

How Destructive is Innovation?

Daniel Garcia-Macia

Stanford University

Chang-Tai Hsieh

University of Chicago and NBER

Peter J. Klenow*

Stanford University and NBER

January 30, 2015

Abstract

Entering and incumbent firms can create new products and displace other firms' 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 firms from 1963 through 2002, we arrive at three main conclusions: First, most growth has seemed to come from incumbents' innovation rather than innovation by entrants. We infer this from the modest market share of entering firms. Second, most growth seems to have come from improvements of existing varieties rather than creation of brand new varieties. We infer this because firm exit rates fall only gradually as firms expand, suggesting they are not accumulating a larger set of products. Third, own-product improvements by incumbents appear to have been more important than creative destruction. We infer this because the distribution of firm growth rates has had thinner tails than implied by a model in which growth is entirely due to creative destruction.

*We thank Ufuk Akcigit 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 market share and 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. See chapter 14 in Acemoglu (2011) for examples.¹ See also Akcigit and Kerr (2013), who provide evidence that firms are more likely to cite their own patents and hence build on them. Still other theories, such as Romer (1990), emphasize the contribution of brand new varieties to growth.

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, in contrast, there are positive knowledge externalities that boost the social return relative to the private return. When incumbents successively improve their own products, business stealing effects and knowledge externalities can be mitigated. Models with expanding varieties, meanwhile, tend to have smaller business-stealing effects but retain knowledge spillovers. See the survey by Jones (2005).

Ideally, one could directly observe the extent to which new products substitute for or improve upon existing products. Broda and Weinstein (2010) is a recent effort 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

¹Another example is Lucas (1988), which emphasizes worker human capital accumulation but can be re-interpreted in terms of firms.

which figure prominently in theories of growth.²

We pursue another 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 overall growth with minimal model assumptions. We consider a specific growth model with a small 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 as far back as 1963 and as recently as 2002. We calculate aggregate TFP growth, the exit rate of plants by age, the size (employment) of plants by age, the exit rate of plants by employment, the distribution of employment growth, and growth in the total number of plants. That parameter values that best fit these moments lead to three conclusions. First, most growth – about 90 percent – seems to come from incumbents rather than entrants. This is because the employment share of entrants is modest. Second, most growth – also about 90 percent – appears to arise through quality improvements rather than brand new varieties. Third, own-variety improvements by incumbents are about twice as important as 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 dataset we exploit. Section 4 presents the parameter values of the model that best fit the moments from the data. Section 5 concludes.

²Gordon (2007) and Greenwood et al. (1997) emphasize the importance of growth embodied in durable goods based on the declining relative price of durables.

2. An Exogenous Growth Model

We adapt the Klette and Kortum (2004) model of quality ladder growth through creative destruction. In this Klette-Kortum model firms produce multiple varieties. Firms 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 varieties are improved upon and 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 is as in Klette-Kortum with the following 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 $\sigma > 1$ (rather than $\sigma = 1$)
- There are brand new varieties (rather than a fixed set of varieties)
- Incumbents can improve the quality of their own varieties (rather than quality improvements only coming from other incumbents or entrants)
- Creative destruction may be directed toward quality levels similar to each firm's existing average quality (rather than being undirected)

Aggregate output

Total output Y in the economy is given by:

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

where M is the total number of varieties and y_j is the output of variety j . The production function for each variety is linear in labor $y_j = q_j l_j$, where q_j is the quality or “process efficiency” of variety j and l_j is labor producing variety 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 arbitrarily small overhead cost of production. This assumption allows the highest quality producer to charge the standard markup over marginal cost. 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 both the market share and the employment of a firm are proportional to the sum of power qualities $q_j^{\sigma-1}$ of the varieties operated by the firm.³ Note that in the special case of $\sigma = 1$ assumed by Klette-Kortum, all varieties have equal market shares (because price is inverse proportional to quality) and employ-

³There is no misallocation of labor whatsoever in this model.

ment, and a firm's size is proportional to the number of varieties it controls. We will find it important to allow $\sigma > 1$, so that firms can be larger because they have higher quality products rather than just a wider array of products.

Aggregate productivity

Labor productivity in the economy is given by

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

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 for each type is given in Table 1. The probabilities shown are per each of the current varieties a firm produces. The probability of a firm improving any given variety it produces is λ_i , and such improvement is associated with step size $s_\lambda \geq 1$. 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 comes with step size $s_\delta \geq 1$.

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 from entrants is drawn at random from the current distribution of qualities (undirected innovation), but a multiplicative step s_κ is added to each quality. The arrival rate of brand new varieties affects growth in the number of firms

Table 1: Channels of Innovation

channel	probability	step size
own-variety improvements by incumbents	λ_i	$s_\lambda \geq 1$
creative destruction by entrants	δ_e	$s_\delta \geq 1$
creative destruction by incumbents	δ_i	$s_\delta \geq 1$
new varieties from entrants	κ_e	s_κ
new varieties from incumbents	κ_i	s_κ

(tied to κ_e and κ_i), while the arrival rate and step size for new varieties (both κ_e and s_κ) will affect the size of new firms.

On top of the eight parameters listed in Table 1, we add two more parameters. Klette-Kortum assumed creative destruction was undirected. We find that, 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 on the losing side of creative destruction can lose their best varieties, leaving them with low quality varieties and steeply negative growth.⁴ To allow some control over the distribution of tail growth rates in the model, we allow for the possibility that creative destruction is directed. We parameterize directedness in the following way. For incumbent firms, we assume that creative destruction occurs within quality quantiles. If there are 100 such quantiles, then creative destruction is random within a quality *percentile*. If there are 10 quantiles of quality, then an incumbent creatively destroys an existing variety in its decile. If there is only a single quantile then creative destruction is undirected.⁵ ρ_i denotes the width of the quantile targeted by incumbents. For entrants, we parameterize the directedness as the market share of the existing varieties targeted by entrant

⁴Of course, there are always entrants and exiters, which are at the extremes in the employment growth distribution.

⁵As all arrival rates are per existing variety, the quantiles are defined for each individual variety. We also allow brand new varieties created by incumbents to be directed in this way.

firms. We denote this market share as ρ_e .⁶

Note that, in our model, each innovation is proportional to an existing quality level. Thus, if innovative effort was endogenous, there would be a positive knowledge 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), Kortum (1997), Klette and Kortum (2004), and Acemoglu et al. (2013).

Output growth

Total output grows at rate

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

The κ components correspond to the creation of new varieties. The $(1 + g_q)$ component reflects growth in average quality per variety. The growth rate of the power mean of quality levels across varieties is:

$$1 + g_q = \frac{s_\kappa^{\sigma-1}\kappa_e + s_\kappa^{\sigma-1}\kappa_i + 1 + (s_\lambda^{\sigma-1} - 1)\lambda_i + (s_\delta^{\sigma-1} - 1)(1 - \lambda_i)(\rho_e\delta_e + \delta_i)}{1 + \kappa_e + \kappa_i}$$

3. U.S. Manufacturing Census Data

We use data from U.S. manufacturing Censuses to quantify dynamics of entry, exit, and survivor growth. We focus on *plants* rather than firms, because mergers and acquisitions can wreak havoc with our strategy to infer innovation from growth dynamics.

We use data as far back as 1963, but often from 1972 onward because there is no capital stock data before 1972. We use data through at most 2002, because the NAICS definitions were changed from 2002 to 2007.

⁶It is convenient to define ρ_e as the market share to obtain an analytic expression for the aggregate productivity growth rate.

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. We take this from the U.S. Bureau of Labor Statistics Multifactor Productivity Growth. Converting their gross output measure to value added, TFP growth averages 2.51 percent per year from 1987–2011 (the timespan of their data, which is not admittedly not ideal for us). The number of manufacturing plants in the U.S. Census of Manufacturing, meanwhile, rose 0.49 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 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 2, from Davis et al. (1998), plots the distribution of 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.

Lastly, we use the standard deviation of firm size in U.S. manufacturing.

4. Indirect Inference

We now compare moments from model simulations to the manufacturing moments we calculated in the previous section. Our aim is to indirectly infer the sources of innovation. Our logic is that plant entry and exit rates, size, size dispersion, and growth are byproducts of innovation. Entrants reflect a combination of new varieties and creative destruction of existing varieties. The better the new varieties, the bigger the market share of entrants. When a plant ex-

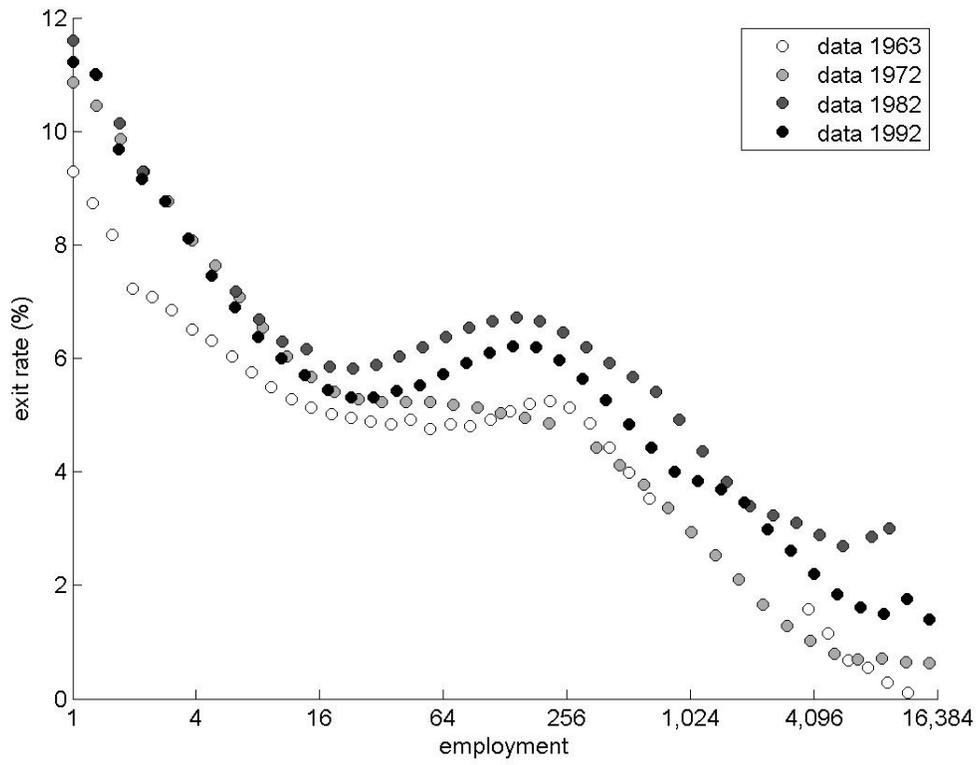


Figure 1: Empirical Exit by Size in the U.S. Census of Manufacturing

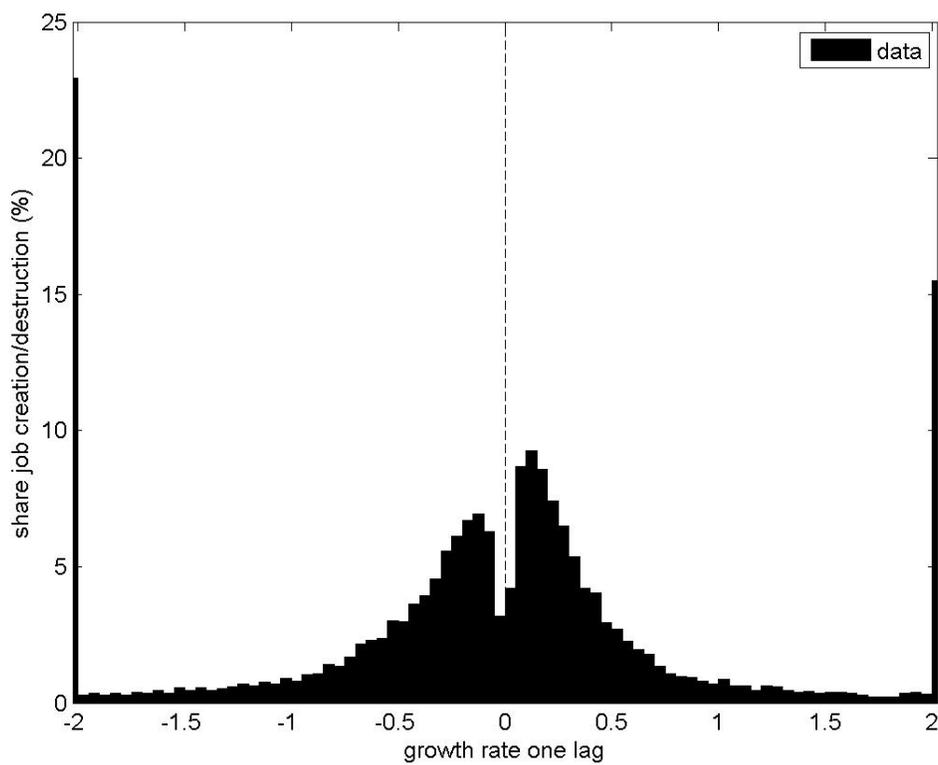


Figure 2: Job Creation and Destruction Rates in U.S. Manufacturing
(via Davis, Haltiwanger and Schuh, 1998)

pands, 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 (and hence real wage 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 to creative destruction. Because creative destruction is independent across a plant's varieties (by assumption), plants with more varieties have much lower exit rates. Plants with more varieties also have less dispersed growth rates. Plants with higher qualities are larger and potentially more protected against exit, as they have a lower likelihood of being captured by entrants, which only target the bottom of the quality distribution.

Simulation algorithm

Though each firm's static maximization problem (specifically, its market share in terms of revenue or labor) can be solved analytically, 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 size heterogeneity in varieties we allow (due to $\sigma > 1$), preclude us from obtaining analytical results. Our numerical simulation algorithm consists of the following steps:

1. Specify the distribution of quality across varieties.
2. Simulate life paths for entering plants such that the total number of plants observed, including incumbents, is the same as in U.S. manufacturing plants from 1992, 1997, and 2002 (321,000 on average, including administrative record plants).
3. Each entrant has one initial variety, captured or newly created. In each year of its lifetime, it faces a probability of each type of innovation occur-

ring per variety it owns, as in Table 1. New varieties by entrants draw a random quality from the distribution in the population (undirected) or from the bottom quantile (directed). For incumbents, newly created varieties and creatively destroyed varieties are either draws at random (undirected) or from the quality quantile of its own variety (directed). A firm's life ends when it loses all of its varieties or when it reaches age 100.

4. Based on an entire population of simulated firms of all ages, compute the joint distribution of quality and variety across firms. Calculate moments of interest (e.g. exit by size).
5. Repeat steps 1 to 4 until all moments converge. In each iteration, update the guess for the distribution of qualities by combining elements from the previous iteration's guess.
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 29 percent arrival rate for own-variety quality improvements by incumbents. Conditional on no own-innovation, quality improvements through creative destruction occur 77 percent of the time by other incumbents, and only 6.2 percent of the time by entrants. As we will explain in more detail below, the distance minimization algorithm would drive us to the corner solution where creative destruction takes the form of pure imitation, i. e. $s_\delta = 1$. However, we find it unrealistic that a firm could entirely lose a variety to another firm that has not made any improvement at all. Therefore, we impose that creative de-

Table 3: Inferred Parameter Values

Channel	Probability	Step Size
Own-variety improvements by incumbents	29.0%	1.06
Creative destruction by entrants	6.2%	1.01
Creative destruction by incumbents	76.6%	1.01
New varieties from entrants	0.5%	1.00
New varieties from incumbents	0.0%	1.00

struction entails at least a one percent jump with respect to the original quality level of a variety. The solution is then $s_\delta = 1.01$. Own-variety improvements are associated with a larger quality jump: 6 percent. New varieties arrive at a slow rate, boosting total variety by 0.5 percent per year, and come entirely from entrants. They do not generate any quality jump, but since entrants are directed to the bottom quality quantile, they are smaller on average.

Table 4 presents the implied sources of growth. About 22 percent of growth comes from creative destruction. Own-variety improvements by incumbents seem more important, at 68 percent. New varieties *a la* Romer (1990) are the remainder at about 10 percent. All three sources matter, but own-variety improvements tend to dominate under the Table 3 parameter values. Incumbents contribute more to growth (88 percent) than do entrants (12 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 $\sigma > 1$, which allows high quality varieties

Table 4: Inferred Sources of Growth

	Entrants	Incumbents	
Creative destruction of existing varieties	0.6%	21.6%	22.2%
Creation of new varieties	9.9%	0.0%	9.9%
Own-variety improvements	-	67.9%	67.9%
	12.1%	87.9%	

Table 5: Simulated Models

	Klette-Kortum KK1	Klette-Kortum KK3	directed	new var.	general
σ	1	3	3	3	3
creative destruction by entr., inc. (partially) directed innovation	✓	✓	✓ ✓	✓ ✓	✓ ✓
new varieties from entr., inc.				✓	✓
own-variety improvements by inc.					✓

to have higher market shares. We next allow for partially directed innovation. 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. It is informative to see how adding layers of generality to the Klette-Kortum model helps us progressively achieve a better fit with the data. Let us start with the baseline Klette-Kortum model. The average exit level, together with the exit by age slope, pin down the arrival rate of creative destruction by entrants and incumbents. Intuitively, 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

Table 6: Parameter Values in the Simulated Models

Parameters		Klette-Kortum		Klette-Kortum 3		directed innov.		new vars.		general	
δ_e	s_δ	4.3%	1.05	4.3%	1.05	3.2%	1.07	3.6%	1.06	6.2%	1.01
$\tilde{\delta}_i$	s_δ	47.9%	1.05	47.9%	1.05	37.0%	1.07	35.0%	1.06	76.6%	1.01
κ_e	s_κ	-	-	-	-	-	-	0.5%	0.84	0.5%	1.00
κ_i	s_κ	-	-	-	-	-	-	0.0%	0.84	0.0%	1.00
λ_i	s_λ	-	-	-	-	-	-	-	-	29.0%	1.06
ρ_i		-		-		1%		1%		1%	
ρ_e		-		-		27%		34%		35%	

improved upon by another firm. In this model, the exit rate of a one-variety firm is given by the following expression:

$$(\delta_e + \delta_i)(1 - \delta_i)$$

This is simply the sum of the probability of creative destruction by an entrant and by an incumbent firm, times the probability that the firm does not creatively destroy a variety from another incumbent. Around half of the varieties are subject to creative destruction. Therefore, the step size associated to this channel has to be roughly twice the aggregate growth rate in order for the model to be consistent with that moment.

Figure 3 plots the exit rate of plants by age in each model vs. that in the data. The fit is visibly accurate for all models (although the models without own innovation understate the exit rate of firms 30 years and older). However, the baseline model is not so successful regarding the other moments that we analyze. Figure 4 plots the empirical size of plants by age against the size of plants by age in each model. In this case, only models with new varieties are able to replicate the amount of growth in size by age seen in the data.

Figures 5, 6, 7, 8 and 9 contrast the exit rates by size in each model with the

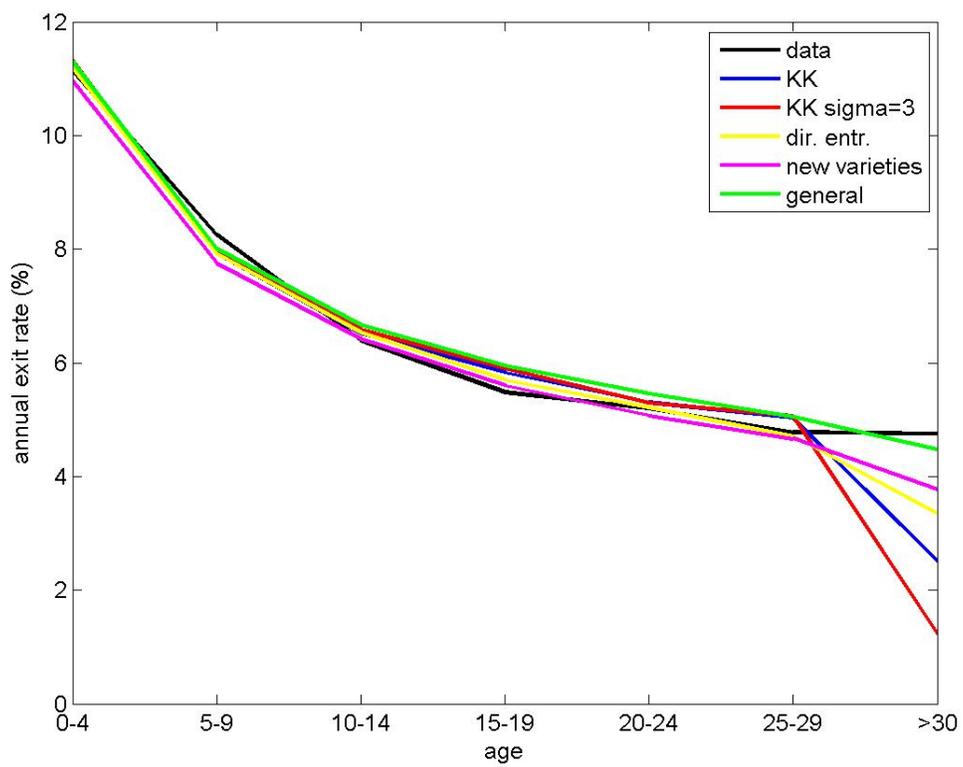


Figure 3: Model Fit, Exit by Age

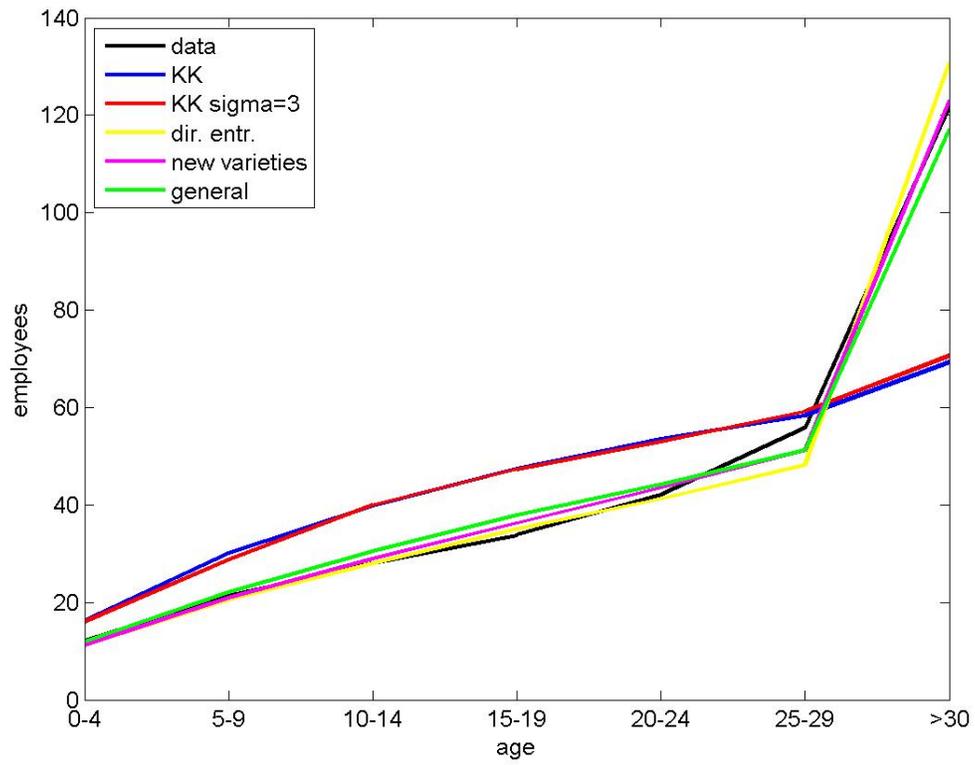


Figure 4: Model Fit, Size by Age

data counterpart. In the original Klette-Kortum model, bigger firms are firms with more 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. On top of that, the original Klette-Kortum model cannot generate the empirical dispersion in firm log employment, which is equal to 2.1 percent. These two failures of the Klette-Kortum model lead us to allow for $\sigma > 1$, so that big firms are also firms with higher quality rather than more varieties. Indeed, allowing for a higher elasticity of substitution in model Klette-Kortum 3 improves the fit of exit by size, shown in Figure 6. In this particular model configuration, it is necessary to assume a non-trivial amount of overhead costs per variety in order to obtain a stationary quality (and size) distribution. We adjust the level of overhead costs to match the empirical size distribution.⁷

Figures 10, 11, 12, 13 and 14 show the job creation and destruction distribution for each model compared to the data. As mentioned above, models with undirected innovation generate thick tails of job creation and destruction. We illustrate this in Figure 11 for the Klette-Kortum model with $\sigma = 3$. The models with directed innovation do much better in terms of matching the empirical distribution of job creation and destruction – see Figures 12 to 14.⁸

A logical next step is to allow for directed innovation by entrants and incumbents. We can set the degree of incumbent directedness to approximately match the shape of the job creation and destruction histogram. Moreover, we can set the level of entrant directedness to generate the empirical growth in plant employment by age, as seen in Figure 4. There are two reasons why a higher degree of entrant directedness increases growth in size by age. First, entrants come in with lower quality compared to the average firm. Second and most important, targeting of low qualities by entrants generates a natural selection effect by which only high quality varieties survive over time. This

⁷Still, the necessary level of overhead costs is too small to generate any visible contribution to growth, i.e. higher than 0.05 percent of the aggregate growth rate.

⁸The reliance of the new varieties model on creative destruction can be seen in the job creation spike around 2/3, which corresponds to a firm going from 1 to 2 varieties.

model converges to a size dispersion lower than half the one in the data. Adding overhead costs, which would reduce the size dispersion, would not be helpful.

Next, we add a new channel for growth: new varieties. The arrival rate of new varieties is pinned down by the growth in the number of entrants over time, because the average amount of varieties per firm does not trend. The step size of this channel is pinned down by the size dispersion. In fact, the role of this channel in keeping the equilibrium size dispersion large enough is the reason why all new varieties are produced by entrants. If new varieties were assigned to incumbents, we would not see enough small firms to match the empirical size dispersion. Yet, the problem with this model is that it generates a very steep fall in exit by size, as we see in Figure 8.

Hence, our final generalization is to add own-variety improvements. We can use the extra parameters associated with this new channel to decrease the exit-by-size slope. This takes us to the corner solution where we attribute a growth contribution as high as possible to own-variety improvements (conditional on the assumption discussed above that creative destruction cannot have an excessively small step size). With own-variety improvements, firms do not need to rely on grabbing more and more varieties in order to grow in size. By decoupling innovation from creative destruction and exit, we can sustain a higher exit rate for large firms. The fact that we hit a corner solution explains why the fit in Figure 9 is still not perfect in the general case.

5. Conclusion

How much of innovation takes the form of creative destruction? Versus firms improving their own products? New varieties? How much of innovation occurs through entrants vs. incumbents? We try to infer the sources of innovation by matching up models with manufacturing plant dynamics in the U.S. We tentatively conclude that creative destruction is important but not the sole source of innovation. Own-product quality improvements by incumbents seem twice

as important. Of lesser but nontrivial importance are new varieties and the contribution of entrants overall.

Our findings could be relevant for innovation policy. The sources of growth we identify have implications for 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 innovation and growth given that surviving incumbents do not expand as much in India as in the U.S. (Hsieh and Klenow (2014)).

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. For example, we assumed that creative destruction was independent across varieties, even within a firm. We plan to explore the possibility that creative destruction is correlated across a family of products (e.g. Apple vs. Samsung smartphones and tablets).

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 (2013) 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. In addition to adjustment costs for capital and labor, it may take plants awhile to build up a cus-

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 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.

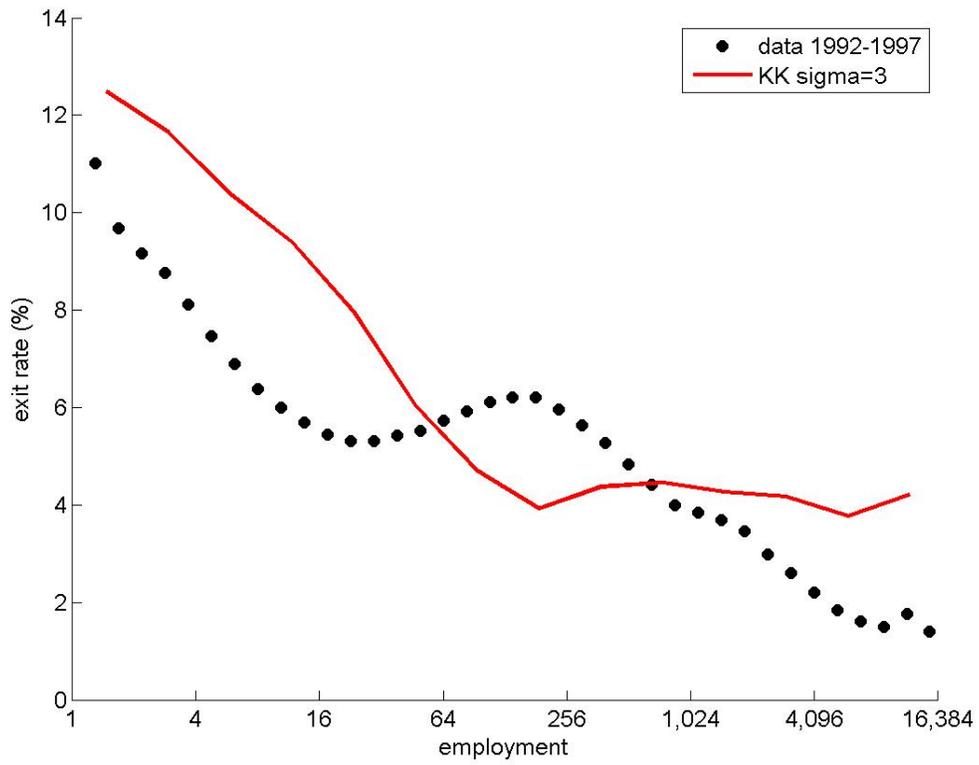


Figure 6: Model Fit, Exit by Size, Klette-Kortum 3

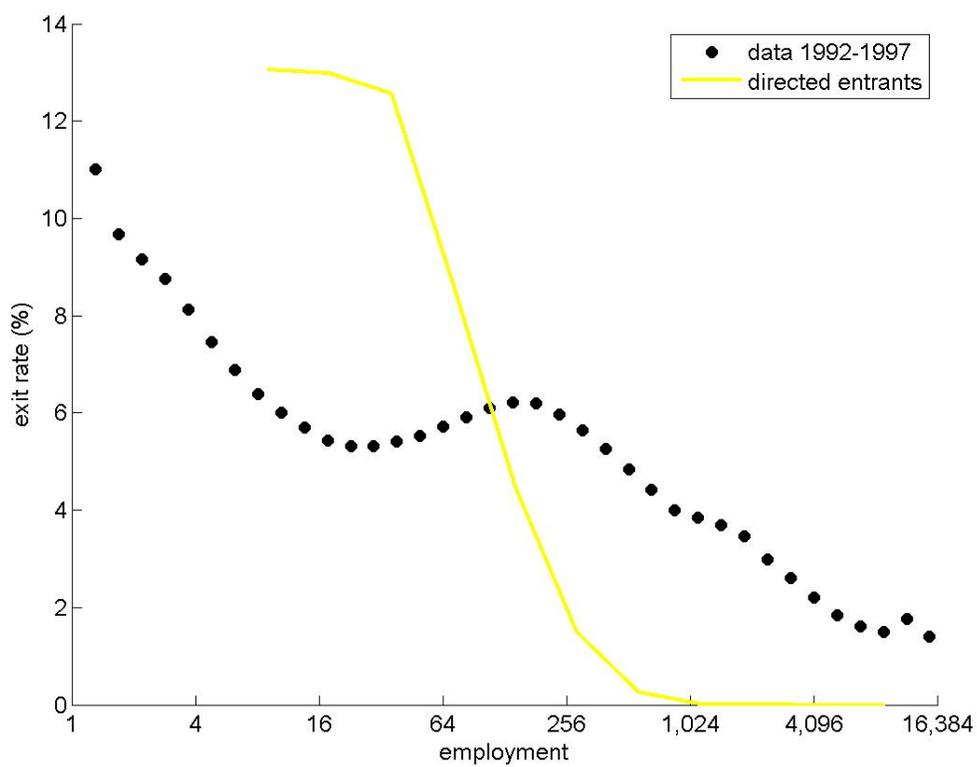
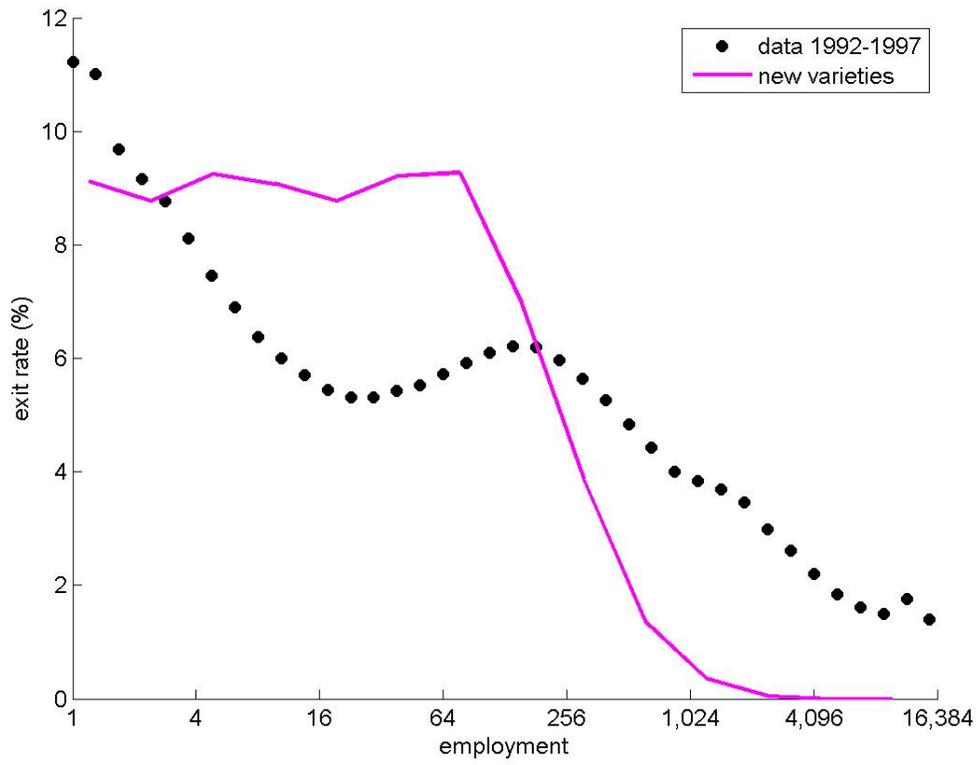
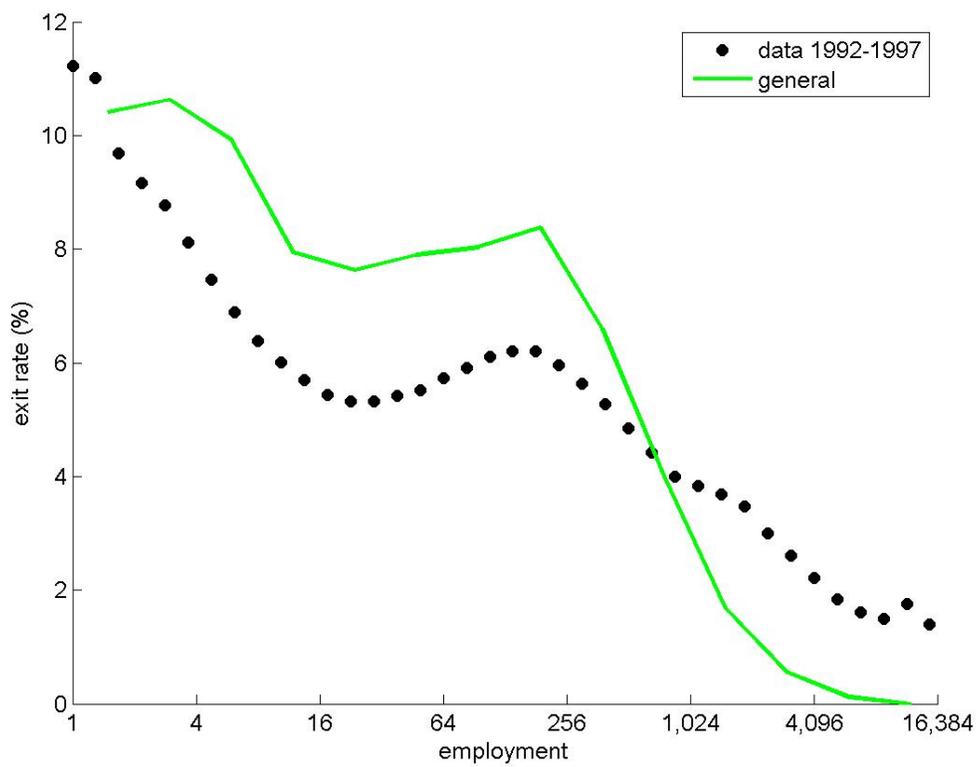


Figure 7: Model Fit, Exit by Size, Directed Innovation





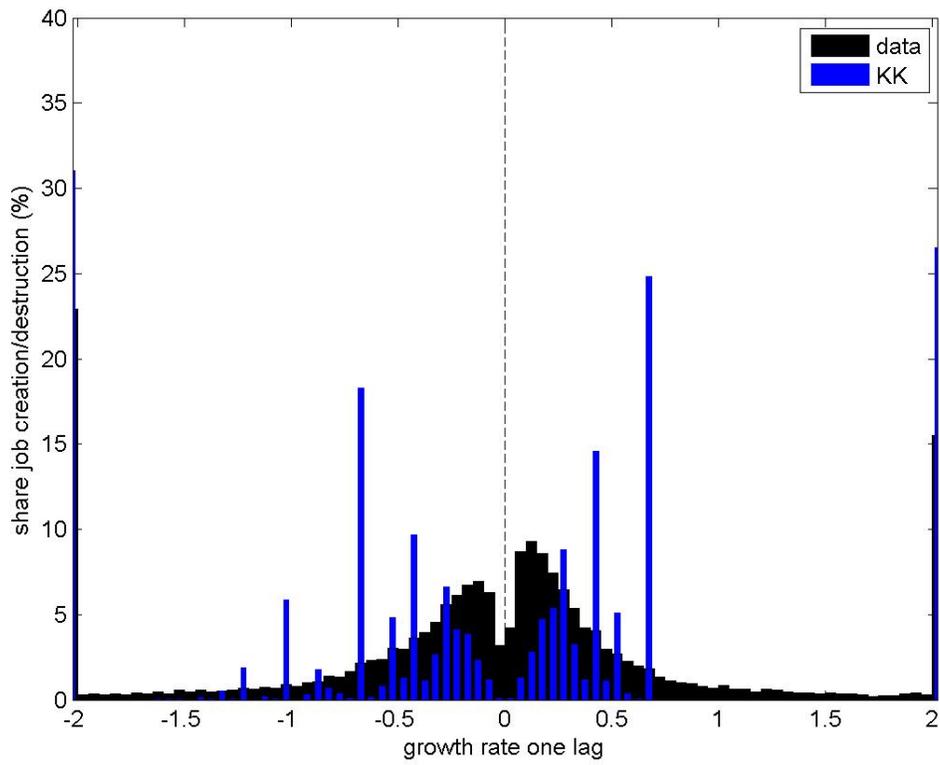


Figure 10: Model Fit, Job Creation and Destruction, Klette-Kortum

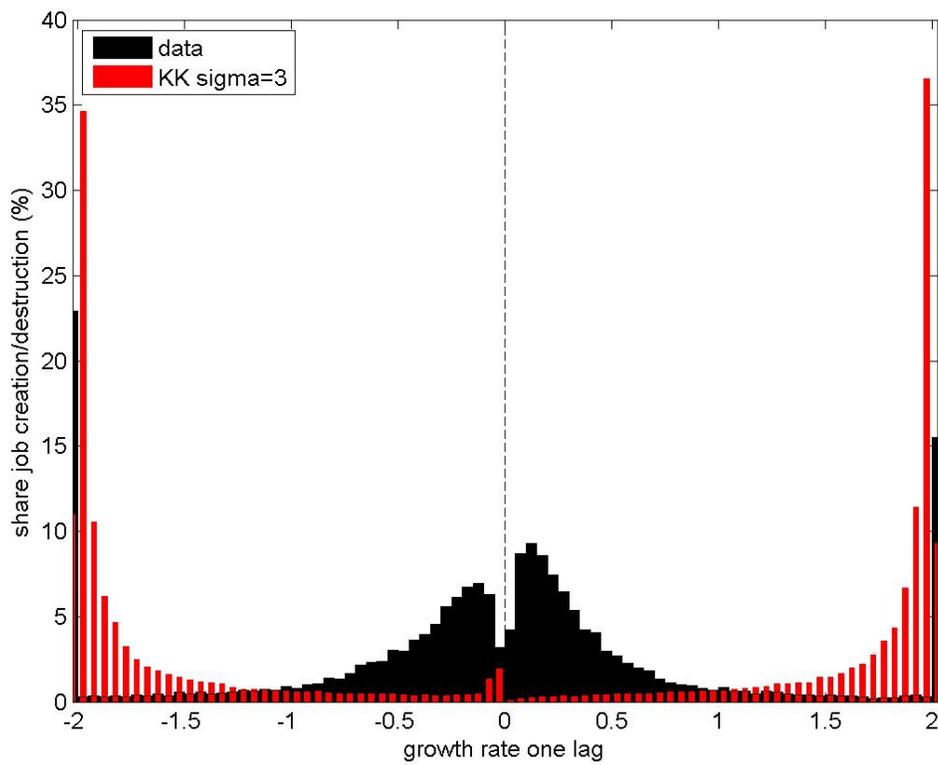


Figure 11: Model Fit, Job Creation and Destruction, Klette-Kortum 3

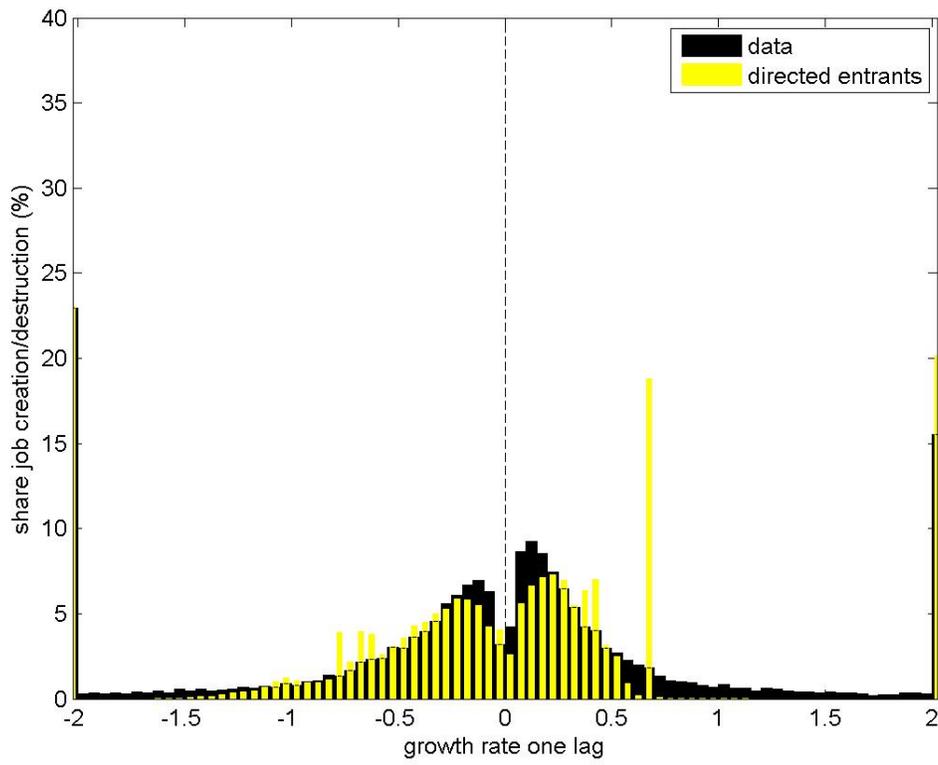


Figure 12: Model Fit, Job Creation and Destruction, Directed Innovation

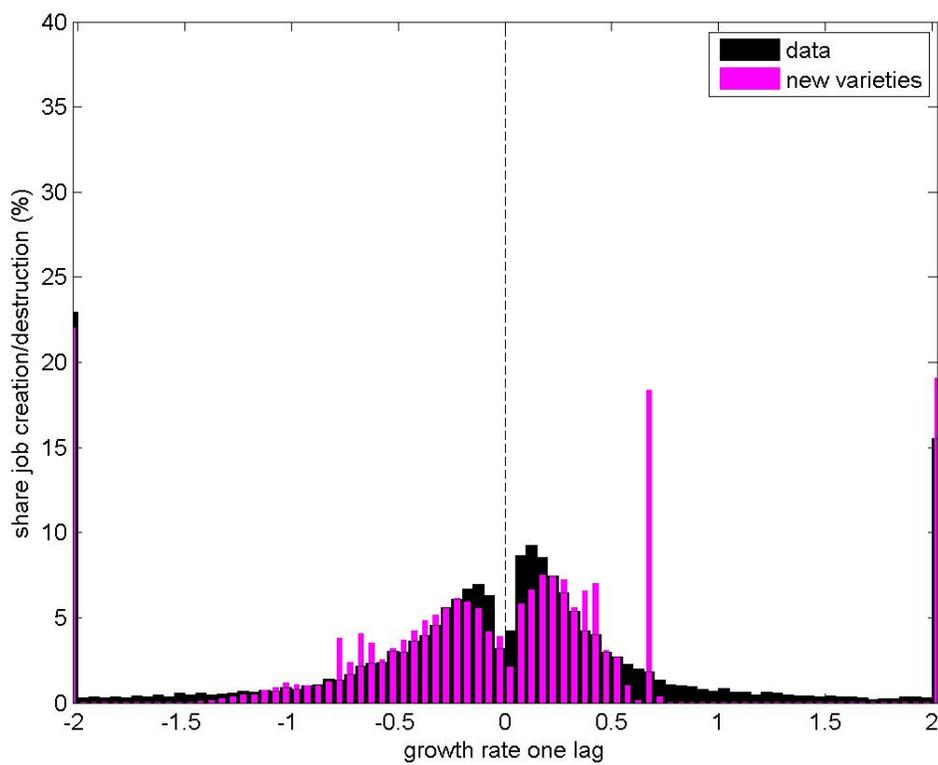


Figure 13: Model Fit, Job Creation and Destruction, New Varieties

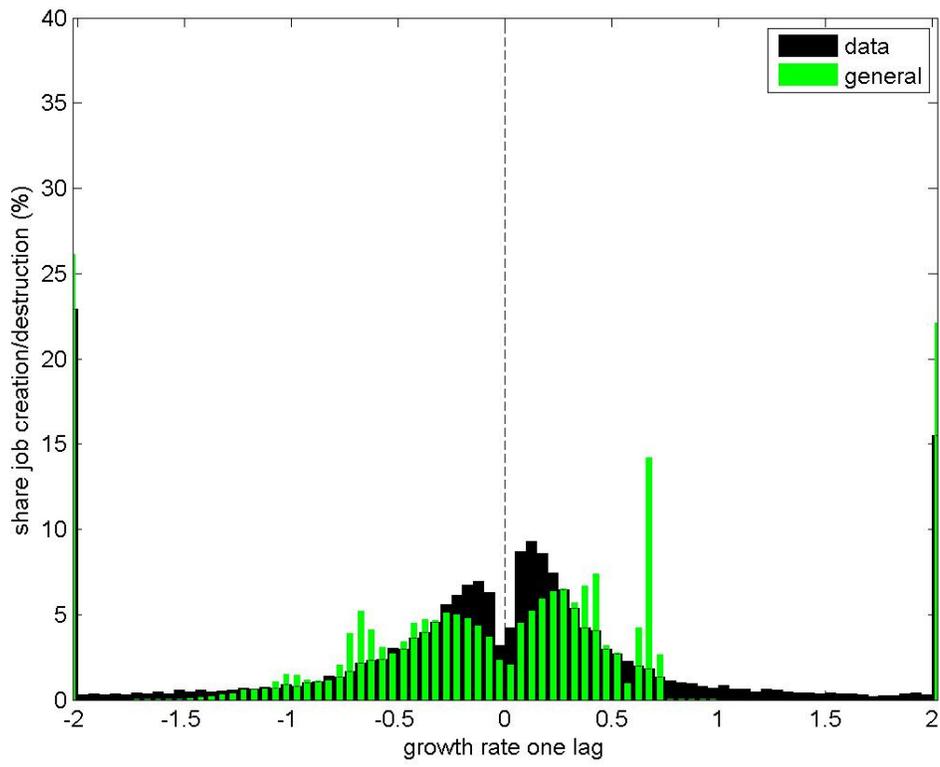


Figure 14: Model Fit, Job Creation and Destruction, General

- , 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,” Research Discussion Paper 28, Bank of Finland November 2013.
- 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*, 2014, 81 (3), 1102–1136.
- 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.
- 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, “Growth and ideas,” *Handbook of Economic Growth*, 2005, 1B, 1063–1111.
- 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 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., “On the mechanics of economic development,” *Journal of monetary economics*, 1988, 22 (1), 3–42.
- Parente, Stephen L and Edward C Prescott, *Barriers to riches*, MIT press, 2002.
- Peters, Michael, “Heterogeneous mark-ups, growth and endogenous misallocation,” Technical Report September 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.