



Working Paper Series

**Aggregate Employment Fluctuations with
Microeconomic Asymmetries**

Jeffrey R. Campbell and Jonas D.M. Fisher

Working Papers Series
Research Department
(WP-96-17)

Federal Reserve Bank of Chicago

Aggregate Employment Fluctuations with Microeconomic Asymmetries

Jeffrey R. Campbell* Jonas D.M. Fisher†

March 5, 1998‡

Abstract

We provide a simple explanation for the observation that the variance of job destruction is greater than the variance of job creation. In our model profit maximization in the presence of proportional plant-level costs of job creation and destruction implies that shrinking plants are more sensitive than growing plants to aggregate shocks. We describe circumstances in which this microeconomic asymmetry is preserved in the aggregate and show that it can account for asymmetries in the variability of job creation and destruction of the kind observed in the U.S. manufacturing sector. This is so even though we abstract from job search and matching frictions, incomplete contracts, and aggregate congestion effects, all of which have been put forward as important for understanding the job creation and destruction evidence.

* University of Rochester and NBER

† Federal Reserve Bank of Chicago

‡ We thank Hal Cole, Tom Cooley, Jeremy Greenwood, Hugo Hopenhayn and Vincenzo Quadrini for helpful conversations, and we are grateful to Ricardo Caballero for his valuable discussion of an earlier version of this paper. We also thank a co-editor and two anonymous referees for their very helpful comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve System or the Federal Reserve Bank of Chicago.

1. Introduction

Consider the problem of choosing employment at a plant facing job creation and destruction costs which are proportional to the number of jobs created or destroyed, as in Bentolila and Bertola (1992), Bertola and Rogerson (1996) and Hopenhayn and Rogerson (1993). When there is a temporary change in the real cost of a job, which includes real wages and the cost of complementary material inputs, the plant's response depends on whether it is creating or destroying jobs. For an expanding plant, the total cost of the last job created is greater than the relevant factor payments because of the expected adjustment costs incurred. These additional costs, which include the job creation cost and the expected cost of destroying that job in the future, lower the elasticity of the total cost of job creation with respect to real factor prices below unity. For a shrinking plant, the total cost of the last job retained is less than the relevant factor payments because by retaining the job the plant avoids the cost of job destruction and the expected cost of recreating that same job in the future. The subtraction of avoided adjustment costs increases the elasticity of the total cost of job retention with respect to real factor prices above unity. Since expanding and contracting plants operate on the margins which equate the costs and benefits of adding and retaining a job, respectively, the asymmetric responses of the total costs of job creation and retention to factor price changes can induce contracting plants to respond by more than expanding plants to fluctuations in the real cost of a job.¹

We apply this insight to study the behavior of aggregate job creation and destruction rates, such as those measured by Davis, Haltiwanger, and Schuh (1996). Their observation that job destruction is more volatile than job creation in the US manufacturing sector has challenged business cycle theory, because standard models based on the one-sector stochastic growth model predict equally volatile job creation and destruction. To show how the addition of adjustment costs to such a model can reconcile business cycle theory with the evidence, we construct a simple model economy with a fixed population of plants which face proportional costs of job creation and destruction. The plants produce a homogeneous good sold in a competitive goods market and purchase inputs in competitive factor markets. Employment varies at plants due to idiosyncratic technology shocks and aggregate shocks which disturb the price of the good or the cost of factor

¹The nondifferentiability of the adjustment costs at the point of zero change implied by proportional adjustment costs is crucial to this result. With strictly convex and everywhere differentiable adjustment costs the job creation and retention margins are identical and so must behave identically in response aggregate shocks.

inputs which are complementary with labor, such as oil. As described above, the adjustment costs induce employment decisions at growing plants to respond by less to aggregate shocks than at shrinking plants. However, as Caballero (1992) has made clear, drawing conclusions about aggregate variables from individual decisions is fraught with ambiguity. To address this, we simulate aggregate data from our model and compare it to empirical evidence on gross job flows. We show how the microeconomic asymmetries can be preserved in the aggregate so that the model reproduces asymmetries in gross job flows of the kind observed in the US manufacturing sector.

The job creation and destruction costs which are central to our analysis are easily motivated without appealing to contracting or search frictions. In our model, the plant pays no adjustment costs if the identities of its employees change but the number of employees remains constant. Therefore, the adjustment costs we study are *net adjustment costs*, as defined by Hamermesh and Pfann (1996). Following their interpretation, we view these costs as arising from the lost output incurred when reorganizing a plant's production process to operate at a larger or smaller scale.² As we show below, the optimal employment policy of a plant facing proportional net adjustment costs involves frequently leaving employment unchanged between periods. Hence, our adjustment cost specification is consistent with Hamermesh's (1989) observation that manufacturing plants adjust their employment only infrequently. Furthermore, using Dutch data, Hamermesh, Hassink, and van Ours (1994) found that many firms kept the total number of jobs constant over a two year period but no firm kept the identity of its employees constant. Firms which did not change their employment replaced about 11% of their work force over these two years.³ These observations suggest the presence of net adjustment costs which do not depend on changing workers' identities. Finally, Lazear (1995) reports that "Human resource managers think in terms of slots or jobs, and think of these slots or jobs as being fundamental to the

²Hamermesh and Pfann (1996), p. 1266, state "For workers (treated now as an input that is homogeneous once it has been trained and whose hours are fixed) the net costs are those of changing the numbers of employees in the firm. These costs include disruptions to production occurring when changing employment causes workers' assignments to be rearranged (implicitly assuming no change in the capital stock) and all other costs that are not related to the identity of the workers but instead depend solely on changing the number of employees ..." Hopenhayn and Rogerson (1993) and Bertola and Rogerson (1996) adopt a specification for job destruction costs that is similar to ours, but they interpret them as costs of firing workers. This interpretation is correct only if the turnover of workers at a plant follows a last in first out policy, so that job destruction is identical to worker separations.

³See the second row of Table 3 in Hamermesh, Hassink, and van Ours (1994).

organization of the firm (p. 77).” This provides further, although indirect, evidence that the costs of creating and destroying jobs are significant.

The optimal employment policy we study can be represented as a function of the gap between a plant’s actual employment and the level it would choose in the absence of adjustment costs. Hence our model is closely related to standard (S, s) models of employment dynamics. As in these models employment is only changed if the employment gap exceeds a destruction threshold or falls below a creation threshold. In these cases employment is adjusted so that the employment gap equals an associated return point, which in our model equals the threshold level crossed. For the reasons described above, aggregate shocks induce asymmetric variation in the job creation and destruction thresholds. The fact that the adjustment thresholds vary asymmetrically with aggregate shocks explains why our results differ from standard (S, s) models. In these models the adjustment thresholds and return points are not disturbed by aggregate shocks, and creation and destruction tend to be equally variable, unless aggregate shocks are asymmetrically distributed, as in Caballero (1992), or there is a trend in the employment gap, as in Foote (1997).⁴

Foote’s (1997) study of employment trends is motivated by his finding, based on unemployment insurance data from Michigan, that job creation is actually more volatile than job destruction in sectors which have expanded over the long run, such as services and trade. In his model there are a fixed number of establishments, so a trend in industry employment is equivalent to a trend in average establishment size. For reasons Foote discusses in detail, a positive trend in average establishment size causes job creation to be more volatile than job destruction in his model, as it is in most non-manufacturing industries. A negative trend, as in the manufacturing sector, implies the opposite. Therefore, his model appears to explain much of the behavior of gross job flows across a wide variety of sectors. However, this explanation is inconsistent with the observed importance of entry in most non-manufacturing industries. In the service, retail trade, and wholesale trade sectors of the Michigan economy, average establishment size hardly changed during the sample period of Foote’s data, although employment in all three sectors grew considerably.⁵ Without a trend in average establishment size, the bunching

⁴The case of symmetrically varying thresholds and return points is indistinguishable from the standard case since it is equivalent to a common aggregate shock which changes the mean of the distribution of plants with respect to the employment gap.

⁵Using the County Business Patterns database, we calculated the average growth rate from 1978 to 1988, Foote’s sample period, of average establishment size for all non-farm business one digit industries in the Michigan economy. Average establishment size grew 0.18% per year in services, while the number of establishments grew

of plants near their employment thresholds, which his model needs to account for the data, could not have occurred. We view our model with a fixed number of producers as appropriate for the study of gross job flows in the manufacturing sector where entry has played a relatively unimportant role, but inappropriate for other sectors where entry has been more important.⁶ These considerations motivate our focus on the manufacturing sector.

The analysis of (S, s) models constitutes one of two broad strategies for understanding the creation and destruction evidence which have been followed in the literature. The other emphasizes the creation and destruction of productive establishments due to technological change and various frictions which hinder this process, such as search and matching, congestion effects in plant creation, and incomplete contracts. Prominent examples of research along these lines are Mortensen and Pissarides (1994) and Caballero and Hammour (1994, 1996). Our analysis differs from these papers along two key dimensions. First, the nature of production at the plant level is different. In the papers just cited production occurs at plants of constant size which combine labor and capital in fixed proportions. Over time plants are created and destroyed because new plants tend to be more productive than older plants and because aggregate shocks affect the profitability of plants. Since employment cannot be varied at individual plants, all variation in job creation and destruction arises from the creation and destruction of plants. Under this view cyclical variation in job creation and destruction is tightly linked to technological advance and obsolescence and disconnected from incentives to vary the utilization of existing capacity. In contrast, our model consists of a fixed number of plants, and variation in gross job flows arises from variation in employment at individual plants only. This seems to be a sensible starting point for taking into account the fact documented by Davis, Haltiwanger, and Schuh (1996) that by far most of the job creation and destruction in the manufacturing sector is due to employment variation at incumbent plants rather than at entering and exiting plants.⁷ Abstracting from plant entry and exit also simplifies our analysis considerably and allows us to focus on the distinct asymmetric plant-level employment adjustment we describe.

3.94% per year. In wholesale trade, these numbers are -0.08% and 1.85% , and in retail trade they are 0.52% and 1.44% .

⁶Caballero and Hammour's (1994) model, where entry accommodates all demand fluctuations so that job creation is more volatile than job destruction, seems appropriate for sectors where net entry has been important.

⁷For example, over the period 1972:II-1988:IV, only 8.4% of total job creation is accounted for by startup plants, while just 11.6% of total job destruction is accounted for by plant shutdowns. See Davis, Haltiwanger and Schuh (1996, figure 2.3, p. 29).

The second key difference is that there are no frictions to hinder the process of job creation and destruction in our model. In the Caballero and Hammour papers a key role is played by congestion in the process of setting up new plants which makes the cost of plant creation at a point in time increasing in the number of currently entering plants.⁸ Under the right conditions, this assumption can lead to the dampening of plant and job creation relative to plant and job destruction. For example, in Caballero and Hammour (1996) this assumption must be combined with job specific physical capital and incomplete contracts between workers and plant owners. Since the number of plants is fixed in our model, there is no role for these congestion effects. In Mortensen and Pissarides (1994) and Caballero and Hammour (1996), search and matching are important. In our model, plants do not have to post a vacancy and wait for a suitable match in order to start production, and workers do not need to wait in an unemployment pool until a suitable match arrives. Therefore, there is no feedback from individual creation and destruction decisions to the rate at which job vacancies are filled or to the wages paid to workers in newly created matches. This feedback is crucial to the dynamics of job creation and destruction in the search and matching models. Instead, in our model plants can hire instantaneously as many workers as they choose to at the prevailing wage.

The remainder of the paper proceeds as follows. The next section contains a simple example which illustrates how optimally chosen job creation and destruction can respond asymmetrically to aggregate shocks. Section 3 presents the model economy and relates it to other models which have been used to study job creation and destruction. In Section 4 we use a calibrated version of our model to investigate its implications for fluctuations in aggregate job creation and destruction rates. In Section 5 we conclude by discussing some implications of our findings for the study of business cycles and directions for future research.

2. An Illustrative Example

To build intuition for our findings, we study a simple example adapted from the steady state model of Bertola and Rogerson (1996). It consists of a perfectly competitive industry populated by a continuum of plants which experience idiosyncratic productivity shocks and must pay job creation and destruction costs whenever they expand or contract employment. We demonstrate

⁸These congestion effects, while external to a plant, are internalized by a firm operating many plants.

the importance of asymmetries in plant-level creation and destruction for the responses of gross job flows by studying the transitional dynamics between steady states induced by a permanent real factor price increase. The analysis is equivalent to one in which a factor price is the numeraire and the real price of the industry's output declines permanently. Due to the adjustment costs, the optimal employment policies of job creating and destroying plants along the transition path are asymmetric. This microeconomic asymmetry causes the industry to adjust employment primarily through job destruction.

2.1. Steady State

Consider a representative plant that produces a homogeneous good for sale in a competitive goods market and purchases inputs in competitive factor markets. The plant uses two variable factors of production, labor, which comes in fixed shift lengths, and materials. These inputs must be combined in fixed proportions. Materials are included in the model to highlight the fact that changes in the cost of a job can reflect changes in the prices of inputs which complement labor as well as changes in the direct price of labor. We refer to an employee plus her accompanying unit of materials as a job. The per-period cost of a job, measured in units of the constant output price, is denoted by W . Although W reflects both labor and materials costs, we refer to it simply as the wage. We initially assume that W is constant.

Let n_t denote the plant's employment at time t and let z_t denote the plant's idiosyncratic productivity level. The latter follows a Markov chain over the two states $z_g > z_b > 0$, the 'good' and 'bad' states, respectively. The probability of z_t changing states is constant and equals p . Output is given by $z_t n_t^\alpha$, where $\alpha \in (0, 1)$. The strict concavity of the production function reflects the presence of a fixed factor at the plant or a limit to a manager's span of control, as in Lucas (1978). When the plant changes the scale of its employment, it incurs an adjustment cost. If it expands employment, it must pay the job creation cost τ_c for every job added. If it contracts employment, it must pay the job destruction cost τ_d for every job lost.⁹ Finally, the plant's manager is risk neutral, discounts future profits with the constant discount factor $0 < \beta < 1$, and maximizes the expected present discounted value of profits.

⁹To avoid the possibility that a plant manager would choose to hoard either materials or labor in excess of the other, we assume that the job creation and destruction costs are incurred whenever the *minimum* of labor or materials changes. This is consistent with our interpretation of these costs as reflecting the reorganization of a plant to accommodate a larger or smaller production scale.

With these assumptions, it is straightforward to show that there exists an ergodic set for the plant's employment which has exactly two points. To do so, suppose that $n_t = n_g$ whenever $z_t = z_g$ and $n_t = n_b$ whenever $z_t = z_b$. Then the expected present value of the marginal product of employment minus the wage follows a two state Markov chain, taking on the values v_g and v_b in the good and bad productivity states, respectively. It follows that v_g and v_b must satisfy

$$v_g = \alpha z_g n_g^{\alpha-1} - W + \beta((1-p)v_g + pv_b) \quad (2.1)$$

$$v_b = \alpha z_b n_b^{\alpha-1} - W + \beta((1-p)v_b + pv_g). \quad (2.2)$$

If $n_g > n_b$, so that over time the plant both creates and destroys jobs, a necessary condition for the optimality of this employment policy is that $v_g = \tau_c$ and $v_b = -\tau_d$. That is, the marginal net revenue of creating or destroying a job must equal the associated marginal cost. Substituting these conditions into (2.1) and (2.2) yields

$$\alpha z_g n_g^{\alpha-1} = W + (1 - \beta(1-p))\tau_c + \beta p \tau_d \quad (2.3)$$

$$\alpha z_b n_b^{\alpha-1} = W - (1 - \beta(1-p))\tau_d - \beta p \tau_c. \quad (2.4)$$

These conditions have a simple interpretation. First consider (2.3). Suppose the plant's productivity switches from bad to good, and the plant considers increasing employment to n . The current net marginal benefit from this decision is $\alpha z_g n^{\alpha-1} - W - \tau_c$, the marginal product of the last job created less the wage and creation costs of this last job. The decision to create jobs in the current period also impacts future profits because of the adjustment costs. If productivity is good in the next period, the plant saves on the costs of creating the last job added in the current period. The discounted expected value of this saving is $\beta(1-p)\tau_c$. If productivity is bad in the next period, the decision to add the last job in the current period involves higher job destruction than otherwise. The discounted expected value of this cost is $\beta p \tau_d$. Equation (2.3) implies that n_g is the level of employment which exactly equates the benefits and costs of adding an additional job when productivity is good. Similarly, (2.4) implies that n_b is the level of employment which exactly equates the benefits and costs of *retaining* an additional job when productivity is bad.

Equations (2.3) and (2.4) can be solved easily for n_g and n_b , and it can be verified that

both n_g and n_b are decreasing in W . We proceed under the assumption that n_g and n_b are well defined and that n_g is strictly greater than n_b . The plant destroys jobs whenever productivity changes from z_g to z_b . It creates jobs when the opposite change occurs, and it neither creates nor destroys jobs if its productivity stays the same. In a steady state, half of the plants will be in the good state and the other half will be in the bad state. Every period, p plants find that their productivity has changed. The absolute value of job growth at such a plant is $\ln(n_g/n_b)$. Therefore, the steady state aggregate rates of job creation and destruction, POS^* and NEG^* , respectively, both equal $(p/2) \ln(n_g/n_b)$.¹⁰

2.2. Transitional Dynamics

To see how adjustment costs impact the volatility of gross job flows, we consider the industry's transition from an initial steady state to a new steady state with a slightly higher wage. Suppose that the real wage, which reflects the costs of both labor and materials due to our assumption of fixed proportions, increases unexpectedly and permanently from W to W' . Because it is a real price, this change in W can reflect both demand and supply shocks. A demand shock which decreases the nominal price of the output good will raise W if the nominal prices of labor and other factors are fixed by long-term contracts or pinned down by their values in other sectors. A supply shock which raises the price of an input which is complementary to labor, such as energy, will also raise W .

To simplify the calculation of a plant's optimal response to this aggregate shock, we assume that it must pay a small fixed cost whenever it changes its employment. This cost is small enough to ensure that the employment policy given by n_g and n_b is still optimal in the steady state and it is large enough so that no plant wishes to change its employment after the wage change unless its productivity has also changed since then.¹¹ Given this auxiliary assumption, the optimal

¹⁰Here, to simplify the example, we approximate the job creation and destruction rates as defined by Davis and Haltiwanger (1990) with the cross sectional averages of $\ln(n_t/n_{t-1}) \times I(n_t > n_{t-1})$ and $-\ln(n_t/n_{t-1}) \times I(n_t < n_{t-1})$, respectively, where $I(\cdot)$ is an indicator function which equals one if its argument is true and equals zero otherwise.

¹¹To be precise, the fixed cost, $f > 0$, must satisfy the following conditions: (i) $v_i - f \geq \tilde{v}_i$, where \tilde{v}_i is the value of delaying by one period the change in employment dictated by the $f = 0$ policy after the productivity state has just switched to i , $i = g, b$; (ii) $v'_i - f \geq \tilde{v}'_i$, where v'_i denotes the value function associated with W' and \tilde{v}'_i is defined analogously to \tilde{v}_i , $i = g, b$; (iii) $\hat{v}'_i \geq v'_i - f$, where \hat{v}'_i denotes the value of delaying by one period the switch to n'_i from n_i after the wage change, $i = g, b$; (iv) $v'_i - f \geq \bar{v}'_i$, where \bar{v}'_i denotes the value of delaying by one period the switch from n_j to n'_i after the wage change and after the productivity state has just switched from j to i , $j \neq i$ and $i, j = g, b$. For the numerical example which we study below to produce Figure 1, the set

employment policy following the wage change is as follows. Let n'_g and n'_b be solutions to (2.3) and (2.4) when W' replaces W . Employment is left at its level before the wage change until the plant experiences a productivity change. Thereafter, $n_t = n'_b$ whenever $z_t = z_b$ and $n_t = n'_g$ whenever $z_t = z_g$.

Because n_g and n_b are decreasing in W , we know that $n'_g < n_g$ and $n'_b < n_b$. To determine which responds by more, we compute the elasticities of n_b and n_g with respect to W using (2.3) and (2.4):

$$\partial \ln n_g / \partial \ln W = \frac{-1}{1 - \alpha} \frac{W}{W + (1 - \beta(1 - p))\tau_c + \beta p \tau_d} \quad (2.5)$$

$$\partial \ln n_b / \partial \ln W = \frac{-1}{1 - \alpha} \frac{W}{W - (1 - \beta(1 - p))\tau_d - \beta p \tau_c}. \quad (2.6)$$

Clearly, $\partial \ln n_b / \partial \ln W$ is greater in absolute value than $\partial \ln n_g / \partial \ln W$, so the policies of shrinking plants respond more than the policies of growing plants to the increase in real factor prices.

The asymmetric responses of n_g and n_b to the wage change cause aggregate job creation and destruction to respond asymmetrically. In the period of the wage change, $p/2$ plants change employment from n_b to n'_g and $p/2$ plants change employment from n_g to n'_b . Hence, the job creation and destruction rates in the period of the wage change are

$$POS_0 = (p/2) \ln(n'_g/n_b)$$

$$NEG_0 = (p/2) \ln(n_g/n'_b).$$

Subtracting the common steady state value of POS^* and NEG^* from both sides of these expressions yields

$$POS_0 - POS^* = (p/2) \ln(n'_g/n_g)$$

$$NEG_0 - NEG^* = (p/2) \ln(n_b/n'_b).$$

These expressions indicate that job creation falls and job destruction rises relative to their previous steady state values. Because $\partial \ln n_b / \partial \ln W$ is greater in absolute value than $\partial \ln n_g / \partial \ln W$, we can conclude that

of f which satisfies all of these constraints is non-empty and equals [0.0113, 0.0457].

$$|NEG_0 - NEG^*| > |POS_0 - POS^*|.$$

That is, job destruction rises more than job creation falls following the increase in input prices. It is important to notice that this implication does not depend on the relative magnitude of the adjustment costs. In particular, it holds if one of the costs is zero while the other is strictly positive, and it holds for any strictly positive pair of values for the costs.

The intuition for the asymmetric responses builds on the intuition underlying (2.3) and (2.4). The current and future adjustment costs associated with creating a job must be added to the wage to determine the marginal cost of job creation, and this causes the marginal cost to respond by less to changes in the wage than it would in the absence of adjustment costs. Hence a smaller change in n_g is required to maintain equality of the marginal costs and benefits of job creation, compared to the no adjustment cost case. In contrast, the current and future adjustment costs associated with destroying a job must be *subtracted* from the wage to determine the marginal cost of job retention. This means a larger change in n_b is required to maintain job destruction efficiency compared to the no adjustment cost case. Since n_g and n_b respond by the same amount in the no adjustment cost case, it follows that job destruction must respond by more than job creation in the presence of adjustment costs.

The asymmetry in the gross job flows' responses to the wage change persists beyond the date of the wage change. Let $P_t = (1 - p)^t$ denote the fraction of plants which have not experienced a productivity change since the wage shock. The job creation and destruction rates t periods after the wage change are

$$\begin{aligned} POS_t &= (p/2) \left(P_t \ln(n'_g/n_b) + (1 - P_t) \ln(n'_g/n'_b) \right) \text{ and} \\ NEG_t &= -(p/2) \left(P_t \ln(n'_b/n_g) + (1 - P_t) \ln(n'_b/n'_g) \right). \end{aligned}$$

Therefore, we can conclude that

$$|NEG_t - NEG^*| > |POS_t - POS^*|, \quad t > 0.$$

Again, this finding does not depend on the relative magnitude of the adjustment costs.

Figure 2.1: $POS_t - POS^*$ and $NEG_t - NEG^*$ on the Transition Between Steady States




Figure 2.1 illustrates the dynamic nature of the gross flow asymmetry. This Figure displays $POS_t - POS^*$ and $NEG_t - NEG^*$ along the transition path following a 1% increase in the wage under the assumption that $\beta = 1.05^{-1/4}$, $\alpha = 2/3$, $p = 1/4$, $\tau_c = \tau_d = 1/2$, $z_g = \exp(0.3)$ and $z_b = \exp(-0.3)$. In the period of the wage change, the job creation rate falls by 0.3 percentage points and the job destruction rate rises by 0.5 percentage points. Thereafter, job creation rises and job destruction descends towards their common, higher level in the new steady state.

This example demonstrates that the presence of proportional job creation and destruction costs can cause the aggregate job creation and destruction rates to respond asymmetrically to a permanent increase in the cost of a job. It is easy to verify that the same asymmetries arise with a permanent wage decrease. However, close inspection of the example reveals an interesting qualification to the basic result: the magnitude of the asymmetry in the responses of n_g and n_b decreases with p , so that the responses of job creating and destroying plants to the aggregate shock become more similar if the idiosyncratic shocks are more persistent. Therefore, the responses of aggregate job creation and destruction become more similar as idiosyncratic shocks become more persistent.

The intuition for this result comes from inspection of the first order conditions which define n_g and n_b , equations (2.3) and (2.4). For an expanding plant, one benefit of creating the marginal job is the future job creation cost avoided if productivity remains good. This benefit partially

offsets the current job creation cost. One cost of creating the marginal job is the job destruction cost incurred if productivity becomes bad. This adds to the total cost of job creation. When productivity is persistent, the offsetting benefit is very large and the additional cost is very small. Because the total adjustment costs incurred for the marginal job are smaller when p is small, n_g responds by more to changes of input prices. The analysis of the job destruction decision is similar. The total adjustment costs avoided from retaining the marginal job are small when productivity is persistent, so n_b responds less to changes of input prices. Because the responses of n_g and n_b to a change in W become more similar when p is smaller, the responses of aggregate job creation and destruction also become more similar.

3. The Model

The example in the previous section illustrates how proportional job creation and destruction costs can give rise to asymmetric responses to aggregate shocks by job creating and destroying plants. To further assess this mechanism for understanding gross job flows, we modify the example in three ways. First, ongoing aggregate fluctuations are added. We assume that the wage, W_t , follows a Markov chain over the grid $W^h > W^l > 0$ with transition matrix Π . Note that, as in the example, changes in the wage can reflect both demand and supply shocks. Second, we assume the distribution of idiosyncratic productivity shocks is continuous rather than discrete. Specifically, $\ln z_t$ follows a random walk with continuously distributed innovation, ε_t . The innovation is *i.i.d.* across time and plants. In practice, ε_t will have a normal distribution truncated at very large and small values with mean μ and standard deviation σ_z .¹² The example suggests that this model is biased towards finding little asymmetry between aggregate job creation and destruction because plants' idiosyncratic productivity levels are very persistent. However, the assumption of a random walk simplifies comparisons between this model and standard (S, s) models of employment dynamics. Finally, we assume that plants do not face fixed costs of changing employment. While such costs simplified the analysis above, here they introduce unnecessary technical difficulties. In all other respects, the model described in this section is the same as the example.

¹²The assumption of a bounded support is a technical condition used in the analysis of the plant's dynamic programming problem.

We begin this section by describing the profit maximization problem of the representative plant manager and the optimal employment policy it implies. After this we discuss the possible implications of the employment policy for aggregate gross job flow dynamics. Finally, we present a general equilibrium interpretation of the model.

3.1. The Plant Manager's Problem

The problem faced by the representative plant manager is to choose employment at all dates to maximize the present discounted value of the plant's profits. To study this problem, we cast it as a dynamic program. Each period, the plant manager chooses current employment after observing current productivity, employment in the previous period, and the wage. Dropping time subscripts, let m and n denote the level of employment in the previous and current period, respectively, and define the *scaled* employment variables, $x = m/z^{1/(1-\alpha)}$ and $y = n/z^{1/(1-\alpha)}$. The dynamic program which summarizes the plant manager's profit maximization problem is then

$$g(z, x, W) = \max_y z^{1/(1-\alpha)} [y^\alpha - Wy - \tau(y, x)(y - x)] + \beta \mathbf{E}[g(z', y/u, W') | z, W]. \quad (3.1)$$

Here we use the $'$ notation in the usual fashion and $u = (z'/z)^{1/(1-\alpha)} = \exp(\varepsilon'/(1-\alpha))$. The conditional expectations operator is $\mathbf{E}[\cdot]$, and we assume that when forming this expectation the plant owner understands the law of motion for z and W . Finally, $\tau(y, x)$ represents the per-job adjustment costs, which, as before, are measured in units of the output good:

$$\tau(y, x) = \tau_c I\{y > x\} - \tau_d I\{y < x\},$$

where $I\{\cdot\}$ is the indicator function that equals unity if its argument is true and zero otherwise.

In the technical appendix, we show that the value function $g(z, x, W)$ that corresponds to the solution to the sequence version of the plant manager's dynamic profit maximization problem is homogeneous of degree $1/(1-\alpha)$ in z . Therefore, we can focus on solutions to (3.1) which satisfy this property and write $g(z, x, W) = z^{1/(1-\alpha)}v(x, W)$, where $v(x, W)$ is defined to equal $g(1, x, W)$. If we substitute this expression for $g(z, x, W)$ into (3.1) and use the fact that $\ln z$ follows a random walk, we can eliminate z from both sides of the equation. The resulting

Bellman equation is

$$v(x, W) = \max_y y^\alpha - Wy - \tau(y, x)(y - x) + \beta \mathbf{E}[uv(y/u, W') | W]. \quad (3.2)$$

The solution to (3.2) characterizes the solution to the plant manager's profit maximization problem. After being adapted to this particular problem, the standard dynamic programming arguments found in Stokey and Lucas (1989) can be used to demonstrate the existence and uniqueness of a function $v(x, W)$ that satisfies (3.2) and to show that this function is concave in x . These arguments are discussed in more detail in the technical appendix.

3.2. The Optimal Employment Policy

We now describe the policy function for y and show how it is used to formulate the plant manager's optimal employment policy. Because $v(x, W)$ is concave in x (and so differentiable in x almost everywhere) and u has a continuous distribution, $\mathbf{E}[uv(y/u, W') | W]$ is a differentiable function of y .¹³ The differentiability of the expected value function implies that the objective on the right hand side of (3.2) is differentiable in y everywhere except at $y = x$. Therefore, the optimal choice of y must satisfy the following first order conditions

$$\alpha y^{\alpha-1} - W - \tau_c + \beta \mathbf{E}[v_1(y/u, W') | W] \leq 0 \text{ if } y \geq x, \text{ with equality if } y > x, \quad (3.3)$$

$$\alpha y^{\alpha-1} - W + \tau_d + \beta \mathbf{E}[v_1(y/u, W') | W] \geq 0 \text{ if } y \leq x, \text{ with equality if } y < x. \quad (3.4)$$

Here, $v_1(x, W)$ is defined as the right (or left) derivative of the value function with respect to its first argument, which exists everywhere because of the concavity of the value function.¹⁴


Define the *scaled employment thresholds* $\underline{y}(W)$ and $\bar{y}(W)$ as the values of y which satisfy the derivative conditions in (3.3) and (3.4) with equality.¹⁵ Notice that $\underline{y}(W) < \bar{y}(W)$ if either τ_c or τ_d is positive. These functions characterize the optimal scaled employment choice associated

¹³This is proven in the technical appendix. Note that the arguments of Benveniste and Scheinkman (1979), which rely on the differentiability of the return function to establish differentiability of the value function, do not hold here.

¹⁴This function equals the associated derivative of the value function everywhere except on a set of measure zero where the derivative of the value function is undefined.

¹⁵Because $\alpha y^{\alpha-1}$ and $\mathbf{E}[v_1(y/u, W') | W]$ are decreasing functions of y , both $\underline{y}(W)$ and $\bar{y}(W)$ are well defined. We exploit the envelope condition to formulate algebraic expressions for $\underline{y}(W)$ and $\bar{y}(W)$, which are straightforward to solve numerically. See the technical appendix for details.

Figure 3.1: Illustration of Optimal Employment Policies



with (3.2), y^* . If $\underline{y}(W) \leq x \leq \bar{y}(W)$ then employment is unchanged, $y^* = x$. If $x < \underline{y}(W)$, then $y^* = \underline{y}(W)$ and if $x > \bar{y}(W)$, then $y^* = \bar{y}(W)$. With this policy and the definitions of x and y in hand, it is straightforward to construct the optimal policy for the *unscaled* level of employment, n^* . This can be described in terms of the *job creation schedule*,

$$\underline{n}(z, W) = \underline{y}(W) z^{1/(1-\alpha)} \quad (3.5)$$


and the *job destruction schedule*,

$$\bar{n}(z, W) = \bar{y}(W) z^{1/(1-\alpha)}. \quad (3.6)$$

If lagged employment, m , is between $\underline{n}(z, W)$ and $\bar{n}(z, W)$, then the plant neither creates nor destroys jobs, $n^* = m$. If $m < \underline{n}(z, W)$ then the optimal policy specifies $n^* = \underline{n}(z, W)$ so that $n^* - m$ jobs are created. Similarly, if $m > \bar{n}(z, W)$, then the optimal policy specifies $n^* = \bar{n}(z, W)$ and $m - n^*$ jobs get destroyed.

Figure 3.1 illustrates the optimal employment policy for a given value of W . The job creation and destruction schedules are both log linear in z with intercepts $\ln \underline{y}(W)$ and $\ln \bar{y}(W)$ and common slope $1/(1 - \alpha)$. Since $\ln \underline{y}(W) < \ln \bar{y}(W)$, the job creation schedule lies below the job destruction schedule. Consider the employment decisions at three plants with identical lagged employment but different realizations of technology in the current period. These plants are

Figure 3.2: Illustration of Optimal Policy Responses to a Change in W_t



denoted A , B , and C in the Figure. Plant A lies above the job destruction schedule. It chooses current employment to be the value implied by the destruction schedule at its current level of technology. Therefore, it destroys jobs at the rate equal to the vertical distance from A to the job destruction schedule. Plant B lies between the two schedules, so it leaves employment unchanged. Plant C lies below the job creation schedule. It creates job at the rate equal to the vertical distance from C to the job creation schedule.

3.3. Gross Job Flow Dynamics

We are interested in the dynamics of gross job flows implied by a continuum of plants following the optimal employment policy just described. Since the wage determines $\ln \bar{y}(W)$ and $\ln \underline{y}(W)$, it follows that variation in the wage leads to fluctuations in the creation and destruction schedules. The example in Section 2 suggests that $\ln \underline{y}(W)$ will respond by less to a wage change than $\ln \bar{y}(W)$, so we expect the creation schedule will be less variable than the destruction schedule. That is, the employment decisions of job creating plants are expected to be less volatile than the employment decisions of job destroying plants.

Figure 3.2 is helpful for understanding how asymmetry in the volatility of the policy schedules can lead to asymmetric volatility of aggregate job creation and destruction. It contains hypothetical job creation and destruction schedules for the two wage states. The top solid line is $\ln \bar{n}(W^l)$ and the bottom solid line is $\ln \underline{n}(W^l)$. Similarly, the top and bottom dashed lines

are $\ln \bar{n}(W^h)$ and $\ln \underline{n}(W^h)$, respectively. The dashed arrows indicate the movement of the schedules when the wage changes from W^l to W^h . As they are drawn, the job destruction schedule responds by more than the job creation schedule to a wage change.

Consider the employment decisions of four plants with employment and technology as indicated by the labelled points in Figure 3.2. Plant *A* destroys jobs regardless of whether the wage changes, but when the wage changes it destroys more jobs than otherwise. In the absence of a wage change, plant *B* keeps its employment constant, but it destroys jobs when the wage changes. Plants *C* and *D* both create jobs in the absence of the wage change. When the wage changes, plant *C* leaves its employment unchanged and plant *D* creates fewer jobs than otherwise. Because the job destruction schedule shifts by more than the job creation schedule, the intensity of job destruction is increased at plants which behave like *A* more than the intensity of job creation is reduced at plants like *D*. Furthermore, it is possible that the number of plants like *B*, which switch from inactivity to job destruction when the wage changes, is greater than the number of plants like *C*, which switch from job creation to inactivity. These considerations combined suggest that we should expect aggregate job destruction to be more variable than aggregate job creation in the model. However, Caballero's (1992) observations alert us to the possibility that variation in the distribution of plants across the state space might undo the microeconomic asymmetries described here. In the quantitative analysis presented in Section 4 we verify that job destruction does tend to be more volatile than job creation when the policy schedules vary asymmetrically as they do here.

To gain further insight into the mechanism which underlies the dynamics of gross job flows in our model, it is helpful to compare the optimal employment policy with the type of policy studied by Caballero (1992) and Foote (1997). These authors study two-sided (S, s) employment policies in which the relevant state variable is the *employment gap*, defined as the difference between a plant's actual employment and its optimal employment in the absence of adjustment costs. A typical employment policy of this type specifies how to control the employment gap in the presence of plant-specific and aggregate shocks which perturb the frictionless optimum. It is completely specified by two *threshold points*, L and U , and two *return points*, l and u , such that $L \leq l < u \leq U$. These threshold and return points are time invariant, and in particular they do not depend on the aggregate shock. Employment is unchanged as long as the shocks leave the employment gap between L and U . If the employment gap falls below L or exceeds U in

the absence of an employment change, jobs are created or destroyed until it equals the nearest return point.

This kind of policy emerges in a version of our model in which the wage is constant, say $W = 1$, and the technology shock is replaced by the product of the idiosyncratic term we have already described, z_t , and an aggregate shock, A_t . Assume $\ln A_t$ follows a random walk, possibly with drift, with continuously distributed innovation, θ_t , on a compact support. By scaling current and lagged employment by $(z_t A_t)^{1/(1-\alpha)}$, the plant manager's problem is identical to before, with the obvious reinterpretation of the stochastic technology terms in (3.1) and (3.2). The optimal employment policy is derived in the same way as well. Analogous to before, the creation schedule is given by

$$\underline{n}^A(z_t, A_t) = \underline{y}^A [z_t A_t]^{1/(1-\alpha)} \quad (3.7)$$

and the destruction schedule is given by

$$\bar{n}^A(z_t, A_t) = \bar{y}^A [z_t A_t]^{1/(1-\alpha)}, \quad (3.8)$$

where \underline{y}^A and \bar{y}^A solve the appropriately modified derivative conditions in (3.3) and (3.4) with equality. Unlike before, \underline{y}^A and \bar{y}^A are constants and do not depend on the aggregate shock.

We now show how the employment policy for this model can be represented as a standard two-sided (S, s) policy. Since there are no fixed costs, the threshold and return points are identical, $L = l$ and $U = u$. Using the fact that the frictionless optimal employment is given by $(\alpha z_t A_t)^{1/(1-\alpha)}$, the employment gap for this model, s_t , is

$$s_t = \ln n_t^* - \frac{1}{1-\alpha} \ln z_t - \frac{1}{1-\alpha} \ln A_t - \frac{1}{1-\alpha} \ln \alpha, \quad (3.9)$$

where, as before, we use n_t^* to denote the optimal employment policy with adjustment costs.¹⁶

¹⁶Here we have measured the employment gap in logs while Caballero (1992) and Foote (1997) measure it in levels.

The employment gap evolves according to

$$s_t = \begin{cases} l & \text{if } s_{t-1} - \eta_t - v_t < L \\ s_{t-1} - \eta_t - v_t & \text{if } L \leq s_{t-1} - \eta_t - v_t \leq U \\ u & \text{if } s_{t-1} - \eta_t - v_t > U \end{cases}, \quad (3.10)$$

where

$$\begin{aligned} L &= l = \ln \underline{y}^A - (1/(1 - \alpha)) \ln \alpha \\ U &= u = \ln \bar{y}^A - (1/(1 - \alpha)) \ln \alpha \\ \eta_t &= \varepsilon_t / (1 - \alpha) \\ v_t &= \theta_t / (1 - \alpha). \end{aligned}$$

Notice that the evolution of s_t is influenced by both idiosyncratic and aggregate shocks and the thresholds and return points do not depend on the aggregate shock.¹⁷

To appreciate the implications of this type of policy for the dynamics of gross job flows, reconsider Figure 3.2. The creation and destruction schedules for the current model, (3.7) and (3.8), are log-linear as in the wage uncertainty version of the model and have the same slope (see (3.5) and (3.6)). The crucial difference for gross job flows is how the schedules shift in response to aggregate shocks. In the wage uncertainty model we expect the schedules to shift asymmetrically. Here they shift symmetrically since, unlike in (3.5) and (3.6), \underline{y}^A and \bar{y}^A do not depend on the aggregate state of technology and the schedules are equally elastic with respect to the aggregate shock, A_t . Consider a reduction in A_t which shifts the schedules down symmetrically (instead of asymmetrically as they are drawn in Figure 3.2). Because the schedules shift symmetrically, the employment gains at job creating plants are reduced, on average, the same amount as the employment losses at job destroying plants are increased. Therefore, in the absence of asymmetrically distributed aggregate shocks and employment trends, we expect aggregate job creation and destruction to be equally variable.

¹⁷The employment policy in the version of our model with wage uncertainty can also be represented in terms of the employment gap. The transition equation for s_t has the same form as (3.10) with $L(W_t) = l(W_t) = \ln \underline{y}(W_t) + (1/(1 - \alpha)) (\ln W_t - \ln \alpha)$, $U(W_t) = u(W_t) = \ln \bar{y}(W_t) + (1/(1 - \alpha)) (\ln W_t - \ln \alpha)$, $v_t = (1/(1 - \alpha)) \ln(W_{t-1}/W_t)$ and $\eta_t = \varepsilon_t / (1 - \alpha)$. Notice that the thresholds (and return points) are not constants in this case, but vary with the wage.

3.4. General Equilibrium

Although our model has been presented in partial equilibrium, it can be reinterpreted within a general equilibrium framework with particular assumptions on tastes and technology. In this subsection, we sketch such a framework for the version of our model with wage uncertainty. We then use this framework to contrast our model with that studied by Caballero and Hammour (1996).

The economy is populated by a continuum of households and a continuum of consumption good producers. The households are identical and their preferences for sequences of the non-storable consumption good, c_t , a nonstorable good called materials, d_t , and leisure, l_t , are given by

$$\mathbf{E} \left[\sum_{t=0}^{\infty} \beta^t [c_t + (W_t - \gamma) d_t + \gamma l_t] \right].$$

The marginal utility of leisure, γ , and the marginal utility of materials, $W_t - \gamma$, are both positive. The variable W_t is an aggregate disturbance to preferences which evolves according to the Markov chain described previously. In each period households are endowed with fixed quantities of time and materials. Labor, consumption goods and materials are all traded in competitive markets and the consumption good is taken as the numeraire. The consumption good producers are endowed with the technology described in Section 4.1. Given that they are owned by households and we assume a complete set of markets in state contingent claims, the goods producers discount future profits with the net interest rate $1/\beta - 1$. Since labor and materials are employed by plants in fixed proportions, the per-period cost of a job (defined above as the combination of a unit of labor and materials as determined by the fixed proportions technology) is given in equilibrium at each date by W_t , as long as the aggregate supply of labor and materials is sufficient to meet the demand of producers in all states of the world and households always consume some materials and leisure. It follows that the dynamics of gross job flows emerging from this model are identical to the partial equilibrium model.

The simplicity of this framework helps us to focus on the features of a plant's profit maximization problem which generate asymmetries in job creation and destruction decisions. Notice that it is similar to the model studied by Caballero and Hammour (1996). In particular, our assumptions on preferences, endowments and the source of aggregate fluctuations correspond

to their model.¹⁸ There are two main differences with their model and these explain why the results we report in the next section are different from their findings.

First, our production technology is different from the one studied by Caballero and Hammour. In their model output is produced in production units (what we call plants) of constant size which combine capital and jobs in fixed proportions. Since jobs are identical to production units, the job creation costs they study are indistinguishable from costs of installing new capital. Other implications of their technological assumptions are that the marginal product of a new job is invariant to the number of jobs created in any period, and variation in aggregate employment only arises due to changes in the number of production units. In our model, there is a constant number of plants each with a fixed stock of capital and total employment varies due to variation in the number of jobs at individual plants. Also, the adjustment costs we study are net costs of adjusting employment at an existing plant. Finally, due to diminishing returns to labor at the plant level, the marginal product of labor averaged across plants is inversely related to the number of jobs created in a given period in our model.

Second, in our model the labor market is competitive while in Caballero and Hammour search and matching frictions and incomplete contracting over match-specific quasi-rents play important roles. Search and matching dampen variation in the cost of a job due to materials price changes and tend to induce creation and destruction to be concentrated in recessions. Incomplete contracting breaks the tight connection between creation and destruction and helps the model to overcome the counterfactual prediction of positively correlated creation and destruction.¹⁹ Without search and matching frictions, the price of labor does not dampen variation in the cost of a job in our model and there is not a strong force for positive co-movement between creation and destruction. Since movements in the price of labor only partially offset movements in the materials price in Caballero and Hammour, the cost of a job is typically countercyclical in their model.²⁰ Hence, it is likely that our main results reported in the next section, which depend on

¹⁸They do not include leisure in the preferences of the representative household. However, they argue that it is left out only for simplicity and incorporating it does not change their results (see Caballero and Hammour (1996), p. 811 footnote 8).

¹⁹The tendency for positive co-movement between job creation and destruction is a common feature of models with search and matching frictions. Other examples include Andolfatto (1996), Mortensen (1994), and Mortensen and Pissarides (1994). For more discussion on this point, see Caballero and Hammour (1996), pp. 819-825.

²⁰When the cost of job creation is constant, there are no search frictions, and there is incomplete contracting over match specific quasi-rents, the endogenous fluctuations in the wage completely offset the exogenous fluctuations in the materials price, leaving the total cost of employment constant. For this result and similar findings

the countercyclicality of the cost of a job, will continue to hold in a richer environment which incorporates search and matching.

4. Implications of the Model for Gross Job Flows

We have argued that the model developed in the previous section has the potential to generate asymmetries in gross job flows. Now we assess this potential by studying the gross job flows implied by various parameterizations of the model. Our main finding is that the model can generate asymmetries in gross job flows of the type observed in the US manufacturing sector. These asymmetries arise due to asymmetric variation in the job creation and destruction schedules and they are strongest when wage changes are not expected to persist for very long. We begin the discussion by describing how we choose parameter values. After this we study several experiments designed to establish the basic qualitative implications of the model. Finally, we analyze statistics on gross job flows implied by simulating the model and compare them to quarterly data on gross job flows for the transportation equipment industry (SIC 37).²¹ We work with industry-level data since we prefer to view the model as characterizing a single industry, instead of an entire economy or sector as its general equilibrium interpretation might suggest.

4.1. Parameter Values

To implement our model we need to specify the following parameters

Plant-level parameters : $\alpha, \beta, \mu, \sigma_z, \tau_c, \tau_d;$

Aggregate Uncertainty : $\Pi, W^h, W^l.$

The elasticity of production with respect to labor input, α , is approximately equal to the share of gross output paid to labor and materials.²² We set α to 0.85, which is a typical value for this share in the transportation industry.²³ The discount factor β is set to $1.05^{-1/4}$ so that the

in Caballero and Hammour's (1994) partial equilibrium model, the assumption that all new jobs are equally productive is essential. We expect that in versions of these models with a decreasing marginal product of labor, the wage will not completely offset fluctuations in the price of materials.

²¹We confine our analysis to the transportation sector to conserve space. We find qualitatively similar results for other manufacturing industries for which data on gross job flows are available.

²²This will not be exact because of the adjustment costs.

²³See, for example, the 1977 Census of Manufactures.

annual real interest rate is 5% and a period in the model corresponds to one quarter.

Average employment growth and the average job reallocation rate (the sum of the rates of job creation and destruction) are used to identify μ and σ_z . Without aggregate uncertainty, the distribution of employment across plants converges to a steady state, and aggregate employment growth is given by the mean of technology, $\exp(\mu/(1 - \alpha) + \sigma_z^2/(2(1 - \alpha)^2))$. With aggregate uncertainty, employment growth may differ from this formula, but given σ_z and the other model parameters, μ still influences its magnitude as the formula suggests. To avoid the effects of trends on gross job flows, we choose μ so that the model implies zero average employment growth. Holding fixed all other model parameters, the job reallocation rate in the model is increasing in the amount of idiosyncratic uncertainty. Therefore we choose σ_z by matching the model to the average job reallocation rate for the transportation industry.²⁴

Now consider the adjustment cost parameters, τ_c and τ_d . Ideally we would like to take these values from a study of micro data. Unfortunately, this option is not available to us. Recall that τ_c and τ_d are costs of changing the number of employees at a plant. These net adjustment costs involve disruptions to production and all other costs that are not related to the identity of the workers but depend solely on changing the number of employees. As Hamermesh and Pfann (1996) emphasize in their review of work on adjustment costs in factor demand, net adjustment costs are intrinsically difficult to measure because usually they are implicit, in that they result in lost output, and thus are not measured and reported by firms.²⁵ Without good estimates of these costs, we choose them on *a priori* grounds. Since we are unaware of compelling evidence that one adjustment cost is significantly larger than the other, we impose the restriction $\tau_c = \tau_d$. Below, we choose the wage so that its average roughly equals unity. This means the adjustment costs can be interpreted as fractions of the flow cost of a job. It follows that the common value of the adjustment costs multiplied by the job reallocation rate equals total industry adjustment costs incurred as a fraction of total variable costs.²⁶ We set $\tau_c = \tau_d = 1/2$. Combined with a calibrated job reallocation rate of 11.3%, this value implies that adjustment costs incurred by

²⁴Note that we measure gross job flows as the population counterparts to the measures defined by Davis, Haltiwanger and Schuh (1996). In particular, the rate of job creation in a given period is the sum of all jobs created at job creating plants divided by the average of current and lagged total industry employment. Similarly, job destruction is the sum of all jobs destroyed at job destroying plants divided by the average of current and lagged total industry employment. For further details, see the technical appendix.

²⁵Our specification of adjustment costs is consistent with the view that net costs to employment adjustment involve lost output.

²⁶We thank Fabio Schiantarelli for this observation.

the industry as a whole average about 5.65% of total variable costs.

Our last choices involve the specification of aggregate uncertainty. We consider several specifications of the transition matrix for the wage, Π , below. In each case Π has the basic structure

$$\Pi = \begin{bmatrix} \frac{1+\rho}{2} & \frac{1-\rho}{2} \\ \frac{1-\rho}{2} & \frac{1+\rho}{2} \end{bmatrix}, \quad (4.1)$$

where ρ denotes the autocorrelation of the wage. The two wage states are expressed as $W^h = \exp(\sigma_w)$ and $W^l = \exp(-\sigma_w)$, where σ_w denotes the standard deviation of the wage. We select σ_w so that the model reproduces the variance of employment growth for the transportation industry. Note that these choices for Π and the two wage states imply the wage is symmetrically distributed. Therefore the mechanism for generating asymmetries in gross job flows through asymmetries in the aggregate driving process described by Caballero (1992) is not at work in our analysis.


4.2. Illustrating the Basic Mechanism

Here we study the dynamics of gross job flows implied by the optimal employment policy for two polar cases for the wage transition matrix. The first case we consider involves specifying $\rho = -1$ so that the off-diagonal elements of Π equal unity. This implies that changes in the wage are known with certainty to last only one period. That is, the wage follows a two-period cycle, which roughly corresponds to the deterministic cycle studied in Caballero and Hammour (1996). We calibrate parameters as described above.²⁷ To assess the impact of adjustment costs, we also consider a benchmark frictionless case in which all parameters are left at their calibrated values except for the adjustment costs, which are set to zero. In this case employment decisions at plants are governed by the usual static labor demand schedule.

Simulated gross job flows implied by the two parameterizations are illustrated in Figure 4.1. This shows that, regardless of whether there are adjustment costs or not, job creation and destruction move in opposite directions every period. Job destruction is highest when the wage equals W^h and job creation is highest when the wage equals W^l . These dynamics are interesting given that standard search and matching models have a tendency for job creation and

²⁷Values for the parameters not given in the text are $\mu = -0.186\%$, $\sigma_z = 7.48\%$, and $\sigma_w = 4.84\%$.

Figure 4.1: Gross Job Flows in the Two-Period Cycle Model



destruction to co-move positively (which is counterfactual, see below). Without adjustment costs the absolute value of the changes in job creation and job destruction are the same every period so that these variables are equally volatile. The adjustment costs act to dampen fluctuations in both job creation and job destruction. Moreover, job creation is dampened by more than job destruction so that job destruction is more variable than job creation when there are adjustment costs.

The symmetry of gross job flows in the frictionless model can be traced directly to the fact that in this model the job creation and destruction schedules are identical. Positive adjustment costs break the tight connection between the creation and destruction decisions in the frictionless model and introduce a significant asymmetry into how they respond to wage changes. In particular, the difference between the intercepts of the two log destruction schedules, $\ln \bar{y}(W^l) - \ln \bar{y}(W^h)$, is 0.220, while the difference between the intercepts of the two log creation schedules, $\ln \underline{y}(W^l) - \ln \underline{y}(W^h)$, is 0.113: the response of the destruction schedule to a wage change is 95% larger than the corresponding response of the creation schedule. This asym-

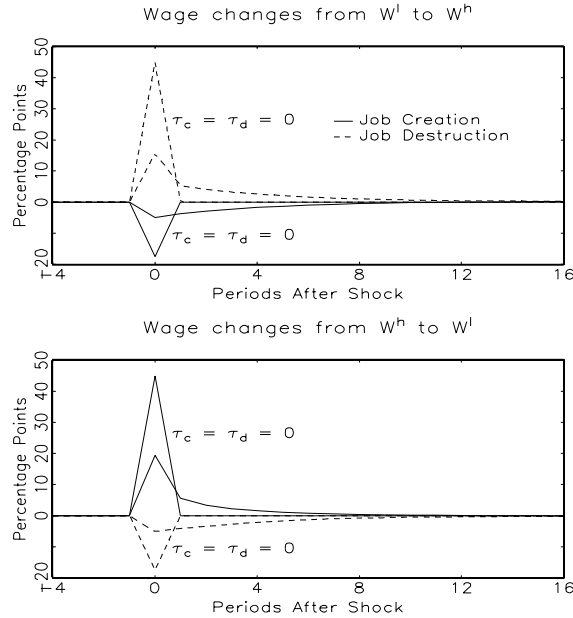
metry induces the intensity of job destruction at shrinking plants to change by more than the intensity of job creation at growing plants when the wage changes and leads to job destruction being more variable than job creation.

The polar opposite of the two-period cycle for the wage is to suppose that when computing their optimal employment policy, plant managers expect the current wage to prevail forever. This is captured by assuming $\rho = 1$ in (4.1). We examine how gross job flows respond to a once and for all unanticipated change in the wage, holding fixed the other parameters at the values use in the two-period cycle examples. This experiment is analogous to the example discussed in Section 2.2. Figure 4.2 plots job creation and destruction for changes in the wage from W^l to W^h (top graph) and from W^h to W^l (bottom graph). In each case the change in the wage occurs after first simulating the model for several periods at the pre-change wage until employment and the distribution of employment across plants converge. Gross job flows under the two assumptions of positive and zero adjustment costs are plotted relative to their pre-wage-change value.

Consider the frictionless cases first. As in the two-period cycle examples, wage changes induce movements in creation and destruction in opposite directions. Job destruction rises and job creation falls following a wage increase, and they reverse their roles following a wage decrease. After one period job creation and destruction return to their pre-wage-change values. In both cases, the magnitude of the initial change in the gross job flow which increases is greater than the corresponding magnitude for the gross flow which decreases. In addition, job destruction's rise when the wage increases is the same as job creation's rise when the wage decreases. Similarly, job creation's fall when the wage increases equals job destruction's fall when the wage decreases. Consequently, the responses of net job growth (not shown) are, as expected, symmetric across the two wage changes.

The responses of job creation and destruction without adjustment costs reflect an interaction between the distribution of idiosyncratic shocks and plants' employment decisions. This interaction plays a significant role in the dynamics of job creation and destruction under positive adjustment costs. Figure 4.2 indicates the adjustment costs dampen the responses of both job creation and destruction in the period of the wage changes. As in the frictionless case, the magnitude of the initial change in the gross job flow which increases is greater than the corresponding magnitude for the gross flow which decreases. In contrast to the two-period cycle example, the adjustment costs appear to dampen destruction by more than creation in the period of the

Figure 4.2: The Dynamics of Gross Job Flows in Response to Permanent Wage Changes



wage changes. The fall in job creation when the wage increases is essentially the same as the fall in job destruction when the wage decreases. However, the rise in job creation when the wage decreases exceeds the rise in job destruction when the wage increases (19% versus 15%). These observations are consistent with empirical evidence reported by Campbell and Kuttner (1996) that employment in the US manufacturing sector adjusts primarily through job creation in response to permanent shocks.

As in the two-period cycle case, positive adjustment costs introduce an asymmetry into how the creation and destruction schedules respond to a wage change. However, the asymmetry is much less pronounced here, and this helps to explain why adjustment costs have a different effect on creation and destruction when wage changes are permanent compared to the two-period cycle case. In the current example, the difference between the intercepts of the two log destruction schedules is 0.663, while the difference between the intercepts of the two log creation schedules is 0.589. Therefore, the shift in the destruction schedule is only 13% larger than the shift in the creation schedule when the wage changes, compared to 95% in the two-period cycle case.

The reason for the reduction in the asymmetry is that when the wage change is permanent the expected savings of job creation costs tomorrow from creating a job today loom larger in the calculation of the effective cost of job creation and the expected cost of reversing the creation decision is less important. Thus, the total expected discounted adjustment costs associated with the last job created fall. The permanent nature of the wage change has a similar impact on the destruction margin by lowering the total expected discounted adjustment costs associated with the last job destroyed. The reduction in total adjustment costs mitigates the asymmetric fluctuations in the job creation and destruction schedules.²⁸

The findings indicated in Figure 4.2 contrast sharply with the permanent wage change examined in Section 2.2. In Section 2.2 the example was constructed so that the distribution of employment across plants was passive in the determination of job creation and destruction. Specifically, a constant fraction of plants increased employment and a constant fraction decreased employment each period regardless of the wage, so that the employment decisions of individual plants were sufficient to understand gross job flows. Here, the population of plants either creating or destroying jobs is influenced by the changes in the employment policy induced by the wage changes. As we described in Section 4.3, this influence can work to offset the impact of the decisions of individual plants on aggregate job flows. The same effect was at work in the two-period cycle case, but since the magnitude of the shifts in the employment schedules were much smaller there, it played a relatively small role. These findings are reminiscent of the fallacy of composition described by Caballero (1992) and highlight the difficulty in drawing conclusions about aggregates from microeconomic decisions.

A final observation worth making about Figure 4.2 is that positive adjustment costs introduce persistence into the responses of gross job flows. The key here is the region of inaction in the state space of the employment policy function induced by proportional adjustment costs. As is well known, this feature of the model provides for an endogenous source of persistence in gross job flows since individual plants are slow to respond to wage changes.

²⁸At this point it may be helpful to reconsider the discussion at the end of Section 2 surrounding the elasticities in (2.5) and (2.6).

4.3. Quantifying the Effects of Adjustment Costs

We now study the statistical properties of the model and compare them to deseasonalized quarterly data on gross job flows for the transportation industry over the sample period 1972:II to 1988:IV.²⁹ For this exercise we specify $\rho = 0$ in (4.1) so that the wage is *i.i.d.* This represents an intermediate case compared to the experiments conducted above, and it seems like a sensible benchmark given the endogenous persistence in gross job flows already present in the model. In Table 4.1 we report statistics on gross job flows for the transportation industry and statistics based on simulating the model. Our baseline case is the calibrated model with positive adjustment costs and we also consider a calibration based on zero adjustment costs.³⁰ Other entries in the table correspond to experiments designed to illustrate various model properties.³¹ We adopt the notation used by Davis, Haltiwanger, and Schuh (1996) to denote the components of gross job flows: job destruction is denoted by *NEG*, job creation by *POS*, job reallocation is denoted by *SUM*, where $SUM = NEG + POS$, and job growth is denoted by *NET*, where $NET = POS - NEG$.

Consider the calibrated frictionless case, $\tau_c = \tau_d = 0$, first. Notice that this model can account for many features of the transportation industry data. The calibration procedure guarantees that the standard deviation of *NET* in the model exactly matches that in the data. In the model and the data *POS* is strongly procyclical, *NEG* is strongly countercyclical, and *POS* and *NEG* are negative correlated. These observations seem in line with the impressions given by Figures 4.1 and 4.2.

The frictionless model has two main shortcomings. First, it underpredicts the variance ratio, $v(NEG)/v(POS)$, the measure of asymmetries in gross job flows used by Davis, Haltiwanger and Schuh (1996). Second, it predicts that job reallocation is acyclical, that is the correlation of *SUM* and *NET* is zero. The failure with respect to the variance ratio derives from underpredicting the volatility of *NEG* and overpredicting the volatility of *POS*. That *SUM* is acyclical is explained by the fact that the covariance between *NET* and *SUM* is equal to the variance of

²⁹The raw data were compiled by Davis, Haltiwanger and Schuh (1996). Our deseasonalization procedure involves removing quarterly means.

³⁰In the baseline case $\mu = -1.85\%$, $\sigma_z = 7.46\%$ and $\sigma_w = 4.31\%$, and in the frictionless case $\mu = -0.14\%$, $\sigma_z = 2.06\%$ and $\sigma_w = 0.37\%$.

³¹All statistics for the model are based on simulating the model for 1050 periods, starting at the steady state distribution of employment across plants, and then discarding the first 50 observations.

Table 4.1: Model Implications for Gross Job Flows

Statistic ⁽ⁱ⁾	Data ⁽ⁱⁱ⁾	Calibrated ⁽ⁱⁱⁱ⁾		Perturbations of the Baseline Calibration ^(iv)				
		Baseline	$\tau_c = 0$ $\tau_d = 0$	$\tau_c = 1$ $\tau_d = 1$	$\tau_c = 0$	$\tau_d = 0$	$E(NEG)$ $= -0.11\%$	$\rho = 0.9$
$\frac{v(NEG)}{v(POS)}$	2.99	1.89	1.00	2.86	1.39	1.38	1.97	0.95
$s(NEG)$	3.55	3.55	3.55	2.30	5.59	5.58	3.56	3.62
$s(SUM)$	2.73	0.65	0.35	0.62	0.82	0.82	0.69	1.25
$s(POS)$	2.73	2.06	1.78	1.45	3.05	3.04	2.09	1.89
$c(POS, NEG)$	-0.30	-0.98	-0.98	-0.99	-0.97	-0.97	-0.98	-0.79
$c(NEG, SUM)$	-0.52	-0.87	0.00	-0.96	-0.56	-0.56	-0.88	0.04
$c(NEG, POS)$	0.67	0.99	0.99	0.99	0.99	0.99	0.99	0.95
$c(NEG, NEG)$	-0.91	-0.99	-0.99	-0.99	-0.99	-0.99	-0.99	-0.94
$\rho(NEG)$	0.08	-0.34	-0.51	-0.35	-0.31	-0.31	-0.34	0.48

Notes: (i) $v(x)$ denotes the variance of variable x , $s(x)$ denotes the standard deviation of variable x , $c(x, y)$ denotes the correlation between variable x and variable y , and $\rho(x)$ is the autocorrelation coefficient of variable x . Standard deviations are reported in per cent. See the text for variable definitions (ii) Based on deseasonalized data for the transportation equipment industry. (iii) Models calibrated to the transportation industry as described in the text. (iv) Models with all parameters identical to the baseline calibration except as indicated.

POS less the variance of NEG . Other weaknesses of the frictionless model are that SUM is too smooth and that NEG is negatively autocorrelated. The former is due to the excessive negative correlation of POS and NEG while the latter follows from the close association of job growth with wage growth (which is negatively autocorrelated here) when employment at the micro-level is determined by a static labor demand schedule.

Table 4.1 shows that the baseline model represents an improvement over the frictionless model. Significantly, the variance ratio rises from unity to 1.89. This is below the value in the data, but still substantial. The improvement in the variance ratio compared to the frictionless case is achieved by reducing the volatility of POS and amplifying fluctuations in NEG , in both cases moving the model closer to the data. These effects are due to asymmetric variation in the employment schedules: the destruction schedule shifts by 71% more than the creation schedule whenever the wage changes. One way to gauge the importance of this asymmetry is to compare the baseline model to a case where we keep the parameters at their baseline values but artificially constrain the log creation and destruction schedules to shift by the same amount in response to

a wage change. Specifically, we fix the schedules as follows:

$$\begin{aligned} \text{Fixed Creation Schedule:} \quad \ln \underline{n}_t &= \ln \underline{y}(1) + \frac{1}{1-\alpha} \ln z_t - \frac{1}{1-\alpha} \ln W_t \\ \text{Fixed Destruction Schedule:} \quad \ln \bar{n}_t &= \ln \bar{y}(1) + \frac{1}{1-\alpha} \ln z_t - \frac{1}{1-\alpha} \ln W_t \end{aligned}$$

where $\underline{y}(1)$ and $\bar{y}(1)$ are the values of the scaled employment thresholds computed in a version of the model without aggregate uncertainty in which the wage is equal to unity. When we simulate this model the variance ratio falls to 0.73.

Given the asymmetry in how *POS* and *NEG* vary, the model also predicts countercyclical job reallocation – the correlation of *NET* with *SUM* is strongly negative as it is in the data. Another improvement over the frictionless model can be seen in terms of the persistence of fluctuations in job growth, which is closer to the data. Finally, incorporating adjustment costs does not improve the model’s predictions with respect to the correlation between *POS* and *NEG*, which is still too strongly negative, and therefore *SUM* remains too smooth relative to the data. However, the fact that *POS* and *NEG* remain negatively correlated as they are in the data is a desirable feature of the model.

For the remainder of this section we discuss how various parameters influence the behavior of the model. In Table 4.1 we report results of three experiments which illustrate how the model’s properties depend on the assumed adjustment costs. For these experiments all parameters are kept at their baseline values except for the adjustment costs. We consider the implications of assuming $\tau_c = \tau_d = 1.0$, $\tau_c = 0$ and $\tau_d = 1/2$, and $\tau_c = 1/2$ and $\tau_d = 0$. The $\tau_c = \tau_d = 1.0$ case reveals that adjustment costs of 11.3% of variable costs give rise to asymmetries in gross job flows approaching that in the data. Except for the reduced volatility expected with larger adjustment costs, this version of the model has the same basic characteristics as the baseline. The $\tau_c = 0$ and $\tau_d = 0$ experiments yield virtually identical results. In particular they both imply a variance ratio significantly greater than unity and, volatility aside, the statistics conform with the baseline. Taken together, the three adjustment cost perturbations demonstrate that it is the overall magnitude of adjustment costs, not whether they are costs of job creation or destruction, which is crucial to determining the size of the variance ratio.

Given Foote’s (1997) results it is interesting to investigate how trends in employment growth influence the model’s predictions. The ‘ $E(NET) = -0.11\%$ ’ experiment in Table 4.1 is designed

for this purpose. Here we have left all parameters at their baseline values except for μ , which has been adjusted so that the model implies average employment growth equal to -0.11% per quarter, the value for the transportation sector over the sample period. As the table shows, the trend has a modest effect on the variance ratio, increasing it from 1.89 to 1.97. The other statistics are not influenced very much either. While larger trends certainly have a greater impact on gross job flows, generally the effects of employment trends seem small relative to the influence of the employment schedules.

The last experiment in Table 4.1 illustrates the impact of highly persistent wage changes. All parameters are again left at their baseline values, except for ρ which is set to 0.9. The asymmetry in gross job flows disappears in this case. This is consistent with the analysis of permanent wage changes in Section 4.2.

5. Conclusion

We have described a model in which asymmetric variation in gross job flows arises naturally from profit maximization in the presence of proportional job creation or destruction costs. The adjustment costs cause employment at shrinking plants to respond by more to aggregate shocks to the cost of a job than employment at growing plants. Our quantitative analysis confirmed that this mechanism can account for asymmetric variation in job creation and destruction of the kind observed in the US manufacturing sector. In addition, we found the mechanism to have the strongest influence on gross job flows when disturbances to the cost of a job are expected to be short lived. Interpreting the cost shocks broadly as being due to changes in energy prices, the terms of trade, or the cost of short term working capital, these findings suggest a new way of understanding the evidence on gross job flows.

The apparent inconsistency of observed gross job flows with standard equilibrium business cycle models has motivated a substantial body of research. This research has emphasized two main departures from the standard framework. First, identical producers have been replaced by heterogeneous producers. Second, frictions in labor markets have been introduced. One conclusion that could be drawn from this literature is that both heterogeneity among producers and labor market frictions are important for understanding the relationship between gross job flows and business cycles. Labor market frictions played no role in our analysis. This suggests

that the second of these modifications may not be important for understanding the evidence on gross job flows.

Although our model has a general equilibrium interpretation, the assumptions of constant marginal utility of consumption, infinitely elastic labor supply, and no saving are too restrictive for quantitative general equilibrium business cycle studies. Accordingly, it would be useful to incorporate our model of plant dynamics into a standard dynamic general equilibrium business cycle framework. Krieger's (1997) general equilibrium analysis of partially reversible investment decisions is suggestive of what we might find by doing this. In his model, a fixed population of heterogeneous plants follow investment policies characterized by investment and disinvestment schedules which have constant and identical slopes, analogous to the job creation and destruction schedules studied here. General equilibrium effects of technology shocks cause the schedules' intercepts to vary over time, which leads plants that scrap capital to be more sensitive to aggregate shocks than plants that expand their capital stock. This finding resembles our result regarding the relative sensitivity of job creating and destroying plants, and so we expect our results will extend to a standard general equilibrium business cycle framework.

Other extensions to our model are suggested by recent empirical work. For example, Caballero, Engel and Haltiwanger (1997) provide support for the view that non-convexities are important for explaining employment variation at manufacturing plants. It would be interesting to determine whether the microeconomic asymmetries we study here extend to environments with non-convexities. Caballero, Engel and Haltiwanger (1996), Campbell and Kuttner (1996) and Davis and Haltiwanger (1997) emphasize the importance of reallocation disturbances. Such shocks could be accommodated in our framework by allowing the variance of the distribution of idiosyncratic shocks to be a random variable. This modification would encourage positive co-movement between job creation and destruction, which would help our model to account for the generally moderate negative correlations between creation and destruction found in the manufacturing data.

References

- Andolfatto, David. 1996. Business Cycles and Labor-Market Search. *American Economic Review*, **86**(1), 112–132.
- Benveniste, L.M. and J.A. Scheinkman. 1979. On the Differentiability of the Value Function in Dynamic Models of Economics. *Econometrica*, **47**(3), 727–732.
- Bertola, Giuseppe and Richard Rogerson. 1996. Institutions and Labor Reallocation. NBER Working Paper 5828.
- Caballero, Ricardo J. and Eduardo M.R.A. Engel. 1991. Dynamic (S, s) Economies. *Econometrica*, **59**(6), 1659–1686.
- Caballero, Ricardo J. 1992. A Fallacy of Composition. *American Economic Review*, **82**(5), 1279–1292.
- Caballero, Ricardo J. and Mohamad L. Hammour. 1994. The Cleansing Effect of Recessions. *American Economic Review*, **84**(5), 1350–1368.
- Caballero, Ricardo J. and Mohamad L. Hammour. 1996. On the Timing and Efficiency of Creative Destruction. *Quarterly Journal of Economics*, **111**, 805–852.
- Caballero, Ricardo J., Eduardo M.R.A. Engel, and John Haltiwanger. 1997. Aggregate Employment Dynamics: Building from Microeconomic Evidence. *American Economic Review*, **87**(1), 115–137.
- Campbell, Jeffrey R. and Jonas D.M. Fisher. 1996. Aggregate Employment Fluctuations with Microeconomic Asymmetries. NBER Working Paper #5767.
- Campbell, Jeffrey R. and Jonas D.M. Fisher. 1997. Technical Appendix to “Aggregate Employment Fluctuations with Microeconomic Asymmetries.”
- Christiano, Lawrence J. and Martin S. Eichenbaum. 1992. Current Real-Business-Cycle Theories and Aggregate Labor-Market Fluctuations. *American Economic Review*, **82**(3), 430–450.
- Davis, S. and J. Haltiwanger. 1990. Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications. *NBER Macroeconomics Annual 1990*, 123–168.
- Davis, S. and J. Haltiwanger. 1992. Gross Job Creation, Gross Job Destruction, and Employment Reallocation. *Quarterly Journal of Economics* **107**, 819–863.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*. MIT Press. Cambridge, Mass.
- Davis, Steven J. and John Haltiwanger. 1996. Driving Forces and Employment Fluctuations. NBER Working Paper #5775.
- Foote, Christopher L. 1997. Trend Employment Growth and the Bunching of Job Creation and Destruction. Unpublished Working Paper. Harvard University
- Hall, Robert E. 1995. Lost Jobs. *Brookings Papers on Economic Activity 1*: 221–273.
- Hamermesh, Daniel S. 1989. Labor Demand and the Structure of Adjustment Costs. *American Economic Review* **79**(4), 674–689.

- Hamermesh, Daniel S., Wolter H.J. Hassink, and Jan C. van Ours. 1994. Job Turnover and Labor Turnover: A Taxonomy of Employment Dynamics. Working Paper, Vrije Universiteit Amsterdam.
- Hamermesh, Daniel S. and Gerard A. Pfann. 1996. Adjustment Costs in Factor Demand. *Journal of Economic Literature* **34**(3) 1264–1292.
- Hopenhayn, Hugo. and Richard Rogerson. 1993. Job Turnover and Policy Analysis: A General Equilibrium Analysis. *Journal of Political Economy*, **60**(5), 1127-1150.
- Krieger, Stefan. 1997. The General Equilibrium Dynamics of Investment and Scrapping in an Economy with Firm Level Uncertainty. Unpublished Manuscript. Yale University.
- Lazear, Edward P. 1995. *Personnel Economics*. MIT Press. Cambridge, Mass.
- Lucas, Robert E. 1978. On the Size Distribution of Business Firms. *Bell Journal of Economics*, **9**(2), 508–523.
- Mortensen, Dale T. 1994. The Cyclical Behavior of Job and Worker Flows. *Journal of Economic Dynamics and Control*, 18, 1121-1142.
- Mortensen, Dale T. and Christopher A. Pissarides. 1994. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, **61**, 397–415.
- Stokey, Nancy L. and Robert E. Lucas Jr. (with Edward C. Prescott). 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press. Cambridge, Mass.