

How Destructive is Innovation?

Daniel Garcia-Macia

Stanford University

Chang-Tai Hsieh

University of Chicago and NBER

Peter J. Klenow*

Stanford University and NBER

May 27, 2014

Preliminary

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? The answer matters for optimal innovation policy because knowledge spillovers and business stealing differ depending on the source of growth. Using U.S. Census data on manufacturing firms from 1963 through 2002, we arrive at three main conclusions: First, most growth comes from incumbents' innovation rather than innovation by entrants. This follows from the modest market share of entering firms. Second, most growth comes from improvements of existing varieties rather than creation of brand new varieties. Third, own-product improvement by incumbents is roughly as important as quality improvements through creative destruction.

*For financial support, Hsieh is grateful to Chicago's Initiative for Global Markets and the Templeton Foundation, 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. (2013) 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 optimal policy toward innovation. 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

¹Although Lucas (1988) highlights individual worker human capital accumulation, his seminal model 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. This is related to, but not the same as, growth decompositions by Baily et al. (1992) and Foster et al. (2001). These influential papers document the contributions of entry, exit, reallocation, and within-plant productivity growth to overall growth in a general fashion with minimal model assumptions. We consider a specific exogenous growth model with a few parameters, and try to pin down realistic parameter values using moments from the data. Acemoglu et al. (2013) also conduct such indirect inference with U.S. manufacturing data, and with a full-blown endogenous growth model. Their model focuses 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 fraction of plants by age, the employment share of plants by age, the exit rate of plants by size (employment), the distribution of employment, the distribution of employment growth, and growth in the total number of plants. To best fit these moments, we arrive at parameter values that lead to three tentative conclusions. First, most growth – about three-quarters – seems to come from incumbents rather than entrants. This is because the employment share of entrants is modest. Second, most growth – about four-fifths – appears to arise through quality improvements rather than brand new varieties. Third, creative destruction (by entrants and incumbents) and own-variety improvements by incumbents are roughly equal in importance.

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 datasets we exploit. Section 4 presents the parameter

²Gordon (1990) and Greenwood et al. (1997) emphasize the importance of growth embodied in durable goods based on the declining relative price of durables.

values of the model in Section 2 that best fit the moments from the Section 3 data. Section 5 concludes.

2. An Exogenous Growth Model

We adapt the Klette and Kortum (2004) model of quality ladder growth through creative destruction. In this KK model firms produce multiple varieties and try to capture each other's varieties by improving on them. Entrants try to improve on existing varieties and take them over in the process. Unlike KK 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.

More exactly, our set-up is as in KK 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 y_j is the output of variety j . The production function for each variety j 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 employed 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, so that there is no need limit pricing with respect to potential producers of a lower quality version of the same product. Each firm can charge the monopoly markup on each of its varieties even if the innovation step size is small. This avoids the markup heterogeneity that is the focus of Peters (2013) for example.

With a competitive labor market, 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

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

in the special case of $\sigma = 1$ assumed by KK, all varieties have equal market shares (because price is inverse proportional to quality) and employment, 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 . If a firm fails to improve on a given variety it produces, then that variety is vulnerable to creative destruction by other incumbent firms (with conditional probability δ_i) or by entrants (with conditional probability δ_e). Creative destruction, like own-variety improvements, comes with a step size $s_q \geq 1$.

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

Table 1: Channels of Innovation

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

arrival rate of brand new varieties affects growth in the number of firms (tied to κ_e), while the arrival rate and step size for new varieties (both κ_e and s_κ) will affect the employment share of new firms.

On top of the seven parameters listed in Table 1, we add one more parameter. KK 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 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. It seems plausible that firms would be able to improve upon products of similar quality to their existing portfolio.⁵

Note that, in our model, each innovation is proportional to an existing qual-

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.

ity 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:

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

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.

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 similarly calculate the share of *employment* by age in U.S. manufacturing. The second column of Table 2 shows that 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, com-




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

prising only 12 percent of plants but employing almost 42 percent of all workers in U.S. manufacturing. Hsieh and Klenow (2014) document that 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: Size Distribution of Plants in the 1997 U.S. Census of Manufacturing

Figure 2 provides the size distribution of plants in the 1997 U.S. Census of Manufacturing. This includes data on administrative plants. The vertical axis is the percent of plants, and the horizontal axis is the level of plant employment (on a natural log scale).

Finally, we will also use statistics from Davis et al. (1998) on the distribution of job creation and destruction rates in U.S. manufacturing from 1973–1988. Figure 3 plots these annual rates, which are bounded between -2 (exit) and +2 (entry) because they represent 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 creation or destruction bin.




Figure 3: Job Creation and Destruction Rates in U.S. Manufacturing
(via Davis, Haltiwanger and Schuh, 1998)

4. Indirect Inference

We now proceed to 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, market share, size distribution, and growth distribution 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 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 (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 but not necessarily more protected against creative destruction (e.g., Apple has a few, high-quality product lines that are vulnerable to innovation by Samsung and others). The innovation and arrival rates for the parameters in Table 1 should therefore be linked to the moments we observed in the U.S. manufacturing data.

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 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 occurring per variety it owns, as in Table 1. New varieties by entrants draw a random quality from the distribution in the population. 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, whichever comes first.
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 in the economy 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 moments in step 5. and the empirical moments.

Table 3: Inferred Parameter Values

Channel	Probability	Step Size
Own-variety improvements by incumbents	71.0%	1.023
Creative destruction by entrants	6.7%	1.023
Creative destruction by incumbents	93.3%	1.023
New varieties from entrants	0.5%	0.90
New varieties from incumbents	0.05%	0.90

Sources of growth

We present our inferred parameter values in Table 3. As we will describe, these parameter values produce simulated moments closer to the empirical moments than the alternatives we consider with no own-variety improvements by incumbents, no new varieties, and/or symmetric market shares for all varieties.

We infer a 71 percent arrival rate for own-variety quality improvements by incumbents. Quality improvements through creative destruction occur the other 29 percent of the time – mostly by other incumbents (93 percent conditional probability) rather than by entrants (7 percent conditional probability). Thus every variety is improved in every year – a corner solution for which we will provide intuition shortly. Quality improvements occur in 2.3 percent jumps. Average quality does not rise by 2.3 percent per year, however, because new varieties average only 90 percent of the average quality of existing varieties. These new varieties arrive at a slow rate, boosting total variety by 0.55 percent per year, and come mostly from entrants (0.50 percent) rather than from incumbents (0.05 percent).

Table 4 presents the implied sources of growth. About 35 percent of growth comes from creative destruction. Own-variety improvements by incumbents seem just as important, at 44 percent. New varieties *a la* Romer (1990) are the

Table 4: Inferred Sources of Growth

	Entrants	Incumbents	
Creative destruction of existing varieties	7.3%	27.2%	34.5%
Creation of new varieties	16.3%	5.2%	21.5%
Own-variety improvements	-	44.1%	44.1%
	23.6%	76.4%	

remainder at about 21 percent. All three sources matter, but no single channel dominates under the Table 3 parameter values. Incumbents contribute more to growth (76 percent) than do entrants (24 percent). Aghion et al. (2013) 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 5. We start with the baseline KK model, which features $\sigma = 1$ and only creative destruction. Then we generalize the KK model to $\sigma > 1$, which allows high quality varieties to have higher market shares. We next 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. The step size of innovation is bigger in the models without own-variety improvements (at 5 or 6 percent) than it is in the model with own-variety improvements (closer to 2 percent). The arrival rates are correspondingly lower. About 43 percent of varieties are subject to creative destruction annually. In contrast, quality improvements occur on all varieties in each year in the model with own-quality improvements. Fitting the share of plants by age requires fitting the average exit rate, which is tied to the rate of creative destruction and the share of firms with few varieties.

Table 5: Simulated Models

	KK	KK 3	New var.	General
σ	1	3	3	3
creative destruction by entr.	✓	✓	✓	✓
creative destruction by inc.	✓	✓	✓	✓
new varieties from entr.s			✓	✓
new varieties from inc.			✓	✓
own-variety improvements by inc.				✓

Survivor growth, meanwhile, is tied to the step size. To maintain a realistic exit rate, the rate of creative destruction cannot be too high. To explain aggregate growth conditional on that pace of creative destruction, the step size must be larger.

Figure 4 plots the empirical fraction of plants by age against the density of firms by age in each model. The fit is similar for all of the models; all of the models overstate the fraction of firms 30 years and older. Figure 5 shows that the fit is better, and equally good for all of the models, for the *employment share* of plants by age. Thus the models cannot be easily distinguished based on entry, exit, and survivor growth rates alone.

Figure 6 contrasts the exit rates by size in the KK models with that in the data. In the original KK 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 sharply with size. We can soften this prediction by generalizing the KK model to $\sigma > 1$, so that big firms are also ones with higher quality rather than more varieties. But Figure 6 shows that very big firms still have low exit rates. This is because, with undirected innovation, the way to get very big is still by accumulating many varieties. We need quality

Table 6: Parameter Values in the Simulated Models

Parameters		KK		KK 3		New varieties		General	
λ_i	s_q	-	-	-	-	-	-	71.0%	1.023
δ_e	s_q	2.4%	1.058	2.3%	1.057	1.9%	1.051	6.7%	1.023
δ_i	s_q	41%	1.058	41%	1.057	41%	1.051	93.3%	1.023
κ_e	s_κ	-	-	-	-	0.5%	1.051	0.5%	0.90
κ_i	s_κ	-	-	-	-	0.05%	1.051	0.05%	0.90




Figure 4: Model Fit, Fraction of Firms by Age




Figure 5: Model Fit, Employment Share of Firms by Age




Figure 6: Model Fit, Exit by Size with Undirected Innovation

variation to be the dominant source of size variation to (realistically) weaken the impact of size on exit.⁶

Figure 7 plots exit by size for the two models with new varieties and directed innovation.⁷ Directed innovation is a powerful force generating quality variation, so it succeeds in flattening the relationship between exit and size. When large firms have high quality varieties more than a diverse array of varieties, they are only partially insulated from the threat of creative destruction. The fit

⁶This is why we inferred little creation of new varieties by incumbents in Table 3. If they grow by creating more varieties, then their exit rate will fall too quickly.

⁷To facilitate comparison with Klette and Kortum (2004), we stick with undirected innovation in the KK models. When we add new varieties, however, we optimize the number of quantiles at 12. When we add own-variety improvements, we infer even more directed innovation with 100 quantiles.




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

is better for the model with own-variety improvements by incumbents. As we shall see, this is because quality dispersion is excessive in the models without own-variety improvements.

As mentioned above, models with undirected innovation generate thick tails of job creation and destruction. We illustrate this in Figure 8 for the KK model with $\sigma = 3$. The models with directed innovation do much better in matching the empirical distribution of job creation and destruction – see Figures 9 and 10.⁸

The new varieties model and the general model do *not* do equally well in

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

terms of the size distribution. Figure 11 compares the percent of plants by employment in the two models to that in the data. Both models match the data's mean workers per plant of 45 by construction. But the dispersion of firm size is too great in both models relative to the data. This is more acute for the new varieties model – which lacks own-variety improvements. As described earlier, when growth is driven by creative destruction, the arrival rate of innovations is lower and the step size is bigger. The combination of an arrival rate closer to 0.5 (43 percent vs. 100 percent with own innovations) and a bigger step size (6 percent vs. 2 percent) generates more size dispersion in the creative destruction model. We could change parameter values to reduce this problem in the creative destruction model, but it would compromise the model's fit on other dimensions of the data, especially the distribution of job creation and destruction. To recap, a model in which there are frequent small quality improvements – just as much through own-improvements as through creative destruction – comes closest to the data.

5. Conclusion

How much of innovation takes the form of creative destruction? Vs. firms improving their own products? Vs. 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 just as important. Of lesser but nontrivial importance are new varieties and the contribution of entrants overall.

Our findings could be relevant for optimal innovation policy. The proximate sources of growth we identify may have different implications for business stealing effects vs. knowledge spillovers and hence the social vs. private return to innovation. The importance of creative destruction relates to political




Figure 8: Model Fit, Job Creation and Destruction




Figure 9: Model Fit, Job Creation and Destruction




Figure 10: Model Fit, Job Creation and Destruction




Figure 11: Model Fit, Firm Size Distribution

economy theories of incumbents blocking entry, 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 Indian manufacturing incumbents may be less important for innovation and growth in India given that surviving incumbents do not expand as much 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. 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 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.

References

- Acemoglu, Daron, *Introduction to Modern Economic Growth*, Vol. 1, Princeton University Press, 2011.
- and James Robinson, “Why Nations Fail: Origins of Power, Poverty and Prosperity,” 2012.
- , Ufuk Akcigit, Nicholas Bloom, and William Kerr, “Innovation, Reallocation and Growth,” April 2013.
- Aghion, Philippe and Peter Howitt, “A Model of Growth Through Creative Destruction,” *Econometrica*, 1992, 60 (2), pp. 323–351.
- , Ufuk Akcigit, and Peter Howitt, “What Do We Learn From Schumpeterian Growth Theory?,” Technical Report, National Bureau of Economic Research 2013.
- Akcigit, Ufuk and William R. Kerr, “Growth through heterogeneous innovations,” Research Discussion Papers 28/2013, Bank of Finland November 2013.
- Baily, Martin Neil, Charles Hulten, and David Campbell, “Productivity Dynamics in Manufacturing Plants,” 1992.
- Broda, Christian and David E Weinstein, “Product Creation and Destruction: Evidence and Implications,” *The 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, 1.
- Foster, Lucia, John C Haltiwanger, and Cornell John Krizan, “Aggregate productivity growth. Lessons from microeconomic evidence,” in “New developments in productivity analysis,” University of Chicago Press, 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,” *National Bureau of Economic Research Books*, 1990.

- Gourio, Francois and Leena Rudanko, "Customer capital," *The Review of Economic Studies*, forthcoming, 2014.
- Greenwood, Jeremy, Zvi Hercowitz, and Per Krusell, "Long-run implications of investment-specific technological change," *The American Economic Review*, 1997, pp. 342–362.
- Grossman, Gene M and Elhanan Helpman, "Quality ladders in the theory of growth," *The Review of Economic Studies*, 1991, 58 (1), 43–61.
- Hsieh, Chang-Tai and Peter J Klenow, "Misallocation and manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- and Peter J. Klenow, "The Life Cycle of Plants in India and Mexico," *The Quarterly Journal of Economics*, forthcoming, 2014.
- Jones, Charles I, "Growth and ideas," *Handbook of economic growth*, 2005, 1, 1063–1111.
- Jovanovic, Boyan, "Selection and the Evolution of Industry," *Econometrica: Journal of the Econometric Society*, 1982, pp. 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, pp. 1389–1419.
- Krusell, Per and Jose-Victor Rios-Rull, "Vested interests in a positive theory of stagnation and growth," *The Review of Economic Studies*, 1996, 63 (2), 301–329.
- 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*, Vol. 3, MIT press, 2002.
- Peters, Michael, "Heterogeneous mark-ups, growth and endogenous misallocation," LSE Research Online Documents on Economics 54254, London School of Economics and Political Science, LSE Library September 2013.

Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, 1990, 98 (5), pp. S71–S102.

Schumpeter, Joseph Alois, *Business cycles*, Vol. 1, Cambridge Univ Press, 1939.

Stokey, Nancy L, "Learning by doing and the introduction of new goods," *The Journal of Political Economy*, 1988, pp. 701–717.