

How Destructive is Innovation?

Daniel Garcia-Macia

Chang-Tai Hsieh

International Monetary Fund

University of Chicago and NBER

Peter J. Klenow *

Stanford University and NBER

May 1, 2018

Abstract

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.

*We thank Ufuk Akcigit, Andy Atkeson, Michael Peters, Paul Romer, Jaume Ventura, Gianluca Violante, and four anonymous referees for very helpful comments. For financial support, Hsieh is grateful to the Polsky Center at Chicago and Klenow to the Stanford Institute for Economic Policy Research (SIEPR). Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau, the IMF, its Executive Board, or its management. All results have been reviewed by the U.S. Census Bureau to ensure no confidential information is disclosed.

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 et al. (2014) survey the theory. Acemoglu and Robinson (2012) provide 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) estimate 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 (2018) analyze 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 find. 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. Such high quality scanner data has 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

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) argues that innovation largely takes the form of creative destruction, and almost always from new firms. Akcigit and Kerr (2018) look 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) point 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

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

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 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 use firm-level data to document the contributions of entry, exit, reallocation, and within-plant productivity growth to overall growth in the manufacturing sector. They use 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, Bloom and Kerr (2017) conduct indirect inference on growth models using firm-level data (from Denmark and the U.S., respectively). They assume the only source of growth is creative destruction, whereas our goal is to infer how much growth comes from creative destruction vs. 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 (2016). 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:

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

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

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

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.

²We normalize the price of aggregate output to 1.

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”). Lastly, 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

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.

³If 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, Akcigit, Bloom and Kerr (2017).

The notation for innovation probabilities is given in Table 1. 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 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 rate κ_e from entrants and at rate 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 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

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

Table 1: Channels of Innovation

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

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

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 1 parameters as follows:

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

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 $\tilde{\delta}_e + \tilde{\delta}_i$ 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 (1) to express growth in terms of contributions by entrants vs. incumbents:

$$\mathbb{E}(1 + g)^{\sigma-1} = 1 + \underbrace{s_{\kappa}\kappa_e + (s_q^{\sigma-1} - 1)\tilde{\delta}_e}_{\text{entrants}} + \underbrace{s_{\kappa}\kappa_i + (s_q^{\sigma-1} - 1)(\lambda_i + \tilde{\delta}_i)}_{\text{incumbents}} - \delta_o\psi \quad (2)$$

Entrants contribute through new varieties κ_e and creative destruction $\tilde{\delta}_e$, with these arrival rates multiplied by their step sizes. Incumbents contribute new varieties κ_i , own innovation λ_i , and creative destruction $\tilde{\delta}_i$, 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 driven by own innovation, creative destruction, and new varieties. These parameters also determine the share of growth driven by incumbents vs. entrants. We will estimate these parameters from the patterns in the LBD micro data. Once we estimate them, we can decompose growth into the three types of innovation as well as the contribution of incumbents vs. entrants.

Firm Dynamics

Firm employment is the outcome of all three types of innovation. A firm's employment is proportional to the number of products the firm produces and the average quality of those products. A firm that is successful in innovating, either by improving the quality of one of its own products, improving the quality of a product made by another firm, or coming up with a brand new variety, will grow in employment. A firm whose products are taken over by another firm, or even fails to innovate, will shrink in employment. A firm will exit when all of its products are creatively destroyed and/or cannot justify their overhead costs.

To convey how we will try to identify the sources of growth, we now highlight the predictions of three polar models, each with only one source of growth.

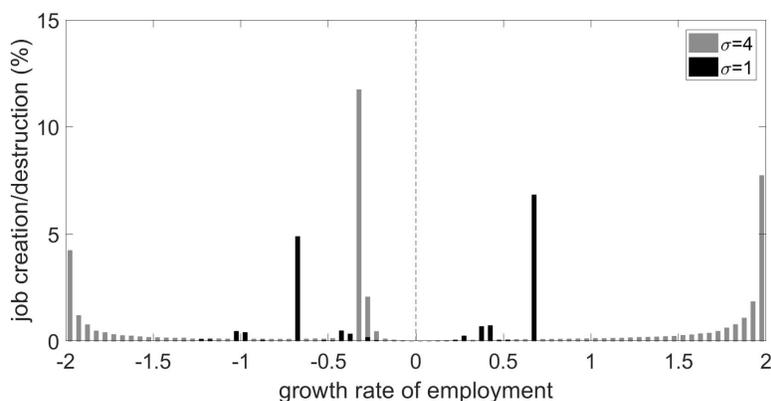
Creative Destruction

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

As can be seen, the distribution of job creation and job 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 polar model because $\sigma = 1$, so firm employment is only a function of the number of varieties the firm produces. The distribution of job creation and job destruction is simply the change in the number of varieties across firms. Figure 1 shows that most

⁵Here we assume $\kappa_i = \kappa_e = \lambda_i = 0$. We use the 1983–1993 data to fix the other parameter values: the Pareto shape parameter for quality draws, θ , to match TFP growth; δ_e to match the entrant share of employment; and δ_i to match the rate of job destruction.

Figure 1: Job Creation and Destruction with only Creative Destruction

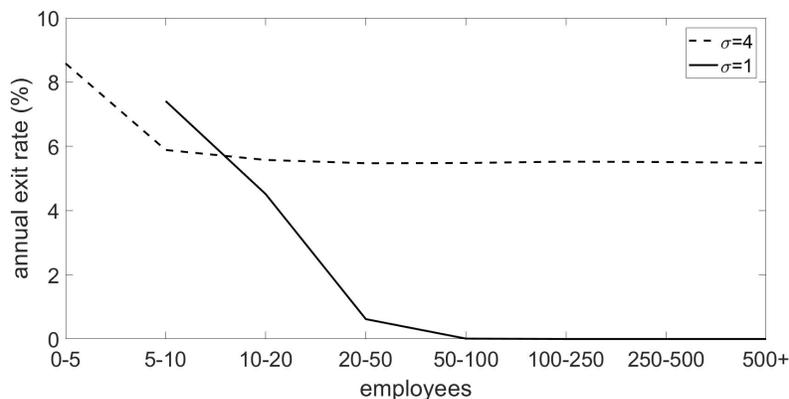


Note: The figure is based on simulating a model with only creative destruction as a source of growth. Employment growth for a firm is defined as the change in employment divided by average employment at the firm at the beginning and end of each period. The vertical axis gives the share of total job creation (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.

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 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

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.

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 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

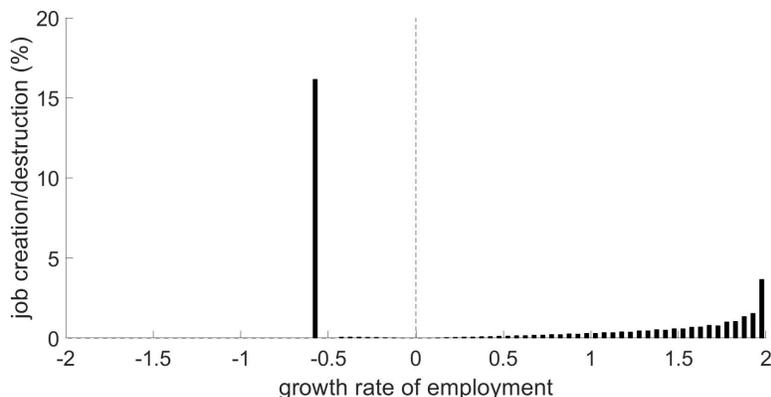
The extreme empirical predictions of a model with only own innovation follow from the absence of creative destruction. Thus, consider a hybrid model in which incumbents improve the quality of their own products and entrants engage in creative destruction.⁸ This hybrid model has the following implications. 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

⁶We assume $\kappa_i = \kappa_e = \delta_i = \delta_e = 0$ and no overhead costs. We set the parameter θ to match 1983–1993 TFP growth and λ_i to the value of δ_i in the creative-destruction only model.

⁷This 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).

⁸We 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.

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.

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 is 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 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.

⁹This is assuming no overhead cost.

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 rate¹⁰
6. Share of job creation due to firm employment growth ≤ 1 ¹¹
7. Minimum firm employment
8. Exit rate for small firms
9. Exit rate for large firms
10. Average employment for incumbents relative to entrants

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).

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

¹¹Recall 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.

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, 2018). 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 (inclusive). 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.¹²

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 2 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

¹²Haltiwanger et al. (2013) show that plants in the LBD grow faster than average until age five.

Table 2: Summary Statistics

	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%

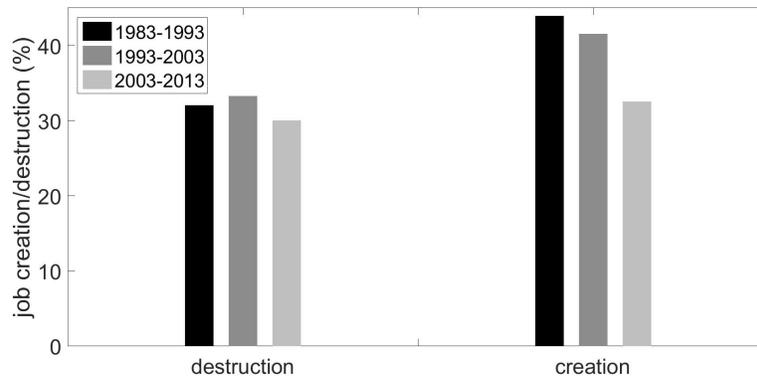
Sources: columns 1-5 are from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. Total employment and number of firms in millions. The last two columns give the average annual growth rates of total employment and TFP, respectively. Data on TFP growth is from the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

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 2 shows the annual TFP growth rate in each ten 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.

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 rate labeled ‘1983–1993’ is the average of the job creation rate 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.

Figure 4: Job Creation and Destruction Rates

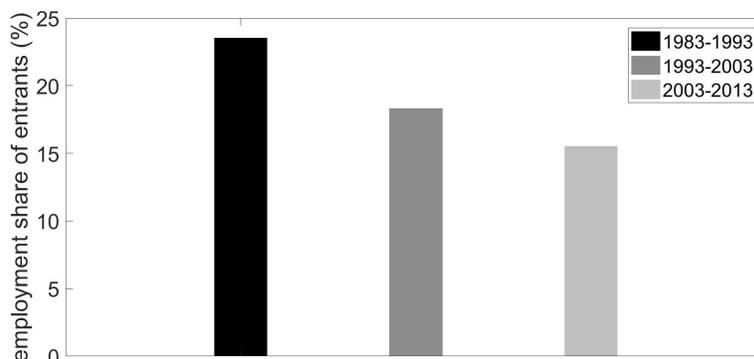


Note: The job creation (destruction) rate is the sum of employment changes at firms with rising (falling) employment divided by the average of aggregate employment in the initial and final year of each 5-year period. This includes entering and exiting firms. The figure shows the average 5-year changes 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.

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 2. 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, Haltiwanger, Jarmin and Miranda (2014) was due to the decline in the job creation rate. The decline in the job destruction rate was much smaller.

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

Figure 5: Employment Share of Entrants



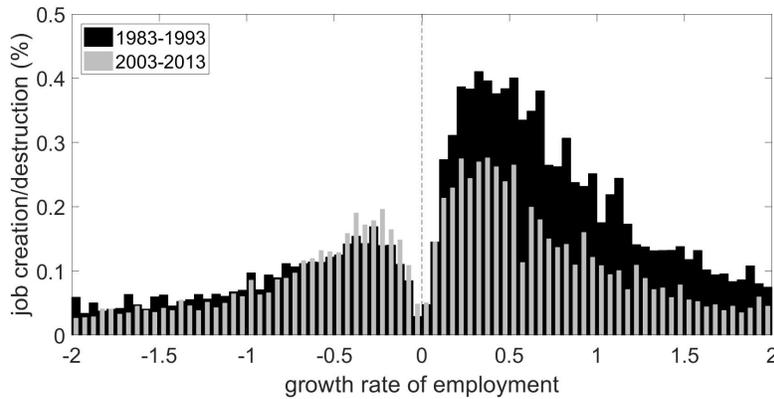
Note: The employment share of entrants for 1983–1993 is the average share of entrant employment in 1988 and 1993. Entrant employment in 1988 (1993) is employment of firms that entered between 1983 and 1988 (1988 and 1993), and the share is relative to aggregate employment in 1988 (1993). The entrant 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.

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 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).

Figure 6: Distribution of Job Creation and Destruction

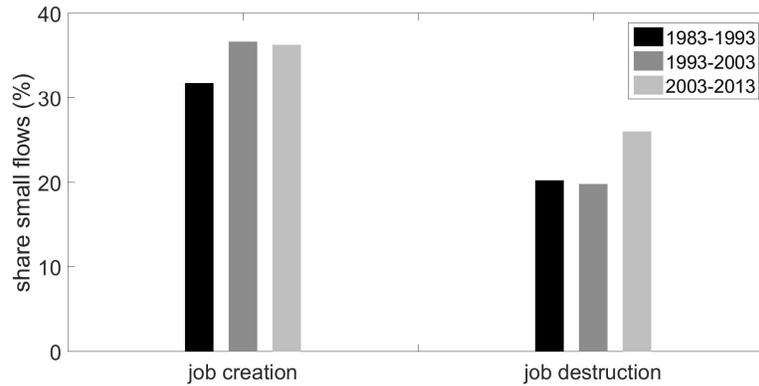


Note: Employment growth for a firm is defined as the change in firm employment over (say) 1988 to 1993 divided by the firm's average employment in 1988 and 1993. The vertical axis gives the share of total job creation (destruction) associated with firms at each given level of employment growth. 1983–1993 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.

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.

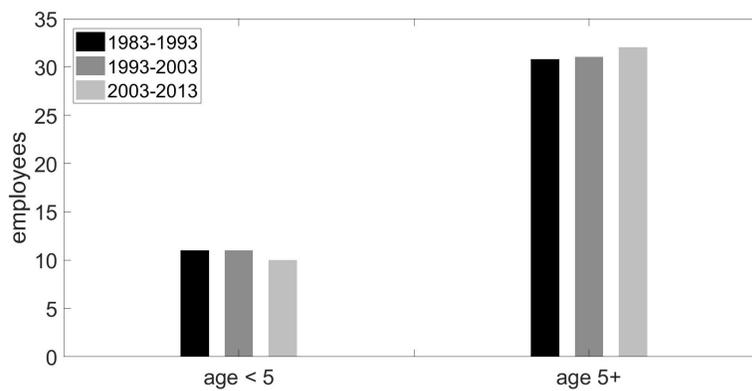
We next present average employment of entrants and incumbents in 1988, 1998, and 2008 (Figure 8). According to Hsieh and Klenow (2014), rapid growth of surviving plants is a robust phenomenon in the U.S. Census of Manufactur-

Figure 7: Share of “Small” Job Flows



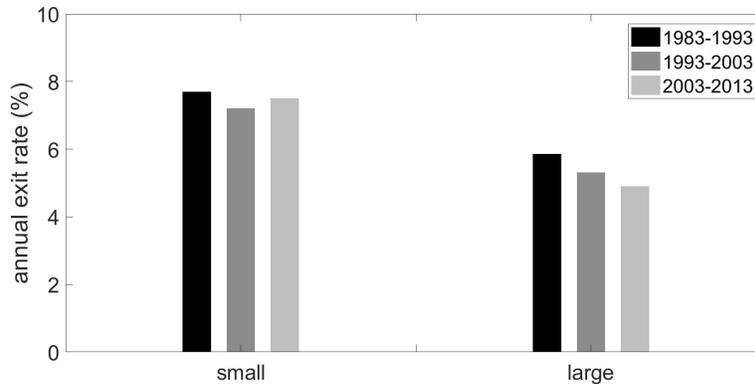
Note: Small flows in job creation is the share of total job creation from firms where employment increased by less than a factor of 3. Small flows in job destruction is share of total job destruction from firms where employment declined by less than a factor of 3. Statistics computed from U.S. Census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

Figure 8: Employment per Firm, Young vs. Old



Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.

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 averages of the exit rates from 1983 to 1988 and 1988 to 1993. The 1993–2003 and 2003–2013 exit rates are defined analogously. Exit rates computed from U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.

ing. 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 vs. small firms, where large is defined as above-mean employment and small as 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. Note that smaller firms have higher exit rates than larger firms. The model can explain this fact if larger firms produce more varieties compared to smaller firms. However, in contrast to the prediction of the polar model with only creative destruction, the decline in exit rates with firm size is gradual. The average size of large firms is about 16 times that of small firms. In the polar model with only creative destruction, the exit rate of large firms should then

equal that of small firms to the 16th power, or essentially zero. In the data, the exit rate of large firms (around 6%) is about 2 percentage points lower than the exit rate of small firms (about 8%).

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 5 innovation rates (δ_i , δ_e , λ_i , κ_i , and κ_e), 2 quality step-size parameters (θ and s_κ), and the overhead cost. So 8 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.¹³

We choose the overhead cost so that the smallest firm has a single production worker. The overhead cost pins down the endogenous rate at which existing varieties disappear due to obsolescence 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 8 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.

¹³We also choose the level of employment in the model to fit employment per firm in the data.

Table 3: Inferred Parameters Values

	1983–1993	1993–2003	2003–2013
Own-variety improvements by incumbents λ_i	69.1%	70.8%	76.9%
Creative destruction by incumbents δ_i	30.3%	48.5%	29.0%
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	11.8%	9.5%	3.2%
Pareto shape of quality draws θ	15.3	11.1	17.4
Relative quality of new varieties s_κ	0.31	0.36	0.31
Average quality of exiting products ψ	0.018	0.024	0.027

Note: Parameter estimates implied by the data moments described in Section 3.

Table 3 presents the 8 parameter values inferred from the data using the procedure described above.¹⁴ Based on the data moments from 1983 to 1993, we infer a 69% arrival rate of own-variety quality improvements per five-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 with probability one. The *unconditional* probability that a given product improves due to creative destruction by an incumbent is thus 9.4%, and the unconditional probability of creative destruction by an entrant is 21.5%.¹⁵ The unconditional probability that a product is improved upon in a five-year period is thus 100%, of which 69% is from own innovation and 31% is from creative destruction (the latter from entrants or incumbents).

¹⁴Table 16 in Appendix B presents bootstrapped standard errors for the parameter estimates.

¹⁵ $(1-0.691) \cdot 0.303 = 0.094$ and $(1-0.691) \cdot (1-0.303) = 0.215$.

The employment-weighted average step size for quality improvements on existing varieties is given by $s_q = \left(\frac{\theta}{\theta - (\sigma - 1)}\right)^{1/(\sigma - 1)}$. Given that $\theta = 15.3$ and $\sigma = 4$, the average improvement in quality (conditional on innovation) is 7.5%. New varieties are only created by entrants, arrive with 11.8% 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 3% 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 11.8% every five years, which matches the growth of total employment and number of firms from 1983–1993.

Table 4 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 (1) for this calculation.¹⁶ 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 0.46% per year. And new varieties generates growth of 0.22% per year.

The rows in Table 5 show the contribution of each source of innovation to aggregate TFP growth. About 28% of the 1.66% growth rate in the 1983–1993 period comes from creative destruction. Own-variety improvements by incumbents account for 59%. New varieties *a la* Romer (1990) are the remainder at around 13%.

The columns in Table 5 also decompose aggregate TFP growth into the percentage contribution of entrants vs. incumbents using equation (2). In the ten years between 1983 and 1993, incumbents account for 68% of aggregate TFP growth, with entrants contributing the remaining 32%. Aghion et al. (2014) provide complementary evidence for the importance of incumbents based on their share of R&D spending and patents.

¹⁶Equation (1) 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 4: TFP Growth from Innovation

	1983–1993	1993–2003	2003–2013
Creative Destruction	0.46%	0.64%	0.30%
New Varieties	0.22%	0.20%	0.06%
Own Innovation	0.98%	1.46%	0.96%

Note: 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 3.

Table 3 shows that the arrival rate of new varieties κ_e fell from 11.8% in 1983–1993 to 9.5% in 1993–2003. We infer this because of 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 0.22% per year in 1983–1993 to 0.20% 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 69% to 71% (Table 3). TFP growth from own innovation rose from 1% to almost 1.5% per year (Table 4). Aggregate TFP growth (from all three sources of innovation) increased by 0.64 percentage points from 1983–1993 to 1993–2003 (Table 2). Therefore, own innovation was responsible for almost 80% of the increase in aggregate TFP growth, and the remaining 20% due to more innovation from creative destruction.

Table 5: Contribution to Aggregate Growth

<u>1983-1993</u>	Entrants	Incumbents	
Creative destruction	19.1%	8.4%	27.5%
New varieties	13.4%	0.0%	13.4%
Own-variety improvements	-	59.1%	59.1%
	32.5%	67.5%	
<u>1993-2003</u>	Entrants	Incumbents	
Creative destruction	14.4%	13.6%	27.9%
New varieties	8.6%	0.0%	8.6%
Own-variety improvements	-	63.5%	63.5%
	23.0%	77.0%	
<u>2003-2013</u>	Entrants	Incumbents	
Creative destruction	16.1%	6.6%	22.7%
New varieties	4.6%	0.0%	4.6%
Own-variety improvements	-	72.7%	72.7%
	20.8%	79.2%	

Note: The table presents the percentage contribution of each source of innovation to aggregate TFP growth. Rows use equation (1) to decompose aggregate TFP into the contribution of creative destruction, new varieties, and own innovation. Columns use equation (2) to decompose aggregate TFP into the contribution of entrants and incumbents.

Turning to the 2003–2013 period, the main fact is the eight 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 vs. 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 3.6%. The residual entrant employment due to new varieties is creative destruction by entrants. Table 3 indicates that the unconditional probability of creative destruction by an entrant fell from 21.5% in 1983–1993 to 16.4% by 2003–2013. The net effect is that TFP growth from innovation by creation of new varieties by entrants fell to 0.05% per year in 2003–2013, down from 0.22% in the 1983–1993 period.

Job creation by incumbents also 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 3, the probability a variety is creatively destroyed by an incumbent firm dropped from 9.4% to 6.7%. Table 4 shows that TFP growth from creative destruction falls from 0.46% in 1983–1993 to 0.3% per year 2003–2013. Since TFP growth from all three types of innovation fell by 0.34 percentage points over this period (Table 2), the decline in creative destruction accounts for almost all the decline in aggregate TFP growth.

So the model’s proximate answer to the question “how much the decline in job creation matters for aggregate TFP growth” is that entrant innovation and incumbent creative destruction have declined. Still, what the aggregate job creation rate misses is the contribution of own innovation. Table 4 shows that growth from own innovation by incumbent firms was roughly the same in 2003–2013 compared to the 1983–1993 period.

As noted, our estimation procedure chose parameter values such that the model exactly matched aggregate TFP growth. Table 6 (column 2) shows the fit of the model in 2003–2013 for the data moments we used for the estimation

Table 6: Model Fit , 2003–2013

	Data	Baseline Model	Directed CD By Entrant
Employment share of entrants	15.5%	16.3%	17.0%
Employment growth rate	0.5%	0.5%	0.8%
Job creation rate	32.5%	28.7%	31.2%
Job destruction rate	30.0%	26.3%	27.0%
Share of job creation < 1	36.3%	21.4%	25.0%
SD(log employment)	1.28	1.25	1.21
Exit rate large/small	0.65	0.81	0.60
Employment incumbents/entrants	3.2	1.5	2.4

Note: 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.

but did not force the model to perfectly match.¹⁷ 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 large vs. small firms, and the average size of incumbent firms vs. entrants.¹⁸

Andrews, Gentzkow and Shapiro (2017) propose a local measure of the relationship between parameter estimates and moments. In this spirit, Table 18 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 2003–2013 estimated parameters.

Table 18 conforms well to our intuition. The entrant employment share is positively related to the pace of creative destruction by entrants (and negatively related to own innovation, which trims the unconditional probability of creative destruction by entrants). The job creation rate is inversely related to own innovation, again because it crowds out creative destruction. The job destruction rate is negatively related to own innovation and new varieties, and positively related to the pace of creative destruction. The share of job creation < 1 is increasing in own innovation and new variety creation by incumbents. Size dispersion is increasing in the rate of new variety creation by incumbents 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 below-average in quality.

The model understates the exit rate and 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

¹⁷We also set the minimum firm size near 1, as in the data.

¹⁸The model fits slightly better in 1983–1993 and 1993–2003 (see Table 17 in Appendix C).

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 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* — i.e., 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 6 shows the fit of the model when we introduce $\rho_e \leq 1$.¹⁹ We estimate $\rho_e = 0.57$, 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 4% 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.²⁰ The share of job destruction due to small changes in employment is 25%, up from 21% 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 82% of growth in the model with $\rho_e \leq 1$, compared to 89% in our baseline model.

¹⁹We target the same 10 data moments.

²⁰The arrival rate (average quality) of new varieties by incumbents is $\kappa_i = 0.04$ ($s_\kappa = 0.47$).

Table 7: Simulated Distribution of Products per Firm

	1983–1993	1993–2003	2003–2013
1 Product	91.7%	84.7%	89.0%
2 Products	7.2%	12.0%	9.4%
3 Products	0.9%	2.5%	1.3%
> 3 Products	0.1%	0.8%	0.2%
Average Products per Firm	1.09	1.20	1.13

Note: Table presents the number of products per firm implied by the innovation parameters in Table 3.

We did not target the share of large employment changes in job destruction to arrive at our baseline innovation parameter estimates in Table 3. 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, an implication is that most firms will own only a small number of products. Table 7 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 7, is around 1.1 to 1.2. We cannot verify this empirically because we do not know of a dataset 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.

5. Robustness

We now examine the robustness of our inference exercise. First, we use the number of establishments as a proxy for the number of products of the firm. Second, we use alternative measures of job flows to infer the sources of growth. Third, we measure the sources of innovation in the fast growing information and communication technology sector.

5.1. 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 8. 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 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, $R^2 = 0.68$).²¹

²¹We 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 with a standard error of 0.016 and an R-squared is 0.52.

Table 8: Alternative Measures of Job Flows in the LBD (2003–2013)

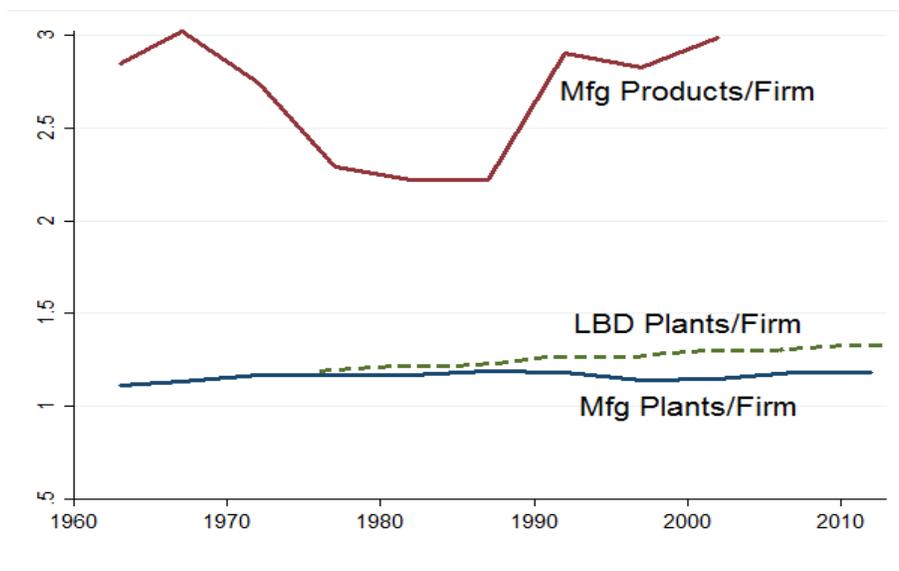
	Full Sample	Drop New & Exiting Estab.	Only Services	Exclude M&A ¹	Exclude M&A ²
Job Creation Rate	32.5%	10.1%	34.2%	25.6%	26.2%
Employment Share of Entrants	15.1%	–	15.1%	14.9%	14.8%
Share of Job Creation < 1	36.3%	78.1%	35.7%	39.9%	39.1%
Job Destruction Rate	30.0%	9.6%	29.2%	23.1%	23.1%

¹ Exclude plants that switch ownership between 2003 and 2008 (and between 2008 and 2013).

² 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.

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 three 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 five 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.

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.

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 ten 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.²²

²²See Bernard, Redding and Schott (2010) for evidence that individual manufacturing plants frequently add and subtract products, even at the 5-digit SIC level.

5.2. Alternative Measures of Job Flows

We now look at job flows in alternative samples. First, we have a closed economy model so 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 non-traded service sectors in column 3 of Table 8. Job flows outside of manufacturing are essentially identical to flows in the full-sample that includes traded sectors.

Recall that we include mergers and acquisitions in our baseline statistics on job creation and job destruction. The last two columns in Table 8 check the effect of dropping those establishments who change ownership. Column 4 entirely drops from the sample 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) to measure innovation. In the model, the change in employment is the same as the change in output, but this is not true in the data. For example, 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.²³ The hypothetical job flows with this imputation are shown in the second column of Table 9. 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 per-

²³We impute employment as the product of the firm's share of industry's value-added and total employment in the industry.

Table 9: Job Flows in Manufacturing (2002–2012)

	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%

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 $TFPQ_i^{\sigma-1} / \sum_j TFPQ_j^{\sigma-1}$.

centage 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 9 imputes the job flows implied by the change in firm *productivity* (“TFPQ”) in the Census of Manufacturing.²⁴ 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 10 shows what these alternate job flow statistics imply for the sources

²⁴Following 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 $TFPQ_i^{\sigma-1} / \sum_j TFPQ_j^{\sigma-1}$.

Table 10: Sources of Growth Using Alternative Measures of Job Flows

	Baseline	No Mergers	Imputed from VA	Imputed from TFPQ
Creative destruction	22.7%	19.7%	21.3%	49.2%
New varieties	4.6%	4.1%	4.6%	9.8%
Own-variety improvements	72.7%	76.2%	74.2%	41.0%

Note: 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.

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 mergers and acquisitions have no effect on TFP growth. Since job creation and destruction rates are lower when we drop mergers and acquisitions (Table 8), we infer a lower rate of creative destruction (and a higher rate of own innovation).²⁵

The next two columns in Table 10 show the effect of using output and productivity on the sources of growth. To arrive at these numbers, we impute job

²⁵We 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.

flows in the LBD from output and productivity data in manufacturing.²⁶ The sources of growth implied by the 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 9. The share of small changes does not change, so the relative importance of creative destruction vs. own innovation also does not change.

The story is quite different when we target job flows implied by productivity growth (column 4 in Table 10). 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 9 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.

Remember we calculate job flows over five years. We present job flows over *one* year in Table 11. An entrant is now defined as a firm created in the last year (instead of the last five years), and job creation between years t and $t + 1$ is now 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.²⁷

²⁶We 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 9 and those in columns 2 and 3 of the same Table. This adjustment increases the job creation and destruction 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 vs. output (or productivity).

²⁷Job destruction between t and $t + 1$ is the sum of employment of firms in year t that exit by year $t + 1$ and the change in employment among firms with shrinking employment between these two years.

Table 11: Job Flows Over One Year in LBD

	1983–1993	2003–2013
Job Creation Rate	16.9%	13.1%
Employment Share of Entrants	5.2%	3.4%
Share of Job Creation by Incumbents Age < 5	22.4%	16.5%
Share of Job Creation < 1	51.2%	54.0%
Job Destruction Rate	14.5%	12.6%
Share of Job Destruction by Incumbents Age < 5	17.7%	13.0%

Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. The job flows are computed over 1 year intervals over each ten year period. Entrants defined as firms established in the last year.

The data moments change in two ways. First, firms between the ages of 1 and 5 previously classified as entrants are now incumbents. As a consequence, the employment share of entrants is lower. Second, job flows are larger. The annual job creation rate from 2003 to 2013 in Table 11 is 13.1%. The equivalent number implied by the job creation rate over five years was 6.5%.²⁸ Intuitively, job creation and destruction now include flows *within* each five year period. For example, Table 11 indicates that incumbent firms with age < 5 account for 16.5% of job creation and 13% of job destruction. Job creation computed over five years only counted *net* job creation by incumbents age < 5 (which we attributed to entrants).

However, remember the *distribution* of job changes is key for distinguishing between 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.²⁹ Our inference exercise will

²⁸The job creation rate over five years displayed in Figure 4 from 2003–2013 was 32.5%.

²⁹In 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.

Table 12: Sources of Growth Implied by Job Flows over 1 year

	1983–1993	2003–2013
Creative destruction	10.6%	10.9%
New varieties	23.1%	5.6%
Own-variety improvements	66.2%	83.4%

Note: Entries are the percent of aggregate TFP growth due to each source of innovation implied by the job flows over one year shown in Table 11.

therefore indicate that the share of growth driven by own innovation vs. creative destruction will now be even higher. Table 12 shows the share of growth due to each source of innovation implied by job flows over 1 year. The contribution of creative destruction to aggregate growth from 2003 to 2013 is now only 11%. In the same time period, the share of growth from creative destruction inferred from job flows over 5 years was 23%. The share of growth from own innovation is, correspondingly, about ten percentage points higher at 83%.

The one-year flows also convey a decline in job creation between 1983–1993 and 2003–2013. The difference now is that the decline is due to the diminishing employment share of entrants and smaller job flows by young firms (incumbents age < 5).³⁰ 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 five-year job flows, own innovation has become more important.

We reiterate that, when innovation can take the form 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 compared to another firm that grows due to creative

³⁰When 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).

Table 13: Job flows vs. Innovation by Age, 2003–2013

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

Note: 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 11.

destruction, even when the innovation step size is the same for the two types of innovation. Table 13 illustrates this point by showing the contribution of young vs. old firms to job creation (column 1) versus TFP growth (column 2). Compare, for example, the contribution of entrants (age < 1) with that of old firms (age > 15). As noted by many authors, a large share of job creation is due to entry — entrants account for 31% of all job creation in our calculation.³¹ Yet entrants only account for 9% of TFP growth. In contrast, old firms (age > 15) are responsible for more than 50% of TFP growth, but closer to 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.

³¹Haltiwanger, Jarmin and Miranda (2013) provide similar evidence on the importance of entrants for job creation.

Table 14: Information Technology & Communications Industry

	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%

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 five-year periods.

5.3. Innovation in the Information Technology Sector

We end with a case study of the fast growing information and communication technology (ICT) sector. Table 14 (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%.³² But the distribution of job flows also changes. The acceleration 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 15 shows that 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 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

³²We revert back to measuring job flows over five years.

Table 15: Sources of Growth in Information Technology

	1983-1993	1993-2003	2003-2013
Creative destruction	17.0%	35.9%	15.3%
New varieties	5.9%	2.9%	4.4%
Own-variety improvements	77.0%	61.0%	78.4%

Note: Table shows the percentage of growth due to each source of innovation in the ICT sector implied by the job flows shown in Table 14.

back to 16%.³³ 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 ten 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.

³³For fitting the ICT facts in our simulations, we allow the scale parameter of the Pareto distribution of quality draws to be greater than 1. Without this additional parameter, 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 (2018), 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 estimate 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 U.S. led by Walmart's expansion. Online retailing has made inroads at the expense of brick-and-mortar stores (Einav, Klenow, Klopach, Levin, Levin and Best, 2017). 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 U.S. (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 1993–2003.

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 et al. (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.

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 spillovers are just as strong for incumbent innovations as for entrant innovations. Young firms might instead generate more knowledge spillovers than old firms do — Akcigit and Kerr (2018) provide 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.³⁴

We assumed no frictions in employment growth or misallocation of labor 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 awhile to build up a customer base, as in work by Gourio and Rudanko (2014) and Foster, Haltiwanger 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 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.

³⁴Smith, Yagan, Zidar and Zwick (2017) use administrative tax records to document that firm profits dive after the death of under-65 owner/managers.

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 1. 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
 - (h) Exit rate for small firms
 - (i) Exit rate for large firms
 - (j) Average employment for incumbents relative to entrants
5. Repeat steps 1-4 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 (1) exactly match TFP growth in the data; (2) set minimum firm employment to 1; and (3) minimize the mean squared percent distance between the simulated and empirical moments for the remaining statistics in step 4.

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:

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

where J is the Jacobian matrix (in levels), W is the weighting matrix, and $\tilde{\Sigma}$ 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 one for all moments except for exit by size and size by age, whose corresponding entries are 0.5. $\tilde{\Sigma}$ is estimated by Monte Carlo simulation, calculating the variance-covariance matrix after 200 model simulations with fixed parameters.

Table 16 shows the estimated coefficients and standard errors for the pa-

Table 16: Bootstrapped parameter estimates, 2003-2013

	Coefficient	Standard error
Creative destruction by incumbents δ_i	0.290	0.0098
New varieties by entrants κ_e	0.030	0.0021
Own-variety improvement by incumbents λ_i	0.769	0.0041
Relative quality of new varieties s_κ	0.31	0.0309

rameters with interior solutions. As can be seen, the standard errors of the estimated parameters are extremely small.

C Model Fit and Sensitivity

Table 17 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 ten year period. And Table 18 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 moments with respect to a 0.1 level increase in each parameter.

Table 17: Model Fit , 1983–1993 and 1993–2003

	1983–1993		1993–2003	
	Data	Model	Data	Model
Entrant employment share	23.5%	23.8%	18.3%	17.6%
Employment growth rate	2.4%	2.3%	1.6%	1.8%
Job creation rate	43.9%	41.5%	41.5%	40.6%
Job destruction rate	32.0%	30.4%	33.2%	31.7%
Share of job creation < 1	31.7%	21.4%	36.6%	24.7%
SD(log employment)	1.25	1.37	1.27	1.32
Exit rate large/small	0.76	0.84	0.74	0.70
Employment incumbents/entrants	2.8	1.7	2.8	2.1

Table 18: Jacobian Matrix, 2003–2013

	δ_i	δ_e	κ_i	κ_e	λ_i	s_κ	ρ_e
Entrant employment share	0%	20%	5%	-6%	-51%	9%	0%
Job creation rate	9%	7%	7%	9%	-33%	0%	-3%
Job destruction rate	10%	8%	-25%	-23%	-36%	0%	-3%
Share of job creation < 1	2%	-6%	16%	4%	32%	-6%	2%
SD(log employment)	-18%	-19%	83%	16%	4%	-15%	-3%
Exit rate of small firms	4%	5%	-18%	1%	-23%	0%	0%
Exit rate of large firms	6%	9%	-38%	-10%	-30%	0%	-3%
Employment entrants	-20%	11%	23%	-63%	-13%	12%	7%
Employment incumbents	3%	0%	-3%	8%	-1%	-1%	-1%

Note: 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 2003–2013 estimated parameters.

References

- Acemoglu, Daron, “Labor- and Capital-Augmenting Technical Change,” *Journal of the European Economic Association*, 2003, 1 (1), 1–37.
- , *Introduction to Modern Economic Growth*, Princeton University Press, 2011.
- and James Robinson, *Why Nations Fail: Origins of Power, Poverty and Prosperity*, Crown Publishers (Random House), 2012.
- , Ufuk Akcigit, Nicholas Bloom, and William R Kerr, “Innovation, Reallocation and Growth,” Working Paper 18993, National Bureau of Economic Research 2017.
- Aghion, Philippe and Peter Howitt, “A Model of Growth Through Creative Destruction,” *Econometrica*, 1992, 60 (2), pp. 323–351.
- and —, *The Economics of Growth*, Massachusetts Institute of Technology (MIT) Press, 2009.
- , Christopher Harris, Peter Howitt, and John Vickers, “Competition, Imitation and Growth with Step-by-Step Innovation,” *The Review of Economic Studies*, 2001, 68 (3), pp. 467–492.
- , Ufuk Akcigit, and Peter Howitt, “What Do We Learn From Schumpeterian Growth Theory?,” *Handbook of Economic Growth*, 2014, 2B, 515–563.
- Akcigit, Ufuk and William R. Kerr, “Growth through Heterogeneous Innovations,” *Journal of Political Economy*, 2018, *forthcoming*.
- Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro, “Measuring the sensitivity of parameter estimates to estimation moments,” *The Quarterly Journal of Economics*, 2017, 132 (4), 1553–1592.
- Atkeson, Andrew and Ariel Burstein, “Aggregate Implications of Innovation Policy,” *Journal of Political Economy*, 2018, *forthcoming*.
- Baily, Martin Neil, Charles Hulten, and David Campbell, “Productivity Dynamics in Manufacturing Plants,” *Brookings Papers: Microeconomics*, 1992, 4, 187–267.

- Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott, "Multi-Product Firms and Product Switching," *American Economic Review*, 2010, 100 (1), 70–97.
- Broda, Christian and David E Weinstein, "Globalization and the Gains From Variety," *The Quarterly Journal of Economics*, 2006, 121 (2), 541–585.
- and —, "Product Creation and Destruction: Evidence and Implications," *American Economic Review*, 2010, 100 (3), 691–723.
- Christensen, Clayton, *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*, Harvard Business Review Press, 1997.
- Davis, Steven J, John C Haltiwanger, and Scott Schuh, "Job Creation and Destruction," *MIT Press Books*, 1996.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, "The Role of Entrepreneurship in US Job Creation and Economic Dynamism," *The Journal of Economic Perspectives*, 2014, 28 (3), 3–24.
- Einav, Liran, Peter J Klenow, Benjamin Klopock, Jonathan D Levin, Larry Levin, and Wayne Best, "Assessing the gains from e-commerce," 2017.
- Feenstra, Robert C, "New product varieties and the measurement of international prices," *The American Economic Review*, 1994, pp. 157–177.
- Foster, Lucia, John C Haltiwanger, and C.J. Krizan, "Aggregate Productivity Growth. Lessons from Microeconomic Evidence," *New Developments in Productivity Analysis*, 2001, pp. 303–372.
- , John Haltiwanger, and Chad Syverson, "The Slow Growth of New Plants: Learning about Demand?," *Economica*, 2016, 83 (329), 91–129.
- Gordon, Robert J, *The Measurement of Durable Goods Prices*, University of Chicago Press, 2007.
- Gourieroux, Christian and Alain Monfort, *Simulation-based econometric methods*, Oxford University Press, 1996.

- Gourio, Francois and Leena Rudanko, "Customer capital," *Review of Economic Studies*, 2014, 81 (3).
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 1997, 87 (3), 342–362.
- Grossman, Gene M and Elhanan Helpman, "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 1991, 58 (1), 43–61.
- Haltiwanger, John, Ron S Jarmin, and Javier Miranda, "Who Creates Jobs? Small versus Large versus Young," *Review of Economics and Statistics*, 2013, 95 (2), 347–361.
- Hottman, Colin J, Stephen J Redding, and David E Weinstein, "Quantifying the sources of firm heterogeneity," *The Quarterly Journal of Economics*, 2016, 131 (3), 1291–1364.
- Howitt, Peter, "Steady Endogenous Growth with Population and R & D Inputs Growing," *Journal of Political Economy*, 1999, 107 (4), pp. 715–730.
- Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- and Peter J. Klenow, "The Life Cycle of Plants in India and Mexico," *Quarterly Journal of Economics*, 2014, 129 (3), 1035–1084.
- Jones, Charles I., "Life and Growth," *Journal of Political Economy*, 2016, 124, 539–578.
- Jovanovic, Boyan, "Selection and the Evolution of Industry," *Econometrica*, 1982, 50 (3), 649–670.
- Klette, Tor Jakob and Samuel Kortum, "Innovating firms and aggregate innovation," *Journal of Political Economy*, 2004, 112 (5), 986–1018.
- Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, 65 (6), 1389–1419.
- Krusell, Per, "How is R&D Distributed Across Firms? A Theoretical Analysis," 1998.

- and Jose-Victor Rios-Rull, “Vested Interests in a Positive Theory of Stagnation and Growth,” *Review of Economic Studies*, 1996, 63 (2), 301–329.
- Lentz, Rasmus and Dale T Mortensen, “An Empirical Model of Growth through Product Innovation,” *Econometrica*, 2008, 76 (6), 1317–1373.
- Lucas, Robert E Jr. and Benjamin Moll, “Knowledge Growth and the Allocation of Time,” *Journal of Political Economy*, 2014, 122 (1).
- Parente, Stephen L and Edward C Prescott, *Barriers to Riches*, MIT press, 2002.
- Peters, Michael, “Heterogeneous Mark-ups, Growth and Endogenous Misallocation,” 2016.
- Romer, Paul M., “Endogenous Technological Change,” *Journal of Political Economy*, 1990, 98 (5), pp. S71–S102.
- Schumpeter, Joseph A., *Business Cycles*, Cambridge University Press, 1939.
- Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick, “Capitalists in the Twenty-first Century,” *UC Berkeley and University of Chicago Working Paper*, 2017.
- Stokey, Nancy L, “Learning by Doing and the Introduction of New Goods,” *Journal of Political Economy*, 1988, 96 (4), 701–717.
- Young, Alwyn, “Growth without Scale Effects,” *Journal of Political Economy*, 1998, 106 (1), pp. 41–63.