Entry, Exit, and Plant-level Dynamics over the Business Cycle*

Yoonsoo Lee
Federal Reserve Bank of Cleveland
yoonsoo.lee@clev.frb.org

Toshihiko Mukoyama
Department of Economics
University of Virginia
and
CIREQ
tm5hs@virginia.edu

December 2007

Abstract

This paper analyzes the implications of plant-level dynamics over the business cycle. We first document basic patterns of entry and exit of U.S. manufacturing plants, in terms of employment and productivity, between 1972 and 1997. We show how entry and exit patterns vary during the business cycle, and that the cyclical pattern of entry is very different from the cyclical pattern of exit. Second, we build a general equilibrium model of plant entry, exit, and employment and compare its predictions to the data. In our model, plants enter and exit endogenously, and the size and productivity of entering and exiting plants are also determined endogenously. Finally, we explore the policy implications of the model. Imposing a firing tax that is constant over time can destabilize the economy by causing fluctuations in the entry rate. Entry subsidies are found to be effective in stabilizing the entry rate and output.

Keywords: plant-level dynamics, entry and exit, business cycles

JEL Classifications: E23, E32, L11, L60

*We thank Gian Luca Clementi, Tim Dunne, Marcelo Veracierto, and the seminar participants at the Federal Reserve Bank of Cleveland, the Federal Reserve Bank of New York, the Korean Econometric Society, the Northeast Ohio Economics workshop, the Rochester Wegmans conference 2006, the Third New York/Philadelphia Workshop on Quantitative Macroeconomics, the University of Tokyo, and Yonsei University for comments and suggestions. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of Cleveland or the Federal Reserve System. The research in this paper was conducted while the authors were Special Sworn Status researchers of the U.S. Census Bureau at the Michigan Census Research Data Center. Research results and conclusions expressed are those of the authors and do not necessarily reflect the views of the Census Bureau. This paper has been screened to insure that no confidential data are revealed. Support for this research at the Michigan RDC from NSF (awards no. SES-0004322 and ITR-0427889) is gratefully acknowledged.

1 Introduction

A growing number of recent studies using plant-level data find a large degree of heterogeneity in the size, productivity, and growth patterns of manufacturing plants.¹ In this paper, we explore the implications of this plant-level heterogeneity for macroeconomic dynamics and policy.

We first document the heterogeneity of U.S. manufacturing plants, using the Annual Survey of Manufactures (ASM) from the U.S. Census Bureau from 1972–1997. While previous studies on the entry and exit of producers document considerable fluctuations in entry and exit rates (e.g., Chatterjee and Cooper, 1993; Campbell, 1998), relatively little is known about how the characteristics of entering and exiting plants vary over the business cycle. We document patterns of entry and exit over the business cycle in terms of entry and exit rates, employment, and productivity. We find that entry rates differ significantly in booms and recessions. In particular, plants entering in booms are very different in terms of productivity and employment from those entering in recessions, but plants that exit in booms and recessions do not differ much. For example, the average employment of entering plants in recessions and their relative size (compared to incumbents) is larger than entering plants in booms. Moreover, the relative productivity of entering plants is also higher in recessions than in booms. Such differences are relatively small for plants exiting in booms or recessions.

It has long been argued that recessions have “cleansing” effects: low-productivity plants are scrapped during recessions, enhancing aggregate efficiency. Many recent papers have suggested alternative views about what is happening in recessions. For example, analyzing a model of creation and destruction of production units, Caballero and Hammour (1994) argue that low-productivity firms can be “insulated” from recessions because fewer new plants are created during recessions. Barlevy (2002) considers a model of on-the-job search and shows that recessions may reduce aggregate efficiency by discouraging the reallocation of workers. In a more recent study, Caballero and Hammour (2005) provide evidence that recessions

¹See Bartelsman and Doms (2000) for a review of the literature.

reduce the amount of cumulative reallocation in the economy.

Focusing on permanent shutdowns, we do not find strong effects of cleansing during recessions. Overall, annual exit rates are similar across booms and recessions. Furthermore, exiting plants in recessions are not very different from those in booms in terms of employment or productivity. Our finding suggests that recessions do not necessarily cause productive plants—those that could have survived in good times—to shut down in large numbers. Rather, strongly procyclical entry rates suggest that the “insulation” effect at the entry margin predominates. In contrast to exiting plants, entrants vary substantially in terms of average size and productivity during the business cycle. Only highly productive plants enter and begin production during recessions. While previous studies on the effects of recessions have focused on selection at the exit (or destruction) margin, our finding suggests that selection at the entry (or creation) margin may be more important.

Based on our observations of heterogeneous plant-level behavior during business cycles, we build a dynamic general equilibrium model and use it to analyze the effects of macroeconomic policies. Our model extends the standard general equilibrium firm dynamics model of Hopenhayn and Rogerson (1993) by incorporating aggregate productivity shocks. To account for entrants’ productivity differences in booms and recessions, we also incorporate self-selection during the entry process. We find that the model performs well in replicating the cyclical-ity of entry and exit rates. However, it is necessary to assume cyclical-ity of entry costs to account for the cyclical patterns of selection in the entry process. When entry costs are acyclical, the effect of aggregate productivity shocks on entrants’ size and productivity is almost completely offset by the general equilibrium effect (cyclical changes in wages). That is, while the increase in aggregate productivity makes entry attractive even for low-productivity plants, this effect is completely washed out by the increase in wages. When we assume a sufficient amount of cyclical-ity in entry costs, the model can account for the observed cyclical plant-level dynamics.

Some other recent papers extend a similar style of a firm dynamics model to incorporate

business cycles. The main difference between this paper’s model and those of the existing literature is that in ours both entry and exit are endogenously determined. Veracierto (2002, 2004) was among the first to incorporate aggregate productivity shocks to the Hopenhayn-Rogerson style model.² Veracierto (2002, 2004), however, assumes exogenous entry and exit, and therefore his models are not directly suitable to explain the cyclical patterns of entry and exit observed in the data. Comin and Gertler (2004) build a model with endogenous innovation and technology adoption over the “medium-term” business cycle. In their intermediate-goods sector, entry (innovation and adoption) is endogenous but exit (obsolescence of technology) is exogenous. In their final-goods sector, profit is a function of the *total* number of firms, and the zero-profit condition pins down the total number of firms in equilibrium. Therefore, their final-goods sector does not have separate gross flows of entry and exit.³ A recent work by Bilbiie, Ghironi, and Melitz (2006) also constructs a model of business cycles with endogenous entry, but their model features firms with homogeneous technology and exogenous exit.

Samaniego (2006) constructs a general equilibrium model of industry dynamics with endogenous entry and exit. Instead of solving a model with aggregate shocks, he characterizes the (deterministic) transition path after the change in aggregate productivity. He finds that both entry and exit respond very little to the change in aggregate productivity. This contrasts with our result—in our model, the entry rate responds strongly to an aggregate productivity shock. A large part of the difference comes from the specification of the entry cost. Samaniego assumes that the marginal cost of building a plant is increasing in the number of existing plants, while we assume a constant marginal cost. Our empirical evidence suggests that the constant marginal cost specification achieves an outcome closer to the behavior of U.S. manufacturing plants.

²As in Veracierto (2001), the models in Veracierto (2002, 2004) also incorporate saving and the capital stock, which are absent in our model.

³Similarly, in other papers that employ monopolistic competition models to explore markup dynamics (e.g., Chatterjee and Cooper, 1993; Devereux, Head, and Lapham, 1996; and Jaimovich, 2007), the zero-profit condition in each period determines the total number of firms, so that it is not possible to address the issues of gross entry and exit separately.

Utilizing our model, we conduct four policy experiments. The first two analyze the firing tax. As in Hopenhayn and Rogerson (1993), a constant firing tax reduces the average level of employment. Interestingly, a constant firing tax *increases* the variance of output. This contrasts with the stabilizing effects of firing taxes in models with exogenous entry and exit (e.g., Veracierto, 2004)—thus, it is in fact important for policy analysis to model entry and exit behavior endogenously. The reason a constant firing tax is destabilizing is that the entry rate becomes more volatile when this tax is imposed. Given the mean-reverting nature of the idiosyncratic productivity process, the firing tax has a greater impact on large plants because they are more likely to contract in the near future. Since entrants are larger during recessions than during booms, the effect of firing taxes on entrants is stronger during recessions. Therefore, the difference between entry rates during booms and entry rates during recessions widens due to the firing tax.

We next consider a countercyclical firing tax that is intended to reduce the amount of firing during recessions. When a firing tax is imposed only during recessions, job destruction rates during recessions are reduced. However, the variance of output increases dramatically.

Experiments three and four analyze entry subsidies. In both experiments, subsidizing entry costs during recessions stabilizes both the entry rate and aggregate output.

The paper is organized as follows. In the next section, we document the empirical facts on entry, exit, and employment in U.S. manufacturing. In Section 3, we build a general equilibrium model of plant-level dynamics and match it to the data. In Section 4, we conduct the policy analysis. Section 5 concludes.

2 Empirical evidence on employment and productivity dynamics

2.1 Measurement and data

We use the ASM portion (from 1972 through 1997) of the Longitudinal Research Database (LRD), which is constructed by the U.S. Census Bureau, to analyze the behavior of plants

during business cycles. Many recent theoretical studies on plant-level dynamics are based on the evidence provided by Dunne, Roberts, and Samuelson (1988, 1989a, 1989b). They utilize the Census of Manufacturers (CM) dataset, which is a part of the LRD. The CM is conducted for the universe of U.S. manufacturing plants, and the evidence from the CM has been used to calibrate stationary equilibrium models describing the entry, exit, and employment dynamics of U.S. plants (e.g., Hopenhayn and Rogerson, 1993). However, because the CM is conducted every five years, it is not suitable for describing plant-level behavior over the business cycle. The ASM, conducted annually for non-census years, overcomes this issue. The ASM utilizes a probability-based sample of plants drawn from the universe of plants identified by the CM. We use ASM sample weights so that the sample is representative of the entire U.S. manufacturing sector.⁴

In this study, entering plants are new plants, which appear in the ASM or CM for the first time with at least one employee (birth). Similarly, exiting plants include only permanent shutdowns (death). We do not include temporary exit and re-entry of plants, in order to exclude possible spurious entries and exits in the ASM panels (See Appendix A for more detail about measuring entry and exit). As discussed in detail in Davis, Haltiwanger, and Schuh (1996), samples in the ASM panels are rotated every five years. Only large “certainty” plants are continuously observed across different ASM panels. In order to avoid measurement errors in entry and exit that are caused by the panel rotations, the results reported in this paper exclude entries and exits measured between two different ASM panels, namely for the years 1973-74, 1978-79, 1983-84, 1988-89, and 1993-94.

In addition to employment dynamics, we also examine the extent to which the productivity of entering and exiting plants varies over the business cycle. The ASM contains data on material inputs, output, and the capital stock in addition to employment at each plant. We construct various measures of productivity.

First we look at total factor productivity (TFP), as in the standard macroeconomic

⁴See Davis, Haltiwanger, and Schuh (1996) for details about the data.

growth-accounting analysis. Our plant-level TFP closely follows Baily, Hulten, and Campbell (1992).⁵ Assuming that the production function is $y_t = s_t k_t^{\alpha_k} n_t^{\alpha_n} m_t^{\alpha_m}$, where y_t is real gross output, s_t is TFP, k_t is real capital stock, n_t is labor input, and m_t is real material inputs, the TFP (s_t) can be measured from

$$\ln(s_t) = \ln(y_t) - \alpha_k \ln(k_t) - \alpha_n \ln(n_t) - \alpha_m \ln(m_t). \quad (1)$$

We measure factor elasticities using 4-digit industry-level revenue shares. Real capital stocks are obtained from the perpetual inventory method. Output and material inputs are measured in 1987 constant dollars using deflators from the NBER manufacturing productivity dataset. Labor input is measured as total hours for production and non-production workers following Baily, Hulten, and Campbell (1992). Appendix A describes the construction of our measures in more detail.

While this measure of TFP follows the practice used in the literature for measuring plant-level TFP, it may be subject to measurement errors of the capital stock. To avoid this issue, we consider the following specification, $y_t = s_t n_t^\theta$. Now we measure y_t by value added, rather than output. Then s_t can be measured from

$$\ln(s_t) = \ln(y_t) - \theta \ln(n_t). \quad (2)$$

We prefer this specification for two reasons. First, the measurement of the capital stock is likely to be less reliable than employment, and thus the measurement based on (1) may suffer from large measurement errors. Second, our model in Sections 3 and 4 does not include capital, and the measure of productivity by (2) directly matches our model specification. The value of s_t in the second measure represents both TFP's and the capital stock's contributions to production.

⁵Without a proper measure of prices for individual plants, it is not possible to measure total factor productivity at the plant level. While we call this measure TFP, it is actually real revenue per unit input and reflects within-industry price variation. See Foster, Haltiwanger, and Syverson (2007) for possible issues involved in using revenue-based productivity measures.

	Continuing	Entering	Exiting
Average size	87.5	50.3	35.3
Relative size (to the industry average of continuing plants)	—	0.60	0.49

Table 1: Average size of plants

	Entering	Exiting
TFP based on (1)	0.96	0.86
Labor productivity (using employment)	1.00	0.92
Labor productivity (using hours)	0.98	0.91
TFP based on (2)	0.75	0.65

Table 2: Average plant productivity

2.2 Results

2.2.1 Average statistics of employment and productivity

First, we document the characteristics of entering and exiting plants in terms of employment and productivity. Those statistics are used to calibrate the steady state of the model. Table 1 documents the average size of the plants, in terms of the number of workers. Entering plants (time- t size of the plants that entered between time $t - 1$ and t) and exiting plants (time- $t - 1$ size of the plants which exited between time $t - 1$ and t) are much smaller than continuing plants (time- t size of the plants which survived from time $t - 1$ to time t). Table 1 also reports the relative size of entering and exiting plants. The relative size of an entering (exiting) plant is obtained by dividing the size of the entrant by the average size of continuing plants in the same four-digit industry.⁶ Entering plants are 40 percent smaller than continuing plants in the same four-digit SIC industry, while exiting plants are about half of the size of continuing plants in the same industry.

These differences in size are partly explained by differences in productivity. Table 2 shows various measures of relative productivity. This finding is one of the contributions of

⁶By dividing by the average size of continuing plants in the same four-digit industry, we control for the effects of changes in the industrial composition of entrants over the cycle, as well as differences in plant size across industries.

our paper, since in the previous literature direct measures of productivity were not available at an annual frequency. Each cell in Table 2 represents the relative productivity (compared to the four-digit SIC industry average of continuing plants) of entering and exiting plants. Two properties are found consistently across different productivity measures. First, entering and exiting plants are less productive than continuing plants (except for one case).⁷ Second, exiting plants are less productive than entering plants. These findings are consistent with the pattern of employment size in Table 1, provided that a productive plant employs more workers.

The first row in Table 2 is the TFP, based on (1). The second and third rows are measures of labor productivity (output divided by labor input). The second row measures labor input by employment, and the third row measures labor input by hours. The last row is the productivity measure based on equation (2). Here, instead of using (2) directly, we control for industry heterogeneity in labor shares by postulating the production function $y_t = s_t n_t^{\theta_I}$. We obtain s_t by calculating $\ln(s_t) = \ln(y_t) - \theta_I \ln(n_t)$. θ_I is obtained from the four-digit SIC industry-level labor share. As is discussed above, the advantage of the measure based on (2) is that the measurements of output and employment are relatively more reliable than capital and material inputs. Moreover, we use this exact form of production function in Sections 3 and 4. Therefore, we mainly utilize this last measure of productivity in calibrating the model.

In measuring the process of plant-level employment, we estimate an AR(1) process

$$\ln(n_{t+1}) = \alpha + \rho \ln(n_t) + \varepsilon_{t+1}, \quad (3)$$

where $\varepsilon_{t+1} \sim N(0, \sigma^2)$. We report estimates of the AR(1) process of employment dynamics in Table 3. Since only large plants are observed across two different ASM panels, the results

⁷Differences in plant-level productivity must be interpreted with caution. Because plant-level prices are not observed, our revenue-based productivity measures reflect price or demand variation within an industry in addition to differences in technical efficiency. In studies focusing on a small number of industries where producer-level prices and quantities are observed separately, Foster, Haltiwanger, and Syverson (2007) argue that the true technological productivity of entrants may be understated when traditional revenue-based measures are used because new plants have lower prices than incumbents.

	ρ	σ
Without any control	0.97	0.40
With year dummy	0.97	0.40
With year dummy, age, industry	0.97	0.39

Table 3: Employment dynamics

	1 - 19	20 - 49	50 - 99	100 - 249	250 +
plants	0.457	0.239	0.131	0.106	0.067
employment	0.049	0.090	0.109	0.194	0.559
hiring	0.082	0.120	0.131	0.207	0.460
firing	0.142	0.143	0.135	0.195	0.385
exit rate	0.080	0.039	0.031	0.025	0.015
plants, by age					
Young	0.073	0.016	0.007	0.005	0.002
Middle	0.184	0.083	0.038	0.024	0.009
Old	0.200	0.140	0.087	0.078	0.055
employment, by age					
Young	0.006	0.006	0.006	0.008	0.014
Middle	0.019	0.031	0.031	0.043	0.064
Old	0.023	0.053	0.072	0.142	0.480

Table 4: Size distribution of plants

exclude samples between two ASM panels in a way that is similar to our calculation of entry and exit statistics. The estimated value of ρ is very close to one—the employment process is very persistent. The results are similar when we correct the selection bias using the maximum likelihood estimates from the Heckman selection model.⁸

Table 4 describes the distribution of plant size (employment). It also reports employment shares, hiring shares, and firing shares for each size class. It shows that employment is skewed

⁸We also calculated this process for the productivity measures in Table 2 and found a smaller ρ and a bigger σ for the TFP based on (2). However, we are concerned with the bias due to measurement errors, which will underestimate the coefficient ρ (i.e., attenuation bias) and overestimate the variance of the residual, σ . Given that we do not know the magnitude of the bias, we do not use the estimate for the productivity process in the model. We utilize the employment statistics extensively, since we believe that they are least subject to the measurement errors. See also Ábrahám and White (2007) for detailed discussions on AR(1) process estimates of plant-level productivity.

towards large plants. Hiring and firing are also concentrated in large plants. While exit rates are higher in smaller plants, some large plants also exit. As we will discuss later, previous models such as Hopenhayn and Rogerson (1993) cannot explain this particular phenomenon. Because productivity and size have a one-to-one relationship, very large plants have high productivity levels and do not exit in their model. The distribution of plant size and the employment shares are also presented for each age category.⁹ This presentation shows that young plants tend to be small.

2.2.2 Business cycle patterns

Here we characterize how entry and exit, employment, and productivity differ in booms and recessions. When considering business cycles, we divide the sample years into two categories, good and bad, based on the growth rate of manufacturing output. If the growth rate of manufacturing output from year $t - 1$ to t is above average, we call year t a good year; if it is below average, we call year t a bad year.¹⁰ The reason we base our distinction on the growth rate rather than the level is twofold. First, the division based on the (HP-filtered) level does not match the conventional boom-recession division. For example, based on the level criterion, 1990 (the only year in the 1990s where more than half of one year is recorded as a “contraction,” according to NBER business cycle dates) is considered a good year, while most years of the mid-1990s are considered bad. Second, we consider the growth rate to be an important indicator since our analysis stresses the cyclical movement of entry and exit rates, which are more related to the “change” than the “level.”¹¹

⁹The age categories follow Davis, Haltiwanger, and Schuh (1996, p.225). In the ASM, panel rotation makes it impossible to determine the exact age of plants. Roughly speaking, “Young” corresponds to 0–1 years in operation, “Middle” corresponds to 2–10 years, and “Old” corresponds to more than 11 years.

¹⁰Good years are 72, 73, 76, 77, 78, 83, 87, 88, 92, 93, 94, 95, 96, 97 and bad years are 74, 75, 79, 80, 81, 82, 84, 85, 86, 89, 90, 91.

¹¹Later we utilize a business-cycle model with two levels of aggregate productivity. It turns out that significant positive/negative growth is realized only when the state switches between different levels of aggregate productivity. We mainly report the result based on the different aggregate productivities, but we will also report the results based on these “switching states” in Appendix D. The results are qualitatively similar. The quantitative differences between booms and recessions are larger when we consider only the switching periods. We view this as a result of having only two levels of productivity—the switch between states is more extreme compared to reality.

	Good	Bad	Total average
Entry (birth)	8.1%	3.4%	6.2%
Exit (death)	5.8%	5.1%	5.5%

Table 5: Entry and exit rates

	Good	Bad	Total average
Job creation from startups	1.76	1.21	1.52
Job creation from continuers	8.20	6.48	7.44
Job destruction from shutdowns	2.52	2.27	2.41
Job destruction from continuers	6.72	8.74	7.61

Table 6: Job creation and job destruction

Table 5 displays the entry and exit rates of plants in good and bad times. The entry rate is much higher in booms than in recessions. In contrast, exit rates are similar throughout the business cycle. The simple correlation between entry rates and the annual growth rates of manufacturing output is 0.413, while the same statistic for the exit rates is 0.240. Our finding that the plant entry rate is more cyclical than the plant exit rate contrasts with the previous findings on the cyclical pattern of job creation and destruction (e.g., Davis, Haltiwanger, and Schuh, 1996; Campbell, 1998). Focusing on employment flows, Davis, Haltiwanger, and Schuh (1996) find that the job destruction rate is more cyclical than the job creation rate. Using the aggregate job flows data from Davis, Haltiwanger, and Schuh (1996), Campbell (1998) finds that employment-weighted entry rates are procyclical, while employment-weighted exit rates are countercyclical. This difference can be explained by the following three facts. First, the sample period between those studies and ours is different. These previous studies cover 1972–1988 and our study includes more recent cohorts of the data and covers 1972–1997. In a recent study, Foster, Haltiwanger, and Kim (2006) find that job creation and destruction trends in the latest panel (1994–1998) are very different than the earlier period. Second, we focus on only first-time plant openings (births) and permanent shutdowns (deaths), whereas Davis, Haltiwanger, and Schuh (1996) and Campbell (1998)

	Good	Bad	Total average
Average size	85.4	89.5	87.5

Table 7: Average employment of continuing plants

	Good	Bad	Average
Average size, entering	45.1	59.2	50.3
Average size, exiting	34.9	35.9	35.3
Relative size, entering	0.53	0.70	0.60
Relative size, exiting	0.50	0.46	0.49

Table 8: Average employment of entering and exiting plants

also count reopened plants and temporary shutdowns as entries and exits. Finally, while our study focuses on the entry and exit margin, the time series pattern for gross job flows can be driven by job creation and destruction from continuing plants. We also analyze job creation (destruction) due to startups (shutdowns) and job creation (destruction) for continuing plants separately. Table 6 presents job creation and job destruction statistics calculated from the published job flows data for our sample period (1972–1997).¹² We find that job creation from startups is much higher during booms, while job destruction from shutdowns is only slightly higher. The simple correlation between the job creation rate due to startups and the percentage change in manufacturing output (annual) is 0.368, while the simple correlation between the job destruction rate due to shutdowns and the percentage change in manufacturing output is -0.006 . Although the job destruction rate for continuing plants is higher during recessions, we do not see the “cleansing” effect in the exit margin during recessions. This finding suggests that adjustment over the business cycle at the entry margin may be more important.

Tables 7 and 8 describe average plant size during booms and recessions. In general, the

¹²The job flows data are available from the webpage of John Haltiwanger, <http://www.econ.umd.edu/~haltiwan/download.htm>. While quarterly job flows data are also available, we choose to use annual job creation and destruction rates because our statistics are measured in annual frequency. Annual job flows data are measured based on all employees (total employment), but quarterly job flows data include only production workers.

	Good	Bad	Total average
Relative productivity, entering	0.69	0.85	0.75
Relative productivity, exiting	0.65	0.65	0.65

Table 9: Productivity relative to continuing plants

average size is larger during recessions. Exiting plants are of similar size across booms and recessions, but the average size of entering plants dramatically changes during recessions.¹³ Compared to entering plants in booms, entering plants in recessions start with about 30% more workers. The simple correlation between the relative size of exiting plants and the percentage change in manufacturing output is 0.034, while the simple correlation between the relative size of entering plants and the percentage change in manufacturing output is -0.241 .

Relative productivity, presented in Table 9, exhibits a similar pattern. The relative productivity of exiting plants is similar across booms and recessions, while the relative productivity of entering plants is substantially different in the two phases of the cycle. The relative productivity of entering plants in recessions are about 20% higher than that of entering plants in booms.¹⁴ The simple correlation between the relative productivity of exiting plants and the percentage change in manufacturing output is -0.054 , while the simple correlation between the relative size of entering plants and the percentage change in manufacturing output is -0.232 .

¹³There were some outliers among entering plants in 1980. Because dropping a few outliers would cause some disclosure issues, we chose to drop the whole year when calculating the average size and productivity of entering plants in Tables 8 and 9. Because those outliers have substantially higher productivity levels, including them results in a much greater difference in entrants' productivity in booms and recessions, adding support to our finding. Although the results on average employment do not vary much with or without the outliers, we also dropped this year in Table 8 for consistency. Since the statistics for exiting plants are not affected by the outliers, we include the 1980 observations in the calculation.

¹⁴Since we use revenue-based productivity measure, a caution is needed in interpreting the finding of higher productivity for entering plants in recessions. While they can indeed be more productive than those entering in booms, it is also possible that they have about the same productivity levels but just produce goods with higher price. Given the substantial difference in the size, we believe the differences in the revenue-based productivity measure would reflect differences in true productivity.

3 Model

In this section, we set up a dynamic general equilibrium model of plant employment, entry, and exit. We base our model on Hopenhayn and Rogerson (1993), departing significantly from their baseline model in four respects.

First, we add aggregate shocks to the economy. This ingredient is essential in analyzing the business cycle implications of the model.

Second, we assume that there is a positive (and stochastic) value of exiting. This modification is necessary for the model to match the exit pattern observed in the data. In Hopenhayn and Rogerson’s model (in the benchmark case), plants compare the value of staying with the value of exiting, which is zero. In their model, there is a threshold productivity, s^* : when the productivity is higher than s^* , the plant stays; if the productivity is lower than s^* , the plant exits. Since productivity and employment have a one-to-one relationship, it means that the exit rate is 100% for plants that are smaller than a certain threshold and zero for plants that are larger than the threshold. The data do not exhibit this type of pattern: even large plants with more than 250 employees have an exit rate of over 1% (see Table 4 in the previous section). Furthermore, with Hopenhayn and Rogerson’s formulation, an annual exit rate of 5.5% (see Table 5) implies that only very unproductive and small plants exit, which is at odds with the employment and productivity evidence in Tables 8 and 9.

Third, we consider entry in two steps—to enter, one first has to pay some cost and come up with an “idea.” Then, after observing the quality of the idea, one decides whether to pay an additional cost to actually enter the market. This “two-step” process introduces the endogenous selection of the entering plants.¹⁵ In the data, we observe that the productivity of entering plants is very different across booms and recessions (see Table 9). In Hopenhayn and Rogerson’s model, the entering plants receive a productivity draw after the decision to enter, so that the productivity distribution of entrants in the economy is always the same and is given exogenously. Our modification makes it possible for the productivity distribution of

¹⁵Melitz (2003) employs a similar selection process in entry.

entrants to vary endogenously across booms and recessions.

Finally, we introduce the cost of adjusting employment. The estimation of the employment process by Cooper, Haltiwanger, and Willis (2004) strongly indicates that there are important adjustment costs in the employment process. It turns out that adding a moderate amount of adjustment cost dramatically changes the amount of job reallocation.

3.1 Plants

The model consists of two kinds of entities: plants and consumers. Plants use labor to produce output. Consumers own plants, supply labor, and consume. There is only one type of good, which is used for entry costs and consumption; we use it as the numeraire. In our model, the only price we have to keep track of is the wage of the workers. We assume that the plants have to pay adjustment costs and the firing tax when labor input is adjusted. The specifics of the adjustment costs and the firing tax are explained later.

Here, we describe the decision of the plants. First, we outline the behavior of the incumbent plants. Then we illustrate the entrant's behavior.

The timing of events for an incumbent plant at period t is as follows. In the beginning of the period, all plants observe the current aggregate state, z_t . An incumbent plant starts a period with the individual state (s_{t-1}, n_{t-1}) . s_{t-1} is the individual plant's productivity level at period $t - 1$. n_{t-1} is the employment level at period $t - 1$. The value function of a plant at this stage is denoted as $W(s_{t-1}, n_{t-1}; z_t)$. Then, it observes its (stochastic) exit value, x_t . Here, x_t can be interpreted as the scrap value of its capital (and owned land), although we do not explicitly model capital stock or land.¹⁶ After observing the exit value, the plant decides whether to stay or exit. If it exits, it has to pay the firing tax, since it has to adjust the employment level from n_{t-1} to zero. If it decides to stay, it observes this period's individual productivity (idiosyncratic shock), s_t . The value function at this point is denoted as $V^c(s_t, n_{t-1}; z_t)$. Then it decides the amount of employment in the current period, n_t , and

¹⁶The entry cost that is introduced later can be interpreted as (partially sunk) investment in new capital and land.

produces. The production function is $z_t f(n_t, s_t)$, where the function $f(n_t, s_t)$ is increasing and concave in n_t . If $n_t \neq n_{t-1}$, it pays adjustment costs (and a firing tax, if $n_t < n_{t-1}$).¹⁷ This concludes the period.

The timing for entrants is as follows. In the beginning of the period, everyone observes z_t . To enter, the first step is to come up with an “idea.” To come up with an idea, one has to pay c_q and receive a random number q_t (quality of the idea). A large q_t indicates that productivity after the entry is high. We call the people with an idea “potential entrants.” We denote the expected value of having an idea, before knowing q_t as $V^p(z_t)$. We denote the value of a potential entrant after paying c_q and receiving q_t as $V^e(q_t; z_t)$. Given q_t , a potential entrant decides whether to enter. To enter, the entry cost c_e is paid. We interpret c_e as (partially sunk) investment in plants. The potential entrant, therefore, compares $V^e(q_t; z_t)$ and c_e . From here, the decision is the same as for the incumbent, except that the productivity s_t will depend on q_t instead of s_{t-1} . The plant observes s_t (its value function is $V^c(s_t, 0; z_t)$ now), then it decides the employment n_t , pays the adjustment cost, and produces.

An incumbent’s value at the beginning of the period is described by the Bellman equation

$$W(s_{t-1}, n_{t-1}; z_t) = \int \max \langle E_s[V^c(s_t, n_{t-1}; z_t) | s_{t-1}], x_t - g(0, n_{t-1}) \rangle d\xi(x_t).$$

Here, $g(n_t, n_{t-1})$ is the firing tax. In the $\max \langle \cdot, \cdot \rangle$, the plant compares the value of staying (the first term) and exiting (the second term). $E_s[\cdot | s_{t-1}]$ denotes the expectation regarding s_t , conditional on s_{t-1} . We assume that the exit value x_t follows an i.i.d. distribution $\xi(x_t)$, and that the exit value distribution does not vary over the business cycle. As we will see later, our model can match the exit pattern in the data without relying on the cyclical exit values. $E_s[V^c(s_t, n_{t-1}; z_t) | s_{t-1}]$ is the expected value of a continuing plant $V^c(s_t, n_{t-1}; z_t)$, and is calculated as

$$E_s[V^c(s_t, n_{t-1}; z_t) | s_{t-1}] = \int V^c(s_t, n_{t-1}; z_t) d\psi(s_t | s_{t-1}),$$

¹⁷The details of the adjustment cost are explained later.

where

$$V^c(s_t, n_{t-1}; z_t) = \max\langle V^a(s_t, n_{t-1}; z_t), V^n(s_t, n_{t-1}; z_t) \rangle,$$

and $\psi(s_t|s_{t-1})$ is the distribution of s_t given s_{t-1} . Here, $V^a(s_t, n_{t-1}; z_t)$ is the value function when the plant adjusts employment, and $V^n(s_t, n_{t-1}; z_t)$ is the value function when it does not adjust employment.

If the plant decides to adjust employment, the current period profit is

$$\pi^a(s_t, n_{t-1}, n_t; z_t) \equiv \lambda z_t f(n_t, s_t) - w_t n_t - g(n_t, n_{t-1}),$$

where $\lambda < 1$ represents the “disruption cost” type of adjustment cost, emphasized by Cooper, Haltiwanger, and Willis (2004). This represents the cost of slowing down the production process when employment is adjusted. In Cooper, Haltiwanger, and Willis’s (2004) estimation, this cost turns out to be the most important type of adjustment cost in explaining employment dynamics observed at the plant level.

If the plant does not adjust employment, the current period profit is

$$\pi^n(s_t, n_{t-1}; z_t) \equiv z_t f(n_{t-1}, s_t) - w_t n_{t-1}.$$

Therefore,

$$V^a(s_t, n_{t-1}; z_t) = \max_{n_t} \pi^a(s_t, n_{t-1}, n_t; z_t) + \beta E_z[W(s_t, n_t; z_{t+1})|z_t],$$

and

$$V^n(s_t, n_{t-1}; z_t) = \pi^n(s_t, n_{t-1}; z_t) + \beta E_z[W(s_t, n_{t-1}; z_{t+1})|z_t].$$

Here, $E_z[\cdot|z_t]$ takes the expectation regarding z_{t+1} , conditional on z_t .

The entrant’s value function is

$$V^e(q_t; z_t) = \int V^c(s_t, 0; z_t) d\eta(s_t|q_t),$$

where $\eta(s_t|q_t)$ is the distribution of s_t given q_t . Only the potential entrant with high enough q_t will actually enter. There is a threshold value of q_t , q_t^* , which is determined by

$$V^e(q_t^*; z_t) = c_e. \tag{4}$$

A potential entrant will enter if and only if $q_t \geq q_t^*$. A potential entrant's value function is

$$V^p(z_t) = \int \max\langle V^e(q_t; z_t) - c_e, 0 \rangle d\nu(q_t),$$

where $\nu(q_t)$ is the distribution of ideas. We impose a free-entry condition for becoming a potential entrant:

$$V^p(z_t) = c_q. \quad (5)$$

3.2 Consumers

The representative consumer maximizes the expected utility:

$$\mathbf{U} = E \left[\sum_{t=0}^{\infty} \beta^t [C_t + Av(1 - L_t)] \right],$$

where $v(\cdot)$ is the increasing and concave utility function for leisure, C_t is the consumption level, L_t is the employment level, $\beta \in (0, 1)$ is the discount factor, and A is a parameter. Here, for simplicity, we consider linear utility for consumption.¹⁸ This simplification enables us to discount the firm's profit by the discount factor β . Since we consider the adjustment of L_t at the extensive margin, the appropriate interpretation of the $v(\cdot)$ function is that it is the result of an aggregation of many consumers who have different preferences over consumption and leisure. The budget constraint in each period is:

$$C_t = w_t L_t + \Pi_t + R_t, \quad (6)$$

where w_t is the wage rate, Π_t is the firm's profit, and R_t is the transfer from the government. The government transfers the firing tax to the consumer in a lump-sum manner every period. We assume that there is no saving. The first-order condition in each period is:

$$Av'(1 - L_t) = w_t. \quad (7)$$

¹⁸Hopenhayn and Rogerson (1993) assume a period utility function that is concave in consumption and linear in leisure.

3.3 General equilibrium

First, consider a situation where z_t is constant. We use the solution of this steady-state situation for the purpose of calibration later. In this case, the definition of the stationary equilibrium is similar to Hopenhayn and Rogerson (1993). In our model, the general equilibrium can be summarized in the labor market.¹⁹ The free-entry condition (5) characterizes the demand side of the labor market. The quantity of the labor demand in the steady state is given by

$$L^d = N \int \phi(s', n) d\mu(s', n), \quad (8)$$

where $\mu(s', n)$ is the stationary distribution of the plants with the state (s', n) when we assume that the mass of potential entry in each period is one. $\phi(s', n)$ is the labor demand for a plant with the state (s', n) . N is the actual mass of potential entry. If (5) is not met, N becomes ∞ or 0. It results in $L^d = \infty$ or $L^d = 0$, which violates the labor-market equilibrium condition. N can take a finite and positive number only if $w = w^*$, where w^* satisfies the free-entry condition (5).

The consumer's first-order condition (7) characterizes the labor supply side. The labor demand side in effect determines the wage level at w^* (with the free-entry condition (5)). Combined with the labor-supply curve (7), the equilibrium level of labor, L^* , is determined. This labor-market equilibrium is depicted in Figure 1. Once L^* is determined, the equilibrium level of N , N^* , is determined by (8).

When we introduce an aggregate shock, L^* and N^* move over time. The labor demand is now characterized by

$$L_t^d = L_{it}^d + N_t L_{et}^d, \quad (9)$$

where L_{it}^d is the labor demand from incumbents at period t and L_{et}^d is the labor demand from the entrant when the mass of potential entry is assumed to be one. The determination of

¹⁹In Hopenhayn and Rogerson (1993), the general equilibrium is determined in the goods market. The difference in our model comes from the fact that we have a utility function that is linear in consumption and concave in leisure, while Hopenhayn and Rogerson have a utility function that is linear in leisure and concave in consumption.

Figure 1: Labor-market equilibrium

the equilibrium is similar: the free-entry condition (5) determines the wage, the labor-supply equation (7) determines L , and the labor-demand equation (9) determines N .

Aggregate profit is given by

$$\Pi_t = Y_t - w_t L_t - R_t - N_t c_q - M_t c_e + X_t,$$

where Y_t is aggregate output, N_t is the number of potential entrants, M_t is the number of actual entrants, and X_t is the total value of exiting. Therefore, combining this with (6), in equilibrium (where labor demand equals labor supply)

$$C_t = Y_t - N_t c_q - M_t c_e + X_t.$$

3.4 Calibration

Our strategy is to use the steady state of the model with constant z (we set $z = 1$) as the benchmark for calibration, and to add the aggregate shocks later on. A large part of our calibration is based on the statistics presented in Section 2. Following Hopenhayn and Rogerson (1993), we normalize the wage rate, w , in the benchmark to 1. As in Hopenhayn and Rogerson (1993), the model exhibits a homogeneity property in the sense that given

prices, all of the aggregate variables (quantities) are proportional to the number of potential entrants, N . We pin down the benchmark value of N by setting aggregate employment, L , to 0.6 (approximate employment rate in the U.S.). The value of A is backed out from (7) and the fact that $w = 1$ and $L = 0.6$ in the benchmark. We set $\beta = 0.94$ and $\theta = 0.7$.

The process for idiosyncratic productivity, s , is chosen so that the model generates the employment process observed in the data (Table 3). First, the process is assumed to be

$$\ln(s') = a_s + \rho_s \ln(s) + \varepsilon_s,$$

where

$$\varepsilon_s \sim N(0, \sigma_s^2).$$

Then, this process is approximated by a Markov process using Tauchen's (1986) method to obtain $\psi(s'|s)$. We set 30 evenly-spaced grids on $\ln(s)$ over the interval $[a_s/(1 - \rho_s) - 3\sqrt{\sigma_s^2/(1 - \rho_s^2)}, a_s/(1 - \rho_s) + 3\sqrt{\sigma_s^2/(1 - \rho_s^2)}]$. The constant a_s is set so that the average value of employment matches the data. ρ_s is set to 0.97. This will bring the value of ρ in equation (3) to the empirical value of 0.97 in Table 3. σ_s is set so that the variance of the growth rate of n is close to the empirical value of 0.14. The resulting values are $a_s = 0.04$ and $\sigma_s = 0.11$.

The adjustment factor λ is set at 0.983, following Cooper, Haltiwanger, and Willis (2004).²⁰ As mentioned above, Cooper, Haltiwanger, and Willis show that this form of adjustment cost is empirically most relevant. In our model, including a large quadratic adjustment cost makes the size of entering plants unrealistically small, and a small quadratic adjustment cost does not alter the quantitative predictions along other dimensions.

The exit value is assumed to be zero with probability x_0 . With probability $(1 - x_0)$, the exit value is uniformly distributed over $[0, \bar{x}]$. We set x_0 and $\bar{x}$ so that the exit rate and the size of the exiting plants are similar to the empirical values. We choose $x_0 = 0.9$ and $\bar{x} = 2500$. We assume the entry transition function to be identical to the transition function

²⁰This is their point estimate with a small quadratic adjustment cost. Their point estimate for λ with no other adjustment cost is 0.988. We prefer the former number since it produces a more reasonable job reallocation rate.

	Data	Model
Average size of continuing plants	87.5	87.6
Average size of entering plants	50.3	49.7
Average size of exiting plants	35.0	35.8
Entry rate	6.2%	5.4%
Exit rate	5.5%	5.4%
AR(1) coefficient ρ for employment	0.97	0.97
Variance of growth rate for n	0.14	0.14
Job reallocation rate	19.4%	23.0%

Table 10: Data and model statistics in the steady-state

for the incumbents: $\eta(s'|q) = \psi(s'|s)$. The entry costs, c_q and c_e , are backed out from the model. Given the value function $V^c(s', n)$, conditions (4) and (5) pin down the values of c_q and c_e , given $\nu(q)$ and the equilibrium value of q^* that we target. We assume that $\nu(q)$ follows $\nu(q) = B \exp(-q)$ over the lower part of the grids on s (B is the scale parameter to make $\nu(q)$ sum up to one).²¹ We select the value of c_e so that the target value of $\ln(q^*)$ is 0.5. As we see below, this choice of $\nu(q)$ and q^* brings the size distribution of young plants close to the data. In the benchmark, we set the firing cost, $g(n', n)$, to zero. For the $v(\cdot)$ function, we use $\ln(\cdot)$.

3.5 Steady-state results

First, we compute the model without aggregate shocks to establish the steady-state behavior of the model. The details of the computation of the steady-state model are described in Appendix B. Table 10 compares the output of our model to the data.²² Everything except for the job reallocation rate is our “target” for calibration, and we can see that we are able to get close to the empirical values. The job reallocation rate is also close to the value in the data.²³

²¹We set 200 grids on x and 25 grids on q . The results do not change when we increase the number of x grids to 1000 or when we use interpolation to approximate the continuous x distribution.

²²The job reallocation rate is taken from Davis, Haltiwanger, and Schuh (1996, Table 2.1).

²³If we assume that $\lambda = 1.0$, the job reallocation rate increases to over 30%.

	Data	Model
1-19	0.457	0.483
20-49	0.239	0.189
50-99	0.131	0.107
100-249	0.106	0.122
250+	0.067	0.099

Table 11: The number of plants for each size class

	Data	Model
1-19	0.049	0.040
20-49	0.090	0.066
50-99	0.109	0.080
100-249	0.194	0.197
250+	0.559	0.617

Table 12: Employment share by size class

The following tables evaluate the distributional performance of the model. Table 11 summarizes the distribution of plant size by the fraction of plants in each size class. Overall, the size distribution of the model matches the data quite well. Table 12 describes the employment share by size class. It also exhibits a good match.

Since much of our focus is on entry and exit behavior, it is critical that the model's properties for entering and exiting plants match the data. Table 13 exhibits the distribution of young plants in the model and the data, which are reasonably similar.²⁴

3.6 Adding aggregate shocks

To analyze business cycles, we assume that z_t fluctuates between two values. Following the standard calibration in the real business cycles tradition, we assume $z_t = 1.01$ in good times and $z_t = 0.99$ in bad times. This results in a 1% standard deviation in z_t . z_t follows a symmetric Markov process. We calibrate the transition probabilities so that the average

²⁴The definition of young plants follow Davis, Haltiwanger, and Schuh (1996) for the data. In the model, young plants are defined as those which are 0, 1, or 2 years old.

	Data	Model
1-19	0.712	0.614
20-49	0.156	0.150
50-99	0.068	0.079
100-249	0.044	0.093
250+	0.020	0.064

Table 13: The number of plants within each size class—within young plants

	Data	Model
1-19	0.080	0.063
20-49	0.039	0.059
50-99	0.031	0.053
100-249	0.025	0.042
250+	0.015	0.011

Table 14: Exit rates for each size class

duration of each state is three years.

The computation turns out to be much simpler than standard heterogeneous-agent models, since the wage depends only on z . From this property, we can perform the optimization by plants and determine $w(z)$ without considering the labor-market equilibrium. After $w(z)$ is determined, the labor-market equilibrium determines the equilibrium quantities, in particular the mass of entrants, N . The details of the computation are in Appendix C.

The results of the model with aggregate shocks are summarized in Table 15. First, notice that the wage fluctuates substantially. While the cyclicalty of wages is empirically controversial, we follow Cooley and Prescott (1995) and consider wages to be procyclical and to have a standard deviation of less than 1% (see their Table 1.1). In our model, a procyclical wage is necessary to make employment procyclical. Figure 2 shows the labor market equilibrium. For the equilibrium L to increase when the labor supply curve stays the same, we need w to increase. Somewhat surprisingly, in Table 15, the equilibrium value of q^* does not change with the change in z . It is because the effect of z and the effect of w offset each other. If z goes up and w does not change, it is profitable for low- q plants to enter.

	Good	Bad
Wage	1.014	0.986
q^*	0.5000	0.5000
Entry rate	6.7%	4.0%
Exit rate	5.3%	5.4%
Average size of all plants	84.6	86.5
Relative size of entrants	0.57	0.57
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.85	0.85
Relative productivity of exiting plants	0.84	0.84

Table 15: Results with aggregate shocks

Therefore q^* goes down. However, since w goes up, the profitability of entry goes down. It turns out that these two effects offset each other almost exactly. Since q^* is the same across the two states, the plants that enter are similar across different aggregate states. This is reflected in the similar relative sizes and relative productivities of entrants. This pattern is at odds with the data—as is shown in Section 2, the data show a strong cyclical pattern in the selection of the entrants.

In the model, exiting plants compare the value of staying with the value of exiting when making exit decisions. Since the exit value distribution is quite wide, a 1% difference in z does not make a large difference for this comparison.²⁵ Thus the exit rate and the size and productivity of exiting plants are similar throughout the business cycle in Table 15. This fits well with the pattern we observe in the data.

The entry rate fluctuates significantly in Table 15, as we see in the data. The mechanism here is simple: since the wage increases during booms, the average size of incumbent plants shrinks (which we observe in the data). The labor that is released from the incumbents can be hired by the entrants—thus, the entry rate goes up. Here, again, the procyclicality of wages plays an important role.

²⁵If the exit value is concentrated at some value and there are many “marginal” plants, it is possible that these plants exit with a small change in z . Thus, it is important for this result that we choose the exit value distribution so that the model matches the cross-sectional patterns of exit.

Figure 2: Labor-market equilibrium when z increases

Summarizing, we found that the model is successful in matching patterns in the data in some respects. The two characteristics that are at odds with the data are the large fluctuations in wages and the lack of selection in entry. These two are linked in the sense that if the wage fluctuates less, we expect that some selection effect will emerge.

How, then, can we reduce the fluctuation of wages? To answer this question, it is useful to first consider why wages are so volatile in the model. As is argued above, the equilibrium wage is determined from the free-entry condition (5), which equates the idea cost to the present value of the profit associated with coming up with an idea. If the idea cost stays the same, the present value of the profit has to stay the same. Since an increase in z will increase the profit, w has to increase to offset the increase in the profit.²⁶

Therefore, it is possible to reduce the volatility of wages in the model by introducing another force that counteracts the change in the profit. In the benchmark model, we assume

²⁶The following example shows that the wage has to increase by about 1.5% to offset the 1% increase in z . For simplicity, consider a one-period model without any adjustment cost. Let the profit function be $\pi = zsn^\theta - wn$. Let $\theta = 2/3$, for computational ease. From the first-order condition, $w = \frac{2}{3}zsn^{-\frac{1}{3}}$. If the wage is constant, when z increases by 1%, n has to increase by 3%. This will increase the profit by 3%. If, instead, z does not change and w increases by 1%, n has to decrease by 3%. The resulting decrease in profit is 2%. Thus, to keep π constant, a 1% increase in z has to be accompanied by a 1.5% increase in w .

that the entry costs, c_e and c_q , are constant. If either c_e or c_q is procyclical, the fluctuation of wages will be smaller. In the following, we explain how changing the entry costs affects the aggregate performance of the model.

First, we consider a procyclical c_e . Table 16 shows the result for the case where c_e is 0.032% larger in booms and 0.032% smaller in recessions. Here, we change c_e so that the wage has roughly the same degree of variation as z . It is clear that even though the fluctuation in w is smaller, the selection effect goes in the wrong direction: it is more difficult for low-productivity plants to enter during booms. This comes from the fact that c_e has a direct effect on the selection process—a large c_e makes the actual entry harder. In fact, empirical studies on investment costs also suggest that c_e may move in a countercyclical direction. Recall that c_e can be thought of as the cost of actual entry—in particular the sunk investment in equipment and structures at entry. It is known that the price of investment goods tends to be lower during booms (see Fisher, 2006), and this evidence suggests that c_e may be lower during booms. In terms of the model, this can be treated as an exogenous shock to the value of c_e , which is negatively correlated with the variation in z . Furthermore, financial costs may also depend on the aggregate state of the economy. If financing new plants is more difficult during recessions, higher financial costs will cause c_e to rise during these periods. We do not explicitly model the financial intermediation process, but this may be an important factor when we consider the role of financial frictions for business cycle propagation.²⁷

Now we turn to another mechanism, procyclical c_q . The interpretation of c_q is the “idea cost.” In reality, it will come up as the cost of R&D to create new ideas (innovation). Idea creation is a human capital intensive process. The cost of hiring a good inventor is particularly higher during booms, partly because the wages for these workers are higher then.²⁸ In addition, there are more entries and idea creations during booms, and the idea creation

²⁷See, for example, Bernanke and Gertler (1989), Carlstrom and Fuerst (1997), and Kiyotaki and Moore (1997).

²⁸The National Science Foundation (NSF) collects various data on R&D expenditures and costs. On average, the cost per R&D scientist or engineer in companies performing R&D was about 8.6% higher during booms (good times) than during recessions (bad times).

	Good	Bad
Wage	1.010	0.990
q^*	0.5062	0.4938
Entry rate	6.1%	4.7%
Exit rate	5.3%	5.4%
Average size of all plants	85.2	85.9
Relative size of entrants	0.57	0.56
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.85	0.84
Relative productivity of exiting plants	0.84	0.84

Table 16: The case of a procyclical c_e

	Good	Bad
Wage	1.010	0.990
q^*	0.4997	0.5003
Entry rate	6.0%	4.7%
Exit rate	5.3%	5.4%
Average size of all plants	85.3	85.9
Relative size of entrants	0.57	0.57
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.85	0.85
Relative productivity of exiting plants	0.84	0.84

Table 17: The case of a procyclical c_q

process may suffer from decreasing returns (due to, for example, the “fishing-out” effect).²⁹ Table 17 shows the result. Here, c_q is 0.165% larger during booms and 0.165% smaller during recessions. We can see that this generates a *qualitatively* successful result. The selection goes in the right direction. However, it does not generate a selection effect *quantitatively* large enough to match the data.

How can we generate a stronger selection effect? Notice that the change in c_e has the

²⁹In this paper, we do not explicitly consider this possibility, since modeling such an effect requires making c_q a function of N , and eventually making w dependent on N . This will add a significant computational burden since N is a function of the current plant distribution. This remains an important extension for future research.

	Good	Bad
Wage	1.010	0.990
q^*	0.3047	0.6335
Entry rate	7.0%	3.9%
Exit rate	5.3%	5.5%
Average size of all plants	79.5	82.4
Relative size of entrants	0.48	0.69
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.78	0.94
Relative productivity of exiting plants	0.84	0.84

Table 18: The case of a countercyclical c_e and procyclical c_q

opposite effect on selection. Then, a countercyclical c_e would help to generate a larger selection effect. Furthermore, as argued above, the empirical evidence seems to suggest that c_e is in fact countercyclical. If we combine a countercyclical c_e with a procyclical c_q (which will counteract the counterfactual wage effect of a countercyclical c_e), we may be able to generate a larger selection effect, since both a countercyclical c_e and a procyclical c_q make the selection effect work in the right direction.

Table 18 describes the result of an experiment where c_e is 0.8% higher during recessions and 0.8% lower during booms, and c_q is 3.3% lower during recessions and 4.1% higher during booms. This generates a large selection effect, and the differences in the relative size and productivity of entrants in booms and recessions is comparable to the data.

Figure 3 depicts a simulated sample path of aggregate output. We can see that most of the changes in output occur during periods where the aggregate state switches. If we look closely, however, there are nontrivial dynamics (although the magnitude is very small) within each aggregate state. Figure 4 magnifies a part of Figure 3, by picking up on one boom. We can see “hump-shaped” dynamics.³⁰ Output increases in the period following switching, and

³⁰Chang, Gomes, and Schorfheide (2002) analyze a propagation of shocks in U.S. business cycles and obtain hump-shaped responses of output and hours in their vector autoregressions analysis. They argue that the hump-shaped responses cannot be obtained by a standard real business cycles model, and suggest that learning-by-doing can be an important propagation mechanism.

then it starts to decline. Since L is constant for a given aggregate state z (the wage w is a function of only z , and the labor supply curve does not shift), this movement comes purely from the change in productivity. Since z is constant within each aggregate state, the source of these dynamics is the change in the composition of the idiosyncratic productivity at different plants. Here, two effects are at work. In general, an increase in the number of plants for a given L increases average productivity, since a plant's production function is subject to decreasing returns to scale. Since the number of entrants is above average during a boom, this increases the number of total plants, and average productivity increases. At the same time, an entering plant is less productive than incumbents, particularly during a boom, so the distribution of plants' productivity worsens as new plants are added. It turns out that, with our calibration, the first effect dominates initially and the second effect dominates later on.³¹ Figure 5 shows that similar dynamics can be observed during recessions. The magnitude of these responses is very small, so that these are likely to be dominated even by a small aggregate productivity shock if we allow for more than two levels of z . However, we believe that it is an interesting theoretical possibility that the change in the distribution of idiosyncratic productivity can serve as a propagation mechanism through entry and exit. In particular, this mechanism would be relevant in sectors where entry and exit rates are large.

4 Policy implications

In this section, we explore the policy implications of our model, using the calibration with countercyclical c_e and procyclical c_q . (In the following, we call the results in Table 18 our “benchmark.”) We consider four experiments. In three of these, we consider a cyclical policy—a particular policy is imposed only during recessions. First, we consider a firing tax, which is constant over time (the tax revenue is given back to consumers in a lump-sum manner, so it is counted in aggregate output). We consider the following specification of the

³¹Interestingly, in the model with constant c_e and c_q (Table 15), the first effect always dominates. This means that the cyclicity of the entrants' productivity is important in generating the hump-shaped dynamics.

Figure 3: A sample path of aggregate output

Figure 4: A sample path of aggregate output (magnified: a boom)

Figure 5: A sample path of aggregate output (magnified: a recession)

firing tax:

$$g(n_t, n_{t-1}) = \tau \max\langle 0, n_{t-1} - n_t \rangle.$$

We set $\tau = 0.1$. Since the wage is set at 1 at the benchmark, this implies that the firing tax per person is 10% of an annual wage. Second, we impose this firing tax only during recessions. Third, entry is subsidized by 0.1% during recessions, in terms of both c_e and c_q . Fourth, entry is subsidized by 0.5% during recessions, but only in terms of c_q . The subsidies are financed by lump-sum taxes.

4.1 Constant firing tax

The consequences of a firing tax on the allocation of employment have been analyzed by many researchers in recent years. For example, Veracierto (2004) analyzes the implications of a constant firing tax in a general equilibrium establishment-level dynamics model. Veracierto's model incorporates saving and the capital stock, but entry and exit are assumed to be exogenous. Samaniego (2006) also conducts a similar exercise.

The results from our model with 10% firing tax are shown in Table 19. Entry and exit behavior do not experience quantitatively large changes compared to the benchmark,

	Good	Bad
Wage	0.999	0.980
q^*	0.3042	0.6337
Entry rate	7.1%	3.8%
Exit rate	5.3%	5.5%
Average size of all plants	80.1	83.2
Relative size of entrants	0.46	0.66
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.77	0.94
Relative productivity of exiting plants	0.83	0.83

Table 19: Results with a constant firing tax (10%)

although we see some differences in the entry threshold, q^* , and in the entry rate. The average size of plants increases, reflecting the reluctance to fire. In terms of the average statistics, we see changes we associated with the firing costs. The job reallocation rate falls from 23.0% (in the benchmark) to 21.3%. Average output falls by 0.9%, and average employment falls by 0.7%.

Interestingly, the variance of output *increases* slightly, by 2.4%.³² This contrasts with Veracierto (2004) and Samaniego (2006), who find that the firing cost is stabilizing. In our experiment, the variance of output by survivors decreases with the firing tax, as does the variance of output by unit mass of entrants. However, the variance of the entry rate increases, which leads to the increase in the variance of total output.³³ Intuitively, the firing tax is a tax on relatively large plants, which are more likely to fire in the near future. During recessions, entering plants are typically larger than entering plants during booms, so they experience a larger tax burden. This works in the direction of reducing the entry rate during recessions relative to booms. This is reflected in the difference in q^* s between Table 18 and Table 19: q^* decreases in good times and increases in bad times. The general equilibrium effects also

³²The coefficient of variation also increases, since the mean becomes smaller with the firing cost.

³³Veracierto does not have this margin, since the entry rate is assumed to be constant in his model. Samaniego's model features endogenous entry. However, in his model, the entry rate reacts very little to the change in aggregate productivity.

	Good	Bad
Wage	1.006	0.983
q^*	0.3046	0.6337
Entry rate	9.1%	1.8%
Exit rate	5.3%	5.7%
Average size of all plants	73.1	80.6
Relative size of entrants	0.52	0.70
Relative size of exiting plants	0.41	0.40
Relative productivity of entrants	0.80	0.97
Relative productivity of exiting plants	0.84	0.83

Table 20: Results with a firing tax, during recessions only (10%)

operate, but it turns out that, with our calibration, the entry rate increases with firing costs during booms and it decreases with firing costs during recessions.

4.2 Firing taxes during recessions

Next, we consider the case where the government imposes the tax only during recessions, so that it can reduce the amount of firing during these times. The results are summarized in Table 20. The government succeeds in its intention—the average job destruction rate during recessions falls to 9.7% versus 11.6% in the benchmark. However, as we can see from the table, the entry and exit rates fluctuate more than in the benchmark, and the variance of output more than doubles as a result.

4.3 Entry subsidies during recessions, both c_e and c_q

Looking at the data, the government might think that entry rates are too low during recessions and decide to subsidize entry costs only during recessions. In our model, there are two types of entry costs: the “idea cost” and the “implementation cost.” First consider the case where both costs are subsidized at the same rate. A subsidy to the “idea cost” can be interpreted as an R&D subsidy, and a subsidy to the “implementation cost” can be interpreted as an investment subsidy. The results are summarized in Table 21. Entry rates are less volatile

	Good	Bad
Wage	1.004	0.997
q^*	0.3053	0.6206
Entry rate	6.1%	4.6%
Exit rate	5.4%	5.4%
Average size of all plants	81.6	83.1
Relative size of entrants	0.47	0.67
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.77	0.93
Relative productivity of exiting plants	0.84	0.84

Table 21: Results with entry subsidies on c_e and c_q only during recessions (0.1%)

compared to the benchmark. The variance of output is substantially reduced—it becomes less than half of the variance found in the benchmark. The selection of entrants during recession is not as stringent as in the benchmark: q^* is smaller now. Since wage volatility is also smaller, the average size of plants is also similar across booms and recessions. If the government’s goal is to stabilize output along with entry and exit, this type of subsidy is more effective than the (cyclical or noncyclical) firing cost.

4.4 Entry subsidies during recessions, only c_q

Second, consider the case where only the “idea cost” is subsidized during recessions—in our interpretation, this corresponds to an R&D subsidy. The results are in Table 22. Again, the government can achieve stability in entry rates. The variance of output is also small—less than half of the benchmark. The difference from the previous experiment is that here the selection of entering plants is more stringent during recessions than in the benchmark.

5 Conclusion

This paper explores the business-cycle implications of plant-level dynamics, particularly the entry and exit behavior of plants. First we documented patterns of plant entry, exit, employment, and productivity in U.S. manufacturing, utilizing the Annual Survey of Manufactures.

	Good	Bad
Wage	1.004	0.998
q^*	0.3053	0.6348
Entry rate	5.9%	4.9%
Exit rate	5.4%	5.4%
Average size of all plants	82.6	83.5
Relative size of entrants	0.47	0.67
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.76	0.93
Relative productivity of exiting plants	0.84	0.84

Table 22: Results with entry subsidies on c_q only during recessions (0.5%)

We found that the entry rate is much more cyclical than the exit rate, and entering plants’ average size and productivity vary significantly over the business cycle. Then we constructed a general equilibrium model of plant dynamics by extending Hopenhayn and Rogerson’s (1993) model. Our model accounts for the properties that we found in the data, when several assumptions are made about the cyclicalities of entry costs. We conducted several policy experiments using our model. Both a constant firing tax and a countercyclical firing tax increase the volatility of the entry rate and aggregate output. Countercyclical entry subsidies stabilize the entry rate and aggregate output over the business cycle.

We found that a countercyclical “implementation cost” and a procyclical “idea cost” are important in matching our model to the data. In this paper, we did not explicitly model why these costs exhibit such cyclical patterns. An important research topic for the future will be to uncover the nature of these costs theoretically (by modeling the microeconomic foundations of these costs³⁴) and empirically (by looking into the microeconomic process of entry).

Finally, we would like to emphasize that our study focuses only on the U.S. manufacturing sector. Investigating whether other U.S. industrial sectors or manufacturing sectors in other

³⁴Explicitly modeling the limited enforceability of contracts, as in Cooley, Marimon, and Quadrini (2004), is one possible direction.

countries exhibit the same patterns is beyond the scope of this paper, but we believe that these are also very important topics for future research.

Appendix

A Data and measurement

A.1 Identifying entry and exit

In this paper, we focus on first-time plant openings (i.e., birth) and permanent shutdowns (i.e., death). The Census Bureau adds new plants from the Company Organization Survey and the Business Register into the ASM panel. We identify startup and shutdown candidates following Davis, Haltiwanger, and Schuh (1996). All startup plants should have at least one employee, while plants with zero employees are considered to be shutdowns (either temporary or permanent). We exclude reopened plants and temporary or indefinite shutdowns using the Census of Manufactures. In order to exclude spurious startup and shutdowns, we mainly use information from the Census of Manufactures. While Davis, Haltiwanger, and Schuh (1996) use the coverage code (CC), we find the coverage code less useful for more recent cohorts. We find the number of startup and shutdown candidates with “CC= 0 (no change in operations)” increases over time. Furthermore, even valid coverage codes are reported with some leads and lags as compared to the timing of startup or shutdown.

It may be possible that the timing of birth in the ASM panel may be earlier or later than actual birth due to the time lag in adding start-ups into the survey. While this may cause a problem with statistics for an individual year (e.g., annual averages for entrants or job creation from startups), it is less likely be a problem in our statistics on booms and recessions because our classification of booms and recessions has an average duration of 3 years. Because plants that enter during this multi-year window of recessions (e.g., recessions between 1979-1983) are all counted as entrants during recessions, a lag of a year or two would not alter our statistics substantially.

A.2 Variables for productivity measures

This appendix documents how variables in the productivity measures used in this paper are constructed.

Capital: We follow Dunne, Haltiwanger, and Troske (1997) closely in constructing the capital stock. For the initial benchmark, we use the book value of structures or equipment, deflated by the two-digit industry capital deflator from the BEA (2-digit). We use the average of beginning-of-year assets and end-of-year assets. While we separately examine structures and equipment, for recent years where the ASM reports only total assets (structures and equipment together), the deflated book value of total assets is used as the initial benchmark. Investment deflators are from the NBER manufacturing productivity database (Bartelsman and Gray, 1996). The depreciation rate for each two-digit industry was also obtained from the BEA. Real capital stocks are obtained by summing up the real value of structures and the real value of equipment constructed from the perpetual inventory method.

Labor input: Labor input for TFP (based on (1)) is measured as total hours for production and nonproduction workers. Because hours for nonproduction workers are not collected, we estimate the value for total hours by following the method in Baily, Hulten, and Campbell (1992), which is to multiply the total hours of production workers by the ratio of the total payroll for all workers to the payroll for production workers. Following the model, we use the total number of workers as labor input for the TFP based on (2).

Materials: Costs of materials are deflated by material deflators from the NBER manufacturing productivity database.

Output: For TFP, we use the total value of shipments (TVS) deflated by the shipments deflator from the NBER manufacturing productivity database. Although it is possible to adjust output for the change in inventories, inventories for some plants (in particular, for small plants) are imputed (Baily, Bartelsman, and Haltiwanger, 2001). To avoid a possible measurement issue, we have chosen to use gross shipments as a simple measure. For the productivity based on (2), value added, deflated by shipment deflators, is used. We also used

deflated shipments ($TVS/PISHIP$) minus the real value of materials, but the results did not change much.

Revenue shares: We use 4-digit industry-level revenue shares as factor elasticities. This procedure implicitly assumes that all plants in the industry operate with the same production technology, a common assumption in studies measuring plant-level productivity. In calculating labor's share of total costs, we follow Bills and Chang (2000), magnifying each four-digit industry's wage and salary payments to reflect other labor payments, such as fringe payments and employer FICA payments. We use information from the National Income and Product Accounts to calculate the ratio of these other labor payments to wages and salaries at the two-digit industry level.

B Computation of the steady state

This section outlines the computation of the model without aggregate shocks. We omit the notation on z since it is constant here.

1. Set discrete grids on n and s . Set the Markov transition matrix for s . Set the distribution of the exit value x .
2. Optimization loop. Objects: $W(s, n)$, $V^a(s', n)$, $V^n(s', n)$, $V^c(s', n)$, $Z(s, n)$, $\phi^a(s', n)$, $\phi(s', n)$, $\zeta(s', n)$, and $\chi(s, n)$. (These functions are defined in the following.)
 - (a) Give the initial value for $W(s, n)$, where s and n are the realizations in the last period. This is the beginning-of-period value for an incumbent.
 - (b) Calculate $V^a(s', n)$ and $V^n(s', n)$ by

$$V^a(s', n) = \max_{n'} \pi^a(s', n, n') + \beta W(s', n')$$

and

$$V^n(s', n) = \pi^n(s', n) + \beta W(s', n),$$

where

$$\pi^a(s', n, n') = \lambda z f(n', s') - w n' - g(n', n)$$

and

$$\pi^n(s', n) = z f(n, s') - w n.$$

Record the decision rule of n' when adjusted: $\phi^a(s', n)$.

(c) Calculate $V^c(s', n)$ by

$$V^c(s', n) = \max\langle V^a(s', n), V^n(s', n) \rangle.$$

Record the decision rule. $\zeta(s', n) = 1$ if adjust, and $\zeta(s', n) = 0$ if not. $\phi(s', n) = \phi^a(s', n)$ if $\zeta(s', n) = 1$ and $\phi(s', n) = n$ if $\zeta(s', n) = 0$.

(d) Calculate $Z(s, n) = E_{s'}[V^c(s', n)|s]$ by

$$Z(s, n) = \int V^c(s', n) d\psi(s'|s).$$

(e) Calculate $W(s, n)$ by

$$W(s, n) = \int \max\langle Z(s, n), x - g(0, n) \rangle d\xi(x).$$

Thus, the ratio of plants that exit with the state (s, n) is

$$\chi(s, n) = \int_{Z(s, n) + g(0, n)}^{\infty} d\xi(x).$$

(f) Update and repeat.

3. Now we can calculate

$$V^e(q) = \int V^c(s', 0) d\eta(s'|q)$$

for each q .

We can set the cut-off for q , q^* , and find c_e by

$$V^e(q^*) = c_e. \tag{10}$$

Then we find V^p by

$$V^p = \int \max\langle V^e(q) - c_e, 0 \rangle d\nu(q)$$

and find c_q by the free-entry condition:

$$V_p = c_q. \quad (11)$$

4. Calculate the stationary measure of plants (survivors from the last period plus this period's entrant, after receiving this period's shock), $\mu(s', n)$, given $N = 1$. From linear homogeneity, the actual measure of survivors will be $N\mu$.

Note that M is a function of N :

$$M = N \int_{q^*}^{\infty} d\nu(q).$$

5. Obtain N by solving

$$Av'(1 - L(N)) = w. \quad (12)$$

In the benchmark, we choose A so that $L = 0.6$ when $w = 1$. Thus, when $v(x) = \ln(x)$,

$$A = w - wL(N) = 1 - 1 \times 0.6 = 0.4.$$

Here, $L(N)$ is

$$L(N) = N \int \phi(s', n) d\mu(s', n)$$

so that N is calculated by

$$N = \frac{0.6}{\int \phi(s', n) d\mu(s', n)}.$$

The total output can be calculated by:

$$Y(N) = N \int [zf(\phi(s', n), s') - (1 - \lambda)\zeta(s', n)zf(\phi(s', n), s')] d\mu(s', n).$$

The total exit value $X(N)$ is calculated by

$$X(N) = N \int \int_{Z(s', \phi(s', n)) + g(0, \phi(s', n))}^{\infty} x d\xi(x) d\mu(s', n).$$

C Computation of the model with aggregate shocks

1. Set discrete grids on n and s . Set the Markov transition matrix for s . Set the distribution of the exit value x .
2. Guess w as a function of z .
3. Optimization loop. Objects: $W(s, n; z)$, $V^a(s', n; z)$, $V^n(s', n; z)$, $V^c(s', n; z)$, $Z(s, n; z)$, $\phi^a(s', n; z)$, $\phi(s', n; z)$, $\zeta(s', n; z)$, and $\chi(s, n; z)$.

(a) Give the initial value for $W(s, n; z)$, where s and n are the realizations in the last period. This is the beginning-of-period value for an incumbent.

(b) Calculate $V^a(s', n; z)$ and $V^n(s', n; z)$ by

$$V^a(s', n; z) = \max_{n'} \pi^a(s', n, n'; z) + \beta E_{z'}[W(s', n'; z')|z]$$

and

$$V^n(s', n; z) = \pi^n(s', n; z) + \beta E_{z'}[W(s', n; z')|z],$$

where

$$\pi^a(s', n, n'; z) = \lambda f(n', s', z) - w(z)n' - g(n', n)$$

and

$$\pi^n(s', n; w, z) = f(n, s', z) - w(z)n.$$

Record the decision rule of n' when adjusted: $\phi^a(s', n; z)$.

(c) Calculate $V^c(s', n; z)$ by

$$V^c(s', n; z) = \max(V^a(s', n; z), V^n(s', n; z)).$$

Record the decision rule. $\zeta(s', n; z) = 1$ if adjust, and $\zeta(s', n; z) = 0$ if not.

$\phi(s', n; z) = \phi^a(s', n; z)$ if $\zeta(s', n; z) = 1$ and $\phi(s', n; z) = n$ if $\zeta(s', n; z) = 0$.

(d) Calculate $Z(s, n; z) = E_{s'}[V^c(s', n; z)|s]$ by

$$Z(s, n; z) = \int V^c(s', n; z) d\psi(s'|s).$$

(e) Calculate $W(s, n; z)$ by

$$W(s, n; z) = \int \max \langle Z(s, n; z), x - g(0, n) \rangle d\xi(x).$$

Thus, the ratio of plants that exit with the state $(s, n; z)$ is

$$\chi(s, n; z) = \int_{Z(s, n; z) + g(0, n)}^{\infty} d\xi(x).$$

(f) Update and repeat.

4. Now we can calculate

$$V^e(q; z) = \int V^c(s', 0; z) d\eta(s' | q)$$

for each q .

We can set the cut-off for q , $q^*(z)$, and find c_e by $V^e(q^*; z) = c_e(z)$.

Then we find $V^p(z)$ by

$$V^p(z) = \int \max \langle V^e(q; z) - c_e(z), 0 \rangle d\nu(q).$$

5. Check if $V_p(z) = c_q$ is satisfied. If not, revise $w(z)$. Repeat until convergence.

6. Simulation.

(a) Let $\delta_t(s, n)$ be the total measure of incumbents at the beginning of period t , before entry and exit.

(b) After observing z , incumbents decide whether to exit. The measure of survivors is

$$\sigma_t(s, n; z) = (1 - \chi(s, n; z)) \delta_t(s, n).$$

The total measure of survivors, with new shocks, is

$$\vartheta_t(s', n; z) = \int \psi(s' | s) \sigma_t(ds, n; z).$$

(c) The measure of entrants for unit N (after observing z) is

$$\varrho_t(s'; z) = \int_{q^*(z)}^{\infty} \eta(s'|q) d\nu(q)$$

(d) Thus the total measure is

$$\mu_t(s', n; z) = N\varrho_t(s'; z) + \vartheta_t(s', n; z).$$

(e) We have to solve for N to actually calculate the total measure. The labor market equilibrium condition is (from the representative consumer's first-order condition)

$$w(z) = A \frac{1}{1 - L}.$$

L is the sum of the incumbent's employment L^i , which can be calculated from $\theta_t(s', n; z)$, and the entrant's employment, NL^e , where L^e can be calculated from $\varrho_t(s'; z)$. Therefore,

$$N = \frac{(w(z) - A)/w(z) - L^i}{L^e}.$$

(f) The beginning-of-next-period measure can be found by

$$\delta_{t+1}(s', n') = \int_{n'=\phi(s', n; z)} \mu_t(s', dn; z).$$

(g) Repeat.

D Results for significant positive and negative growth states

In this appendix, we report the results of the experiments originally reported in Tables 15 to 18, when we use a different categorization of good and bad times. Here, we categorize the good times as the times when output growth rates are more than 0.1%, and the bad times as the times when output growth rates are less than -0.1% .³⁵ This is closer to the spirit of the categorization in the data, which is based on the output growth rate. The main properties of the results are unaltered from the main text, although the difference between good and bad states is more pronounced quantitatively.

	Good	Bad
Wage	1.014	0.986
q^*	0.5000	0.5000
Entry rate	8.9%	1.7%
Exit rate	5.2%	5.6%
Average size of all plants	84.8	86.2
Relative size of entrants	0.56	0.58
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.85	0.85
Relative productivity of exiting plants	0.84	0.84

Table 23: Results with aggregate shocks

	Good	Bad
Wage	1.010	0.990
q^*	0.5062	0.4938
Entry rate	7.5%	3.1%
Exit rate	5.2%	5.5%
Average size of all plants	85.2	86.1
Relative size of entrants	0.57	0.57
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.85	0.85
Relative productivity of exiting plants	0.84	0.84

Table 24: The case of a procyclical c_e

³⁵We do not set zero as a threshold, so that we do not capture the small movements of output shown in Figures 4 and 5.

	Good	Bad
Wage	1.010	0.990
q^*	0.4997	0.5003
Entry rate	7.4%	3.3%
Exit rate	5.2%	5.5%
Average size of all plants	85.2	86.0
Relative size of entrants	0.57	0.57
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.85	0.85
Relative productivity of exiting plants	0.84	0.84

Table 25: The case of a procyclical c_q

	Good	Bad
Wage	1.010	0.990
q^*	0.3047	0.6335
Entry rate	8.8%	2.6%
Exit rate	5.2%	5.6%
Average size of all plants	80.4	80.8
Relative size of entrants	0.47	0.71
Relative size of exiting plants	0.41	0.41
Relative productivity of entrants	0.77	0.95
Relative productivity of exiting plants	0.84	0.84

Table 26: The case of a countercyclical c_e and procyclical c_q

References

- [1] Ábrahám, Á. and T. K. White (2007). “The Dynamics of Plant-level Productivity in U.S. Manufacturing,” mimeo. University of Rochester and U.S. Bureau of the Census.
- [2] Baily, M. E. Bartelsman, and J. Haltiwanger, (2001). “Labor Productivity: Structural Change and Cyclical Dynamics,” *Review of Economics and Statistics* 83, 420–433.
- [3] Baily, M., C. Hulten, and D. Campbell (1992). “Productivity Dynamics in Manufacturing Plants.” *Brookings Papers on Economic Activity: Microeconomics* 187-249.

- [4] Barlevy, G. (2002). “The Sullyng Effect of Recessions,” *Review of Economic Studies* 69, 41–64.
- [5] Bartelsman, E. and M. Doms, (2000). “Understanding Productivity: Lessons from Longitudinal Micro Datasets,” *Journal of Economic Literature* 38, 569–594.
- [6] Bartelsman, E. J. and W. B. Gray (1996). “The NBER Manufacturing Productivity Database,” NBER Working Paper T0205.
- [7] Bernanke, B. and M. Gertler (1989). “Agency Costs, Net Worth, and Business Fluctuations,” *American Economic Review* 79, 14–31.
- [8] Bilbiie, F. O. F. Ghironi, and M. J. Melitz (2006). “Endogenous Entry, Product Variety, and Business Cycles,” mimeo. Oxford University, Boston College, and Princeton University.
- [9] Bills, M. and Chang, Y. (2000). “Understanding How Price Responds to Costs and Production,” *Carnegie-Rochester Conference Series on Public Policy* 52, 33–78.
- [10] Caballero, R. J. and M. L. Hammour (1994). “The Cleansing Effect of Recessions,” *American Economic Review* 84, 1350–1368.
- [11] Caballero, R. J. and M. L. Hammour (2005). “The Cost of Recessions Revisited: A Reverse-Liquidationist View,” *Review of Economic Studies* 72, 313–341.
- [12] Campbell, J. (1998). “Entry, Exit, Embodied Technology, and Business Cycles,” *Review of Economic Dynamics* 1, 371–408.
- [13] Carlstrom, C. T. and T. S. Fuerst (1997). “Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis,” *American Economic Review* 87, 893–910.
- [14] Chang, Y., J. Gomes, and F. Schorfheide, (2002). “Learning-by-Doing as a Propagation Mechanism,” *American Economic Review* 92, 1498–1520.

- [15] Chatterjee, S. and R. W. Cooper (1993). “Entry and Exit, Product Variety and the Business Cycle,” NBER Working Paper 4562.
- [16] Comin, D. and M. Gertler (2006). “Medium-Term Business Cycles,” *American Economic Review*, 96, 523–551.
- [17] Cooley, T.; R. Marimon; and V. Quadrini (2004). “Aggregate Consequences of Limited Contract Enforceability,” *Journal of Political Economy*, 112, 817–847.
- [18] Cooley, T. F. and Prescott, E. C. (1995). “Economic Growth and Business Cycles,” in Cooley, T. F., (ed.) *Frontiers of Business Cycle Research*, Chapter 1. Princeton University Press, Princeton.
- [19] Cooper, R., J. C. Haltiwanger, and J. Willis (2004). “Dynamics of Labor Demand: Evidence from Plant-level Observations and Aggregate Implications,” NBER Working Paper 10297.
- [20] Davis, S. J., J. C. Haltiwanger, and S. Schuh (1996). *Job Creation and Destruction*, Cambridge, MIT Press.
- [21] Devereux, M. B., A. C. Head, and B. J. Lapham (1996). “Aggregate Fluctuations with Increasing Returns to Specialization and Scale,” *Journal of Economic Dynamics and Control* 20, 627–656.
- [22] Dunne, T., J. C. Haltiwanger, and K. Troske (1997). “Technology and Jobs: Secular Changes and Cyclical Dynamics,” *Carnegie-Rochester Conference Series on Public Policy* 46, 107–178.
- [23] Dunne, T., M. J. Roberts, and L. Samuelson (1988). “Patterns of Firm Entry and Exit in U.S. Manufacturing Industries,” *RAND Journal of Economics* 19, 495–515.
- [24] Dunne, T., M. J. Roberts, and L. Samuelson (1989a). “Plant Turnover and Gross Employment Flows in the U.S. Manufacturing Sector,” *Journal of Labor Economics* 7, 48–71.

- [25] Dunne, T., M. J. Roberts, and L. Samuelson (1989b). “The Growth and Failure of U.S. Manufacturing Plants,” *Quarterly Journal of Economics* 104, 671–698.
- [26] Foster, L., J. Haltiwanger, and N. Kim (2006). “Gross Job Flows for the U.S. Manufacturing Sector: Measurement from the Longitudinal Research Database,” Center for Economics Studies Working Paper 06-30.
- [27] Foster, L., J. Haltiwanger, and C. Syverson, (2007). “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *American Economic Review*, forthcoming.
- [28] Hopenhayn, H. and R. Rogerson (1993). “Job Turnover and Policy Evaluation: A General Equilibrium Analysis,” *Journal of Political Economy* 101, 915–938.
- [29] Jaimovich, N. (2007). “Firm Dynamics, Markup Variations, and the Business Cycle,” mimeo. Stanford University.
- [30] Kiyotaki, N. and J. Moore (1997). “Credit Cycles,” *Journal of Political Economy* 105, 211–248.
- [31] Melitz, M. J. (2003). “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica* 71, 1695–1725.
- [32] Samaniego, Roberto M. (2006). “Entry, Exit and Business Cycles in a General Equilibrium Model,” mimeo. George Washington University.
- [33] Tauchen, G. (1986). “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters* 20, 177–181.
- [34] Veracierto, M. L. (2001). “Employment Flows, Capital Mobility, and Policy Analysis,” *International Economic Review* 42, 571–595.
- [35] Veracierto, M. L. (2002). “Plant-Level Irreversible Investment and Equilibrium Business Cycles,” *American Economic Review* 92, 181–197.

- [36] Veracierto, M. L. (2004). “Firing Costs and Business Cycle Fluctuations,” mimeo. Federal Reserve Bank of Chicago.