flows are related to both employment growth and the unemployment rate, and quits and hires are particularly tightly related to each other. Quits and hires interact to form a hiring chain—hires beget quits through on-the-job search, and quits beget hires to replace quitters. High unemployment crowds out quits, shortens the hiring chain, reduces the number of hires, and also results in an elevated layoff rate. The simple hiring chain model fits the data surprisingly well.

Keywords: Job flows, worker flows, heterogeneity, on-the-job search, hiring chain.

JEL classification: E32, J63, J21.

Christopher Phillip Reicher
Kiel Institute for the World Economy
24100 Kiel, Germany
Telephone: +49 (0)431 8814 300
E-mail: christopher.reicher@ifw-kiel.de

Hiring chains and the dynamic behavior of job and worker flows

Christopher Phillip Reicher [a]
Kiel Institute for the World Economy

This version: June 14, 2011

Abstract:

In this paper, I discuss three sets of links which I uncover in the data on aggregate US job and worker flows. Job flows are strongly related to aggregate employment growth, while worker flows are strongly related to employment growth and the unemployment rate. I show that a simple frictionless business cycle model with heterogeneity and a simple form of on-the-job search can explain these links. Job flows respond simply to the cross-section of firm growth, which responds to aggregate employment growth. Worker flows are related to both employment growth and the unemployment rate, and quits and hires are particularly tightly related to each other. Quits and hires interact to form a hiring chain—hires beget quits through on-the-job search, and quits beget hires to replace quitters. High unemployment crowds out quits, shortens the hiring chain, reduces the number of hires, and also results in an elevated layoff rate. The simple hiring chain model fits the data surprisingly well.

Hiring chains and the dynamic behavior of job and worker flows

1. Introduction

In this paper I document a number of basic facts regarding job and worker flows over the cycle for the United States and show that a simple frictionless business cycle model with heterogeneous firm growth and a simple form of on-the-job search can generate these facts. Understanding the behavior of the cross-sectional distribution of firm growth is the key to understanding aggregate job flows, while worker flows are more complicated. To understand the behavior of worker flows, it is necessary to understand the behavior of quits (or employer-employer flows), since quits form a large portion of hires. In the presence of on-the-job search, quits and hires form a positive feedback loop called a "hiring chain" whereby hires beget quits, and quits beget hires to replace quit workers. By introducing hiring chains into the model, it is possible to match the theoretical and actual behavior of worker flows over the cycle with a surprising degree of accuracy.

In both the data and in the hiring-chain model, job flows and worker flows exhibit a number of strong statistical links with each other and with macroeconomic aggregates.[1] First of all, job flows are strongly related to employment growth but are little affected by the unemployment rate; this finding is compatible with the idea that aggregate job flows simply move with the desired distribution of firm growth, with job creation rising during expansions and job destruction falling. Secondly, worker flows are systematically related to both employment growth and the unemployment rate, with hires and quits rising during expansions and falling during periods of high unemployment, and with layoffs falling during expansions and rising during periods of high unemployment.

Thirdly, quits show a particularly strong relationship with hires and the unemployment rate, which is compatible with the idea that high unemployment crowds out job-to-job transitions while hires crowd in quits. This third link is the key to understanding the second. When unemployment rises, quits (which consist mostly of employer-employer flows) fall; attrition falls; and hiring to make up for attrition falls. Since hiring induces quits, quits fall by even more. The mutual dependence between quits and hires forms a hiring chain, and this chain

---

[1] Davis, Faberman, and Haltiwanger (2011) use the term "iron link" to refer to the microeconomic relationships which connect job and worker flows with aggregate conditions; this paper argues that their "iron links" show up as statistical links in the macroeconomic data as well.

shortens when unemployment is high. In addition, high unemployment results in more layoffs since more firms must lay off workers who they could have shrunken through attrition before. In the data, hires, quits, and layoffs all depend on each other and on unemployment, and the model offers clear intuition as to why.

Even though hiring frictions exist in reality, I use a standard frictionless business cycle model to which I add heterogeneity and a simple form of on-the job search but no meaningful hiring frictions. It turns out that ignoring search costs makes it possible to tractably study the consequences of heterogeneity in the context of large firms. Looking at large firms makes the distinction between job flows and worker flows absolutely clear, since many firms hire new workers to replace quitters even if they destroy jobs on the net. A natural extension of this paper would be to combine a rich array of hiring frictions, large firms, and realistic heterogeneity into a large-firm version of the Diamond-Mortensen-Pissarides (Mortensen and Pissarides, 1994, hereafter DMP) model, though to do so would be a major computational undertaking which is best left for future research. In contrast with the DMP model, the indivisible-labor RBC model has a reduced-form linearized solution which is easy to discuss in the presence of heterogeneity and can offer a great deal of intuition as to why job and worker flows behave the way that they do over the cycle. If search costs are small, then a large amount of intuition from the simple RBC model should be expected to carry over into the more complicated DMP model.

The rest of this paper follows the traditional format. Section 2 discusses the cyclical behavior of job and worker flows and in particular discusses the three sets of links. Section 3 lays out the basic model of firm growth. Section 4 discusses the calibration strategy, and Section 5 discusses the implications of the model in more detail. Section 6 concludes.

2. The facts about job flows, worker flows, vacancies, and firm heterogeneity
2.1 Job and worker flows over the cycle

There are a number of data sources which, when taken together, paint a detailed picture of aggregate job and worker flows in the United States. The most comprehensive and longest-spanning dataset contains the annual job creation and destruction data from the Census Department's Business Dynamics Statistics (BDS) series, which run from 1977 through 2009. These relatively new series cover a much longer horizon than the Business Employment

Dynamics (BDM) series, which begin in the early 1990s. Unlike the Longitudinal Research Database, the BDS covers the entire economy, but the data are only at an annual frequency. The BDM program provides quarterly data on job flows, so the BDM data are better suited to looking at business cycles than the BDS despite the shorter sample.

There are two major datasets on worker flows from the establishment side. The Job Openings and Labor Turnover Survey (JOLTS) program goes back to late 2000 and provides monthly data on the number of job openings and worker flows from the establishment perspective. While these data only cover one and a half business cycles, they have provided invaluable information on worker flows which had been lacking since the end of the old manufacturing Labor Turnover Survey (LTS) in 1981. The LTS runs from 1930 to the end of 1981 and covers only manufacturing. LTS data in some form go back to 1919, so the LTS provides a valuable historical picture of an important though somewhat unrepresentative sector. Finally, from the household side, the BLS publishes a CPS-based gross flow series going back to 1990. The CPS series covers transitions among employment, unemployment, and labor force nonparticipation.

In the discussion which follows, I HP-detrend all of the quarterly data using a smoothing coefficient of 10,000 and all of the annual data using a smoothing coefficient of 100. Doing so allows me to concentrate my discussion on the short-run and medium-run components of these series, by which I mean everything except the lowest-frequency components. I prefer to use a higher smoothing parameter to avoid overfiltering the data, which is an especially important issue when looking at short datasets with infrequent recessions. Since the data are already in rates, they should exhibit very little variation in their stochastic trends. Detrending the data using an HP filter with a high smoothing parameter has the most effect on the measured contribution of new establishments to net employment growth, since both rates have exhibited a downward trend over time (as shown in Figure 1). My results are highly insensitive to the choice of a smoothing parameter.

These data together tell an interesting story about the behavior of job and worker flows over the cycle. Table 1 contains the results of regressions of the components of net employment growth rates on net employment growth using data from these different datasets.[2] Looking at the annual BDS data (also shown in Figure 1), job creation in expanding establishments

---

[2] Regressing one component of job flows on total job flows will yield that component's contribution to the variance of net job flows. Appendix A contains a simple proof of this proposition.

contributes about 31% of the variance of net job flows; job destruction in contracting establishments contributes 51%; job creation in entering establishments contributes 8%; and job destruction in exiting establishments contributes approximately 10% of the cycle. The quarterly Business Employment Dynamics (BDM, formerly BED) data paint a similar picture. In those series, expanding establishments contribute about 35% of the variation in employment growth over the cycle, while contracting establishments contribute about 53%. Entering and exiting establishments contribute very little to the cycle. In both datasets, job destruction contributes 60% of the cycle and job creation contributes 40% of it. Furthermore, almost all of the cyclical variation in job flows is accounted for by incumbent establishments. In discussing the cyclical behavior of job flows, it is reasonable to treat establishment entry and exit as exogenous.

Looking at worker flows from the establishment side, the picture becomes more nuanced. Using quarterly data from JOLTS on separations, hires, layoffs and discharges, and quits and other separations between the beginning of 2001 and the beginning of 2011, separation rates seem to be rather acyclical, contributing approximately 3% of the variation in net employment growth. Fluctuations in hiring rates account for 97% of the cycle. However, within separations, layoffs and discharges account for 60% of the cycle, while quits and other separations almost exactly offset layoffs and discharges. It is misleading to say anything simple about separations as an aggregate since they are rather heterogeneous and occur for different reasons. Furthermore, it is particularly important to take the nature of quits into account if one is to make a meaningful statement about worker flows. Hires account for almost all of the cycle, but that does not mean that layoffs do not account for most of the cycle either.

The rough qualitative pattern in the JOLTS data seems to be robust. The old Labor Turnover Survey also provides some evidence which corroborates the evidence provided by JOLTS. The LTS runs from 1919 to 1981, and a fairly consistent series begins in 1930. The LTS contains monthly data on worker turnover rates in the manufacturing sector on the old SIC basis, which are reproduced in print form by the BLS (1978 for the data through 1970, 1982 for the revised data from 1971 through 1981). I use the data since 1947 aggregated to a quarterly frequency so as to exclude the effects of the Depression and World War II. Regressions using the LTS data provide similar qualitative results to the results from the JOLTS data. Compared with the JOLTS, the LTS indicates a slightly reduced role for

accessions, a slightly less procyclical quit rate, a slightly more countercyclical layoff rate, and a somewhat countercyclical overall separation rate. These data are not representative of the entire economy, but they tell a roughly similar story to the story told by the JOLTS data. Unfortunately, comprehensive seasonally adjusted JOLTS data on manufacturing (on an NAICS basis) are not currently available from the BLS due to the excessively irregular behavior of layoffs and discharges, so it is currently difficult to see how the post-2001 manufacturing data and aggregate data compare with each other.

Looking at worker flows from the household side tends to support the story told by the worker flow data from the establishment side, if one thinks of cyclical outflows from employment to unemployment or out of labor force as primarily occurring because of layoffs and discharges. The BLS provides research series on worker flows from the CPS from early 1990 onward, which I aggregate to a quarterly frequency to reduce the influence of measurement error. In the BLS series, employment inflows account for about 43% of the cycle and employment outflows account for 57% of the cycle; the majority of cyclical employment outflows go to unemployment but the majority of cyclical employment inflows come from out of the labor force. There is some lingering doubt as to the accuracy of these results with regard to the distinction between the unemployed and those out of the labor force. Blanchard and Diamond (1990) apply an Abowd-Zellner (1985) correction for misclassification based on reinterview data, and their estimates indicate that flows in and out of the labor force are much smaller than those measured in the raw tabulations. Since misclassification will have a large effect on the measured importance of gross flows affecting the nonemployment states, the numbers in this paper should be taken with a grain of salt. The CPS data do indicate an important role for employment outflows over the cycle; outflows from employment appear to be cyclically similar to layoffs and discharges as reported by JOLTS or the LTS.

2.2 Three sets of statistical links

The first set of statistical links links job flows with aggregate employment growth. Table 2 shows the results of projecting the components of job and worker flows using net employment growth and the unemployment rate, with Newey-West standard errors calculated using a Bartlett kernel with lag length one. These projections are intended to summarize the comovement between job and worker flows, on one hand, and cyclical conditions on the other. It is hazardous to give these regressions a strong structural interpretation since the

5

structural residual for each regression (e.g. a blip in job destruction unrelated to the cycle) would be correlated with net employment growth through the accounting identity linking gross flows with net flows. The projections merely give a reduced-form picture of the comovement between different components of job and worker flows as seen in the data; they could be treated mathematically as a generalization of the variance-decomposition regressions in Table 1 conditional on the unemployment rate.

The results of these projections are nonetheless interesting. Job flows track net employment growth very closely, while they only very weakly track the unemployment rate. The coefficients on net job growth are very similar to the coefficients given by the accounting decomposition in Table 1. The coefficients on unemployment also indicate a very small effect of unemployment on job flows, which in the case of expanding and contracting establishments is statistically significant but economically significant at about -0.04. From a macroeconomic perspective it seems that a strong link exists between job flows and net employment growth. The R-squared values from the regressions in Table 2 are 0.823 for job creation in expanding establishments, 0.058 for job creation in new establishments, 0.902 for job destruction in contracting establishments, and 0.244 for job destruction in exiting establishments. Restricting the coefficient on unemployment to equal zero, these R-squared values go to 0.800, 0.054, 0.887, and 0.243, respectively. Omitting unemployment from the regression has basically no effect on the regression's fit. There seems to be a very tight link in the data between net employment growth and job flows, and this link holds most strongly in continuing firms. There is at most a weak link between unemployment and job flows which appears to be economically unimportant.

There is a second set of links which link worker flows to both net employment growth and the unemployment rate. A simple link between worker flows and employment growth most definitely does not hold, but a more complex link among worker flows, employment growth, and the unemployment rate does appear to hold. In the worker flow regressions in Table 2, the coefficients on unemployment are all large and statistically distinguishable from zero. Hires and quits respond strongly and negatively to unemployment, while layoffs respond positively to unemployment. Hires and layoffs also respond to net employment growth, while quits do not respond much. The R-squared values of the worker flow regressions in Table 2 are 0.865, 0.865, and 0.730 for hires, quits, and layoffs, respectively; macroeconomic conditions explain these flows extremely well.

The behavior of quits at the aggregate level is particularly interesting and deserves more attention; looking at quits and hires gives the third link. A regression of quits on hires and the unemployment rate yields coefficient estimates of 0.492 and -0.271 with standard errors of 0.051 and 0.033, respectively. The regression has an R-squared of 0.926. The relationship between quits and hires is so strong that quits and hires have a correlation of +0.91 in the data as shown in Table 3, while quits and layoffs have a correlation of -0.62. The behavior of quits is so predictable that it is possible to say that there is a link which relates quits to hires and unemployment. As this paper will show, this link can emerge naturally from a simple model of on-the-job search, and the quantitative magnitude of that link matches what one would expect given the fragmentary data on employer-employer flows. Looking further at Table 3, the link between quits and hires appears to be particularly special; the next strongest statistical link is between job creation and job destruction among continuing firms with a contemporaneous correlation of -0.72.

It is possible to distill the empirical results into a few robust points. Continuing establishments contribute the majority of job flows to the cycle even though net establishment entry is highly procyclical. Job destruction is somewhat more important than job creation over the cycle but both job creation and job destruction matter. Layoffs and employment outflows contribute the majority of job and worker flows to the cycle. This is not incompatible with saying that hiring contributes the majority of cyclical worker flows, because quits are highly procyclical. Additionally, there is a set of links which primarily link job flows to net employment growth, and there is an additional set of links among worker flows, employment growth, and the unemployment rate. Finally, there is a particularly strong relationship which links quits to hires and unemployment. Worker flows in general are affected by the unemployment rate to a degree that job flows are not; job flows and worker flows seem to show substantially different medium-run behavior.

2.3 Vacancies and the Beveridge Curve

Figure 2 shows the Beveridge Curve from 1951 through 2011 for the United States using monthly data obtained by extending the JOLTS data with a help-wanted index. I use Barnichon's (2010) help wanted index which is adjusted to account for the presence of online help-wanted advertising after 1995, and the Conference Board's index before that. First I

divide help wanted and vacancies by CES employment. I splice the help-wanted index to the JOLTS vacancy data by multiplying the help wanted index by the ratio between geometric mean vacancies and geometric mean help wanted (which is valid assuming that they are cointegrated). Inspection reveals an almost perfect fit at the splice between November and December 2000, so I do not make any further adjustments to the series.

For graphical exposition, I divide the series into subseries based on troughs in the unemployment rate. After doing this, the cyclical dynamics of vacancies and unemployment become clear. The Beveridge Curve is not really a curve; it would be more accurate to refer to it as a "Beveridge Cigar" or, as Blanchard and Diamond (1990) do, the "Beveridge Loop". During a recession, unemployment rises and vacancies fall sharply. During the early stages of a recovery, vacancies rise while unemployment takes time to fall. As the recovery continues, unemployment falls and vacancies continue to rise. In unemployment-vacancy space, the business cycle traces out a long thin cigar-shaped counterclockwise loop.

3. The basic indivisible labor model with heterogeneity

As one takes the limit of a well-specified DMP model in the absence of frictions, one ends up with an indivisible-labor RBC model. In a sense, the indivisible-labor RBC model takes the logic of Hagedorn and Manovskii (2008) to the extreme by ignoring search costs entirely.[3] There is additional justification for ignoring search costs. Carlsson, Eriksson, and Gottfries (2008) examine Swedish firm-level data and find no evidence that the rate of unemployment contributes to job creation at the firm level, casting doubt upon the relative importance of search and matching frictions. Barron, Berger, and Black (1997), whom Hagedorn and Manovkii cite, show that search costs represent a small share of hours worked, though training costs are somewhat larger but still small.

I use the indivisible-labor RBC model because its aggregate behavior is well-understood and it is an easy model onto which to append heterogeneity. Workers work either a set number of hours or not at all; this behavior can be generated by introducing a nonconvexity into the

---

[3] If search costs are zero, then firms will post vacancies until the surplus to be earned from posting a vacancy equals zero. In that event, workers and firms both earn their outside option; the marginal product of labor will have to equal the disutility from work (in product terms) for each firm-worker combination. In the context of Pissarides (2000), this would involve setting the vacancy cost c in equation (1.7) to zero, implying that the profits from a job J also equal zero. By equations (1.8) and (1.9) which link the surplus to period-by-period profits, profits must be zero and workers must earn their marginal product.

extensive-intensive margin tradeoff. Since eighty percent of US labor market movements happen on the extensive margin, this seems to be a good starting point. The model has the property that workers are indifferent among working, changing jobs, or not working. Given the widespread view that unemployment is not voluntary, this assumption is undoubtedly too extreme, but quantitatively it approximates the behavior that one would expect to see in a model with a very small bargaining surplus, and it is much more tractable.

3.1 Households

Households are identical and have preferences which are balanced-growth compatible and separable in consumption and labor input. The labor force and the measure of firms grow at a constant rate Γ^{LF}. Lifetime expected utility takes the form:

$$E_t \sum_{i=0}^{\infty} \beta^{t+i} \left[\ln(C_{t+i}) - A N_{t+i} \right].$$

The period by period budget constraint relates end-of-period bondholdings, capitalholdings, and personal consumption, with beginning-of-period bondholdings (gross of interest at a rate r_{t-1}), capitalholdings (net of depreciation at a rate δ), capital income (based on a rental rate ρ_t), labor income (based on a wage rate W_t), and firms' accounting profits which are remitted to households. It takes the form:

$$B_{t+1} + K_{t+1} + C_t = (1 + r_{t-1})B_t + \rho_t K_t + (1 - \delta)K_t + W_t N_t + \Pi_t.$$

Households are not allowed to run Ponzi schemes. The first-order conditions for consumption, bonds (which are in zero net supply), capital, and labor inputs are the usual ones:

$$\lambda_t = 1 / C_t; \tag{1}$$

$$\lambda_t = (1 + r_t)\beta E_t \lambda_{t+1}; \tag{2}$$

$$\lambda_t = \beta E_t \lambda_{t+1}(1 - \delta + \rho_{t+1}); \tag{3}$$

9

and

$$\lambda_t W_t = A, \tag{4}$$

with a transversality condition which rules out bubbles in capital accumulation:

$$\lim_{T \to \infty} \beta^T E_t \frac{\lambda_{t+T}}{\lambda_t} K_{t+T+1} = 0. \tag{5}$$

There is no government spending. Since the government does not have any debt, market clearing in debt markets ensures that total bondholdings B_{t+1} equal zero in every period.

3.2 Production and factor demand within firms and in the aggregate

The number of firms is proportional to the number of participants in the labor force. Each firm is located in a competitive sector, and each sector is located somewhere on the unit interval. The distribution of firms among sectors is uniform. Firms seek to maximize period by period accounting profits, and they can adjust inputs immediately. They are price takers in output and factor markets. They produce according to a Cobb-Douglas production function with sector-specific productivity a_{it}, economywide productivity z_t, and factor inputs k_{it} and n_{it}:

$$y_{it} = a_{it} k_{it}^{\alpha} (z_t n_{it})^{1-\alpha}. \tag{6}$$

Profits depend on each firm's own output price, productivity, input choices, and factor prices:

$$\Pi_{it} = p_{it} a_{it} k_{it}^{\alpha} (z_t n_{it})^{1-\alpha} - W_t n_{it} - \rho_t k_{it}. \tag{7}$$

Firm-level output is aggregated into economywide output according to the CES aggregator:

$$Y_t = \left[\int_0^1 y_{jt}^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}}. \tag{8}$$

Total expenditure (with the price level normalized to one) also equals the integral of firm-level expenditure:

$$Y_t = \int_0^1 p_{jt} y_{jt} dj .$$ (9)

The first-order conditions at the micro level are standard. Appendix B goes through the derivations underlying the relationship between firm-level equilibrium and aggregate equilibrium. In the aggregate, income shares are constant:

$$W_t N_t = (1 - \alpha) Y_t ,$$ (10)

and

$$\rho_t K_t = \alpha Y_t .$$ (11)

The unemployment rate is given by the share of people who are not working at a given moment. I abstract from flows in and out of the labor force because of the severe degree of classification error in the CPS between the categories of workers who are unemployed and out of the labor force. Blanchard and Diamond (1990) document the behavior of gross flows from 1968 through 1986 and they find that the actual number of flows in and out of the labor force appear to be just over half of those reported by a simple tabulation of the CPS data, and that many of those flows are seasonal flows of 16 to 19 year-olds in and out of the labor force over the summer. They use data corrected according to the corrections of Abowd and Zellner (1985) based on reinterview data which are not available on a consistent basis. Much more work needs to be done to determine the effect of misclassification on the measurement of labor force flows and stocks, which is why I am reluctant to emphasize such movements in this paper.

I treat unemployment therefore as the complement of employment, with a concept of an acyclical true or unmeasured labor force lurking in the background:

$$U_t = 1 - N_t .$$ (12)

Accounting profits Π_t are zero because of free entry and constant returns. The economy ends up producing according to an economywide production function with an additional productivity shifter term l_t which reflects the contribution of firm-level productivity dispersion toward aggregate productivity:

$$Y_t = l_t^{1-\theta} K_t^{\alpha} (z_t N_t)^{1-\alpha}, \tag{13}$$

where the productivity dispersion term is given by:

$$l_t = \left[\int_0^1 a_{it}^{\theta-1} dj \right]^{\frac{1}{\theta-1}}. \tag{14}$$

Aggregate labor N_t and capital K_t in (13) are simple arithmetic sums across sectors.

Market clearing gives the economywide resource constraint:

$$K_{t+1} + C_t = (1-\delta)K_t + Y_t. \tag{15}$$

Aggregate productivity follows a loglinear random walk with drift Γ^z:

$$\ln(z_t) = \ln(\Gamma^z) + \ln(z_{t-1}) + \varepsilon_t^z. \tag{16}$$

An equilibrium in this economy is a situation where all households maximize intertemporal expected utility; all firms maximize profits; households and firms do not run Ponzi schemes; and all markets clear. This economy in the aggregate is a standard RBC economy with an additional term reflecting the aggregate effects of productivity dispersion. The RBC economy is well-understood and provides a useful vehicle to deliver intuition surrounding the effect of dispersion in labor demand.

3.3 Firm-level labor demand, job creation, and job destruction

Based on aggregate real allocations, one can back out the distribution of firm growth and shrinkage given the distribution of firm-level productivity growth. Each firm's share of

overall employment is given by $l_t^{1-\theta}a_{it}^{\theta-1}$, so job growth at the individual firm level is given by:

$$\frac{n_{it}}{n_{it-1}} = \frac{a_{it}^{\theta-1}}{a_{it-1}^{\theta-1}} \frac{l_t^{1-\theta}N_t}{l_{t-1}^{1-\theta}N_{t-1}}.$$ (17)

Assumptions about firm growth may turn this problem into a tractable or an intractable one. At the establishment and firm levels, there is a large literature on the topic of firm size. In general, that literature has found that the growth rate of a firm is independent of that firm's size and age, at least for firms that are not very small or young. This finding is known as Gibrat's Law. Sutton (1997) gives a summary of some of the earlier literature on Gibrat's Law. More recently, Lotti, Santarelli, and Vivarelli (2003, 2009) and Contini and Revelli (1989) show that Gibrat's Law reasonably characterizes the behavior of larger and better-established firms in Italian data, though Gibrat's Law is less accurate when it comes to describing new entrants and small firms. Contini and Revelli summarize the literature on US firm sizes as well, finding the same conditional adherence to Gibrat's Law. Rogers, Helmers, and Koch (2010) look at more recent British data and come to the same conclusion; Gibrat's law holds up fairly well for the third and higher deciles of firm size, which account for the lion's share of employment fluctuations.

Using the high-quality Longitudinal Business Database (LBD) data upon which the BDS series are based, Haltiwanger, Jarmin, and Miranda (2010) discuss the validity of Gibrat's Law in detail for the United States. They find that it holds extremely well among continuing firms and moderately well for small firms, depending on the weighting scheme used. For entering, very young, and very small firms, Gibrat's Law does not hold quite as well. Haltiwanger et al. attribute this failure of Gibrat's Law to a "regression to the mean" effect in firm growth, but there is another issue at work: For a very small continuing firm, there is nowhere to go but up as a statistical matter. It is impossible for a one-person firm to shrink by half, but it is easy for it to double. Haltiwanger et al. also uncover an interesting age-growth pattern, whereby very new firms grow the fastest and then show an "up or out" dynamic. This paper does not concentrate on the behavior of small firms since they represent a small share of aggregate employment, but it is an interesting pattern in its own right. For the purposes of modeling aggregate firm dynamics, it is reasonable to claim that for those firms which employ the majority of workers, Gibrat's Law holds.

13

I assume that firm growth follows Gibrat's Law in a tractable way. One example of a tractable formulation which is broadly consistent with the observed firm age-size profile would be to say that a fraction γ of firms closes exogenously in a given period and new firms exogenously open up, drawing from some initial well-behaved productivity distribution with average end-of-period employment μ_0. In order to accommodate new entrants to the labor force, which grows at a gross rate Γ^{LF}, an additional measure of new firms must enter, so the number of new firms as a proportion of the previous period's labor force equals $\gamma + \Gamma^{LF} - 1$. For those firms who survive the wave of exogenous destruction, idiosyncratic productivity grows according to an iid geometric random walk which obeys Gibrat's Law. Since the cross-sectional distribution of relative firm productivity is constant over time, one could treat l_t as a constant, so that firm-level employment growth follows a lognormal distribution with some drift and dispersion parameter:

$$\ln(g_{it}) = \ln\left(\frac{a_{it}^{\theta-1}}{a_{it-1}^{\theta-1}} \frac{N_t}{N_{t-1}}\right) \sim N\left(\mu_g + \ln(N_t) - \ln(N_{t-1}), \sigma_g^2\right). \tag{18}$$

Some previous models of Gibrat's Law suffer from the criticism that under a lognormal random walk, average firm size may explode. Firm exit mitigates that problem if exit happens quickly enough. To illustrate this, average firm size is given by the recursive equation:

$$\overline{m}_t = \frac{\mu_0\left(\gamma + \Gamma^{LF} - 1\right)}{\Gamma^{LF}} + \frac{(1-\gamma)}{\Gamma^{LF}} e^{\left(\mu_g + \frac{1}{2}\sigma_g^2\right)} \overline{m}_{t-1}. \tag{19}$$

A necessary and sufficient condition for average firm size not to explode is that some combination of attrition and labor force growth happen faster than the growth of continuers, so that in particular:

$$\ln(1-\gamma) - \ln(\Gamma^{LF}) + \mu_g + \frac{1}{2}\sigma_g^2 < 0. \tag{20}$$

In the calibration used in this paper, condition (20) holds. The upward growth of firm size for continuing firms does not outrun the continual replenishment of firms in the smaller size categories, so the mean firm size exists.

Job creation and destruction are defined as in Davis, Faberman, and Haltiwanger (2006), except that in the dynamic model I use the previous period's aggregate employment N_{t-1} as the denominator when calculating rates, while Davis et al. use the average between the previous period's employment level and the current period's. Such a distinction in a quarterly business cycle is numerically unimportant. Job creation and destruction are defined at the firm level as net positive and net negative employment growth in a given period, respectively. Since economywide job flows and worker flows equal the aggregate over each firm, it is possible to come up with the economywide job creation and destruction rates by simply accounting for exogenous job creation and destruction and then integrating over g_{it} to obtain the endogenous portion of job and worker flows.

3.4 Deriving aggregate job flows from firm-level flows

Job destruction encompasses jobs destroyed for either of two reasons. A fraction γ of firms are destroyed every quarter exogenously at the end of the period, after production takes place. The remainder of job destruction is given by those firms who wish to voluntarily contract given idiosyncratic and aggregate conditions—this is shown in Figure 3 as the integral of negative firm growth given the distribution of net firm growth rates at a particular time. Job flows and worker flows are registered by comparing the end-of-period employment of each firm with its end-of-period employment from the previous period.

The job destruction rate is given by the sum of these two components divided by the previous period's employment:

$$jd_t = \gamma \frac{N_t}{N_{t-1}} + (1-\gamma) \int_{g_{it}=0}^{1} (1-g_{it})dF(g_{it}).$$

(21)

Job creation is the positive portion of net firm growth over the period. The first part of job creation consists of those jobs accounted for by new firms, some of which are exogenously created at a rate γ and others which are created in order to allow for labor force growth. The

rest of job creation consists of jobs created by surviving firms who wish to expand, shown on the right hand side of Figure 3. The job creation rate consists of exogenous and endogenous job creation expressed as a proportion of the previous period's employment:

$$jc_t = \mu_0\left(\gamma + \Gamma^{LF} - 1\right)\frac{N_t}{N_{t-1}} + (1-\gamma)\int_{g_{it}=1}^{\infty}(g_{it}-1)dF(g_{it}). \tag{22}$$

Plugging in the formulas for a lognormal distribution (where Φ is the normal cdf) yields the links which link aggregate job creation and job destruction to aggregate employment growth in a one-to-one way:

$$jc_t = \mu_0\left(\gamma + \Gamma^{LF} - 1\right)\frac{N_t}{N_{t-1}} + (1-\gamma)S_t, \tag{23}$$

and

$$jd_t = \gamma\frac{N_t}{N_{t-1}} + (1-\gamma)\left[S_t + 1 - \frac{N_t}{N_{t-1}}m\right], \tag{24}$$

where

$$S_t = \frac{N_t}{N_{t-1}}m\Phi\left(\frac{\mu_g + \ln(N_t) - \ln(N_{t-1}) + \sigma_g^2}{\sigma_g}\right) - \Phi\left(\frac{\mu_g + \ln(N_t) - \ln(N_{t-1})}{\sigma_g}\right),$$

and

$$m = e^{\mu_g + \frac{1}{2}\sigma_g^2}.$$

The term S_t is common to both job creation and destruction and encompasses the contribution of idiosyncratic firm growth to job creation as shown in the right hand side of Figure 3. During a boom, the entire distribution of firm growth rates shifts rightward, raising job creation rates and reducing job destruction rates. Fewer firms find themselves in the job destruction region and those that do find themselves there tend to destroy fewer jobs, because there are fewer firms sitting on the left hand tail of growth rates. More firms find themselves in the job creation region and those who find themselves there create more jobs than before.

The intuition behind strongly procyclical job creation and strongly countercyclical job destruction is relatively straightforward when viewed through the lens of Figure 3; it is simply a matter of seeing how the entire distribution of firm growth rates evolves with the cycle. Inspecting (23) and (24) shows that the behavior of job flows is entirely driven by the effect of aggregate employment growth on the distribution of desired firm growth.

3.5 Passive on-the-job search and the behavior of worker flows

New vacancies V_t are posted on a bulletin board in the town square each morning after aggregate shocks are realized but before job and worker flows actually occur. Each firm's total desired vacancy posting and hiring is common knowledge at that point. I model on-the-job search in an extremely simple manner. All of the previous day's unemployed people walk past the bulletin board, and a fixed proportion κ of the previous day's employed people do. Those who find jobs that they mildly prefer to their existing job (or lack thereof) begin working immediately thereafter. Hires (given by a rate h_t) occur according to a Cobb-Douglas matching function which reflects these two sources of new hires:

$$ h_t N_{t-1} = M \left(U_{t-1} + \kappa N_{t-1} \right)^{\zeta} V_t^{1-\zeta} . \tag{25} $$

Vacancies are posted at little to no cost; it is useful to think of them as determined residually by the matching function. They represent the amount of help wanted advertising necessary to fill jobs at a desired rate.

To model separations, it is necessary to think about the distinction among job destruction, layoffs, quits, and total separations (the latter given in rates as l_t, q_t, and s_t). In terms of the model, job destruction is very straightforward. Total separations are more complicated—these consist of layoffs and quits. Quits consist of flows between employers (ee_t) in addition to other types of exogenous quits such as retirement or death (q_0)—Akerlof, Rose, and Yellen (1988) and Blanchard and Diamond (1990) discuss the relationship between quits and employer-to-employer flows, and they find that quits into new employment comprise the majority of all quits based on micro evidence. Quits into unemployment are relatively uncommon, so I abstract from quits into unemployment and concentrate on quits into new employment and quits out of the labor force. The quit rate is given by the sum of its two components:

$$q_t = ee_t + q_0. \tag{26}$$

Krause and Lubik (2010) and others have found that including on-the-job search can substantially improve the otherwise-poor cyclical fit of the standard DMP model. Here I show that on-the-job search can interact with heterogeneity in interesting ways, even if it does not affect such things as aggregate labor input. I model quits much more simply than they do; even a simple model of quits helps the model to fit the data on worker flows surprisingly well.

The share of job to job transitions among total hires is given by the share of employed searchers out of all searchers. Previously employed workers and unemployed workers are seen by potential employers as interchangeable. Combined with a provision for the number of exogenous quits, this proportion multiplied by hires yields the total quit rate:

$$q_t = \frac{\kappa N_{t-1}}{\left(U_{t-1} + \kappa N_{t-1}\right)} h_t + q_0 = \frac{\kappa(1 - U_{t-1})}{\left(U_{t-1} + \kappa(1 - U_{t-1})\right)} h_t + q_0. \tag{27}$$

A simple model of on-the-job search automatically gives a link which links quits, unemployment, and hires which corresponds very closely with the one observed in the data. Equation (27) immediately offers some intuition as to why quits should be related to hires and unemployment. Since a proportion of hires comes directly from quits based on raw accounting, the two should be linked as a statistical matter as well. That proportion should vary depending on the ratio of the number of employed searchers to the number of total searchers, which should itself vary strongly based on the unemployment rate.

Firms have no control over which workers quit their jobs and take on those new jobs, since these workers find jobs that they prefer by a tiny amount, and the old firm cannot match the new firm's offer. Firms will lose a certain share of workers to quits which they take as exogenous. Firms which do not undergo total destruction in a given period face a worker survival rate ω_t such that:

$$q_t = (1 - \gamma)(1 - \omega_t). \tag{28}$$

Firms wishing to grow at a gross rate less than ω_t must do so by laying off workers, while those which do not shrink radically still continue to hire workers to replace workers lost through attrition. Figure 4 graphically shows the difference between layoffs and job destruction at the firm level. Job destruction consists of all net negative firm growth. Firms that shrink by only a little bit will do so through attrition without layoffs, while those that wish to shrink by more than $1 - \omega_t$ will have to do so by laying people off. It is possible for a firm to simultaneously destroy jobs and hire people, if its desired job destruction is not as severe as the attrition which would naturally take place due to quits.

Layoff dynamics have some of the same basic intuition associated with them as job destruction dynamics; most dynamics come from changes in the distribution of firm growth rates. There is an added "punch" which comes from quit dynamics. As hiring and quits fall during a recession and the worker survival rate ω_t rises, an extra measure of firms needs to maintain its desired relative growth through layoffs rather than attrition. Since (27) also implies that high unemployment shrinks the quit rate, hiring and quits should remain low and layoffs should remain high when unemployment is high, while job creation and destruction should not vary by much.

Layoffs and quits are the two components which add up to total separations:

$$s_t = l_t + q_t. \tag{29}$$

Total hiring and separation rates are given by integrating over those realizations where a firm would want to hire and fire in relation to the level of employment that it would naturally have through the attrition of workers. Hires consist of hires which occur because of firm entry and hires which are used to keep a measure of surviving firms above the layoff threshold, expressed as a rate:

$$h_t = \mu_0 \left(\gamma + \Gamma^{LF} - 1\right)\frac{N_t}{N_{t-1}} + (1 - \gamma) \int_{g_{it}=\omega}^{\infty} (g_{it} - \omega_t)dF(g_{it}). \tag{30}$$

Separations consist of quits and the exogenous and endogenous portions of layoffs, respectively, expressed as a rate:

$$s_t = (1-\gamma)(1-\omega_t) + \gamma \frac{N_t}{N_{t-1}} + (1-\gamma) \int_{g_{it}=0}^{\omega} (\omega_t - g_{it}) df(g_{it}).$$ (31)

Plugging in the formulas for a lognormal distribution gives aggregate worker flow rates gives analytical expressions linking worker flows to attrition and aggregate employment growth:

$$h_t = \mu_0 (\gamma + \Gamma^{LF} - 1) \frac{N_t}{N_{t-1}} + (1-\gamma) J_t,$$ (32)

and

$$s_t = \gamma \frac{N_t}{N_{t-1}} + (1-\gamma) \left[J_t + 1 - \frac{N_t}{N_{t-1}} m \right],$$ (33)

where

$$J_t = \frac{N_t}{N_{t-1}} m\Phi\left(\frac{\mu_g + \ln(N_t) - \ln(N_{t-1}) - \ln(\omega_t) + \sigma_g^2}{\sigma_g} \right),$$
$$- \omega_t \Phi\left(\frac{\mu_g + \ln(N_t) - \ln(N_{t-1}) - \ln(\omega_t)}{\sigma_g} \right).$$

Unlike job flows, worker flows are related to both the behavior of the cross-sectional distribution of firm growth and to attrition coming from quits. Since quits depend to some degree on unemployment, all of the other worker flows will depend some degree on unemployment as well.

4. The calibration and its implications

Most of the parameters on the RBC side are chosen to fit a standard calibration based on the first moments of US data, some of which are shown in Table 4. Capital's share in the production function is 0.3; depreciation is 2% per quarter; real interest rates are 1% per quarter; trend productivity growth is 0.25% per quarter; trend labor force growth is 0.23% per quarter (given by net BDM net job creation). Government consumption is 16% of output with the rest of government spending lumped into investment, and the unemployment rate is 6%,

while the vacancy rate is 3.61% from the composite help wanted / JOLTS series spanning the period 1951-2011.

The BDS and BDM data provide a good conceptual fit with the model in that both are intended capture medium-run as opposed to short-run firm dynamics. The BDS data do a better job at capturing long-run dynamics, while the BDM data do a better job at capturing short-run dynamics. The BDS data suffer from a large degree of time aggregation bias since many job flows get reversed over the year, so I use the BDM data to calibrate the model. Taking the resulting quarterly volatility and backing out job creation and destruction at annual rates gives something very similar to BDS averages, using equations (21) and (22). Table 4 shows the first moments of each of the flow rates. From the BDM data, job destruction is 7.07% per quarter, and job creation is 7.30% per quarter, the difference accounting for labor force growth. The portion of net job creation accounted for by surviving establishments is 0.13% per quarter. Establishment closings γ are 1.42% of employment per quarter, which gives a value for the establishment growth rate for surviving establishments m of 1.0013 / (1 – 0.0142) = 1.0157. This also means that new establishments are 92% the size of the average establishment, which gives the value of μ_0. The aggregate of idiosyncratic productivity l is normalized to 1. Based on the JOLTS data for 2001-2011, separations are 10.99 percent per quarter. These rates deliver a value of the survival rate ω of 0.931.

Fallick and Fleischman (2004) calculate a series of employer-employer flows using CPS data from 1994 through 2003. They find that hires from employment are 1.6% of the labor force per month, while hires from out of employment are 2.6% of the labor force. This would yield a value of κ of 0.039, though the exact number is subject to the usual caveats related to the CPS flow data. It is possible to cross-check the implications of this parameter choice. The linearized version of the link (27) linking quits with hires and the measured unemployment rate, omitting constants, is as follows:

$$q_t = \frac{\kappa N}{(U + \kappa N)} h_t - \frac{\kappa h}{(U + \kappa N)^2} u_t.$$

(34)

Using the Fallick-Fleischman numbers, the first coefficient should equal 0.381 and the latter coefficient should equal -0.467, while the OLS coefficients are about 0.492 and -0.271, respectively. The link among quits, hires, and unemployment in the simple model and in the

data look remarkably similar to each other. Using the calibrated value of κ, it is also possible to pin down the parameters of the matching function. Blanchard and Diamond (1989, 1991) estimate unemployment's share in the matching function as about 0.4. They do not take job to job transitions into account, and they base their estimates on CPS gross flow data. Using the JOLTS series on hires and vacancies and the calibrated value of κ derived from the CPS gross flow series, a nonlinear OLS regression of the log vacancy filling rate on log inverse labor market tightness with AR(1) errors yields a coefficient of 0.374, which is similar to their estimate. The regression has an R-squared of 0.869.[4]

The calibration strategy produces a quarterly dispersion parameter σ_g of 0.1615 and a drift parameter μ_g of 0.0026 which can be calculated from the job flow equations. There is external evidence on this issue. Davis, Faberman, Haltiwanger, Jarmin, and Miranda (2010) calculate that the cross-sectional dispersion (weighted standard deviation) of establishment-level growth rates across all firms (including opening and closing firms) is 61% per year. They also present a measure of volatility that is cleansed of firm-specific effects, of 46% per year. Both measures include establishment openings and closures, however, and these measures will vastly overstate what happens to those firms which do not open or close. According to their formula which weights openings and closings at one half, the mean and variance of measured idiosyncratic firm growth approximately equal:

$$\frac{\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + (1-\gamma)m - \gamma}{.5\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + 1 + .5\gamma},$$

and

$$\frac{.5\mu_0\left(\gamma + \Gamma^{LF} - 1\right)(2 - mean^*)^2 + (1-\gamma)(m - 1 - mean^*)^2 + .5\gamma(2 + mean^*)^2}{.5\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + 1 + .5\gamma},$$

respectively.

---

[4] Using a value of κ of 0 as in the standard case gives a coefficient of 0.296. Blanchard and Diamond base their estimates on CPS gross flows not including EE flows, while I use the JOLTS data to arrive at my estimates. Approximating their methodology by using log(hires-quits)-log(v/n) as my left hand side variable with a value of κ of 0 gives a coefficient of 0.459.

As one can see, these measures are very sensitive to tail events such as firm and establishment opening and closure. Using the calibrated parameters adjusted to annual rates, the Davis et al. measure of firm growth dispersion would equal approximately 53% per year, which is toward the middle of the published Davis et al. measures. The model seems to do a reasonably good job at matching the firm-level volatility observed in the annual data even when the model is calibrated using quarterly aggregate sources.

Since I am only looking at conditional movements, I set the standard deviation of productivity shocks to 0.01, with a persistence of one. This can be done without loss of generality since I am not concerned with the unconditional second moments of the data.

5. Hiring chains, worker flows, and simulation results
5.1 The behavior of job flows

Figure 5 shows what happens after a negative one percent productivity shock is fed into the model. Aggregate variables move in their expected manner. Investment and labor input fall sharply. Job creation falls and job destruction spikes. During the recovery, job creation slightly outpaces job destruction, but not by nearly the same pace that job destruction outpaces job creation during the downturn. Visually, job creation and destruction show "spiky" behavior relative to worker flows; they do not have much persistence associated with them, and they move with net employment growth according to the logic of Figure 3.

Appendix C goes through the derivation of the linearized model. Job creation and destruction rates can be described in relation to observed employment growth:

$$jc_t = \frac{\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + (1-\gamma)a}{\Gamma^{LF}} n_t^g, \tag{35}$$

and

$$jd_t = \frac{(1-\gamma)(a-m) + \gamma}{\Gamma^{LF}} n_t^g, \tag{36}$$

where

$$a = m\Phi\left(\frac{\mu_g + \sigma_g^2}{\sigma_g}\right) + \frac{m}{\sigma_g}\phi\left(\frac{\mu_g + \sigma_g^2}{\sigma_g}\right) - \frac{1}{\sigma_g}\phi\left(\frac{\mu_g}{\sigma_g}\right).$$

Using the calibrated model, the coefficient of job creation on employment growth is 0.585 and the coefficient of job destruction on employment growth is 0.415. The coefficients in the data are 0.390 and 0.611, respectively. The model and data both agree that job creation and destruction are highly important over the cycle, but they do not exactly agree as to the relative importance of the two flows. The model predicts job creation which is somewhat too responsive to the cycle, which is possibly an artifact of the lack of frictions. The model does predict that job creation and job destruction are both important contributors to the cycle.

5.2 Hiring chains and the dynamics of worker flows

The lower left panel of Figure 5 shows what happens to the various components of worker flows after a productivity shock. The sharp fall in employment has a large effect on all worker flows, with worker flows showing much more persistence than job flows as unemployment takes some time to return to normal. In the short term, a fall in employment requires fewer hires and more layoffs based on the cross-sectional logic of Figure 4. In the medium term, quits remain low because higher unemployment crowds out the transition of workers between employers. Fewer quits imply less attrition and more layoffs; and fewer quits also imply fewer hires to replace workers lost through attrition. The intuition is roughly similar to that underlying the "vacancy chain" model used by Akerlof, Rose, and Yellen (1988) and Contini and Revelli (1990) to study job satisfaction and mobility. In this model, hires beget quits which beget hires, leading to a hiring chain; vacancies are filled instantly.

To investigate the hiring chain idea a bit more, it is possible to write the reduced-form linearized relationship between hires and quits as:

$$h_t = \frac{\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + (1 - \gamma)b}{\Gamma^{LF}}n_t^g + \frac{c}{\omega}q_t, \tag{37}$$

where the constants b and c are positive numbers which are derived in Appendix C, and n_t^g is observed employment growth. Quits depend on hires and unemployment based on (27), and hires depend on quits. Substituting the linearized version of (27) into (37) gives hires as a

function of aggregate employment growth and aggregate unemployment. The denominator of the resulting equation is less than one and greater than zero, so dividing by it gives the length of the hiring chain relative to the original event which triggered the spree of hiring:

$$h_t = \frac{\dfrac{\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + (1-\gamma)b}{\Gamma^{LF}} n_t^g - \dfrac{c}{\omega} \dfrac{\kappa h}{\left(U + \kappa N\right)^2} u_t}{1 - \dfrac{c}{\omega} \dfrac{\kappa N}{\left(U + \kappa N\right)}}.$$ (38)

Comparing (38) with (37) implies that every movement in employment growth is met by an increase in hires, which causes an increase in quits by passive searchers, which causes further hiring to replace those quits, and so on. Furthermore, high unemployment reduces the number of hires coming from employment, which in turn reduces the number of quits, which reduces the number of hires, and so on. High unemployment should therefore result in a shorter hiring chain and fewer overall hires, holding employment growth constant. The model-based coefficients of hires on employment growth and unemployment work out to 0.996 and -0.420, respectively. The model fits the data rather well in this respect; the differences are small compared with the coefficients of 0.776 and -0.385 found in the data. The model does predict that hires are slightly more procyclical than they actually are, but overall it does a good job at describing the overall cyclical behavior of hires.

It is also possible to derive statistical links which link quits and layoffs to employment growth and unemployment. Substituting (38) back in to (34) gives the link linking quits to employment growth and unemployment:

$$q_t = \frac{\dfrac{\kappa N}{\left(U + \kappa N\right)}\left[\dfrac{\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + (1-\gamma)b}{\Gamma^{LF}}\right]n_t^g - \dfrac{\kappa h}{\left(U + \kappa N\right)^2} u_t}{1 - \dfrac{c}{\omega} \dfrac{\kappa N}{\left(U + \kappa N\right)}}.$$ (39)

The model-based coefficients of quits on employment growth and unemployment come out to 0.379 and -0.627, versus 0.331 and -0.465 in the data. As with hires, the model matches the behavior of quits relatively well. Both quits and hires are highly negatively related to unemployment in a way that makes sense based on the hiring chain concept. During periods of unemployment, a higher share of hires comes from the unemployment pool than from the

employment pool, so quits remain low so long as unemployment remains high. Basically, high unemployment crowds out quits, which results in a reduction in hires, reducing quits by more, and so on.

The linearized condition for layoffs relates layoffs to net employment growth and quits:

$$l_t = \frac{(1-\gamma)(b-m)+\gamma}{\Gamma^{LF}} n_t^g - \left(1 - \frac{c}{\omega}\right) q_t . \tag{40}$$

Layoffs are decreasing in quits because as quits fall, more firms will want to shrink by more than attrition will allow. Combining (40) with (39) yields a link linking layoffs with employment growth and unemployment:

$$l_t = dn_t^g + \frac{\left(1-\dfrac{c}{\omega}\right)\dfrac{\kappa h}{(U+\kappa N)^2}}{1-\dfrac{c}{\omega}\dfrac{\kappa N}{(U+\kappa N)}} u_t , \tag{41}$$

where

$$d = \frac{(1-\gamma)(b-m)+\gamma}{\Gamma^{LF}} - \left(1-\frac{c}{\omega}\right)\frac{\dfrac{\kappa N}{(U+\kappa N)}\left[\dfrac{\mu_0\left(\gamma+\Gamma^{LF}-1\right)+(1-\gamma)b}{\Gamma^{LF}}\right]}{1-\dfrac{c}{\omega}\dfrac{\kappa N}{(U+\kappa N)}} .$$

The coefficient d is less than zero which is perhaps not obvious from inspection; during recessions a large wave of layoffs typically occurs as the distribution of firm growth moves leftward. The coefficient on unemployment is positive. Higher unemployment causes the number of quits to fall, and the lower number of quits increases the number of layoffs through decreased attrition. The coefficients on employment growth and unemployment are -0.384 and 0.206, which are roughly consistent with the coefficients estimated from the data (-0.555 and 0.081). The model explains the behavior of layoffs extremely well, and in particular it explains why layoffs remain relatively elevated during periods of high unemployment.

In the simulated data, regressions of hires, total separations, quits, and layoffs on net worker flows give variance contributions of 80%, 20%, -8%, and 29%, respectively, to the cycle. The model predicts that hires should account for most of the changes in employment, as is true in the data, and it predicts mildly countercyclical separations. The model predicts quits which are not quite as procyclical as they are in reality, while it predicts layoffs which are somewhat less countercyclical. The model matches the qualitative cyclical behavior of the various worker flows relatively well; in particular, it explains why total separations are not highly countercyclical, while layoffs and flows out of employment are highly countercyclical.

In general, it seems that the simple model does a surprisingly good job of explaining the comovement among worker flows, employment growth, and unemployment. Since unemployed job searchers tend to crowd out employed job searchers, unemployment will have an effect on worker flows which it does not have on job flows. Since there is also a link between quits and the other flows which comes from attrition, elevated unemployment will tend to keep hires and quits low and layoffs high though a shortened hiring chain.

5.3 The Beveridge Curve

Figure 6 shows what happens in unemployment-vacancy rate space after a productivity shock. The initial impulse moves the economy far to the southeast in this space; employment growth falls; unemployment rises and the initial crash in hiring requires fewer vacancies to support it. As the economy recovers, it moves back to the northwest in a counterclockwise fashion (as hiring slightly exceeds its long-run average), forming a long thin loop. Contrary to what Akerlof, Rose, and Yellen (1988) say in their opening paragraph, a market-clearing model can yield a realistic Beveridge Curve and realistic worker flows. While the exact shape of the curve is affected by time aggregation, the behavior of vacancies predicted by the model broadly fits their behavior in the data.

Based on the link which links hires with employment growth and unemployment, it is possible to discuss additional link among vacancies, employment growth, and unemployment. Based on the matching function, vacancies are increasing in hiring and decreasing in unemployment. Hiring in turn is increasing in employment growth and decreasing in unemployment. Vacancies must therefore be increasing in employment growth and decreasing in unemployment, with the exact coefficients determined with some algebra.

Rising vacancies during the period of high unemployment which follows a recession should be viewed as a positive indicator of employment growth to come, since they reflect an uptick in hiring activity compared with the contraction phase of a recession.

6. Conclusion

The data on job and worker flows exhibit three sets of strong links. Job flows are tied very closely to employment growth and little else, while worker flows are tied both to the rate of employment growth and the unemployment rate. Within worker flows, there appears to be a particularly strong link among quits, hires, and unemployment. The basic behavior of job flows is a natural outcome of heterogeneity. As the cross-section of desired firm growth shifts, so do aggregate job creation and job destruction. Worker flows are more complicated because of the behavior of quits. The economy exhibits a positive-feedback hiring chain through which hires beget quits, which beget more hires, and so on, which causes quits and hires to vary procyclically and layoffs to vary countercyclically. The model and data agree on the basic behavior of job and worker flows over the cycle; an RBC model without frictions can fit the data surprisingly well.

The simple model has offered some valuable intuition as to the dynamics underlying worker and job flows. A natural next step might be to embed a hiring chain into a large-firm search model with an array of frictions. The technical barriers to integrating search and matching with serious heterogeneity are relatively large. Cooper, Haltiwanger, and Willis (2007) and Fujita and Nakajima (2009) have made some progress on adding realistic frictions into a large-firm search model with heterogeneity, but much more work remains to be done. A more realistic array of frictions might improve the fit between theory and data (particularly by dampening the behavior of job creation and hires), and it would be interesting to see how well the intuition from the simple model survives the addition of such frictions.

Appendix A: An accounting-based variance decomposition

Imagine a vector of time series y_t with a covariance matrix Σ with an accounting identity linking it to an aggregate x_t. The goal is to attribute movements in the aggregate to movements in the original series. The original series y_t is a column vector of length J. Let b equal an accounting identity which links the rows of y to the scalar aggregate x. Then one could write:

$$x_t = by_t. \tag{A1}$$

For each i, regressing y_i on x and then multiplying by b_i gives the coefficient:

$$c_i = \left(\sum_{t=1}^{T} x_t x_t' \right)^{-1} \left(\sum_{t=1}^{T} x_t y_{it}' \right) b_i, \tag{A2}$$

which converges in probability to

$$c_i \xrightarrow{p} \left(E(x_t x_t') \right)^{-1} \left(E(x_t y_{it}') \right) b_i. \tag{A3}$$

Writing out the variance terms, where subscripts under the covariance matrix denote columns of that matrix,

$$c_i \xrightarrow{p} \left(b\Sigma b' \right)^{-1} \left(b\Sigma_i b_i \right). \tag{A4}$$

The right hand side is series i's contribution to the overall variance of the aggregate series (that is, the variance of the aggregate series conditional on series i), while the left hand side is the overall variance of the aggregate series. Furthermore, the c_i coefficients all sum up to one. It is in this sense that the elements of c_i could be thought of as an accounting-based variance decomposition.

Appendix B: Factor demand and output across heterogeneous sectors

This appendix goes through some of the mathematics relating firm-level optimization with aggregate conditions. As a reminder, firm-level output follows the production function:

$$y_{it} = a_{it} k_{it}^{\alpha} (z_t n_{it})^{1-\alpha} .$$

(B1)

Profits depend on prices, output, and factor costs:

$$\Pi_{it} = p_{it} a_{it} k_{it}^{\alpha} (z_t n_{it})^{1-\alpha} - W_t n_{it} - \rho_t k_{it} ;$$

(B2)

and firm-level output is aggregated according to the CES aggregator to produce aggregate output:

$$Y_t = \left[\int_0^1 y_{jt}^{\frac{\theta-1}{\theta}} \, dj \right]^{\frac{\theta}{\theta-1}} .$$

(B3)

Deriving factor shares is very straightforward based on the Cobb-Douglas structure of production. As a result of profit maximization, income shares are constant at the firm level:

$$W_t n_{it} = (1-\alpha) p_{it} y_{it} ,$$

(B4)

and

$$\rho_t k_{it} = \alpha p_{it} y_{it} .$$

(B5)

Summing across i yields the aggregate factor shares, which are the same as the factor shares at the firm level:

$$W_t N_t = (1-\alpha) Y_t ,$$

(B6)

and

$$\rho_t K_t = \alpha Y_t.$$ (B7)

Price-taking behavior by the aggregators yields the isoelastic demand curve:

$$y_{it} = p_{it}^{-\theta} Y_t.$$ (B8)

Firms may enter freely, ensuring that accounting profits are zero and that:

$$p_{it} a_{it} = p_{jt} a_{jt} \equiv l_t,$$ (B9)

for all firms i and j and some constant l_t. Substituting this free entry condition into each firm's demand curve yields:

$$y_{it} = l_t^{-\theta} a_{it}^{\theta} Y_t,$$ (B10)

and substituting this into the output aggregator yields:

$$l_t = \left[\int_0^1 a_{it}^{\theta-1} dj \right]^{\frac{1}{\theta-1}}.$$ (B11)

To get relative factor demand functions, one can divide firm-specific factor demand by economywide factor demand to get:

$$n_{it} = \frac{p_{it} y_{it}}{Y_t} N_t,$$ (B12)

and

$$k_{it} = \frac{p_{it} y_{it}}{Y_t} K_t.$$ (B13)

Substituting the firm-level production function into the aggregator gives the expression:

$$Y_t = \left[\int_0^1 \left(a_{it} k_{it}^\alpha (z_t n_{it})^{1-\alpha} \right)^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}}.$$

(B14)

Substituting relative factor demand back in yields:

$$Y_t = \left[\int_0^1 (p_{it} y_{it} a_{it})^{\frac{\theta-1}{\theta}} dj \right]^{\frac{\theta}{\theta-1}} \frac{K_t^\alpha (z_t N_t)^{1-\alpha}}{Y_t}.$$

(B15)

Substituting the free entry equation (B9) into the integral in (B15) gets rid of the terms involving idiosyncratic prices and productivity, and further cancellation yields the economywide Cobb-Douglas production function:

$$Y_t = l_t K_t^\alpha (z_t N_t)^{1-\alpha},$$

(B16)

where l_t can be thought of as an aggregate productivity term which represents the contribution of the dispersion of idiosyncratic productivity to aggregate productivity:

$$l_t = \left[\int_0^1 a_{it}^{\theta-1} dj \right]^{\frac{1}{\theta-1}}.$$

(B17)

As the dispersion of a increases, for any θ greater than one (a situation where firm-level output is grossly substitutable), aggregate productivity increases. One can think of productivity dispersion as offering an option toward shifting production into higher value-added sectors. An economy with one sector with productivity 1.5 and a sector with productivity 0.5 will be more productive than an economy with two sectors with productivity 1, because firms will tend to shift inputs toward the sector with productivity 1.5. Firm-level labor demand is related to aggregate labor demand and idiosyncratic productivity.

Dividing (B4) by (B6) implies that firm-level labor demand is proportionate to firm-level output:

$$\frac{n_{it}}{N_t} = \frac{p_{it}\, y_{it}}{Y_t}.$$

(B18)

Substituting in the free entry condition (B9) and its aggregate version (B10) give the firm's relative labor demand as a function of its relative productivity:

$$\frac{n_{it}}{N_t} = l_t^{1-\theta} a_{it}^{\theta-1}.$$

(B19)

Appendix C: Linearized indivisible-labor model with job, worker, and vacancy flows

Here are the linearized versions of the model equations:

$$\hat{\lambda}_t + \hat{C}_t = 0 ; \tag{C1}$$

$$\hat{\lambda}_t - \hat{r}_t = E_t \hat{\lambda}_{t+1} ; \tag{C2}$$

$$\hat{r}_t = E_t \frac{\rho}{1 - \delta + \rho} \hat{\rho}_{t+1} ; \tag{C3}$$

$$\hat{\lambda}_t + \hat{W}_t = 0 ; \tag{C4}$$

$$\hat{Y}_t - \alpha \hat{K}_t - (1-\alpha)\hat{z}_t - (1-\alpha)\hat{N}_t = 0 ; \tag{C5}$$

$$\hat{W}_t + \hat{N}_t - \hat{Y}_t = 0 ; \tag{C6}$$

$$\hat{\rho}_t + \hat{K}_t - \hat{Y}_t = 0 ; \tag{C7}$$

$$(1-\delta)\frac{K}{Y}\hat{K}_t + \hat{Y}_t - \frac{C}{Y}\hat{C}_t - \frac{G}{Y}\hat{G}_t = \Gamma^z \Gamma^{LF} \frac{K}{Y}\hat{K}_{t+1}. \tag{C8}$$

Technology follows a loglinear random walk with drift:

$$\hat{z}_t = \hat{z}_{t-1} + \varepsilon_t^z , \tag{C9}$$

and autonomous demand G_t does not move:

$$\hat{G}_t = 0. \tag{C10}$$

Job creation and destruction are both linearized as a proportion of employment:

$$jc_t = \frac{\mu_0 \left(\gamma + \Gamma^{LF} - 1\right) + (1-\gamma)a}{\Gamma^{LF}} \left(\hat{N}_t - \hat{N}_{t-1}\right), \tag{C11}$$

and

$$jd_t = \frac{(1-\gamma)(a-m) + \gamma}{\Gamma^{LF}} \left(\hat{N}_t - \hat{N}_{t-1}\right), \tag{C12}$$

where

$$a = e^{\mu_g + \frac{1}{2}\sigma_g^2} \Phi\left(\frac{\mu_g + \sigma_g^2}{\sigma_g}\right) + \frac{e^{\mu_g + \frac{1}{2}\sigma_g^2}}{\sigma_g} \phi\left(\frac{\mu_g + \sigma_g^2}{\sigma_g}\right) - \frac{1}{\sigma_g} \phi\left(\frac{\mu_g}{\sigma_g}\right).$$

Unemployment is linearized as follows:

$$\hat{U}_t = -\frac{N}{1-N} N_t, \tag{C13}$$

and the matching function is given by:

$$\frac{1}{h} h_t - (1-\zeta)\hat{V}_t = \zeta \frac{U}{U + \kappa N} \hat{U}_{t-1} + \left(\zeta \frac{\kappa N}{U + \kappa N} - 1\right)\hat{N}_{t-1}. \tag{C14}$$

Hiring and separations are both linearized as a proportion of employment:

$$h_t = \left[\mu_0\left(\gamma + \Gamma^{LF} - 1\right) + (1-\gamma)b\right]\left(\hat{N}_t - \hat{N}_{t-1}\right) - (1-\gamma)c\hat{\omega}_t, \tag{C15}$$

and

$$s_t = (1-\gamma)[b-m]\left(\hat{N}_t - \hat{N}_{t-1}\right) - (1-\gamma)c\hat{\omega}_t, \tag{C16}$$

where

$$b = e^{\mu_g + \frac{1}{2}\sigma_g^2}\Phi\left(\frac{\mu_g - \ln(\omega) + \sigma_g^2}{\sigma_g}\right) + \frac{e^{\mu_g + \frac{1}{2}\sigma_g^2}}{\sigma_g}\phi\left(\frac{\mu_g - \ln(\omega) + \sigma_g^2}{\sigma_g}\right) - \frac{\omega}{\sigma_g}\phi\left(\frac{\mu_g - \ln(\omega)}{\sigma_g}\right),$$

and

$$c = \frac{e^{\mu_g + \frac{1}{2}\sigma_g^2}}{\sigma_g}\phi\left(\frac{\mu_g - \ln(\omega) + \sigma_g^2}{\sigma_g}\right) + \omega\Phi\left(\frac{\mu_g - \ln(\omega)}{\sigma_g}\right) - \frac{\omega}{\sigma_g}\phi\left(\frac{\mu_g - \ln(\omega)}{\sigma_g}\right).$$

Quits are related to hires and the ratio of active searchers to passive searchers:

$$(U + \kappa N)q_t - \kappa N h_t = \kappa N\left(\frac{U}{U + \kappa N}\right)h\hat{N}_{t-1} - \left(\frac{\kappa N}{U + \kappa N}\right)hU\hat{U}_{t-1}. \tag{C17}$$

Layoffs are related to separations and quits:

$$l_t = s_t - q_t. \tag{C18}$$

Quits are also related inversely to firms' ability to retain workers:

$$(1 - \gamma)\omega\hat{\omega}_t = -q_t. \tag{C19}$$

To relate the log-linearized model to observables as discussed in the main body of the paper, it is necessary to specify a link between beginning-of-period unemployment in the model and in the data. That link is given as follows:

$$u_t = -N\hat{N}_{t-1} = U\hat{U}_{t-1}. \tag{C20}$$

Expressed as a function of observables, the quit function (C17) becomes:

$$q_t = \frac{\kappa N}{U + \kappa N}h_t - \frac{\kappa h}{(U + \kappa N)^2}u_t. \tag{C21}$$

Within the text, (C19) combined with (C15) yields (37). (C21) combined with (37) yields (38), which when recombined with (C21) yields (39).

The linearized condition for layoffs can be derived from the condition for separations minus quits:

$$l_t = s_t - q_t = \frac{(1-\gamma)(b-m)+\gamma}{\Gamma^{LF}} n_t^g - (1-\gamma)c\hat{\omega}_t - q_t. \tag{C22}$$

When combined with (C19), (C22) yields (40), which combined with (39) yields (41).

References

Abowd, John M., and Arnold Zellner, 1985. "Estimating Gross Labor Force Flows." Journal of Business and Economic Statistics 3, pages 254-283.

Akerlof, George A., Andrew K. Rose, and Janet L. Yellen, 1988. "Job Switching and Job Satisfaction in the U.S. Labor Market." Brookings Papers on Economic Activity 1988(2), pages 495-594.

Barnichon, Regis, 2010. "Building a Composite Help Wanted Index." Economics Letters 109(3), pages 175-178.

Barron, John M., Berger, Mark C., and Black, Dan A., 1997. "Employer Search, Training, and Vacancy Duration." Economic Inquiry 35(1), pages 167-192.

Blanchard, Olivier Jean, and P. Diamond, 1989. "The Beveridge Curve." Brookings Papers on Economic Activity 1989(1), pages 1-60.

Blanchard, Olivier Jean, and P. Diamond, 1990. "The Cyclical Behavior of the Gross Flows of U.S. Workers." Brookings Papers on Economic Activity 1990(2), pages 85-155.

Blanchard, Olivier Jean, and P. Diamond, 1991. "The Aggregate Matching Function." NBER Working Paper 3175.

Bureau of Labor Statistics (BLS), 1978. "Employment and Earnings, United States, 1909-78." BLS Bulletin 1312-11.

Bureau of Labor Statistics (BLS), 1982. "Employment and Earnings, March 1982." Employment and Earnings 29(3).

Carlsson, Mikael, Stefan Eriksson, and Nils Gottfries, 2008. "How Important are Search Frictions for Employment Dynamics?" Working Paper, Uppsala University.

Contini, Bruno, and Ricardo Revelli, 1989. "The Relationship between Firm Growth and Labor Demand." Small Business Economics 1, pages 309-314.

Contini, Bruno, and Ricardo Revelli, 1990. "Job creation and labour mobility: the vacancy chain model and some empirical evidence." in F. Padoa Schioppa (ed.), Squilibri e rigidit nel mercato del lavoro italiano, F. Angeli, Milan.

Cooper, Russell, John Haltiwanger, and Jonathan L. Willis, 2007. "Search frictions: Matching aggregate and establishment observations." Journal of Monetary Economics 54 (Supplement), pages 56-78.

Davis, Steven J., R. Jason Faberman, John Haltiwanger, 2006. "The Flow Approach to Labor Markets: New Data Sources and Micro–Macro Links." Journal of Economic Perspectives 20(3), pages 3-26.

Davis, Steven J., R. Jason Faberman, John Haltiwanger, 2011. "Labor Market Flows in the Cross Section and Over Time." Manuscript.

Davis, Steven J., R. Jason Faberman, John Haltiwanger, Ron Jarmin, and Javier Miranda, 2010. "Business Volatility, Job Destruction, and Unemployment." American Economic Journal: Macroeconomics 2(2), pages 259-287.

Fallick, Bruce, and Fleischman, Charles, A., 2004. "Employer-to-employer flows in the U.S. labor market: the complete picture of gross worker flows." Finance and Economics Discussion Series 2004-34, Board of Governors of the Federal Reserve System.

Fujita, Shigeru, and Makato Nakajima, 2009. "Worker Flows and Job Flows: A Quantitative Investigation." Federal Reserve Bank of Philadelphia Working Paper 09-33.

Hagedorn, Marcus, and Iourii Manovskii, 2008. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited." American Economic Review 98(4), pages 1692-1706.

Haltiwanger, John D., Ron S. Jarmin, and Javier Miranda, 2010. "Who Creates Jobs? Small vs. Large vs. Young." NBER Working Paper 16300.

Krause, Michael J., and Lubik, Thomas, 2010. "On-the-job search and the cyclical dynamics of the labor market." Federal Reserve Bank of Richmond Working Paper 10-12.

Lotti, Francesca, Enrico Santarelli, and Marco Vivarelli, 2003. "Does Gibrat's Law hold among young, small firms?" Journal of Evolutionary Economics 13, pages 213-235.

Lotti, Francesca, Enrico Santarelli, and Marco Vivarelli, 2009. "Defending Gibrat's Law as a Long-Run Regularity." Small Business Economics 32(1), pages 31-44.

Mortensen, Dale T., Christopher A. Pissarides, 1994. "Job Creation and Job Destruction in the Theory of Unemployment." Review of Economic Studies 61(3), pages 397-415.

Pissarides, Christopher A., 2000. *Equilibrium Unemployment Theory*. Cambridge, MA: MIT Press.

Rogers, Mark, Christian Helmers, and Christoffer Koch, 2010. "Firm growth and firm size." Applied Economics Letters 2010, pages 1-4.

Sutton, John, 1997. "Gibrat's Legacy." Journal of Economic Literature 35, pages 40-59.

**Table 1: Percent variance contributions of job and worker flow rates
to net job and worker flows based on OLS regressions**

Source	BDS	BDM	JOLTS	LTS	CPS	Model
Period	1977-2009	1992.III-2010.III	2001.I-2011.I	1947.I-1981.IV	1990.II-2011.I	
Job creation	39.3%	88.0%				58.5%
Expanding establishments	31.2%	66.7%				
Entering establishments	8.1%	21.3%				
Job destruction	60.8%	61.1%				41.5%
Contracting establishments	50.6%	53.1%				
Exiting establishments	10.2%	8.0%				
Hires / accessions			97.2%	60.4%		79.8%
Separations			2.7%	39.6%		20.2%
Quits and other separations			-56.9%	-27.9%		-8.5%
Layoffs and discharges			59.6%	67.5%		28.7%
Employment inflows					43.2%	71.3%
From unemployment					12.5%	
From not in the labor force					30.7%	
Employment outflows					56.9%	28.7%
To unemployment					39.6%	
To not in the labor force					17.3%	

Source: Census Bureau, BLS, and author's calculations. The monthly JOLTS, LTS, and CPS series are aggregated to a quarterly frequency to mitigate the role of measurement error and temporary blips in the data. I omit the reconciling item "OE flows" from the CPS calculations since they consist primarily of adjustments to the BLS's census-based population controls and have little economic meaning. The LTS data cover manufacturing on an SIC basis only; no comparable seasonally adjusted JOLTS data currently exist for the modern period due to concerns by the BLS about the spiky nature of the data. For the LTS, I use accessions as my hiring measure. Quarterly data are HP-filtered with a smoothing parameter of 10,000 and annual data with a smoothing parameter of 100.

**Table 2: Linear projections of job and worker flows on net job growth
and the unemployment rate**

Estimate	R^2	Net	U rate
Job creation			
Expanding establishments	0.823	0.340 (0.020)	-0.035 (0.013)
Entering establishments	0.058	0.033 (0.009)	-0.006 (0.009)
Job destruction			
Contracting establishments	0.902	-0.548 (0.023)	-0.039 (0.013)
Exiting establishments	0.244	-0.080 (0.011)	-0.002 (0.009)
Hires	0.865	0.776 (0.119)	-0.385 (0.043)
Separations			
Quits and other separations	0.865	0.331 (0.069)	-0.465 (0.042)
Layoffs and discharges	0.730	-0.555 (0.089)	0.081 (0.024)

Source: BLS and author's calculations. The monthly JOLTS series are aggregated to a quarterly frequency to mitigate the role of measurement error and temporary blips in the data. Quarterly data are HP-filtered with a smoothing parameter of 10,000. The system was estimated using ordinary least squares, with Newey-West autocorrelation-consistent standard errors based on a lag length of one. Standard errors are listed in parentheses. The R-squared values for the job creation and destruction regressions restricting the coefficient on unemployment to equal zero are 0.800, 0.054, 0.887, and 0.243, respectively.

Table 3: Contemporaneous correlations of job and worker flow rates

| Object	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) JC, Expanding estabs.	1						
(2) JC, Entering estabs.	0.33	1					
(3) JD, Contracting estabs.	-0.72	-0.05	1				
(4) JD, Exiting estabs.	-0.26	0.46	0.53	1			
(5) Hires	0.67	0.19	-0.56	-0.49	1		
(6) Quits and other seps.	0.50	0.20	-0.31	-0.33	**0.91**	1	
(7) Layoffs and discharges	-0.65	-0.18	0.71	0.59	-0.63	-0.62	1

Source: BLS and author's calculations. The monthly JOLTS series are aggregated to a quarterly frequency to mitigate the role of measurement error and temporary blips in the data. Data are detrended by an HP filter using a smoothing parameter of 10,000.

Table 4: First moments of job and worker flow rates as a share of employment

| Source	BDS	BDM	JOLTS	LTS	CPS
Period	1977-2009	1992.III-2010.III	2001.I-2011.I	1947.I-1981.IV	1990.II-2011.I
Job creation	17.2%	6.7%			
Expanding establishments	10.8%	5.3%			
Entering establishments	6.4%	1.3%			
Job destruction	15.2%	6.8%			
Contracting establishments	9.9%	5.5%			
Exiting establishments	5.4%	1.3%			
Hires / accessions			10.9%	12.9%	
Separations			10.9%	13.1%	
Quits and other separations			6.8%	8.3%	
Layoffs and discharges			4.1%	4.9%	
Employment inflows					12.4%
From unemployment					4.7%
From not in the labor force					7.7%
Employment outflows					12.4%
To unemployment					4.3%
To not in the labor force					8.1%

Source: Census Bureau, BLS, and author's calculations. The monthly JOLTS, LTS, and CPS series are aggregated to a quarterly frequency to mitigate the role of measurement error and temporary blips in the data. I omit the reconciling item "OE flows" from the CPS calculations since they consist primarily of adjustments to the BLS's census-based population controls and have little economic meaning. The LTS data cover manufacturing on an SIC basis only; no comparable seasonally adjusted JOLTS data exist for the modern period due to concerns by the BLS about the spiky nature of the data. For the LTS, I use accessions as my hiring measure since that measure is available throughout the whole sample going back to 1930.

Figure 1: Annual March-March BDS job flows as share of employment, 1977-2009

Source: Census Bureau Business Dynamics Statistics (BDS) and author's calculations. All series are measured as a percent of nonfarm employment over the intervening year.

Figure 2: Beveridge Curve by employment cycle

Source: Help wanted data from the Conference Board, Barnichon (2010), nonfarm employment and vacancies from the BLS, and the author's calculations.

Figure 3: Job creation and destruction in the cross section

New job destruction region
(firm growth < 0)

Old job destruction region
(firm growth < 0)

Job creation region
(firm growth > 0)

-30% -20% -10% 0% 10% 20% 30%

Firm growth rate (%)

Old job destruction region New job destruction region
Old distribution of growth rates New distribution of growth rates

This chart shows what happens to job flows when the entire distribution of firm-level growth shifts rightward by 2%. Job creation rises and job destruction falls.

Figure 4: Job destruction versus layoffs in the cross section

Job destruction region
(firm growth < 0)

Layoff region
(firm growth < $1-\omega_{it}$)

Job creation region
(firm growth > 0)

-30% -20% -10% 0% 10% 20% 30%

Firm growth rate (%)

Job destruction region Layoff region Distribution of growth rates

This chart shows the difference between layoffs and job destruction—layoffs only hit workers in rapidly contracting firms.

Figure 5: Impulse response to a -1% productivity shock, model
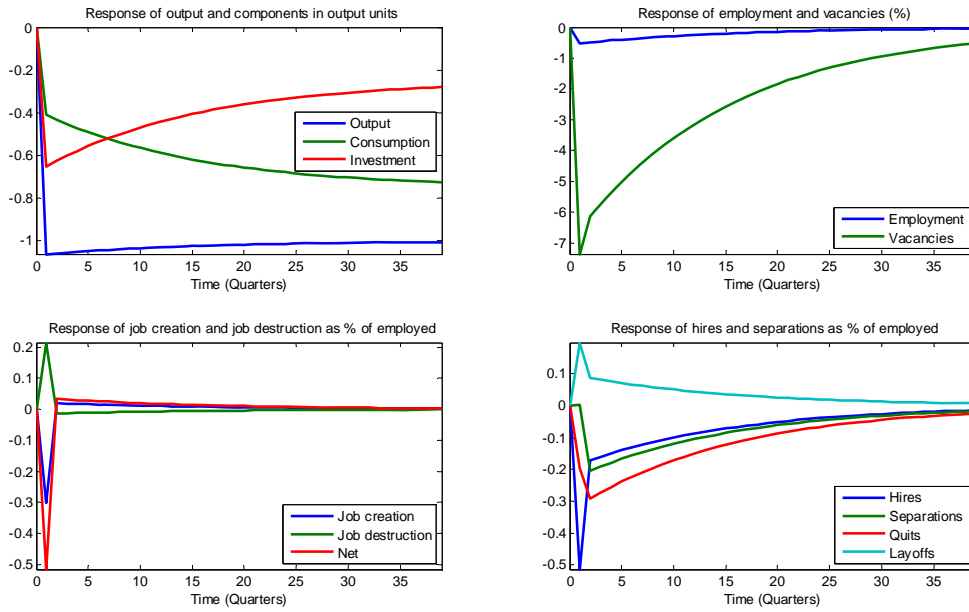
Source: Author's calculations using linearized model.

Figure 6: Response of vacancies and unemployment to a -1% productivity shock, model
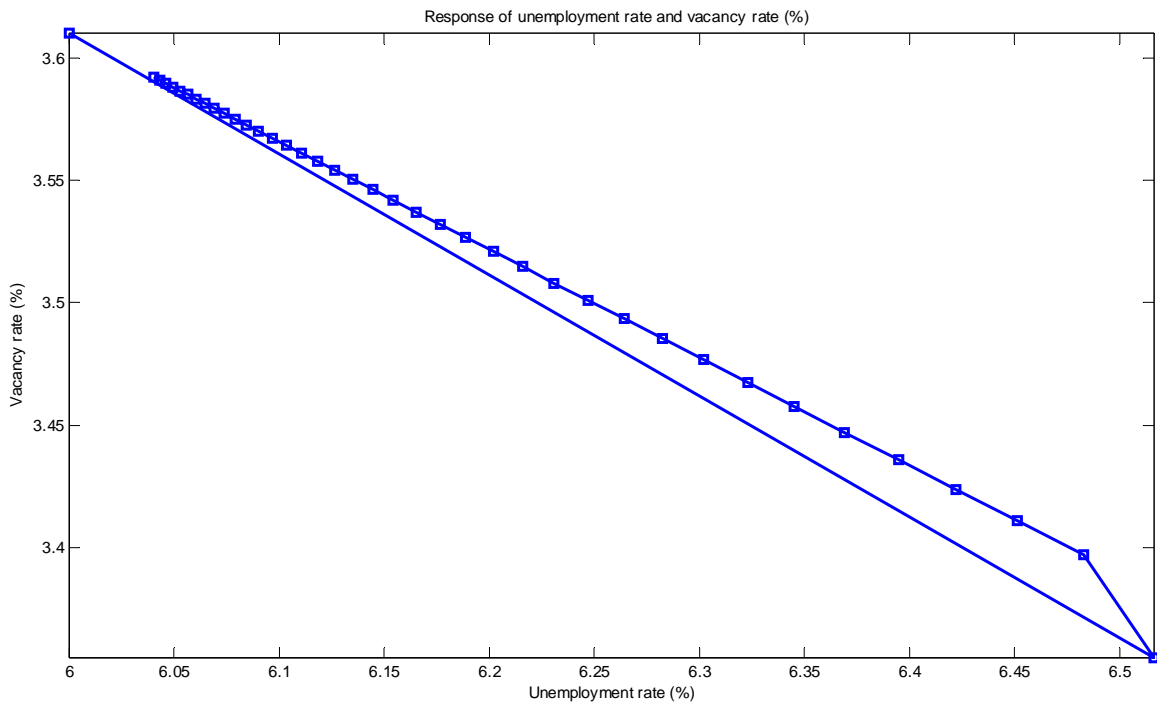
Source: Author's calculations using linearized model. The dynamics move in a counterclockwise direction.