HOW DESTRUCTIVE IS INNOVATION?

DANIEL GARCIA-MACIA
International Monetary Fund

CHANG-TAI HSIEH
Booth School of Business, University of Chicago and NBER

PETER J. KLENOW
Department of Economics, Stanford University and NBER

Entrants and incumbents can create new products and displace the products of competitors. Incumbents can also improve their existing products. How much of aggregate productivity growth occurs through each of these channels? Using data from the U.S. Longitudinal Business Database on all nonfarm private businesses from 1983 to 2013, we arrive at three main conclusions: First, most growth appears to come from incumbents. We infer this from the modest employment share of entering firms (defined as those less than 5 years old). Second, most growth seems to occur through improvements of existing varieties rather than creation of brand new varieties. Third, own-product improvements by incumbents appear to be more important than creative destruction. We infer this because the distribution of job creation and destruction has thinner tails than implied by a model with a dominant role for creative destruction.

KEYWORDS: Economic growth, innovation, creative destruction.

1. INTRODUCTION

INNOVATING FIRMS can improve on existing products made by other firms, thereby gaining profits at the expense of their competitors. Such creative destruction plays a central role in many theories of growth. This goes back to at least Schumpeter (1939), carries through to Stokey (1988), Grossman and Helpman (1991), and Aghion and Howitt (1992), and continues with more recent models such as Klette and Kortum (2004). Aghion, Akcigit, and Howitt (2014) surveyed the theory. Acemoglu and Robinson (2012) provided historical accounts of countries that stop growing when creative destruction is blocked.

Other growth theories emphasize the importance of firms improving their own products, rather than displacing other firms' products. Krusell (1998) and Lucas and Moll (2014) are examples. Some models combine creative destruction with such "own innovation" by existing firms on their own products—see Chapter 12 in Aghion and Howitt (2009) and Chapter 14 in Acemoglu (2011). A recent example is Akcigit and Kerr (2018).

Still other theories emphasize the contribution of brand new varieties. Romer (1990) is the classic reference. Acemoglu (2003) and Jones (2016) are some of the many follow-ups. Studies such as Howitt (1999) and Young (1998) combine variety growth with quality

growth. Feenstra (1994) and Broda and Weinstein (2006) estimated the importance of variety growth for U.S. imports.

How much each channel contributes to growth matters for innovation policy. Knowledge externalities are a force pushing up the social return to research relative to its private return in all of these models. But business stealing, a powerful force in creative destruction models, boosts the private return relative to the social return. Atkeson and Burstein (forthcoming) analyzed the welfare effects of increasing research in a model with own innovation, creative destruction, and new varieties. The more growth involves business stealing (creative destruction), the smaller the welfare gains they found. Thus, to determine the welfare effects of innovation policy, it is important to know the extent to which growth comes from creative destruction (which entails a lot of business stealing) as opposed to own innovation and new variety creation (which do not).

Ideally, one could directly observe the extent to which new products substitute for existing products. Broda and Weinstein (2010) and Hottman, Redding, and Weinstein (2016) are important efforts along these lines for nondurable consumer goods. Scanner data have not been analyzed in the same way for consumer durables, producer intermediates, or producer capital goods—all of which figure prominently in theories of growth.[1] And scanner data are simply not suitable for comparing the quality of service provided by entering versus exiting establishments in business services, retail trade, and so on.

Likewise, when a new product replaces an existing product, one would like to identify whether the new product is owned by another firm ("creative destruction") or the same firm ("own innovation"). Based on case studies, Christensen (1997) argued that innovation largely takes the form of creative destruction, and almost always from new firms. Akcigit and Kerr (2018) looked at whether patents cite earlier patents by the same firm or by other firms. The case studies and the sample of patenting firms, however, may not be representative of firms in the broader economy. Many innovative firms, particularly outside of manufacturing, do not patent.

In the absence of more direct evidence, we try to infer the sources of growth indirectly from the patterns of job creation and job destruction among all private sector firms in the U.S. nonfarm economy. We use data from the U.S. Longitudinal Business Database (LBD) from 1983 to 2013. The seminal work of Davis, Haltiwanger, and Schuh (1996) documents the magnitude of job flows within U.S. manufacturing, and these flows are commonly used as proxies for the intensity of creative destruction. For example, Decker, Haltiwanger, Jarmin, and Miranda (2014) pointed to the decline in U.S. job reallocation since the 1970s as evidence of a decline in the rate of creative destruction.

We view the LBD data through the lens of an exogenous growth model featuring creative destruction, own innovation, and new varieties. For industries such as manufacturing, the object of innovation may be products. For services and retail, which make up the bulk of the LBD data, innovation may take the form of new and improved establishments. For example, Walmart opening a new store may be akin to adding a new product. A new Walmart store arguably gains market share by offering a distinct variety (the store format, including all the items for sale within it) and/or by offering low prices (due to high process efficiency) relative to existing stores in the local market.

We reach four conclusions from our indirect inference based on LBD data. First, most growth appears to come from incumbents rather than entrants. This is because the employment share of entrants is modest. Second, most growth seems to occur through quality

---

[1]Gordon (2007) and Greenwood, Hercowitz, and Krusell (1997) emphasized the importance of growth embodied in durable goods based on the declining relative price of durables.

improvements rather than brand new varieties. Third, own-variety improvements by incumbents loom larger than creative destruction (by entrants and incumbents). The contribution of creative destruction is around 25 percent of growth, with the remainder mostly due to own innovation by incumbent firms. Fourth, the contribution of entrants and creative destruction declined from 1983–1993 to 2003–2013, while the contribution of incumbent firms, particularly through own innovation, increased.

Influential papers by Baily, Hulten, and Campbell (1992) and Foster, Haltiwanger, and Krizan (2001) likewise used firm-level data to document the contributions of entry, exit, reallocation, and within-plant productivity growth to overall growth in the manufacturing sector. They used accounting frameworks without any model assumptions. In contrast, we analyze the data through the lens of a specific model of growth. Like us, Lentz and Mortensen (2008) and Acemoglu, Akcigit, Alp, Bloom, and Kerr (2018) conducted indirect inference on growth models using firm-level data (from Denmark and the U.S., respectively). They assumed the only source of growth is creative destruction, whereas our goal is to infer how much growth comes from creative destruction versus other sources of innovation (own-variety improvements by incumbents and creation of new varieties). Their results are for manufacturing, whereas ours are for the entire nonfarm business sector.

The rest of the paper proceeds as follows. Section 2 lays out the parsimonious exogenous growth model we use. Section 3 presents the data moments from the LBD that we exploit to infer the sources of innovation. Section 4 presents the model parameter values which best match the moments from the LBD data. Section 5 reports how our results change under alternative measures of job flows. Section 6 concludes.

2. MODELS OF INNOVATION

This section lays out a model in which growth occurs through a combination of creative destruction, own innovation, and new varieties. Although all three types of innovation can contribute to aggregate growth, the goal is to illustrate how they might leave different telltale signs in the micro-data.

Static Equilibrium

Aggregate output is a CES combination of quality-weighted varieties:

$$Y = \left[\sum_{j=1}^{M} (q_j y_j)^{1 - \frac{1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}},$$

where y_j denotes the quantity and q_j the quality of variety j. Labor is the only factor of production. Output of variety j is given by $y_j = l_j$, where l_j is labor used to produce variety j. As in Klette and Kortum (2004), a firm may produce multiple varieties. We assume an overhead cost of production that must be expended before choosing prices and output. The overhead cost allows the highest quality producer to charge the monopoly markup $\frac{\sigma}{\sigma-1}$, as the next lowest quality competitor will be deterred by zero ex post profits under Bertrand competition. Without this assumption, firms would engage in limit pricing and markups would be heterogeneous across varieties as in Peters (2018). With common markups, there is no misallocation of labor across varieties.

Firms face a common wage in a competitive labor market, so the profit maximizing quantity of labor employed in producing variety j is

$$l_j = \left(\frac{\sigma - 1}{\sigma}\right)^{\sigma-1} LW^{1-\sigma} q_j^{\sigma-1}.$$

Here, W is the real wage and L is aggregate employment.[2] Employment of a firm L_f is then given by

$$L_f \equiv \sum_{j=1}^{M_f} l_j = \left(\frac{\sigma - 1}{\sigma}\right)^{\sigma-1} LW^{1-\sigma} \sum_{j=1}^{M_f} q_j^{\sigma-1}, \tag{1}$$

where M_f denotes the number of varieties produced by firm f. As shown, firm employment is proportional to $\sum_{j=1}^{M_f} q_j^{\sigma-1}$. Larger firms control more varieties and produce higher quality products. In the special case of $\sigma = 1$ analyzed by Klette and Kortum (2004), an equal amount of labor is used to produce each variety and a firm's employment is proportional to the number of varieties it controls. We will find it important to allow $\sigma > 1$ so that firms can be larger when they have higher quality products rather than just a wider array of products.

After imposing the labor market clearing condition, the real wage is proportional to aggregate labor productivity:

$$W \propto Y/L = M^{\frac{1}{\sigma-1}} \left[\sum_{j=1}^{M} \frac{q_j^{\sigma-1}}{M}\right]^{\frac{1}{\sigma-1}}.$$

To the right of the equality, the first term captures the benefit of having more varieties, and the second term is the power mean of quality across varieties.

Innovation

Aggregate growth in the model comes from the creation of new varieties and from quality improvements on existing varieties. In Klette and Kortum (2004), the number of varieties is constant and quality growth only occurs when a firm innovates upon and takes over a variety owned by another firm ("creative destruction"). We allow quality growth to also come from innovation by firms on the products they own ("own innovation"). Last, we allow growth to come from the creation of brand new varieties.

We make the following assumptions about innovation. First, we assume a constant exogenous arrival rate for each type of innovation. Second, we assume that arrivals are in proportion to the number of products owned by a firm. For example, a firm with two products is twice as likely to creatively destroy another firm's variety compared to a firm with one product. Third, in the case of an existing product, we assume that innovation builds on the existing quality level of the product.[3]

---

[2]We normalize the price of aggregate output to 1.

[3]If innovation was endogenous, there would be a positive externality to research unless all research was done by firms on their own products. Such knowledge externalities are routinely assumed in the quality ladder literature, such as Grossman and Helpman (1991), Aghion and Howitt (1992), Kortum (1997), and Acemoglu et al. (2018).

TABLE I

CHANNELS OF INNOVATION[a]

Channel	Probability
Own-variety improvements by incumbents	λ_i
Creative destruction by entrants	δ_e
Creative destruction by incumbents	δ_i
New varieties from entrants	κ_e
New varieties from incumbents	κ_i

[a]The average step size for quality improvements for own innovation and creative destruction, weighted by employment, is $s_q = (\frac{\theta}{\theta-(\sigma-1)})^{1/(\sigma-1)} \geq 1$. The quality of a new variety is drawn from the quality distribution of existing products multiplied by s_κ.

The quality drawn by an innovation follows a Pareto distribution with shape parameter θ and a scale parameter equal to the existing quality level. Thus, the proportional step size of innovation on a given variety, $\tilde{q}$, follows a Pareto distribution with shape parameter θ and scale parameter 1. The average proportional improvement in quality on an existing variety, conditional on innovation and weighted by employment, is thus

$$s_q = \left(\frac{\theta}{\theta - (\sigma - 1)} \right)^{1/(\sigma-1)} > 1.$$

More precisely, $s_q \equiv (\mathbb{E}\,\tilde{q}^{\sigma-1})^{1/(\sigma-1)}$.

Finally, we assume that entrants have one product, which they obtain by improving upon an existing variety or by creating a brand new variety.

The notation for innovation probabilities is given in Table I. Time is discrete and innovation rates are per existing variety. The arrival rate of each type of innovation increases linearly with the number of varieties owned by the firm. The probability an existing variety is improved upon by the firm that currently owns the product is λ_i. If a firm fails to improve on a variety it produces, then that variety is vulnerable to creative destruction by other firms. Conditional on not being improved by the incumbent, δ_i is the probability the product is improved by another incumbent. Conditional on not being improved by any incumbent, δ_e is the probability the product will be improved by an entrant.

In short, a given product can be improved upon by the current owner of the product, another incumbent firm, or an entrant. The probability a product will be improved upon by the owner is λ_i. The *unconditional* probability of innovation by another incumbent is $\tilde{\delta}_i \equiv \delta_i(1 - \lambda_i)$. The *unconditional* probability the product will be improved upon by an entrant is $\tilde{\delta}_e \equiv \delta_e(1 - \delta_i)(1 - \lambda_i)$.[4] The probability an existing product is improved upon by any firm is thus $\lambda_i + \tilde{\delta}_i + \tilde{\delta}_e$. And conditional on innovation, the average improvement in quality is $s_q > 1$.

Brand new varieties arrive at rates κ_e from entrants and κ_i from incumbents. These arrival rates are per existing variety and independent of other innovation types. The quality of each new variety is drawn from the quality distribution of existing products multiplied by s_κ, which can be greater or less than 1.

The last parameter we introduce is overhead labor, which pins down the minimum firm size. We set overhead labor so that the smallest firm has 1 unit of labor for production. The overhead cost determines the cutoff quality—varieties below the threshold

---

[4]The ordering of conditional/unconditional probabilities is inconsequential scorekeeping.

have negative present discounted value (even taking into account the arrival of innovations associated with owning a variety), and therefore exit endogenously. The cutoff rises endogenously with wage growth, and ψ denotes the average quality of varieties (relative to the mean quality of all existing varieties) that exit due to the overhead cost. The overhead cost ensures that the distribution of quality across varieties is stationary. We denote the endogenous exit rate of existing varieties due to overhead as δ_o. The net growth rate of varieties is therefore $\kappa_e + \kappa_i - \delta_o$.

The expected growth rate is a function of Table I parameters as follows:

$$\mathbb{E}\big[(1+g)^{\sigma-1}\big] = 1 + \underbrace{s_\kappa(\kappa_e + \kappa_i)}_{\text{new varieties}} + \underbrace{(s_q^{\sigma-1} - 1)\lambda_i}_{\text{own innovation}} + \underbrace{(s_q^{\sigma-1} - 1)(\widetilde{\delta}_e + \widetilde{\delta}_i)}_{\text{creative destruction}} - \delta_o\psi. \qquad (2)$$

The contribution of new varieties is given by $s_\kappa(\kappa_e + \kappa_i)$, which is increasing in the arrival rates κ_e and κ_i and in the quality of new varieties as determined by s_κ. The contribution of own innovation is the product of the probability of own innovation λ_i and the quality improvements $s_q^{\sigma-1}$ associated with them. The contribution of creative destruction is the product of the probability of creative destruction $\widetilde{\delta}_e + \widetilde{\delta}_i$ and the corresponding quality increases. Finally, the loss from the exit of low quality varieties due to overhead costs is captured by $\delta_o\psi$, the product of the frequency and quality of varieties lost.

We can rearrange (2) to express growth from entrants versus incumbents:

$$\mathbb{E}(1+g)^{\sigma-1} = 1 + \underbrace{s_\kappa\kappa_e + \big(s_q^{\sigma-1} - 1\big)\widetilde{\delta}_e}_{\text{entrants}} + \underbrace{s_\kappa\kappa_i + \big(s_q^{\sigma-1} - 1\big)(\lambda_i + \widetilde{\delta}_i)}_{\text{incumbents}} - \delta_o\psi. \qquad (3)$$

Entrants contribute through new varieties κ_e and creative destruction $\widetilde{\delta}_e$, with these arrival rates multiplied by their step sizes. Incumbents contribute new varieties κ_i, own innovation λ_i, and creative destruction $\widetilde{\delta}_i$, where again the arrival rates are multiplied by their corresponding step sizes (s_κ in the case of new varieties and $s_q^{\sigma-1}$ for own innovation and creative destruction).

In sum, the innovation probabilities (κ_i, κ_e, λ_i, δ_i, and δ_e) along with the quality steps (θ and s_κ) pin down the share of growth from own innovation, creative destruction, and new varieties. These parameters also determine the share of growth driven by incumbents versus entrants. We will estimate these parameters from patterns in the LBD micro-data to infer the sources of growth.

Firm Dynamics

A firm's employment is proportional to the number of products the firm produces and the average quality of those products, as shown in equation (1) above. Thus, for a random firm, we expect

$$L_f \propto \sum_{j=1}^{M_f} q_j^{\sigma-1} = M_f \left(\frac{1}{M_f} \sum_{j=1}^{M_f} q_j^{\sigma-1} \right) = M_f E\big[q^{\sigma-1}\big].$$

A firm's employment *growth* is the outcome of all three types of innovation. However, the magnitude of the employment growth depends on the precise type of innovation.

Suppose a random firm "own-innovates" on all its M_f products with the expected step size s_q. Employment of this firm will be equal to

$$L_f^{\mathrm{OI}} \propto M_f E\big[q^{\sigma-1}\big] \times s_q,$$

where "OI" stands for own innovation. This firm's employment growth is $s_q - 1$.

Conversely, if a random firm creatively destroys one additional product *per existing product*, its employment will become

$$L_f^{\mathrm{CD}} \propto \underbrace{M_f E\big[q^{\sigma-1}\big]}_{\text{old products}} + \underbrace{M_f E\big[q^{\sigma-1}\big] \times s_q}_{\text{new products}} = M_f E\big[q^{\sigma-1}\big](1 + s_q),$$

where "CD" stands for creative destruction. This firm's employment growth is equal to s_q. The firm's employment from its new products is proportional to that from its old products (in expectation) because creative destruction is simply a random draw from the existing distribution of qualities.

The corresponding rates of firm-level job creation in these hypothetical scenarios (holding the wage W_t constant) are given by

$$JC^{\mathrm{OI}} = \frac{L_f^{\mathrm{OI}} - L_f}{\frac{1}{2}\big(L_f^{\mathrm{OI}} + L_f\big)} = 2\frac{s_q - 1}{s_q + 1},$$

$$JC^{\mathrm{CD}} = \frac{L_f^{\mathrm{CD}} - L_f}{\frac{1}{2}\big(L_f^{\mathrm{CD}} + L_f\big)} = 2\frac{s_q}{s_q + 2}.$$

Crucially, $JC^{\mathrm{CD}} > JC^{\mathrm{OI}}$ even though the aggregate productivity effect is *exactly* the same: M_f products in the economy saw their quality increase by factor s_q. Below, we will find that job creation from creative destroyers will be an order of magnitude larger than job creation from own innovators, for a given s_q. For the same innovation, creative destruction will show up in the right tail of job creation (and job destruction), whereas own innovation will show up as firms experiencing modest rates of job creation.

To recap, a firm that is successful in innovating will grow in employment, and the *magnitude* of the employment growth depends on whether the firm improved the quality of its own products or the products made by another firm. To convey how we will use this idea in our data inference, we now highlight the predictions of three polar models, each with only one source of innovation.

Creative Destruction

Consider a polar model where the *only* source of innovation is creative destruction.[5] Further assume $\sigma = 1$ so quality has no effect on firm employment. This is simply Klette and Kortum (2004). In this polar model, incumbent firms grow when they take over another firm's product and shrink when another firm innovates on their products. The rate

---

[5]Here, we assume $\kappa_i = \kappa_e = \lambda_i = 0$. To fix the other parameter values, we use the 2003–2013 simulated values we obtain in the next section (shown in Table III): the Pareto shape parameter for quality draws $\theta = 17.3$ and the unconditional arrival rates $\widetilde{\delta}_i = (1-0.779)0.283 = 6.2\%$ and $\widetilde{\delta}_e = (1-0.779)(1-0.283) = 15.8\%$. A period is 5 years in our calibration.
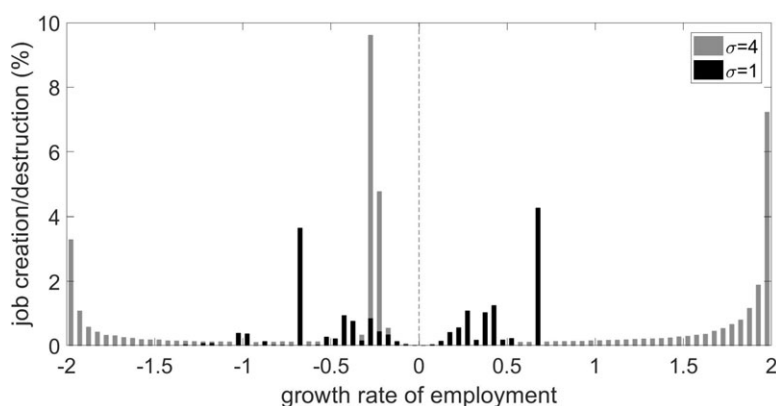
FIGURE 1.—Job creation and destruction with only creative destruction. *Note*: The figure is based on simulating a model with only creative destruction as a source of growth. Employment growth for a firm is defined as the change in employment divided by average employment at the firm at the beginning and end of each period. The vertical axis gives the share of total job creation (destruction) associated with firms at each given level of employment growth. Entry (+2) and exit (−2) are omitted in the figure. σ is the elasticity of substitution across varieties.

of job destruction is tied to the rates of creative destruction by incumbent firms δ_i and entrants δ_e. The other side of job destruction is job creation, which can come from entrants or from incumbents. The rate of creative destruction by incumbents δ_i determines the job creation rate by incumbents, and the unconditional rate of innovation by entrants $(1 - \delta_i)\delta_e$ pins down the job creation rate by entrants. The two parameters δ_e and δ_i collectively determine the aggregate rate of job creation and job destruction.

This polar model has specific predictions for the *distribution* of job creation and job destruction across firms. See the bars labeled $\sigma = 1$ in Figure 1. Following Davis, Haltiwanger, and Schuh (1996), the percent change in firm employment on the horizontal axis is measured as the change in firm employment between time t and $t + 1$ divided by the average of the firm's employment at time t and $t + 1$. The rates are thus bounded between −2 (exit) and +2 (entry). The density on the vertical axis is the percent of all job creation or destruction contributed by firms in each bin of employment growth. For visual clarity, the figure omits firm exit (−2) and firm entry (+2).

The distribution of job creation and destruction in this polar model is concentrated at a small number of discrete bins of employment growth. Product quality has no effect on firm employment in this model because $\sigma = 1$; firm employment is only a function of the number of varieties the firm produces. The distribution of job creation and job destruction therefore reflects the change in the number of varieties across firms. Figure 1 shows that most expanding firms double their number of varieties (the bin with employment growth = 0.67). Conditional on shrinking, the majority of shrinking firms lose half of their varieties (the bin with employment growth = –0.67). The density of *small* changes in job creation and destruction depends on the fraction of firms with multiple products that either gain or lose a small number of these products. For example, the bin with employment growth = –0.4 are firms that lose a third of their products.

The model has implications for two additional moments in the data. First, growth in firm employment by age is driven by the accumulation of varieties. Life cycle growth is therefore determined by the rate at which incumbent firms improve upon the varieties of other firms, δ_i. Second, the model predicts that firm exit rates will fall sharply with firm
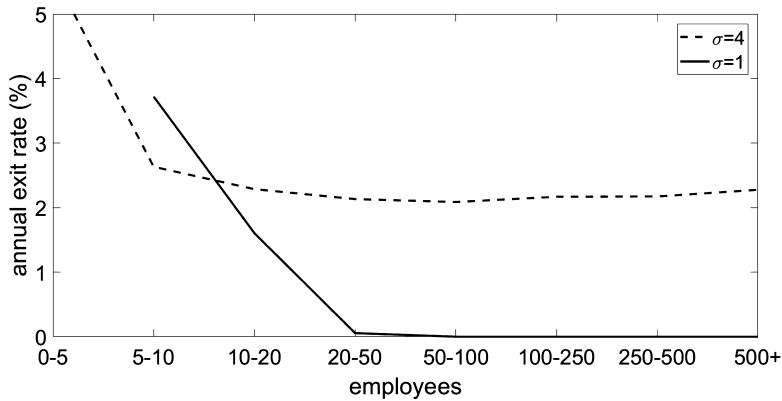
FIGURE 2.—Exit by size with only creative destruction. *Note*: The figure is based on simulating a model with only creative destruction as a source of growth. The exit rate is annualized from a model with 5-year periods. σ is the elasticity of substitution across varieties.

size. To see this, note that a firm with n varieties exits when other firms innovate upon and take over all n of its varieties. The probability that a firm with n varieties exits is thus given by the exit probability of a one-variety firm to the *power* of n. Since the employment of a firm is proportional to the number of varieties it produces, the model predicts that an n-fold difference in firm employment will be associated with a change in the exit rate to the power of n. Figure 2 illustrates the relationship between firm exit and firm employment predicted by this polar model.

To get quality to matter for firm employment, we need to drop the assumption that $\sigma = 1$. Figure 2 shows that raising σ from 1 to 4 flattens the exit-size slope. This change also makes the distribution of job creation and job destruction more continuous, as shown in Figure 1. Employment growth rates are now a function of the change in average quality as well as the change in the number of varieties. In addition, changing σ to 4 makes the tails of the distribution of employment growth thicker. Firms will experience large increases in employment when they grab a high quality variety from another firm, and sharp drops in employment when they lose their high quality varieties. An empirical challenge for this creative-destruction-only model, even with $\sigma > 1$, will be generating small increases in employment at many firms.

Own Innovation

We next consider a model where the *only* source of growth is own innovation.[6] This polar model has its own stark properties. The share of entrants is zero because there are neither new varieties nor creative destruction from entrants. The exit rate is zero because there is no creative destruction. Figure 3 plots the distribution of job creation and destruction. Firms grow only when they innovate on their products. The distribution of job creation is only a function of the heterogeneity across firms in quality improvements.[7]

---

[6]We assume $\kappa_i = \kappa_e = \delta_i = \delta_e = 0$, $\lambda_i = 1$, and no overhead costs. The scale parameter is as in the 2003–2013 simulation (described later in Section 4): $\theta = 17.3$.

[7]This model does not have a stationary quality or employment distribution across firms. But the distribution of job creation and destruction is stationary, as it only depends on the Pareto distribution of quality improvements (steps s_q).
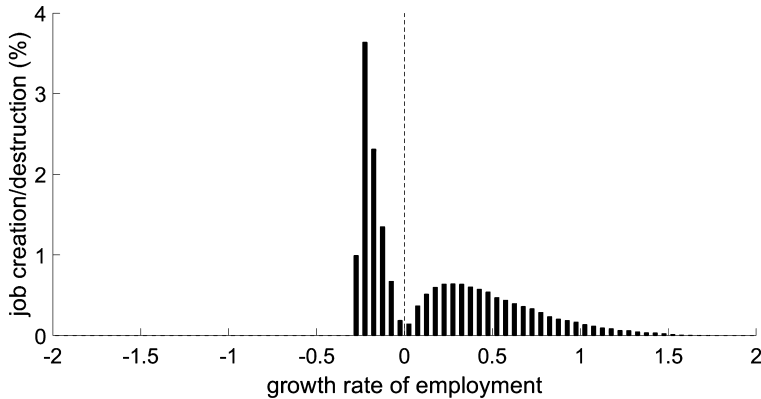
FIGURE 3.—Job creation and destruction with own innovation. *Note*: The figure is based on simulating a model with only incumbent improvement of their own products (own innovation). Employment growth for a firm is defined as the change in employment divided by average employment at the firm at the beginning and end of each period. The vertical axis gives the share of total job creation (destruction) associated with firms at each given level of employment growth. We used $\sigma = 4$ for the elasticity of substitution between varieties.

Firms that do not improve on their products (at all or enough) shrink due to the general equilibrium effect of a rising real wage, and this is the only force that generates job destruction in the model. This effect can be seen in the spike of employment declines by 25 percent. If all firms innovated, then this spike would not be present. The model predicts that no firms experience an employment decline in excess of 25 percent. Those firms who innovate but draw small steps shrink more modestly.

Own Innovation + Creative Destruction by Entrants

The extreme empirical predictions of a model with only own innovation follow from the absence of creative destruction. Thus, consider a hybrid model in which incumbents improve the quality of their own products and entrants engage in creative destruction.[8] This hybrid model has the following implications. First, the employment share of entrants is positive and pinned down by the rate of creative destruction by entrants δ_e. Second, the tail of job creation is thin because there is no creative destruction by incumbent firms. The tail of job destruction is thin because job destruction is pinned down by the employment share of entrants. Third, entrants are slightly larger than incumbents on average. This is because all entrants improve on incumbent quality, whereas only a subset of surviving incumbents improve their quality. Finally, since incumbent firms can only innovate by improving their one product, larger firms are larger only because of higher quality, not because they produce more varieties. As a result, the probability that a firm exits is the same regardless of its size.

If the data are inconsistent with the predictions of this hybrid model, then it might help to add creative destruction from *incumbents*. Creative destruction from incumbent firms will thicken the tails of job creation and destruction. It will also generate heterogeneity in the number of varieties across firms. Older firms will tend to have more varieties than young firms. This implies that average employment will tend to be higher in older firms

---

[8]We keep the same parameters as in the own innovation model and we use the same δ_e as in the model with only creative destruction.

relative to young firms. Larger firms will also tend to exit less than small firms do, because larger firms will have more varieties on average.

More creative destruction will thicken the tail of job destruction. But the frequency of *small* employment declines depends on the fraction of varieties that do not improve or that innovate but draw small steps. Employment falls in such firms due to the general equilibrium effect of a rising real wage.

New Varieties

We now consider the effect of allowing firms to create new varieties. A model where firms *only* create new varieties has stark predictions that are not likely to hold empirically. As in the model where growth is only driven by own innovation, there is no exit and the only source of job destruction is the general equilibrium effect of a rising real wage on firms that do not create new varieties.[9] So new varieties will need to be combined with other sources of innovation.

How might we infer new variety creation ($\kappa_e + \kappa_i$) from the data moments we have? Constant arrival rates per variety turn out to imply a constant steady-state ratio of total varieties to the total number of firms. The total number of firms is by definition the product of average employment per firm and total employment. As we will see later, average firm size is fairly stable in the data so we will infer growth in the number of varieties from the growth of total employment.

How do we know whether new varieties come from entrants or incumbents (κ_e vs. κ_i)? Total innovation from entrants will be disciplined by the employment share of entrants. If new varieties come from incumbents, this will be a source of life cycle employment growth (i.e., firm size increasing with age).

Finally, how will we infer how *good* new varieties are? Suppose new varieties are of lower quality than existing qualities (i.e., $s_\kappa < 1$). This will be a force increasing the dispersion of quality and firm size. If entrants make them, this will tend to make young firms smaller than old firms. If incumbents create these low quality new varieties, that will increase the mass of job creation at lower values of employment growth.

Recap on Innovation and Job Flows

We will use data on job flows to speak to the sources of innovation. Motivated by the preceding discussion of how to discriminate between sources, we will examine the following 10 data moments:

1. Aggregate TFP growth rate.
2. Standard deviation of log employment across firms.
3. The employment share of entrants.
4. Job creation rate.
5. Job destruction rate.[10]
6. Share of job creation due to firm employment growth ≤ 1.[11]
7. Minimum firm employment (1).
8. Exit rate for small firms (firms with below-average employment).
9. Exit rate for large firms (firms with above-average employment).

---

[9]This is assuming no overhead cost.

[10]In fitting the job creation and destruction rates, we also fit employment growth.

[11]Recall that firm employment growth rate is defined as the ratio of the change in employment to the average of initial and final employment. A growth rate of 1 is therefore a three-fold increase in employment relative to initial employment.

10. Average employment for incumbents relative to entrants.

These moments are a mix of life cycle (3, 8–10), job flow (4–6), cross-sectional (2, 7), and aggregate (1) moments.[12]

3. U.S. LONGITUDINAL BUSINESS DATABASE

We use firm-level data on employment from the U.S. Census Longitudinal Business Database (LBD). The LBD is based on administrative employment records of every nonfarm private establishment in the U.S. economy. The advantages of the LBD are its broad coverage and its quality (e.g., the Census uses it to identify and correct for measurement error in its quinquennial Census surveys).

The establishment-level variables we use are employment, industry (4-digit SIC or 6-digit NAICS), the year the establishment appears in the LBD for the first time, the establishment's ID, and the ID of the firm that owns the establishment. We use the year the establishment appears in the LBD to impute the establishment's age (the LBD does not provide the establishment's age directly). We restrict the sample to 1983–2013 and drop establishments in the public, educational, agricultural, and mining sectors.

We focus on firms rather than establishments because business stealing has clear implications for innovation policy (Atkeson and Burstein (forthcoming)). In our baseline sample, we aggregate the data of establishments within a firm. Firm employment is the sum of employment at the establishments owned by a firm. Firm age is the age of the oldest establishment owned by the firm. An "entrant" is a firm for which the oldest establishment was created within the last five years. An "incumbent" is a firm for which the oldest establishment was created five or more years earlier. A firm "exits" when it loses all its establishments in the next five years. We calculate 5-year moments because various adjustment costs might suppress an entrant's market share relative to the quality of its products.[13]

Our firm employment dynamics include the direct positive and negative effects of mergers and acquisitions. Such M&A activity can be the result of the innovation forces we model in this paper. A firm may acquire another firm or establishment to implement an improvement on the target firm's products. Still, to check the robustness of our estimates, we also use an alternative sample where we drop establishments that undergo ownership changes, and calculate job creation and destruction rates based on this sample.

Table II presents some summary statistics for our LBD sample. Total employment and the number of firms increased from 1983 to 2013, but average employment per firm and the dispersion of firm size were fairly stable. The growth rate of aggregate employment was much slower from 2003–2013 than from 1983–1993. Since average employment per firm was roughly constant, the growth rate of the number of firms also fell.

The last column in Table II shows the annual TFP growth rate in each 10-year period: 1.66% from 1983 to 1993, 2.30% from 1993 to 2003, and 1.32% in the 2003–2013 period. So 1993–2003 was the high growth period and 2003–2013 the low growth period. These estimates are from the U.S. Bureau of Labor Statistics and cover the nonfarm business sector, like the LBD.

---

[12]We think the life cycle, job flow, and aggregate moments are critical for identification. In future work, it would be useful to explore whether identification can be achieved without relying on the two cross-sectional moments.

[13]Haltiwanger, Jarmin, and Miranda (2013) showed that plants in the LBD grow faster than average until age five.

TABLE II

SUMMARY STATISTICS[a]

	Employment	Firms	Average Employment	S.D. of log Employment	Employment Growth	TFP Growth
1983–1993	93.6	4.287	22	1.25	2.4%	1.66%
1993–2003	112.0	4.857	23	1.27	1.6%	2.30%
2003–2013	125.2	5.301	24	1.28	0.5%	1.32%

[a]Sources: Columns 1–5 are from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. Total employment and number of firms in millions. The last two columns give the average annual growth rates of total employment and TFP, respectively. Data on TFP growth are from the U.S. Bureau of Labor Statistics.

Next, Figure 4 presents the overall rates of job creation and destruction, calculated over five years. We remove net job creation at the industry level (either 4-digit SIC or 6-digit NAICS) so the job flows in Figure 4 are mostly *within* industries. If we do not net out employment changes at the industry level, the aggregate job creation and destruction rates are about 1.3 percentage points higher. The job creation (destruction) rate labeled '1983–1993' is the average statistic from 1983–1988 and 1988–1993. The job creation (destruction) rate is the sum of employment changes at firms with rising (falling) employment divided by the average of aggregate employment in the initial and final year of each period. This includes entering and exiting firms. The job creation and destruction rates labeled '1993–2003' and '2003–2013' are defined analogously.

The job creation rate fell by 12 percentage points from 1983–1993 to 2003–2013, while the job destruction rate fell by only 2 percentage points. This is consistent with the decline in aggregate employment growth from 1983–1993 and 2003–2013 shown in Table II. Since we control for net employment changes at the industry level, the decline in job creation is driven by the change in job flows across firms in the same industry. Most of the decline in the aggregate job reallocation rate highlighted by Decker et al. (2014) was due to the decline in the job creation rate. The decline in the job destruction rate was much smaller.
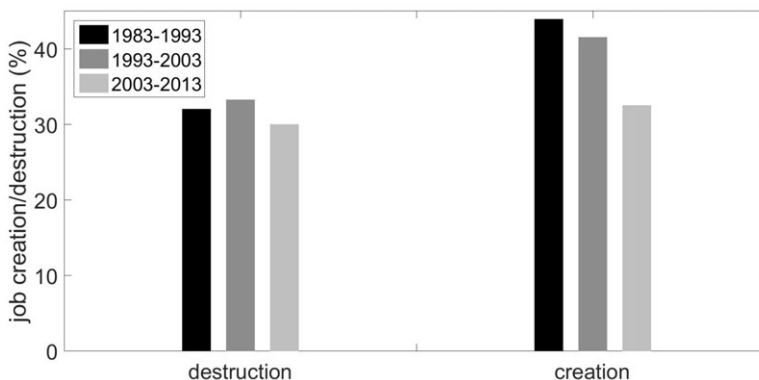

FIGURE 4.—Job creation and destruction rates. *Note*: The job creation (destruction) rate is the sum of employment changes at firms with rising (falling) employment divided by the average of aggregate employment in the initial and final year of each 5-year period. This includes entering and exiting firms. The figure shows the average 5-year changes for 1983–1988 and 1988–1993, 1993–1998 and 1998–2003, and 2003–2008 and 2008–2013. Statistics computed from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.
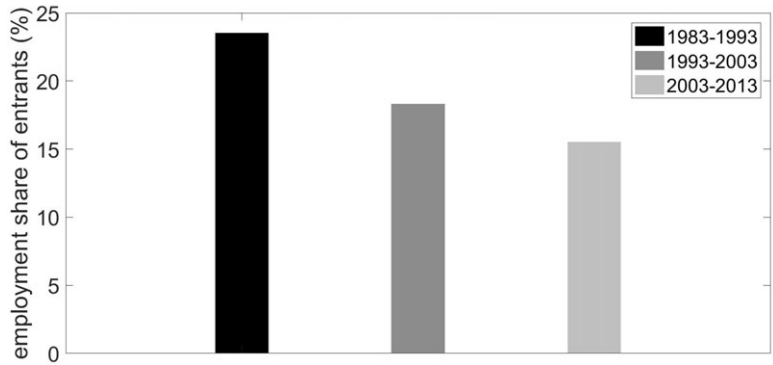
FIGURE 5.—Employment share of entrants. *Note*: The employment share of entrants for 1983–1993 is the average share of entrant employment in 1988 and 1993. Entrant employment in 1988 (1993) is employment of firms that entered between 1983 and 1988 (1988 and 1993), and the share is relative to aggregate employment in 1988 (1993). The entrant employment shares for 1993–2003 and 2003–2013 are defined analogously. Statistics computed from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.

Aggregate job creation is the sum of job creation by incumbent firms and job creation by entering firms. Figure 5 presents the job creation rate due to entrants—the employment share of firms that did not exist five years earlier. The employment share of entrants labeled '1983–1993' is the average of employment of entrants in 1988 as a share of total employment in 1988 and employment of entrants in 1993 as a share of total employment in 1993. Similarly, the label '1993–2003' ('2003–2013') refers to the average of the employment share of entrants in 1998 and 2003 (2008 and 2013).

The employment share of entrants fell by 8 percentage points between 1983–1993 and 2003–2013, as displayed in Figure 5. The aggregate job creation rate, shown previously
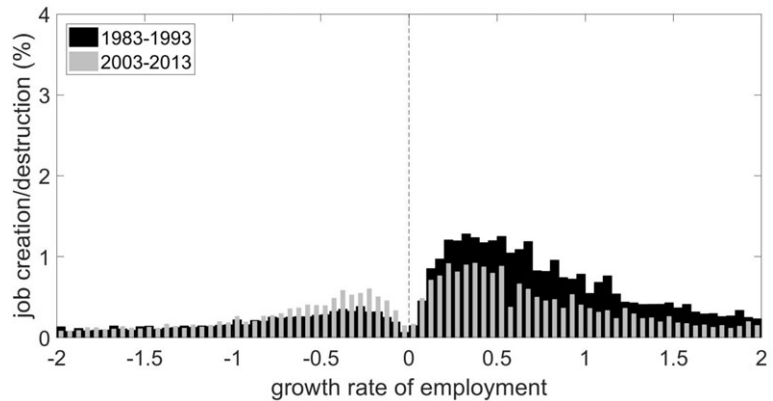
FIGURE 6.—Distribution of job creation and destruction. *Note*: Employment growth for a firm is defined as the change in firm employment over (say) 1988 to 1993 divided by the firm's average employment in 1988 and 1993. The vertical axis gives the share of total job creation (destruction) associated with firms at each given level of employment growth. 1983–1993 refers to averaging these job creation and destruction rates in the two periods. 2003–2013 entries are defined analogously. Entry ($+2$) and exit (-2) are omitted in the figure. Statistics computed from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.

FIGURE 7.—Share of "small" job flows. *Note*: Small flows in job creation is the share of total job creation from firms where e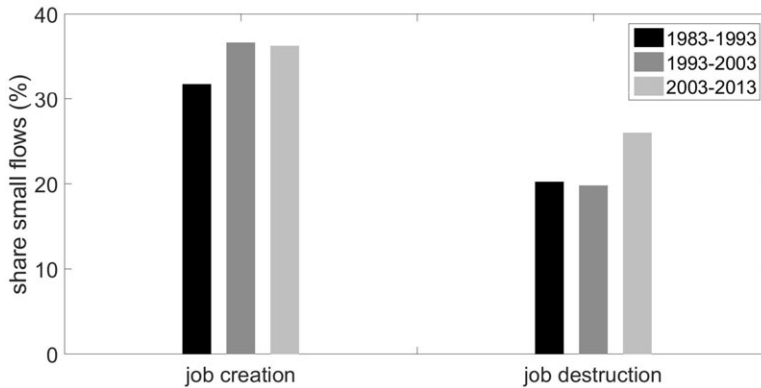mployment increased by less than a factor of 3. Small flows in job destruction is share of total job destruction from firms where employment declined by less than a factor of 3. Statistics computed from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.

in Figure 4, declined by 12 percentage points. So about two-thirds of the decline in the aggregate job creation rate was due to the decline in the employment share of entrants.

Figure 6 plots the *distribution* of job creation and destruction in the LBD. We plot averages from 1983–1988 and 1988–1993 (labeled as '1983–1993') and averages from 2003–2008 and 2008–2013 (labeled as '2003–2013'). The growth of firm employment on the x-axis is measured as the change in firm employment divided by the average of the firm's employment in the initial and final years. These rates are bounded between -2 (exit) and $+2$ (entry). The vertical axis shows the percent of all creation or destruction contributed by firms in each bin. These definitions follow Davis, Haltiwanger, and Schuh (1996).

The empirical distribution of job creation and destruction in Figure 6 looks very different from the distribution in the models with only creative destruction (Figure 1) or own innovation (Figure 3). There is much more mass on smaller changes in employment in the data compared to the polar model with only creative destruction. And there is far greater mass in the tail of job destruction in the data than in a model of incremental growth through own innovation. The empirical moment we use from Figure 6 is the share of job creation at firms where employment increases by less than a factor of 3. This number, shown in Figure 7, is 32% in 1983–1993 and 36% in both 1993–2003 and 2003–2013. The share of job *destruction* at firms where employment declines by less than a factor of 3 averages 23% over the thirty years of our data.

We next present average employment of entrants and incumbents in 1988, 1998, and 2008 (Figure 8). Hsieh and Klenow (2014) documented rapid growth of surviving plants in the U.S. Census of Manufacturing. Figure 8 suggests that the same is true for the entire U.S. private sector. The model can explain this fact if older firms have more and better products compared to young firms.

Figure 9 shows the exit rate of large versus small firms, where large firms are defined as those with above-mean employment and small firms as those with below-mean employment. Here, the exit rate is the annualized probability the firm exits within the next *five* years. The label '1983–1993' refers to the average of the annualized exit rates from 1983–1988 and 1988–1993. The labels '1993–2003' and '2003–2013' are the corresponding average of the exit rates from 1993–1998 and 1998–2003 and from 2003–2008 and 2008–2013. The figure shows that smaller firms have higher exit rates than larger firms. Our
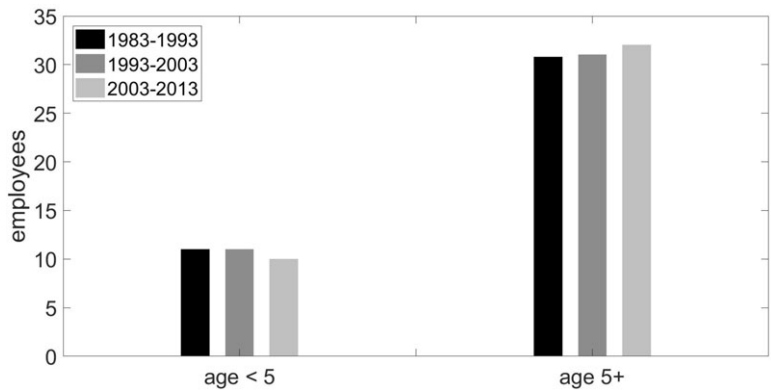
FIGURE 8.—Employment per firm, young versus old. *Source*: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.

model setup can match this qualitative pattern if larger firms produce more varieties than do smaller firms.

In contrast to the prediction of the polar model with only creative destruction, the decline in exit rates with firm size in Figure 9 is gradual. On average, large firms in the data are about 16 times larger than small firms. With only creative destruction, therefore, the exit rate of large firms should be equal to that of small firms to the 16th power, or essentially zero. In the data, the exit rate of large firms (around 6%) is only 2 percentage points lower than the exit rate of small firms (about 8%).

It is worth stressing that *annual* exit rates differ more starkly between small firms and large firms. From 2003–2004 to 2012–2013 in the LBD, the 1-year exit rate averages 12.0% for small firms and 1.0% for large firms. By looking at whether firms exit after 5 years, we allow more time for big firms to shrink (perhaps lose multiple products) and exit. Also,
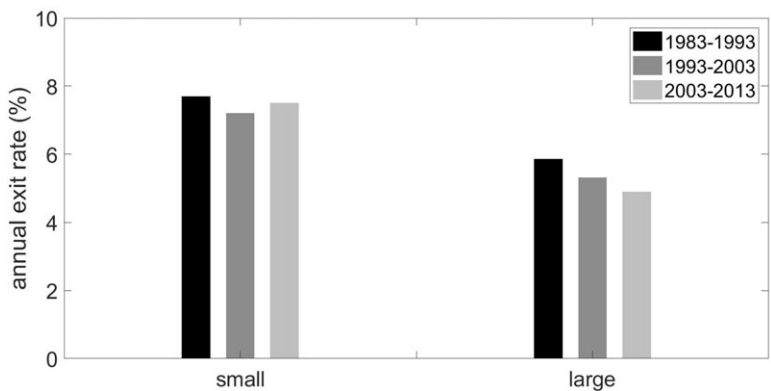
FIGURE 9.—Exit Rate, Large vs. Small Firms. *Note*: The exit rate is the annualized fraction of firms that operated in (say) 1983 but not in 1988. The data for 1983–1993 are annual averages of the 5-year exit rates from 1983 to 1988 and 1988 to 1993, respectively. The 1993–2003 and 2003–2013 exit rates are defined analogously. Exit rates are computed from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. Small firms are defined as those with below-mean employment, and large firms as those with above-mean employment.

TABLE III

INFERRED PARAMETER VALUES[a]

	1983–1993	1993–2003	2003–2013
Own-variety improvements by incumbents λ_i	70.1%	71.0%	77.9%
Creative destruction by incumbents δ_i	29.6%	49.0%	28.3%
Creative destruction by entrants δ_e	100.0%	100.0%	100.0%
New varieties from incumbents κ_i	0.0%	0.0%	0.0%
New varieties from entrants κ_e	12.3%	8.9%	3.0%
Pareto shape of quality draws θ	15.4	11.0	17.3
Relative quality of new varieties s_κ	0.31	0.36	0.30
Average quality of exiting products ψ	0.018	0.043	0.025

[a]Parameter estimates implied by the data moments described in Section 3.

our annualized 5-year exit rates are less subject to a partial-year problem which might make exiting firms seem smaller than they were in their last full year of production.

4. SOURCES OF GROWTH

We now estimate parameters to match moments from model simulations to the corresponding moments in the U.S. LBD. We define a period in the model as *five years*. We need to estimate five innovation rates (δ_i, δ_e, λ_i, κ_i, and κ_e), two quality step size parameters (θ and s_κ), and the overhead cost. So eight parameters. In our base case, we target the aggregate rate of TFP growth, minimum employment per firm of 1, the cross-sectional standard deviation of log firm employment, the aggregate rates of job creation and destruction (the difference is the growth rate of employment), the employment share of entrants, the share of job creation < 1, average employment for old relative to young firms, and exit rates for both small and large firms. So 10 moments.[14]

We choose the overhead cost so that the smallest firm has a single production worker. The overhead cost endogenously pins down the rate at which varieties disappear (δ_o) and their average quality (ψ). Conditional on the overhead cost, we make sure the combination of δ_i, δ_e, λ_i, κ_i, κ_e, θ, and s_κ are such that expected TFP growth in the model exactly equals average TFP growth in the data. We choose the individual parameters to best fit the remaining data moments.

Appendix A provides more detail on the simulated method of moments we deploy. In short, we choose parameter values to minimize the mean squared percent distance between the simulated and empirical moments for the eight moments other than aggregate TFP growth and minimum firm size. We weight moments equally because, given the large number of firms in the LBD, sampling error is a minor consideration for all of the moments.

Table III presents the eight parameter values inferred from the data using the procedure described above.[15] Based on the data moments from 1983 to 1993, we infer a 70% arrival rate of own-variety quality improvements per 5-year period. Conditional on no own innovation, quality improvements through creative destruction occur 30% of the time by other incumbents. Conditional on no own innovation and creative destruction by another incumbent, quality improvement through creative destruction by entrants occurs

---

[14]We also choose the level of employment in the model to fit employment per firm in the data.

[15]Table XVII in Appendix B presents bootstrapped standard errors for the parameter estimates.

TABLE IV

TFP GROWTH FROM INNOVATION, IN %[a]

	1983–1993	1993–2003	2003–2013
Creative destruction	0.44%	0.64%	0.29%
New varieties	0.23%	0.19%	0.06%
Own innovation	0.99%	1.47%	0.98%
All sources	1.66%	2.30%	1.32%

[a]The table presents the product of the arrival rate of each source of innovation and the average improvement in quality (conditional on innovation) implied by the innovation parameters in Table III.

with probability 1. The *unconditional* probability that a given product improves due to creative destruction by an incumbent is thus 8.9%, and the unconditional probability of creative destruction by an entrant is 21.0%.[16] The unconditional probability that a product is improved upon in a 5-year period is thus 100%, of which 69% is from own innovation and 31% is from creative destruction (the latter from entrants or incumbents).

The employment-weighted average step size for quality improvements on existing varieties is given by $s_q = (\theta/(\theta - (\sigma - 1)))^{1/(\sigma-1)}$. Given that $\theta = 15.4$ and $\sigma = 4$, the average improvement in quality (conditional on innovation) is 7.5%.[17] New varieties are only created by entrants, arrive with 12.3% probability per existing variety, and have an average quality that is 31% of the average quality of existing varieties. Overhead costs imply that the average quality of exiting products ψ is 2% of the average quality of existing varieties, and that the probability a variety exits due to overhead cost δ_o is essentially zero. The net number of varieties thus grows by 12.3% every five years, which matches the growth of total employment and number of firms from 1983 to 1993.

Table IV presents TFP growth due to creative destruction (row 1), new varieties (row 2), and own innovation (row 3). TFP growth due to each source of innovation is the product of arrival rate of innovation and the quality improvement conditional on innovation. We use equation (2) for this calculation.[18] The first column shows that TFP growth due to own innovation was about 1% per year in 1983–1993. Growth due to creative destruction was about half that, at 0.44% per year. And new varieties generate growth of 0.23% per year.

The rows in Table V show the contribution of each source of innovation to aggregate TFP growth. About 27% of the 1.66% growth rate in the 1983–1993 period comes from creative destruction. Own-variety improvements by incumbents account for 60%. New varieties à la Romer (1990) are the remainder at around 14%.

The columns in Table V also decompose aggregate TFP growth into the percentage contribution of entrants versus incumbents using equation (3). In the ten years between 1983 and 1993, incumbents account for 68% of aggregate TFP growth, with entrants contributing the remaining 32%. Aghion, Akcigit, and Howitt (2014) provided complementary evidence for the importance of incumbents based on their share of R&D spending and patents.

---

[16]$(1 - 0.701) \cdot 0.296 = 0.089$ and $(1 - 0.701) \cdot (1 - 0.296) = 0.210$.

[17]Interestingly, our estimate of θ is lower at around 11 for 1993–2003, which implies a higher average step size of around (coincidentally) 11% in the fast-growth middle decade. When growth faltered in the 2003–2013 period, our estimate of θ rose to around 17 and the corresponding average step size fell to 6.5%.

[18]Equation (2) is nonlinear so there is no unique decomposition. For the table, we calculate growth from each source in isolation. The contributions are very similar if we instead subtract each source individually from overall growth.

TABLE V

CONTRIBUTION TO AGGREGATE GROWTH[a]

	Entrants	Incumbents	
1983–1993			
Creative destruction	18.6%	7.9%	26.5%
New varieties	13.7%	0.0%	13.7%
Own-variety improvements	–	59.8%	59.8%
	32.3%	67.7%	
1993–2003			
Creative destruction	14.2%	13.6%	27.8%
New varieties	8.3%	0.0%	8.3%
Own-variety improvements	–	63.9%	63.9%
	22.4%	77.6%	
2003–2013			
Creative destruction	15.6%	6.2%	21.9%
New varieties	4.2%	0.0%	4.2%
Own-variety improvements	–	74.0%	74.0%
	19.8%	80.2%	

[a]The table presents the percentage contribution of each source of innovation to aggregate TFP growth. Rows use equation (2) to decompose aggregate TFP into the contribution of creative destruction, new varieties, and own innovation. Columns use equation (3) to decompose aggregate TFP into the contribution of entrants and incumbents.

Table III shows that the arrival rate of new varieties κ_e fell from 12.3% in 1983–1993 to 8.9% in 1993–2003. We infer this because the employment growth rate fell from 2.4% in 1983–1993 to 1.6% in 1993–2003. Since average firm size—and thus varieties per firm—was roughly constant, we infer from the drop in employment growth that the growth rate in the total number of varieties must also have fallen. So TFP growth from new variety creation fell modestly from 0.23% per year in 1983–1993 to 0.19% per year in 1993–2003.

The share of job creation < 1 also increased to 37% in the high growth period (from 32% in 1983–1993). The increase in the importance of small changes in employment implies that the arrival rate of own-variety improvement increased relative to the rate of creative destruction. Remember that own-variety innovation only results in small changes in employment, whereas creative destruction generates large changes in employment. Comparing the 1983–1993 to the 1993–2003 period, the own innovation rate increased modestly from 70% to 71% (Table III). TFP growth from own innovation rose from 1% to almost 1.5% per year (Table IV). Aggregate TFP growth (from all three sources of innovation) increased by 64 basis points from 1983–1993 to 1993–2003 (Table II). Therefore, own innovation was responsible for almost 80% of this increase in aggregate TFP growth, with the remaining 20% due to more innovation from creative destruction.

Turning to 2003–2013, the main fact is the 8 percentage point decline in the employment share of entrants between 1983–1993 and 2003–2013 seen in Figure 5. The model interprets the decline as reflecting less innovation by entrants. How much is less new variety creation versus creative destruction by entrants? New variety creation is pinned down by the growth rate of employment, which fell to 0.5% in 2003–2013. The implied arrival rate of new varieties thus fell to 3.0%. The residual entrant employment due to new varieties is creative destruction by entrants. Table III indicates that the unconditional probability of creative destruction by an entrant fell from 21.0% in 1983–1993 to 15.8% by 2003–2013. The net effect is that TFP growth from innovation by creation of new varieties by entrants fell to 6 basis points per year in 2003–2013, down from 23 basis points

TABLE VI

Job Flows versus Innovation by Age, 2003–2013[a]

	% Job Creation	% TFP Growth
Age < 1	31%	9%
Age 1–5	13%	14%
Age 5–10	11%	14%
Age 10–15	9%	12%
Age > 15	36%	51%

[a]Entries are the share of job creation and aggregate TFP growth due to firms in each age group implied by the innovation parameters estimated from the job flows over 1 year in Table XIII.

in 1983–1993. Job creation by incumbents fell by 4 percentage points between 1983–1993 and 2003–2013, which the model interprets as a decline in the arrival rate of creative destruction by incumbents. According to Table III, the probability a variety was creatively destroyed by an incumbent firm dropped from 8.9% to 6.2%. Table IV says TFP growth from creative destruction fell from 44 basis points per year over 1983–1993 to 29 basis points from 2003–2013. TFP growth overall fell by 34 basis points over this period. A reduced flow of new varieties also contributed to the decline, while own innovation was stable. Slowing creative destruction played a bigger role in the sharper decline in growth from 1993–2003 to 2003–2003, as shown in Table IV.

So the model's proximate answer to the question "how much does the decline in job creation matter for aggregate TFP growth?" is that entrant innovation and incumbent creative destruction indeed declined. Still, what the aggregate job creation rate misses is the contribution of own innovation. Table IV shows that growth from own innovation by incumbent firms was roughly the same in 2003–2013 compared to the 1983–1993 period.

We reiterate that, because of own innovation, employment growth of a firm can be a misleading indicator of the firm's contribution to aggregate growth. Employment will grow by less in a firm that improves its own product than in another firm that grows due to creative destruction, even when the innovation step size is the same. Table VI illustrates this using the contribution of young versus old firms to job creation (column 1) versus TFP growth (column 2).[19] Take the contribution of new firms (age < 1) with that of old firms (age > 15). As noted by many authors, a large share of job creation is due to new firms—firms with age < 1 account for 31% of all job creation in our calculation.[20] Yet new firms only account for 9% of TFP growth. In contrast, old firms (age > 15) are responsible for more than 50% of TFP growth, but only a third of total job creation. The reason for this discrepancy is that entrant innovation largely takes the form of creative destruction, while a substantial share of innovation by older firms takes the form of own innovation. Again, own innovation entails less job reallocation than when the innovator displaces an existing firm.

As noted, our estimation procedure chose parameter values such that the model exactly matched aggregate TFP growth. Table VII (column 2) shows the fit of the model in 2003–2013 for the data moments we used for the estimation but did not force the model

---

[19]For this calculation, we use innovation rates and step sizes implied by job flows computed over *one* year instead of five years. So a new firm here is defined as a firm created in the last year (not the last five years). See Section 5.3 for details.

[20]Haltiwanger, Jarmin, and Miranda (2013) provided similar evidence on the importance of firms with age < 1 for job creation.

TABLE VII

MODEL FIT, 2003–2013[a]

	Data	Baseline Model	Directed CD by Entrant
Employment share of entrants	15.5%	15.9%	15.3%
Employment growth rate	0.5%	0.5%	0.5%
Job creation rate	32.5%	27.6%	28.8%
Job destruction rate	30.0%	25.2%	26.1%
Share of job creation < 1	36.3%	22.7%	26.5%
SD(log employment)	1.28	1.37	1.21
Exit rate large/small	0.65	0.81	0.62
Employment incumbents/entrants	3.2	1.5	2.3

[a]The table presents the predicted data moments we target in the inference exercise. Column 2 shows the fitted data moments for the baseline model. Column 3 does the same for the model where we allow the quality of products entrants obtain via creative destruction to differ from the quality of the average product. For the exit rate calculation, small firms are defined as those with below-average employment and large firms are those with above-average employment.

to perfectly match.[21] The model comes close on the entrant employment share, the job creation and destruction rates, and the standard deviation of log employment, but understates the share of job creation < 1, the exit rate of small versus large firms, and the average size of incumbent firms versus entrants.[22]

Andrews, Gentzkow, and Shapiro (2017) proposed a local measure of the relationship between parameter estimates and moments. In this spirit, Figure 11 in Appendix C presents the Jacobian matrix of moments with respect to parameter values. The table shows the percentage change in each moment with respect to a 0.1 level increase in each parameter. The results are based on linear extrapolation of the numerical derivative, evaluated at the parameter values we estimate for 2003–2013.

Figure 11 conforms well to our intuition. The entrant employment share is negatively related to own innovation, since more own innovation implies less room for creative destruction by entrants. Overall job creation and destruction rates are also inversely related to own innovation, again because they entail less job reallocation than does creative destruction. Job creation is positively related to the arrival of new varieties, of course.

The share of job creation < 1 is most sensitive to the arrival of new varieties, but is also positively related to the arrival of own innovation. Size dispersion is increasing in the rate of new variety creation by entrants because new varieties are drawn from the entire quality distribution, only with lower average quality. Own innovation and new variety creation provide protection against exit. Entrants are smaller if they enter with new varieties, which tend to be of below-average quality.

Now, the model understates the size of large firms because it generates too little heterogeneity in the number of varieties. A higher rate of creative destruction by incumbents would increase the average number of varieties among incumbent firms. This would increase the average size of incumbents relative to entrants, and would also make the exit rate fall more steeply with firm size. The rate of creative destruction is constrained, however, by the need to fit the share of job creation driven by firms with small changes in employment. A higher rate of creative destruction by incumbents would lower the share

---

[21]We also set the minimum firm size near 1, as in the data.

[22]The model fits slightly better in 1983–1993 and 1993–2003 (see Table XVIII in Appendix C).

TABLE VIII

SIMULATED DISTRIBUTION OF PRODUCTS PER FIRM[a]

	1983–1993	1993–2003	2003–2013
1 Product	91.9%	84.1%	89.3%
2 Products	7.1%	12.3%	9.1%
3 Products	0.8%	2.6%	1.3%
>3 Products	0.1%	0.9%	0.2%
Average Products per Firm	1.09	1.20	1.12

[a]The table presents the number of products per firm implied by the innovation parameters in Table III.

of job creation among firms with small employment growth, which is already too low in the model relative to the data.

We need to increase the size and number of products of incumbent firms without increasing the frequency of creative destruction by incumbents. Remember we assume creative destruction by incumbents and entrants is *undirected*—that is, uniform across all existing varieties—so the average quality of a variety improved upon by an incumbent is the same as that of an entrant. We now relax this assumption. Specifically, we suppose entrants innovate over lower quality products, where $\rho_e \leq 1$ denotes the average (employment-weighted) quality of products creatively destroyed by entrants relative to the quality of the average product.

The third column in Table VII shows the fit of the model when we introduce $\rho_e \leq 1$.[23] We estimate $\rho_e = 0.59$, which lowers the average quality and size of entrants relative to incumbents. Allowing $\rho_e \leq 1$ also increases the arrival rate of new varieties by incumbents to 3% from 0% previously. This decreases the exit rate of large firms, who now have more varieties. Despite having more varieties, these firms are no larger because their new varieties are of low quality.[24] The share of job destruction due to small changes in employment is 27%, up from 23% in the baseline model. The model with $\rho_e \leq 1$ still under-predicts the share of small job changes (in the data, the share is 36%), but the fit is now better. The sources of growth do not materially change when we introduce ρ_e. For example, innovation by incumbents is responsible for 77% of growth in the model with $\rho_e \leq 1$, compared to 80% in our baseline model.

We did not target the share of large employment changes in job destruction to arrive at our baseline innovation parameter estimates in Table III. For 2003–2013, our model predicts that the share of job destruction < -1 is 80.5%. This is higher than in the data, which is 72.5%. On this dimension, our baseline estimates would seem to be overstating the importance of creative destruction.

Finally, creative destruction by incumbents will result in some firms owning a large number of products. Since we find a limited role for creative destruction, most firms own a small number of products in our simulations. Table VIII shows that the innovation parameters we estimate imply that about 90% of firms will have a single product, and less than 1% will have more than 3 products. The average number of products per firm, shown in the last row of Table VIII, is around 1.1 to 1.2. We cannot verify this empirically because we do not know of a data set which contains the number of products and services sold that covers all firms in the private sector. Perhaps such data will become available in the future.

---

[23]We target the same 10 data moments.

[24]The arrival rate (average quality) of new varieties by incumbents is $\kappa_i = 0.03$ ($s_\kappa = 0.49$).

TABLE IX

SOURCES OF GROWTH USING DIFFERENT σ AND JC CUTOFF[a]

	Baseline			
	$\sigma = 4$	$\sigma = 3$	$\sigma = 5$	$JC < 2/3$
Creative destruction	21.9%	28.6%	20.8%	22.2%
New varieties	4.2%	4.2%	6.1%	4.3%
Own-variety improvements	74.0%	67.3%	73.1%	73.5%

[a]The estimated growth contributions are for 2003–2013. Apart from the last column, the cutoff for "small" job creation is equal to $JC < 1$.

5. ROBUSTNESS

We now examine the robustness of our inference exercise. First, we explore how our estimates would change if we entertain a different degree of substitutability between products. Second, we use the number of establishments as a proxy for the number of products of the firm. Third, we use alternative measures of job flows to infer the sources of growth (e.g., excluding M&A activity). Fourth, we measure the sources of innovation in the fast growing information and communication technology sector.

5.1. *Elasticity of Substitution and "Small" Job Creation*

We set the elasticity of substitution across products in the CES aggregator to $\sigma = 4$ based on outside evidence for our baseline estimates. In Table IX, we entertain different values for σ. The share of growth we infer from creative destruction is almost 29% with $\sigma = 3$ and less than 21% with $\sigma = 5$, versus 22% in the baseline. With a lower σ, more creative destruction is needed to generate a realistic amount of job reallocation. The shares of growth from new varieties and own innovation generally move in the same direction as σ.

Separately, we gauge robustness to a more modest definition of "small" job creation. For our baseline estimates, we set the "small" threshold at 1, or less than a tripling of employment. We re-estimated parameter values with a threshold of 2/3, which corresponds to less than a doubling of employment. The last column of Table IX presents the resulting growth decomposition. Compare the first and last column (the first uses $JC < 1$ and the last uses $JC < 2/3$; both use $\sigma = 4$). There is little effect: the growth contributions change by less than 0.5%.

5.2. *Establishments as Proxies for Products*

Ideally, we could directly observe firms discontinuing products due to creative destruction. The LBD contains no data on products produced, but it does include the number of establishments in a firm. Under this proxy, if creative destruction is an important reason for job reallocation, then changes in the number of establishments should have a large effect on aggregate job flows.

For comparison, we reproduce our baseline estimates of job flows in the first column in Table X. The second column then calculates job flows after dropping from the sample new and exiting establishments. The job creation rate falls from 32.5% in the full sample to only 10%. The job destruction rate falls from 30% to about 10%. The share of "small" changes jumps from 36% to 78% of job creation and from 27.5% to almost 80% of job

TABLE X

ALTERNATIVE MEASURES OF JOB FLOWS IN THE LBD (2003–2013)[a]

	Full Sample	Drop New & Exiting Estab.	Only Services	Exclude M&A[b]	Exclude M&A[c]
Job creation rate	32.5%	10.1%	34.2%	25.6%	26.2%
Employment share of entrants	15.5%	–	17.3%	14.9%	12.3%
Share of job creation < 1	36.3%	78.1%	35.7%	39.9%	39.1%
Job destruction rate	30.0%	9.6%	29.2%	23.1%	23.1%

[a]*Source*: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. The share of job creation <1 is the fraction of job creation at firms that expanded by less than a factor of 3 in employment. Column 2 removes from the sample establishments created between 2003 and 2008 (and 2008 and 2013) and establishments that died between the two 5-year periods. Column 3 excludes manufacturing establishments. Columns 4 and 5 remove the direct effect of changes in plant ownership. Column 5 attributes employment change among plants that switch ownership within each period to the acquiring firm.

[b]Exclude plants that switch ownership between 2003 and 2008 (and between 2008 and 2013).

[c]Exclude direct effect of changes in plant ownership between 2003 and 2008 (and between 2008 and 2013), but attribute employment change in such plants to *acquiring* firm.

destruction. Thus, the creative destruction we infer does, indeed, show up largely in plant exit and entry.

We also examine the correlation of job creation and job destruction across industries. If job reallocation is mostly driven by creative destruction, then industries should have simultaneously high (or low) rates of job creation and job destruction. To examine this, we regress the job destruction rate in a 6-digit NAICS from 2003 to 2008 on the job creation rate in the same sector. This regression yields a coefficient of 0.629 (s.e. = 0.016, $R^2 = 0.68$).[25]

Next, recall the model's implication that the average number of products per firm is constant. Figure 10 plots the average number of establishments per firm in the LBD and in the Census of Manufacturing. There is an upward trend in the average number of establishments in the LBD, but the magnitude is very small. There is no trend in average number of establishments in manufacturing.

The Census of Manufacturing also provides detailed information (at the 10-digit level) on the products made by each establishment. After aggregating this information to the firm level, Figure 10 plots the average number of *products* per firm in manufacturing. It is almost three times larger than the number of establishments. This suggests caution in using establishments as a proxy for products in the LBD. Nevertheless, there is no trend in the average number of products in the manufacturing census.[26]

5.3. *Alternative Measures of Job Flows*

In our closed economy model, creative destruction can only come from domestic firms. This is obviously not true in traded sectors where a domestic firm can innovate on a product owned by a foreign firm, and vice versa. So we calculate job flows only in the nontraded service sectors in column 3 of Table X. Job flows outside of manufacturing are similar to flows in the full sample that includes traded sectors.

[25]We obtain similar results in other time periods. For example, for 2008–2013, the coefficient from regressing the job destruction rate on the job creation rate is 0.583 (s.e. = 0.016, $R^2 = 0.52$).

[26]See Bernard, Redding, and Schott (2010) for evidence that individual manufacturing plants frequently add and subtract products, even at the 5-digit SIC level.
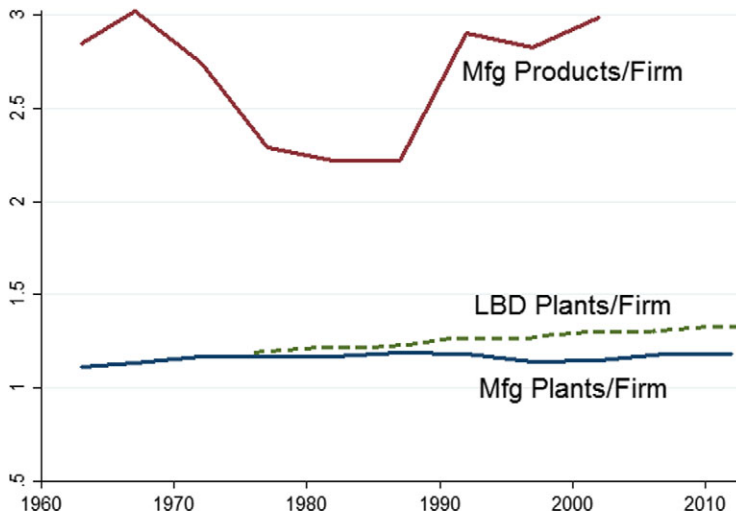
FIGURE 10.—Plants and products per firm. *Note*: Figure plots average number of establishments per firm in the LBD and average number of establishments and 10-digit products per firm in the U.S. Manufacturing Census.

Recall that we include mergers and acquisitions in our baseline statistics on job creation and job destruction. The last two columns in Table X check the effect of dropping establishments who change ownership. Column 4 entirely drops establishments that change ownership between 2008 and 2013. Column 5 only omits the direct effect of establishments that switch ownership; it does add the employment change in such establishments to the employment change of the acquiring firm. The resulting job creation and destruction rates drop by 6 to 7 percentage points. The employment share of entrants and the share of small employment changes in job creation are about the same in the new samples.

Another alternative is to use changes in firm output instead of employment. In the model, the change in employment is the same as the change in output, but this is not true in the data. Innovation may take the form of labor saving technologies. We do not have output data in the LBD, but we do in the Census of Manufacturing. We therefore calculate the growth in employment implied by growth in the firm's value-added.[27] The hypothetical job flows with this imputation are shown in the second column of Table XI. For comparison, we show the job flows computed from firm employment in the same manufacturing sample (in column 1). Imputing employment from firm value-added increases job creation and destruction rates by about 2 percentage points. Intuitively, firm output grows by more than firm employment among growing firms, and declines by more than employment among shrinking firms. But note that the share of job creation < 1 is about the same, so adjusting employment with output data has the same effect on large changes as on small changes in employment.

Even output may be a biased measure of innovation in the presence of adjustment frictions and changes in the regulation, taxes, or factor costs. The last column in Table XI imputes the job flows implied by the change in firm *productivity* ("TFPQ") in the Census

---

[27]We impute employment as the product of the firm's share of industry's value-added and total employment in the industry.

TABLE XI

Job Flows in Manufacturing (2002–2012)[a]

	Employment	Imputed From VA	Imputed From TFPQ
Job creation rate	14.9%	16.8%	30.1%
Employment share of entrants	5.6%	5.4%	3.8%
Share of job creation < 1	50.1%	53.1%	31.9%
Job destruction rate	42.7%	44.7%	57.9%

[a]*Source*: U.S. Manufacturing Census. The second column imputes firm employment as the product of industry employment and the firm's share of the industry's value-added. The third column imputes employment of firm i as the product of industry employment and $TFPQ_i^{\sigma-1}/\sum_j TFPQ_j^{\sigma-1}$.

of Manufacturing.[28] Intuitively, these hypothetical job flows abstract from all other forces behind firm employment other than firm productivity. The hypothetical job creation rate is roughly twice as large and the share of job creation from small changes is about 20 percentage points *lower*. When using productivity data to impute job flows, we find that a much smaller share of overall job creation is due to small changes.

Table XII shows what these alternate job flow statistics imply for the sources of growth. The second column estimates the model using the data moments from the sample where we drop establishments that undergo ownership changes. This exercise assumes the reallocation associated with M&A activity has no effect on TFP growth. Since extreme job reallocations diminish (Table X), we infer a lower rate of creative destruction and a higher rate of own innovation.[29]

The next two columns in Table XII show the effect of using output and productivity on the sources of growth. To arrive at these numbers, we impute job flows in the LBD from output and productivity data in manufacturing.[30] The sources of growth implied by the

TABLE XII

Sources of Growth Using Alternative Measures of Job Flows[a]

	Baseline	No Mergers	Imputed From VA	Imputed From TFPQ
Creative destruction	21.9%	18.6%	21.3%	49.2%
New varieties	4.2%	3.8%	4.6%	9.8%
Own-variety improvements	74.0%	77.5%	74.2%	41.0%

[a]The column "No Mergers" targets data excluding M&A job flows from 2008 to 2013 in the LBD. "Imputed from VA" targets job flows in the manufacturing census computed by imputing firm employment from firm value-added. We add the difference between job flows based on employment and the flows based on value-added to the job flows in the LBD sample. "Imputed from TFPQ" does the same for the job flows based on employment imputed from firm TFPQ in manufacturing.

[28]Following Hsieh and Klenow (2009), we impute firm productivity as $TFPQ = Y^{1/(\sigma-1)} \cdot Y/L$ where Y and L denote firm value-added and *measured* employment. Employment of firm i implied by firm productivity is then given by the product of industry employment and $TFPQ_i^{\sigma-1}/\sum_j TFPQ_j^{\sigma-1}$.

[29]We include M&A activity in our baseline estimates for two reasons. First, job reallocation associated with such activity may be a byproduct of innovation. Second, as we find a smaller role for creative destruction than the existing literature, this is a conservative assumption.

[30]We make three adjustments to the LBD data to mimic the manufacturing data. First, we adjust the job creation and destruction rates in the LBD by the difference between the job flows in column 1 of Table XI and those in columns 2 and 3 of the same table. This adjustment increases the job creation and destruction

TABLE XIII

JOB FLOWS OVER 1, 5, AND 10 YEARS IN THE LBD 2003–2013[a]

	Baseline		
	5 Years	1 Year	10 Years
Job creation rate	32.5%	13.1%	46.4%
Employment share of entrants	15.5%	3.4%	26.3%
Share of job creation by incumbents age < 5	n.a.	16.5%	n.a.
Share of job creation < 1	36.3%	54.0%	28.2%
Job destruction rate	30.0%	12.6%	41.3%
Share of job destruction by incumbents age < 5	n.a.	13.0%	n.a.

[a]*Source*: U.S. Census LBD data on firms in the nonfarm business sector.

output data (column 3) are about the same as when we target employment data (column 1). The reason can be seen by comparing the raw data in columns 1 and 2 in Table XI. The share of small changes does not change, so the relative importance of creative destruction versus own innovation also does not change.

The story is quite different when we target job flows implied by productivity growth (column 4 in Table XII). Here, creative destruction accounts for almost 50% of growth, and own innovation for only 40%. Creative destruction is more important now because the share of overall job creation from small changes is lower. Table XI shows that the share of job creation < 1 is 20 percentage points lower when we use productivity growth (column 3) compared to the calculations based on employment data alone (column 1). Because own innovation generates small changes in employment, we infer a smaller contribution from own innovation and a larger contribution from creative destruction.

We next present job flows over *one* year in Table XIII to compare with our baseline 5-year estimates. For the 1-year estimates, an entrant is defined as a firm created in the last year (not the last five years), and job creation between years t and $t+1$ is the sum of the employment of entrants in year $t+1$ and the employment change among incumbents with growing employment between these two years. The job destruction rate over one year is defined analogously.

For 1-year data moments, we classify firms between ages 1 and 5 as incumbents. As a consequence, the employment share of entrants is lower (3.4%) than in the 5-year baseline (15.5%). The annual job creation rate from 2003 to 2013 in Table XIII is 13.1%. The equivalent number implied by the job creation rate over five years was 6.5%.[31] With 1-year flows, job creation and destruction includes flows *within* each 5-year period. For example, Table XIII indicates that incumbent firms with age < 5 account for 16.5% of job creation and 13% of job destruction with 1-year flows. Job creation computed over five years counted only *net* job creation by incumbents age < 5 (which we attributed to entrants).

rates in the LBD by about 2 to 3 percentage points (for the output adjustment) and 15 percentage points (for the productivity adjustment). Second, we multiply the share of job creation due to small changes by the ratio of the share of job creation < 1 calculated from manufacturing employment and the share computed from output and productivity data. The second adjustment has no effect when we use output data, but lowers the share of small employment changes by almost 50% when we use productivity data. Third, we adjust size by age and exit by size in the LBD by the ratio of these two statistics in the Census of Manufacturing computed using employment versus output (or productivity).

[31]The job creation rate over five years displayed in Figure 4 from 2003–2013 was 32.5%.

TABLE XIV

Sources of Growth Implied by Job Flows Over 1, 5, and 10 Years[a]

	Baseline		
	5 Years	1 Year	10 Years
Creative destruction	21.9%	10.9%	27.2%
New varieties	4.2%	5.4%	7.7%
Own-variety improvements	74.0%	83.4%	65.1%

[a]Entries are the share of aggregate TFP growth due to each source of innovation implied by the job flows over 5 years, 1 year, and 10 years shown in Table XIII. Growth contributions are estimated for 2003–2013.

However, remember the *distribution* of job changes is key for distinguishing own innovation and creative destruction. The share of small employment changes in job creation is now *higher* compared to the share when we calculate employment changes over five years.[32] Our inference exercise will therefore indicate that the share of growth driven by own innovation versus creative destruction will be even higher with 1-year flows. Table XIV shows the share of growth due to each source of innovation implied by job flows over one year. The contribution of creative destruction to aggregate growth from 2003 to 2013 is only 11% with 1-year flows, versus 22% with 5-year flows. The share of growth from own innovation is correspondingly higher at 83% (versus the baseline 74%).

One-year flows also suggest a decline in job creation between 1983–1993 and 2003–2013. The difference is that the decline is due to the diminishing employment share of entrants and smaller job flows by young firms (incumbents age < 5).[33] When we use job flows over one year to calculate the innovation parameters, creative destruction plays a larger role and own innovation a smaller role in 1983–1993 than in 2003–2013. So whether we calculate annual or 5-year job flows, own innovation has become more important.

We can go in the opposite direction and define entrants based on 10-year job flows rather than 1-year or 5-year flows. Tables XIII and XIV present job flow rates and growth contributions with 10-year flows. The employment share of entrants is naturally larger (46% compared to the 33% baseline with 5-year flows), and the share of "small" job creation is understandably smaller (28% versus the baseline 36%). As the result, we infer a higher share of growth from entrants and creative destruction (27% versus the baseline 22%) and a lower share of growth from own innovation (65% versus the baseline 74%). But our conclusion that the bulk of growth comes from own innovation by incumbents remains.

5.4. *Innovation in the Information Technology Sector*

We end with a case study of the fast-growing information and communication technology (ICT) sector. Table XV (row 1) shows that TFP growth in the ICT sector increased from 11.8% in 1983–1993 to 26.4% per *year* in 1993–2003. The job creation rate in ICT increased by almost 20 percentage points to 53% and the job destruction rate increased by 12 percentage points to 39%.[34] But the distribution of job flows also changes. The ac-

---

[32]Table XIII shows that, for 2003–2013, the share of job creation < 1 is 54% when calculated over one year, compared to 36% when job flows were computed over five years.

[33]When we measured job flows over five years, the decline in job creation was entirely due to the fall in the entrant employment share (firms entering in the last five years).

[34]We revert back to measuring job flows over five years. We follow Shackelford and Janowski (2016) in classifying ICT industries—see http://klenow.com/ICT.xlsx for the precise list.

TABLE XV

INFORMATION TECHNOLOGY & COMMUNICATIONS INDUSTRY[a]

	1983–1993	1993–2003	2003–2013
TFP growth rate	11.8%	26.4%	10.0%
Job creation rate	35.4%	53.4%	38.9%
Employment share of entrants	15.6%	18.4%	13.6%
Share of job creation < 1	34.9%	33.2%	36.9%
Job destruction rate	27.2%	39.1%	34.3%

[a]*Source*: TFP growth in row 1 from BLS. Rows 2–5 from U.S. Census Longitudinal Business Database (LBD) on firms in the ICT sector. Job flows computed over 5-year periods.

celeration of TFP growth in the ICT sector in 1993–2003 was associated with a *decline* in the share of job creation < 1 (from 35% to 32%). This fact suggests that the growth acceleration was driven by an increase in the rate of creative destruction.

Table XVI says the contribution of creative destruction to TFP growth in the ICT sector increased from 17% in 1983–1993 to 36% in 1993–2003. So the 15 percentage point increase in TFP growth growth from 1983–1993 to 1993–2003 was mostly due to creative destruction. When TFP growth came back down in 2003–2013 (to "only" 10%), the contribution of creative destruction fell back to 15%.[35] We caution that this inference is based only on employment data. Employment growth may understate the extent of productivity growth, particularly in fast growing sectors like ICT. In addition, our inference exercise assumes that each 10-year period is a steady state, which may be a particularly bad approximation for sectors with accelerating then decelerating growth rates like ICT.

6. CONCLUSION

How much innovation takes the form of creative destruction versus new varieties versus firms improving their own products? How much occurs through entrants versus incumbents? We try to infer the sources of innovation from the employment dynamics of U.S. firms in the nonfarm private sector from 1983 to 2013. We conclude that creative destruction is vital for understanding job destruction and accounts for around one-fourth of growth. Own-product quality improvements by incumbents appear to be the biggest source of growth. Net variety growth contributes much less than quality improvements do.

TABLE XVI

SOURCES OF GROWTH IN INFORMATION TECHNOLOGY[a]

	1983–1993	1993–2003	2003–2013
Creative destruction	17.0%	35.9%	15.3%
New varieties	6.0%	3.1%	6.3%
Own-variety improvements	77.0%	61.0%	78.4%

[a]The table shows the percentage of growth due to each source of innovation in the ICT sector implied by the job flows shown in Table XV.

[35]For fitting the ICT facts in our simulations, we allow the scale parameter of the Pareto quality draws to be greater than 1. Without this, we can only fit the extraordinary TFP growth in ICT with a dispersion parameter θ that implies a quality distribution with infinite variance.

Our findings are relevant for innovation policy. According to Atkeson and Burstein (forthcoming), the consumption-equivalent welfare gain from devoting about 1% more of GDP to research every year is between 0.17 and 0.73 percent in a model calibrated similarly to ours (but with endogenous research investments). They estimated the gains would be larger—0.26 to 2.01 percent—if creative destruction played no role in U.S. growth. Creative destruction is a force raising the private return relative to the social return to research, diminishing the gains from promoting research.

Creative destruction ties into political economy theories in which incumbents block entry and hinder growth and development, such as Krusell and Rios-Rull (1996), Parente and Prescott (2002), and Acemoglu and Robinson (2012). And creative destruction underscores that some growth brings employment dislocation as a byproduct.

It would be interesting to extend our analysis to individual sectors, other time periods, and other countries. Retail trade experienced a big-box revolution in the United States led by Walmart's expansion. Online retailing has made inroads at the expense of brick-and-mortar stores (Dolfen et al. (2018)). In Chinese manufacturing, private enterprises have entered and expanded while state-owned enterprises have closed (Hsieh and Klenow (2009)). In India, incumbents do not expand as much as in the United States (Hsieh and Klenow (2014)) and therefore contribute less to growth. We analyzed non-manufacturing separately, and reached similar conclusions. For the ICT sector, we found that creative destruction played a larger role during its rapid growth phase from 1993 to 2003.

Our accounting is silent on how the types of innovation interact. In Klette and Kortum (2004), more entrant creative destruction discourages R&D by incumbents. Alternatively, as stressed by Aghion, Harris, Howitt, and Vickers (2001), a greater threat of competition from entrants could stimulate incumbents to "escape from competition" by improving their own products. Creative destruction and own innovation could be strategic complements, rather than substitutes. We hope future research on endogenous innovation will capture these interactions, enabling one to do counterfactuals and optimal policy calculations that we cannot.

Our conclusions are tentative in part because they are model-dependent. We followed the literature in several ways that might not be innocuous for our inference. We assumed that step sizes and spillovers are just as big for incumbents as for entrants. Young firms might instead generate more radical innovations with larger knowledge spillovers than old firms do—Akcigit and Kerr (2018) provided evidence for this hypothesis from patent citations.

We assumed all fluctuations in employment across firms (within industries) are a byproduct of innovation. By looking within detailed industries, we hoped to mitigate the influence of non-homothetic preferences and demographic shifts. We demonstrated robustness to looking at value-added and productivity rather than employment, at least within manufacturing where we could see revenue and capital. If exit or downsizing occurs for reasons other than innovation, then our methodology might overstate the contribution of creative destruction.[36]

We assumed no frictions in allocating employment across firms. In reality, the market share of young firms could be suppressed by adjustment costs, financing frictions, and uncertainty. On top of adjustment costs for capital and labor, firms may take a

---

[36]Smith, Yagan, Zidar, and Zwick (2017) used administrative tax records to document that firm profits dive after the death of under-65 owner/managers.

while to build up a customer base, as in Gourio and Rudanko (2014) and Foster, Halti-wanger, and Syverson (2016). Irreversibilities could combine with uncertainty about the firm's quality to keep young firms small, as in the Jovanovic (1982) model. We defined young firms as those younger than five or ten years, but these dynamics could play out for longer. Meanwhile, markups could vary across varieties and firms. All of these would create a more complicated mapping from firm employment growth to firm innovation.

APPENDIX A: SIMULATION ALGORITHM

1. Specify an initial guess for the distribution of quality across varieties.
2. Simulate life paths for a large number of entering firms.
3. Each entrant has one initial variety, captured from an incumbent or newly created. In every period of its lifetime, it faces a probability of each type of innovation per variety it owns, as in Table I. A firm's life ends when it loses all of its varieties to other firms or when 40 periods have passed.
4. Based on the population of simulated firms, calculate 10 moments:
    (a) TFP growth rate.
    (b) Standard deviation of log firm employment.
    (c) Employment share of entrants.
    (d) Job creation rate.
    (e) Job destruction rate.
    (f) Share of job creation where employment growth ≤ 1.
    (g) Minimum firm employment (1).
    (h) Exit rate for small firms (firms with below-average employment).
    (i) Exit rate for large firms (firms with above-average employment).
    (j) Average employment for incumbents relative to entrants.
5. Repeat steps 1 to 4 until all moments and the joint distribution of quality and variety across firms converge. In each iteration, take the quality distribution across varieties from step 4 as the starting point and update the overhead labor requirement to target minimum employment of 1.
6. Repeat steps 1 to 5, searching for parameter values to (i) exactly match TFP growth in the data; (ii) set minimum firm employment to 1; and (iii) minimize the mean squared percent distance between the simulated and empirical moments for the remaining statistics in step 4.

APPENDIX B: ESTIMATES WITH STANDARD ERRORS

We calculate the standard error in the parameters of the model in the following steps. First, we draw a thousand samples with replacement of all firm ids in the LBD data from 1983 to 2013. We then calculate the moments in Section 3 for each sample. This gives us the mean and the variance-covariance matrix of the 10 data moments, which we denote Σ.

We then calculate the variance-covariance matrix of the parameters V. We use the SMM formula from Gourieroux and Monfort (1996). This captures both the variance in data moments and the variance induced by simulation of small samples:

$$V = (J'WJ)^{-1}J'W(\Sigma + \widetilde{\Sigma})WJ(J'WJ)^{-1},$$

TABLE XVII

BOOTSTRAPPED PARAMETER ESTIMATES, 2003–2013

	Coefficient	Standard Error
Creative destruction by incumbents δ_i	0.283	0.0091
New varieties by entrants κ_e	0.030	0.0026
Own-variety improvement by incumbents λ_i	0.779	0.0036
Relative quality of new varieties s_κ	0.30	0.0252

where J is the Jacobian matrix (in levels), W is the weighting matrix, and $\widetilde{\Sigma}$ is the variance-covariance matrix of the model-simulated moments (for a given set of parameters). The Jacobian J is estimated as the numerical gradient of the model-simulated moments with respect to small changes in the parameters. J is only defined for parameters with an interior solution. Parameters with corner solutions (δ_e and κ_i) and exactly-identified parameters (θ, δ_o, and ψ) are assumed to have a standard error of zero. The moment-weighting matrix W is a diagonal matrix with entries equal to 1 for all moments except for exit by size and size by age, whose corresponding entries are 0.5. $\widetilde{\Sigma}$ is estimated by Monte Carlo simulation, calculating the variance-covariance matrix after 200 model simulations with fixed parameters.

Table XVII shows the estimated coefficients and standard errors for the parameters with interior solutions. As can be seen, the standard errors of the estimated parameters are extremely small.

APPENDIX C: MODEL FIT AND SENSITIVITY

Table XVIII shows the fit of the model for 1983–1993 and 1993–2003. For ease of comparison, columns 1 and 3 replicate the data for the corresponding 10-year period. And Figure 11 presents the Jacobian matrix with local derivatives of the data moments with respect to our estimated parameter values for 2003–2013. The table shows the percentage change in each moment with respect to a 0.1 level increase in each parameter. Figure 12 shows the model distribution of job creation and destruction, which can be compared to its empirical counterpart in Figure 6.

See Garcia-Macia, Hsieh, and Klenow (2019) for replication materials.

TABLE XVIII

MODEL FIT, 1983–1993 AND 1993–2003

	1983–1993		1993–2003	
	Data	Model	Data	Model
Entrant employment share	23.5%	22.8%	18.3%	17.5%
Employment growth rate	2.4%	2.3%	1.6%	1.7%
Job creation rate	43.9%	40.6%	41.5%	40.2%
Job destruction rate	32.0%	29.1%	33.2%	31.9%
Share of job creation < 1	31.7%	21.3%	36.6%	25.0%
SD(log employment)	1.25	1.38	1.27	1.28
Exit rate large/small	0.76	0.85	0.74	0.68
Employment incumbents/entrants	2.8	1.8	2.8	2.1

FIGURE 11.—Jacobian matrix, 2003–2013. *Note*: The figure shows the percentage change in each moment with respect to a 0.1 level increase in each parameter. The results are based on linear extrapolation of the numerical derivative, evaluated at the 2003–2013 estimated parameters.
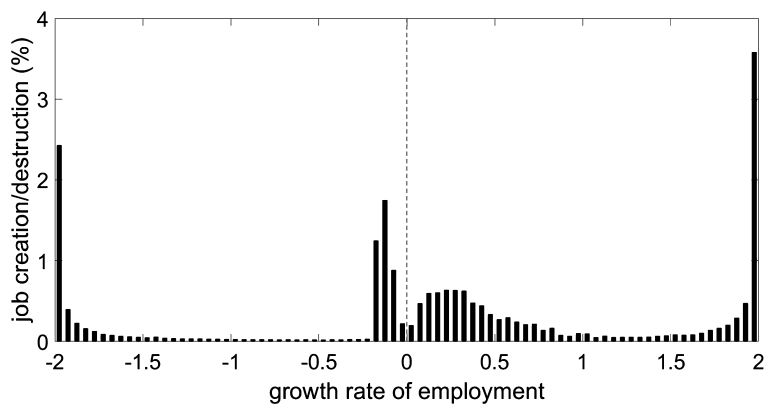

FIGURE 12.—Job creation and destruction, model, 2003–2013. *Note*: The figure is based on the 2003–2013 model simulation. Employment growth for a firm is defined as the change in employment divided by average employment at the firm at the beginning and end of each period. The vertical axis gives the share of total job creation (destruction) associated with firms at each given level of employment growth.

REFERENCES

ACEMOGLU, D. (2003): "Labor- and Capital-Augmenting Technical Change," *Journal of the European Economic Association*, 1 (1), 1–37. [1507]

——— (2011): *Introduction to Modern Economic Growth*. Princeton University Press. [1507]

ACEMOGLU, D., AND J. ROBINSON (2012): *Why Nations Fail: Origins of Power, Poverty and Prosperity*. Crown Publishers (Random House). [1507,1536]

ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): "Innovation, Reallocation, and Growth," *American Economic Review*, 108 (11), 3450–3491. [1509,1510]

AGHION, P., AND P. HOWITT (1992): "A Model of Growth Through Creative Destruction," *Econometrica*, 60 (2), 323–351. [1507,1510]

———— (2009): *The Economics of Growth*. Massachusetts Institute of Technology (MIT) Press. [1507]

AGHION, P., U. AKCIGIT, AND P. HOWITT (2014): "What Do We Learn From Schumpeterian Growth Theory?" *Handbook of Economic Growth*, 2B, 515–563. [1507,1524]

AGHION, P., C. HARRIS, P. HOWITT, AND J. VICKERS (2001): "Competition, Imitation and Growth With Step-by-Step Innovation," *Review of Economic Studies*, 68 (3), 467–492. [1536]

AKCIGIT, U., AND W. R. KERR (2018): "Growth Through Heterogeneous Innovations," *Journal of Political Economy*, 126 (4), 1374–1443. [1507,1508,1536]

ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2017): "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, 132 (4), 1553–1592. [1527]

ATKESON, A., AND A. BURSTEIN (forthcoming): "Aggregate Implications of Innovation Policy," *Journal of Political Economy*. [1508,1518,1536]

BAILY, M. N., C. HULTEN, AND D. CAMPBELL (1992): "Productivity Dynamics in Manufacturing Plants," *Brookings Papers: Microeconomics*, 4, 187–267. [1509]

BERNARD, A. B., S. J. REDDING, AND P. K. SCHOTT (2010): "Multi-Product Firms and Product Switching," *American Economic Review*, 100 (1), 70–97. [1530]

BRODA, C., AND D. E. WEINSTEIN (2006): "Globalization and the Gains From Variety," *Quarterly Journal of Economics*, 121 (2), 541–585. [1508]

———— (2010): "Product Creation and Destruction: Evidence and Implications," *American Economic Review*, 100 (3), 691–723. [1508]

CHRISTENSEN, C. (1997): *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business Review Press. [1508]

DAVIS, S. J., J. C. HALTIWANGER, AND S. SCHUH (1996): *Job Creation and Destruction*. MIT Press Books. [1508, 1514,1521]

DECKER, R., J. HALTIWANGER, R. JARMIN, AND J. MIRANDA (2014): "The Role of Entrepreneurship in U.S. Job Creation and Economic Dynamism," *The Journal of Economic Perspectives*, 28 (3), 3–24. [1508,1519]

DOLFEN, P., L. EINAV, P. J. KLENOW, B. KLOPACK, J. D. LEVIN, L. LEVIN, AND W. BEST (2018): "Assessing the Gains From E-Commerce." [1536]

FEENSTRA, R. C. (1994): "New Product Varieties and the Measurement of International Prices," *The American Economic Review*, 157–177. [1508]

FOSTER, L., J. HALTIWANGER, AND C. SYVERSON (2016): "The Slow Growth of New Plants: Learning About Demand?" *Economica*, 83 (329), 91–129. [1537]

FOSTER, L., J. C. HALTIWANGER, AND C. KRIZAN (2001): "Aggregate Productivity Growth. Lessons From Microeconomic Evidence," *New Developments in Productivity Analysis*, 303–372. [1509]

GARCIA-MACIA, D., C.-T HSIEH P. J KLENOW (2019): "Supplement to 'How Destructive Is Innovation?'," *Econometrica Supplemental Material*, 87, https://doi.org/10.3982/ECTA14930. [1538]

GORDON, R. J. (2007): *The Measurement of Durable Goods Prices*. University of Chicago Press. [1508]

GOURIEROUX, C., AND A. MONFORT (1996): *Simulation-Based Econometric Methods*. Oxford University Press. [1537]

GOURIO, F., AND L. RUDANKO (2014): "Customer Capital," *Review of Economic Studies*, 81 (3), 1102–1136. [1537]

GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87 (3), 342–362. [1508]

GROSSMAN, G. M., AND E. HELPMAN (1991): "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 58 (1), 43–61. [1507,1510]

HALTIWANGER, J., R. JARMIN, AND J. MIRANDA (2013): "Who Creates Jobs? Small versus Large versus Young," *Review of Economics and Statistics*, 95 (2), 347–361. [1518,1526]

HOTTMAN, C. J., S. J. REDDING, AND D. E. WEINSTEIN (2016): "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, 131 (3), 1291–1364. [1508]

HOWITT, P. (1999): "Steady Endogenous Growth With Population and R & D Inputs Growing," *Journal of Political Economy*, 107 (4), 715–730. [1507]

HSIEH, C.-T., AND P. J. KLENOW (2009): "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124 (4), 1403–1448. [1532,1536]

——— (2014): "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics*, 129 (3), 1035–1084. [1521,1536]

JONES, C. I. (2016): "Life and Growth," *Journal of Political Economy*, 124, 539–578. [1507]

JOVANOVIC, B. (1982): "Selection and the Evolution of Industry," *Econometrica*, 50 (3), 649–670. [1537]

KLETTE, T. J., AND S. KORTUM (2004): "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112 (5), 986–1018. [1507,1509,1510,1513,1536]

KORTUM, S. S. (1997): "Research, Patenting, and Technological Change," *Econometrica*, 65 (6), 1389–1419. [1510]

KRUSELL, P. (1998): "How Is R&D Distributed Across Firms? A Theoretical Analysis". [1507]

KRUSELL, P., AND J.-V. RIOS-RULL (1996): "Vested Interests in a Positive Theory of Stagnation and Growth," *Review of Economic Studies*, 63 (2), 301–329. [1536]

LENTZ, R., AND D. T. MORTENSEN (2008): "An Empirical Model of Growth Through Product Innovation," *Econometrica*, 76 (6), 1317–1373. [1509]

LUCAS, R. E. J., AND B. MOLL (2014): "Knowledge Growth and the Allocation of Time," *Journal of Political Economy*, 122 (1). [1507]

PARENTE, S. L., AND E. C. PRESCOTT (2002): *Barriers to Riches*. MIT press. [1536]

PETERS, M. (2018): "Heterogeneous Mark-ups, Growth and Endogenous Misallocation." [1509]

ROMER, P. M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98 (5), S71–S102. [1507,1524]

SCHUMPETER, J. A. (1939): *Business Cycles*. Cambridge University Press. [1507]

SHACKELFORD, B., AND J. JANOWSKI (2016): "Information and Communications Technology Industries Account for $133 Billion of Business R&D Performance in the United States in 2013," Technical Paper 16-309, National Center for Science and Engineering Statistics. [1534]

SMITH, M., D. YAGAN, O. ZIDAR, AND E. ZWICK (2017): "Capitalists in the Twenty-First Century," Working Paper, UC Berkeley and University of Chicago. [1536]

STOKEY, N. L. (1988): "Learning by Doing and the Introduction of New Goods," *Journal of Political Economy*, 96 (4), 701–717. [1507]

YOUNG, A. (1998): "Growth Without Scale Effects," *Journal of Political Economy*, 106 (1), 41–63. [1507]