

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

Yoonsoo Lee
School of Economics
Sogang University
and
Federal Reserve Bank of Cleveland
ylee@sogang.ac.kr

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

February 2012

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 between 1972 and 1997. We find that the entry rate is more cyclical than the exit rate. We also find that the differences in productivity and employment between booms and recessions are particularly larger for entering plants than for exiting plants. Second, we build a general equilibrium model of industry dynamics and compare its predictions to the data. Finally, we explore the policy implications of the model.

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

JEL Classifications: E23, E32, L11, L60

*We thank Shutao Cao, Gian Luca Clementi, Tim Dunne, Jae Won Lee, Maggie Levenstein, Richard Rogerson, Sarah Tulman, Marcelo Veracierto, and the seminar participants at the Federal Reserve Bank of Cleveland, the Federal Reserve Bank of New York, the Federal Reserve Bank of Richmond, the IIOC 2008, the Korean Econometric Society, the Northeast Ohio Economics workshop, the Rochester Wegmans conference, the Third New York/Philadelphia Workshop on Quantitative Macroeconomics, UCSB, the University of Tokyo, Yonsei University, Sogang University, AEA 2009, and Workshop on Firm Dynamics and the Usage of T2LEAP 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. In particular, we focus on the plant-level dynamics over the business cycle.

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 the patterns of entry and exit over the business cycle in terms of rate, employment, and productivity. We find that entry rates differ significantly in booms and recessions. Furthermore, the differences in productivity and employment in booms and recessions are particularly larger for entering plants than the exiting plants. For example, the average size of entering plants (relative to the incumbents) during booms is about 25% smaller than the one during recessions. Moreover, plants entering in booms are about 10–20% less productive (in terms of the relative productivity to the incumbents) than those entering in recessions. Such differences are relatively small for plants exiting in booms or recessions.

The characteristics of entrants are among the important determinants of the size distribution of firms and establishments in an industry. Recent studies utilizing establishment-level data find that entry is an important source of aggregate productivity growth (see, e.g., Foster, Haltiwanger, and Krizan, 2001 and 2002). The fact that the plants that enter in recessions are different from those that enter in booms indicates that there is a much larger barrier to entry during recessions. Such a barrier may hurt the long-run growth of the economy.

It has long been argued that recessions have “cleansing” effects: low-productivity plants are scrapped during recessions, enhancing aggregate efficiency. Many recent papers have

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

provided alternative to this prevailing view. 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 reduce the amount of cumulative reallocation in the economy.

Focusing on the permanent shutdown, we do not find strong effects of cleansing from exit 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 rate suggests that the “insulation” effect at the entry margin predominates. In contrast to the finding on the exiting plants, the average size and productivity of entrants vary substantially 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 the selection at the exit margin, our new finding suggests that the selection at the entry (or “creation”) margin may be more important than the selection at the exit (or “destruction”) margin.

Based on our observations of heterogenous 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 industry dynamics model of Hopenhayn and Rogerson (1993) by incorporating aggregate productivity shocks. In order 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 cyclicalities of entry rates and exit rates. However, it turns out that it is necessary to assume cyclicalities of entry costs in order to account for the cyclical patterns of selection

in entry process. This is because of the following two reasons. First, when entry costs are constant over the cycle, 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. Second, even in a specification where the wage effect is partially shut down by the limited supply of entrants, the aggregate productivity shocks measured by Solow residuals do not have a sufficiently strong effect in order to generate *quantitatively* significant amount of selection. We show that the model with cyclical entry costs can account for the observed cyclical plant-level dynamics.

Some other recent papers extend a similar style of an industry 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, 2008) was among the first to incorporate aggregate productivity shocks to the Hopenhayn-Rogerson style model.² Veracierto (2002, 2008), however, assumes exogenous entry and exit, and therefore these 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

²As in Veracierto (2001), the models in Veracierto (2002, 2008) 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 and Floetotto, 2008), 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.

technology and exogenous exit.

Samaniego (2008) 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 the aggregate productivity. He finds that both entry and exit respond very little to the change in the aggregate productivity. This is in contrast to our result—in our model, the entry rate responds strongly to the 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 newly created plants, while we assume a constant marginal cost in our benchmark model.⁴ Our empirical evidence suggests that the constant marginal cost specification achieves an outcome closer to the behavior of the U.S. manufacturing plants in terms of the entry rates.

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, 2008)—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 the entry rates during booms and the 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

⁴A modified version of our model, explored in Appendix G can be viewed as an extreme case of increasing marginal cost.

rates during recessions are reduced. However, the variance of output increases dramatically.

The third experiment and the fourth experiment 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 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 the business cycle. 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.⁵

In this study, entering plants are new plants, which appear in the ASM or CM for the

⁵See Appendix A and Davis, Haltiwanger, and Schuh (1996) for details about the data.

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 details 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 growth-accounting analysis. Our plant-level TFP measurement closely follows Baily, Hulten, and Campbell (1992).⁶ 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 the growth accounting equation

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

We measure factor elasticities (α_k , α_n , and α_m) 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

⁶Without 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 (2008) for possible issues involved in using revenue-based productivity measures.

Table 1: Average size and productivity of plants

	Continuing	Entering	Exiting
Average size	87.5	50.3	35.0
Relative size	–	0.60	0.49
TFP based on (1)	–	0.96	0.86
TFP based on (2)	–	0.75	0.64
Labor productivity (using employment)	–	1.00	0.92
Labor productivity (using hours)	–	0.98	0.91

Note: The first row reports average employment (number of workers) for continuing, entering, and exiting plants. From the second through the sixth row, size (employment) and various measures of productivity relative to the industry average of continuing plants are reported.

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

$$\ln(s_t) = \ln(y_t) - \theta \ln(n_t). \tag{2}$$

In the model of Sections 3, the output y_t is interpreted as value added, and the measure of productivity from (2) directly matches our model specification.⁷

2.2 Employment and productivity of entering and exiting plants

2.2.1 Average employment and productivity statistics

First, we document the employment and productivity characteristics of entering and exiting plants. Those statistics are used to calibrate the steady state of the model. The first row of Table 1 documents the average size of the plants, in terms of the number of workers. Entering plants (using the time- t size of the plants which entered between time $t - 1$ and time t) and

⁷The production function $y_t = s_t n_t^\theta$ can be considered to be the relationship between y_t and n_t after all of the other variable inputs are taken into account. Suppose, for example, that the “true” production function is $y_t = \tilde{s}_t (x_t^\alpha n_t^{1-\alpha})^\phi$, where $\alpha \in (0, 1)$, $\phi \in (0, 1)$ and x_t is a variable input. Suppose that the price of x_t is r . Then, optimally choosing x_t and plugging the optimal solution into the “true” production function yields the relationship $y_t = s_t n_t^\theta$, where $\theta \equiv \phi(1 - \alpha)/(1 - \alpha\phi)$ and s_t is a function of \tilde{s}_t , α , ϕ , and r .

exiting plants (using the time- $(t - 1)$ size of the plants which exited between time $t - 1$ and time t) are much smaller than continuing plants (using the time- t size of the plants which survived from time $t - 1$ to time t). The second row of Table 1 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 SIC industry.⁸ 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. The third through sixth rows of Table 1 show various measures of relative productivity. This finding is one of the main contributions of our paper, because direct measures of productivity were not available at an annual frequency in the previous literature. Each cell in these rows represents the relative productivity (as compared to the four-digit SIC industry average of continuing plants) of entering and exiting plants. Two properties are consistently found across different productivity measures. First, entering and exiting plants are less productive than continuing plants (except for one case).⁹ Second, exiting plants are less productive than entering plants. These findings are consistent with the pattern of employment size in the first two rows of Table 1, provided that a productive plant employs more workers.

The third row in Table 1 is the TFP, based on (1). The fourth 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-

⁸By 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.

⁹Differences 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 a study focusing on a small number of industries where producer-level prices and quantities are observed separately, Foster, Haltiwanger, and Syverson (2008) 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.

Table 2: Size distribution of plants

	1 - 19	20 - 49	50 - 99	100 - 249	250 +
plants (numbers)	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
number of 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

Note: The first four rows in the upper panel are the shares of each size category in the sample (i.e., total number of establishments). Each row sums to one. The exit rate is the number of exiting establishments divided by the number of total establishments in each size category. In the lower panel, “number of plants, by age” and “employment, by age” respectively report the share of each category in the sample. The numbers in each group of 15 cells sum to one.

level labor share. The advantage of the measure based on (2) is that the measurements of output and employment are relatively more reliable than those for capital and material inputs. Moreover, we use this exact form of production function in Sections 3. Therefore, we mainly utilize this last measure of productivity in calibrating the model. The fifth and sixth rows are measures of labor productivity (output divided by labor input). The fifth row measures labor input by employment, and the sixth row measures labor input by hours.

Table 2 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 towards large plants. Hiring and firing are also concentrated in large plants. Although exit rates are higher in smaller plants, some large plants also exit. As we will discuss later, standard models of plant-level dynamics such as Hopenhayn and Rogerson (1993) cannot explain

this particular phenomenon. Because productivity and size have a one-to-one relationship (when adjustment costs and frictions are absent), very large plants have high productivity levels and do not exit in their model. The distribution of plant size and employment shares are also presented for each age category.¹⁰ 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 during 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.¹¹ The reason why 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 for which more than half of one year was 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 because our analysis stresses the cyclical movement of entry and exit rates, which are more related to the “change” than the “level.”

Table 3 displays the entry and exit rates of plants in good and bad times. The entry (exit) rate is measured by the number of entering (exiting) establishments as a percentage of the total number of establishments each period. The entry rate is much higher during booms than recessions. In contrast, exit rates are similar throughout the business cycle.¹²

¹⁰The 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 11 years or more.

¹¹Good 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. The years in parenthesis are not used because of the ASM panel rotation.

¹²In order to check the robustness of our result, we also consider an alternative division of booms and recessions—we follow the NBER business cycle dates. Because the NBER dates are specified in months, we

Table 3: Entry and exit rates

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

Note: Entry (exit) rate is measured by the number of entering (exiting) establishments as percentage of the total number of establishments each period.

Table 4: Job creation and job destruction 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
Total job creation	9.96	7.69	8.96
Job destruction from shutdowns	2.52	2.27	2.41
Job destruction from continuers	6.72	8.74	7.61
Total job destruction	9.24	11.01	10.02

Note: Job creation (destruction) rate is measured by the number of jobs created (destroyed) in each category of establishments (i.e., startups, continuers, shutdowns, and all establishments (total)) as percentage of total employment.

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 0.240.¹³

We also analyze cyclical patterns in annual job creation due to startups and job destruction due to shutdowns, which can be interpreted as employment-weighted entry and exit rates. Table 4 presents job creation and job destruction rates calculated from the published job flows data for our sample period (1972–1997).¹⁴ Job creation rate from startups is measured with weights according to the number of months that are in booms/recessions. For example, because there is a trough in March 1991, we count 2.5 months of 1991 as in a recession and 9.5 months of 1991 as in a boom. The statistics for 1991 receive the weights of 2.5/12 for bad times and 9.5/12 for good times. The pattern of Table 3 remains similar with this alternative division—the entry rates are 6.3% (Good) and 4.1% (Bad), while the exit rates are 5.5% (Good) and 5.0% (Bad).

¹³Using firm-level data from the Statistics of Canada, Hyunh, Petrunia, and Voia (2008) find a similar pattern of entry and exit rates. During a recession period for the Canadian economy between 1990 and 1993, the entry rate went down to its lowest level. However, the exit rate did not move as much during this time period.

¹⁴The job flows data are available from the webpage of John Haltiwanger, <http://www.econ.umd.edu/~haltiwan/download.htm>.

sured by the number of jobs created in entering establishments as percentage of the total employment in manufacturing in each period. Job creation rate from continuers is measured by the number of jobs created in continuing (expanding) establishments as percentage of the total employment. Total job creation rate is simply the sum of job creation rate from startups and job creation rate from continuers. Job destruction rates from shutdowns and continuers are measured in a similar way. We find that job creation rate from startups is much higher during booms, while job destruction rate 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 .¹⁵ Focusing on employment flows, Davis, Haltiwanger, and Schuh (1996) find that the job destruction rate is more cyclical than the job creation rate. Although we also find that 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, if we consider a plant as a production unit, the adjustment over the business cycle at the entry margin may be more important.

The first three rows of Table 5 describe average plant size (employment) of continuing, entering, and exiting plants during booms and recessions. In general, the 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.¹⁶ Compared to

¹⁵Using the aggregate job flows data from the earlier period (Davis, Haltiwanger, and Schuh, 1996), Campbell (1998) finds that *labor-weighted* entry rates (i.e., job creation rates from startups) are procyclical, whereas *labor-weighted* exit rates (i.e., job destruction rates from shutdowns) are countercyclical. Overall, *quarterly* job destruction rates from shutdown are negatively correlated with the percentage change in output (i.e., manufacturing output or real GDP), as discussed in Campbell (1998). However, in the latest panel (1994–1998) used in Davis, Haltiwanger, and Kim (2006), quarterly job destruction rates are positively correlated with the percentage change in output. Because we use annual data and also drop years between the ASM panels, we cannot directly compare our results to the previous studies using the quarterly data. When we examined the annual job creation and destruction data, we find that cyclical property of employment-weighted birth and death rates (measured as correlation with industry output) may change depending on the sample periods. See Table 17 in Appendix A.

¹⁶There were some outliers among entering plants in 1980. Because dropping a few outliers would cause disclosure issues, we chose to drop the whole year when calculating the average size and productivity of

Table 5: Average and relative size (employment) of continuing, entering, and exiting plants

	Good	Bad	Average
Average size, continuing	85.4	89.5	87.5
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

Note: Each column represents the average during good times, bad times, and the entire period. The relative size is obtained by dividing the average size of entering (exiting) establishments by the average size of continuing establishments in the same four-digit SIC industry.

entering plants in booms, entering plants in recessions start with about 30% more workers. In the fourth and fifth row, we report the size of entering and exiting plants, relative to continuing plants in the same four-digit SIC industry. In relative terms, the entering plants are about 25% smaller in recessions than in booms. 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 of entering and exiting plants, presented in Table 6, exhibits a similar pattern. In Table 6, we also examine the extent to which such fluctuations in relative productivity changes as the assumed returns to scale vary. In particular, in addition to the specifications (1) and (2) reported in the first two rows, we consider the case where the production function exhibits decreasing returns to scale:

$$\ln(s_t) = \ln(y_t) - \gamma(\alpha_k \ln(k_t) + \alpha_n \ln(n_t) + \alpha_m \ln(m_t)). \quad (3)$$

Here, γ is the returns to scale parameter. The specification (1) in the first row corresponds to the case in which $\gamma = 1$. Although the magnitude of the relative productivity became smaller

entering plants in Table 5. Because those outliers have substantially higher productivity levels, including them results in a much greater difference in entrants' productivity between booms and recessions, adding support to our finding. Although the results for average employment do not vary much with or without the outliers, we also dropped this year in Table 5 for consistency. Because the statistics for exiting plants are not affected by the outliers, we include the 1980 observations in the calculation.

Table 6: Relative productivity of entering and exiting plants

TFP based on:	Relative TFP, entering			Relative TFP, exiting		
	Good	Bad	Average	Good	Bad	Average
(1)	0.93	1.02	0.96	0.88	0.84	0.86
(2)	0.69	0.85	0.75	0.65	0.65	0.65
(3), $\gamma = 0.99$	0.92	1.01	0.96	0.87	0.83	0.86
(3), $\gamma = 0.95$	0.89	0.99	0.93	0.85	0.81	0.84
(3), $\gamma = 0.90$	0.86	0.96	0.90	0.82	0.79	0.80
(3), $\gamma = 0.85$	0.83	0.94	0.87	0.80	0.76	0.78

Note: The first row reports the relative TFP based on (1). The second row reports relative TFP based on (2). Last four rows report the relative TFP based on (3), with various assumed values of returns to scale (γ). Relative productivity of entering (exiting) plant is obtained by dividing the productivity of the entering (exiting) plant by the average productivity of continuing plants in the same four-digit industry.

as we reduced the value of the returns to scale parameter, the pattern of cyclical fluctuations in the relative productivity of entering and exiting plants remains the same: the relative productivity of exiting plants is similar across booms and recessions, whereas the relative productivity of entering plants is substantially different in the two phases of the cycle.¹⁷ The relative productivity of entering plants in recessions are about 10–20% higher than that of entering plants in booms.¹⁸ The simple correlation between the relative productivity of exiting plants and the percentage change in manufacturing output is -0.054 , whereas the simple correlation between the relative size of entering plants and the percentage change in manufacturing output is -0.232 .¹⁹

¹⁷We also examined cyclical changes in the relative TFP based on (2) with various assumed values of returns to scale. The observed pattern is similar to Table 6.

¹⁸Because we use a revenue-based productivity measure, 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.

¹⁹If we utilize the NBER business cycle dates as in footnote 12, the results are qualitatively similar but quantitatively less pronounced. The relative size of entering plants is 0.58 (Good) and 0.65 (Bad), while the relative size of exiting plants is 0.50 (Good) and 0.45 (Bad). The relative productivity (based on (2)) of entering plants is 0.74 (Good) and 0.79 (Bad), while the relative productivity of exiting plants is 0.64 (Good) and 0.63 (Bad).

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 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 their benchmark case), plants compare the value of staying with the value of exiting, which is zero. In their model, there is a threshold idiosyncratic 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 without frictions and adjustment costs, 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 2 in the previous section). Furthermore, with Hopenhayn and Rogerson’s formulation, an annual exit rate of 5.5% (see Table 3) implies that only very unproductive and small plants exit, which is at odds with the employment and productivity evidence in Table 5.²⁰

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.²¹ In the data, we observe that the productivity

²⁰Samaniego (2006) is the first to point out this problem. Samaniego (2006, 2008) assumes a stochastic continuation value, rather than a stochastic exit value that we employ, to cope with this problem. We have also experimented with a model that assumes a stochastic continuation value. The results are essentially the same.

²¹Melitz (2003) employs a similar selection process in entry. Here, the interpretation of this process is slightly different from Melitz (2003).

of entering plants is very different across booms and recessions (see Table 6). 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 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.²² After observing the exit value,

²²The entry cost that is introduced later can be interpreted as (partially sunk) investment in new capital and land.

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 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 .²³ If $n_t \neq n_{t-1}$, it pays adjustment costs (and a firing tax, if $n_t < n_{t-1}$).²⁴ 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 $\xi(x_t)$,²⁵

²³We abstract from the capital stock (aside from the entry cost) in the production function. This abstraction makes the computation of the model easier, and this formulation is consistent with our measurement of the plant-level productivity.

²⁴The details of the adjustment cost are explained later.

²⁵Formulating the exit decision using i.i.d. stochastic scrap values is popular in the empirical industrial organization literature—see, for example, Doraszelski and Pakes (2007) and Weintraub, Benkard, and Van

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

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

where

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

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

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

Roy (2008). Ramey and Shapiro (2001) analyze the resale prices of displaced capital. The resale prices vary substantially—see their Figure 2.

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

$$V^e(q_t^*; z_t) = c_e. \quad (4)$$

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.²⁶ 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:

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

²⁶Hopenhayn and Rogerson (1993) assume a period utility function that is concave in consumption and linear in leisure.

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:

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

3.3 General equilibrium

Now we analyze the general equilibrium of the model. The general equilibrium is defined as a situation where (i) consumers and firms (plants) optimize and (ii) the markets clear. For (ii), it is sufficient to ensure that the labor market clears.

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.²⁷ 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) . (Here, s' is the plant-level productivity at the current period and n is the plant-level employment one period before.) N is the actual mass of potential entry at each period.

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. Once L^* is determined, the equilibrium level of N , N^* , is determined by (8).

²⁷In 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.

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 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)

$$C_t = Y_t - N_t c_q - M_t c_e + X_t. \quad (10)$$

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. We set one period as one year. 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 18 in Appendix B). 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, which matches the autocorrelation parameter for the AR(1) process for employment (simulated in the model) to the empirical value of 0.97.²⁸ 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).²⁹ 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. Thus we only include the “disruption cost” type of adjustment cost.

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 for the incumbents: $\eta(s'|q) = \psi(s'|s)$. The entry costs, c_q and c_e , are backed out from the

²⁸See equation (11) and Table 18 in Appendix B. In the Appendix, it is shown that the parameter estimates for the employment AR(1) process can differ depending on which econometric method we employ. In the model, we also experimented with lower values of ρ_s . The problem with lower values of ρ_s is it is impossible to replicate the steady-state distribution of plant size (in particular, there are too few large plants). One remedy for this would be to incorporate a plant-level fixed effect in s_t , reflecting the heterogeneity in the “planned size” of the plants. We did not explore this direction due to computational complexity.

²⁹This is their point estimate with a small quadratic adjustment cost. Their point estimate for λ with no other adjustment cost is 0.988. Although we do not have a quadratic adjustment cost in our model, we prefer the former number because it produces a more reasonable job reallocation rate.

Table 7: Benchmark parameters

β	θ	a_s	ρ_s	σ_s	λ	c_e	c_q
0.94	0.7	0.04	0.97	0.11	0.983	872.9	103.1

Table 8: Data and model statistics in the steady-state

	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%

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).³⁰ 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 $\ln(\cdot)$. Table 7 lists the main parameter values for the benchmark case. Note that c_e and c_q are measured in annual wages, because we normalized to $w = 1$ in equilibrium.

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 C. Table 8 compares the output of our model to the data.³¹ Everything except for the job reallocation rate is our “target” for calibration, and we can see that we are able

³⁰We 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.

³¹The job reallocation rate is taken from Davis, Haltiwanger, and Schuh (1996, Table 2.1).

Table 9: Size distribution of plants, data and model statistics

	Number		Employment		Number (Young)		Exit rate	
	Data	Model	Data	Model	Data	Model	Data	Model
1-19	0.457	0.478	0.049	0.038	0.712	0.612	0.080	0.067
20-49	0.239	0.187	0.090	0.064	0.156	0.151	0.039	0.062
50-99	0.131	0.107	0.109	0.077	0.068	0.079	0.031	0.056
100-249	0.106	0.155	0.194	0.284	0.044	0.117	0.025	0.039
250+	0.067	0.072	0.559	0.536	0.020	0.041	0.015	0.004

Note: The first two columns report the number of plants in each size class. Each cell is the fraction of the total number of plants (each column sums to one). The third and fourth columns report employment share by size class. Each cell represents the share of employment for each size category in total employment (each column sums to one). The fifth and sixth columns report the number of young plants within each size class. The last two columns report exit rates for each size class.

to get close to the empirical values. The job reallocation rate is also close to the value in the data.³²

Table 9 evaluates the distributional performance of the model. The first two columns summarize the distribution of plant size by the fraction of plants in each size class. Although there are a few cells that do not exhibit a perfect match with the data, overall the size distribution of the model matches the data quite well. The third and fourth columns describe the employment share by size class. It also exhibits a good match.

Because 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. The fifth and sixth columns show the distribution of young plants in the model and the data, which are reasonably similar.³³ The last two columns show that the exit rates of the model also exhibit a good match to the data.

³²If we assume that $\lambda = 1.0$, the job reallocation rate increases to over 30%.

³³The 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.

Table 10: Results with aggregate shocks

	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

3.6 Adding aggregate shocks

To analyze business cycles, we assume that z_t fluctuates between two values.³⁴ We assume that z_t takes either 1.01 or 0.99. This results in a 1% standard deviation in z_t .³⁵ z_t follows a symmetric Markov process. We calibrate the transition probabilities so that the average duration of each state is three years.³⁶

The computation turns out to be much simpler than standard heterogeneous-agent models, such as Krusell and Smith (1998), since the wage depends only on z (this is thanks to the utility function that is linear in consumption and the free entry assumption). 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 E.

The results of the model with aggregate shocks are summarized in Table 10. Here, “Good” corresponds to the periods with $z_t = 1.01$ and “Bad” corresponds to the periods

³⁴Appendix D shows that the results below are robust to having more than two possible values of z_t .

³⁵Appendix D shows that the estimated unconditional standard deviation of $\ln(z_t)$ during our sample period (derived from the AR(1) regression of the Solow residuals) is about 1.1%.

³⁶The average duration of the post-war (1945-2009) NBER contraction (peak to trough) is 11 months and NBER expansion (trough to peak) is 59 months. Thus, overall, the average duration of each state is 35 months.

with $z_t = 0.99$.³⁷ First, notice that the wage fluctuates substantially. While the cyclicity of wages is empirically controversial, in Cooley and Prescott (1995) the wages are procyclical and 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—in (7), for L to increase when A stays the same, we need w to increase. Somewhat surprisingly, in Table 10, the equilibrium value of q^* does not change with the change in z . One reason for this result is that 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. 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 cyclicity 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 distributions of the value of staying and the value of exiting are both quite dispersed, a 1% difference in z does not make a large difference for this comparison.³⁸ Thus the exit rate and the size and productivity of exiting plants are similar throughout the business cycle in Table 10. This fits well with the pattern we observe in the data.

The entry rate fluctuates significantly in Table 10, as we see in the data. The mechanism here is simple: since the wage increases during booms, the average size of incumbent plants shrinks (we observe in the data as well). The labor that is released from the incumbents can be hired by the entrants. At the same time, the labor supply increases because of the wage increase. As a result, the entry rate goes up. Here, again, the procyclicity of wages plays

³⁷In Appendix F, we consider a different method of dividing good times and bad times—there, 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% . The following results are robust to this alternative categorization.

³⁸If both values are concentrated around one value and there are many “marginal” plants around that value, it is possible that these plants exit with a small change in z .

an important role.

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.

Thus, it seems natural to consider a modification of the model that reduces the fluctuations of the wages as a first step. How 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.³⁹

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 initial model, we assumed that the entry costs, c_q and c_e , are constant. If either c_q or c_e is cyclical, the fluctuation of wages can be smaller. In the following, we explain how changing the entry costs affects the aggregate performance of the model.⁴⁰

First, we consider a 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

³⁹The 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 .

⁴⁰An alternative modeling strategy is to remove the free-entry assumption. Veracierto (2001) and Samaniego (2006) consider a convex cost of producing new plants. Appendix G explores the opposite extreme of free entry along this line: the pool of potential entrants is fixed at each period. (A similar assumption is employed by Clementi and Palazzo (2010) who extend our model and incorporate capital stock.) This can be seen as the extreme case of convex cost for becoming a potential entrant—the cost is zero until a fixed amount \bar{N} , and infinity after \bar{N} . It turns out that while this modification reduces the volatility of wages, this variant of the model is also incapable of generating sufficient degree of selection over the business cycle.

Table 11: The case of a procyclical c_q

	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

booms, partly because the wages for these workers are higher then.⁴¹ In addition, there are more entries and idea creations during booms, and the idea creation process may suffer from decreasing returns (due to, for example, the “fishing-out” effect). Table 11 shows the result. Here, c_q is 0.165% larger during booms and 0.165% smaller during recessions, targeting 1% cyclicity of the wages. We can see that this generates a *qualitatively* successful result. The selection goes in the right direction— q^* is countercyclical. However, it does not generate a selection effect *quantitatively* large enough to match the data.

How can we generate a stronger selection effect? From the potential entrant’s entry condition (4), we know that c_e has a direct effect on the selection process—a large c_e makes the actual entry harder. Then, a countercyclical c_e would help to generate a larger selection effect.⁴² In fact, empirical studies on investment costs 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

⁴¹The 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).

⁴²We also experimented with procyclical c_e . 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.

Table 12: The case of a countercyclical c_e and procyclical c_q

	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

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.⁴³

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 target the wage to have 1% standard deviation), 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 12 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 are comparable to the data.

Figure 1 draws 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

⁴³See, for example, Bernanke and Gertler (1989), Carlstrom and Fuerst (1997), and Kiyotaki and Moore (1997).

Figure 1: A sample path of aggregate output






Figure 2: A sample path of aggregate output (magnified: a boom)



closely, however, there are nontrivial dynamics (although the magnitude is very small) within each aggregate state. Figure 2 magnifies a part of Figure 1, by picking up on one boom. We can see “hump-shaped” dynamics.⁴⁴ Output increases in the period following switching, and 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

⁴⁴Chang, 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.

Figure 3: A sample path of aggregate output (magnified: a recession)



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. Figure 3 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 more important 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 12 our “baseline case.”) 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 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.

As we discuss below, the effect of a firing tax has been examined in the existing literature. We conduct the firing tax analysis in order to analyze how the (endogenously) cyclical entry of plants affects the aggregate effect of a firing tax. The analysis of an entry subsidy is conducted in order to examine the aggregate implications of cyclical policies that affect the entry process directly.

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 (2008) 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 (2008) also conducts a similar exercise.

The results from our model with 10% firing tax are shown in Table 13. Entry and

Table 13: Results with a constant firing tax (10%)

	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

exit behavior do not experience quantitatively large changes compared to the baseline case, 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 in statistics that we usually associate with the firing costs. The job reallocation rate falls from 23.0% (in the baseline case) 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%.⁴⁵ This contrasts with Veracierto (2008) and Samaniego (2008), 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.⁴⁶ Intuitively, the firing tax is a tax on relatively large plants, which are more likely to fire workers in the near future. During recessions, entering plants are typically larger than entering plants during booms, so they experience a larger expected 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

⁴⁵The coefficient of variation also increases, since the mean becomes smaller with the firing cost.

⁴⁶Veracierto 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.

Table 14: Results with a firing tax, during recessions only (10%)

	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 12 and Table 13: q^* decreases in good times and increases in bad times. The general equilibrium effects also 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 14. The government succeeds in its intention—the average job destruction rate during recessions falls to 9.7% versus 11.6% in the baseline case. However, as we can see from the table, the entry and exit rates fluctuate more than in the baseline case, 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

Table 15: Results with entry subsidies on c_e and c_q only during recessions (0.1%)

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

	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

an investment subsidy. The results are summarized in Table 15. Entry rates are less volatile compared to the baseline case. The variance of output is substantially reduced—it becomes less than half of the variance found in the baseline case. The selection of entrants during recession is not as stringent as in the baseline case: 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 16. Again, the government can achieve stability in entry rates. The variance of output is also small—less than half of the baseline case. The difference from the previous experiment is that here the selection of entering plants is more stringent during recessions than in the baseline case.

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. 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 certain assumptions are made about the cyclical nature 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 costs) and empirically (by looking into the microeconomic process of entry).⁴⁷

⁴⁷Explicitly modeling the limited enforceability of contracts, as in Cooley, Marimon, and Quadrini (2004), is one possible direction.

We employed a stationary model and abstracted from the secular productivity growth. Foster, Haltiwanger, and Krizan (2001) show that a significant part of the productivity growth in the U.S. manufacturing sector comes from entry and exit of plants. Extending our analysis to incorporate secular productivity growth may open up a new possible link between aggregate fluctuations and aggregate growth, as in Barlevy (2004).

Our finding that the productivity of entrants vary over the business cycles may have important asset pricing implications. In a recent paper, Gourio (2007) argues that, in his putty-clay investment model, the relative labor productivity (which is determined by the capital intensity) of new production units has to be countercyclical in order to account for the procyclical stock prices. This cyclical pattern of new production units is consistent with our finding.

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

References

- [1] Abraham, 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] Anderson, T. and Hsiao, C. (1981). “Estimation of Dynamic Models with Error Components,” *Journal of the American Statistical Association* 76, 598–606.
- [3] Arellano, M. (2003). *Panel Data Econometrics*, Oxford University Press: Advanced Texts in Econometrics, Oxford.
- [4] Arellano, M. and Bond, S. (1991). “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations,” *Review of Economic Studies* 58, 277–297.
- [5] Baily, M. E. Bartelsman; and J. Haltiwanger, (2001). “Labor Productivity: Structural Change and Cyclical Dynamics,” *Review of Economics and Statistics* 83, 420–433.
- [6] Baily, M.; C. Hulten; and D. Campbell (1992). “Productivity Dynamics in Manufacturing Plants.” *Brookings Papers on Economic Activity: Microeconomics* 187-249.
- [7] Barlevy, G. (2002). “The Sullyng Effect of Recessions,” *Review of Economic Studies* 69, 41–64.
- [8] Barlevy, G. (2004). “The Cost of Business Cycles under Endogenous Growth,” *American Economic Review* 94, 964–990.
- [9] Bartelsman, E. and M. Doms, (2000). “Understanding Productivity: Lessons from Longitudinal Micro Datasets,” *Journal of Economic Literature* 38, 569–594.
- [10] Bartelsman, E. J. and W. B. Gray (1996). “The NBER Manufacturing Productivity Database,” NBER Working Paper T0205.
- [11] Bernanke, B. and M. Gertler (1989). “Agency Costs, Net Worth, and Business Fluctuations,” *American Economic Review* 79, 14–31.

- [12] 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.
- [13] Bils, M. and Chang, Y. (2000). “Understanding How Price Responds to Costs and Production,” *Carnegie-Rochester Conference Series on Public Policy* 52, 33–78.
- [14] Blundell, R. and Bond, S. (1998). “Initial conditions and moment restrictions in dynamic panel data models,” *Journal of Econometrics* 87, 115–143.
- [15] Bond, S. (2002). “Dynamic panel data models: a guide to microdata methods and practice,” *Portuguese Economic Journal* 1, 141–162.
- [16] Caballero, R. J. and M. L. Hammour (1994). “The Cleansing Effect of Recessions,” *American Economic Review* 84, 1350–1368.
- [17] Caballero, R. J. and M. L. Hammour (2005). “The Cost of Recessions Revisited: A Reverse-Liquidationist View,” *Review of Economic Studies* 72, 313–341.
- [18] Campbell, J. (1998). “Entry, Exit, Embodied Technology, and Business Cycles,” *Review of Economic Dynamics* 1, 371–408.
- [19] 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.
- [20] Chang, Y.; J. Gomes; and F. Schorfheide, (2002). “Learning-by-Doing as a Propagation Mechanism,” *American Economic Review* 92, 1498–1520.
- [21] Chatterjee, S. and R. W. Cooper (1993). “Entry and Exit, Product Variety and the Business Cycle,” NBER Working Paper 4562.
- [22] Clementi, G. L. and Palazzo, D. (2010). “Entry, Exit, Firm Dynamics, and Aggregate Fluctuations,” mimeo. New York University and Boston University.

- [23] Comin, D. and M. Gertler (2006). “Medium-Term Business Cycles,” *American Economic Review*, 96, 523–551.
- [24] Cooley, T.; R. Marimon; and V. Quadrini (2004). “Aggregate Consequences of Limited Contract Enforceability,” *Journal of Political Economy*, 112, 817–847.
- [25] 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.
- [26] 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.
- [27] Davis, S. J.; J. C. Haltiwanger; and S. Schuh (1996). *Job Creation and Destruction*, Cambridge, MIT Press.
- [28] 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.
- [29] Doraszelski, U. and A. Pakes (2007). “A Framework for Applied Dynamic Analysis in IO,” in M. Armstrong and R. Porter (eds.) *Handbook of Industrial Organization, Volume 3*, Amsterdam: North-Holland.
- [30] 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.
- [31] 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.

- [32] 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.
- [33] 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.
- [34] 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.
- [35] Foster, L.; J. Haltiwanger; and C. J. Krizan (2001). “Aggregate Productivity Growth: Lessons from Microeconomic Evidence,” in C. R. Hulten, E. R. Dean, and M. J. Harper (eds.) *New Developments in Productivity Analysis*, Chicago: University of Chicago Press.
- [36] Foster, L.; J. Haltiwanger; and C. J. Krizan (2002). “The Link Between Aggregate and Micro Productivity Growth: Evidence from Retail Trade,” Working Papers 02-18, Center for Economic Studies, U.S. Census Bureau.
- [37] Foster, L.; J. Haltiwanger; and C. Syverson, (2008). “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?” *American Economic Review* 98, 394–425.
- [38] Gourio, F. (2007). “Putty-Clay Technology and Stock Market Volatility,” mimeo. Boston University.
- [39] Hopenhayn, H. and R. Rogerson (1993). “Job Turnover and Policy Evaluation: A General Equilibrium Analysis,” *Journal of Political Economy* 101, 915–938.
- [40] Hyunh, K.; R. Petrunia; and M. Voia (2008). “The Impact of Initial Financial State on Firm Duration Across Entry Cohorts,” mimeo. Indiana University, Lakehead University, and Carleton University.

- [41] Jaimovich, N. and M. Floetotto (2008). “Firm Dynamics, Markup Variations, and the Business Cycle,” *Journal of Monetary Economics*, forthcoming.
- [42] Kiyotaki, N. and J. Moore (1997). “Credit Cycles,” *Journal of Political Economy* 105, 211–248.
- [43] Krusell, P. and A. A. Smith, Jr. (1998), “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy* 106, 867–896.
- [44] Melitz, M. J. (2003). “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica* 71, 1695–1725.
- [45] Ramey, V. A. and M. D. Shapiro (2001). “Displaced Capital: A Study of Aerospace Plant Closings,” *Journal of Political Economy* 109, 958–992.
- [46] Samaniego, R. M. (2006). “Do Firing Costs affect the Incidence of Firm Bankruptcy?” *Macroeconomic Dynamics* 10, 467–501.
- [47] Samaniego, R. M. (2008). “Entry, Exit and Business Cycles in a General Equilibrium Model,” *Review of Economic Dynamics* 11, 529–541.
- [48] Tauchen, G. (1986). “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters* 20, 177–181.
- [49] Veracierto, M. L. (2001). “Employment Flows, Capital Mobility, and Policy Analysis,” *International Economic Review* 42, 571–595.
- [50] Veracierto, M. L. (2002). “Plant-Level Irreversible Investment and Equilibrium Business Cycles,” *American Economic Review* 92, 181–197.
- [51] Veracierto, M. L. (2008). “Firing Costs and Business Cycle Fluctuations,” *International Economic Review* 49, 1–39.
- [52] Weintraub, G. Y.; C. L. Benkard; and B. Van Roy (2008). “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica* 76, 1375–1411.

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 in the timing of startup or shutdown.⁴⁸

It is 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.⁴⁹

⁴⁸Foster, Haltiwanger, and Kim (2006) also uses information from the Longitudinal Business Database (LBD) to identify births and deaths.

⁴⁹We also constructed job creation and destruction rates based on our measure of birth and death and compared those to Foster, Haltiwanger, and Kim (2006). Our job creation and destruction rates are highly correlated with Foster, Haltiwanger, and Kim (2006) and show a cyclical pattern that is similar to theirs.

Recently the Census Bureau has developed the Longitudinal Business Database (LBD), which is a longitudinal version of the Census business register. While the LBD contains more accurate information about the timing of births and deaths, we were not authorized to access the LBD during the course of this project.

ASM sample weights: We use ASM sample weights to measure representative entry and exit rates as well as average size and productivity statistics. The sample weights in the ASM are related to the probability of selection and are intended to create a representative sample of establishments in terms of size (i.e., shipments). We apply sample weights to birth and death in the manner that has been used in gross job flow statistic (Davis, Haltiwanger, and Schuh, 1996). While it is a more reliable way of using the ASM than simply counting the number of births and deaths, appropriate caution is needed because sample weights applied to birth and death rates are activity-weights and may result in birth and death rates that are less *strictly* representative of the manufacturing sector than average statistics of employment or productivity are.

Job creation from startups and job destruction from shutdowns: Campbell (1998) used Davis, Haltiwanger, and Schuh (1996) measures of job creation at entering plants and destruction at exiting plants to study cyclical patterns of employment-weighted entry and exit rates. Using the updated job creation and destruction data, we examine the cyclicity of employment-weighted entry and exit rates across different sample periods. Table 17 reports simple correlations between manufacturing output and annual job creation (destruction) rate from startup (shutdown). The correlation between the job destruction rate due to shutdowns and the percentage change in manufacturing output is -0.048 in Campbell's (1998) sample period (i.e., 1972–1988) while that has a positive value of 0.184 between 1989 and 1998. The correlation between the job creation rate due to startups and the percentage change in manufacturing output is higher for the sample period between 1989 and 1998 than the one for the sample period used in Campbell (1998). While annual job destruction rates due

Table 17: Cyclicalities of Job Creation from Startups and Job Destruction from Shutdowns – Correlations with Industry Output

	Job Creation from Startups	Job Destruction from Shutdowns
1972–1998	0.382	-0.025
1972–1988	0.247	-0.048
1989–1998	0.752	0.184

Note: The correlation between annual manufacturing output growth rate and annual job creation (destruction) rates from startups (shutdowns) for three different sample periods. Campbell’s (1998) results are based on the sample period between 1972 and 1988 (second row).

to shutdowns are positively correlated with output in the later time period, we find that quarterly job destruction rates due to shutdowns are negatively correlated with output. It is worth noting that these entry and exit rates, based on job creation and destruction data, are employment-weighted measures. Simple entry and exit rates (that are not employment-weighted) are not available in quarterly frequency.

Manufacturing output: We use the growth rate of manufacturing output to divide sample years into good and bad years. To measure the changes in the deflated output, we aggregate the 4-digit SIC level real value of shipments, deflated by the 4-digit SIC level price deflator. We use industry-level price because aggregate price deflator may reflect changes in the composition of industries with different prices.

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 for which 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 by 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 measure of productivity based on (2), we use value added deflated by shipment deflators. 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 Estimating employment dynamics

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

$$\ln(n_{i,t+1}) = \alpha + \rho \ln(n_{it}) + \varepsilon_{i,t+1}, \quad (11)$$

where $\varepsilon_{i,t+1} \sim N(0, \sigma^2)$. We report estimates of the AR(1) process of employment dynamics in Table 18. Because only large plants are observed across two different ASM panels, the results exclude samples between two ASM panels in a way that is similar to our calculation of entry and exit statistics. The Ordinary Least Squares (OLS) estimator of ρ is very close to one—the employment process is very persistent. The results were similar when we corrected for selection bias using the maximum likelihood estimates from the Heckman selection model.⁵⁰ However, in the presence of permanent individual effects in the error term (i.e., $\varepsilon_{i,t+1} = \eta_i + v_{i,t+1}$), the OLS estimation of (11) leads to inconsistent estimation of ρ .⁵¹ In particular, standard results for omitted variable bias indicate that the OLS estimator is likely to be biased upwards. We estimate the persistence parameter by applying various econometric methods that are known to eliminate such inconsistency (see Arellano (2003) or Bond

⁵⁰We also calculated this process for the productivity measures in Table 1 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, because 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.

⁵¹For example, different plants may have been built with different intended scales. Such differences may add a permanent plant-specific effect to the error term. We thank one of the referees for pointing out this issue.

(2002) for a detailed discussion of estimation).

Table 18: Employment dynamics, AR(1)

	[1]	[2]	[3]	[4]	[5]
ρ	0.972	0.493	0.548	0.324	0.992
σ	0.379	0.267	0.409	1.05	0.380
	OLS levels	Within Groups	2SLS levels	GMM Diff	GMM System
m1	-33.20	-26.84	-6.33	-13.82	-38.82
m2	-1.14	-38.70	2.93	2.94	3.00
Instruments			n_{t-3}	n_{t-3} n_{t-4}	n_{t-3} n_{t-4}

Note: Table reports the estimate of ρ and σ in $\ln(n_{t+1}) = \alpha + \rho \ln(n_t) + \varepsilon_{t+1}$, where $\varepsilon_{t+1} \sim N(0, \sigma^2)$. m1 and m2 are test statistics for first-order and second-order serial correlation, asymptotically $N(0,1)$. All estimations include year, plant age, and industry dummies.

The Within Groups estimator (or fixed effect estimator) in Column [2] is much smaller than the OLS estimator. But it is also known that the Within Groups estimator is likely to be biased downwards, especially when the panel is short (i.e., T is small). One remedy to fix the bias is to use an instrumental variables estimator with a lagged variable (Anderson and Hsiao, 1981). We report the Two Stage Least Squares (2SLS) estimator using n_{t-3} as an instrument. The basic first-differenced 2SLS estimator in Column [3] is close to the Within Groups estimator in Column [2]. The first-differenced GMM estimator, known as the Arellano-Bond Estimator in the microeconometrics literature, improves the 2SLS estimator by using additional lags of the dependent variable as instruments. In Column [4], we report the estimate when we used n_{t-3} and n_{t-4} as instruments. The estimates were similar when more and longer lags were used as instruments. While the estimate is even lower than the Within Groups estimator, it is known that the first-differenced GMM is likely to be subject to a serious bias and weak instruments problem when the estimate is close to 1 and T is small (which is very likely in our case). Blundell and Bond (1998) show that using additional moment conditions may improve the estimator substantially in such case. The system GMM

estimator in Column [5] is very close to 1 and similar to the OLS estimate in Column [1].

Overall we believe that the system GMM estimator is the most robust estimator if the specification is correct and the employment process follows an AR(1) process. However, the results of the specification test developed by Arellano and Bond (1991) suggest that a critical assumption is violated in certain estimations in Table 18. We assume that disturbances excluding the fixed effect ($v_{i,t+1}$) are serially uncorrelated in each estimation. However, the test statistics for first-order (m1) and second-order (m2) serial correlation in Table 18 suggest that the null hypotheses of no first-order correlation and no second-order serial correlation are rejected in every case with the exception of OLS. The rejection of first-order serial correlation is consistent with the assumption of AR(1) and is not an issue. The rejection of second-order serial correlation does not necessarily indicate that the AR(1) model is mis-specified in the case of the Within Groups estimator, because the Within Groups estimator is biased and the estimates of the first-differenced residuals are likely to be biased. However for the rest of the specifications, significant second-order serial correlation suggests that the disturbances are serially correlated and the AR(1) model may be mis-specified.⁵²

In Table 19 we report the system GMM estimator when the employment process is assumed to follow an AR(2) process. We report the estimates when we used n_{t-5} and n_{t-6} (Column 1) or n_{t-6} and n_{t-7} (Column 2) as instruments. The second-order serial correlation test was rejected when shorter lags (e.g. n_{t-4} and n_{t-5}) were used as instruments. The estimator for ρ_1 is smaller now at around .9 than the OLS estimator or the system GMM estimator in Table 18. But the combined persistence measures ($\rho_1 + \rho_2$) are 0.981 and 0.982, respectively. These are very close to the OLS estimator in the AR(1) case.

⁵²The p-value of the m2 statistics for the first-differenced GMM is 0.003 and that of the system GMM is 0.003. Sargan tests of overidentification restriction are rejected for all specifications and the results are not reported here. According to Arellano and Bond (1991), their simulation results suggest that those identification tests reject too often in the presence of heteroskedasticity, which is very likely in establishment-level data.

Table 19: Employment dynamics, AR(2)

	[1]	[2]
ρ_1	0.901	0.910
ρ_2	0.080	0.072
σ	0.358	0.359
m1	-7.72	-7.83
m2	-0.63	-0.53
Instruments	n_{t-5}	n_{t-6}
	n_{t-6}	n_{t-7}

Note: Table reports the estimates of AR(2) process, $\ln(n_{t+1}) = \alpha + \rho_1 \ln(n_t) + \rho_2 \ln(n_{t-1}) + \varepsilon_{t+1}$, where $\varepsilon_{t+1} \sim N(0, \sigma^2)$. m1 and m2 are test statistics for first-order and second-order serial correlation, asymptotically $N(0,1)$. All estimations include year, plant age, and industry dummies.

C 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

$$\pi^n(s', n) = zf(n, s') - wn.$$

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

$$V^e(q^*) = c_e. \tag{12}$$

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:

$$V_p = c_q. \quad (13)$$

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

$$Av'(1 - L(N)) = w. \quad (14)$$

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

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

D The results with more than two points of z

In the main text, we assumed that the aggregate productivity z_t follows a two-point Markov chain. In this appendix, we check the robustness of the results in terms of this assumption. Our strategy is to estimate an AR(1) process for Solow residual as in the standard Real Business Cycles literature, and approximate it by a 10-point Markov chain using Tauchen’s (1986) method.

The AR(1) regression is

$$\ln(z_{t+1}) = a_0 + a_1 t + \rho \ln(z_t) + \varepsilon_{t+1},$$

where $\varepsilon_{t+1} \sim N(0, \sigma_\varepsilon^2)$. z_t here is computed from the growth accounting equation in annual frequency

$$\ln(z_t) = \ln(Y_t) - 0.7 \ln(L_t),$$

where Y_t is the real GDP, L_t is the index of aggregate weekly hours for production and nonsupervisory employees in total private industries. When the data from $t + 1 = 1966$ to $t + 1 = 2010$ is used, the point estimate for ρ is calculated as $\hat{\rho} = 0.754$ and the point estimate for σ_ε is calculated as $\hat{\sigma}_\varepsilon = 0.0094$.⁵³ We put 10 equally-spaced grids on $\ln(z_t)$ over the interval $[-1.5\hat{\sigma}_\varepsilon/\sqrt{1-\hat{\rho}^2}, 1.5\hat{\sigma}_\varepsilon/\sqrt{1-\hat{\rho}^2}]$. Then we construct the Markov transition matrix following Tauchen (1986).

Table 20 is the benchmark result (that is, when c_e and c_q are constant). Here, “Good” time is the period t in which z_t is above average and “Bad” time is the period t in which z_t is below average. The results are very similar (also in the experiment later) when we define “Good” time as the period t in which the output increased by more than 0.1% from $t - 1$ to t and define “Bad” time as the period t in which the output decreased by more (in absolute value) than 0.1% from $t - 1$ to t . The message is the same as the body of the paper: the

⁵³This implies that implied unconditional standard deviation of $\ln(z)$, $\sigma_\varepsilon/\sqrt{1-\hat{\rho}^2}$, is 0.014. When the data from $t + 1 = 1972$ to $t + 1 = 1997$ is used, the corresponding value is 0.011. These are in line with the value of the standard deviation of the process used in the body of the paper (0.01).

Table 20: Results with aggregate shocks

	Good	Bad
Wage	1.015	0.986
q^*	0.5000	0.5000
Entry rate	6.1%	4.6%
Exit rate	5.3%	5.4%
Average size of all plants	84.2	86.8
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 21: The case of a countercyclical c_e and procyclical c_q

	Good	Bad
Wage	1.010	0.992
q^*	0.3287	0.6082
Entry rate	6.1%	4.1%
Exit rate	5.5%	5.5%
Average size of all plants	74.8	81.2
Relative size of entrants	0.54	0.69
Relative size of exiting plants	0.40	0.41
Relative productivity of entrants	0.82	0.94
Relative productivity of exiting plants	0.83	0.84

model successfully replicates the patterns of the entry rate and the exit rate, but fails to generate any fluctuations in the selection of the entrant plants.

Table 21 is the outcome with countercyclical c_e (1.4% smaller than the steady-state value when z_t is the largest and 1.4% larger than the steady-state value when z_t is the smallest) and procyclical c_q (7% larger than the steady-state value when z_t is the largest and 7% smaller than the steady-state value when z_t is the smallest). c_e and c_q move together with z_t , and the grids for c_e and c_q are placed with the equal spacing. Again, the message is the same as the body of the paper: with countercyclical c_e and procyclical c_q , the model can replicate the behavior of the plant entry and exit in the data, including the size and productivity patterns

of the entrants over the business cycle.

E 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 in order 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.

F Results for significant positive and negative growth states

In this section (Tables 22, 23, and 24), we report the results of the experiments originally reported in Tables 10 to 12, 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% .⁵⁴ 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.

⁵⁴We do not set zero as a threshold, so that we do not capture the small movements of output shown in Figures 2 and 3.

Table 22: Results with aggregate shocks

	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: The case of a procyclical c_q

	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

G An alternative entry assumption: Fixed pool of potential entrants

Here, we consider a model with alternative entry process. In this model, we assume that the number of potential entrants in each period is fixed at a level \bar{N} . That is, in every period there are a fixed number of potential entrants, each with its own q_t . Because the distribution of q_t is unchanged, entry and selection are inversely related—when q_t^* is high, there are more stringent selection into entry (the average productivity of entrants is high) and the entry rate is low. When q_t^* is low, the average productivity of entrants is low and the entry rate is high.

Table 24: The case of a countercyclical c_e and 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

This model is more difficult to analyze as compared to the model in Section 3. (We call the model with constant entry costs in Section 3 as “the basic model” hereafter.) This is because the wage w_t is now simultaneously determined with L_t in labor market equilibrium. Here, we cannot separate the determination of w_t and q^* from the determination of equilibrium quantities as in Section 3. This, in turn, implies that the plants which are making dynamic decisions have to predict what will happen to the equilibrium prices and quantities in the future. Because the equilibrium quantities in the future depend on the distribution of plants in the current period, the entire distribution of the plants enters into the set of relevant information for the plant’s decision problem. We utilize a variant of Krusell and Smith’s (1998) method to solve for the general equilibrium of this model.

G.1 Model

The model primitives are almost identical to the basic model. The only difference is the entry process. Instead of free entry with a fixed payment c_q , here we assume that there are a fixed number of potential entrants with an idea every period. Every period, these potential entrants decide whether to enter by paying c_e . All other timings and aspects are the same as in the basic model.

The consumer’s optimization problem is identical to the basic model. Equation (7) char-

acterizes the labor supply.

The plant side of the model becomes entirely different as a consequence of the modification in the entry process. In particular, because the wage is now determined by the labor market equilibrium, in order to predict the future wage each plant owner has to predict where the labor demand curve lies. Because the future labor demand depends on the current productivity distribution of plants, the plant owner has to incorporate a lot more information when making decisions in the current period.

An incumbent plant's value at the beginning of the period is described by the Bellman equation (the firing tax is omitted here)

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

Ω_{t-1} is the information that the plant owner possesses at the end of period $t - 1$. This includes the distribution of plants over productivity and employment in period $t - 1$. The first term on the right hand side is calculated as

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

where

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

The current period profit of a plant for which employment is adjusted is

$$\pi^a(s_t, n_t; z_t, L_t) \equiv \lambda z_t f(n_t, s_t) - w(L_t)n_t.$$

Note that now we denote the wage w_t as $w(L_t)$, using the relationship (7).

If employment is not adjusted, the current period profit is

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

Therefore, the value functions are

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

and

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

Here, $E_z[\cdot | z_t]$ takes the expectation regarding z_{t+1} , conditional on z_t . Ω_t evolves following the law of motion: $\Omega_t = \Gamma(z_t, \Omega_{t-1})$. Employment L_t is the sum of n_t over all plants, and therefore it is a part of Ω_t .

The entrant's value function is

$$V^e(q_t; z_t, \Omega_{t-1}) = \int V^c(s_t, 0; z_t, \Omega_{t-1}) d\eta(s_t | q_t).$$

As in Model 1, only potential entrants with sufficiently high q_t actually enter. There is a threshold value of q_t , q_t^* , which is determined by

$$V^e(q_t^*; z_t, \Omega_{t-1}) = c_e.$$

A potential entrant will enter if and only if $q_t \geq q_t^*$. The total mass of potential entrants is fixed at \bar{N} .

In equilibrium, the combination (w_t, L_t) has to clear the labor market. The labor supply is given by (7), and the labor demand is given by (9) with $N_t = \bar{N}$. The equilibrium (w_t, L_t) is determined by these two equations. Once (w_t, L_t) is given, the values of the other variables can be determined in the same manner as in the basic model. The only difference is that, in (10), $N_t c_q$ does not exist. This only affects the value of consumption.

G.2 Computation

The computation of this model is potentially much more complex than the basic model, since the optimization (potentially) involves many more state variables. To overcome this difficulty, we follow Krusell and Smith (1998) in using limited information instead of the entire state variables to perform optimization, and check whether the “forecast” using this limited information is accurate by simulation. In particular, what is necessary for optimization is to forecast the value of L_t . In the following, we omit the time subscript and represent the $t - 1$ variable by subscript -1 and $t + 1$ variable by a prime ($'$).

1. Set discrete grids on n and s . Set the Markov transition matrices for s and z . Set the distribution of the exit value x .
2. Guess the “prediction rule” for L . After some experimentation, we found that the following formulation works:

$$\log(L) = a_0 + a_1 \log(L_{-1}) + a_2 \log(z) + a_3 \log(z) \times \log(L_{-1}) + a_4 \mathcal{I}(z \neq z_{-1}). \quad (15)$$

$\mathcal{I}(\cdot)$ is an indicator function which takes the value 1 if the statement in the parenthesis is true and 0 if it is false. By adopting this formulation, we are reducing the aggregate state variable from (z, Ω_{-1}) to (z, z_{-1}, L_{-1}) . Let $\Upsilon \equiv (z, z_{-1}, L_{-1})$. Make a guess on the coefficients $(a_0, a_1, a_2, a_3, a_4)$.

3. Optimization loop. Objects: $W(s, n; \Upsilon)$, $V^a(s'; \Upsilon)$, $V^n(s', n; \Upsilon)$, $V^c(s', n; \Upsilon)$, $Z(s, n; \Upsilon)$, $\phi^a(s'; \Upsilon)$, $\phi(s', n; \Upsilon)$, $\zeta(s', n; \Upsilon)$, and $\chi(s, n; \Upsilon)$.
 - (a) Give the initial value for $W(s, n; \Upsilon)$, where s and n are the realizations in the last period. This is the beginning-of-period value for an incumbent.
 - (b) Calculate $V^a(s'; \Upsilon)$ and $V^n(s', n; \Upsilon)$ by

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

and

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

where

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

and

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

Where $w(L)$ function is from (7), and L is calculate by the forecasting rule (15).

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

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

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

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

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

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

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

(e) Calculate $W(s, n; \Upsilon)$ by

$$W(s, n; \Upsilon) = \int \max\langle Z(s, n; \Upsilon), x \rangle d\xi(x).$$

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

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

(f) Update and repeat.

4. Now we can calculate

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

for each q .

We can calculate the cut-off for q , $q^*(\Upsilon)$ by $V^e(q^*(\Upsilon); \Upsilon) = c_e$.

5. Simulation.

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

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

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

The total measure of survivors, with new shocks, is

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

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

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

(d) Thus the total measure is

$$\mu_t(s', n; \Upsilon) = \bar{N}\varrho_t(s'; \Upsilon) + \vartheta_t(s', n; \Upsilon).$$

(e) L is the sum of the incumbent's employment L^i , which can be calculated from $\theta_t(s', n; \Upsilon)$, and the entrant's employment, $\bar{N}L^e$, where L^e can be calculated from $\varrho_t(s'; \Upsilon)$.

The wage can be solved from

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

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

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

(g) Repeat.

6. From simulation, we have the time series data of z_t and L_t . Check if the initial guess on the coefficients of (15) is correct. If not, adjust these coefficients and go back to Step 3.

The resulting forecast rule is

$$\log(L) = -0.31 + 0.39 \log(L_{-1}) - 3.44 \log(z) - 7.19 \log(z) \times \log(L_{-1}) + 0.06\mathcal{I}(z \neq z_{-1}),$$

and

$$R^2 = 0.968.$$

Thus, this forecasting rule is fairly accurate.

Table 25: Results with aggregate shocks

	Good	Bad
Wage	1.004	0.996
q^*	0.4994	0.5004
Entry rate	5.4%	5.4%
Exit rate	5.4%	5.4%
Average size of all plants	85.7	85.2
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

G.3 Results

Here, we use the same steady-state model as in the basic model, and treat the value of N as fixed (\bar{N} in the current model) instead of treating it as endogenously arising from free entry. The only necessary modification is to remove the resource cost c_q . As is mentioned earlier, this affects only the value of consumption in (10).

The results are summarized in Table 25. There are four notable differences from the basic model. First, the wage fluctuates much less than in the benchmark outcome in the main body of the paper (with constant c_e and c_q). This magnitude is in line with the data. Second, there are some fluctuations in q^* . In particular, q^* is larger during the recession—that is, there is more selection of entrants during recessions. Third, the fluctuations in entry rates are very small (5.61% in good times and 5.60% in bad times). Fourth, average plant size is now procyclical.

The wage fluctuates less because the wage determination mechanism is different. Now the labor demand function is downward-sloping, because there is only a limited number of potential entrants that can enter the market. In considering the shift in the labor demand curve, now the labor demand from all plants (not only entrants) matters. As a result, the equilibrium wage does not fluctuate as much as in the basic model. Because the wage does not fluctuate much, there is less offsetting force against the selection of entrants. This generates

a more cyclical q^* in this model. However, the magnitude of the selection is too small to produce the fluctuations in the size and productivity of entrants that we see in the data. In fact, the relative size and productivity of entrants are still very similar across good times and bad times.

The small fluctuations in wage have a direct consequence on the average size of the plants. Now the direct effect of z on the plant size is stronger than the effect of the wage, and the average plant size is procyclical. This, in turn, reduces the quantity of labor that is released for the entrants. As a result, the fluctuations in the entry rate are very small.

In this modified model, the exit rate and the characteristics of exiting plants still remain acyclical. This feature of the model seems to be robust to the change in the entry process.