

Working Paper Series

Organizational Flexibility and Employment Dynamics at Young and Old Plants

Jeffrey R. Campbell and Jonas D. M. Fisher

Working Papers Series
Research Department
Federal Reserve Bank of Chicago
December 1998 (WP-98-24)

FEDERAL RESERVE BANK
OF CHICAGO

Organizational Flexibility and Employment Dynamics at Young and Old Plants

Jeffrey R. Campbell*

Jonas D.M. Fisher[†]

October, 1998[‡]

Abstract

There are significant differences in the dynamics of employment over the business cycle between young and old manufacturing plants. Young plants are more sensitive to aggregate disturbances, and they respond to them along different margins. We interpret these differences as reflecting greater organizational flexibility at young plants due to the changing nature of a plant's environment as it ages. In the presence of aggregate uncertainty, differences between young and old plants' organizational flexibility allows the model to reproduce their distinct cyclical characteristics. Previous empirical studies show that small firms generally respond by more to aggregate shocks than do large firms. To the extent that small firms tend to operate young plants, our analysis suggests an alternative to conventional explanations of this evidence which appeal to imperfections in credit markets.

JEL Classification: E24, E32, L16.

*University of Rochester and NBER

[†]Federal Reserve Bank of Chicago.

[‡]For their helpful comments, the authors would like to thank conference participants at the 1998 Society for Economic Dynamics and Econometric Society North American Summer Meetings and seminar participants at the 1998 NBER Summer Institute. We especially thank Glenn MacDonald for his discussion of an earlier version of this paper at the 1998 Econometric Society North American Winter Meetings. The National Science Foundation provided research support through grant SBR-9730442. This paper was completed while Campbell visited the Bradley Center for Policy Research at the University of Rochester's William E. Simon Graduate School of Business, and we thank this institution for its generous research support. Any views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1. Introduction

The well known stylized fact that in the aggregate the rate of job creation is less variable than the rate of job destruction is not a feature of job creation and destruction at young manufacturing plants. The job creation and destruction rates tabulated separately for young and old plants by Davis, Haltiwanger, and Schuh (1996), which Section 2 of this paper considers in more detail, reveal that the job creation and destruction rates for young plants are approximately equally variable. In contrast, for older plants the job creation rate fluctuates much less than the job destruction rate. Other significant differences between the two groups of plants emerge from our analysis. The average job creation and destruction rates for young plants are higher than the averages for old plants. Moreover, young plants' job creation, job destruction, and net employment rates are more variable than their counterparts for old plants. Evidently, gross job flows at young plants behave quite differently than those at old plants.

Davis and Haltiwanger's (1990, 1992) original finding regarding the relative variability of job creation and destruction challenged standard business cycle theory because it is difficult to account for in models where establishments can be aggregated into a single representative producer, as in Campbell and Fisher (1997). Similarly, the evidence on young and old plants challenges the body of theory that has emerged to account for the behavior of aggregate job creation and destruction. In this paper we address this challenge by exploring the implications of employment adjustment costs, which have been considered by Caballero and Hammour (1994) and Campbell and Fisher (1998) to account for the aggregate evidence on gross job flows.¹ Our analysis is founded on an organizational interpretation of employment adjustment costs. In this we follow Lucas (1967), who interprets costs of adjusting a firm's capital stock as reflecting the necessity of planning the integration of new capital goods into the production process, and

¹Other theoretical approaches have been taken to address the empirical shortcomings of the representative producer paradigm and these are also challenged by the empirical evidence on young and old plants. One approach, exemplified by Mortensen and Pissarides (1994) and Caballero and Hammour (1996), focuses on search and contracting frictions between workers and employers. A second approach shows how asymmetries in aggregate driving processes can cause job destruction to be more volatile than job creation. Foote (1997) Caballero (1992), Davis and Haltiwanger (1996) and Campbell and Kuttner (1996) are representative of that literature. For a further discussion of these approaches and their relationship to the adjustment cost approach, see Campbell and Fisher (1998).

Hamermesh and Pfann (1996) and Campbell and Fisher (1998) who interpret costs of adjusting employment in a similar organizational light. An explanation for the many observed differences between the employment dynamics at young and old plants suggests itself when one considers the organizational nature of adjustment costs: plants find it optimal in their youth to employ an organization which is more flexible and entails lower adjustment costs than when they are mature.

The changing nature of a plant's environment as it ages motivates the organizational differences between young and old plants in our model. Two features of this change are the learning by young plants documented by Bahk and Gort (1993) and the decline in the cross sectional variance of the growth rates of plants as they age, as shown by Dunne, Roberts, and Samuelson (1989). The theoretical literature on organizational choice predicts that the greater opportunities for innovation which young plants possess and the greater risk they face should both induce their managers to choose more flexible organizations than older plants. For example, Athey and Schmutzler (1995) show that the return to an investment in organizational flexibility increases with an enterprise's opportunities for innovation, while de Groote (1994) shows that the degree of flexibility of the optimal production process is increasing in the risk an enterprise faces. Young plants learning through either production or direct experimentation about an uncertain production process are therefore more likely to invest in organizational flexibility than mature plants that have exhausted the learning curve and operate a stable technology.² The forms of organizational flexibility which Athey and Schmutzler (1995) cite as useful for implementing new process and product innovations, such as using more educated workers or using few job classifications, can also lower the cost of changing the scale of a plant's activity.³ These considerations suggest that, when they are confronted with shocks to their environments, young and flexible plants will respond by more and along different margins than old and inflexible plants.

²Choices regarding organizational flexibility are distinct from choices regarding the technological flexibility of plant design as studied by Chang (1993), He and Pindyck (1992), Fuss and McFadden (1978), Gilchrist and Williams (1998) and Stigler (1939). The ex-ante plant design decision may influence the scope for later changes in the organization of production at a plant. However, the incentives to change the organization of production at a plant over its life cycle should be present regardless of the initial plant design decision.

³Simon (1991) observes that one function of organizations is to coordinate the expectations of many agents participating in a complicated production process. The extent to which the plant exploits this opportunity is another margin upon which a manager can vary the flexibility of a plant's organization.

We develop our explanation for the pattern of gross job flows at young and old plants by constructing a simple model of job creation and destruction over the plant life-cycle in which young plants' organizations tend to be more flexible than those of old plants. Our analysis builds on the employment adjustment cost model of Campbell and Fisher (1998). In that model aggregate gross job flows are explained by the choices made at a fixed population of plants. Since plant entry and exit are not responsible for the empirical results which motivate our analysis, we introduce a plant life-cycle into this model in the simplest way by assuming that plants exit at a constant rate and are replaced at that same rate by entrants. The life-cycle of a plant consists of two stages. After a plant is born, it experiments with the production process. During this *learning* stage, the plant can costlessly change its employment between periods. This assumption embodies the idea that learning plants employ a flexible organization. In any given period a learning plant's manager can decide to end the learning stage and enter the *mature* stage by adopting the process used in the most recent experiment. This decision is irreversible and it leads to generally higher productivity which is less uncertain from period to period than in the learning stage. Changing the employment of a mature plant requires paying adjustment costs. This assumption reflects the idea that mature plants do not invest in organizational flexibility. As in Bentolila and Bertola (1992), Bertola and Rogerson (1997) and Hopenhayn and Rogerson (1993), these costs are proportional to the number of jobs created or destroyed. We interpret them as costs of reorganizing the production process to accommodate a larger or smaller scale.

In the presence of aggregate uncertainty, the difference between learning and mature plants' organizational flexibility allows the model to reproduce the observed cyclical differences between young and old plants. The source of aggregate fluctuations we consider is changes to the real cost of a job, which includes real wages and the cost of complementary materials inputs, such as oil. As Campbell and Fisher (1998) show, adjustment costs of the kind considered here tend to induce contracting plants to respond by more than expanding plants to aggregate shocks. This microeconomic asymmetry can cause the job destruction rate to fluctuate by more than the job creation rate in the aggregate. Here, this kind of behavior will only be observed at mature plants, which tend to be older. Learning plants, which tend to be younger, do not face adjustment costs, and because of this their job creation and destruction rates are roughly

equally variable. In the model, job creation and destruction among young and old plants reflect the differences between learning and mature plants. Hence, the model is capable of reproducing the facts that job creation and destruction rates for young plants are equally volatile while for old plants the job destruction rate is more variable than the job creation rate. Moreover, the model is also successful at reproducing the other main differences in gross job flows between young and old plants emphasized above.

The empirical differences between young and old plants which motivate this paper are similar to, but distinct from, those which motivated the theoretical literature on industry dynamics and the plant life-cycle. Using the same data which underlies Davis, Haltiwanger, and Schuh's (1996) time series on job creation and destruction, Dunne, Roberts, and Samuelson (1989) found significant cross sectional differences between young and old plants. Young plants' grow faster, their growth rates have a higher cross sectional variance, and they exit more frequently than their older counterparts. Jovanovic (1982) explains the higher cross sectional variance of young plants and their larger exit rate as the outcome of a selection process in which young plants learn about their plant specific productivity as they age. In our model, learning through experimentation causes young plants' fast growth. Our work also complements Jovanovic's (1982) explanation for young plants greater cross sectional variance by documenting and modeling their greater time series variance.

Our analysis is also related to the macroeconomics literature which studies the behavior of small and large firms. Gertler and Gilchrist (1994) and Bernanke, Gertler and Gilchrist (1996) document differences in the responses of small and large firms to monetary shocks. These empirical studies show that output is generally more responsive to monetary shocks at small firms than at large firms, due to a relatively large decline in inventories at small firms and essentially no decline in inventories at large firms. A common interpretation of this pattern of responses is that small firms are impeded in their access to credit markets while large firms are not.⁴ In our analysis there are no differences in credit market access across plants. All differences in the cyclical behavior of plants are due to differences in their organizational flexibility. To the extent that small firms tend to operate young plants, our analysis suggests an alternative

⁴See Fisher (1998) for a formal articulation and evaluation of this view.

explanation for the greater adjustment of inventories by small firms. That is, organizational flexibility at small firms reduces any motive to smooth production.

The remainder of the paper is organized as follows. The next section presents a study of Davis, Haltiwanger, and Schuh's (1996) job creation and destruction data broken down by plant age. Section 3 describes the model economy. Section 4 studies various parameterized versions of the model, and Section 5 contains concluding remarks.

2. Evidence on Gross Job Flows by Plant Age

In this section we document how the cyclical behavior of gross job flows differs across plants in the US manufacturing sector depending on their age. Our analysis is based on quarterly and annual data constructed by Davis, Haltiwanger, and Schuh (1996) (hereafter DHS).⁵ The sample period is 1972:II to 1988:IV for the quarterly data and 1973-1988 for the annual data. We find that there are significant differences in the behavior of gross job flows for young versus old plants.⁶ As a basis for comparison, we begin our discussion by reviewing the cyclical properties of gross job flows for the US manufacturing sector as a whole. This is followed by our analysis of plants by age.

2.1. The US Manufacturing Sector

We follow DHS by studying gross job flows in terms of rates of job creation, destruction, growth and reallocation. For a given population of plants, the rate of job creation (destruction) in a period, is defined as the total number of jobs added (lost) since the previous period at plants which have higher (lower) employment in the current period compared to the previous quarter, divided by the average of total employment in the current and previous period. The rate of job growth is given by the difference between the rates of job creation and destruction, and the rate of job reallocation is given by their sum. Figures 1A and 1B plot these variables at the quarterly frequency for the US manufacturing sector as a whole.⁷ The vertical lines in Figure 1 indicate

⁵See the technical appendix of DHS for detailed information about the construction of this data.

⁶DHS briefly discuss some of the differences between young and old plants described here. See Table 4.5, p.77, Table 5.4, p.97, and the surrounding text of DHS.

⁷All quarterly series studied here have been seasonally adjusted by removing quarterly means.

business cycle peaks and troughs, as defined by the NBER.

These figures illustrate several well known facts about gross job flows (see, for example DHS). First of all, the indicated rates of job creation and destruction illustrate that on an ongoing basis relatively large numbers of jobs are either created or destroyed. In fact, as Figure 1B shows, between 8 and 14 percent of all jobs are created or destroyed every quarter. Second, the creation rate is less variable than, and tends to co-move negatively with, the destruction rate. Third, the reallocation rate co-moves negatively with the rate of job growth, that is job reallocation is countercyclical.

The statistics (standard errors in parenthesis) reported in the first column of Table 1 confirm these impressions. For the four variables of interest this table includes means, standard deviations and selected correlation coefficients. Also included in the table is the ratio of the variances of the rates of job destruction and creation, a statistic emphasized by Davis and Haltiwanger (1992). Notice that the excess volatility of the destruction rate relative to the creation rate is statistically significant, as is the negative correlation between these two variables and the negative correlation between job growth and job reallocation.⁸ Not surprisingly, job creation is significantly procyclical and job destruction is significantly countercyclical.

2.2. Young and Old Plants

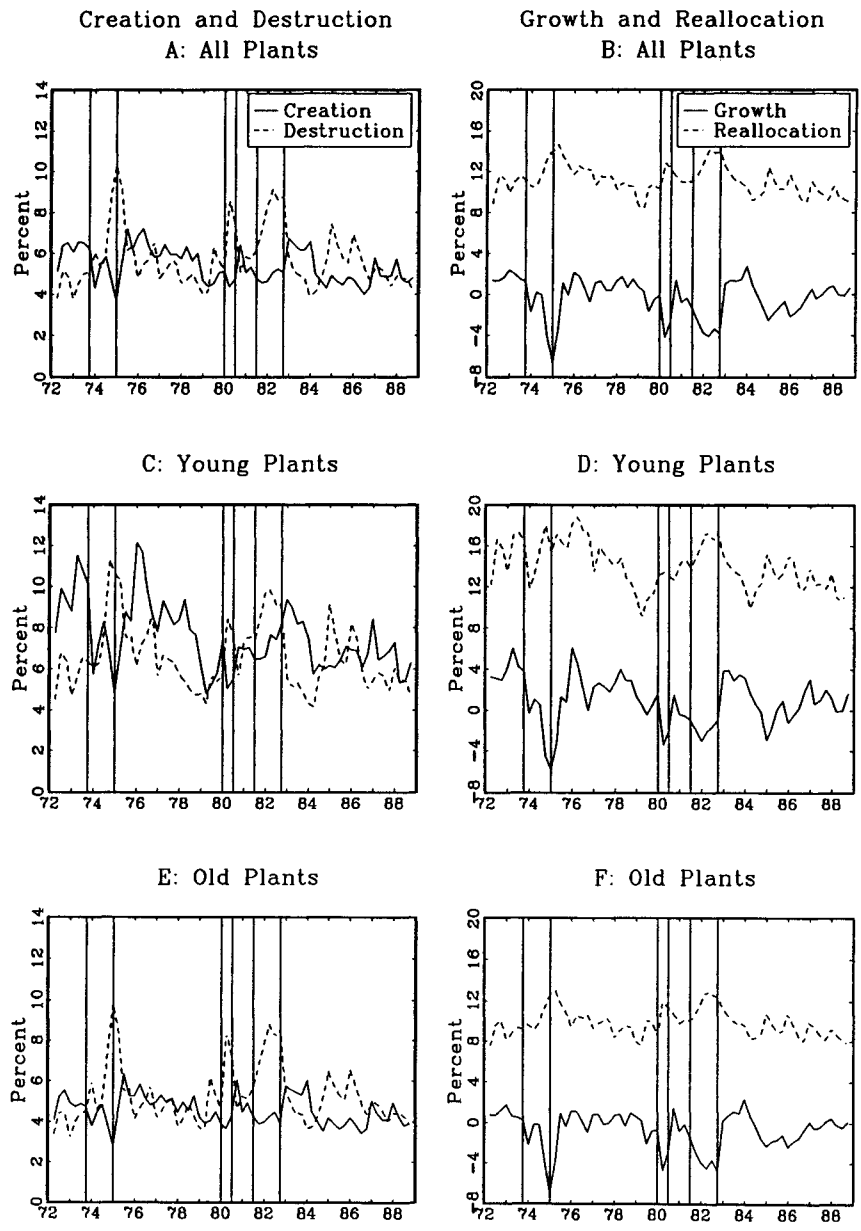
To study the relationship between establishment age and gross job flows, DHS measure job creation and destruction rates for plants in three different age categories. DHS (p. 225) recommend aggregating the two categories which include the youngest establishments and we do this here. We refer to this combination as 'young' plants. These plants are typically less than about 10 years old and account for 22.5% of total employment on average over the sample period.⁹ We refer to the remaining plants as 'old' plants.¹⁰ The bottom two rows of Figure 1 display gross

⁸Note that the larger variability of job destruction relative to job creation and the definitions of the job growth rate and the job reallocation rate imply that the covariance between the latter two variables is negative.

⁹We aggregate the two youngest categories by adding their job creation and destruction rates after weighting them by their total employment shares. We then divide the results by the sum of the two groups' employment shares.

¹⁰Because of the sample design of the underlying data source the threshold between young and old changes over time. The minimum age of an old establishment is at least 9 years and at most 13 years. These changes only occur at the end of a year.

Figure 1: Gross Job Flows in the Manufacturing Sector, 1972:II-1988:IV



job flows for the young and old plants.

Figures 1B and 1D reveal striking differences between the creation and destruction rates of young and old plants. The creation and destruction rates appear to be larger and more volatile for young plants than for old plants. Among young plants, the creation rate varies about the same amount as the job destruction rate, while for old plants job destruction is much more variable than job creation, as in the manufacturing sector as a whole. These differences in variability between young and old plants are particularly evident during non-recessionary periods. During the recoveries of 1975 – 76 and 1982 – 83, the creation rate for young plants reached the levels the destruction rate did during the preceding recession, while job creation at older plants rose only modestly. Figures 1C and 1E confirm that the differential behavior of job creation and destruction between young and old plants carries through to job growth and reallocation. The reallocation rate is higher and more volatile for younger plants. In addition the rate of job growth at young plants is more volatile than at old plants. Generally, job growth at young plants varies consistently throughout the business cycle, while job growth at old plants seems to vary most around recessions and is relatively smooth outside of recessions.

The last two columns of Table 1 summarize this evidence. The sample statistics verify that gross job flows are more volatile at young plants. Furthermore, they reveal that job creation and destruction at young plants do not conform to the well-known regularities from the manufacturing as a whole. We cannot reject the hypotheses for young plants that job creation and destruction are equally variable, that these variables are uncorrelated, and that job growth and reallocation are also uncorrelated. In contrast, the qualitative pattern of statistics for the old plants seem in line with the evidence for all of the manufacturing sector.

Table 2 reports the results of formal tests of these differences between young and old plants. The first row displays Wald-type statistics which test for the equality between young and old plants of the summary statistics indicated in the column headings. This table indicates that the average rate of job reallocation, the excess variability of job destruction compared to job creation, the variability of job growth, the correlation between creation and destruction, and the correlation between job growth and reallocation for young plants are all significantly different than the analogous statistics for old plants, at very high levels of confidence.

2.3. Robustness

The foregoing discussion is based on data which include plant startups and shutdowns. Due to idiosyncracies in how the data are collected, the contributions of these plants to aggregate measures of gross job flows at the quarterly frequency is not based on direct observations but is derived from the information on continuing plants (plants which continue to operate in successive sample periods.) Hence, it is reasonable to be concerned about whether the differences in the cyclical behavior of gross job flows between young and old plants just discussed are genuine and not artifacts of the data construction procedure implemented by DHS. To address this concern we have analyzed the analogous annual data on gross job flows by plant age, in which direct observations on startups and shutdowns are available. We find that all the differences highlighted above appear in the annual data as well and that these differences continue to be statistically significant at very high levels of confidence (not shown).

It is also of interest to examine the behavior of continuing plants separately from startups and shutdowns. This provides some guidance on how the differences between young and old plants might be explained. For example, if the differences were primarily due to job reallocation at plant startups and shutdowns, then an explanation based on the determinants of plant entry and exit would seem promising. On the other hand if the young-old plant differences are present among continuing plants then an explanation based on the determinants of employment decisions at incumbent plants would seem more appropriate. We have analyzed gross job flows at continuing plants using both quarterly and annual data. With two exceptions all the differences highlighted above are reflected in the gross job flows of continuing plants. First, at the quarterly frequency there does not appear to be a significant difference between the job growth rates of young and old plants. However, this difference remains statistically significant at the annual frequency.¹¹ Second, at the annual frequency the differences in the correlation between job creation and job destruction evident in the quarterly and annual data when startups and shutdowns are included in the data disappear when continuing plants are studied separately. The difference is still

¹¹At the quarterly frequency the point estimate of the standard deviation of the rate of job growth is slightly lower for young plants than for old plants – 1.68 (0.21) versus 1.70 (0.24). At the annual frequency the young-old comparison is 5.05 (0.15) versus 4.17 (0.13).

apparent, though less pronounced, in the quarterly data for continuing plants.

To summarize, there appear to be significant differences in the behavior of gross job flows between young and old plants. Furthermore, these differences generally do not depend on the frequency of the data considered or on whether startups and shutdowns are included in the analysis. The main differences between young and old plants, which we will emphasize in what follows, are (i) the rate of job reallocation is larger at young plants, (ii) job destruction and job creation are equally variable at young plants but job destruction is more variable than job creation at old plants, (iii) the rate of job growth is more variable at young plants, and (iv) job reallocation is acyclical at young plants but countercyclical at old plants.¹²

3. The Model Economy

Addressing the observed differences between young and old plants' aggregate employment dynamics requires a model with ongoing aggregate uncertainty as well as heterogeneity of plants across and within age cohorts. In this section we describe a competitive equilibrium industry dynamics model which incorporates these features in a simple way. As in Caballero (1992) and Campbell and Fisher (1998), a fixed measure of plants which experience both idiosyncratic and aggregate shocks populates the industry. To introduce an age distribution, we assume that plants exit at a constant rate and are replaced at the same rate. New plants learn through experimentation about the production process. They operate in the context of a flexible organization to facilitate the learning process. This flexibility manifests itself in it being costless for a learning plant to change its scale of production. In any period the manager of a learning plant can choose to adopt its most recent experiment and take her plant into the mature phase of its life-cycle in which no more learning takes place. In the absence of learning, mature plants operate with a less flexible organization in which it is costly to change the scale of production.

¹²It is interesting to note that, with one exception, the differences we emphasize here regarding young and old plants also hold if the manufacturing data is divided by plant size (small plants with less than 100 employees, and large plants) and by ownership characteristics (plants owned by single-plant firms versus plants owned by multi-plant firms). The one exception is that there is much less evidence in favor of differences in job growth rates when the manufacturing sector is divided along these lines. The similarity in the findings for plant size compared to plant age are consistent with the Dunne, Roberts, and Samuelson (1989) finding that young plants tend to be smaller than average.

The formal description of the model proceeds in three steps. We first describe the environment facing individual plants. Then we characterize and discuss the solution to the problem faced by the manager of a *mature* plant, that is a plant which has stopped learning. Here we make extensive use of results in Campbell and Fisher (1998) (hereafter CF). Finally, we consider the problem of a *learning* plant's manager who must decide whether to take her plant into the mature phase of its life-cycle by adopting the production process from its latest experiment.

3.1. The Environment

The industry is composed of a continuum of plants which produce an homogeneous good for sale in a competitive goods market and purchase inputs in competitive factor markets. We assume that the industry is small relative to these markets, so that its product demand and factor supply curves are all infinitely elastic. As a consequence, equilibrium prices do not depend on the industry's scale of production and aggregate quantities are irrelevant for the decisions of individual plants. This simplifies the analysis considerably since it means that gross job flows, which are the focus of our analysis, can be calculated easily by aggregating the decisions of individual plants that have been computed by taking prices as given.¹³

Plants are operated by risk neutral managers who value future profits with the constant discount factor β . After every period, a fraction δ of all operating plants exit. They are replaced at the beginning of the next period by an identically sized mass of new plants. Each plant uses two variable factors of production which must be combined in fixed proportions: labor, which comes in fixed shift lengths, and materials. We refer to an employee plus her accompanying unit of materials as a job. The per-period cost of a job at time t , measured in units of the constant output price, is denoted by W_t . Although W_t reflects both labor and materials costs, we refer to it simply as the wage. It is the only source of aggregate uncertainty and follows a Markov chain over the set $\{W^1, W^2, \dots, W^N\}$ with transition matrix Π . All time t decisions by plant managers are made after observing W_t . Note that we include materials in the model to highlight the fact that changes in W_t can reflect changes in the prices of inputs which complement labor

¹³Although we present our model in partial equilibrium, it can be reinterpreted in a general equilibrium framework with particular assumptions on tastes and technology along the lines considered by Campbell and Fisher (1998).

as well as changes in the direct price of labor.

Let the output, idiosyncratic productivity level, and employment at date t of a representative plant be denoted by y_t , z_t , and n_t , respectively. Output is produced according to

$$y_t = z_t n_t^\alpha,$$

where α is assumed to be strictly between 0 and 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).

The process governing z_t depends on whether the plant is in its learning phase or has progressed to the mature phase. As long as a plant is in its learning phase, realizations of z_t are an *i.i.d.* sequence of random variables. However, a learning plant's manager does not observe z_t before choosing how many workers to employ, n_t , and production takes place. Instead, she views a signal of z_t , z_t^s , and learns z_t only after production takes place. Let z_t^p denote the part of z_t learned through production and define $z_t = z_t^s z_t^p$. We assume that z_t^s and z_t^p are independently distributed log-normal random variables:

$$\ln z_t^s \sim N(\mu_s, \sigma_s^2) \quad \text{and} \quad \ln z_t^p \sim N(\mu_p, \sigma_p^2).$$

Hence,

$$\ln z_t \sim N(\mu_l, \sigma_l^2),$$

where $\mu_l = \mu_s + \mu_p$ and $\sigma_l^2 = \sigma_s^2 + \sigma_p^2$.

At the beginning of the next period, with knowledge of z_t in hand, the manager of an incumbent learning plant must decide whether to continue learning or to adopt the most recent experiment and cease learning, that is take the plant into the mature phase of its life-cycle. If the manager decides to continue learning about the production process, then it draws a new value of z_t^s and chooses a new level of employment. The flexibility of learning plants' organizations implies that there are no costs to changing the scale of production at these plants. If the manager decides to take the plant into its mature phase, then the value of z_{t+1} will equal z_t . In

all subsequent periods the process for z_t is a random walk in logarithms:

$$\ln z_{t+1} = \ln z_t + \varepsilon_{t+1}.$$

The innovation, ε_{t+1} , is assumed to be *i.i.d.* over time and have a normal distribution truncated above and below by very large numbers, with mean, μ_m , and variance, σ_m^2 .¹⁴ At the time of adoption, the manager also chooses employment. Because learning plants use flexible organizations, this choice is not constrained by costs of adjusting the plant's scale. However, it has dynamic consequences since thereafter it is costly for the plant to change scale between periods. If a mature plant expands, it incurs the job creation cost τ_c for every job added, and if it contracts, it incurs the job destruction cost τ_d for every job lost.¹⁵

3.2. Mature Plants

The manager's problem is to choose employment at all dates to maximize the present discounted value of plant profits. To study this problem, we cast it as a dynamic program. There are three state variables: the plant's current productivity, its employment in the previous period, and the wage. Each period, the manager observes the state variables and chooses current employment. Using m to denote employment in the previous period, and dropping time subscripts, the dynamic program is


$$g(z, m, W) = \max_n zn^\alpha - Wn - \tau(n, m)(n - m) + \beta(1 - \delta) \mathbf{E}[g(z', n, W') | z, W]. \quad (1)$$

Here $\mathbf{E}[\cdot]$ is the conditional expectation operator, and we use the " notation in the usual fashion. In forming the expectation the manager understands the law of motion for z and W . Also, $\tau(n, m)$ represents the per-job adjustment costs, which are measured in units of the output

¹⁴Truncation is a technical condition necessary for the analysis of the plant's dynamic programming problem.

¹⁵To 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 organizational.

Figure 2: Illustration of Optimal Employment Policies for Mature Plants



good:

$$\tau(n, m) = \tau_c I\{n > m\} - \tau_d I\{n < m\}.$$

CF show that two functions, a *job creation schedule*, $\underline{n}(z, W) = \underline{y}(W) z^{1/(1-\alpha)}$ and a *job destruction schedule*, $\bar{n}(z, W) = \bar{y}(W) z^{1/(1-\alpha)}$, characterize the optimal employment policy. If lagged employment, m , is between $\underline{n}(z, W)$ and $\bar{n}(z, W)$, then the plant neither creates nor destroys jobs. If m is less than $\underline{n}(z, W)$ then the optimal policy specifies employment equals $\underline{n}(z, W)$ so that $\underline{n}(z, W) - m$ jobs are created. Similarly, if m is greater than $\bar{n}(z, W)$ then the optimal policy specifies employment equals $\bar{n}(z, W)$ and $m - \bar{n}(z, W)$ jobs get destroyed.

Figure 2 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 $1/(1 - \alpha)$.¹⁶ Since $\underline{y}(W) < \bar{y}(W)$, 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 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

¹⁶The fact that these schedules are log linear in technology is due to the random walk assumption on the evolution of technology at mature plants. A mean reverting process would introduce nonlinearities into the schedules which would complicate the analysis considerably.

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 in the *region of inaction*, so it leaves employment unchanged. Plant C lies below the job creation schedule. It creates jobs at the rate equal to the vertical distance from C to the job creation schedule. Because adjustment costs penalize employment changes, we expect that increasing their size will widen the region of inaction and reduce average gross job flows at mature plants.

We are interested in the dynamics of gross job flows generated by fluctuations in the wage. 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 of mature plants. The analysis of CF 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 mature plants are expected to be less volatile than the employment decisions of job destroying mature plants. CF show, in an economy comprised entirely of mature plants, how this microeconomic asymmetry translates into greater variability of aggregate job destruction compared to aggregate job creation. Therefore, we expect that aggregate job destruction at mature plants will be more variable than aggregate job creation at mature plants here.

Since the microeconomic asymmetry is crucial for the dynamics of gross job flows at mature plants it is helpful to understand how it arises. To this end, consider the effects of a temporary change in the wage. For an expanding plant, the total cost of the last job created is greater than the wage 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 the wage below unity. For a shrinking plant, the total cost of the last job retained is less than the wage 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 the wage 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 wage

changes can induce contracting plants to respond by more than expanding plants to fluctuations in the wage.¹⁷

3.3. Learning Plants

Each period, the manager of a learning plant has one or two decisions to make, depending on whether the plant is an entrant or an incumbent. If it is an entrant then the only choice to make is how many workers to employ. If it is an incumbent then the manager must also decide whether or not to adopt the technology realized from experimenting in the previous period. Recall that this adoption decision is made before the manager observes a new productivity signal and decides how many workers to employ in the period. Below we first describe the employment decision of an entrant and an incumbent that chooses to continue learning. After this we describe the employment decision of a plant which has made the decision to adopt its most recent experiment. Finally, we describe the adoption decision.

If a plant is an entrant or an incumbent that chooses to continue learning, its current employment choice has no dynamic consequences: optimal employment maximizes only expected profits for the period, conditional upon the productivity signal of the current experiment, z^s . Let $\pi^l(z^s, W)$ denote the expected current profits of a learning plant which receives the productivity signal z^s when the wage equals W . The associated employment policy is characterized by a familiar static labor demand schedule. Therefore we expect the creation and destruction decisions of entrants and incumbents which choose to continue learning to be equally variable. Moreover, since there is no region of inaction in the employment policy for these plants we expect their gross job flows generally to be larger on average and more variable than those at mature plants.

Now consider the employment decision of a plant which has already decided to adopt its most recent experiment. Due to the adjustment costs faced by a plant in its mature phase, the manager of an adopting plant must consider the consequences of its current employment choice on future profits. The profit maximization problem for a plant which is adopting its most recent

¹⁷The 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.

experiment, z , when the wage is W is

$$\max_n zn^\alpha - Wn + \beta(1 - \delta) \mathbf{E}[g(z', n, W') | z, W], \quad (2)$$

where $g(z, n, W)$ is the solution to (1). CF show that $g(z, n, W)$ can be written as

$$g(z, n, W) = z^{1/(1-\alpha)} v\left(n/z^{1/(1-\alpha)}, W\right),$$

where $v(\cdot, \cdot)$ is strictly concave in its first argument. Let

$$x = \frac{n}{z^{1/(1-\alpha)}}, \quad (3)$$

and rewrite (2) as

$$\max_x z^{1/(1-\alpha)} (x^\alpha - Wx + \beta(1 - \delta) \mathbf{E}[uv(x/u, W') | W]), \quad (4)$$

where the random variable u is defined to equal $(z'/z)^{1/(1-\alpha)} = \exp(\varepsilon/(1-\alpha))$. CF also demonstrate that $\mathbf{E}[uv(x/u, W') | W]$ is a strictly concave and differentiable function of x . It follows that (4) has a unique solution, characterized by its first order condition. Denote this optimal choice of x with $y^a(W)$, where the superscript 'a' stands for 'adoption'. Rescaling $y^a(W)$ as specified by (3) yields $y^a(W) z^{1/(1-\alpha)}$ as the optimal employment for a plant which is adopting a technology with productivity z . This policy resembles the policy for nonadopting learning plants in that there is no region of inaction. Nevertheless, in our simulation analysis described below we find that the configuration of adjustment costs at mature plants is important for determining the impact that this policy has on the dynamics of gross job flows. Finally, note that the present discounted value of following this policy can be written as $z^{1/(1-\alpha)} v^a(W)$. That is, the value function associated with (4) is an increasing, homogeneous of degree $1/(1-\alpha)$ function of z .

We now characterize the optimal adoption decision. Let $v^l(z, W)$ denote the value of a learning plant at the beginning of a period with productivity from its most recent experiment given by z and current wage W . The Bellman equation associated with the optimal adoption

decision is then given by

$$v^l(z, W) = \max \left\{ z^{1/(1-\alpha)} v^a(W), \mathbf{E} \left[\pi^l(z^{s'}, W) + \beta(1-\delta) v^l(z', W') \mid W \right] \right\}. \quad (5)$$

If W is constant, then (5) is formally identical to the problem of optimal wage search without recall presented by Sargent (1987, Chapter 2). In the search problem, $z^{1/(1-\alpha)} v^a(W)$ is the ‘wage’ and $\mathbf{E} \left[\pi^l(z^{s'}, W) \mid W \right]$ is the ‘unemployment benefit.’ Sargent’s arguments can be applied directly in the constant wage case to demonstrate that the $v^l(z, W)$ which solves (5) is unique and weakly increasing in z . Therefore, analogous to the search problem, a simple threshold rule characterizes the optimal adoption decision: adopt any experiments z which exceed the threshold. In the case where W is stochastic, the optimality of a time varying threshold rule can be more easily demonstrated by using the contraction mapping theorem to characterize $v^l(z, W)$.¹⁸ According to this threshold rule, a plant will adopt a production process if its productivity from experimentation exceeds a lower bound which depends on the current wage, $z(W)$. A key implication of the optimal adoption policy is that the older a plant is, the more likely it is to have made the transition from the learning phase to the mature phase of its life-cycle.

4. Aggregate Fluctuations

Due to the form of the adoption decision, mature plants will tend to be older than average, so we expect their behavior to have a disproportionate influence on the gross job flows of old plants. Similarly, learning plants will tend to be younger than average, and so we expect their behavior will dominate the gross job flows of young plants. With this in mind, the intuition regarding learning and mature plants’ employment policies described in the previous section strongly suggests that the model can reproduce many of the differences between young and old plants’ gross job flow dynamics. However, the model contains several potentially confounding factors which the intuition does not address. Caballero’s (1992) observations alert us to the possibility

¹⁸The proof of this assertion is available from the authors upon request.

that variation in the distribution of plants across the state space might undo the microeconomic asymmetries described here. Furthermore, systematic variation in plants' optimal adoption decisions could have a significant impact on gross job flows. To demonstrate that the differences between learning and mature plants' employment policies drive young and old plants' observed differences, we present in this section evidence from simulating several parameterized versions of the model.

Our discussion is focused around a baseline calibration of the model in which the wage is assumed to be *i.i.d.* over time. The baseline model reproduces the salient features of the gross job flow observations discussed in Section 2. The rate of job reallocation is larger at young plants than at old plants, job destruction and job creation are equally variable at young plants while job destruction is more variable than job creation at old plants, the rate of job growth is more variable at young plants than at old plants, and job reallocation is acyclical at young plants but countercyclical at old plants. Because the baseline model's driving process is *i.i.d.* over time, its fluctuations reflect cycles at all possible frequencies equally. To determine how high and low frequency fluctuations separately contribute to the model's properties, we also consider several versions of the model in which W_t is specified to follow a deterministic cycle. We also consider how the magnitude of adjustment costs at mature plants influences employment dynamics in the model. These two sets of experiments highlight how the adoption decision can have a significant influence on the dynamics of gross job flows in the model.

4.1. Parameter Values

To implement our model we need to specify the following parameters

Plant-level parameters : $\alpha, \beta, \delta, \mu_m, \sigma_m, \mu_s, \sigma_s, \mu_p, \sigma_p, \tau_c, \tau_d$

Aggregate Fluctuations : $\Pi, W_1, \dots, W_N.$

A desire for our model to mimic a "representative" manufacturing industry guides our choice of parameter values. The elasticity of production with respect to labor input, α , is approximately

equal to the share of gross output paid to labor and materials.¹⁹ We set α to 0.85. This is a typical value for the transportation equipment sector, which is the largest two digit manufacturing industry.²⁰ The discount factor β is set to $1.05^{-1/4}$ so that the annual real interest rate is 5% and a period in the model corresponds to one quarter. Because the event of a plant's exit is independent of its characteristics, the model's exit rate, δ , equals the fraction of industry employment at plants which exit between periods. This is set to equal its average value in the DHS (1996) sample for the manufacturing sector, 0.83%.²¹

We adopt several different specifications for the transition matrix and support of W_t . In the baseline case, $\ln W_t$ approximates a Gaussian *i.i.d.* process with a standard deviation of σ_W . Our calibration procedure chooses σ_W so that the standard deviation of net employment growth in the baseline model roughly matches its empirical counterpart from the manufacturing sector (see Table 1). For this case, the support of W_t and its transition matrix are chosen using Tauchen's (1984) method for approximating autoregressive processes with Markov chains.²² In the deterministic cases, the support of W_t and the transition matrix are chosen so that $\ln W_t$ follows cycles of various frequencies, in each case restricted so that the unconditional standard deviation matches σ_W from the baseline case. In all examples we study, the average value of W_t equals one.

Now consider the six parameters which characterize the learning process and the evolution of idiosyncratic productivity. Two of these parameters are redundant, and we utilize four substantial restrictions to identify the remaining ones. First, note that both μ_s and μ_p only affect the scale of individual plants, so they have no impact on model statistics which are scale free, such as those considered in Table 1. We set both of these so that the average employment of a learning plant equals one.

Our first substantial restriction is that the employment of mature plants does not tend to increase or decrease. In the long run, the average growth rate of employment at a mature

¹⁹This will not be exact because of the adjustment costs paid by mature plants.

²⁰See, for example, the 1977 Census of Manufactures.

²¹In DHS's notation, this equals the average value of *NEGD*, job destruction at exiting plants, over their sample period.

²²The approximation uses seven possible states which are equispaced over three standard deviations of the wage's logarithm.

plant equals $\exp(\mu_m/(1 - \alpha) + \sigma_m^2/(2(1 - \alpha)^2))$. Given a value for σ_m^2 , we choose μ_m so that this expression equals one. Holding fixed all other model parameters, the job reallocation rates for young and old plants increase in the amount of idiosyncratic uncertainty. Our second and third restrictions exploit this by constraining the average reallocation rates for old and young plants in the baseline model to equal their empirical averages reported in Table 1. Note that throughout this section we define the set of young plants to be all those plants less than 40 quarters old, which corresponds roughly to the age cut-offs applicable in the construction of the data for young plants underlying Table 1.²³ Our final restriction is that productivity is more certain among mature plants than among learning plants. We operationalize this assumption by constraining $\sigma_l^2 = 12\sigma_m^2$. That is, three years of productivity innovations have the same variance as one experimental draw. Our main findings are unaffected by increasing the ratio of these variances from 12 to 20.

Finally, we need to specify 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.²⁴

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 $\tau_c = \tau_d$. Since we choose the wage so that its average equals unity, the adjustment costs can be interpreted as fractions of the flow cost of a job. It follows that the

²³For both groups and for the group comprising all plants in the industry, we measure gross job flows as the population counterparts to the measures defined by DHS. In particular, the rate of job creation in a given period and for a given group of plants is the sum of all jobs created at job creating plants in the group divided by the average of current and lagged total group employment. Similarly, job destruction is the sum of all jobs destroyed at job destroying plants in the group divided by the average of current and lagged total group employment.

²⁴Our specification of adjustment costs is consistent with the view that net costs to employment adjustment involve lost output.

common value of the adjustment costs multiplied by the aggregate job reallocation rate provides an upper bound for total industry adjustment costs incurred as a fraction of total variable costs. We set $\tau_c = \tau_d = 1/2$ in our baseline model. Combined with a steady state job reallocation rate of about 11%, this value implies that adjustment costs incurred by the industry are no more than 5.5% of total variable costs. Because our choices of τ_c and τ_d are somewhat arbitrary, we do report results from several experiments which use different values for these costs.

Table 3 summarizes our parameter choices. The only surprising result from our calibration exercise is the small value of σ_s , the standard deviation of the experimental productivity signal. In the absence of this signal, learning plants only change their employment in response to aggregate disturbances. The small calibrated value of σ_s indicates that employment at learning plants responds so much to aggregate shocks that little idiosyncratic uncertainty is required for them to reproduce the average job creation and destruction rates for young plants. Indeed, if we set $\sigma_W = 0$ and recalibrate, the calibrated value of σ_s triples to 0.021 while the other parameter values change very little.

4.2. The Baseline Model

Table 4 reports statistics from simulating the baseline calibrated model.²⁵ The first moments show that the model captures the salient features of the empirical evidence on average gross job flows. By construction, the calibrated model exactly matches the average job reallocation rates for young and old plants, however it also nearly reproduces the overall job reallocation rate. The model's steady state value of job reallocation equals 10.31%, while the average of this statistic in the data equals 11.11%.²⁶ In the data, young plants' average job creation rate exceeds their average job destruction rate, and the opposite is true for old plants. The model's first moments also display this pattern, although young plants' grow too quickly and old plants

²⁵All the model statistics in this paper are based on simulating the model for 1100 quarters, beginning from the steady state. Before calculating sample statistics, we disposed of the first 100 observations. We used the same seed for all experiments involving draws from a random number generator.

²⁶This underprediction of the industry-wide job reallocation rate reflects the under representation of young plants in the model relative to the U.S. manufacturing sector. The average share of employment at young plants is about 11% in the model, while the analogous share measured by Davis, Haltiwanger, and Schuh is about 21%. Therefore, the baseline model may also underpredict young plants' influence on other overall statistics.

shrink too slowly relative to the evidence reported in Table 1. The reasons for the positive net growth rate for young plants and negative net growth rate for old plants are twofold. First, exit effects all plants equally, while entry contributes to employment at young plants only. Second, while learning plants which continue to experiment with the production process do not grow on average, the transition from the learning phase to the mature phase typically involves plants expanding their employment. This effect contributes to positive net growth at young plants.

The model is also successful at reproducing many of the cyclical asymmetries between young and old plants. By construction, the standard deviation of overall employment growth in the model is about the same as in the data, about 1.9%. As in the data, the standard deviation of employment growth among old plants is a small amount below this figure. The analogous statistic for young plants is considerably higher. The ratio of the variance of job destruction to the variance of job creation is 1.61 for the industry as a whole. This ratio is lower in the model than it is for the manufacturing sector as a whole (see Table 1), but it is close to the variance ratio for several two digit manufacturing industries.²⁷ The statistics for old plants reflect this asymmetry, but those for young plants do not. The variance ratio for young plants is 1.12 and the analogous statistic for old plants is 1.79. Finally, notice that job reallocation for young plants is close to being acyclical, while job reallocation is strongly countercyclical at old plants.

In contrast to these successes, the model counterfactually predicts a nearly perfect negative correlation between job creation and destruction for both types of plant. In the data this correlation is essentially zero among young plants, and among old plants, it equals -0.33 . This value is not atypical for the correlation between overall job creation and destruction in two digit manufacturing industries.²⁸ The excessively strong negative correlation between the two gross job flows in the model reflects the assumption that a single variable which changes the industry's labor supply curve is the sole source of aggregate fluctuations. Including other sources of aggregate uncertainty which cause job creation and destruction to comove positively, such as the investment specific technology shocks in Campbell's (1998) model of entry and exit, would probably alter this feature of the model.

²⁷In their Table 5.2, DHS report variance ratios for the Lumber industry, the Stone, Clay and Glass industry, the Transportation Equipment industry, and the Instruments industry of 1.54, 1.61, 1.71 and 1.85, respectively.

²⁸See Table 5.2 of DHS.

Further insight into the sources of the model's behavior can be gained by inspecting the learning and mature plants' optimal policies. The static labor demand schedule which characterizes the optimal employment choice of learning plants is log-linear in z^s with slope $1/(1-\alpha)$ and intercept equal to $\ln y^s(W) = \ln(W/\alpha)/(1-\alpha)$.²⁹ The creation and destruction schedules of mature plants and the employment policy of plants just entering the mature phase are also log-linear with the same slope as the static labor demand schedule and intercepts given by $\ln \underline{y}(W)$, $\ln \bar{y}(W)$, and $\ln y^a(W)$, respectively. These four intercepts along with the optimal technology adoption threshold, $\underline{z}(W)$, completely characterize plants' optimal policies. The first column of Table 5 reports some useful statistics summarizing their behavior in the baseline model. To aid with their interpretation, it also reports statistics for $\Pr[z_i \geq \underline{z}(W)]$, the probability of a learning plant adopting its latest experiment. The remaining columns of Table 5 report the analogous statistics from variations on the baseline model which we discuss below.

The means of the employment schedules reveal that $\ln y^s(W)$ tends to lie between $\ln \underline{y}(W)$ and $\ln \bar{y}(W)$, so it is impossible to unambiguously sign the impact of job creation and destruction costs on an individual plant's employment. However, the fact that $\ln y^a(W)$ is generally less than $\ln y^s(W)$ indicates that the prospect of future adjustment costs tends to lower a plant's optimal employment. Because $\ln y^a(W)$ also lies between $\ln \underline{y}(W)$ and $\ln \bar{y}(W)$, plants which just adopted a technology are located in the region of inaction and so will contribute very little to aggregate employment fluctuations.

The volatility statistics highlight the features of the plant's decision problems which drive the results in Table 4. As in CF, $\ln \underline{y}(W)$ is less variable than $\ln \bar{y}(W)$, which indicates that the job creation schedule responds less to aggregate shocks than the job destruction schedule. This policy asymmetry is the direct cause of the variance asymmetry between the job creation and destruction rates for old plants. High variation in the static labor demand schedule, which learning plants use, causes most of the employment variation at young plants. Indeed, the standard deviation of $\ln y^s(W)$ is more than double that of $\ln \bar{y}(W)$. Thus, it appears that even moderately sized adjustment costs can greatly decrease the variance of employment. In contrast, changes in the policies of adopting plants contribute very little to aggregate fluctuations. The

²⁹The superscript 's' in $y^s(W)$ stands for 'static'.

employment policy of adopting plants is less variable than the job creation schedule, while the technology adoption threshold barely moves at all. The decision to end a plant's learning phase and bring it into maturity has inherent long run consequences, so the invariance of the technology adoption threshold to temporary factor price changes is not surprising. When plants do choose to adopt a technology, the employment level they choose tends to lie between the job creation and destruction schedules for mature plants. This implies that adjustment on either margin is not likely in the short to medium run so that transitory disturbances should have little weight in this decision. In contrast, expanding and contracting mature plants are likely to be adjusting employment again soon because they are already close to one of the adjustment schedules. Because the decisions by these plants to delay job creation or destruction can be very easily reversed in the near future, their optimal employment decisions are more sensitive to transitory disturbances.

4.3. Deterministic Cycles

Because the baseline model's aggregate driving process is *i.i.d.* over time, its fluctuations reflect cycles at all possible frequencies equally. By simulating versions of the model in which W_t follows a deterministic cycle, we can determine how high and low frequency fluctuations separately contribute to the model's properties. Table 6 reports statistics from the model's gross job flows for two such experiments. The first three columns report statistics for all plants, young plants, and old plants from solving and simulating the model when W_t follows an eight quarter deterministic cycle. The second three columns report the analogous statistics when W_t follows a twenty quarter cycle.³⁰ The second and third columns of Table 5 contain the statistics describing the plants' optimal policies. Although we have solved the model using many other deterministic cycles, these two simulations are representative of our findings.

Table 6 indicates that the gross job flow asymmetries which adjustment costs induce are strongest when the driving process has relatively high frequency fluctuations. In particular, the gross job flow statistics from the model with an eight quarter cycle are broadly similar to those

³⁰Recall that our calibration procedure chooses the unconditional variance of $\ln W_t$ in these experiments to equal its value in the baseline model.

from the baseline model and so they accord well with the data. As we noted in Section 2, the high frequency fluctuations in the data which are not obviously associated with any particular business cycle episode reveal most acutely the differences between young and old plants' gross job flows. This seems consistent with the eight quarter cycle model, since the frequency of the driving process in this model lies somewhere between seasonal and business cycle frequencies. The statistics describing the optimal policies in this case are also very similar with those from the baseline model. The average values of the employment schedules' intercepts and the technology adoption threshold are identical in the two parameterizations. The job creation and destruction schedules' intercepts have standard deviations which are slightly higher than those from the baseline model. It is unsurprising that these forward looking employment decisions are more responsive to W_t in this parameterization because the process for W_t exhibits more persistence. For the same reasons, the employment schedule for plants that are entering maturity is also more variable. By construction, the first and second moments of the frictionless employment schedule are identical across all of the model parameterizations we consider.

Figure 3 provides additional insight into establishments' optimal policies by graphing $\ln W_t$ and the log intercepts to the four employment schedules over the cycle. So that they will all fit on the same scale, each of the plotted series has had its average value removed. Naturally, the employment schedules all comove negatively with $\ln W_t$. The intercepts of the forward looking employment policies, the job creation and destruction schedules and the employment schedule of adopting plants, all lead $\ln W_t$ by one quarter. This lead is economically interesting but not quantitatively significant because the forward looking employment policies' intercepts change very slowly near their peaks and troughs. To confirm this, Figure 4 plots net employment growth for both young and old plants for a 24 quarter period from the simulation. It is clear from the figure that employment growth at young and old plants move together.

When a twenty quarter cycle drives the model's fluctuations, its behavior contrasts sharply with that for higher frequency cycles. As CF noted, the job creation and destruction schedules of plants facing proportional adjustment costs display little asymmetry when aggregate fluctuations primarily reflect low frequency movements. The policy statistics in the third column of Table 5 confirm that this effect operates here as well. The job creation and destruction schedules are

Figure 3: Employment Policy Intercepts and the Wage in the Model with an Eight Quarter Cycle

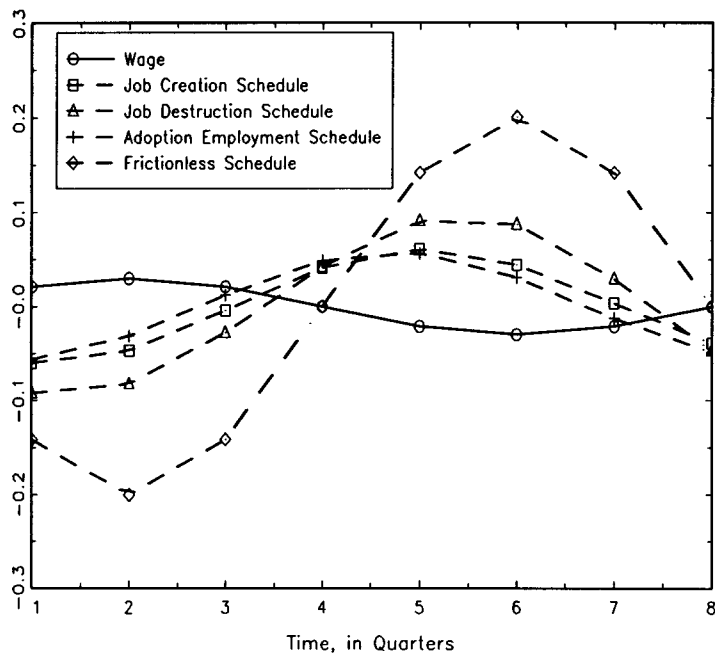
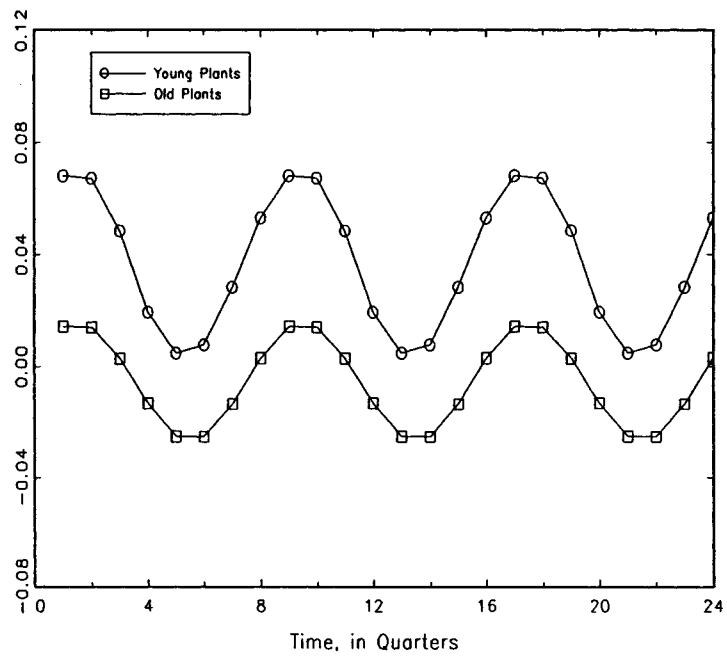


Figure 4: Employment Growth at Young and Old Plants in the Model with an Eight Quarter Cycle



both much more volatile when W_t follows a twenty quarter cycle than when it follows an eight quarter cycle, because W_t is much more persistent. The job destruction schedule still moves more than the job creation schedule, but the difference between the two standard deviations is small relative to their absolute magnitudes. Therefore, old plants' gross job flows display no asymmetry in this case. In contrast to the old plants, job creation at young plants fluctuates substantially *more* than job destruction. This reflects the fluctuations in the employment of plants which are adopting their most recent experiment (see Table 5). These fluctuations are more pronounced when low frequency fluctuations change W_t because this decision is intrinsically forward looking. Plants which are adopting a technology tend to be young and growing, so these fluctuations primarily effect young plants' job creation.

4.4. Adjustment Costs

The existing empirical literature provides very little guidance regarding the values of τ_c and τ_d , the per-job creation and destruction costs. Therefore, it is desirable to assess the sensitivity of the model with respect to these choices. We do so with two sets of experiments. In the first set, we vary both τ_c and τ_d together between 0 and 1.5 and resimulate the model using the baseline specification for the other parameters. The results of these experiments (not shown) are entirely consistent with the intuition developed above. To conserve space, we focus on what these experiments reveal about asymmetries in gross job flows. When both adjustment costs equal zero, there is little substantial difference between young and old plants. The overall job creation and destruction rates have equal variances, as do those for old plants. Job creation at young plants is slightly more variable than job destruction. If the adjustment costs are both set above their baseline values of 0.5, then the ratio of job destruction's variance to job creation's increases for old plants and for all plants. Young plants' gross job flows, however, still exhibit substantially more symmetry than those of old plants. When both τ_c and τ_d equal 1.5, the ratio of job destruction's variance to job creation's equals 2.81 for old plants and 1.14 for young plants.

The second set of experiments were designed to examine the separate contributions of job creation and destruction costs to the model's behavior. Here, we simulate the model twice. Each

time, we set one of the adjustment costs equal to the extreme value of zero while maintaining the other at its baseline value. Table 7 reports the job flow statistics from these two simulations. The first three columns report statistics for the $\tau_c = 0$, $\tau_d = 0.5$ case. Eliminating job creation costs from the baseline model causes three main changes in its behavior. First, the gross job flows are larger on average with the lower adjustment costs. Second, both the gross and net employment flows respond more to aggregate shocks. Finally, the ratio of job destruction's variance to job creation's falls, regardless of whether one considers all plants together or either age group. The fall of the variance ratio for old plants is not surprising, because decreases in the magnitude of the adjustment costs tend to decrease the microeconomic asymmetries in mature plants' employment policies. What is surprising about this result is that the variance ratio for young plants falls considerably, from 1.12 to 0.72. The last three columns of Table 7, which report statistics from the case where $\tau_c = 0.5$ and $\tau_d = 0$, shed some light on this finding. The most noticeable change between these statistics and those of the previous case is in the variance ratios. The variance ratios for both groups of plants are significantly above one, and that for young plants is *higher* than that for old plants.

The difference between these two experiments is surprising since CF find that the results from their model do not depend on whether adjustment costs are particular to either job creation or destruction. Indeed, the fourth and fifth columns of Table 5, which report statistics describing plants' optimal policies, indicate that mature plants' job creation and destruction schedules behave almost identically in the two parameterizations. To explain the difference in their gross job flow statistics, note that the particular configuration of adjustment costs has a significant impact on the level of employment chosen by plants adopting their most recent experiment. Recall that such plants face no adjustment costs for this initial choice of scale. Because of this assumption, the employment choice of an adopting plant coincides exactly with that from an expanding mature plant with the same productivity level when $\tau_c = 0$. Therefore, young plants which have just adopted a technology are bunched near the job creation schedule. As Foote (1997) explained, such bunching increases the variability of job creation for these plants while decreasing it for job destruction. This explains the sizeable fall in the variance ratio for young plants when $\tau_c = 0$. The opposite happens when $\tau_d = 0$ and τ_c is positive. In this case, the

employment choice of an adopting plant coincides exactly with that from a *shrinking* mature plant with the same productivity level. Young plants are bunched near the destruction schedule in this case, so the variance of job destruction increases and that of job creation decreases.

These last two experiments reveal that the bunching of plants in the state space which occurs over the plant life-cycle can impact aggregate statistics as significantly as the aggregate trends which Foote (1997) considered. Furthermore, explicitly adding plant life-cycle considerations to the sensitivity analysis of Campbell and Fisher (1998) overturns their finding that the ability of adjustment costs to reproduce the job creation and destruction evidence does not depend on whether those costs are incurred with job creation or with job destruction. While their conclusion still holds if one only considers the aggregate evidence, the model with only job creation costs cannot reproduce the aggregate dynamics of young plants.

5. Conclusions

Young plants' failure to mimic the employment dynamics of their older counterparts presents a direct challenge to the body of macroeconomic theory which arose to explain aggregate gross job flows. Our model of organizational differences over the plant life-cycle addresses this challenge for one class of models, those which rely on organizational frictions to motivate plant level employment adjustment costs. In our model, the changing nature of a plant's environment as it ages underlies organizational differences between young and old plants, and it is these organizational differences which give rise to the observable differences in their employment dynamics. Because young plants maintain flexible organizations which entail no adjustment costs, they respond more to aggregate shocks and on different margins than do old plants which are less flexible and face costs of job creation and destruction. The microeconomic asymmetries identified by Campbell and Fisher (1998) cause gross job flows at old plants to respond asymmetrically to aggregate shocks. Without adjustment costs, employment at young plants responds on both the job creation and destruction margins equally. Therefore, the model reproduces the finding that young plants' job creation and destruction rates are roughly equally variable, but job destruction at old plants varies considerably more than job creation. Adjustment costs tend to reduce job

flows and dampen plants' responses to aggregate disturbances, so the model also reproduces the fact that the volume and variability of gross job flows at young plants are greater than at old plants. The features of the model we emphasize are most pronounced in the presence of high frequency aggregate disturbances, so differences between young and old plants in our model most resemble those in the U.S. data in this case.

Our work has implications for macroeconomics beyond this paper's focus on the gross job flows measured by DHS, because it represents the first articulation of an economic rationale for differences in how groups of producers respond to business cycle shocks which does not rely on imperfections in credit markets. We have focused on aggregate differences among plants to connect our work to the existing evidence on gross job flows. However, our plants could be easily reinterpreted as firms, which may be a more appealing object of study when considering organizational choice. Under this interpretation, our model implies that young firms respond more to business cycle shocks than do old firms. Further research into the organizational approach to adjustment costs should reveal additional implications which will allow us to further differentiate it from existing credit market theories. For example, the connection we have emphasized between technological learning and plants' organizational choices suggests that the pace and form of *aggregate* technological change could influence the distribution of flexibility among producers and therefore the magnitude and character of an economy's response to business cycle shocks.

Although our model can reproduce the salient differences between employment dynamics at young and old plants, the further development of an organizational theory of adjustment costs will require a richer framework. One microeconomic limitation of our model is its direct connection of the activity of learning through experimentation with flexibility in scale. Existing models of plants' organizational flexibility choices, such as Milgrom and Roberts (1990) and Athey and Schumtzler (1995), endogenize similar connections by demonstrating mutual complementarity between investments in flexibility along several dimensions of the firm's problem. Extending these models to production problems with an infinite horizon and life-cycle considerations is therefore an important step in developing a more complete macroeconomic theory of organizationally based adjustment costs.

Two macroeconomic limitations of our model are the assumptions that the relevant labor

supply curve is infinitely elastic and that all managers use a constant real interest rate to discount future profits. These were useful for gaining intuition regarding plants' optimal employment decisions and their aggregation, but they rule out any important general equilibrium feedback to producers' behavior from consumers' labor supply, consumption and investment decisions. To properly understand how investments in organizational structure and the resulting adjustment costs contribute to business cycle dynamics, the employment choice problem with endogenous organizational change should be integrated into a standard business cycle framework.

References

- Athey, Susan and Armin Schumtzler. 1995. Product and Process Flexibility in an Innovative Environment. *RAND Journal of Economics*, 26(4), 557–574.
- Bahk, Byong-Hyong and Michael Gort. 1993. Decomposing Learning by Doing in New Plants. *Journal of Political Economy*, 101(4), 561–583.
- Bentolila, Samuel and Guiseppe Bertola. 1990. Firing Costs and Labour Demand: How Bad is Eurosclerosis. *Review of Economic Studies*, 57(3), 381–402.
- Bernanke, Ben, Mark Gertler and Simon Gilchrist. 1996. The Financial Accelerator and the Flight to Quality. *Review of Economics and Statistics* 78(1), 1-15.
- Bertola, Guiseppe and Richard Rogerson. 1997. Institutions and Labor Reallocation. *European Economic Review*, 41(6), 1147–1171.
- 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(3), 805–852.
- Campbell, Jeffrey R. 1998. Entry, Exit, Embodied Technology, and Business Cycles. *Review of Economic Dynamics*, 1(2), 371–408.
- Campbell, Jeffrey R. and Jonas D. M. Fisher. 1997. Understanding Aggregate Job Flows. *Economic Perspectives* 21(5), 19-37.
- Campbell, Jeffrey R. and Jonas D. M. Fisher. 1998. Aggregate Employment Fluctuations with Microeconomic Asymmetries. Unpublished Manuscript. University of Rochester.
- Campbell, Jeffrey R. and Kenneth N. Kuttner. 1996. Macroeconomic Effects of Employment Reallocation. *Carnegie-Rochester Conference Series on Public Policy*, 44, 87–116.

- Chang, M.H. 1993. Flexible Manufacturing, Uncertain Consumer Tastes, and Strategic Entry Deterrence. *Journal of Industrial Economics* 41, 77-90
- Christiano, Lawrence J. and Martin Eichenbaum. 1992. Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations. *American Economic Review* 82(3), 430-50.
- Davis, Steven J. and John Haltiwanger 1990. Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications. *Pages 123-186 of: Blanchard, Olivier Jean, & Fischer, Stanley (eds). NBER Macroeconomics Annual*
- Davis, Steven J. and John Haltiwanger. 1992. Gross Job Creation, Gross Job Destruction, and Employment Reallocation. *Quarterly Journal of Economics*, 107(3), 819-863.
- Davis, Steven J. and John Haltiwanger. 1996. Driving Forces and Employment Fluctuations. NBER Working Paper # 5775.
- Davis, Steven J., John Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*, MIT Press, Cambridge, Mass.
- de Groot, Xavier. 1994. The Flexibility of Production Processes: A General Framework. *Management Science*, 40(7), 933-945.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. The Growth and Failure of U.S. Manufacturing Plants. *Quarterly Journal of Economics*, 104(4), 671-698.
- Fisher, Jonas D.M. 1998. Credit Market Imperfections and the Heterogeneous Response of Firms to Monetary Shocks. *Journal of Money, Credit and Banking*, forthcoming.
- Foote, Christopher L. 1997. Trend Employment Growth and the Bunching of Job Creation and Destruction. Working Paper, Harvard University.
- Fuss, M.C. and D. McFadden. 1978. Flexibility versus Efficiency in Ex Ante Plant Design. In M.C. Fuss and D. McFadden, eds., *Production Economics: A Dual Approach to Theory and Applications*. Amsterdam: North-Holland.
- Gertler, Mark and Simon Gilchrist. 1994. Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms. *Quarterly Journal of Economics*, 109(2), 311-340.

- Gilchrist, Simon and John C. Williams. 1998. Putty–Clay Investment: A Business Cycle Analysis. Unpublished Working Paper, Boston University.
- Hamermesh, Daniel S. and Gerard A. Pfann. 1996. Adjustment Costs in Factor Demand. *Journal of Economic Literature*, **34**(3) 1264–1292.
- He, H. and R.S. Pindyck. 1992. Investments in Flexible Production Capacity. *Journal of Economic Dynamics and Control*, **16**, 575–599.
- Hopenhayn, Hugo and Richard Rogerson. 1993. Job Turnover and Policy Evaluation: A General Equilibrium Analysis. *Journal of Political Economy*, **101**(5), 915–938.
- Jovanovic, Boyan. 1982. Selection and the Evolution of Industry. *Econometrica*, **50**(2), 649–670.
- Lucas, Robert E. 1967. Adjustment Costs and the Theory of Supply. *Journal of Political Economy*, **75**(4), 321–334.
- Lucas, Robert E. 1978. On the Size Distribution of Business Firms. *Bell Journal of Economics*, **9**(2), 508–523.
- Milgrom, Paul and John Roberts. 1990. The Economics of Modern Manufacturing: Technology, Strategy, and Organization. *American Economic Review*, **80**(3), 511–528.
- Mortensen, Dale T. and Christopher A. Pissarides. 1994. Job Creation and Destruction in the Theory of Unemployment. *Review of Economic Studies*, **61**, 397–415.
- Sargent, Thomas J. 1987. *Dynamic Macroeconomic Theory*, Harvard University Press, Cambridge, Mass.
- Simon, Herbert A. 1991. Organizations and Markets. *Journal of Economic Perspectives*, **5**(2), 25–44.
- Stigler, George. 1939. Production and Distribution in the Short Run. *Journal of Political Economy*, **47**(3), 305–327.
- Tauchen, George E. 1984. Finite State Markov–Chain Approximations to Univariate and Vector Autoregressions. Working Paper #84–09, Duke University Department of Economics.

Table 1: Gross Job Flows in the Manufacturing Sector by Plant Type

Statistic	All Plants	Plant Age	
		Young	Old
$E(\text{creation})$	5.39 (0.16)	7.52 (0.33)	4.54 (0.12)
$E(\text{destruction})$	5.72 (0.28)	6.56 (0.32)	5.24 (0.27)
$E(\text{growth})$	-0.33 (0.36)	0.95 (0.46)	-0.70 (0.33)
$E(\text{reallocation})$	11.11 (0.28)	14.08 (0.45)	9.77 (0.26)
$S(\text{creation})$	0.85 (0.06)	1.66 (0.19)	0.71 (0.06)
$S(\text{destruction})$	1.45 (0.21)	1.65 (0.20)	1.43 (0.20)
$S(\text{growth})$	1.90 (0.23)	2.32 (0.25)	1.81 (0.24)
$S(\text{reallocation})$	1.42 (0.16)	2.26 (0.20)	1.35 (0.15)
$\frac{V(\text{destruction})}{V(\text{creations})}$	2.89 (0.87)	0.99 (0.37)	4.06 (1.12)
$\rho(\text{creation,destruction})$	-0.33 (0.10)	-0.07 (0.13)	-0.36 (0.09)
$\rho(\text{creation,growth})$	0.70 (0.05)	0.73 (0.06)	0.68 (0.05)
$\rho(\text{destruction,growth})$	-0.91 (0.03)	-0.73 (0.09)	-0.93 (0.02)
$\rho(\text{reallocation,growth})$	-0.51 (0.12)	-0.002 (0.19)	-0.63 (0.09)

Notes: $V(x)$ denotes the variance of variable x , $S(x)$ denotes the standard deviation of variable x in percent, $\rho(x, y)$ denotes the correlation between variable x and variable y , and $E(x)$ the mean of variable x . Numbers in parentheses are standard deviations, computed using the procedure described in Christiano and Eichenbaum (1992). For estimation of the relevant zero-frequency spectral density, a Bartlett window truncated at lag three was used. See the text for variable definitions and the data sources.

Table 2: Differences in Gross Job Flows by Plant Age

Test	$E(\text{realloc.})$	$\frac{V(\text{destruction})}{V(\text{creation})}$	$V(\text{growth})$	$\rho(\text{creation, destruction})$	$\rho(\text{realloc., growth})$
Statistic	156.1	10.0	11.1	11.0	16.9
<i>p</i> -value	(0)	(0.002)	(0.001)	(0.001)	(0.0001)

Notes: $V(x)$ denotes the variance of variable x , $\rho(x, y)$ denotes the correlation between variable x and variable y , and $E(x)$ the mean of variable x . Table entries not in parenthesis are values of the Wald-type test statistic for the null hypothesis that the statistic indicated in the column heading is identical for young and old plants. The test statistics are distributed $\chi^2(1)$. The number in parenthesis is the probability that a chi-square random variable with one degree of freedom exceeds the reported value of the associated test statistic.

Table 3: Baseline Parameter Values

Parameter	Value
α	0.85
β	$1.05^{-1/4}$
δ	0.0083
μ_m	-0.013
σ_m	0.062
μ_s	-0.152
σ_s	0.007
μ_p	-0.023
σ_p	0.214
τ_c	0.5
τ_d	0.5
σ_W	0.021

Table 4: Gross Job Flows in the Baseline Model

Statistic	All Plants	Young	Old
$E(\text{creation})$	5.15	8.86	4.62
$E(\text{destruction})$	5.15	5.14	5.16
$E(\text{growth})$	0.00	3.71	-0.53
$E(\text{reallocation})$	10.31	14.00	9.78
$S(\text{creation})$	0.84	1.57	0.74
$S(\text{destruction})$	1.06	1.67	0.99
$S(\text{growth})$	1.88	3.11	1.72
$S(\text{reallocation})$	0.32	0.91	0.29
$\frac{V(\text{destruction})}{V(\text{creations})}$	1.61	1.12	1.79
$\rho(\text{creation,destruction})$	-0.97	-0.84	-0.99
$\rho(\text{creation,growth})$	0.99	0.96	1.00
$\rho(\text{destruction,growth})$	-0.99	-0.96	-1.00
$\rho(\text{reallocation,growth})$	-0.70	-0.11	-0.86

Notes: $V(x)$ denotes the variance of variable x , $S(x)$ denotes the standard deviation of variable x in percent, $\rho(x, y)$ denotes the correlation between variable x and variable y , and $E(x)$ the mean of variable x .

Table 5: Model Policy Statistics

Statistic	Baseline Model	Cycle of		No Costs of	
		8 Quarters	20 Quarters	Creation	Destruction
$E(\ln y^s(W))$	-1.08	-1.08	-1.08	-1.08	-1.08
$E(\ln y^a(W))$	-1.39	-1.39	-1.39	-1.73	-0.73
$E(\ln \underline{z}(W))$	0.26	0.26	0.27	0.26	0.26
$E(\ln \underline{y}(W))$	-2.06	-2.06	-2.06	-1.73	-1.79
$E(\ln \bar{y}(W))$	-0.57	-0.57	-0.57	-0.66	-0.73
$E(\Pr[z_t \geq \underline{z}(W)])$	0.02	0.02	0.02	0.02	0.02
$S(\ln y^s(W))$	14.73	14.17	14.17	14.73	14.73
$S(\ln y^a(W))$	2.76	4.09	8.62	5.02	7.02
$S(\ln \underline{z}(W))$	0.04	0.89	1.00	0.02	0.07
$S(\ln \underline{y}(W))$	3.76	4.30	8.27	5.01	4.99
$S(\ln \bar{y}(W))$	6.16	6.79	10.30	7.08	7.02
$S(\Pr[z_t \geq \underline{z}(W)])$	0.01	0.20	0.22	0.01	0.02

Notes: $S(x)$ denotes the standard deviation of variable x in percent, and $E(x)$ denotes the mean of variable x .

Table 6: Gross Job Flows in the Model with Deterministic Cycles

Statistic	8 Quarter Cycle			20 Quarter Cycle		
	All Plants	Young	Old	All Plants	Young	Old
$E(\text{creation})$	5.11	8.62	4.61	5.11	8.46	4.63
$E(\text{destruction})$	5.12	4.91	5.15	5.11	4.72	5.16
$E(\text{growth})$	0.00	3.71	-0.54	0.00	3.73	-0.53
$E(\text{reallocation})$	10.23	13.53	9.75	10.21	13.18	9.79
$S(\text{creation})$	0.74	1.21	0.68	0.98	1.42	0.91
$S(\text{destruction})$	0.90	1.25	0.85	0.94	1.00	0.93
$S(\text{growth})$	1.62	2.39	1.51	1.89	2.39	1.82
$S(\text{reallocation})$	0.33	0.61	0.30	0.31	0.58	0.30
$\frac{V(\text{destruction})}{V(\text{creations})}$	1.47	1.07	1.58	0.94	0.50	1.05
$\rho(\text{creation,destruction})$	-0.94	-0.88	-0.95	-0.95	-0.94	-0.95
$\rho(\text{creation,growth})$	0.98	0.97	0.98	0.99	0.99	0.99
$\rho(\text{destruction,growth})$	-0.99	-0.97	-0.99	-0.99	-0.98	-0.99
$\rho(\text{reallocation,growth})$	-0.49	-0.07	-0.59	0.10	0.72	-0.08

Notes: $V(x)$ denotes the variance of variable x , $S(x)$ denotes the standard deviation of variable x in percent, $\rho(x, y)$ denotes the correlation between variable x and variable y , and $E(x)$ the mean of variable x .

Table 7: Gross Job Flows in the Model with Asymmetric Adjustment Costs

Statistic	No Creation Costs			No Destruction Costs		
	All Plants	Young	Old	All Plants	Young	Old
$E(\text{creation})$	6.21	9.53	5.77	6.76	11.77	5.99
$E(\text{destruction})$	6.21	5.62	6.29	6.76	8.21	6.54
$E(\text{growth})$	0.00	3.90	-0.52	0.00	3.56	-0.55
$E(\text{reallocation})$	12.43	15.15	12.07	13.52	19.98	12.52
$S(\text{creation})$	1.33	2.14	1.23	1.31	1.95	1.22
$S(\text{destruction})$	1.45	1.82	1.41	1.59	2.44	1.47
$S(\text{growth})$	2.76	3.86	2.62	2.88	4.30	2.67
$S(\text{reallocation})$	0.34	0.94	0.32	0.42	1.00	0.36
$\frac{V(\text{destruction})}{V(\text{creations})}$	1.19	0.72	1.33	1.45	1.56	1.46
$\rho(\text{creation,destruction})$	-0.98	-0.90	-0.98	-0.98	-0.92	-0.98
$\rho(\text{creation,growth})$	0.99	0.98	0.99	0.99	0.98	0.99
$\rho(\text{destruction,growth})$	-0.99	-0.97	-1.00	-1.00	-0.98	-1.00
$\rho(\text{reallocation,growth})$	-0.38	0.34	-0.58	-0.66	-0.50	-0.70

Notes: $V(x)$ denotes the variance of variable x , $S(x)$ denotes the standard deviation of variable x in percent, $\rho(x, y)$ denotes the correlation between variable x and variable y , and $E(x)$ the mean of variable x .