

ARTICLES

PLANT LIFE CYCLE AND AGGREGATE EMPLOYMENT DYNAMICS

MIN OUYANG

School of Economics and Management, Tsinghua University

Past empirical studies have repeatedly found that plant age matters for aggregate employment dynamics. This paper develops a model of plant life cycle to capture this empirical regularity. In the model, plants differ by vintage and by idiosyncratic productivity. The idiosyncratic productivity is not directly observable, but can be learned over time. This setup gives rise to a learning effect and a creative-destruction effect, under which labor flows from plants with low idiosyncratic productivity to those with high idiosyncratic productivity and from old vintages to new vintages. When calibrated to the U.S. manufacturing job flow series, our model of plant life cycle delivers the observed link between plant age and aggregate employment dynamics.

Keywords: Plant Life Cycle, Aggregate Employment Dynamics, Heterogeneous Businesses

1. INTRODUCTION

With access to longitudinal business data bases becoming more available in the past 20 years, an assortment of studies have emerged exploring the micro foundations of aggregate employment dynamics. It has been documented that economies across time and regions are characterized by large and pervasive job flows. Maybe surprisingly, the majority of this process turns out to occur within industries.¹ In other words, labor is being reallocated across heterogeneous businesses within the same industry, rather than across different industries. Such findings have stimulated theoretical work on models with heterogeneous businesses, as they point out that the traditional representative-agent paradigm does not hold in the real world.²

Although many factors contribute to business heterogeneity, this paper proposes a model that emphasizes business age. We are motivated by an age pattern that stands out in many empirical studies: job-flow magnitude declines with business

The author thanks John Haltiwanger, Michael Pries, John Rust, John Shea, and Deniel Vincent for their support and encouragement. I am especially indebted to John Haltiwanger for his providing the 1972–1998 job flow series by plant age. Address correspondence to: Min Ouyang, Department of Economics, School of Economics and Management, Tsinghua University, Beijing 100084, PR China; email: ouyangm@sem.tsinghua.edu.cn.

age, even controlling for other business characteristics.³ Moreover, business age matters for the cyclical dynamics of various job-flow margins. According to Davis and Haltiwanger (1998), in the U.S. manufacturing sector, job creation and job destruction are approximately equally volatile for young plants whereas, for old plants, job destruction appears more volatile than job creation.

In this paper we present a model to incorporate the age pattern of job flows, emphasizing two dimensions in the productivity dynamics over a plant's life cycle. Intuitively, old plants are more experienced, as they have survived longer; but their technologies are usually outdated, and their products often flag in popularity. In contrast, young plants, although embodying the leading technology, tend to be inexperienced. If more productive plants hire more labor, then the productivity dynamics over a plant's life cycle should generate multiple margins to hire or fire labor, which, when interacted with the business cycle, give rise to an age pattern in aggregate employment dynamics.

To capture this story, we build a model that combines passive learning in the spirit of Jovanovic (1982) with a vintage model by Aghion and Howitt (1992). In our model, plants differ in two productivity components: vintage and idiosyncratic productivity. The idiosyncratic productivity is not directly observable, but can be learned over time. Young plants have newer vintages, but are unsure about their true idiosyncratic productivity. Old plants perceive themselves to have higher idiosyncratic productivity, but have older vintages. Two reallocation effects arise with this setup: a learning effect, under which a plant creates jobs when it learns its true idiosyncratic productivity to be high or destroys jobs when it learns it to be low, and a creative-destruction effect, under which entrants with new technology create jobs to enter and incumbents destroy jobs as they become technologically outdated over time and eventually exit the industry.

This model of plant life cycle can incorporate the observed age pattern of job flows. Because marginal learning diminishes as plants grow old, learning-driven reallocation declines with plant age, so that job flows at young plants are of higher magnitude. Because creative-destruction-driven job destruction is concentrated at old plants, the cyclical dynamics in job destruction dominates those in job creation. We calibrate our model to the job flow series from the U.S. manufacturing sector. Simulation of the calibrated model delivers, at least qualitatively, the observed age pattern in the magnitude and in the cyclical dynamics of job flows.

The age pattern of cyclical job flows has been explored by Campbell and Fisher (2004), who model adjustment cost as proportional to the number of jobs created or destroyed. In their model, a plant currently adjusting employment is more likely to do so again in the immediate future. Because plants have to create jobs to enter, young plants adjust employment more frequently. However, their baseline model fails to deliver the high variance ratio of job destruction to job creation for old plants (p. 349). Our model emphasizes the interaction of the learning effect with the creative destruction effect. The learning effect causes the job-flow magnitude to decline with plant age, and the creative-destruction effect drives job destruction to be more responsive at old plants. Nonetheless, our calibration exercises suggest

TABLE 1. Quarterly gross job flows by plant birth, plant death, and continuing operating plants in the U.S. manufacturing sector: 1972 to 1998

A. Means						
Plant type	$E(Cb)$	$E(Cc)$	$E(C)$	$E(Dd)$	$E(Dc)$	$E(D)$
All	0.39	4.60	4.99	0.61	4.70	5.31
Young	1.65	5.98	7.63	1.27	5.23	6.50
Old	0.12	4.27	4.40	0.45	4.58	5.03
B. Variance ratio of job destruction to job creation						
Plant type	$\sigma(D)^2/\sigma(C)^2$		$\sigma(Dc)^2/\sigma(Cc)^2$			
All	3.27		2.99			
Young	1.01		1.47			
Old	3.75		3.36			

Notes: Young plants are plants younger than 40 quarters. Cb denotes job creation by plant birth, Dd is job destruction by plant death, and Cc and Dc are job creation and destruction by continuing operating plants. C and D are gross job creation and destruction. $C = Cc + Cb$, $D = Dd + Dc$. All numbers are in percentage points.

that, to deliver the observed high variance ratio of job destruction to job creation for old plants, our model requires a creative-destruction effect probably stronger than the existent empirical evidence suggests. We discuss possible ways to improve the model's quantitative performance.

The paper proceeds as follows. Section 2 summarizes the age pattern in employment dynamics with data from the U.S. manufacturing sector. Section 3 presents the model. Section 4 analyzes the dynamics of the learning and creative destruction effects at the steady-state equilibrium. Section 5 calibrates the model to the data, and discusses results from the simulation exercises. We conclude in Section 6.

2. PLANT AGE AND EMPLOYMENT DYNAMICS

Our data source is the gross job flow series for the U.S. manufacturing sector compiled by Davis et al. (1996). Job flows are separated into two components: job creation as the number of jobs created at expanding and newly born plants, and job destruction as the number of jobs lost at declining and closing plants. The age dynamics of job creation and destruction are displayed in Table 1 and Figure 1. The sample covers quarterly job creation and destruction for plants of three different age categories from the second quarter of 1972 to the fourth quarter of 1998. Following Davis et al. (1996), we define young plants as those younger than 10 years; accordingly, old plants are those that have stayed in operation for 10 years or more.

Table 1A shows that the *magnitude* of job flows declines with plant age. From 1972 to 1998, the quarterly job creation rate averages 7.63% and the job destruction




FIGURE 1. Gross job flows at young and old plants, 1972–1998. Dashed lines represent the job creation series; solid lines represent the job destruction series.

rate averages 6.50% for young manufacturing plants; the corresponding statistics are 4.40% and 5.03% for old plants. Table 1A further decomposes job creation into jobs contributed by plant births (Cb) and by continuing operating plants that are expanding (Cc); job destruction is divided into job losses contributed by plant deaths (Dd) and by continuing operating plants that are contracting (Dc). Not surprisingly, Cb is much higher at young plants, as Cb at old plants mainly arises from a few plants that resume operation after temporary closure. Interestingly, Dd , Cc , and Dc are all higher at young plants: for example, young plants' Dd averages 1.27%, about three times of Dd at old plants. In summary, during early years of operation, plants not only face higher chances of closing, as implied by Dd , but also experience much more turnover even if they manage to survive, as suggested by Cc and Dc .

Table 1B shows that plant age matters for the *cyclical dynamics* of job flows. For young plants, job creation and job destruction are equally volatile, as the ratio of the variance of job destruction to that of job creation equals 1.01. For old plants, however, job destruction is much more volatile than job creation, with a variance ratio of 4.18. This age pattern persists even after plant births and deaths are excluded: this ratio for continuing operating plants equals 3.36 for old plants but only 1.47 for young plants. In summary, Table 1B suggests that a plant features similar volatility in job destruction and job creation during its early years

of operation; however, as it grows older, the volatility in job destruction dominates that in job creation.

The link between plant age and employment dynamics summarized in Table 1 has been documented by many other authors. For example, Evans (1987) finds that firm growth decreases with firm age. Dunne et al. (1989) report that employment growth rates and failure rates both decline with plant age. Similar patterns are documented by Aw et al. (2001) for Taiwanese firms. The asymmetry in the volatilities of job destruction versus job creation is discussed in Foote (1998). The influence of plant age on the cyclical dynamics of job flows is also reported by Faberman (2007) with the 1992–2000 data of Business Employment Dynamics.

3. A MODEL OF LEARNING AND CREATIVE DESTRUCTION

To incorporate the age pattern of employment dynamics documented in Section 2, we present a model of plant life cycle. Consider an industry where labor and capital are combined in fixed proportions to produce a single output. Each production unit is called a plant. Each plant consists of

1. a group of machines that embody some vintage;
2. unobservable idiosyncratic productivity;
3. a group of employees.

Let A be the leading technology. Exogenous technological progress drives A to grow over time at rate $\gamma > 0$. When entering the market, a new plant adopts the leading technology at the moment, which becomes its vintage and remains constant throughout its life cycle. $A(a)$ is the leading technology a periods ago, and is also the vintage of a plant of age a :

$$A(a) = A(1 + \gamma)^{-a}. \quad (1)$$

At the time of entry, a plant is endowed with idiosyncratic productivity θ . Many factors can contribute to θ : it can be the talent of the manager, as in Lucas (1978), the location of the store, the organizational structure of the production process, or the plant's fitness to the embodied technology. The key assumption regarding θ is that its value, although fixed at the time of entry, is not directly observable.

Production takes place through a group of workers. Let n be the number of employees. The output of a plant with vintage $A(a)$ and n employees in period t equals

$$A(a)x_t n_t^\alpha, \quad (2)$$

where $0 < \alpha < 1$. The variable $x_t = \theta + \varepsilon_t$: the true idiosyncratic productivity θ is covered by a random noise ε_t . The noise ε_t is an i.i.d. draw from a fixed distribution that masks the influence of θ on output. A plant knows its vintage and can directly observe output. Thus, it can infer the value of x . A plant uses its information on current and past x 's to learn about θ .

3.1. “All-or-Nothing” Learning

Plants are price takers and profit maximizers. They attempt to resolve the uncertainty of θ . The random component ε represents transitory factors that are independent of the idiosyncratic productivity θ . Assuming mean zero for ε , $E(x) = E(\theta) + E(\varepsilon) = E(\theta)$.

Given knowledge of the distribution of ε , a sequence of observations on x allows a plant to learn about its θ . Although a continuum of potential values for θ is more realistic, for simplicity it is assumed that there are only two values: θ_g for a good plant and θ_b for a bad plant. Furthermore, ε is assumed to be distributed uniformly on $[-\omega, \omega]$. Therefore, a good plant will have x each period as a random draw from a uniform distribution over $[\theta_g - \omega, \theta_g + \omega]$, whereas the x of a bad plant is drawn from an uniform distribution over $[\theta_b - \omega, \theta_b + \omega]$. Finally, θ_g, θ_b and ω satisfy $0 < \theta_b - \omega < \theta_g - \omega < \theta_b + \omega < \theta_g + \omega$.

Pries (2004) shows that the preceding assumptions give rise to an “all-or-nothing” learning process. With an observation of x within $(\theta_b + \omega, \theta_g + \omega]$, the plant learns with certainty that it is a good plant; conversely, an observation of x within $[\theta_b - \omega, \theta_g - \omega)$ indicates that it is a bad plant. However, an x within $[\theta_g - \omega, \theta_b + \omega]$ does not reveal anything, because the probability of falling in this range as a good plant and that as a bad plant both equal $\frac{2\omega + \theta_b - \theta_g}{2\omega}$.

This all-or-nothing learning process simplifies the model considerably. Three values of θ^e correspond to three types of plants: those that know they are good with $\theta^e = \theta_g$, those that know they are bad with $\theta^e = \theta_b$, and those that remain unsure about their true idiosyncratic productivity with $\theta^e = \theta_u$ (the prior mean of θ). We call the third type “unsure plants.”

Let φ be the unconditional probability of $\theta = \theta_g$, and p the probability of true idiosyncratic productivity being revealed every period. According to the all-or-nothing learning process, $p = \frac{\theta_g - \theta_b}{2\omega}$. This setup generates the following plant life cycle: a flow of new plants enter the market as unsure; thereafter, every period they stay unsure with probability $1 - p$, learn that they are good with probability $p\varphi$, and learn that they are bad with probability $p(1 - \varphi)$. The evolution of θ^e from the time of entry is a Markov process with values $(\theta_g, \theta_u, \theta_b)$, an initial probability distribution $(0, 1, 0)$, and a transition matrix

$$\begin{bmatrix} 1 & 0 & 0 \\ p\varphi & 1 - p & p(1 - \varphi) \\ 0 & 0 & 1 \end{bmatrix}. \tag{3}$$

If plants were to live forever, eventually all uncertainty would be resolved, as the limiting probability distribution as a goes to ∞ is $(\varphi, 0, (1 - \varphi))$. Because there is a continuum of plants, it is assumed that the law of large numbers applies, so that φ and p are not only the probabilities but also the fractions of good plants and of plants that learn their true idiosyncratic productivity each period. Hence, *ignoring plant exit for now*, the fractions of three plant groups in a cohort of age

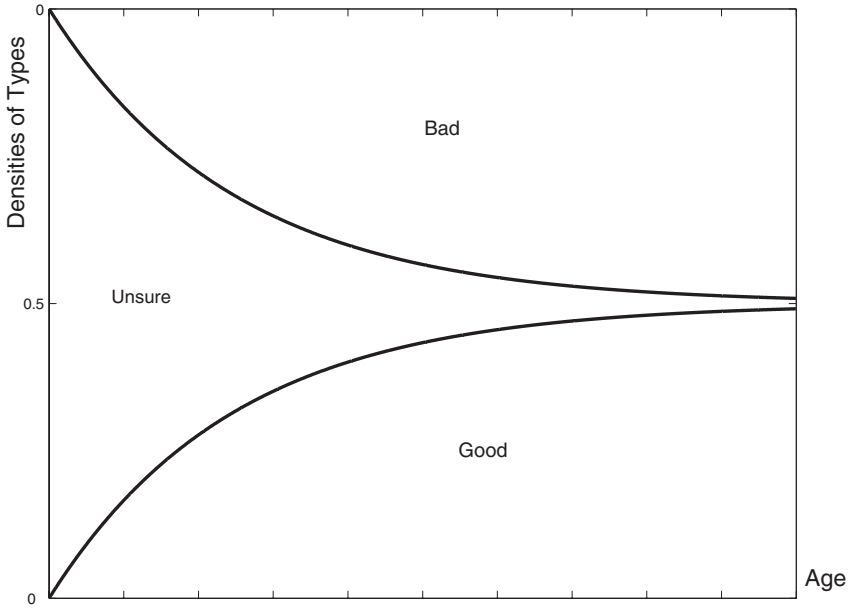


FIGURE 2. Dynamics of a birth cohort by learning: the distance between the concave curve and the bottom axis measures the density of plants with $\theta^e = \theta_g$; the distance between the convex curve and the top axis measures the density of plants with $\theta^e = \theta_b$; and the distance between the two curves measures the density of unsure plants (plants with $\theta^e = \theta_u$).

a are

$$(\varphi[1 - (1 - p)^a], (1 - p)^a, (1 - \varphi)[1 - (1 - p)^a]). \tag{4}$$

Figure 2 presents the evolution of plant distribution across three values of θ^e within a birth cohort. The horizontal axis depicts the cohort age over time. The densities of plants that are certain about their idiosyncratic productivity, whether good or bad, grow as a cohort ages. Moreover, the two “learning curves” (depicting the evolution of the densities of good plants and bad plants) are concave. This feature is the decreasing property of marginal learning with age in Jovanovic (1982), which, in my model, is reflected as the marginal number of learners decreasing with cohort age. The convenient feature of the all-or-nothing learning process is that, on one hand, any single plant learns suddenly, which allows us to keep track of the cross-section distribution easily, whereas, on the other hand, any cohort learns gradually as a group.

3.2. Plant Decisions and Industry Equilibrium

We consider a recursive competitive (partial) equilibrium definition, which includes as a key component the law of motion of the aggregate state of the industry. The aggregate state is (F, D) . F denotes the distribution of plants across vintages

and expected idiosyncratic productivity. D is an exogenous demand parameter; it captures the aggregate condition and is fully observable. The law of motion for D , denoted H_D , is exogenous. The law of motion for F , denoted H_F , is endogenous and defined as $F' = H_F(F, D)$. The element in F that measures the number of plants with belief θ^e and age a is $f(\theta^e, a)$. H_F captures the influence of entry, exit, and learning based on the following sequence of events:

At the beginning of a period, a plant make entry or exit decision after observing the aggregate state. If it decides to enter or stay in operation, a plant adjusts its employment and produces. The equilibrium price is realized. Plants observe their revenue and update their beliefs. Another period begins.

The employment decision. A plant adjusts its employment to solve a static profit maximization problem. The wage rate is normalized as 1; θ^e is a plant's current belief of its true idiosyncratic productivity; P represents the equilibrium output price; Ψ denotes a fixed operation cost each period. A plant's optimal employment is

$$\begin{aligned}
 n(\theta^e, a) &= \arg \max_{n \geq 0} E[PA(a)xn^\alpha - n - \Psi] \\
 &= \left[\alpha \frac{PA\theta^e}{(1 + \gamma)^a} \right]^{\frac{1}{1-\alpha}}, \tag{5}
 \end{aligned}$$

where θ^e can take on the three values θ_g, θ_b , and θ_u . (5) suggests that plant-level employment depends positively on θ^e and P but negatively on a . Thus, a plant hires more employees if the output price is higher, the expected idiosyncratic productivity is higher, or its vintage is newer. Put differently, a plant creates jobs if the output price rises, or if it learns its true idiosyncratic productivity to be good; it destroys jobs as its vintage becomes older.

The corresponding output of a plant with θ^e and a is

$$q(\theta^e, a) = (\alpha P)^{\frac{\alpha}{1-\alpha}} \left[\frac{A\theta^e}{(1 + \gamma)^a} \right]^{\frac{1}{1-\alpha}}, \tag{6}$$

and the corresponding profit is

$$\pi(\theta^e, a) = (\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}) \left[\frac{PA\theta^e}{(1 + \gamma)^a} \right]^{\frac{1}{1-\alpha}} - \Psi. \tag{7}$$

The exit decision. A plant exits if and only if its expected value of staying is below zero. Let $V(\theta^e, a; F, D)$ to be the expected value of staying in operation for one more period and optimizing afterward for a plant with age a and belief θ^e when the aggregate state is (F, D) . V satisfies

$$V(\theta^e, a; F, D) = \pi(\theta^e, a; F, D) + \beta E\{\max[0, V(\theta^e, a + 1; F', D')]\mid\theta^e, F, D\} \tag{8}$$

subject to

$$F' = H_F(F, D), D' = H_D(D) \quad (9)$$

and the all-or-nothing learning that determines the evolution of θ^e .

To further simplify the model, it is assumed that the parameter values are such that $V(\theta_b, a; F, D) < 0$ for any a, F and D , so that the expected value of staying for plants with $\theta^e = \theta_b$ is always negative. Therefore, plants that know their true idiosyncratic productivity to be bad always exit.

According to (4), $\pi(\theta^e, a; F, D)$ decreases in a . It follows that, $V(\theta^e, a; F, D)$ also decreases in a holding other parameters constant. Therefore, there exists a maximum plant age $\bar{a}(\theta^e; F, D)$ for each pair of θ^e and (F, D) , so that plants with θ^e and $a > \bar{a}(\theta^e; F, D)$ choose to exit.

The entry size. New plants are free to enter as long as they bear an entry cost c . c is modeled as a linear function of the entry size. According to the sequence of events, entry and exit update F to F' after plants observe F . Let $f(\theta_u, 0)'$ to be the element of F' that measures the number of new entrants as plants with $\theta^e = \theta_u$ and $a = 0$:

$$c = c_0 + c_1 f(\theta_u, 0)', \quad c > 0, c_1 \geq 0. \quad (10)$$

We let the entry cost to depend positively on the entry size to capture the idea that, for the industry as a whole, *fast* entry is costly and adjustment cannot take place instantaneously. This can arise from a limited amount of land available to build production sites or an upward-sloping supply curve for the industry's specific capital (Goolsbee, 1998). Entrants drive up the price of land or other capital goods, so that entry becomes more costly as more plants enter.

With this setup, new plants keep entering as long as the expected value of entry is above the entry cost. As entry size rises, the entry cost is driven up until it reaches the expected value of entry:

$$V(\theta_u, 0; F, D) = c_0 + c_1 f(\theta_u, 0)'. \quad (11)$$

At this point, entry stops.

The industry equilibrium. Let $Q(F, D)$ represent the equilibrium output. The recursive competitive industry equilibrium constitutes a law of motion H_F , a value function V , and a pricing function P such that

1. V satisfies (8);
2. $F' = H_F(F, D)$ is generated by appropriate summing up of plant entry, exit, and learning. Let $f(\theta^e, a)'$ denote the element of F' that measures the number of plants

with belief θ^e and age a , and $f(\theta^e, a)$ be that of F . $H_F(F, D)$ is such that

$$f(\theta_u, a)' = \left\{ \begin{array}{l} \frac{V(\theta_u, 0; F, D) - c_0}{c_1} \text{ for } a = 0 \\ (1 - p)f(\theta_u, a - 1) \text{ for } 1 \leq a \leq \bar{a}(\theta_u) \\ 0 \text{ for } a > \bar{a}(\theta_u) \end{array} \right\}, \tag{12}$$

$$f(\theta_g, a)' = \left\{ \begin{array}{l} f(\theta_g, a - 1) + p\phi f(\theta_u, a - 1) \text{ for } 1 \leq a \leq \bar{a}(\theta_g) \\ 0 \text{ for } a = 0 \text{ or } a > \bar{a}(\theta_g) \end{array} \right\},$$

where $\bar{a}(\theta_g)$ is the maximum a that satisfies $V(\theta_g, a; F, D) \geq 0$ and $\bar{a}(\theta_u)$ is the maximum a that satisfies $V(\theta_u, a; F, D) \geq 0$.

3. $Q(F, D)$ equals the sum of operating plants' output.⁴
4. $P(F, D)$ satisfies

$$P(F, D) = \frac{D}{Q(F, D)}. \tag{13}$$

The key component of this definition, H_F , is characterized by three essential elements: the entry size, good plants' maximum age, and the unsure plants' maximum age. These three elements, together with all-or-nothing learning, update F to F' ; F' determines the equilibrium output and price by equilibrium conditions 3 and 4 and serves as the aggregate state for the next period.

4. THE STEADY STATE

In the model, new plants embodied with the latest technology keep entering the industry; the employment levels of incumbents grow or shrink, depending on what they learn and how fast the technology updates; those that learn their true idiosyncratic productivity to be bad or whose vintages are too old cease operation. Hence, plant entry, exit, and learning generate a reallocation process in which labor flows from bad plants to good plants, and from old vintages to new vintages. This reallocation process is driven by two forces: learning and creative destruction.

Before allowing for stochastic demand fluctuations, this section examines the link between plant life cycle and aggregate employment dynamics at the steady-state equilibrium with constant demand. The next section will turn to stochastic demand fluctuations and explore whether the intuition from this section carries over.

4.1. The Steady State

We define a steady state as a recursive competitive equilibrium with time-invariant aggregate states. D is and is perceived as time-invariant: $D = H_D(D)$; F is also time-invariant: $F = H_F(F, D)$. Because H_F is generated by entry, exit, and learning, a steady state must feature time-invariant entry and exit for $F = H_F(F, D)$ to hold. Thus, it can be summarized by $\{f^{ss}(0, D), \bar{a}_u^{ss}(D), \bar{a}_g^{ss}(D)\}$: $f^{ss}(0, D)$ is the steady-state entry size when demand equals D , $\bar{a}_g^{ss}(D)$ is the maximum age for good plants, and $\bar{a}_u^{ss}(D)$ is the maximum age for unsure plants. Proposition 1 establishes the existence of a unique steady-state equilibrium.




FIGURE 3. The steady state.

PROPOSITION 1. *With constant D , there exists a unique $\{f^{ss}(0, D), \bar{a}_u^{ss}(D), \bar{a}_g^{ss}(D)\}$ that satisfies conditions 1–4. Moreover, the steady-state PA remains time-invariant: $PA = \bar{PA}^{ss}(D)$.*

See Appendix A.1 for proof. Proposition 1 suggests that, in a steady state, the plant distribution across vintages and idiosyncratic productivity stays time-invariant; so does the product of the price and the leading technology, PA . Because A grows at γ , P must be declining at the same rate for the steady-state PA to be constant. Put intuitively, continuous entry and exit drive the leading technology A as well as the industry output Q to grow at γ . With constant D , growing Q causes P to decline at the same rate.

Figure 3 displays the steady-state evolution of plant distribution within a representative cohort, with the horizontal axis depicting the cohort age across time. Upon entry, the cohort contains unsure plants only. As it ages, bad plants keep learning and exiting. At a certain age, all unsure plants exit with their vintage too old to survive by being unsure; but those plants that have learned their idiosyncratic productivity to be good stay. Afterward, the cohort contains good plants only; its size remains constant. Eventually, all plants exit because their vintage is too old for even good plants to survive.

Figure 3 can be interpreted differently. Note that, in a steady state, all cohorts enter with the same size and experience the same dynamics afterward. Thus, at any time point, different life stages of various birth cohorts overlap, giving rise to the steady-state plant distribution across ages and idiosyncratic productivity, as in

Figure 3 with the horizontal axis as cohort age *across section*. The entering cohort is of size $f^{ss}(0, D)$. Older cohorts are of smaller sizes, because plants that learned their idiosyncratic productivity were bad have exited. No unsure plants are older than $\bar{a}_u^{ss}(D)$. No good plants are older than $\bar{a}_g^{ss}(D)$. Cohorts older than $\bar{a}_u^{ss}(D)$ contain no unsure plants, and thus are of the same size.

4.2. The Plant Life Cycle

To assess how plant life cycle impacts aggregate employment dynamics, we begin with the steady-state plant-level employment n^{ss} as a function of θ^e , a , and D . (5) can be written as

$$n^{ss}(\theta^e, a; D) = \left[\alpha \theta^e \frac{\overline{PA}(D)}{(1 + \gamma)^a} \right]^{\frac{1}{1-\alpha}}, \tag{14}$$

$$\theta^e = \theta_u \text{ or } \theta_g,$$

$$a \leq \bar{a}_u^{ss}(D) \quad \text{if } \theta^e = \theta_u,$$

$$a \leq \bar{a}_g^{ss}(D) \quad \text{if } \theta^e = \theta_g.$$

According to (13), three forces affect the plant-level employment. D affects n^{ss} through its influence on \overline{PA} . θ^e affects n^{ss} positively: a plant creates jobs when learning its idiosyncratic productivity to be good and destroys jobs when learning it to be bad. a affects n^{ss} negatively: a plant destroys jobs as it ages, and exits completely once it reaches $\bar{a}_g^{ss}(D)$ or $\bar{a}_u^{ss}(D)$. We call the impact of θ^e *the learning effect* and that of a *the creative-destruction effect*.

Assuming $\frac{\theta_g}{1+\gamma} > \theta_u$, Figure 4 plots the steady-state plant life cycle driven by the learning and creative-destruction effects. The horizontal axis depicts plant age across time. It presents the employment dynamics of three representative plants, two of which happen to learn their true idiosyncratic productivity at age 50. Plant one stays unsure and exits at 71 as the exit age for unsure plants; plant two learns its true idiosyncratic productivity to be good at age 50, stays in operation afterward, and exits at age 150 as the exit age for good plants; plant three learns its idiosyncratic productivity to be bad at age 50, and exits immediately. Figure 4 shows that, before any plant learns, all three keep shrinking because of the creative-destruction effect. At age 50, plant two creates but plant three destroys jobs by exiting. After age 50, plant two again destroys jobs over time, and it exits completely at age 150.

4.3. The Cohort Life Cycle

The cohort life cycle is given by aggregating the plant life cycle to the cohort level using the evolution of within-cohort plant distribution. Propositions 2 and 3 capture the dynamics of the learning and creative-destruction effects.




FIGURE 4. Three representative plants in a steady state. Solid line plots the employment dynamics of an unsure plant that never learns and exits after 71 periods; dotted line plots the employment dynamics of a plant that learns it is a bad plant after 50 periods and exits immediately; dashed line plots the employment dynamics of a plant that learns it is a good plant after 50 periods and exits after 150 periods.

PROPOSITION 2. *The steady-state job creation and destruction of a cohort of age a driven by the learning effect, $jc_1^{ss}(a)$ and $jd_1^{ss}(a)$, both decrease weakly with cohort age.*

The magnitudes of $jc_1^{ss}(a)$ and $jd_1^{ss}(a)$ are determined by the fraction of plants that learn and how many of the learning plants are good or bad. Combining (4) with (13) gives

$$\begin{aligned}
 jc_1^{ss}(a) &= \frac{\varphi p (1-p)^a \left[\left(\frac{\theta_g}{1+\gamma} \right)^{\frac{1}{1-\alpha}} - \theta_u^{\frac{1}{1-\alpha}} \right]}{\theta_g^{\frac{1}{1-\alpha}} \varphi + (1-p)^a \left(\theta_u^{\frac{1}{1-\alpha}} - \theta_g^{\frac{1}{1-\alpha}} \varphi \right)} && \text{for } 0 < a \leq \bar{a}_u^{ss}(D); \\
 jd_1^{ss}(a) &= \frac{\theta_u^{\frac{1}{1-\alpha}} p (1-p)^a (1-\varphi)}{\theta_g^{\frac{1}{1-\alpha}} \varphi + (1-p)^a \left(\theta_u^{\frac{1}{1-\alpha}} - \theta_g^{\frac{1}{1-\alpha}} \varphi \right)} && \text{for } 0 < a \leq \bar{a}_u^{ss}(D); \quad (15) \\
 jc_1^{ss}(a) &= jd_1^{ss}(a) = 0 && \text{otherwise.}
 \end{aligned}$$

Apparently, $jc_1^{ss}(a)$ and $jd_1^{ss}(a)$ both decrease weakly in a . This is because of the decreasing property of marginal learning captured at the cohort level by the all-or-nothing learning process. For $a > \bar{a}_u^{ss}(D)$, learning has stopped, so there is no learning-driven job creation or destruction. For $0 < a \leq \bar{a}_u^{ss}(D)$, the fraction

of learning plants per period equals $\frac{p(1-p)^a}{(1-\varphi)(1-p)^a+\varphi}$, out of which proportion φ create jobs and $1 - \varphi$ destroy jobs. Because $\frac{p(1-p)^a}{(1-\varphi)(1-p)^a+\varphi}$ decreases in a , the learning-driven job creation and destruction also decrease in cohort age.

PROPOSITION 3. *Under the creative-destruction effect, the steady-state job destruction by aging plants $jd_{cd}^{ss}(a)$, as plants who do not learn or have learned already, increases weakly with cohort age; moreover, $jd_{cd}^{ss}(a)$ rises in γ .*

According to Proposition 1, the steady-state PA remains time-invariant. Therefore, (13) suggests that n^{ss} declines as a grows as long as θ^e does not change. Put intuitively, plants that have learned their idiosyncratic productivity to be good and those that remain unsure destroy jobs as they grow older; moreover, the higher γ is, the more jobs they destroy while they age. Thus, $jd_{cd}^{ss}(a)$ is determined by the within-cohort fraction of aging plants and by how fast the technology updates. Combining (4) and (14) gives

$$\begin{aligned}
 jd_{cd}^{ss}(a) &= \left[1 - (1 + \gamma)^{-\frac{1}{1-\alpha}} \right] \\
 &\times \left[\frac{\theta_g^{\frac{1}{1-\alpha}} \varphi [1 - (1 - p)^a] + \theta_u^{\frac{1}{1-\alpha}} (1 - p)^{a+1}}{\theta_g^{\frac{1}{1-\alpha}} \varphi + (1 - p)^a \left(\theta_u^{\frac{1}{1-\alpha}} - \theta_g^{\frac{1}{1-\alpha}} \varphi \right)} \right] \quad \text{for } 0 \leq a < \bar{a}_u^{ss}(D); \\
 & \hspace{15em} (16) \\
 jd_{cd}^{ss}(a) &= 1 - (1 + \gamma)^{-\frac{1}{1-\alpha}} \quad \text{for } \bar{a}_u^{ss}(D) < a < \bar{a}_g^{ss}(D).
 \end{aligned}$$

Apparently, $jd_{cd}^{ss}(a)$ increases weakly in a and is bigger with higher γ . As a cohort ages, the fraction of learning plants declines and that of aging plants rises. Hence, creative-destruction-driven job destruction by aging plants rises with cohort age.

Two remarks should be made. First, in addition to job destruction by aging plants, the creative-destruction effect also drives job creation by entry as well as job destruction by exit at $\bar{a}_u^{ss}(D)$ and $\bar{a}_g^{ss}(D)$. Because job creation by entry occurs before the plant life cycle starts, it does not contradict the intuition that the creative-destruction effect strengthens with plant age. But job destruction by exit at $\bar{a}_u^{ss}(D)$ and $\bar{a}_g^{ss}(D)$ adds to this intuition, as those exits occur only at older ages. Put intuitively, the creative-destruction effect drives more and more plants to shrink as a cohort grow older, and eventually causes the entire cohort to exit.

Second, the learning effect and the creative-destruction effect interact. It is the decreasing property of the marginal learning effect that causes the creative-destruction effect to strengthen. Moreover, according to (15), the magnitude of the learning-driven job creation equals $\left[\left(\frac{\theta_g}{1+\gamma} \right)^{\frac{1}{1-\alpha}} - \theta_u^{\frac{1}{1-\alpha}} \right]$. Intuitively, a plant’s vintage grows one period older when it learns; this is captured by $\frac{\theta_g}{1+\gamma}$. Therefore, the creative-destruction effect dampens the learning-driven job creation.

Figure 5 displays the life cycle of a representative cohort. It presents the employment dynamics of an “average” plant, as the sum of a good plant’s




FIGURE 5. Employment of an “average” plant: the total employment of a cohort divided by the number of plants. In this simulation case, the exit age of unsure plants is 70. See the caption to Figure 4 for details.

employment weighted by the within-cohort fraction of good plants and an unsure plant’s employment weighted by the within-cohort fraction of unsure plants. When the “average” plant is very young, it keeps creating jobs because the learning effect dominates the creative-destruction effect. As it grows older, the learning effect weakens and the creative-destruction effect strengthens. Once the creative-destruction effect dominates, its employment begins to fall. Its employment jumps at $\bar{a}_u^{ss}(D)$ because of the exit of all unsure plants. After $\bar{a}_u^{ss}(D)$, it keeps destroying jobs because of the creative destruction effect, and it exits at $\bar{a}_g^{ss}(D)$.

4.4. Aggregate Employment Dynamics

The aggregate employment dynamics is given by aggregating the cohort life cycle to the industry level. It reflects the number of plants choosing to adjust employment and the magnitude of their adjustment. Demand has significant influences on both dimensions.

First, demand influences the magnitude of employment adjustment. According to (13), D affects n^{ss} through $\bar{P}A$. $\bar{P}A$ is canceled out in (15) and (16) by staying constant at the steady state. However, with stochastic demand variations, $\bar{P}A$ would fluctuate over time. In that case, an increase in $\bar{P}A$ amplifies job creation but dampens job destruction; an increase in $\bar{P}A$ does the opposite. In extreme cases, an increase in $\bar{P}A$ may dominate the creative-destruction effect, so that the aging plants create jobs even if their vintages grow older; or a decrease in $\bar{P}A$ may dominate the learning effect, so that the learning plants destroy jobs even if they learn their idiosyncratic productivity to be good.

Second and more importantly, demand influences the number of plants that choose to adjust employment. From a pure accounting point of view, the industry

features four job-flow margins: the entry margin, the learning margin, the aging margin, and the exit margin at $\bar{a}_u^{ss}(D)$ and $\bar{a}_g^{ss}(D)$. Demand affects the number of plants at all these margins through its influence on $f^{ss}(0, D)$, $\bar{a}_u^{ss}(D)$, and $\bar{a}_g^{ss}(D)$:

PROPOSITION 4. *At a steady-state equilibrium, $f^{ss}(0, D)$, $\bar{a}_u^{ss}(D)$, and $\bar{a}_g^{ss}(D)$ all weakly increase in D .*

See Appendix A.2 for proof. Proposition 4 suggests that, if we compare two steady-state equilibria, one with a high demand and the other with a low demand, the low-demand steady state features less entry and younger exit ages. Intuitively, lower demand lowers the number of plants at the entry margin but raises that at the exit margin. Moreover, younger exit ages imply a shorter plant life cycle, less time for unsure plants to learn, and thus a smaller number of plants at the learning margin.

If this intuition carried over when demand fluctuates stochastically, then a drop in demand would cause less entry, suggesting a drop in job creation by young plants, and more exit, implying a rise in job destruction by old plants. Moreover, if demand varied stochastically, birth cohorts would start with different sizes, so that, at any time point, the size of a cohort of age a would be affected by the demand level a periods ago. Therefore, past demand variations determine current plant distribution and, consequently, the number of plants at the learning and aging margins. The next section explores quantitatively our model's implications for the aggregate employment dynamics with stochastic demand.

5. QUANTITATIVE IMPLICATIONS

This section applies numerical techniques to analyze a stochastic version of our model, in which demand follows a two-state Markov process with values $[D_h, D_l]$ and transition probability μ . Throughout this section, plants expect the current demand to persist for the next period with probability μ , and to change with probability $1 - \mu$.

We first calibrate our model to data presented in Section 2. Then we approximate value functions and the laws of motion using the approach by Krusell and Smith (1998). With the approximated laws of motion, we simulate the calibrated model with a random initial plant distribution to examine whether our model can deliver the observed properties of aggregate employment dynamics.

5.1. Baseline Calibration

Table 2 presents the parameter values in our baseline calibration. Some of the parameter values are pre-chosen. We allow a period to represent one quarter and set the quarterly discount factor $\beta = 0.99$. μ is chosen to equal 0.95, so that demand switches between a high level and a low level with a constant probability 0.05 per quarter. Ψ is set to 1 as it matters only as a scale for D . We set α , the elasticity of production with respect to labor input, to 0.66 as the share of return to labor.

TABLE 2. Baseline calibration

Calibrated parameters	Value
Productivity of bad plants: θ_b	1
Productivity of good plants: θ_g	2.4
Quarterly technological pace: γ	0.003
Quarterly discount factor: β	0.99
The elasticity of production w.r.t. labor: α	0.66
Entry cost parameter: c_0	0.158
Entry cost parameter: c_1	0.728
Persistence rate of demand: μ	0.95
Prior probability of being a good plant: φ	0.017
Quarterly pace of learning: p	0.083
Operation cost per period: Ψ	1
High demand: D_h	925.98
Low demand: D_l	873.90

Some other parameter values are chosen according to the empirical evidence. The elasticity of entry cost with respect to entry size, c_1 , is chosen based on Goolsbee (1998), who estimates that a 10% increase in demand for equipment raises equipment price by 7.284% (p. 143, Table VII). Accordingly, $c_1 = 0.7284$. The relative productivity of good and bad plants is chosen following Davis and Haltiwanger (1998), who assume a ratio of high to low productivity of 2.4, based on between-plant productivity differentials reported by Bartelsman and Doms (2000). We normalize the productivity of bad plants as 1 and set the productivity of good plants to 2.4. The technological pace, γ , determines the strength of the creative-destruction effect. The technology pace, γ , is set according to Basu et al. (2001), who estimate a quarterly technological pace of 0.0037 for durable manufacturing and a pace of 0.0027 for nondurable manufacturing. We set $\gamma = 0.003$.

The strength of the learning effect is determined by the prior probability of being a good plant (φ) and the quarterly pace of learning (p). Because the learning effect is strong at young plants, φ and p are calibrated to the observed job creation and destruction rates for young plants. Table 1 shows that, in the U.S. manufacturing sector, young plants' mean job creation rate equals 7.63% and their mean job destruction rate equals 6.50%. Let jc_y denote young plants' job creation rate, jd_y their job destruction rate. Young plants are defined as those younger than 40 quarters as in Table 1. This sets the following two restrictions on φ and p : $jc_y = 7.52\%$ and $jd_y = 6.56\%$. With other parameter values given by Table 2, jc_y and jd_y are functions of φ and p only as $\bar{P}A(D)$ cancels out. Using a search algorithm, we find $p = 0.083$ and $\varphi = 0.017$.

The remaining parameters are high demand D_h , low demand D_l , and the fixed component of the entry cost c_0 . Proposition 4 suggests that demand variations cause exit ages to vary, reflected in the data as fluctuations in job destruction. Thus, the values of D_h , D_l , and c_0 are chosen to match the observed moments

of job destruction. From 1972 to 1998, the U.S. manufacturing job destruction rate fluctuates between 2.96% and 11.43% with a mean of 5.31%. This puts the following restrictions on our calibrated model.

First, its implied long-run job destruction rate must be around 5.31%. We let \bar{a}_g and \bar{a}_u represent the maximum ages of good plants and unsure plants in the high-demand steady state and \bar{a}_g' and \bar{a}_u' represent the exit ages in the low-demand steady state. The steady-state job destruction rate implied by either pair, has to be around 5.31%.

Second, we match the peak in job destruction that occurs at the onset of a recession. Our model suggests that the jump in the job destruction rate at the beginning of a recession comes from the shift of exit margins to younger ages. We assume that when demand drops, the exit margins shift from \bar{a}_g and \bar{a}_u to \bar{a}_g' and \bar{a}_u' immediately, and the job destruction rate at this moment must not exceed 11.43%.⁵

Third, we match the trough in job destruction that occurs at the onset of a boom. Our model suggests that when demand goes up, the exit margins extend to older ages, so that for several subsequent periods job destruction comes only from the learning margin, implying a trough in the job destruction rate. To match the data, the job destruction rate at this moment has to be around 3%.

Additionally, (\bar{a}_g, \bar{a}_u) and (\bar{a}_g', \bar{a}_u') must satisfy the steady state conditions on the gap between the exit ages of good and unsure plants. Using a search algorithm, we find that these conditions are satisfied for the following combination of parameter values: $\bar{a}_g = 78$, $\bar{a}_u = 64$, $\bar{a}_g' = 73$, $\bar{a}_u' = 59$, and $c_0 = 0.158$. By applying these values of \bar{a}_g , \bar{a}_u , \bar{a}_g' and \bar{a}_u' to the steady state industry structure, we find $D_h = 925.98$ and $D_l = 873.90$.

5.2. Simulated Employment Dynamics

With all the parameter values assigned, we approximate the value functions. Our key computational task is to map F , the plant distribution across ages and idiosyncratic productivity, given demand D , into a set of value functions $V(\theta^e, a; F, D)$. Unfortunately, the endogenous state variable F is a high-dimensional object. It is known that the numerical solution of dynamic programming problems becomes increasingly difficult as the size of the state space increases. Our computational strategy follows Krusell and Smith (1998) by shrinking the state space into a limited set of variables and showing that these variables' laws of motion can approximate the equilibrium behavior of plants in the simulated time series. Details are presented in Appendix A.3. Our simulation exercises suggest that the dynamic system is globally stable: the industry structure eventually settles down with constant entry and exit along any sample path with constant demand. With the corresponding decision rules and an initial plant distribution, we then investigate the aggregate time series properties of employment fluctuations in the calibrated model.

We start with a random plant distribution, simulate the model for 5,000 periods according to the approximated value functions, and discard the first 500 periods to investigate the properties of the stationary region of the simulated time

TABLE 3. Simulation statistics under the baseline calibration

A. Means						
Plant type	$E(Cb)$	$E(Cc)$	$E(C)$	$E(Dd)$	$E(Dc)$	$E(D)$
All	3.07	1.59	4.66	2.40	2.83	5.22
Young	5.06	2.26	7.32	3.00	2.21	5.21
Old	0	0.55	0.55	1.45	3.25	4.65

B. Variance ratio of job destruction to job creation		
Plant type	$\sigma(D)^2/\sigma(C)^2$	$\sigma(Dc)^2/\sigma(Cc)^2$
All	1.21	1.28
Young	1.20	1.24
Old	1.85	1.58

Note: See Table 1 for notations.

series. Table 3 reports the statistics of the simulated job flows under the baseline calibration.

Table 3 suggests that the calibrated model does deliver, at least qualitatively, the relationship between plant age and employment dynamics: the means of job flow rates decline with plant age, and the relative volatilities of job destruction to creation rises with age. These patterns remain evident after plant birth and plant death are excluded.

However, two remarks should be made. Table 3 shows that the simulated magnitudes of old plants' job flows are significantly smaller than those in the data. For example, the mean job creation for old plants is 4.40% in the data, but it is only 0.55% in our simulation. Two factors contribute to this difference. First, the simulated Cb , job creation by plant birth, is zero for old plants, because newborn plants cannot be old by definition. However, in the data, old plants display a job creation rate of 0.12% by plant birth. This is because, in reality, some old plants cease operation temporarily and reopen later. Such temporary exit behavior is not captured by our model. Second, the all-or-nothing learning process, although greatly simplifying our analysis, imposes very little learning on old plants, as many old cohorts contain no unsure plants. Because of such limited learning, old aging plants experience much less job creation and destruction in our calibrated model.

Moreover, the difference between the variance ratio of job destruction over job creation for old plants and that for young plants is rather small under our baseline calibration. For example, this variance ratio rises from 1.01 for young plants to 3.75 for old plants in the data, but from 1.20 only to 1.85 under our baseline calibration.

We further explore our model's quantitative performance in matching the relative volatility of job destruction and creation. In the baseline calibration, we set the quarterly technological pace γ to 0.003 according to Basu et al. (2001). An alternative estimate of γ is provided by Caballero and Hammour (1994), who set γ to 0.007 by attributing all manufacturing employment growth to technological

TABLE 4. Simulation statistics under the baseline calibration and additional calibration

Plant type	Mean(C)	Mean(D)	Emp. share	$\frac{\text{Var}(D)}{\text{Var}(C)}$
A. Additional calibration $p = 0.075, \varphi = 0.030, \gamma = 0.007$				
All	4.66	5.22	100	1.66
Young	7.31	6.20	75.4	1.51
Old	0.55	3.70	24.6	3.57
B. Baseline calibration $p = 0.083, \varphi = 0.017, \gamma = 0.003$				
All	4.33	4.27	100	1.22
Young	7.66	6.36	65.8	1.20
Old	0.62	2.22	34.2	1.85

progress. To explore how the value of γ impacts the simulation results, we re-calibrate our model assuming $\gamma = 0.007$, approximate the corresponding laws of motion, and simulate the re-calibrated model. The new simulation results are summarized in Table 4. Key results from the baseline calibration are included for comparison.

Table 4 shows that, with $\gamma = 0.007$, our model needs a slower learning pace ($p = 0.075$) and a higher prior probability of being good ($\varphi = 0.030$) to match the data. Moreover, $\gamma = 0.007$ magnifies the difference in the destruction-to-creation variance ratio between young and old plants: it equals 1.51 for young plants and 3.57 for old plants, compared to 1.20 and 1.85 in the baseline calibration.

This result can be explained as follows. A lower p implies that plants learn at a slower pace and that the learning effect is weak. A higher γ suggests that technology updates at a faster pace and that the creative destruction effect is strong. The change in the relative strength of the two effects causes the creative-destruction-driven job destruction to respond more but the learning-driven job creation to respond less to demand variations, so that the destruction-to-creation variance ratio rises for both young and old plants in Table 4. Moreover, $\gamma = 0.007$ shortens plant life cycle, so that fewer plants become old. This is shown, in Table 4, as the drop in old plants' employment share with $\gamma = 0.007$. With higher γ but lower p and, additionally, fewer old plants, Table 4 implies that the volatility of job destruction relative to job creation not only rises for old plants, but also rises by more than it does for young plants.

However, 0.007 likely overestimates the technological pace in reality, according to Basu et al. (2001). Therefore, the results in Table 4 imply that our model requires a creative-destruction effect stronger than what the data suggest to deliver the high variance ratio of job destruction to creation for old plants.

There are possible ways to improve our model's quantitative performance in this aspect. For example, a more general learning process can be introduced, assuming ε , the noise that covers θ , to be distributed normally instead of uniformly. In that case, any realized ε value would change the perceived probability of being good, so that θ^e can take on any values between θ_b and θ_g . This would generate more than two exit margins for old plants, and job destruction by old plants may respond more to demand variations, even with a quarterly technological pace slower than 0.007.

6. CONCLUSION

We propose a model of plant life cycle, inspired by the observed age pattern in aggregate employment dynamics. Our model combines the learning effect with the creative-destruction effect. As plants grow older, job flows under the learning effect diminish, but those driven by the creative-destruction effect grow. When calibrated to the U.S. manufacturing data, our model generates, at least qualitatively, the observed age pattern in aggregate employment dynamics.

This model of plant life cycle provides a useful framework for future research in many directions. For example, Ouyang (2009) applies this framework to examine the impact of cyclical job flows on aggregate productivity and argues that recessions bring a "scarring" effect by killing off unsure firms that are truly good. Foote (1998) examines the variance ratio of job destruction to job creation across different industries and finds that this ratio declines in the trend employment growth of the industry. Can the relative strength of learning and creative destruction modeled in this paper explain the cross-industry differences in trend employment growth documented by Foote (1998)? We leave the question for future research.

NOTES

1. According to Davis and Haltiwanger (1998), about 87% of the excess reallocation takes place within each of the four-digit Standard Industry Classification (SIC) industries in the U.S. manufacturing sector. The dominance of within-industry reallocation persists under even narrower industry definitions such as four-digit industries by specific regions.

2. See, for example, Caballero and Hammour (1994), Mortensen and Pissaridies (1994), Gomes et al. (2001), and Barlevy (2002).

3. This has been documented, for example, by Evans (1987), Dunne et al. (1989), and Aw et al. (2001). Davis et al. (1996) summarize the age pattern of job flows as surprisingly robust regardless of major differences among studies in measurement, country, sectoral coverage, and data.

4. Although industry-level output should equal the sum of realized plant-level output, it can be shown that the expectation error and the random noise cancel out within each age cohort so that the sum of expected plant output equals the sum of realized output.

5. The simulation exercises suggest that, when a negative demand shock strikes, the exit margins shift past \bar{a}_g' and \bar{a}_u' . The bigger shift implies a bigger jump in job destruction. This is why we require the maximum job destruction to be below 11.43%.

REFERENCES

- Aghion, P. and P. Howitt (1992) A model of growth through creative destruction. *Econometrica* 60(2), 323–351.
- Aw, B. Y., X. Chen, and M. J. Roberts (2001) Firm-level evidence on productivity differentials, turnover, and exports in Taiwanese manufacturing. *Journal of Development Economics* 66(1), 51–86.
- Barlevy, G (2002) The sully effect of recessions. *Review of Economic Studies* 69(1), 65–96.
- Bartelsman, E. and M. Doms (2000) Understanding productivity: Lessons from longitudinal microdata. *Journal of Economic Literature* 38(3), 569–594.
- Basu, S., J. Fernald, and M. Shapiro (2001) Productivity Growth in the 1990s: Technology, Utilization, or Adjustment? NBER working paper 8359.
- Caballero, R. and M. Hammour (1994) The cleansing effect of recessions. *American Economic Review* 84(5), 1350–1368.
- Campbell, J. and J. Fisher (2004) Idiosyncratic risk and aggregate employment dynamics. *Review of Economic Dynamics* 7(2), 331–353.
- Davis, S. and J. Haltiwanger (1998) Gross job flows. In O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, pp. 2711–2805. Amsterdam: North-Holland.
- Davis, S., J. Haltiwanger and S. Schuh (1996) *Job Creation and Job Destruction*. Cambridge, MA: MIT Press.
- Den Haan, W. and A. Marcat (1994) Accuracy in simulations. *Review of Economic Studies* 61, 3–18.
- Dunne, T., M. Roberts and L. Samuelson (1989) The growth and failure of U.S. manufacturing plants. *Quarterly Journal of Economics* 104(4), 671–698.
- Evans, D (1987) The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries. *Journal of Industry Economics* 35(4), 567–581.
- Faberman, J. (2007) The relationship between the establishment age distribution and urban growth. Working Paper, Federal Reserve Bank of Philadelphia.
- Foote, C. (1998) Trend employment growth and the bunching of job creation and destruction. *Quarterly Journal of Economics* 113(3), 809–834.
- Gomes, J., J. Greenwood and S. Rebelo (2001) Equilibrium unemployment. *Journal of Monetary Economics* 48(1), 109–152.
- Goolsbee, A. (1998) Investment tax incentives, prices, and the supply of capital goods. *Quarterly Journal of Economics* 113(1), 121–148.
- Jovanovic, B. (1982) Selection and the evolution of industry. *Econometrica* 50(3), 649–670.
- Krusell, P. and A. Smith (1998) Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106(5), 867–895.
- Lucas, R. On the size distribution of business firms. *Bell Journal of Economics* 9(2), 508–523.
- Mortensen, D. and C. Pissarides (1994) Job creation and job destruction in the theory of unemployment. *Review of Economic Studies* 61(3), 397–415.
- Ouyang, M. (2009) The scarring effect of recessions. *Journal of Monetary Economics* 56, 184–199.
- Pries, M. (2004) Persistence of employment fluctuations: A model of recurring job loss. *Review of Economic Studies* 71(1), 193–215.

APPENDIX

A.1. PROOF OF PROPOSITION 1

We first prove that PA stays time-invariant at a steady state. According to steady-state demand condition,

$$D = (PA)^{\frac{1}{1-\alpha}} \left\{ \begin{array}{l} \sum_{a=0}^{\bar{a}_u^{ss}(D)} f^{ss}(\theta_u, a; D) \alpha^{\frac{\alpha}{1-\alpha}} \left[\frac{\theta_u}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} + \\ \sum_{a=0}^{\bar{a}_g^{ss}(D)} f^{ss}(\theta_g, a; D) \alpha^{\frac{\alpha}{1-\alpha}} \left[\frac{\theta_g}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} \end{array} \right\}, \quad (\text{A.1})$$

where $f^{ss}(\theta^e, a; D)$ is the steady-state measure of plants with age a and the expected idiosyncratic productivity θ^e . More specifically,

$$f^{ss}(\theta_u, a; D) = f^{ss}(0, D) (1-p)^a, \quad (\text{A.2})$$

$$f^{ss}(\theta_g, a; D) = f^{ss}(0, D) \varphi [1 - (1-p)^a]. \quad (\text{A.3})$$

By definition, a steady state features time-invariant distribution of plants across a and θ^e . This implies that PA has to be time-invariant for (A.1) to hold.

Now we prove that there exists a unique steady state equilibrium. The following three conditions, together with (A.1) characterize a steady-state equilibrium. The exit condition for a good plant is

$$\left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right) \left[\frac{PA\theta_g}{(1+\gamma)\bar{a}_g^{ss}(D)} \right]^{\frac{1}{1-\alpha}} - \Psi = 0. \quad (\text{A.4})$$

The exit condition for an unsure plant is

$$\begin{aligned} & \left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right) \left[\frac{PA\theta_u}{(1+\gamma)\bar{a}_u^{ss}(D)} \right]^{\frac{1}{1-\alpha}} - \Psi \\ & + p\varphi \sum_{a=\bar{a}_u^{ss}(D)+1}^{\bar{a}_g^{ss}(D)} \beta^a - \bar{a}_u^{ss}(D) \left\{ \left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right) \left[\frac{PA\theta_g}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - \Psi \right\} = 0. \end{aligned} \quad (\text{A.5})$$

The free entry condition is

$$\begin{aligned} c_0 + c_1 f^{ss}(0, D) &= \sum_{a=0}^{\bar{a}_u^{ss}(D)} \beta^a (1-p)^a \left\{ \left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right) \left[\frac{PA\theta_u}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - \Psi \right\} \\ &+ \sum_{a=0}^{\bar{a}_g^{ss}(D)} \beta^a \varphi [1 - (1-p)^a] \left\{ \left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right) \left[\frac{PA\theta_g}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - \Psi \right\}. \end{aligned} \quad (\text{A.6})$$

Furthermore, (A.4) suggests

$$(PA)^{\frac{1}{1-\alpha}} = \frac{\Psi}{\left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}} \right)} \left[\frac{(1+\gamma)\bar{a}_g^{ss}(D)}{\theta_g} \right]^{\frac{1}{1-\alpha}}. \quad (\text{A.7})$$

Plugging (A.7), (A.2), and (A.3) into (A.1) gives

$$D = \frac{\Psi}{\left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}\right)} \left[\frac{(1+\gamma)^{\bar{a}_g^{ss}(D)}}{\theta_g} \right]^{\frac{1}{1-\alpha}} f^{ss}(0, D) \\ \times \left\{ \begin{aligned} &\sum_{a=0}^{\bar{a}_u^{ss}(D)} (1-p)^a \alpha^{\frac{\alpha}{1-\alpha}} \left[\frac{\theta_u}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} + \\ &\sum_{a=0}^{\bar{a}_g^{ss}(D)} \varphi [1 - (1-p)^a] \alpha^{\frac{\alpha}{1-\alpha}} \left[\frac{\theta_g}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} \end{aligned} \right\}. \tag{A.8}$$

Plugging (A.7) into (A.5) gives

$$\frac{1 - \beta \bar{a}_s^{ss}(D) - \bar{a}_u^{ss}(D) + 1}{1 - \beta} = \left(\frac{\theta_u}{\theta_g} \right)^{\frac{1}{1-\alpha}} (1+\gamma)^{\frac{\bar{a}_g^{ss}(D) - \bar{a}_u^{ss}(D)}{1-\alpha}} \\ + p\varphi \frac{\frac{\beta}{(1+\gamma)^{\frac{1}{1-\alpha}}}}{1 - \frac{\beta}{(1+\gamma)^{\frac{1}{1-\alpha}}}} \left[(1+\gamma)^{\frac{\bar{a}_g^{ss}(D) - \bar{a}_u^{ss}(D)}{1-\alpha}} - \beta \bar{a}_s^{ss}(D) - \bar{a}_u^{ss}(D) \right]. \tag{A.9}$$

Notice that all the conditions become recursive. D does not enter (A.9), so that (A.9) determines a unique value for $\bar{a}_g^{ss}(D) - \bar{a}_u^{ss}(D)$. Then, (A.6) and (A.8) jointly determine $\bar{a}_g^{ss}(D)$ and $f^{ss}(0, D)$, after replacing PA in (A.6) as a function of $\bar{a}_g^{ss}(D)$ and replacing $\bar{a}_u^{ss}(D)$ in (A.8) with $\bar{a}_g^{ss}(D) - (\bar{a}_g^{ss}(D) - \bar{a}_u^{ss}(D))$. It can be proven that, for (A.9) to reveal a unique solution for $\bar{a}_g^{ss}(D) - \bar{a}_u^{ss}(D)$, it requires that $\theta_u < \theta_g$, which holds by definition. This proves Proposition 1.

A.2. PROOF OF PROPOSITION 4

Plugging (A.7) into (A.6) gives

$$f^{ss}(0, D) \\ = \left(\begin{aligned} &\sum_{a=0}^{\bar{a}_u^{ss}(D)} \beta^a (1-p)^a \Psi \left\{ \left[\frac{(1+\gamma)^{\bar{a}_g^{ss}(D)}}{\theta_g} \right]^{\frac{1}{1-\alpha}} \left[\frac{\theta_u}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - 1 \right\} + \\ &\sum_{a=0}^{\bar{a}_g^{ss}(D)} \beta^a \varphi [1 - (1-p)^a] \Psi \left\{ \left[\frac{(1+\gamma)^{\bar{a}_g^{ss}(D)}}{\theta_g} \right]^{\frac{1}{1-\alpha}} \left[\frac{\theta_g}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - 1 \right\} - c_0 \end{aligned} \right) / c_1. \tag{A.10}$$

Combining (A.10) and (A.8) gives

$$D = \frac{\Psi}{\left(\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}\right)} \left[\frac{(1+\gamma)^{\bar{a}_g^{ss}(D)}}{\theta_g} \right]^{\frac{1}{1-\alpha}} \left\{ \begin{aligned} &\sum_{a=0}^{\bar{a}_u^{ss}(D)} (1-p)^a \alpha^{\frac{\alpha}{1-\alpha}} \left[\frac{\theta_u}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} + \\ &\sum_{a=0}^{\bar{a}_g^{ss}(D)} \varphi [1 - (1-p)^a] \alpha^{\frac{\alpha}{1-\alpha}} \left[\frac{\theta_g}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} \end{aligned} \right\},$$

$$\left(\begin{array}{l} \sum_{a=0}^{\bar{a}_u^{ss}(D)} \beta^a (1-p)^a \Psi \left\{ \left[\frac{(1+\gamma)\bar{a}_g^{ss}(D)}{\theta_g} \right]^{\frac{1}{1-\alpha}} \left[\frac{\theta_u}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - 1 \right\} + \\ \sum_{a=0}^{\bar{a}_g^{ss}(D)} \beta^a \varphi [1 - (1-p)^a] \Psi \left\{ \left[\frac{(1+\gamma)\bar{a}_g^{ss}(D)}{\theta_g} \right]^{\frac{1}{1-\alpha}} \left[\frac{\theta_g}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - 1 \right\} - c_0 \end{array} \right) / c_1. \quad (\text{A.11})$$

where $\bar{a}_u^{ss}(D) = \bar{a}_g^{ss}(D) - (\bar{a}_g^{ss}(D) - \bar{a}_u^{ss}(D))$ with $(\bar{a}_g^{ss}(D) - \bar{a}_u^{ss}(D))$ determined by (A.9) independently. Apparently, the right-hand side of (A.11) increases monotonically in $\bar{a}_g^{ss}(D)$. This implies that higher D leads to higher $\bar{a}_g^{ss}(D)$ and $\bar{a}_u^{ss}(D)$. Moreover, the right-hand side of (A.10) also increases monotonically in $\bar{a}_g^{ss}(D)$, which suggests that, by causing higher $\bar{a}_g^{ss}(D)$, higher D will also give higher $f^{ss}(0, D)$. This proves Proposition 4.

A.3. APPROXIMATING VALUE FUNCTIONS WITH KRUSELL AND SMITH (1998) APPROACH

The key computational task is to map F , the plant distribution across ages and idiosyncratic productivity, given demand level D , into a set of value functions $V(\theta^e, a; F, D)$. To make the state space tractable, we define a variable X such that

$$X(F) = \sum_a \sum_{\theta^e} f(\theta^e, a) \cdot q(\theta^e, a), \quad (\text{A.12})$$

where $f(\theta^e, a)$, as a component of F , measures the mass of plants with expected idiosyncratic productivity θ^e and age a . Apparently,

$$P(F, D) \cdot A = \frac{D}{X(F')} = \frac{D}{X(H(F, D))}. \quad (\text{A.13})$$

F' is the updated plant distribution after entry and exit and $F' = H(F, D)$; $P(F, D)$ is the equilibrium price in a period with initial aggregate state (F, D) . Plugging (A.20) into (A.3) gives

$$\pi(a, \theta; F, D) = (\alpha^{\frac{\alpha}{1-\alpha}} - \alpha^{\frac{1}{1-\alpha}}) \left[\frac{D}{X(H(F, D))} \right]^{\frac{1}{1-\alpha}} \left[\frac{\theta^e}{(1+\gamma)^a} \right]^{\frac{1}{1-\alpha}} - \Psi. \quad (\text{A.14})$$

Thus, the aggregate state (F, D) and its law of motion help plants to predict future profitability by suggesting sequences of X 's from today onward under different paths of demand realizations. The question then is: what is the plant's critical level of knowledge of F that allows it to predict the sequence of X 's over time? Although plants would ideally have full information about F , this is not computationally feasible. Therefore we need to find an information set Ω that delivers a good approximation of plants' equilibrium behavior, yet is small enough to reduce the computational difficulty.

We look for an Ω through the following procedure. In step 1, we choose a candidate Ω . In step 2, we postulate perceived laws of motion for all members of Ω , denoted H_Ω , such that $\Omega' = H_\Omega(\Omega, D)$. In step 3, given H_Ω , we calculate plants' value functions on a grid of points in the state space of Ω applying value function iteration, and obtain the corresponding industry-level decision rules – entry sizes and exit ages across aggregate states. In step 4, given such decision rules and an initial plant distribution. We simulate the

TABLE A.1. The estimated laws of motion and measures of fit

Ω	$\{X\}$
H_Ω	$H_x(X, D_h): \log X' = 0.0415$ $+ 0.9937 \log X$ $H_x(X, D_l): \log X'$ $= 0.0495 + 0.9924 \log X$
R^2	for D_h : 0.9996 for D_l : 0.9989
Standard forecast error	for D_h : $4.4 \cdot 10^{-7}\%$ for D_l : $4.7 \cdot 10^{-7}\%$
Maximum forecast error	for D_h : $1.76 \cdot 10^{-6}\%$ for D_l : $1.78 \cdot 10^{-6}\%$
Den Haan and Marcet test statistic (χ^2_7)	0.4228

behavior of a continuum of plants along a random path of demand realizations, and derive the implied aggregate behavior—a time series of Ω . In step 5, we use the stationary region of the simulated series to estimate the *implied* laws of motion and compare them with the *perceived* H_Ω ; if different, we update H_Ω , return to step 3 and continue until convergence. In step 6, once H_Ω converges, we evaluate the fit of H_Ω in terms of tracking the aggregate behavior. If the fit is satisfactory, we stop; if not, we return to step 1, make plants more knowledgeable by expanding Ω , and repeat the procedure.

We start with $\Omega = \{X\}$ —plants observe X instead of F . We further assume that plants perceive the sequence of future coming X 's as depending on nothing more than the current observed X and the state of demand. The perceived law of motion for X is denoted H_x so that $X' = H_x(X, D)$. We then apply the procedure described above and simulate the behavior of a continuum of plants over 10,000 periods. The results are presented in Table A.1.

As shown in Table A.1, the estimated H_x is log-linear. The fit of H_x is quite good, as suggested by the high R^2 , the low standard forecast error, and the low maximum forecast error. The good fit when $\Omega = \{X\}$ implies that plants perceiving these simple laws of motion make only small mistakes in forecasting future prices. To explore the extent to which the forecast error can be explained by variables other than X , we implement the test by Den Haan and Marcet (1994) using instruments $[1, X, \mu_a, \sigma_a, \gamma_a, \kappa_a, r_u]$, where $\mu_a, \sigma_a, \gamma_a, \kappa_a, r_u$ are the mean, standard deviation, skewness, and kurtosis of the age distribution of plants, and the fraction of unsure plants, respectively. The test statistic is 0.4228, well below the critical value at the 1% level. This suggests that given the estimated laws of motion, we do not find much additional forecasting power contained in other variables. Nevertheless, we expand Ω further to include σ_a , the standard deviation of the age distribution of firms. The results when $\Omega = \{X, \sigma_a\}$ are shown in Table A.2.

TABLE A.2. The estimated laws of motion with two moments and measures of fit

Ω	$\{X, \sigma_a\}$
H_Ω	booms ($\log X$): $\log X' = -1.0406 + 0.9954 \log X + 0.1262\sigma_a$ booms(σ_a): $\sigma'_a = 0.2785 - 0.0068 \log X + 0.9754\sigma_a$ recessions($\log X$): $\log X' = -1.0371 + 0.9963 \log X + 0.8988\sigma_a$ recessions(σ_a): $\sigma'_a = 0.2775 - 0.0065 \log X + 0.9751\sigma_a$
R^2	booms ($\log X$): 0.9999 recessions($\log X$): 0.99999 booms (σ_a): 0.9989 recessions(σ_a): 0.9990
Standard forecast error	booms ($\log X$): $1.1 \cdot 10^{-8}\%$ recessions($\log X$): $1.2 \cdot 10^{-8}\%$ booms (σ_a): $6.4 \cdot 10^{-9}\%$ recessions(σ_a): $6.25 \cdot 10^{-9}\%$
Maximum forecast error	booms ($\log X$): $4.87 \cdot 10^{-8}\%$ recessions($\log X$): $5.05 \cdot 10^{-8}\%$ booms (σ_a): $1.48 \cdot 10^{-8}\%$ recessions(σ_a): $1.51 \cdot 10^{-8}\%$
Den Haan and Marcet test statistic (χ^2_7)	0.4375