NBER WORKING PAPER SERIES

HOW DESTRUCTIVE IS INNOVATION?

Daniel Garcia-Macia
Chang-Tai Hsieh
Peter J. Klenow

Working Paper 22953
http://www.nber.org/papers/w22953

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2016

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, its management, or the National Bureau of Economic Research. All results have been reviewed by the U.S. Census Bureau to ensure no confidential information is disclosed.

How Destructive is Innovation?
Daniel Garcia-Macia, Chang-Tai Hsieh, and Peter J. Klenow
NBER Working Paper No. 22953
December 2016
JEL No. E24,O3,O4,O5

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 non-farm private businesses from 1976–1986 and 2003–2013, 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.

Daniel Garcia-Macia
Research Department
International Monetary Fund
700 19th St NW
Washington, DC 20431
dgarciamacia@imf.org

Chang-Tai Hsieh
Booth School of Business
University of Chicago
5807 S Woodlawn Ave
Chicago, IL 60637
and NBER
chsieh@chicagoBooth.edu

Peter J. Klenow
Department of Economics
579 Serra Mall
Stanford University
Stanford, CA 94305-6072
and NBER
Klenow@Stanford.edu

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 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 (2015).

Still other theories emphasize the contribution of brand new varieties to growth. 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 empirically for U.S. imports.

These theories have different implications for innovation policy. Business stealing is a force pushing up the private return to innovation relative to the social return. To the extent firms build on each other's innovations, there are positive knowledge externalities to innovation. When incumbents improve their own products, business stealing and knowledge externalities are mitigated. Models with expanding varieties, meanwhile, tend to have smaller business-stealing effects but retain knowledge spillovers. Burstein and Atkeson (2015) drive home the distinct policy implications of these competing theories.

Ideally, one could directly observe the extent to which new products substitute for existing products. Broda and Weinstein (2010) and Hottman et al. (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]

Similarly, 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 a small number of case studies, Christensen (1997) argues that innovation largely takes the form of creative destruction and that such innovation almost always comes from new firms. Akcigit and Kerr (2015) use patent citations to measure whether the citations are to patents of the same firm or of other firms. The limitation is that the case studies and the sample of patenting firms may not be representative of firms in the broader economy. For example, 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. non-farm economy. We use data from the U.S. Longitudinal Business Database (LBD) from 1976–1986 and 2003–2013. The seminal work of Davis et al. (1998) 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 et al. (2014) point to the decline in U.S. job reallocation since the 1970s as evidence of a decline in the rate of creative destruction.

Viewing the LBD data through the lens of an exogenous growth model featuring creative destruction, own innovation, and new varieties, we reach four conclusions. First, most growth appears to come from incumbents rather than

[1] Gordon (2007) and Greenwood et al. (1997) emphasize the importance of growth embodied in durable goods based on the declining relative price of durables.

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 1976–1986 to 2003–2013, while the contribution of incumbent firms, particularly through own innovation, increased.

Influential papers by Baily et al. (1992) and Foster et al. (2001) use similar 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 lenses of a specific (albeit somewhat general) model of growth. Like us, Lentz and Mortensen (2008) and Acemoglu et al. (2013) conduct indirect inference on growth models with manufacturing data (from Denmark and the U.S., respectively). They assume the only source of growth is creative destruction, however, 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).

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 U.S. LBD that we exploit to infer the sources of innovation. Section 4 presents the model parameter values which best match the moments from the data. Section 5 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 further 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 within varieties. Without this assumption, firms would engage in limit pricing and markups would be heterogeneous as in Peters (2013).

Firms face the same 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 wage and L is aggregate employment.[2] Employment of a firm L_f is then given by:

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

[2]We normalize the price of aggregate output to 1. Note the model features no misallocation.

where M_f denotes the number of varieties produced by firm f. 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 $\sigma = 1$ assumed by Klette and Kortum (2004), the same 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 wage is proportional to aggregate labor productivity, which is given by:

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

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

Innovation

Aggregate growth in the model comes from the creation of new varieties and from growth in average quality. In Klette and Kortum (2004), the number of varieties is constant and quality growth can only come when a firm innovates upon and takes over a variety owned by another firm ("creative destruction"). We will also allow quality growth to come from innovation by firms to improve the quality of the products they own ("own innovation"). Lastly, we also 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] Specifically, the quality drawn by an innovation follows a Pareto distribution with shape parameter θ and scale parameter equal to the existing quality level. The average improvement in quality of an existing variety (conditional on innovation), weighted by employment, is thus $s_q = \left(\frac{\theta}{\theta-(\sigma-1)}\right)^{1/(\sigma-1)} > 1$. Finally, we assume that entrants have one product, which they obtain by improving upon an existing variety or by creating a brand new variety. Incumbent firms, on the other hand, potentially produce many varieties.

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 thus 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 $\tilde{\delta}_e \equiv \delta_e(1-\delta_i)(1-\lambda_i)$. 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 $s_q > 1$.

The rate at which new varieties are created is governed by κ_e and κ_i. 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

[3]If innovation was endogenous, there would be a positive externality to research unless all research was done by firms on their own products. Such knowledge externalities are routinely assumed in the quality ladder literature, such as Grossman and Helpman (1991), Aghion and Howitt (1992), Kortum (1997), and Acemoglu et al. (2013).

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 new variety is drawn from the quality distribution of existing products multiplied by s_κ.

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 such that the minimum size firm has 1 unit of labor for production plus overhead. The overhead cost determines the cutoff quality of varieties as products with a quality below the threshold have negative present discounted value of profits (even taking into account the arrival of innovations associated with owning a variety), and therefore exit endogenously. This cutoff, which we denote as ψ percent of employment-weighted mean quality of existing varieties, rises endogenously with wage growth which ensures that the distribution of quality across varieties is stationary. We denote the endogenous exit rate of existing varieties due to overhead as δ_o. The net growth rate of varieties is therefore $\kappa_e + \kappa_i - \delta_o$.

We can now express the expected growth rate of the wage and output per worker as a function of the parameters in Table 1:

$$1+g=\left(1+\underbrace{s_\kappa\left(\kappa_e+\kappa_i\right)}_{\text{new varieties}}+\underbrace{\left({s_q}^{\sigma-1}-1\right)\lambda_i}_{\text{own innovation}}+\underbrace{\left({s_q}^{\sigma-1}-1\right)\left(\tilde{\delta}_e+\tilde{\delta}_i\right)}_{\text{creative destruction}}-\delta_o\psi\right)^{\frac{1}{\sigma-1}} \tag{1}$$

The contribution of new varieties is given by $s_\kappa\left(\kappa_e+\kappa_i\right)$ which is increasing in the arrival rates κ_e and κ_i and 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 also rearrange (1) to express growth in terms of contributions by entrants vs. incumbents:

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

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 ($\kappa_i, \kappa_e, \lambda_i, \delta_i,$ 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. The same 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 in the model 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. At the extreme, a firm will exit when all of its products are creatively destroyed by other firms or are too low in quality to cover 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.[4] 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 (entrant or incumbent) innovates on their products. The rate of job destruction is therefore pinned down by the rate 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

[4]Here we assume $\kappa_i = \kappa_e = \lambda_i = 0$. We use the 1976–1986 data to fix the other parameter values: the Pareto shape parameter for quality draws, θ, to match TFP growth; δ_e to match the entrants share of employment; and δ_i to match the level of job destruction.

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, shown in Figure 1 (see the bars labeled $\sigma = 1$). Following Davis et al. (1998), the percent change in firm employment on the horizontal axis is measured as the change in firm employment divided by the average of the last period and the current period employment of the firm. 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 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 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 varieties. The probability that a firm with n varieties exits is thus given by the exit probability of a one-variety firm to the *power* of n. Since the employment of a firm is proportional to the number of varieties it produces, the model predicts that an n-fold difference in firm employment will be associated with a change in the

Figure 1: Job Creation and Destruction with only Creative Destruction

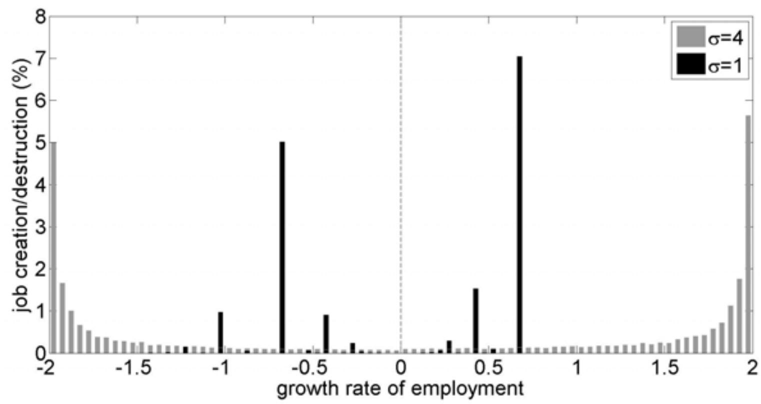

Note: The figure is based on simulating a model with only creative destruction as a source of growth. Employment growth for a firm is defined as the change in employment divided by average employment at the firm across the two periods. The vertical axis gives the sum of total job creation (destruction), divided by aggregate employment in the initial period, 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.

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 $\sigma = 1$. Figure 2 shows that changing 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. More quality heterogeneity across firms implies that more firms will experience large changes in employment when they grab a high quality variety from another firm. Similarly, firms will experience a sharp drop in employment when they lose their high quality varieties.

Figure 2: Exit by Size with only Creative Destruction

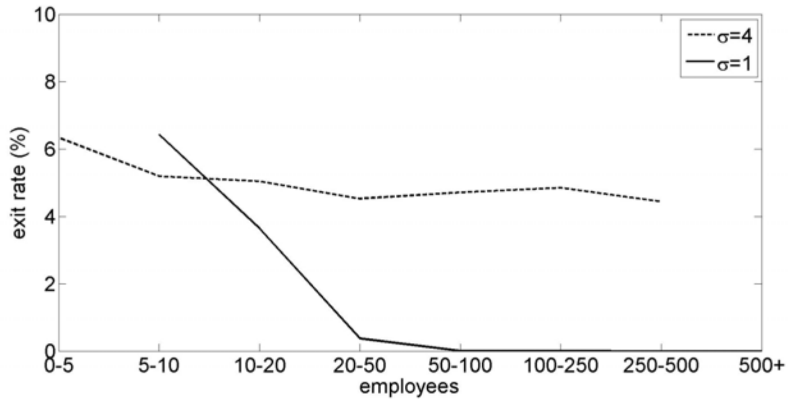

Note: The figure is based on simulating a model with only creative destruction as a source of growth. The exit rate is annualized from a model with 5-year periods. σ is the elasticity of substitution across varieties.

Own Innovation

We next consider a model where the *only* source of innovation is own innovation.[5] 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. 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. The model predicts that there are no firms that experience an employment decline of more than 25 percent. Those firms who innovate but draw small steps shrink more modestly.

[5]We assume $\kappa_i = \kappa_e = \delta_i = \delta_e = 0$ and no overhead costs. We set the parameter θ to match 1976–1986 TFP growth and λ_i to the value of δ_i in the creative-destruction only model.

Figure 3: Job Creation and Destruction with Own Innovation

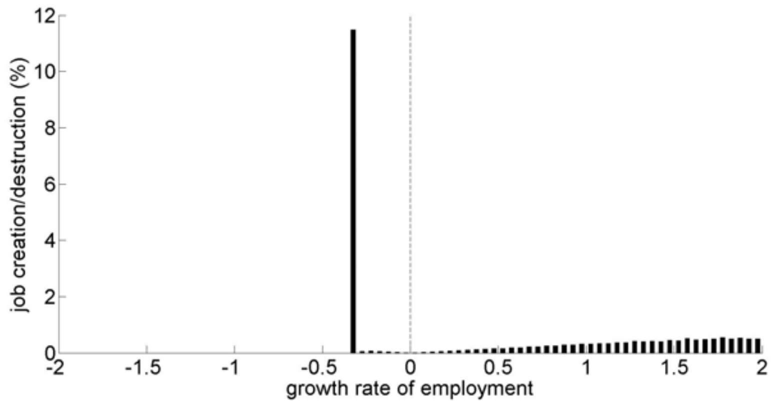

Note: The figure is based on simulating a model with only incumbent's improvement of their own products (own innovation) as a source of growth. Employment growth for a firm is defined as the change in employment divided by average employment at the firm across the two periods. The vertical axis gives the sum of total job creation (destruction), divided by aggregate employment in the initial period, associated with firms at each given level of employment growth. We used $\sigma = 4$ for the elasticity of substitution between varieties.

Own Innovation + Creative Destruction by Entrants

Clearly, what is driving the extreme empirical predictions of a model of own innovation is 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.[6] This hybrid model has the following implications. First, the employment share of entrants is positive and pinned down by the rate of creative destruction by entrants δ_e. Second, the tail of job creation is thin because there is no creative destruction by incumbent firms. The tail of job destruction is thin because job destruction is pinned down by the employment share of entrants in this model. Third, entrants are slightly larger than incumbents on average. This is because all entrants improve on incumbent quality

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

whereas only a subset of surviving incumbents improve their quality. Finally, since incumbent firms can only innovate by improving their one product in this hybrid model, 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 thicken the tail of job destruction It will also generate heterogeneity in the number of varieties across firms. Older firms will tend to have more varieties than young firms. Larger firms will also tend to have more varieties which will imply that they will have lower exit rates compared to small firms.

New Varieties

We now consider the effect of allowing firms to also create new varieties. A model where firms *only* create new varieties also has stark predictions that are not likely to hold empirically. Perhaps the most important is that, as in the model where growth is only driven by own innovation, there is no exit and the only source of job destruction are job losses due to the general equilibrium effect of a rising real wage among firms that do not create new varieties. 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 stationary distribution of varieties per firm. This makes the total number of varieties proportional to the total number of firms. We will therefore infer growth in the number of varieties equal to growth in the number of firms in the data.

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 hence firm size. If incumbents create these low quality new varieties, then they will increase the mass of job creation at lower values of employment growth. If instead entrants create these low quality new varieties, they will tend to make entrants smaller than incumbents (i.e., increase the slope of size by age).

Recap on Innovation and Job Flows

In short, we will use data on job flows to speak to the sources of innovation. Motivated by preceding discussion of how to discriminate between sources, the specific data moments we examine will be:

1. Aggregate TFP Growth
2. Standard Deviation of Employment across Firms
3. Job Creation and Job Destruction[7]
4. Job Creation due to Entry (the employment share of entrants)
5. Job Creation due to Firm Employment Growth ≤ 1.[8]
6. Employment by Age across Firms
7. Exit Rate by Firm Employment

The next section presents the data moments on the above list.

[7] Specifically, the aggregate rates of job creation and destruction. Fitting both implies fitting the difference between them, which is aggregate employment growth.

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

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* non-farm private establishment in the U.S. economy. The key advantages of the LBD are its broad coverage of the U.S. economy and its quality (e.g., the Census uses it to identify and correct for measurement error in its quinquennial Census surveys). The establishment-level variables we use are employment, the year the establishment appears in the LBD for the first time, 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 drop establishments in the public, educational, and agricultural & mining sectors, and restrict the sample to 1976–1986 and 2003–2013 (the first and last ten years of the LBD microdata).

We focus on firms rather than establishments because business stealing has policy implications for creative destruction vs. other sources of growth (e.g. Burstein and Atkeson, 2015). In our baseline sample, we aggregate establishment data of *all* establishments within a firm (using the firm identifier). Firm employment is the sum of employment at all the establishments owned by a firm. Firm age is defined as 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 within the next five years. We incorporate a five-year lag in light of adjustment costs that might keep an entrant's market share initially low relative to its product quality.[9]

Our firm employment dynamics include the direct positive and negative effects of mergers and acquisitions. Such M&A activity may also be the result of

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

Table 2: Summary Statistics in the LBD

	Employment	Firms	Average Employment	S.D. of log Employment	Employment Growth Rate
1981	78.9	3.307	23.9	1.26	2.54%
2008	125.6	5.134	24.5	1.28	0.85%

Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. Total employment and number of firms are in millions. The last column gives average annual growth of total employment from 1976 to 1986 and from 2003 to 2013, respectively.

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 job destruction based on this sample.

Table 2 presents some summary statistics for our LBD sample. Total employment and the number of firms increased from 1981 to 2008, but average employment per firm and the dispersion of firm size were fairly stable. The last column shows that the growth rate of aggregate employment was much lower from 2003–2013 than from 1976–1986. Since average employment per firm was roughly constant, the growth rate of the number of firms also fell.

Figure 4 plots the eleven-year rolling average of labor-augmenting TFP growth rates. The growth rate of aggregate TFP *increased* from 1.03% in the 1976–1986 period to 1.44% in the 2003–2013 period. These estimates are from the U.S. Bureau of Labor Statistics and cover the nonfarm business sector, like the LBD.

Next, Figure 5 presents the overall rate of job creation and destruction, where the rates are calculated over five years. The job creation rate labeled '1976–1986' is the average of the job creation rates from 1976 to 1981 and from 1981

Figure 4: U.S. TFP Growth

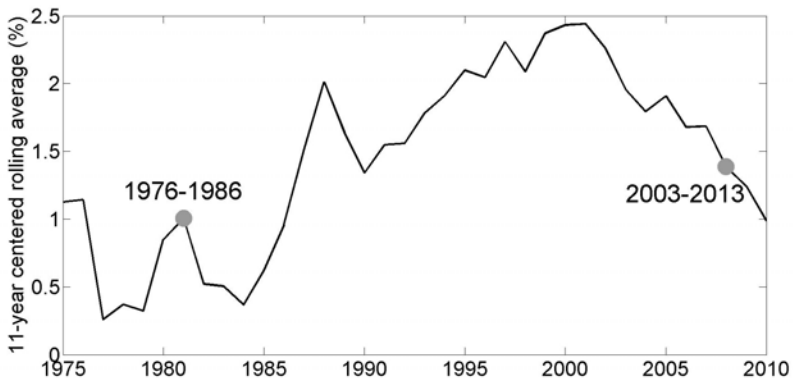

Note: Labor-augmenting TFP growth is from the U.S. Bureau of Labor Statistics. Each point plots the 11-year rolling average of the growth rate of TFP.

to 1986. The job creation (destruction) rate is the sum of employment changes at firms with (rising) falling employment divided by aggregate employment in the initial period. This includes entering and exiting firms. The job creation and destruction rates labeled '2003–2013' are defined analogously.

The overall job creation rate fell by 12 percentage points between 1976–1986 and 2003–2013, while the job destruction rate only fell by 3 percentage points. Most of the decline in the aggregate job reallocation rate highlighted by Decker et al. (2014) was due to the decline in the job creation rate. The decline in the job destruction rate was much smaller. Also, consistent with the fact in Table 2 that the growth rate of aggregate employment fell between 1976–1986 and 2003–2013, the gap between the rate of job creation and job destruction fell as well.

Aggregate job creation is the sum of job creation by incumbent firms and job creation by entering firms. Figure 6 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 '1976–1986' is the average of employment of entrants in 1981 as a share of total employment five years earlier

Figure 5: Job Creation and Destruction Rates

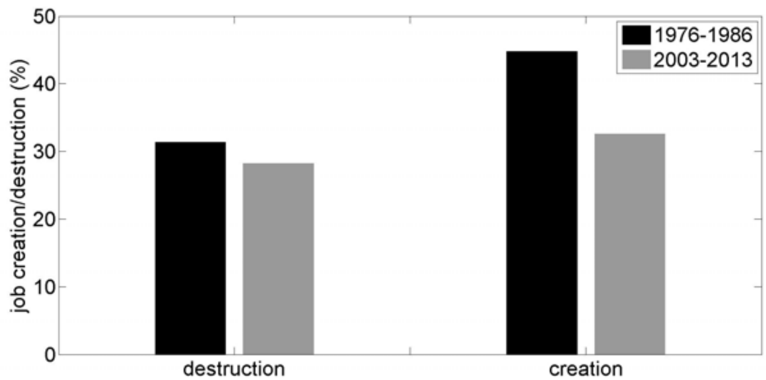

Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. The job creation (destruction) rate is the sum of employment changes at firms with rising (falling) employment divided by aggregate employment in the initial period. This includes entering and exiting firms. The Table shows the average 5-year changes for 1976–1981 and 1981–1986, and for 2003–2008 and 2008–2013.

(1976) and employment of entrants in 1986 as a share of total employment five years earlier (1981). Likewise, the label '2003–2013' refers to the average of the employment share of entrants in 2003–2008 and in 2008–2013.

Note that the employment share of entrants fell by 7 percentage points between 1976–1986 and 2003–2013. The aggregate job creation rate, shown previously in Figure 5, declined by 12 percentage points. So more than half of the decline in the aggregate job creation rate was due to the decline in the employment share of entrants.

Figure 7 plots the *distribution* of job creation and destruction rates in the LBD. As in Figures 5 and 6, we plot the average from 1976–1981 and 1981–1986 (labeled as '1976–1986') and the average 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 year. These rates are bounded between -2 (exit)

Figure 6: Employment Share of Entrants

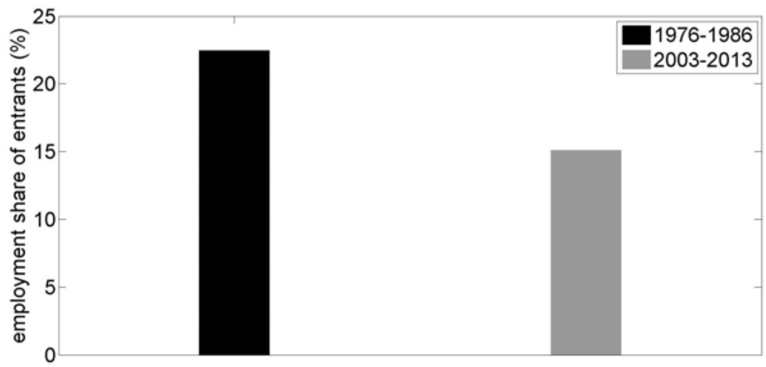

Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. The employment share of entrants for 1976–1986 is the average share of entrant employment in 1981 and 1986. Entrant employment in 1981 (1986) is defined as employment at firms that entered between 1976 and 1981 (1981 and 1986), and the share is defined relative to aggregate employment in 1976 and 1981. The entrant employment share for 2003-2013 is defined analogously.

and +2 (entry). The vertical axis shows the percent of all creation or destruction contributed by firms in each bin. For visual clarity, we omit job creation due to entry (+2) and job destruction due to exit (-2).

The empirical distribution of job creation and destruction in Figure 7 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, of course, 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 7 is the share of job creation at firms whose growth rate of employment is lower than 1, which is 35% and 39% in the two time periods.

We next present average employment of entrants and incumbents in 1981 and 2008 (Figure 8). According to Hsieh and Klenow (2014), rapid growth of

Figure 7: Distribution of Job Creation and Destruction

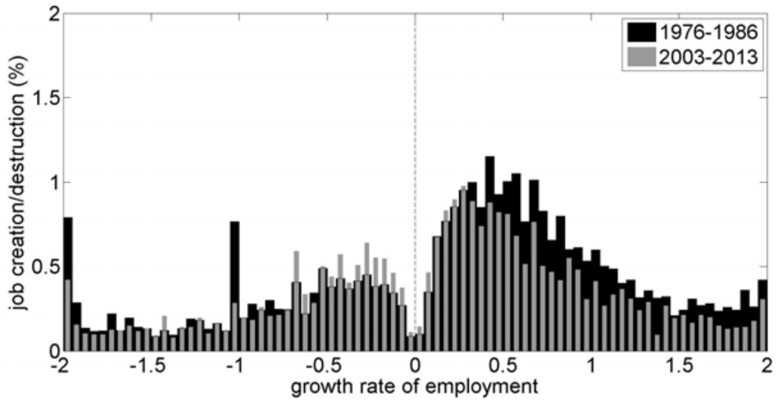

Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. Employment growth for a firm is defined as the change in firm employment over (say) 1981 to 1986 divided by the firm's average employment in 1981 and 1986. The vertical axis gives the sum total job creation (destruction), divided by average aggregate employment in 1981, associated with firms at each given level of employment growth. 1976–1986 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.

Figure 8: Employment per Firm, Young vs. Old

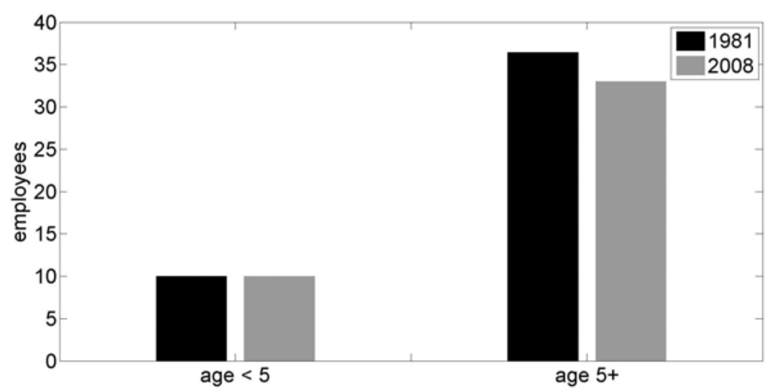

Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector.

Figure 9: Exit Rate, Large vs. Small Firms

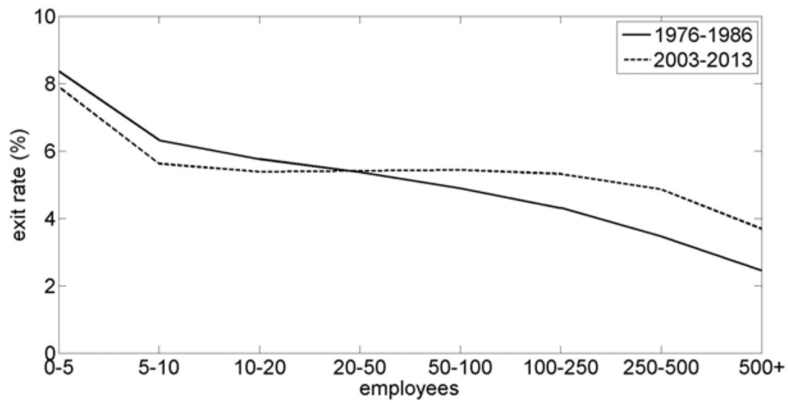

Source: U.S. Census Longitudinal Business Database (LBD) on firms in the nonfarm business sector. The exit rate is the annualized fraction of firms that operated in (say) 1976 but not in 1981. The data for 1976–1986 are averages of the exit rates from 1976 to 1981 and 1981 to 1986. The 2003–2013 exit rates are defined analogously.

surviving plants is a robust phenomenon in the U.S. Census of Manufacturing. Figure 8 suggests that the same is true for the entire U.S. private sector. The model can explain this fact if older firms have more products compared to young firms.

Figure 9 shows the exit rate of large vs. small firms. Here the exit rate is the annualized probability that the firm exits within the next *five* years. The label '1976–1986' refers to the average of the annualized exit rates from 1976–1981 and 1981–1986. The label '2003–2013' is the corresponding average of the exit rates from 2003–2008 and 2008–2013. Note that smaller firms have higher exit rates compared to larger firms. The model can explain this fact if larger firms produce more varieties compared to smaller firms. Note as well that, in contrast to the prediction of the polar model with only creative destruction, the decline in exit rates with firm size is gradual. A ten-fold increase in firm size is associated with only a one percentage point decrease in the exit rate.

Recall that we include mergers and acquisitions in our baseline statistics on

Table 3: Firm Dynamics dropping Mergers and Acquisitions (2003–2013)

	Full Sample	No Mergers & Acquisitions
Job Creation Rate	32.6%	27.7%
Employment Share of Entrants	15.1%	15.1%
Share of Job Creation < 1	39.3%	38.2%
Job Destruction Rate	28.2%	23.9%

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 then a factor of three in employment. We dropped mergers and acquisitions by excluding from the calculations plants whose firm changed between 2003 and 2008 (or 2008 and 2013).

job creation and job destruction. We now check the effect of dropping from the sample those establishments who change ownership. Table 3 presents the key moments of job creation and destruction for the 2003-2013 period in this new sample. For comparison, the first column presents the same moments for the full sample (where we keep establishments that switch ownership). The main change when we drop mergers and acquisitions is that the job creation and job destruction rates both drop by roughly 4 percentage points. The employment share of entrants is about the same in the new sample.

4. Sources of Growth

We now estimate parameters to match moments from model simulations to 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, the growth rate of aggregate employment, minimum employment per firm of 1, the cross-sectional standard deviation of log firm employment, the aggregate rates of job creation and destruction, the employment share of entrants, and the share of job creation < 1. So 8 moments.[10]

The overhead cost is chosen to make the minimum firm size one. The overhead cost pins down the cutoff quality threshold ψ and the endogenous rate at which existing varieties disappear due to obsolescence δ_o. 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 observed growth rate of aggregate employment, standard deviation of firm employment, aggregate rates of job creation and destruction, entrant share of employment, and share of job creation from employment growth less than 1. We do *not* use average employment by age or the exit rate by employment in our baseline estimates, though we provide robustness checks where we include these moments to estimate the model's parameters.

The Appendix 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 7 moments other than aggregate TFP growth. We weight moments equally because, given the large number of firms in the LBD, sampling error is a minor consideration for all of the moments.

Table 4 presents the parameter values inferred from the data using the procedure described above. Based on the data moments from 1976 to 1986, we infer a 71% arrival rate per period (five years) for own-variety quality improvements by incumbents. Conditional on no own-innovation, quality improvements through creative destruction occur 30% of the time by other incumbents. And conditional on no own-innovation and creative destruction by another in-

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

Table 4: Inferred Parameters Values

	1976–1986	2003–2013
Own-variety improvements by incumbents λ_i	71.0%	81.0%
Creative destruction by incumbents δ_i	30.0%	34.0%
Creative destruction by entrants δ_e	100.0%	100.0%
New varieties from incumbents κ_i	13.4%	4.3%
New varieties from entrants κ_e	0.0%	0.4%
Pareto shape of quality draws θ	22.5	16.1
Relative quality of new varieties s_κ	0.093	0.22
Cutoff quality relative to average quality ψ	0.068	0.057

cumbent, quality improvement through creative destruction by entrants occurs with probability one. The *unconditional* probability that quality of a given product improves due to creative destruction by an incumbent is thus 8.7%, and the corresponding probability of creative destruction by an entrant is 20.3%.[11] The unconditional probability a product is improved upon in a period is thus 100%, of which 71% is from own innovation and 29% from creative destruction.

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 $\theta = 22.5$ and $\sigma = 4$, the average improvement in quality (conditional on innovation) is 4.9%. New varieties are only created by incumbents, arrive with 13.4% probability per existing variety, and have an average quality that is only 9.3% of the average quality of existing varieties. Overhead costs imply that the cutoff quality threshold ψ is 6.8% of the average quality of existing varieties. This cutoff implies that the

[11] $(1 - 0.71) \cdot 0.3 = 0.087$ and $(1 - 0.71) \cdot (1 - 0.3) = 0.203$.

probability a variety exits due to overhead cost δ_o is essentially zero.[12] The net number of varieties thus grows by 13.4% every five years, which matches the growth of both total employment and the number of firms from 1976–1986.

Table 5 presents the sources of growth in the 1976 to 1986 period implied by the parameters in Table 4. The rows decompose aggregate TFP growth into the contribution of growth from creative destruction, new varieties, and own-variety improvements using equation (1).[13] About 27% of growth comes from creative destruction. Own-variety improvements by incumbents account for 65%. New varieties *a la* Romer (1990) are the remainder at around 8%.

The columns in Table 5 decompose aggregate TFP growth into the contribution of entrants vs. incumbents using equation (2). Incumbents account for 81% of aggregate TFP growth, with entrants contributing the remaining 19%. Aghion et al. (2014) provide complementary evidence for the importance of incumbents based on their share of R&D spending and patents.

Table 4 also presents parameter estimates for the 2003–2013 period. Recall the seven percentage point decline in the employment share of entrants between 1976–1986 and 2003–2013 in Figure 6. The model interprets the decline as reflecting less innovation by entrants. The estimates in Table 4 imply that the unconditional probability of creative destruction by an entrant fell from 20% to 12%. This swamped a minor uptick in new varieties from entrants.

Job creation by incumbents also fell by 5 percentage points, which the model interprets as a decline in the arrival rate of creative destruction by incumbents. According to Table 4, the probability a variety is creatively destroyed by an incumbent firm dropped from 8.4% to 6.5%.

Table 6 shows that the contribution of entrants to aggregate TFP growth fell to 12.8% in 2003–2013, down from 19.1% in the 1976–1986 period. The contribution of creative destruction by incumbent firms also fell, from 8.2% to 6.4%.

[12]The exact number is $\delta_o = 0.0078\%$.

[13]Equation (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 5: Sources of Growth, 1976–1986

	Entrants	Incumbents	
Creative destruction	19.1%	8.2%	27.3%
New varieties	0.0%	7.6%	7.6%
Own-variety improvements	-	65.1%	65.1%
	19.1%	80.9%	

Table 6: Sources of Growth, 2003–2013

	Entrants	Incumbents	
Creative destruction	12.5%	6.4%	18.9%
New varieties	0.3%	4.1%	4.4%
Own-variety improvements	-	76.7%	76.7%
	12.8%	87.2%	

This is the model's proximate answer to the question of how much the decline in job creation matters for aggregate TFP growth.

Recall that the growth rate of aggregate TFP actually *increased* from 1.03% a year during 1976–1986 to 1.44% a year during the 2003–2013 period. How could the growth rate of aggregate TFP increase when innovation by entrants and creative destruction by incumbents fell? The model's answer is that own innovation by incumbents must have increased and more than offset lower innovation by entrants and less creative destruction by incumbents. The arrival rate of quality improvement via own innovation by incumbent firms increased from 71% to 81%, and the share of TFP growth due to this channel increased from 65% to 77%.

As noted, our estimation procedure chose parameter values such that the

Table 7: Model Fit , 1976–1986

	Data	Model
Employment share of entrants	22.5%	21.5%
Employment growth rate	2.5%	2.5%
Job creation rate	44.8%	40.0%
Job destruction rate	31.4%	26.6%
Share of job creation < 1	34.6%	28.9%
SD(log employment)	1.26	1.18

model exactly matched aggregate TFP growth. Table 7 shows the fit of the model in 1976–1986 for the data moments that we used for the estimation but did not force the model to fit exactly.[14] The model comes close on the entrant employment share but understates the aggregate rates of job creation and destruction, the share of job creation where employment growth is less than 1, and the standard deviation of log employment. Although we used 8 parameters to try to fit 8 moments, we did not succeed more exactly because the model moments are nonlinear functions of the parameters. Table 10 in the Appendix provides the model fit for 2003–2013, which is similar.

For these baseline estimates, we did not target exit by firm size in the data. This moment does depend, however, on the sources of innovation. The exit rate does not fall so fast with firm employment under own innovation as with creative destruction by incumbents. Figure 10 presents the relationship between firm exit and employment predicted by the parameters in Table 4 for the 1976–1986 period. The model's exit rate rate is below that in the data, and the model's exit rate does not fall as fast with employment as it does in the data. The model

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

Figure 10: Exit by Employment, Model vs. Data (1976–1986)

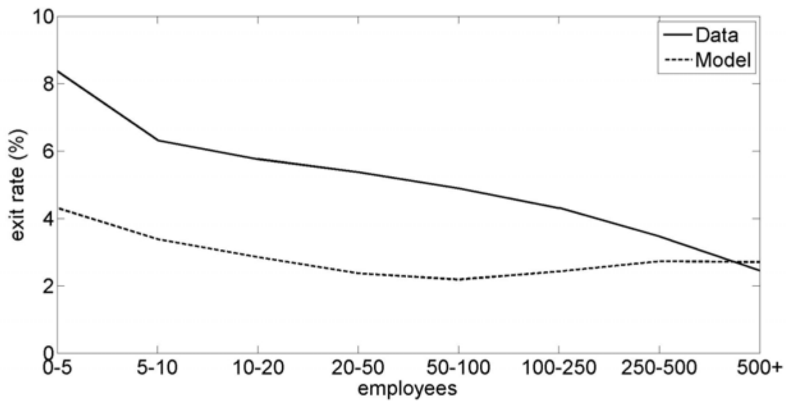

generates too little heterogeneity in the number of varieties across firms to fit this fact. A higher rate of creative destruction by incumbents would increase the exit rate and its decline with firm size, but this parameter is constrained by the need to fit the share of job creation driven by firms 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 relative to the data.)

Another moment we did not use is how average employment varies with age. In the data, the ratio of average employment of incumbent firms (age $\geq$ 5) to that of entrants (age $\leq$ 4) is 3.6. In our simulated model, the ratio is 1.6. So although the model generates rising employment with age, the magnitude is smaller than what we observe in the data. To generate steeper employment by age, we need more creative destruction by incumbents or lower quality new varieties created by entrants. The quality of new varieties by entrants is constrained by size dispersion. The rate of creative destruction by incumbents is constrained by also trying to fit job creation patterns. More creative destruction will *lower* the share of job creation around small changes in employment, which is already too low in the model.

We end this section with a series of robustness checks. We first add aver-

Table 8: Robustness Checks: Sources of Growth, 2003–2013

	Baseline	Emp. by age	Exit by emp.	No M&A
Creative destruction	18.9%	23.2%	17.1%	15.8%
New varieties	4.4%	2.7%	3.4%	5.6%
Own-variety improvements	76.7%	74.1%	79.5%	78.6%

Note: The column "Emp. by age" includes growth in employment by age as a target moment. "Exit by emp." includes exit by employment bin as a target moment. The column "No M&A" targets data excluding M&A job flows.

age employment of incumbents vs. entrants (Figure 8) as an additional target moment. The resulting parameter estimates, in the second column in Table 8, show that the contribution of creative destruction to TFP growth is about 4 percentage points higher than the baseline.[15] We reproduce the baseline estimates in the first column of Table 8 for comparison. The baseline model understates employment by age because δ_i is lower than what is needed to match employment by age in the data. When we also try to fit employment by age, we get a larger estimate of δ_i and thus a larger role for creative destruction.

We next add exit by size (employment) as an additional data target (but no longer target employment by age). These estimates are shown in the third column of Table 8 and are similar to the baseline estimates.

Finally, we estimate the model using the data moments from the sample where we drop establishments that undergo ownership changes. We revert to targeting the baseline data moments, except now calculated on the sample that excludes mergers and acquisitions. This assumes mergers and acquisitions have no effect on TFP growth. The effect, shown in the last column of Table 8, is to lower the contribution of creative destruction and slightly increase the contri-

[15]The Appendix presents the estimated parameters and the fit for the models in Table 8.

bution of own innovation to TFP growth. Since job creation and destruction rates are lower when we drop mergers and acquisitions (Table 3), we infer a lower rate of creative destruction (and a higher rate of own innovation).[16]

5. Conclusion

How much innovation takes the form of creative destruction versus firms improving their own products versus new varieties? How much innovation occurs through entrants versus incumbents? We try to infer the sources of innovation from the employment dynamics of U.S. firms in the non-farm private sector from 1976–1986 and 2003–2013. 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.

Our findings could be relevant for innovation policy because the sources of growth we identify have different business stealing effects versus knowledge spillovers. The importance of 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 the employment dislocations that come along with some growth.

It would be interesting to extend our analysis to individual sectors, other time periods, and countries. Retail trade experienced a big-box revolution in the U.S. led by Wal-Mart's expansion. Online retailing has made inroads at the expense of brick-and-mortar stores. In Chinese manufacturing private enterprises have entered and expanded while state-owned enterprises have closed

[16]We include mergers and acquisitions in the baseline estimates for two reasons. First, because job reallocation associated with such activity may be a byproduct of innovation. Second, as we are finding a smaller role for creative destruction than the existing literature, this is a conservative assumption.

(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.

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. Or, 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 (2015) provide evidence for this hypothesis from patent citations.

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 et al. (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.

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 these moments:

   (a) TFP growth rate

   (b) Aggregate employment growth rate

   (c) Standard deviation of log firm employment

   (d) Employment share of entrants

   (e) Job creation rate

   (f) Job destruction rate

   (g) Share of job creation where employment growth ≤ 1

   (h) Minimum firm employment

   In the simulations for Table 8, calculate two additional data moments:

   (a) Exit by Employment

   (b) Employment of Entrants and Incumbents

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 Robustness Estimates

Table 9 presents the inferred parameter values and Tables 10 and 11 the model fit for the robustness estimates shown in Table 8.

Table 9: Inferred Parameters Values, 2003–2013

	Baseline	Emp. by age	Exit by emp.	No M&A
OI by incumbents λ_i	81.0%	77.0%	83.0%	84.0%
CD by incumbents δ_i	34.0%	58.0%	4.5%	16.5%
CD by entrants δ_e	100.0%	100.0%	100.0%	100.0%
NV from incumbents κ_i	4.3%	0.0%	0.0%	4.3%
NV from entrants κ_e	0.4%	5.1%	4.8%	0.6%
Pareto scale θ	16.1	15.9	16.0	16.2
Relative quality NV s_κ	0.22	0.13	0.17	0.27

Note: “Emp. by age” and “Exit by emp.” include growth in employment by age and exit by employment, respectively, as target moments. ‘’‘No M&A” targets data excluding M&A job flows.

Table 10: Model Fit, 2003–2013

	Data	Baseline	Emp. by age	Exit by emp.
Employment share of entrants	15.1%	13.6%	12.1%	17.4%
Employment growth rate	0.85%	0.85%	0.89%	0.85%
Job creation rate	32.6%	27.0%	29.7%	25.8%
Job destruction rate	28.2%	22.7%	25.1%	21.5%
Share of job creation < 1	39.3%	30.5%	29.3%	29.5%
SD(log employment)	1.28	1.26	1.31	1.26
Incumbent emp. / Entrant emp.	3.30	1.50	2.08	1.40
Exit rate by employment bin	5.5%	2.4%	2.6%	3.3%

Note: Moments targeted but not fit exactly. The estimation also used TFP growth rates, which the model fits exactly. "Emp. by age" includes growth in employment by age as a target moment. "Exit by emp." includes exit by employment as a target.

Table 11: Model Fit, 2003–2013 excluding M&A

	Data	Model
Employment share of entrants	15.1%	14.2%
Employment growth rate	0.85%	0.85%
Job creation rate	27.7%	24.6%
Job destruction rate	23.9%	20.3%
Share of job creation < 1	38.2%	32.2%
SD(log employment)	1.28	1.26
Incumbent emp. / Entrant emp.	3.30	1.42
Exit rate by employment bin	5.5%	2.4%

Note: Table presents predicted values for the moments used for estimation that were not forced to exactly fit the data. The estimation also used TFP growth rates, which we force the model to match exactly.

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 2013.

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," 2015.

Baily, Martin Neil, Charles Hulten, and David Campbell, "Productivity Dynamics in Manufacturing Plants," *Brookings Papers: Microeconomics*, 1992, *4*, 187–267.

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.

Burstein, Ariel and Andrew Atkeson, "Aggregate Implications of Innovation Policy," in "2015 Meeting Papers" number 640 Society for Economic Dynamics 2015.

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*, 1998.

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.

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.

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, Stephen J Redding, and David E Weinstein, "Quantifying the Sources of Firm Heterogeneity," *forthcoming in The Quarterly Journal of Economics*, 2016.

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,” 2013.

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.

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.