How Destructive is Innovation?

Daniel Garcia-Macia
Stanford University

Chang-Tai Hsieh
University of Chicago and NBER

Peter J. Klenow*
Stanford University and NBER

July 17, 2015

Abstract

Entering and incumbent plants can create new products and displace existing products. Incumbents can also improve their existing products. How much of aggregate growth occurs through each of these channels? Using U.S. Census data on manufacturing plants from 1992, 1997 and 2002, we arrive at three main conclusions: First, most growth appears to come from incumbents. We infer this from the modest employment share of entering plants. Second, most growth seems occur through improvements of existing varieties rather than creation of brand new varieties. We infer this because of modest net entry of plants and gently falling exit rates as plants expand (the latter suggesting bigger plants produce better products more than a wider array of products). Third, own-product improvements by incumbents appear to have been more important than creative destruction. We infer this because the distribution of job creation and destruction has thinner tails than if growth mostly came from creative destruction.

*We thank Ufuk Akcigit, Andy Atkeson, Michael Peters, and Paul Romer for very helpful comments. For financial support, Hsieh is grateful to Chicago's Initiative for Global Markets (IGM) 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. 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 those 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) provide a recent survey.

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 and quality improvement 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), who also provide evidence that firms are more likely to cite their own patents and hence build on them.

Still other theories emphasize the contribution of brand new varieties to growth. Romer (1990) is the classic reference, and Acemoglu (2003) and Jones (2014) are some of the many follow-ups. Studies such as Howitt (1999) and Young (1998) combine variety growth with quality growth.

Ideally, one could directly observe the extent to which new products substitute for or improve upon existing products. Broda and Weinstein (2010) and Hottman et al. (2014) are important efforts along these lines for consumer nondurable goods. Such high quality scanner data has not been available or analyzed in the same way for consumer durables, producer intermediates, or producer capital goods – all of which figure prominently in theories of growth.[1]

We pursue a complementary approach. We try to infer the sources of growth indirectly from empirical patterns of firm and plant dynamics. The influential papers by Baily et al. (1992) and Foster et al. (2001) document the contributions of entry, exit, reallocation, and within-plant productivity growth to over-

[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.

all growth with minimal model assumptions. We consider a specific growth model with a limited set of parameters. Like us, Lentz and Mortensen (2008) and Acemoglu et al. (2013) conduct indirect inference on growth models with manufacturing data (from the U.S. and Denmark, respectively). They fully endogenize growth, whereas we consider exogenous growth models. The trade-off is that they focus on creative destruction, whereas we further incorporate new varieties and own-variety improvements by incumbents.

We use data on plants from U.S. manufacturing censuses in 1992, 1997 and 2002. Over this period, we calculate aggregate TFP growth, the exit rate of plants by age and employment, employment by age, job creation and destruction rates across plants, growth in the total number of plants, and moments of the employment distribution across plants (the min, median, and mean). The parameter values which best fit these moments lead to three conclusions. 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 rest of the paper proceeds as follows. Section 2 lays out the parsimonious exogenous growth model we use. Section 3 briefly describes the U.S. manufacturing census data we exploit. Section 4 presents the model parameter values which best match the moments from the data. Section 5 concludes.

2. An Exogenous Growth Model

We adapt the Klette and Kortum (2004) model of quality ladder growth through creative destruction. Firms produce multiple varieties and grow when they improve upon and capture the varieties produced by other firms. Entrants try to improve on existing varieties and take them over in the process. Incumbent firms die when their last varieties are captured by other firms (incumbents or

entrants). Unlike Klette-Kortum we treat the arrival rates of creative destruction from entrants and incumbents as exogenously fixed parameters, rather than being endogenously determined by underlying preferences, technology, and market structure. This allows us to keep the model parsimonious while adding exogenous arrival rates of new varieties from entrants, new varieties from incumbents, and own-variety quality innovations by incumbents.

Our set-up follows Klette-Kortum with these differences:

- Time is discrete (rather than continuous)
- There are a finite number of varieties (rather than a continuum)
- Innovation is exogenous (rather than endogenous)
- Demand for varieties is CES with elasticity $\sigma \geq 1$ (rather than $\sigma = 1$)
- The number of varieties is growing (rather than fixed)
- Incumbents can improve the quality of their own varieties
- Creative destruction by incumbents may be directed toward quality levels near the firm's existing quality levels (rather than being undirected)
- Creative destruction by entrants may be directed toward lower quality levels (rather than being undirected)

Aggregate output

Total output Y in the economy is given by:

$$Y_t = \left[\sum_{j=1}^{M_t} \left(q_{j,t} y_{j,t} \right)^{1-1/\sigma} \right]^{\frac{\sigma}{\sigma-1}}$$

where M is the total number of varieties, q_j is the quality of variety j, and y_j is the quantity produced of variety j. The production function for each variety is linear in labor $y_j = l_j$.

Static problem of the firm

Firms control multiple varieties, but we assume they are still monopolistic competitors for each variety. We assume further that there is an overhead cost of production that must be expended before choosing prices and output. This assumption allows the highest quality producer to charge the standard markup over marginal cost, as the next lowest quality competitor will be deterred by zero ex post profits. Without this assumption, firms would engage in limit pricing and markups would be heterogeneous as in Peters (2013).

Assuming firms face the same wage, revenue generated by variety j is

$$p_j y_j = \left(\frac{\sigma - 1}{\sigma}\right)^{\sigma-1} P^{\sigma} Y W^{1-\sigma} q_j^{\sigma-1} \propto q_j^{\sigma-1}$$

where P is the aggregate price level, Y is aggregate output, and W is the wage. Labor employed in producing variety j is also proportional to $q_j^{\sigma-1}$:

$$l_j = \left(\frac{\sigma - 1}{\sigma}\right)^{\sigma} \left(\frac{P}{W}\right)^{\sigma} Y q_j^{\sigma-1} \propto q_j^{\sigma-1}$$

Thus a firm's revenue and employment are both proportional to the sum of power qualities $q_j^{\sigma-1}$ of the varieties operated by the firm.[2] In the special case of $\sigma = 1$ assumed by Klette-Kortum, all varieties have the same employment, 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.

Aggregate productivity

Labor productivity in the economy is given by

$$Y_t/L = M_t^{\frac{1}{\sigma-1}} \left[\sum_{j=1}^{M_t} \frac{q_{j,t}^{\sigma-1}}{M_t}\right]^{\frac{1}{\sigma-1}}.$$

[2]There is no misallocation of labor whatsoever in this model.

where L is total labor across all varieties (which is exogenously fixed in supply). The first term captures the benefit of having more varieties, and the second term is the power mean of quality across varieties.

Exogenous innovation

There is an exogenous arrival rate for each type of innovation. The notation is given in Table 1. The probabilities are per variety a firm produces. The probability of a firm improving any given variety it produces is λ_i, and such improvement is associated with multiplicative step size $s_q \geq 1$. If a firm fails to improve on a given variety it produces, then that variety is vulnerable to creative destruction by other (incumbent or entrant) firms. A fraction δ_i of vulnerable varieties is creatively destroyed by another incumbent, and a fraction δ_e by an entrant. Creative destruction also comes with step size s_q.

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 step size for quality improvements is $s_q \geq 1$. New varieties are drawn from a Pareto distribution with shape θ.

Brand new varieties arrive at rate κ_e from entrants and at rate κ_i from incumbents – again per existing variety produced by an incumbent. These arrivals are independent of other innovation types. The quality of each new variety is drawn from a Pareto distribution with shape parameter θ. The arrival rate of brand new varieties affects growth in the number of firms (tied to κ_e and κ_i), while the quality of new varieties (tied to θ) affects the size of new firms.

On top of the seven parameters listed in Table 1, we add three more parameters. Klette-Kortum assumed creative destruction was undirected. When combined with $\sigma > 1$, undirected creative destruction leads to a thick-tailed distribution of employment growth rates. Firms can capture much better varieties than their own, growing rapidly in the process. Incumbents can lose their best varieties, leaving them with low quality varieties and steeply negative growth. To allow some control over the tails of employment growth, we allow for the possibility that creative destruction is directed. We set ρ_i to be quality quantiles in which incumbent creative destruction is directed.[3] If $\rho_i = 0.1$ then incumbent creative destruction occurs within quality deciles. If $\rho_i = 1$ then incumbent creative destruction is undirected.

For entrants, we assume creative destruction targets the lowest quality varieties with collective employment per variety equal to ρ_e of average employment per variety across all varieties. Defining ρ_e in this way preserves an analytical expression for aggregate productivity growth.

The last parameter is overhead labor, which pins down the minimum firm size. We set overhead labor to $1/\sigma$ so that the minimum size firm has 1 unit of labor for production plus overhead. The overhead cost endogenously determines the scale parameter for the Pareto distribution of quality for new varieties, as lower quality draws do not produce. The scale parameter rises with wage growth. We denote the scale parameter relative to the CES power mean of quality as ψ. We denote the exit rate of existing varieties due to overhead as δ_o. Net growth of varieties is therefore $\kappa_e + \kappa_i - \delta_o$.

Note that each innovation is proportional to an existing quality level. 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).

[3] As all arrival rates are per existing variety, the quantiles are for each individual variety.

Productivity growth

Total output grows at rate

$$1 + g_Y = \left[\left(1 + \kappa_e + \kappa_i - \delta_o\right)\left(1 + g_q\right)\right]^{\frac{1}{\sigma-1}}$$

The κ components correspond to the creation of new varieties. The $(1 + g_q)$ component reflects growth in average quality per variety, which is:

$$1 + g_q = \frac{\left(\left(\frac{\theta}{\theta-1}\right)^{\sigma-1}\left(\kappa_e + \kappa_i\right) - \delta_o\right)\psi + 1 + \left(s_q^{\sigma-1} - 1\right)\left(\lambda_i + \left(1 - \lambda_i\right)\left(\rho_e\delta_e + \delta_i\right)\right)}{1 + \kappa_e + \kappa_i - \delta_o}$$

3. U.S. Manufacturing Census Data

We use data from U.S. manufacturing Censuses to quantify dynamics of entry, exit, and survivor growth. We mostly use the 1992, 1997 and 2002 Censuses, stopping at 2002 because the NAICS definitions were changed from 2002 to 2007. We focus on *plants* rather than firms, because mergers and acquisitions can wreak havoc with our strategy to infer innovation from growth dynamics.

We are ultimately interested in decomposing the sources of TFP growth into contributions from different types of innovation. We therefore start by calculating manufacturing-wide TFP growth from the U.S. Bureau of Labor Statistics. Converting their gross output measure to value added, TFP growth averages 3.28 percent per year from 1987–2011 (the timespan of their data). The number of manufacturing plants in the U.S. Census of Manufacturing, meanwhile, rose 0.5 percent per year from 1972 to 2002.

Since the Census does not ask about a plant's age directly, we infer it from the first year a plant shows up in the Census going back to 1963. We therefore have more complete data on the age of plants in more recent years. We use this data to calculate exit by age from 1992 to 1997. We then combine it with the assumption of 0.5 percent per year growth in the number of entering plants to calculate the share of plants by age brackets of less than 5 years old, 5 to 9 years

old, 10 to 14 years old, and so on until age 30 years and above. See the first column of Table 2 for the resulting density. About one-third of plants are less than 5 years old, and about one-eighth of plants are 30 years or older.

Table 2: Plants by Age

Age	Fraction	Employment Share
< 5	.358	.124
5-9	.189	.115
10-14	.128	.102
15-19	.091	.088
20-24	.068	.081
25-29	.046	.073
≥ 30	.120	.418

Note: Author calculations from U.S. Census of Manufacturing plants in 1992 and 1997.

We report the share of *employment* by age in U.S. manufacturing in the second column of Table 2. Young plants are much smaller on average, as their employment share (12 percent) is much lower than their fraction of plants (36 percent). Older surviving plants are much larger, comprising only 12 percent of plants but employing almost 42 percent of all workers in U.S. manufacturing. According to Hsieh and Klenow (2014), rapid growth of surviving plants is a robust phenomenon across years in the U.S. Census of Manufacturing.

We plot how a plant's exit rate varies with its size in Figure 1. The exit rate is annualized based on successive years of the Census. The dots labeled "1992" are based on exit from 1992 to 1997, those labeled "1982" are based on exit from 1982 to 1987, and so on back to "1963". As shown, the annual exit rate is about 10 percent for plants with a single employee, declines to about 6 percent for plants with several hundred employees, then falls further to about 1 percent for plants with thousands of workers.

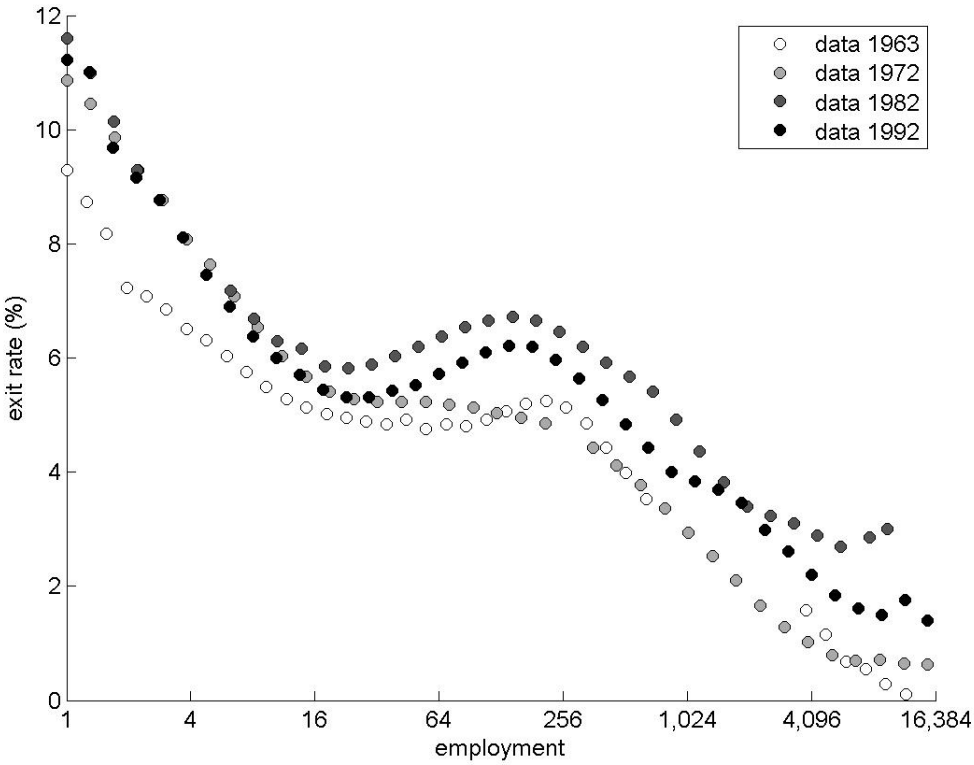

Figure 1: Empirical Exit by Size in the U.S. Census of Manufacturing

Figure 2, from Davis et al. (1998), plots the distribution of annual job creation and destruction rates in U.S. manufacturing from 1973–1988. These rates are bounded between -2 (exit) and +2 (entry) because they are the change in employment divided by the average of last year's employment and current year's employment. The distribution on the vertical axis is the percent of all creation or destruction contributed by plants in each bin. We also target the absolute rates of job creation (11.1 percent) and job destruction (11.0 percent), and the absolute rates of job creation from entrants (2.8 percent) and job destruction from exiters (3.4 percent).

Lastly, we target several moments of the employment distribution across U.S. manufacturing plants: the minimum (1 employee), the mean (51 employees), and the plant size of the median worker (900).

4. Indirect Inference

We now choose parameter values to match moments from model simulations to the U.S. manufacturing moments. Our aim is to indirectly infer the sources of innovation from plant dynamics. Entrants reflect a combination of new varieties and creative destruction of existing varieties. The more entrants innovate, the bigger the employment share of entrants. When an existing plant expands, it does so because it has innovated on its own varieties, created new varieties, or captured varieties previously produced by other incumbents. When a plant contracts it is because it has failed to improve its products or add products to keep up with aggregate growth, or because it lost some of its varieties to creative destruction from entrants or other incumbents. Outright exit occurs, as in Klette and Kortum (2004), when a plant loses all of its varieties.

Because creative destruction is independent across a plant's varieties (by assumption), plants with more varieties should have lower exit rates. Plants with higher qualities should be larger and potentially more protected against exit, as they have a lower likelihood of being captured by entrants who target

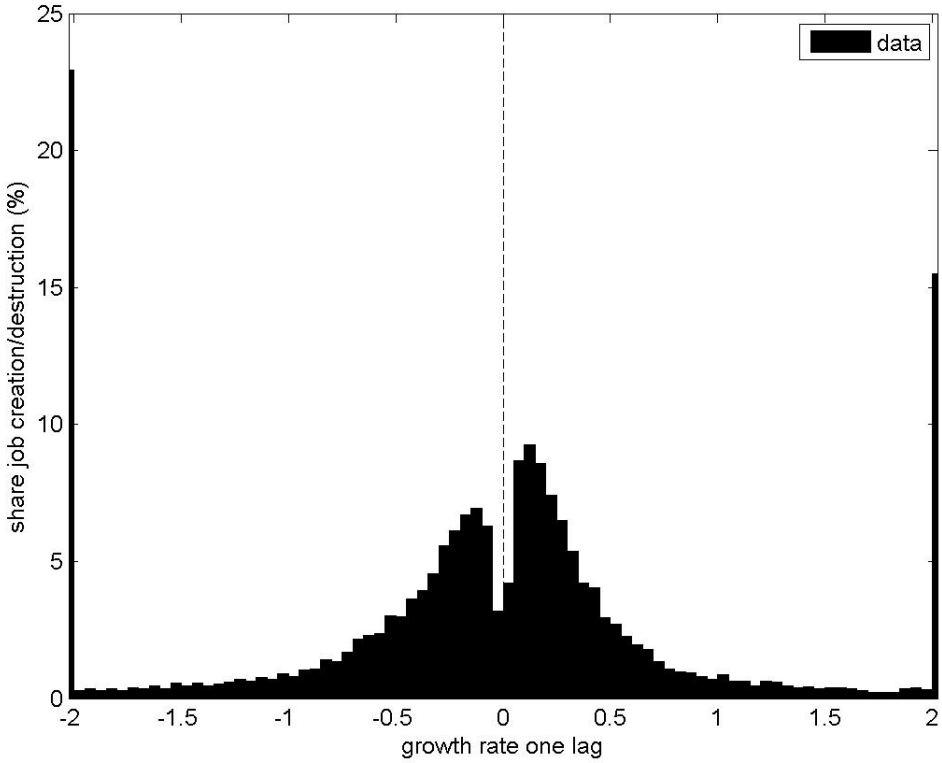

Figure 2: Job Creation and Destruction Rates in U.S. Manufacturing (via Davis, Haltiwanger and Schuh, 1998)

the bottom of the quality distribution.

Simulation algorithm

Each firm's static price and employment can be solved analytically, but we have to numerically compute the firm-level quality distribution. Compared to the Klette and Kortum (2004) environment, our additional channels of innovation, as well as the elastic demand for quality ($\sigma > 1$), preclude analytical results. Our simulation algorithm consists of the following steps:

1. Specify an initial guess for the distribution of quality across varieties.

2. Simulate life paths for entering firms such that the total number of firms equals the number of plants in the U.S. data over 1992–2002 (321,000 on average, including administrative record plants, and growing 0.5 percent per year).

3. Each entrant has one initial variety, captured from an incumbent or newly created. In every year of its lifetime, it faces a probability of each type of innovation per variety it owns, as in Table 1. In every year, each of an incumbent's varieties can be captured by another firm (entrant or incumbent), it can be improved by its owner, or it can stay with the same quality. If $\rho_e < 1$, only lower-quality varieties are subject to creative destruction by entrants. Moreover, each variety can lead to the addition of an extra variety captured from an incumbent and/or to a newly created variety. Therefore, firms can grow by improving the quality of their own varieties, capturing varieties from other firms, and adding brand new varieties. A firm's life ends when it loses all of its varieties to other firms or when it reaches age 200.

4. Based on an entire population of simulated firms, calculate moments of interest (e.g. exit by employment across firms). Simulating life paths for many entering firms produces a joint distribution of firm age, number of

varieties, quality of each variety, employment and exit probability. Take into account growth in average quality over time when computing relative qualities of varieties for a firm at different ages. Set aggregate employment in each year so that employment per firm in the simulation is always 51, to match average employment per plant in the U.S. data from 1992–2002.[4]

5. Repeat steps 1 to 4 until all moments converge. Step 1 takes the quality distribution across varieties from step 4 as the starting point. The converged moments imply a stationary joint distribution of all the variables of interest (age, varieties, quality, employment and exit) across firms.

6. Repeat steps 1 to 5, searching for parameter values to minimize the absolute distance between the simulated and empirical moments.

Sources of growth

We present our inferred parameter values in Table 3. We will first discuss the inferred parameter values and the implications of these values for the sources of growth. We will then examine why the data fitting exercise yields the parameters it does by shutting down each source of innovation.

We infer a 45 percent arrival rate for own-variety quality improvements by incumbents. Conditional on no own-innovation, quality improvements through creative destruction occur 15.5 percent of the time by other incumbents, and 10 percent of the time by entrants. The step size for all quality improvements is 5.8 percent.[5] Incumbent creative destruction is estimated to be within deciles ($\rho_i = 0.10$), and entrant creative destruction is targeted toward the lowest quality varieties whose collective employment per variety is 41 percent of average employment per variety ($\rho_e = 0.41$). The latter lowers entrant quality relative to

[4]Given the number of firms in the model is set to grow 0.5 percent per year, aggregate employment in the model must also grow at 0.5 percent per year to obtain a fixed average employment per plant across simulation years.

[5]We can allow a different step size for own innovation than for creative destruction. When we do so, the step size for creative destruction becomes very small.

Table 3: Inferred Parameter Values

Channel	Probability
Own-variety improvements by incumbents λ_i	45.0%
Creative destruction by incumbents δ_i	15.5%
Creative destruction by entrants δ_e	10.0%
New varieties from incumbents κ_i	0.0%
New varieties from entrants κ_e	2.0%
Step size for quality improvements s_q	1.058
Pareto shape θ	18
Pareto scale ψ	0.035
Directedness of $\delta_i \to \rho_i$	0.10
Directedness of $\delta_e \to \rho_e$	0.41
Overhead-related exit δ_o	1.5%

Note: The Pareto scale parameter is relative to average quality.

average incumbent quality, and raises the exit rate for smaller firms (which tend to have lower quality). New varieties arrive at a 2 percent annual rate and come entirely from entrants ($\kappa_e = 0.02$ and $\kappa_i = 0$). Overhead costs kill off 1.5 percent of varieties every year ($\delta_o = 0.015$) so that net growth in varieties matches the 0.5 percent per year growth in the number of plants. The Pareto quality draws for new varieties are set to match the minimum plant size (1 worker) and the plant size of the median worker (900) in the data.[6] The estimated shape parameter is 18 and the scale parameter (relative to average quality) is 0.035.

Table 4 presents the implied sources of growth. About 19 percent of growth comes from creative destruction. Own-variety improvements by incumbents account for 80 percent. Net new varieties *a la* Romer (1990) are the remainder

[6]Since we simulate the same number of plants as in the data, we fit average employment per plant of 51 by construction.

Table 4: Inferred Sources of Growth

	Entrants	Incumbents	
Creative destruction	4.1%	15.3%	19.4%
Net new varieties	0.3%	0.0%	0.3%
Own-variety improvements	-	80.2%	80.2%
	4.4%	95.5%	

at less than 1 percent.[7] Incumbents are the dominant source of growth (over 95 percent), with entrants contributing under 5 percent. Aghion et al. (2014) provide complementary evidence for the importance of incumbents based on their share of R&D spending.

At this point, a few key questions arise: What empirical moments suggest the presence of own-variety improvements? And how well does a model with only creative destruction fit the data? To help answer these questions, we examine a sequence of models as listed in Table 5. We start with the baseline Klette-Kortum model, which features $\sigma = 1$ and only creative destruction. Then we generalize the Klette-Kortum model to $\sigma > 1$, so that sales and employment are higher for high quality varieties. We next allow for directed creative destruction. Then we add the creation of brand new varieties by entrants and incumbents. Finally, we add own-variety improvements by incumbents.

Table 6 reports the parameter values we infer for each model. By adding layers to the Klette-Kortum model, we progressively achieve a better fit with the data. The models from KK3 onwards have entry of new varieties to replace low-quality varieties exiting due to overhead costs.[8] Only the last two models have positive net variety growth.

[7] Less than 0.1 percent of growth comes from the endogenous exit of low quality varieties.

[8] Overhead costs would not generate exit of varieties in KK1, since all varieties generate the same profits when $\sigma = 1$.

Table 5: Simulated Models

	Klette-Kortum KK1	Klette-Kortum KK3	Directed Creative Destr.	New Varieties	Own Innov.
σ	1	3	3	3	3
Creative destruction	√	√	√	√	√
Directed Creative Destruction			√	√	√
New varieties				√	√
Own-variety improvements					√

To provide some intuition, consider the baseline Klette-Kortum model. In this model, the exit and job destruction rates are pinned down by the rates of creative destruction from entrants and incumbents. The exit rate of the smallest firm in the baseline Klette-Kortum model is simply the probability that the single variety owned by the smallest firm is improved upon by another firm. The exit rate of a one-variety firm is $(\delta_e + \delta_i)(1 - \delta_i)$, i.e., the sum of the probability of creative destruction by an entrant or an incumbent firm, times the probability that the firm does not creatively destroy a variety from another incumbent. Around 11.5 percent of the varieties are subject to creative destruction each year, consistent with the empirical exit rate and overall job destruction rate. A large step size (26 percent) is therefore needed to match the annual growth rate of 3.3 percent. This pattern holds for all of the models without own-variety quality improvements.

Figure 3 plots the exit rate of plants by age in each model vs. the data. The fit is similar for all models. Figure 4 plots the size (average number of employees) of plants by age. Only models with directed innovation by entrants are able to replicate the amount of growth in size by age seen in the data. As shown in Table 6, entrants target the lower third or so of the quality distribution. This makes quality grow with age, helping to fit the size growth with age in the data.

Table 6: Parameter Values in the Simulated Models

Parameters	KK	KK 3	Directed CD	New Varieties	Own Innov.
δ_e	6.5%	5.5%	5.8%	5.8%	10.0%
δ_i	5.0%	3.7%	4.5%	6.0%	15.5%
κ_e	-	6.3%	6.5%	8.7%	2.0%
κ_i	-	-	-	0.0%	0.0%
λ_i	-	-	-	-	45.0%
s_q	1.257	1.313	1.422	1.359	1.058
ρ_i	1	1	0.10	0.10	0.10
ρ_e	-	1	0.35	0.32	0.41
θ	10	80	80	80	18
ψ	-	0.024	0.022	0.026	0.035
δ_o	-	6.3%	6.5%	8.2%	1.5%

Also helping is the high Pareto shape parameter for the quality of new varieties. The high value $\theta = 18$ implies a narrow, low-mean distribution of quality for new varieties, helping to make younger firms smaller than older firms.

Figures 5 through 9 contrast the exit rate by size in each model with the data counterpart. In the original Klette-Kortum model with $\sigma = 1$ (KK), firm size is proportional to its number of varieties. Firms with many varieties are unlikely to lose them all at once, as creative destruction is independent across varieties. Thus exit falls too sharply with size, as displayed in Figure 5. This failure of the original Klette-Kortum model leads us to consider $\sigma > 1$, so that higher quality varieties employ more workers in their production. As a result, big firms tend to have higher quality varieties rather than just more varieties.[9] A higher elasticity of substitution ($\sigma = 3$ in KK 3) flattens out the exit rate by size, as shown in Figure 6. When we allow for directed creative destruction, which is helpful to

[9]In the KK 3 model, overhead costs are crucial for obtaining a stationary quality (and therefore size) distribution. The overhead costs kill off lower quality varieties, offsetting the spreading impact of random arrival of quality innovations.

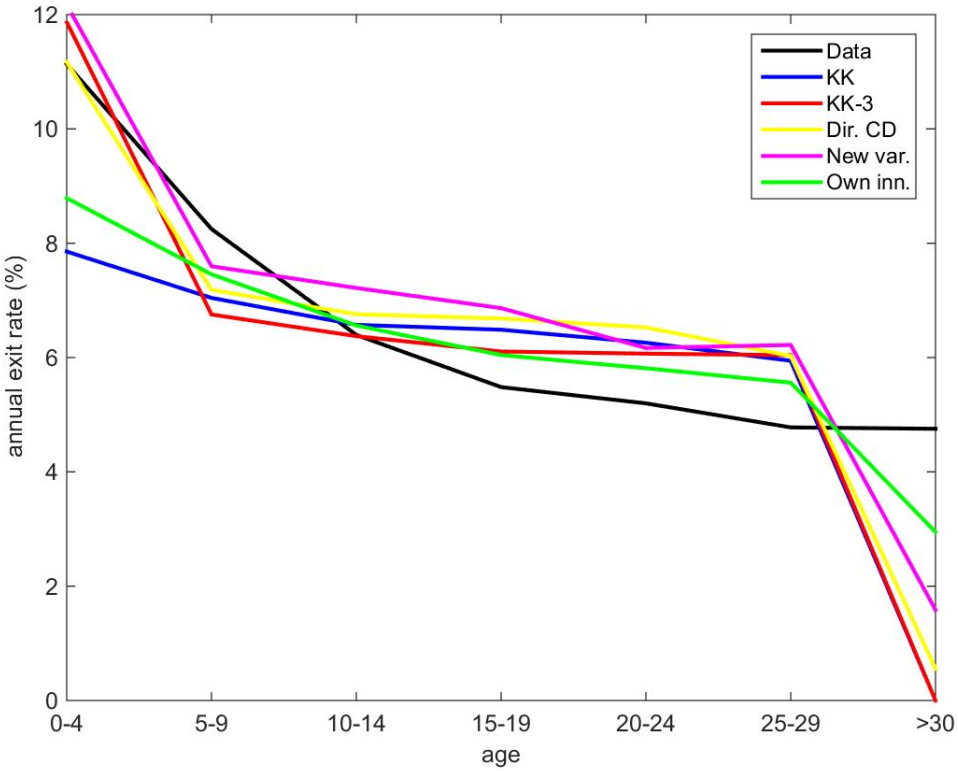

Figure 3: Model Fit, Exit by Age

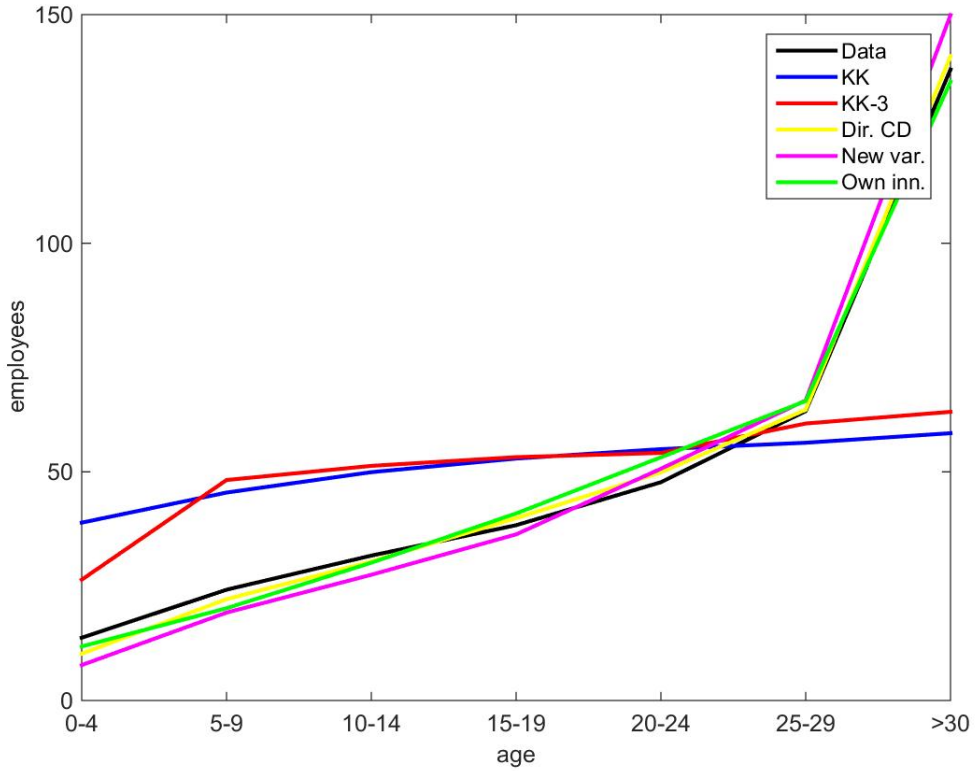

Figure 4: Model Fit, Size by Age

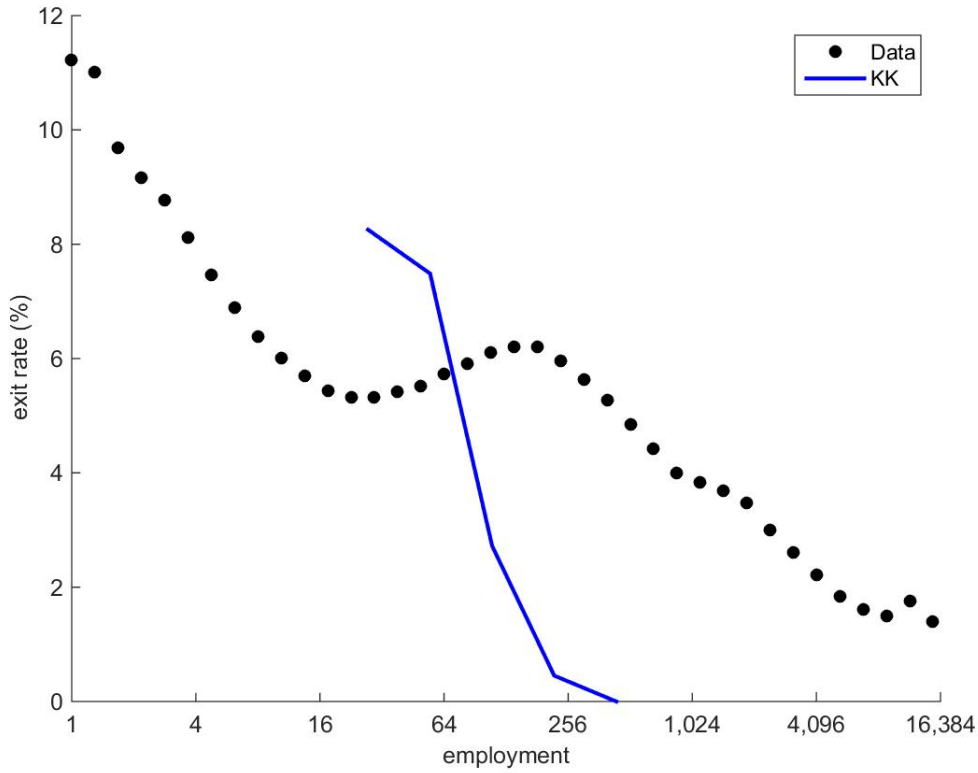

Figure 5: Model Fit, Exit by Size, Klette-Kortum

better fit size by age, the exit rate falls more with size (Figure 7). Adding new varieties to match growth in the number of plants has little effect (Figure 8). The fit with incumbents improving their own varieties is in Figure 7.

Figures 10 through 14 plot job creation and destruction for each model compared to the data. These Figures display the *share* of job creation and destruction due to firms exiting, contracting by various amounts, expanding by various amounts, and entering. All of the models come close to fitting the overall job creation and destruction rates of around 11 percent each.

As mentioned earlier, models with undirected innovation generate unrealistically thick tails of job creation and destruction. This is most striking in Figure 11 for the Klette-Kortum model with $\sigma = 3$. When firms can acquire much better

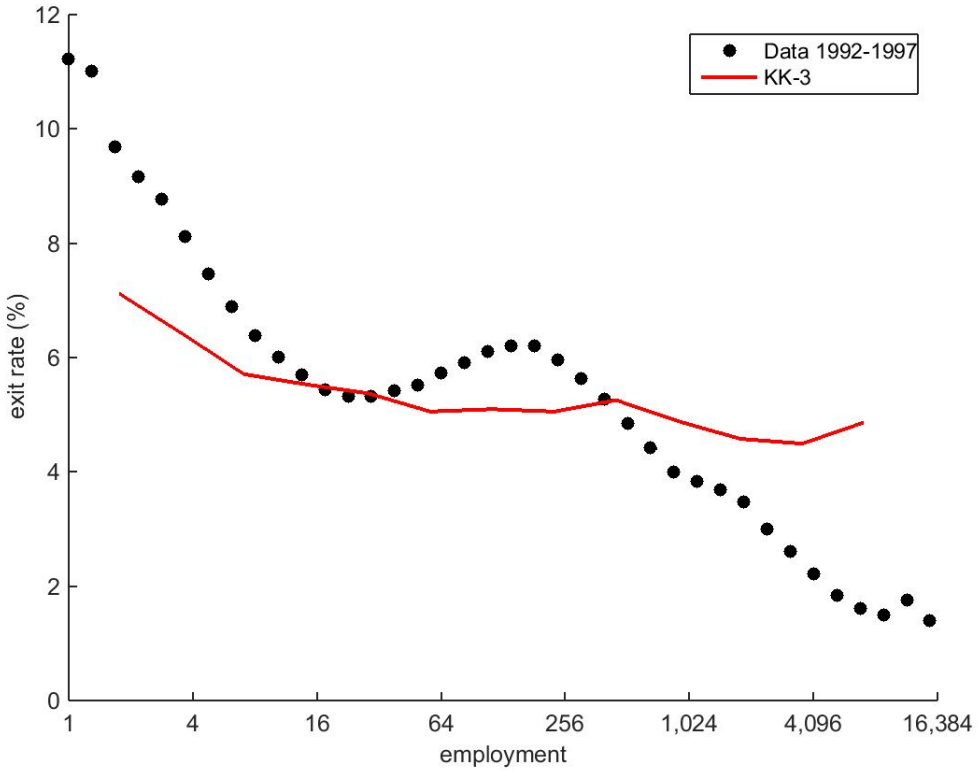

Figure 6: Model Fit, Exit by Size, Klette-Kortum 3

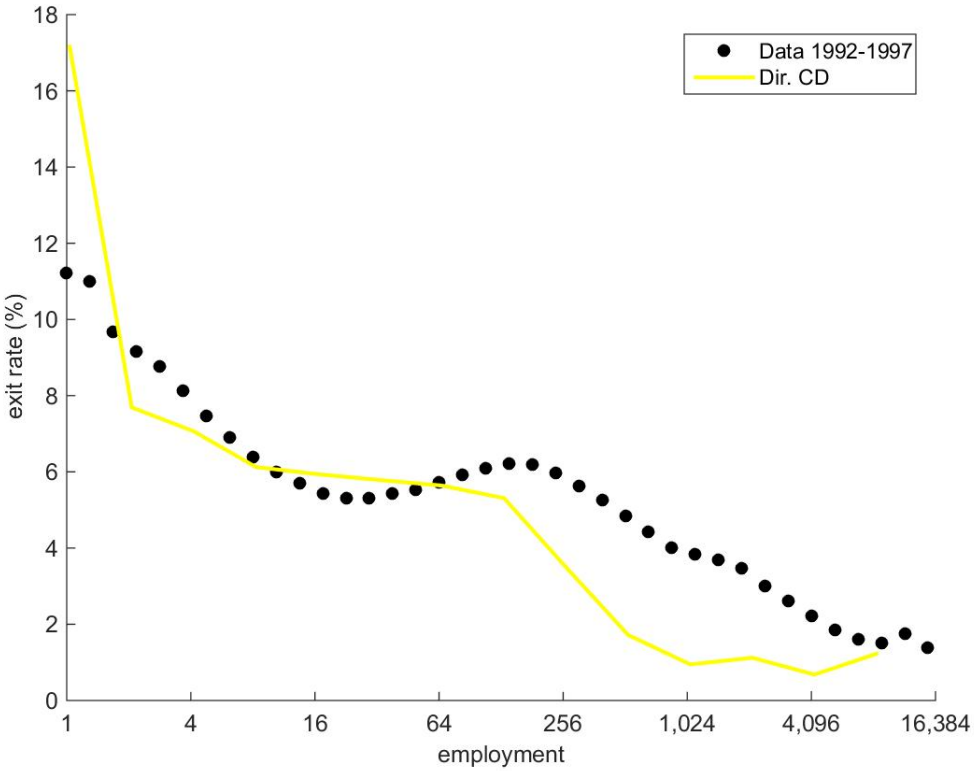

Figure 7: Model Fit, Exit by Size, Directed Creative Destruction

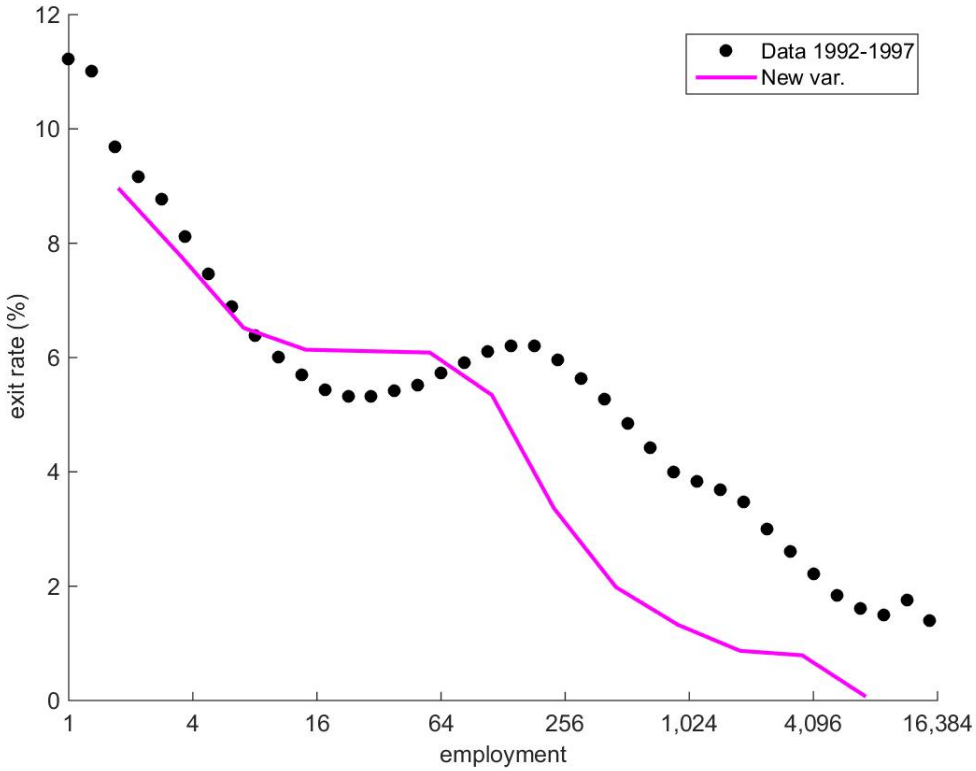

Figure 8: Model Fit, Exit by Size, New Varieties

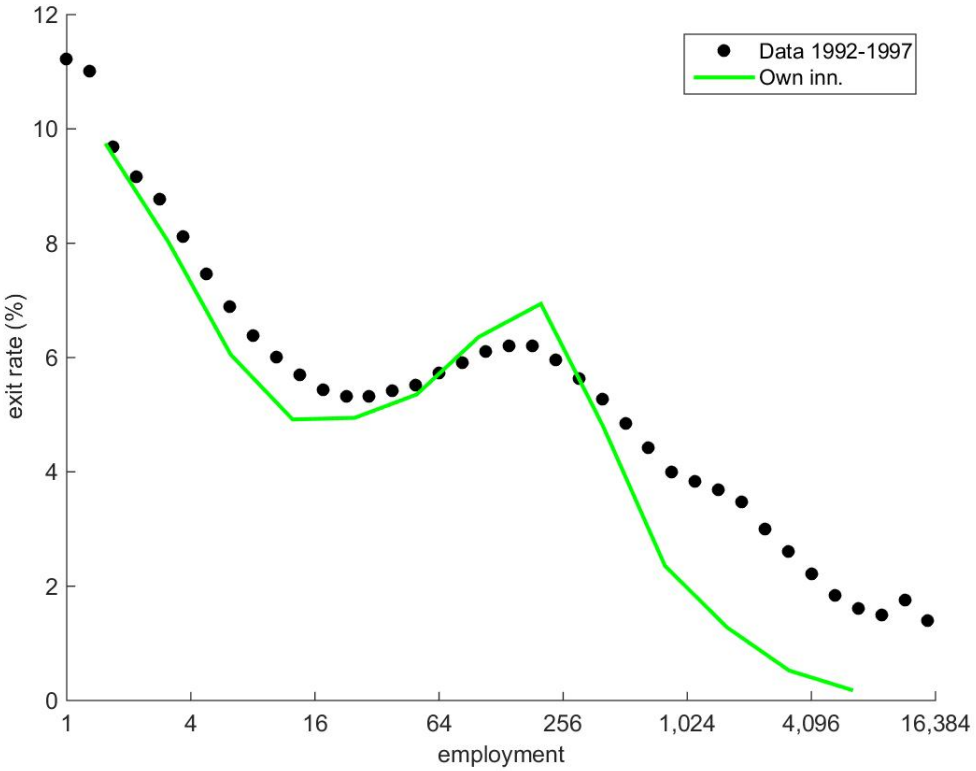

Figure 9: Model Fit, Exit by Size, Own Innovation

varieties they routinely grow by a lot; when they can lose their best varieties and retain their worst ones they can shrink by a lot.

Models with directed innovation do much better at matching the empirical distribution of job creation and destruction. This can be seen in Figures 12 to 14. When incumbents target their own quality decile for creative destruction, the varieties they acquire and lose are similar to the varieties they start with and retain, making job creation and destruction less extreme.

Adding own-variety improvements fits the job creation and destruction picture better. With frequent own-variety improvements, the step size can be much smaller: 6 percent rather than 26 to 42 percent in Table 6. Firms experience many modest increases in their size, and fewer extreme increases or decreases. See Figure 14.

We need creative destruction to generate the left tail of job destruction, including exit. Overhead costs generate exit of only the smallest plants. Creative destruction produces nontrivial exit rates for bigger plants. There is not enough job creation from entrants in the data to be the flip side for all job destruction, so creative destruction has to come from incumbents as well as entrants. Also, there is too little aggregate productivity growth in the data for all the life-cycle growth of incumbents to come from own-variety quality improvements; much of incumbent growth has to come from creative destruction by incumbents. We infer that incumbents do not grow by adding brand new varieties, as their exit rates would fall too quickly with size. We get modest creation of net new varieties by entrants because growth in the number of plants in the data is limited, and varieties/plant do not trend in our model.[10] Finally, when incumbents improve their own varieties we obtain a more realistic (but far from perfect) picture for job creation, with many plants expanding by modest amounts. Creative destruction alone makes the tails of the job growth distribution too thick.

[10]This is tied to the assumption that new plants produce a single variety. We plan to explore whether varieties per plant are indeed fixed over time in the U.S. Census data.

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 vs. incumbents? We try to infer the sources of innovation from manufacturing plant dynamics in the U.S. We conclude that creative destruction is vital for understanding job destruction and accounts for something like 20 percent of growth. Own-product quality improvements by incumbents appear to be the source of 80 percent of growth. Net variety growth contributes little, as creation of new varieties is mostly offset by exit of low quality varieties.

Our findings could be relevant for innovation policy. The sources of growth we identify can alter business stealing effects vs. knowledge spillovers, and hence the social vs. private return to innovation. 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).

It would be interesting to extend our analysis to other sectors, 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. Chinese manufacturing has seen entry and expansion of private enterprises at the expense of state-owned enterprises (Hsieh and Klenow (2009)). In India, manufacturing incumbents may be less important for growth given that surviving incumbents do not expand anywhere near as much in India as in the U.S. (Hsieh and Klenow (2014)).

Our accounting is silent on how the types of innovation interact. In Klette and Kortum (2004) a faster arrival of entrant creative destruction discourages R&D by incumbents. But, 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 innovation as for entrant innovation. Young firms might instead generate more knowledge spillovers than old firms do – Akcigit and Kerr (2015) provide evidence for this hypothesis in terms of patent citations by other firms.

We assumed no frictions in employment growth or misallocation of labor across firms. In reality, the market share of young plants could be suppressed by adjustment costs, financing frictions, and uncertainty. On top of adjustment costs for capital and labor, plants may take awhile to build up a customer base, as in work by Foster et al. (2013) and Gourio and Rudanko (2014). Irreversibilities could combine with uncertainty about the plant's quality to keep young plants small, as in the Jovanovic (1982) model. Markups could vary across varieties and firms. All of these would create a more complicated mapping from plant employment growth to plant innovation.

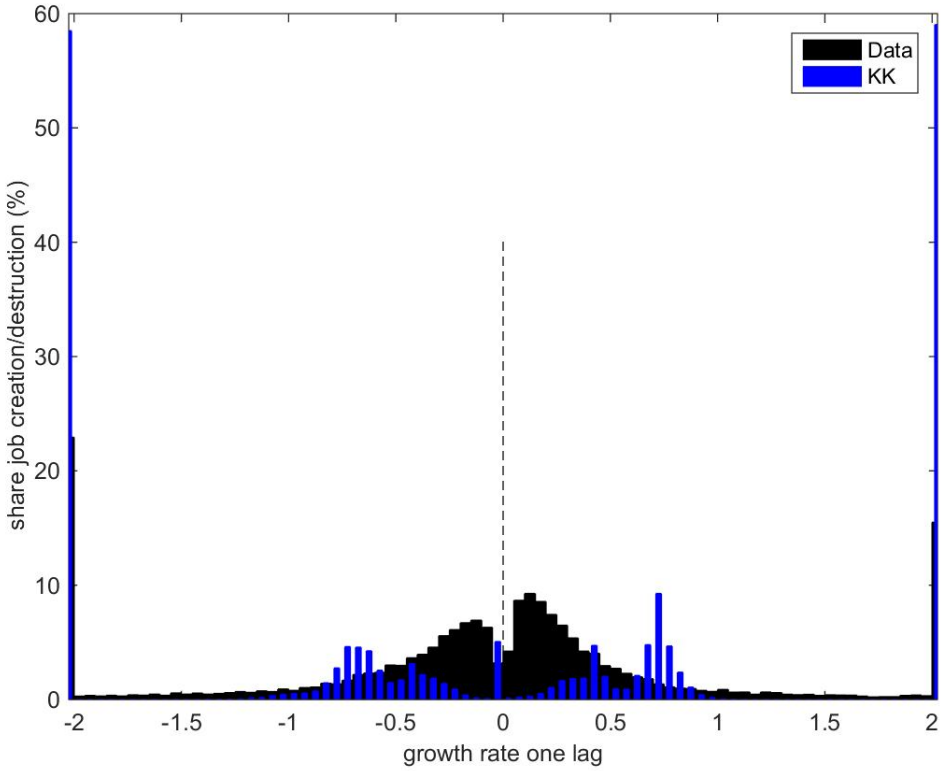

Figure 10: Model Fit, Job Creation and Destruction, Klette-Kortum

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.

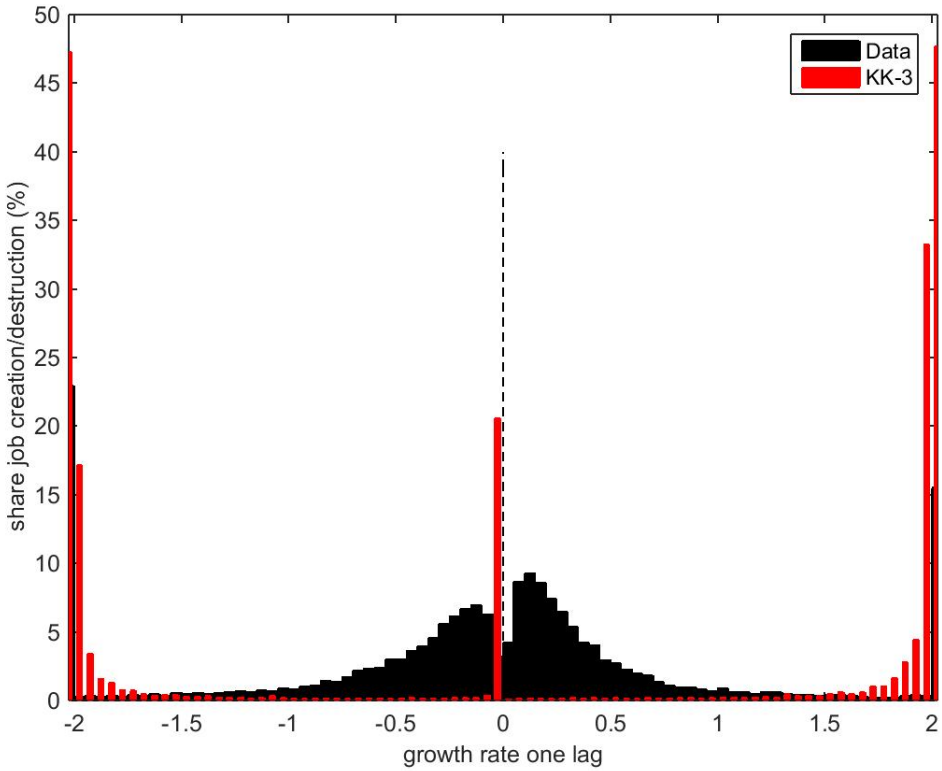

Figure 11: Model Fit, Job Creation and Destruction, Klette-Kortum 3

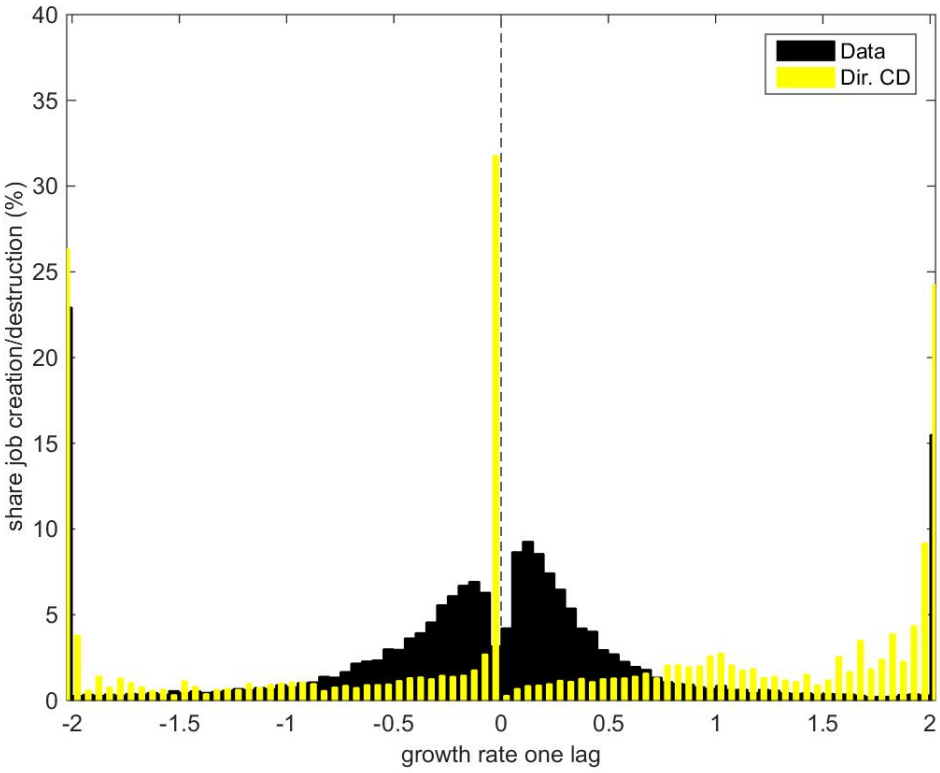

Figure 12: Model Fit, Job Creation and Destruction, Directed Creative Destruction

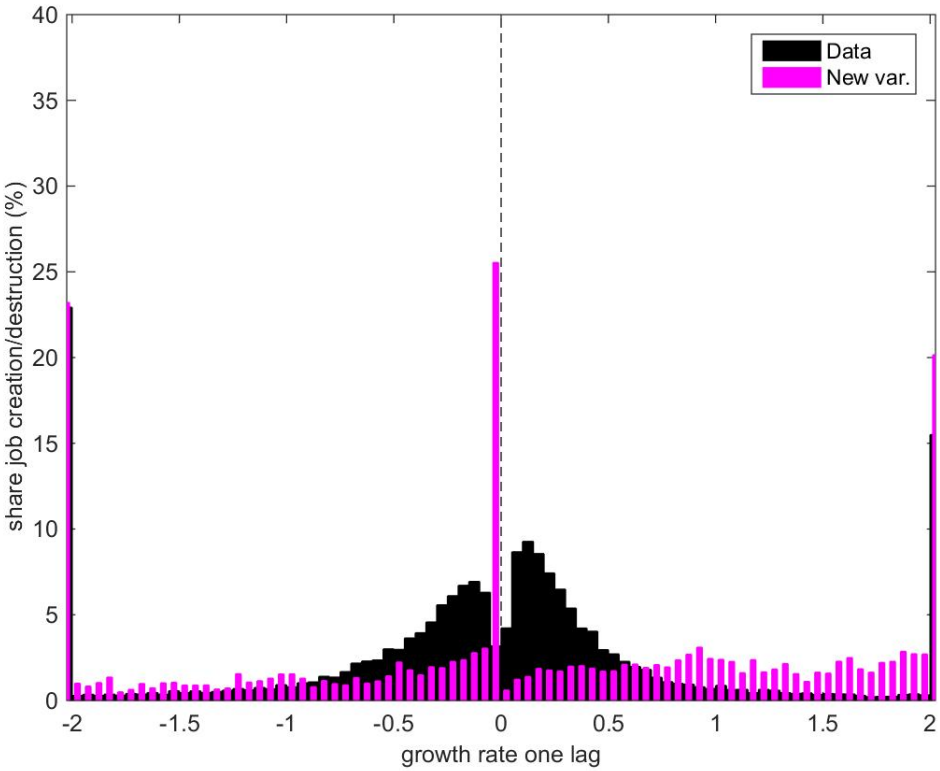

Figure 13: Model Fit, Job Creation and Destruction, New Varieties

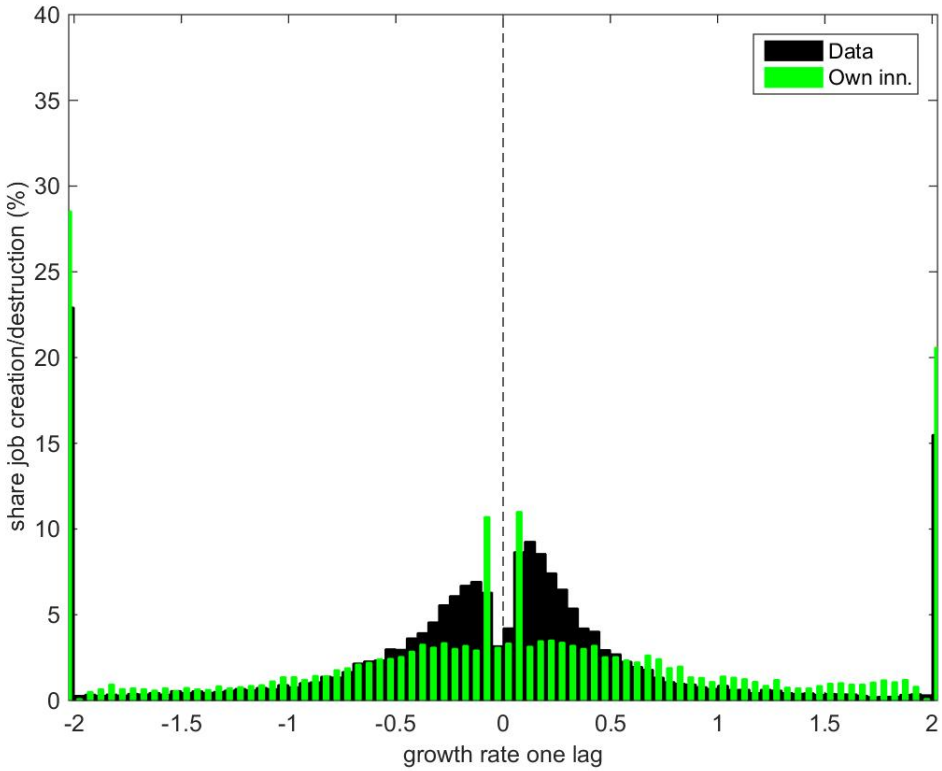

Figure 14: Model Fit, Job Creation and Destruction, Own Innovation

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

Davis, Steven J, John C Haltiwanger, and Scott Schuh, "Job creation and destruction," *MIT Press Books*, 1998.

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?," 2013.

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, forthcoming*, 2014.

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.

Hottman, Colin, Stephen J Redding, and David E Weinstein, "What is' Firm Heterogeneity'in Trade Models? The Role of Quality, Scope, Markups, and Cost," Technical Report, National Bureau of Economic Research 2014.

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, forthcoming*, 2014.

Jones, Charles I., "Life and Growth," *Journal of Political Economy, forthcoming*, 2014.

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.