

Growth through Heterogeneous Innovations

Ufuk Akcigit

University of Chicago, National Bureau of Economic Research, and Centre for Economic Policy Research

William R. Kerr

Harvard University and National Bureau of Economic Research

We build a tractable growth model in which multiproduct incumbents invest in internal innovations to improve their existing products, while new entrants and incumbents invest in external innovations to acquire new product lines. External and internal innovations generate heterogeneous innovation qualities, and firm size affects innovation incentives. We analyze how different types of innovation contribute to economic growth and the role of the firm size distribution. Our model aligns with many observed empirical regularities, and we quantify our framework with Census Bureau and patent data for US firms. Internal innovation scales moderately faster with firm size than external innovation.

I. Introduction

Innovations differ substantially in their qualities, from major breakthroughs to small incremental refinements. Innovations also differ in

We thank Daron Acemoglu, John Haltiwanger, Dirk Krueger, Rasmus Lentz, Matt Mitchell, anonymous referees, and many seminar participants for their insights. We especially thank Sam Kortum for his comments and guidance on the work. Kerr is a research associate of the Bank of Finland and thanks the bank for hosting him during a portion of this research. Selman Erol, Kaushik Ghosh, and especially Harun Alp provided excellent research assistance on this project. The research in this paper was conducted while the authors were Special Sworn Status researchers of the US Census Bureau at the Boston Census

Electronically published July 18, 2018

[*Journal of Political Economy*, 2018, vol. 126, no. 4]

© 2018 by The University of Chicago. All rights reserved. 0022-3808/2018/12604-0001\$10.00

their types: Many innovations help firms improve their existing portfolio of products or technologies, while others expand the portfolios of firms and enable them to enter into new markets. How do innovation qualities and types relate to firm characteristics? Some accounts emphasize the many great breakthroughs of independent entrepreneurs, while others describe the financial might and longer investment horizons that large companies can take toward innovation. Either way, a Silicon Valley start-up will behave very differently from the R&D laboratory of General Electric. These observations lead to important questions: Are there innovation differences between large and small firms and, if so, how substantial are the gaps? How do firms change their innovation strategies over their life cycles? What are the aggregate implications of different-sized firms producing heterogeneous innovations and spillovers?

This paper is a major attempt to answer these questions empirically, theoretically, and quantitatively using a fully specified endogenous growth model. Despite many advances, growth theory mostly provides frameworks that include a single type of innovation, perhaps drawn from a distribution, but not the variation in types that empirical work has uncovered. Similarly, the firm size distribution is rarely important for how these growth models function. Our framework allows for heterogeneity along both dimensions and links them together. We describe an economy with firms of multiple sizes that pursue different types of innovations and affect growth in different ways. As a result, the model allows for different-sized firms to generate multiple forms of innovation that have different spillovers.

The model of Klette and Kortum (2004) provides a first step in this effort. Their framework allows firms to own multiple product lines that are added or lost on the basis of innovation and creative destruction forces. Klette and Kortum (2004) and Lentz and Mortensen (2008) show that this setup exhibits many behaviors consistent with the applied micro literature (e.g., skewness of the firm size distribution, greater growth volatility of small firms). Following Lentz and Mortensen (2008), many researchers use this powerful platform for applied growth theory, and we use it ourselves in Acemoglu et al. (2016, forthcoming). This framework does not, however, incorporate heterogeneous types of innovation, and innovation decisions are uniform across the firm size distribution (indeed, the model's perfect scaling of innovation choices with firm size underlies the framework's analytical beauty).

Research Data Center. We gratefully acknowledge financial support from the Alfred P. Sloan Foundation, the Ewing Marion Kauffman Foundation, Harvard Business School, the Innovation Policy and the Economy forum, and the National Science Foundation. Research results and conclusions expressed are the authors' and do not necessarily reflect the views of the Census Bureau or NSF. This paper has been screened to ensure that no confidential data are revealed. Data are provided as supplementary material online.

We introduce into this framework new heterogeneity in the types of innovations undertaken by firms, which in turn shapes how the firm size distribution can matter for the economy. We distinguish two types of innovation that firms undertake: *external* and *internal*. Firms undertake external innovations to create new products and capture markets from others, while internal innovations improve product lines that firms currently own. This heterogeneity in the forms of innovation and the step sizes of associated advances is central in accounts of the differences in innovation for large and small firms and yet is not included in prior growth models.¹

Our paper makes three key advances. The first is to build a growth model that incorporates multiple forms of innovation, a direct connection from firm size to choices over types of innovations, and multiple step sizes in the impact of innovations that are endogenously determined. Our baseline model analyzes a setting in which internal innovations scale up with firm size, while external innovations do not. The tractable model yields analytical solutions and stark predictions about how the innovations of new entrants and small firms will differ from large firms. The model provides microfounded explanations for small firms experiencing faster average growth and contributing disproportionately to major innovations.

The second contribution is to incorporate patents and patent citations into our endogenous growth framework, which allows us to connect endogenous growth theory to the empirical innovation literature (e.g., Griliches 1990). Using findings from this empirical literature, we characterize how patent citations would look in our economy and show how citation patterns hold information relevant to the model. While these additions do not affect the model's economy directly (e.g., firms do not block rivals with patents), citations provide greater depth to the results that we can characterize. For example, we derive tests that employ patent citations to compare the growth spillover effects from external and internal innovations. Moreover, distributions of patent citations contain much of the information that we need to quantify the model.

Our third contribution is a generalized framework that allows an arbitrary amount of scaling for external innovation with firm size (internal innovation always scales fully). At the extremes of this generalized framework are the extended Klette and Kortum (2004) framework (perfect scaling) and our baseline model (no scaling). We quantify the model using indirect inference with Census Bureau data on all patenting firms

¹ We mostly use the terms R&D and innovation interchangeably, favoring the latter. Strictly speaking, firms make R&D investments and realize innovation outcomes, and these are not perfectly correlated because of randomness in achieving results. Nevertheless, our model makes similar predictions for both objects given their tightly coupled nature.

during the 1982–97 period. We observe moderate departures from the Klette and Kortum world for the United States. However, we also find that this departure could generate some sizable cost increases for large firms. In particular, according to our estimates, it costs 25 percent more for a firm that is at the 90th percentile of the size distribution to produce a major innovation than the median innovative firm in the economy.

Our analysis helps inform long-standing debates about the role of small versus large firms for innovation. For instance, we show that the relative rate of major inventions is higher in small firms. We demonstrate that these distributional differences are not due to differences in research capabilities or technologies, but are instead an outcome of innovation investment choices by firms. We also decompose the aggregate growth due to innovation and find that 19.8 percent is due to internal efforts of incumbents, 54.5 percent to external efforts of incumbents, and 25.7 percent to new entrants.

In terms of the literature, we most clearly build on the efforts of Klette and Kortum (2004), Lentz and Mortensen (2008), and Akgigit (2010) to incorporate more insights from the empirical literature on innovation into workhorse theoretical models.² These papers in turn depend on the long endogenous growth literature.³ Our work on spillover benefits builds on contributions such as Spence (1984) and Griliches (1992), with Caballero and Jaffe (1993) and Eeckhout and Jovanovic (2002) being rare examples that connect patent citations to a growth model. We are the first to do so at a firm level and with a focus on identifying varieties of innovation. Finally, we are deeply connected to the empirical literature on firm size and innovation that we review in the next section as a prelude to our model.⁴

II. Empirics of Innovation

This section reviews prior work on the differences in innovation across the firm size distribution. We then document three empirical regulari-

² More recent contributions are Lentz and Mortensen (2014), Garcia-Macia, Hsieh, and Klenow (2016), and Akgigit, Hanley, and Serrano-Velarde (2016).

³ Classics include Romer (1986, 1990), Grossman and Helpman (1991), Aghion and Howitt (1992), Jones (1995), Aghion, Howitt, and Vickers (1997), Kortum (1997), Howitt (1999), and Aghion et al. (2001). Barro and Sala-i-Martin (1995), Acemoglu (2008), Aghion, Akgigit, and Howitt (2014), and Akgigit (2017) provide full reviews.

⁴ Our work likewise relates to the economics literatures on innovation and industry structure and evolution. Examples include Gilbert and Newbery (1982), Jovanovic (1982), Reinganum (1983), Hausman, Hall, and Griliches (1984), Dunne, Roberts, and Samuelson (1988), Klepper and Graddy (1990), Rosen (1991), Hopenhayn (1992), Jovanovic and MacDonald (1994), Cohen (1995), Lerner (1997), Gans, Hsu, and Stern (2002), Cabral and Mata (2003), Hopenhayn, Llobet, and Mitchell (2006), Duranton (2007), Luttmer (2007, 2011), Cai (2010), Kerr (2010), Arkolakis (2011), Lamoreaux, Sokoloff, and Sutthiphisal (2011), Acemoglu and Akgigit (2012), Bloom, Schankerman, and Van Reenen (2013), Nicholas (2014), and Acemoglu and Cao (2015).

ties that motivate our model and are used to discipline its quantitative analysis. In later sections, we provide additional results when comparing the quantified model and empirical data on untargeted dimensions. Appendixes A–D and our NBER working paper, Akcigit and Kerr (2010), contain many empirical extensions.

A. *Innovation across the Firm Size Distribution*

A large empirical literature debates whether small or large firms contribute disproportionately as the source of radical innovations or achieve a greater innovation return per R&D dollar invested.⁵ Our model attempts to address these questions using a novel approach. Our framework would be extremely uninteresting if we endowed firms of various sizes with capabilities not available to others (e.g., assuming that small firms could achieve new breakthrough improvements not possible for larger firms). Instead, we trace out why large and small firms might invest at different rates in the same set of potential innovation approaches, with the heterogeneous innovations being an outcome rather than an assumption.

We focus on internal versus external innovation as it aligns with many important empirical insights and it is the type of heterogeneity that we can measure most directly with data. At an extreme, external versus internal differences must exist. Entering entrepreneurs do not have products to improve on and so by definition are different from incumbents. The literature further suggests this difference is pervasive, rather than confined to the entry margin, and usually emphasizes a greater internal focus for large firms.

Large firms might invest more in internal improvements since they can derive a better return from these investments than small firms. In situations in which innovations are useful for enhancing a company's operations but are otherwise hard to protect/sell, large companies achieve a greater return for the same investment due to their larger base of operations. These incentive differences are frequently discussed for process innovations (e.g., Klepper 1996), and Cohen and Klepper (1996) show that process R&D is more tightly linked to firm size than product R&D. While these patterns are consistent with internal innovation scaling more directly with firm size, at least one counterexample exists. Basic R&D is also more likely to be conducted by large firms because of the fixed costs of basic R&D laboratories and the ability to realize resulting discoveries across a

⁵ See, e.g., Nelson and Winter (1982), Acs and Audretsch (1987, 1988, 1991), Rosen (1991), Peretto (1998), Zucker, Darby, and Brewer (1998), Baumol (2002), Thomke (2003), Rausch (2010), Samila and Sorenson (2011), Kerr, Nanda, and Rhodes-Kropf (2014), Kueng, Yang, and Hong (2014), and Acemoglu et al. (2016). Of the efforts to quantify these claims, the best known is the Kortum and Lerner (2000) finding that venture capital dollars invested in small start-ups are three times more potent for generating patented innovations than corporate R&D expenditures.

range of products. To the extent that basic R&D also provides serendipitous advances that aid entry into new industries outside of the firm's current span, larger companies garner more external innovation.

An additional class of explanations for why large companies may pursue proportionately less external innovation relates to organizational frictions and managerial capabilities. Under the Lucas (1978) span-of-control model, there are limits to the number of operations that the world's best managers can effectively guide, and thus large companies might endogenously invest more in improving their existing products versus further expansion. These limits to optimal firm size would effectively give a comparative advantage to small firms for pursuing the acquisition of new lines.⁶ Related models in Gromb and Scharfstein (2002) and Hellmann and Perotti (2011) emphasize situations in which the internal resources of large companies can be necessary for completing innovations. On the other hand, the management literature frequently stresses organizational rigidities that inefficiently inhibit the external innovation efforts of large companies (e.g., Henderson and Clark 1990; March 1991; Henderson 1993; Christensen 1997).⁷

External environments also shape innovation incentives for large companies, with financial markets being a well-studied example. Bernstein (2015) finds that being a publicly listed firm reduces the novelty of a firm's innovations by 40 percent and shifts work toward more conventional and internal projects, while perhaps offering additional funds for acquisitions. Lerner, Sorensen, and Strömberg (2011) reach similar conclusions when examining the impact of private equity firms on the innovation rates of the firms that they remove from public markets. Other studies find that conglomerate firms frequently trade at a discount and that managers often reduce R&D budgets to meet short-term return targets. Thus, while deep capital markets may provide valuable resources to public companies, they appear to create environments less attractive for external innovation.⁸

⁶ Akcigit, Alp, and Peters (2015) introduce a span-of-control limitation for firms into the Klette and Kortum (2004) framework to study the firm dynamics in the Indian manufacturing sector. Their model emphasizes how these managerial conditions create limits to the scaling of firms and their pursuit of new product lines, connecting to the empirical work of Hsieh and Klenow (2014).

⁷ Galasso and Simcoe (2011) identify how CEO personality traits shape innovation investments, and Lerner (2012) further reviews the recent literature on the advantages and liabilities of large companies for pursuing new innovation areas compared to start-ups (e.g., compensation constraints).

⁸ Differences beyond financial markets also exist. Agrawal, Cockburn, and Rosell (2010) consider how large companies may be located in more isolated cities that limit the diversity of external ideas that they receive and can build on. Some industries are also characterized by a market for ideas (e.g., Gans et al. 2002) that shifts the organization of innovation for external work. Finally, policies with firm-size-dependent components like labor regulations may make external innovation less attractive for large companies to the extent that policies make the labor adjustments associated with risky activities more costly for larger employers.

This brief literature discussion highlights why the internal versus external distinction is likely to be important. Several data sources are consistent with this observation:⁹

- Using the 2008 Business R&D and Innovation Survey, we observe a $-.16$ correlation between firm size and the share of R&D that the firm reports is directed toward business areas and products where the company does not have existing revenues. Similar negative correlations are found for questions about the share of firm R&D being directed to technologies new to markets.
- Using the 1979–89 National Science Foundation (NSF) R&D Surveys that recorded product versus process R&D expenditures, we observe a $.22$ correlation between firm size and the share of R&D that the firm reports is process oriented. This accords with Cohen and Klepper (1996), and we find similar results for indicator variables about the firm conducting any process-focused R&D.
- Using the citations that firms make on the patents they file, we observe a $.11$ correlation between firm size and the share of citations given that are to a firm's own prior patents. Firms with larger past patent portfolios are mechanically more likely to self-cite, and Section A of appendix C reports Monte Carlo simulations that measure the expected likelihood of self-citations given the technology and years that a firm cites in its patents. Larger firms are more likely to show abnormal rates of self-citations compared to these counterfactuals, with the correlation to firm size of being out of the simulated 95th percentile bound being $.23$.

These correlations point toward a consistent picture of heterogeneity in innovation behavior by firm size. An advantage of our model is its capacity to place these data pieces into context and use indirect inference for more general statements.

B. Data Development

Our project employs the Longitudinal Business Database (LBD) and the NBER Patent Database. The LBD is a business registry for the United States that contains annual observations for every private-sector establishment with payroll from 1976 onward (Jarmin and Miranda 2002). The Census Bureau data are an unparalleled laboratory for studying the firm size distribution, entry/exit rates, and life cycles of US firms. Sourced from US

⁹ All reported correlations measure firm size through log employment and are statistically significant at a 5 percent level. The correlations are taken over reported data in each survey, and some of these sources have incomplete coverage for small R&D producers, as described in our working paper (Akcigit and Kerr 2010). These sample constraints likely weaken the observed correlations to firm size.

tax records and Census Bureau surveys, the micro records document the universe of establishments and firms rather than a stratified random sample or published aggregate tabulations. We aggregate establishment-level records into firm-year observations using parent firm identifiers.

We next match into the LBD the individual records of all patents granted by the US Patent and Trademark Office (USPTO) from January 1975 to May 2008. Each patent record provides information about the invention and the inventors submitting the application. Hall, Jaffe, and Trajtenberg (2001) provide extensive details about these data, and Griliches (1990) surveys the use of patents as economic indicators of technology advancement. We employ only patents (1) filed by inventors living in the United States at the time of the patent application and (2) assigned to industrial firms. In 1997, this group comprised about 77,000 patents (40 percent of the total USPTO patent count in 1997, with most of the residual being patents to foreign inventors). We match these patent data to the LBD using firm name and location matching algorithms that build on Kerr and Fu (2008) and Balasubramanian and Sivadasan (2011).¹⁰

Our final sample is the universe of patenting US firms with employees, composed of 23,927 firms that have been granted at least one patent by the USPTO over the 1982–97 period. We use earlier and later time periods for calculating some of our metrics on these firms. This data set is the foundation for our empirical estimates in this section and also our quantitative analysis in later sections. There are several important features about this data set to highlight.

First, our sample includes only innovative firms, which have a different firm size distribution than the economy as a whole. In our sample, for example, 14 percent of firms have more than 500 employees at some point in their life span (12 percent for all observations of the firm), while this share is about 0.3 percent for the whole economy. This tilt toward larger firms is not surprising, as the majority of small firms do not seek new innovations or to grow from their current size (Hurst and Pugsley 2011). This is often connected to nonpecuniary motivations for starting a business (e.g., to be one's own boss). We thus exclude large numbers of noninnovative firms from our sample (e.g., restaurants, beauty salons, grocery stores) to be in keeping with the model of innovative firms.¹¹

Second, only a few innovative firms patent in every year, and the same is true in our model with respect to realizing an innovation. These consider-

¹⁰ Our NBER working paper (Akcigit and Kerr 2010) describes this matching procedure and the data employed more extensively. The working paper also provides complementary evidence from the NSF's R&D Survey that supports the patent-based results provided here. The NSF survey subsamples R&D performers that conduct less than \$1 million in R&D annually, and thus our focus on patenting allows us greater confidence for capturing the complete firm size distribution for innovative firms.

¹¹ Approximate 25th, 50th, and 75th percentile levels of employment in our sample are 17, 70, and 370 employees. These are "fuzzy" averages around these points in order to satisfy Census Bureau disclosure requirements. The mean employment level is about 1,805 workers.

ations lead us to use our data in two ways. In some cases (e.g., Gibrat's law estimations), we conduct an annual analysis as the necessary data elements are continually observed in both the data and the model. In other cases (e.g., quality distributions of realized innovations), we focus on 5-year periods and the firms achieving innovations as depicted below. Our quantitative model exactly mirrors each data development step undertaken to ensure that we precisely align the model with the data. This mirroring technique has the powerful advantage of allowing us to select the approach that best suits each prediction, accounting for the nuances of the data assembled.

Sample selection is very important, and we align the data and model as much as possible and subject to the same treatment. Many parts of this effort are quite straightforward. First, we define metrics the same in both data sets (e.g., how exit is coded). Second, we ensure in both data sets that we measure moments in comparable groups. For example, most innovation-related moments are measured across continuously innovative firms in both data sets to ensure comparability. Third, our model is one in which every firm invests in R&D and attempts to grow, and thus we earlier noted the boundary condition that we are not attempting to model firms that do not seek to develop new ideas or expand (e.g., Hurst and Pugsley 2011; Akcigit et al. 2015). This too is true in both the empirical work and model estimation.

There is one selection margin, however, that is more challenging and worth additional comment and checking. While the model is built on firms always investing in R&D, success in these efforts is stochastic and thus some R&D firms do not realize innovations. In the empirical data, by contrast, the best and most comprehensive approach to identifying innovative firms is to use patenting, as this does not encounter truncation biases common to R&D surveys that subsample firms below a threshold of R&D expenditure. However, our selection process does mean that every firm in our sample has achieved at least one patent in our sample period, which is not strictly true in the model (e.g., a firm might have achieved an innovation in the distant past that allowed it to enter the market but it has not been successful in further efforts).

Several auxiliary tests suggest that this is not a first-order concern. First, when we impose the equivalent of a patent-like selection criterion in our simulation period, we retain almost 99 percent of our sample, and hence our simulated moments are essentially unchanged. Second, as most moments are aggregated or calculated over our continuously innovative samples, the greatest potential sensitivity to this feature is the upcoming estimation of Gibrat's law in Section II.C. An additional form of assurance comes in that we observe a very similar growth-size relationship when expanding our sample to the broader manufacturing sector without imposing any selection criteria—specifically, our sample's coefficient of -0.0351 (0.0013) becomes -0.0495 (0.0004) in the full sectorwide estimation—suggesting broad stability on this particular margin.



FIG. 1.—Firm growth by firm size. Color version available as an online enhancement.

C. Firm Growth by Firm Size

We first document the empirical regularity that small firms grow faster than large firms.¹² We test this prediction using annual employment growth patterns for US innovative firms. Following Lentz and Mortensen (2008), we define for firm f the employment growth of $\text{EmpGr}_{f,t} = (\text{Emp}_{f,t+1} - \text{Emp}_{f,t})/\text{Emp}_{f,t}$. We model employment growth without conditioning on survival and thus retain $\text{EmpGr}_{f,t} = -1$ for businesses that close between t and $t + 1$ (the LBD measures employment in March of each year). This metric is unbounded upwardly, and we impose a 1,000 percent growth cap. With this winsorization, the mean of $\text{EmpGr}_{f,t}$ is 0.0745.

Dividing our sample into 20 roughly equal-sized bins in terms of numbers of firms by current employment levels, figure 1 displays the average forward growth rate across firms in each size bin. The horizontal axis provides the average employment level for firms in the bin. The declines in

¹² The empirical deviation from Gibrat's law of proportionate growth is extensively documented in surveys such as Sutton (1997), Caves (1998), and Geroski (1998) and is among the stylized facts in Klette and Kortum (2004). The Klette and Kortum model yields Gibrat's law. Lentz and Mortensen (2008) show that the addition of firm heterogeneity into the Klette and Kortum model is consistent with deviations from proportionate growth observed in Danish firm-level data.

average forward growth are substantial until about 50 employees, and they are again strong at the largest firm sizes. Section D of appendix C shows a negative relationship when using the establishment counts of firms to generate firm size bins. Size bins based on establishment counts are substantially coarser than employment but provide a complementary approach.

To provide a single estimate and also control for industry-year fixed effects $\eta_{i,t}$ we estimate

$$\text{EmpGr}_{f,t} = \eta_{i,t} - \underset{\text{(s.e. 0.0013)}}{0.0351} \cdot \ln(\text{Emp}_{f,t}) + \epsilon_{f,t}.$$

This coefficient finds that a 10 percent increase in firm employment is associated with a 0.35 percent reduction in forward employment growth, or about 5 percent of the sample mean. The growth impact of the interquartile range of firm size (approximately 17 to 370 employees) is 10.8 percent, somewhat larger than the mean.¹³ This relationship is robust to alternative measures of firm size, weighting observations, or considering panel variation, reflecting the many settings in which it has been observed in prior research. Conditional growth estimations that exclude exiting firms yield a steeper negative relationship, as does raising the maximum growth rate (discussed further below in model robustness checks). When using the Davis, Haltiwanger, and Schuh (1996) formula that compares growth to the average of the two periods, conditional estimations also yield a consistent negative relationship across the many specification variants discussed, while unconditional estimations do not exhibit a clear pattern.

D. Innovation Intensity by Firm Size

We next consider the innovation intensity to firm size relationship. Our model will consider this intensity in terms of firm-level inputs (e.g., R&D-to-sales ratios) and realized outputs (e.g., rate of realized innovations per product line). We can discipline the model through either relationship, and for several data quality reasons we pursue the realized rate of innovation outputs.¹⁴

We study this prediction through patents per employment $\text{Patent}/\text{Empl}_{f,t}$, where the timing of patents is shown by their application year. The largest innovative firms such as Microsoft or Boeing apply for many patents each year, but most innovative firms are irregular and lumpy in their patent filings. We thus analyze this prediction with 5-year periods that extend

¹³ The regression sample includes 146,678 observations. We assign industries to firms at the two-digit level of the Standard Industrial Classification system using industries in which firms employ the most workers. Regressions are unweighted and cluster standard errors at the firm level.

¹⁴ Our NBER working paper (Akcigit and Kerr 2010) provides complementary tabulations of R&D expenditures per sales or per employee across the firm size distribution using the NSF R&D Survey.

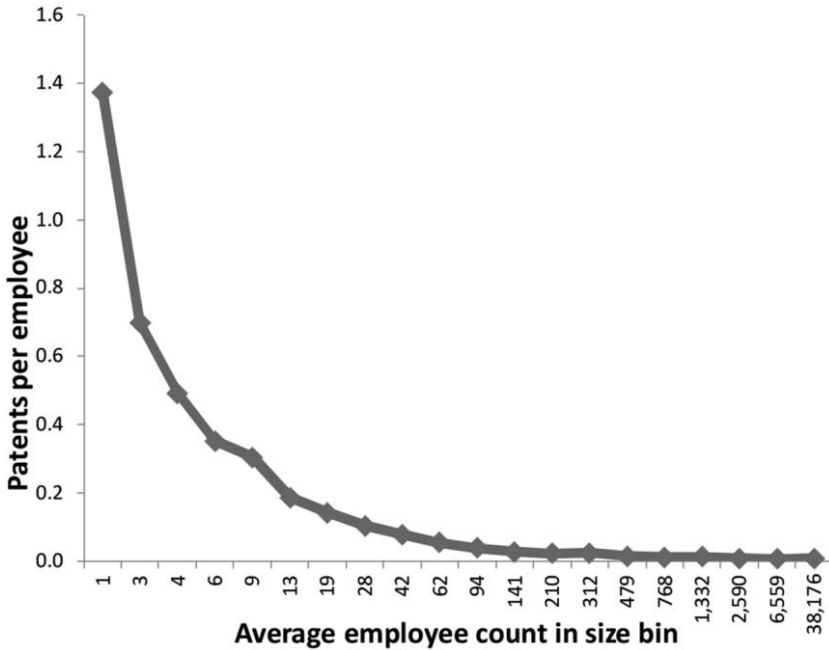


FIG. 2.—Innovation intensity by firm size. Color version available as an online enhancement.

1982–86, 1987–92, and 1992–96. (With some abuse of notation, we continue to use t to represent time periods.) We focus this exercise on “continually innovative firms” in the sense that included firms file at least one patent in each 5-year period that they are observed to be in operation. This data set includes 16,818 firm-period observations. The continuous sample approach keeps a consistent definition with respect to nonzeros and facilitates a sharper match with the model, where we also impose this requirement for included firms to be continually innovative over 5-year periods.

Figure 2 shows the empirical relationship in which we again divide our sample into 20 size bins. There is a substantial decline in innovation intensity with firm size among the continuously innovative firms. Section D of appendix C again shows a similar relationship when using establishment counts to develop size bins.¹⁵

¹⁵ The restriction to continuously innovative firms will influence this relationship. For example, the sample excludes firms that attempt to innovate but fail to achieve a patent, be they small firms or larger ones that are only marginally innovative. Our quantitative analysis handles this feature by treating the simulated data in an identical way for this relationship and selecting firms that achieve innovations in each time interval we use for equivalent estimations.

To prepare for the future matching of our data moment to the model, we transform Patent/Empl_{*f,t*} to be of mean zero and unit standard deviation during each period. We use the transformed series because the exact level of US patenting per employee does not have a direct meaning or counterpart to the levels of a theoretical model. By placing both data and model outcomes into unit standard deviations, we are able to match and compare them. Our key estimation is

$$\text{Patent/Empl}_{f,t} = \eta_{i,t} - \underset{\text{(s.e. 0.0058)}}{0.1816} \cdot \ln(\text{Empl}_{f,t}) + \epsilon_{f,t}.$$

This coefficient finds that a 10 percent increase in firm employment is associated with a reduction of 0.018 standard deviations in patents per employee among innovative firms. Across the interquartile range of firm sizes, the impact is 0.561 standard deviations. If we relax the continuous innovator sample restriction, the coefficient is very similar at -0.164 . We also find robust results with the many regression variants discussed above with the employment growth specifications.

E. Fraction of Major Innovations by Firm Size

Our model's structure allows for internal and external innovations to have different average impacts in terms of realized improvements on existing technologies, and the model does not require one form of innovation to be larger than the other. Nevertheless, if external innovations have a larger average impact than internal innovations, then our baseline model makes some important predictions regarding small innovative firms and new entrants having a comparative advantage for achieving major advances. If internal innovations have the larger average impact, then larger firms will hold this advantage for achieving major advances.

To investigate, figure 3 provides some empirical evidence regarding the relative impact of external versus internal innovations using patent citations. The sample is restricted to patents of US industrial firms that have all inventors located in the United States. Similarly to academic papers, patents give citations to prior patents on which the current invention builds. Going forward, the impact of a patent is often measured in terms of the citations it subsequently receives. Examining the citations given to prior patents at the time of the patent filing, we classify patents into external versus internal innovations. Internal patents are those in which 50 percent or more of the given citations are to the prior inventions of the firm filing the patent (termed self-citations). External patents are those in which self-citations represent less than 50 percent of the citations given at filing.

We measure forward impact through external citations received by the patent in the future (i.e., excluding future self-citations made by the

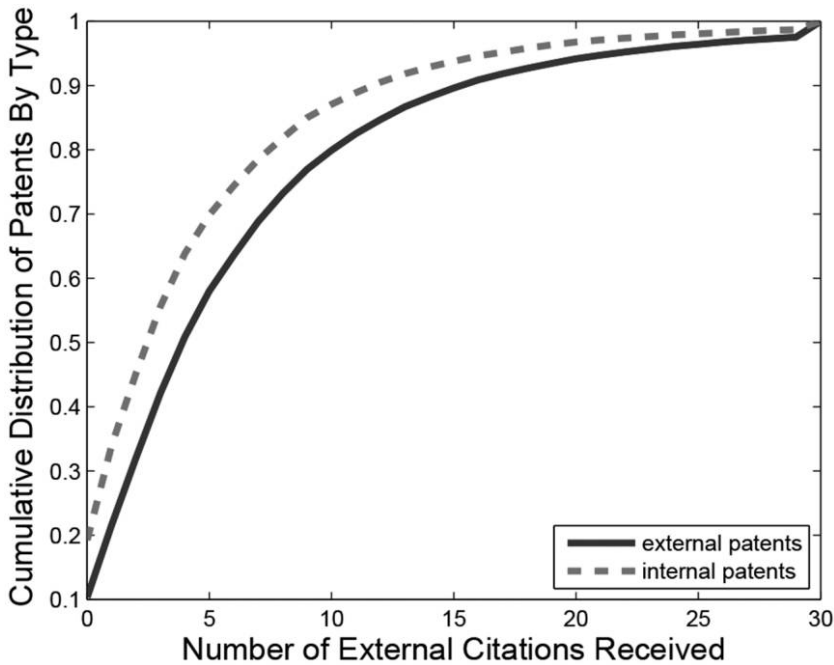


FIG. 3.—Citation distribution by patent type. Color version available as an online enhancement.

firm). The lighter dashed line in figure 3 provides the distribution of future citations received for internal patents filed between 1975 and 1984. The darker solid line provides the distribution for external patents that make no citations to prior patents of the firm. There is no mechanical reason for these two series to be different from each other as the citations given and received are distinct from each other. Both series display a large number of patents with no external citations and a skewed distribution, which are predictions of our framework. More important, the comparison of the external and internal distributions shows that the former exceeds the latter in a form akin to first-order stochastic dominance.¹⁶

With this background, we next verify that small innovative firms and new entrants have a comparative advantage for achieving major advances. We first identify the quality of each patent in terms of its external citations

¹⁶ The differences are statistically significant and hold in regressions that control for a variety of traits about the patents (e.g., technology-year fixed effects) or firm fixed effects. The omitted, middle group (i.e., patents for which backward self-citations are present but are not a majority) behaves similarly to the no-self-citation group and is excluded for visual clarity; later, we will group them with external patents for our model quantification.

TABLE 1
FIRM SIZE AND PATENT QUALITY DISTRIBUTION

	SHARE OF FIRM'S PATENTS IN QUALITY DISTRIBUTION RANGE			
	[0, 25)	[25, 50)	[50, 75)	[75, 100]
Log firm employment _{<i>t</i>}	.0027 (.0009)	.0048 (.0010)	.0000 (.0010)	-.0074 (.0012)

NOTE.—Estimates include 16,818 observations, are unweighted, and cluster standard errors by firm.

compared to its peers from the same technology class and application year. Constructing an indicator variable for the patent being in the top decile in terms of these external citations, we calculate Top Patent Share_{*f,t*} as the average of these patent-level indicators across a time period for a firm. Not surprisingly, the average of this variable is about 0.10. We then estimate this firm-level measure as a function of firm size as

$$\text{Top Patent Share}_{f,t} = \eta_{i,t} - \underset{\text{(s.e. 0.0008)}}{0.0034} \cdot \ln(\text{Emp}_{f,t}) + \epsilon_{f,t}. \quad (1)$$

This estimation finds that a 10 percent increase in firm employment is associated with a reduction of 0.034 percent in the fraction of a firm's patents among the top decile of the patent quality distribution. Relative to the sample mean, this effect is 0.34 percent. Across the interquartile range of firm sizes, the impact is 0.011, or a tenth of the sample mean.

Table 1 broadens the lens and repeats specification (1) for each quartile of the patent quality distribution using our continuous innovation sample. The first column documents the lowest-quality quartile, while the last column is the highest one; coefficients across the four specifications naturally sum to zero. Estimations again control for industry-period fixed effects. Larger firms are associated with a systematic shift in the quality of their patents out of the top quartile and into the bottom half of the distribution.¹⁷

III. Baseline Theoretical Framework

We begin with a baseline model that incorporates the empirical regularity that external R&D does not scale as fast as internal R&D with firm size. Our goal is to study the implications of this heterogeneity on the R&D, innovation, and growth dynamics of firms. To allow for analytical solutions and to build intuition, we first consider a stark environment in which external R&D does not scale at all with firm size. We then gener-

¹⁷ Our working paper (Akcigit and Kerr 2010) further uses this framework to confirm our model's hypothesis that firm size differentials weaken with more stringent citation quality thresholds because of the increasing relative importance of the stochastic nature of realized inventions.

alize the theoretical framework in Section IV to allow scaling of external R&D, with this baseline model and Klette and Kortum (2004) being extremes of the general framework. On top of this general framework, we also overlay patent citation behavior in Section V. Within this framework we can interpret data on patent citations, allowing us in Section VI to estimate parameters of R&D scaling within firms.

A. Preferences and Final Good Technology

Consider the following continuous time economy. The world admits a representative household with a logarithmic utility function

$$U = \int_0^{\infty} \exp(-\rho t) \ln C(t) dt. \quad (2)$$

The term $C(t)$ is consumption at time t , and $\rho > 0$ is the discount rate. The household is populated by a continuum of individuals with measure one. Each member is endowed with one unit of labor that is supplied inelastically.

Individuals consume a final good $Y(t)$, which is also used for R&D as discussed below. The final good is produced by labor and a continuum of intermediate goods $j \in [0, 1]$ with the production technology

$$Y(t) = \frac{L^\beta(t)}{1 - \beta} \int_0^1 q_j^\beta(t) k_j^{1-\beta}(t) dj. \quad (3)$$

In this specification, $k_j(t)$ is the quantity of intermediate good j , and $q_j(t)$ is its quality. We normalize the price of the final good Y to be one in every period without loss of generality. The final good is produced competitively with input prices taken as given. Henceforth, the time index t will be suppressed when it causes no confusion.

There is a set of firms that are producing intermediate goods and their measure, $F \in (0, 1)$, will be determined in equilibrium. Each intermediate good j is owned by a firm f . A firm is characterized by the collection of its product lines $\mathcal{J}_f = \{j : j \text{ is owned by firm } f\}$. Similarly, we denote the product (quality) portfolio of firm f by a *multiset* $\mathbf{q}_f = \{q_j : j \in \mathcal{J}_f\}$ and denote the cardinality by n_f .¹⁸ Figure 4 illustrates two firms. Firm $f = 1$ has five product lines and $f = 2$ has three product lines (i.e., $n_1 = 5$ and $n_2 = 3$).

Each intermediate good $j \in [0, 1]$ is produced with a linear technology

$$k_j = \bar{q} l_j, \quad (4)$$

¹⁸ A multiset is a generalization of a set that can contain more than one instance of the same member.




FIG. 4.—Example of firms

where l_j is the labor input and $\bar{q} \equiv \int_0^1 q_j dj$ is the average quality in the economy.

In addition to the variable cost, production requires also a fixed cost of operation $\Phi\bar{q}$ at the firm level in terms of the final good. As we will discuss later, this fixed cost avoids any nonlinearities in the firm's value function.¹⁹

Individuals work in two capacities: final good production (L) and intermediate good production (\tilde{L}). In each period, the labor market has to satisfy the constraint

$$L + \tilde{L} \leq 1. \quad (5)$$

The variable R is the total R&D spending, K is the total fixed cost paid by firms, and therefore the resource constraint of the economy is $Y = C + R + K$.

B. Research and Development

The last innovator in each product line owns the leading patent and has monopolist pricing power until being replaced by another firm. Inter-

¹⁹ See proposition 1 and the text above it for details.

mediate producers have profit incentives to improve the technologies for their existing products, thereby increasing associated quality. In addition, both incumbents and potential entrants have incentives to add new products to their portfolios through R&D competition. We now describe the innovation types, which are also illustrated in figure 5.

Internal R&D.—Incumbent firms undertake internal R&D (or innovation) to improve their existing products. To improve an existing product $j \in \mathcal{J}_f$, firm f spends

$$R_z(z_j, q_j) = \hat{\chi} z_j^{\hat{\psi}} q_j \tag{6}$$

units of the final good, where $\hat{\chi} > 0$ and $\hat{\psi} > 1$. Internal innovations are realized with the instantaneous Poisson flow rate of $z_j \geq 0$. Cost (6) is proportional to the quality of the good that the firm is improving. First, this implies that a more advanced technology has higher R&D costs. Second, as will be shown in the next section, equilibrium returns to internal innovations are linear in q_j . Therefore, the linear effects in return and cost cancel out and yield an internal innovation effort that is independent of the quality of the product line. When internal R&D is successful, the current quality improves by a multiplicative factor $\lambda > 0$ such that $q_j(t + \Delta t) = (1 + \lambda)q_j(t)$.

External R&D.—External R&D (or innovation) is undertaken by incumbents and potential new entrants to obtain technology leadership

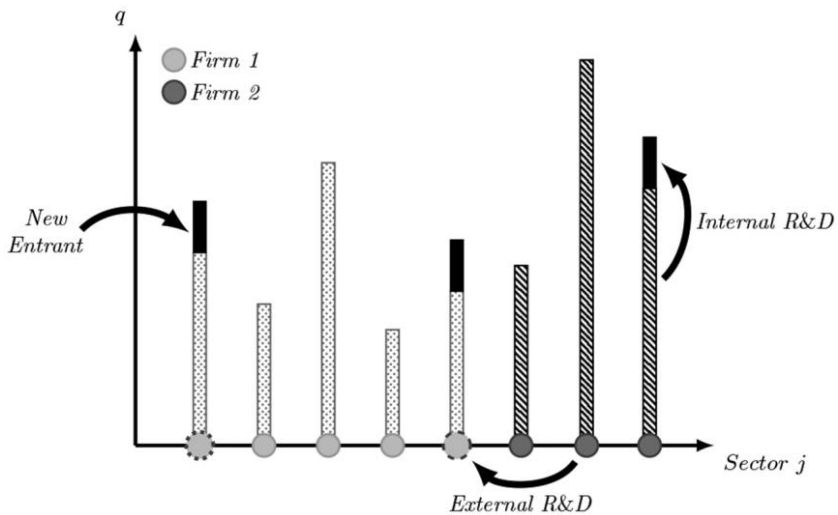


FIG. 5.—Innovation types

over products that they do not currently own. A firm with $n > 0$ produces a flow rate x by paying R_x in terms of the final good according to the following cost function:²⁰

$$R_x(x, \bar{q}) = \tilde{\chi} x^{\tilde{\psi}} \bar{q}, \quad (7)$$

where $\tilde{\chi} > 0$ and $\tilde{\psi} > 1$. Cost (7) is proportional to the average quality level \bar{q} in the economy, which again removes the dependence of innovation efforts on average quality since the returns to external innovations will be proportional to \bar{q} and ensures that the R&D spending is a constant fraction of the total output Y .

External R&D efforts are undirected in the sense that resulting innovations are realized in any product line $j \in [0, 1]$ with equal probability. This model structure has two main implications. First, firms do not innovate over their own product lines through external R&D since this event has zero probability. Second, there is no strategic interaction among firms. In addition to stochastic arrival rates, the sizes of realized quality improvements are randomly determined (see fig. 6):

- i. With probability $\theta \in (0, 1)$, the innovation is a *major advance* that substantially shifts forward the latest quality level by a size $\eta \bar{q}$ such that $q_j(t + \Delta t) = q_j(t) + \eta \bar{q}(t)$. This generates a new technology cluster with an associated wave of subsequent follow-on innovations. Prominent examples include the transistor and the map of the human genome, but the step functions need not be so profound. The conceptual construct is that these major advances define a wave of innovation and product development until another major advance starts a new wave.
- ii. With probability $1 - \theta$, the innovation is a *follow-up improvement* to the current technology level of the product line that does not generate a new technology cluster. The size of the follow-up improvement declines with the number of follow-up inventions since the last major advancement. If the last major innovation in product line j occurred k_j innovations ago, the new step size is $s_j \bar{q}$, where $s_j = \eta \alpha^{k_j}$ with $\alpha \in (0, 1)$.

²⁰ Note that the cost function in (7) corresponds to the following production function: $x = [R_x / \tilde{\chi} \bar{q}]^{1/\tilde{\psi}} \mathbf{1}_{n>0}$, where $\mathbf{1}_{n>0}$ is an indicator function. This specification implies that past innovation, i.e., $n > 0$, affords firms capacities to innovate in the future. This structure is in the same spirit as the Klette and Kortum (2004) model that assumes a Cobb-Douglas functional form: $x = R^{1/\tilde{\psi}} n^{1-(1/\tilde{\psi})}$. For now, we shut down the dependence on n at the intensive margin to prevent any scaling and just keep the dependence on the extensive margin via the indicator function.

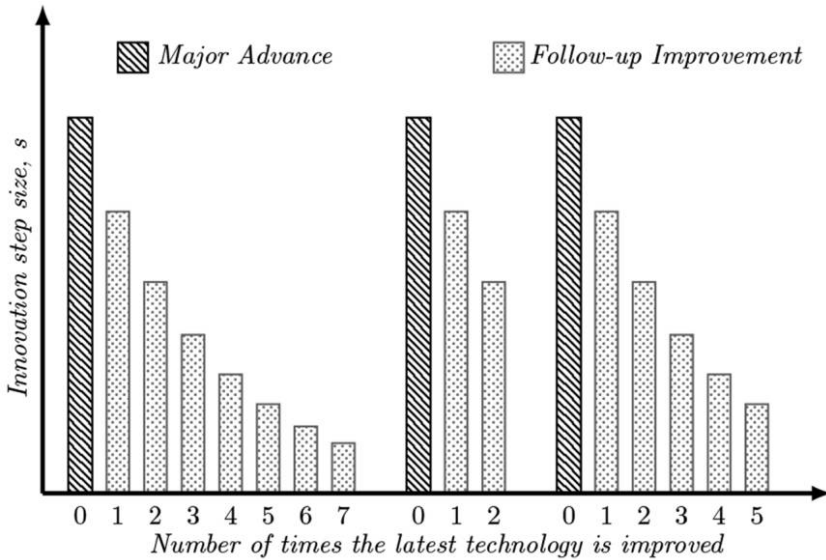


FIG. 6.—Examples of external innovations

Technology clusters and evolution.—The economywide arrival rate of new products, denoted by τ , is endogenously determined by external R&D efforts of incumbents and potential entrants and is characterized in detail below. With τ determined, the probabilistic evolution of the quality level q_j after a short interval Δt is

$$q_j(t + \Delta t) = q_j(t) + \begin{cases} \eta \bar{q}(t) & \text{with probability } \theta \tau \Delta t \\ \eta \alpha^h \bar{q}(t) & \text{with probability } (1 - \theta) \tau \Delta t \\ \lambda q_j(t) & \text{with probability } z_j \Delta t \\ 0 & \text{with probability } 1 - z_j \Delta t - \tau \Delta t. \end{cases} \quad (8)$$

The first line represents a major advance that results from external R&D with probability θ . The second line represents a follow-up innovation that results from external R&D with probability $1 - \theta$. The third line shows an internal improvement of size λ by the current owner of product line j through internal R&D. The final line represents the case in which no quality improvement is realized during Δt , which results in stagnant technology quality.

The following example illustrates a possible evolution of innovations in a random product line, with the top row of numbers representing the step size of each innovation:

Example 1.

An Example of a Sequence of Innovations in a Product Line

$\underbrace{\frac{\eta\bar{q}}{P_{1,f_1}} \frac{\eta\alpha\bar{q}}{P_{2,f_2}} \frac{\eta\alpha^2\bar{q}}{P_{3,f_3}} \frac{\lambda q_j}{P_{4,f_4}} \frac{\lambda q_j}{P_{5,f_5}} \frac{\eta\alpha^3\bar{q}}{P_{6,f_6}}}_{\text{Tech Cluster 1}}$	$\underbrace{\frac{\eta\bar{q}}{P_{7,f_7}} \frac{\lambda q_j}{P_{8,f_8}} \frac{\eta\alpha\bar{q}}{P_{9,f_9}}}_{\text{Tech Cluster 2}}$	$\underbrace{\frac{\eta\bar{q}}{P_{10,f_{10}}} \frac{\eta\alpha\bar{q}}{P_{11,f_{11}}} \frac{\eta\alpha^2\bar{q}}{P_{12,f_{12}}} \dots}_{\text{Tech Cluster 3}}$
--	--	--

Here $P_{m,f}$ denotes that the m th patent is obtained by firm f . The example starts with a major innovation that opens a new technology cluster by firm f_1 . Firms f_2 and f_3 then produce follow-up external innovations. Firm f_3 further improves its own product twice. Firm f_4 then produces a further follow-up external innovation. Next, this technology cluster is replaced by a new leading innovation by firm f_5 , which is patented as P_7 . The second cluster is then replaced by another leading innovation by firm f_7 . This new cluster is further improved by patents 11 and 12, and so on.

We later analytically solve for an expected step size \bar{s} from external innovations. For now, it is important to note that this theoretical structure does not depend on \bar{s} being greater or smaller than λ , and in fact this comparison may differ substantially depending on the country and time period studied. The baseline model framework is very general with respect to the relative sizes of internal versus external improvements.

C. Entry and Exit

As in Klette and Kortum (2004), a mass of entrants invest in R&D in order to become intermediate producers on a successful innovation. Entrants choose an innovation flow rate $x_e > 0$ with an R&D cost $C_e(x_e, \bar{q}) = x_e \nu \bar{q}$ in terms of the final good, where $\nu > 0$ is a constant scale parameter. The value V_0 of being an outside entrepreneur is the expected value from innovating successfully and entering the market. This value is determined according to

$$rV_0 - \dot{V}_0 = \max_x \{x_e [\mathbb{E}_j V(\{q_j + s_j \bar{q}\}) - V_0] - \nu x_e \bar{q}\}, \tag{9}$$

where $V(\{q_j\})$ denotes the value of a firm that owns a single product line with quality q_j and $\dot{V}_0 \equiv \partial V_0 / \partial t$ denotes the partial derivative of the outside value with respect to time. The expected value $\mathbb{E}_j V(\{q_j + s_j \bar{q}\})$ of a new innovation is an expectation over both quality level q_j and innovation size s_j . When there is positive entry, the equilibrium is such that

$$\mathbb{E}_j V(\{q_j + s_j \bar{q}\}) = \nu \bar{q}. \tag{10}$$

Incumbent firms produce intermediate inputs and invest in R&D. As a result, firms simultaneously expand into new product lines and lose some of their current product lines to other firms in the economy through competition. Each product line faces the same aggregate endogenous creative destruction rate τ . A firm that loses all product lines to competitors exits the economy.

D. Equilibrium

We now characterize the Markov perfect equilibria of the economy that make strategies a function of payoff-relevant states only. We focus on the steady state in which aggregate variables (Y, C, R, K, w, \bar{q}) grow at the constant rate g .

1. Production

The standard maximization problem of the representative household yields the Euler equation

$$\frac{\dot{Y}}{Y} = \frac{\dot{C}}{C} = r - \rho. \tag{11}$$

The maximization problem of the final goods producer generates the inverse demand $p_j = L^\beta q_j^\beta k_j^{1-\beta}$ for all $j \in [0, 1]$. The constant marginal cost of producing each intermediate variety is w/\bar{q} .

The profit maximization problem of the monopolist j is thus

$$\pi(q_j) = \max_{k_j \geq 0} \left\{ L^\beta q_j^\beta k_j^{1-\beta} - \frac{w}{\bar{q}} k_j \right\} \quad \forall j \in [0, 1]. \tag{12}$$

The first-order condition for (12) yields an optimal quantity and price for intermediate good j

$$k_j = \left[\frac{(1 - \beta)\bar{q}}{w} \right]^{1/\beta} L q_j \quad \text{and} \quad p_j = \frac{w}{(1 - \beta)\bar{q}}. \tag{13}$$

The realized price is a constant markup over the marginal cost and is independent of the individual product quality. Thus, the profit for each active good is $\pi(q_j) = \pi q_j$, where $\pi \equiv L(\bar{q}/w)^{(1-\beta)/\beta} (1 - \beta)^{(1-\beta)/\beta} \beta$. In order to avoid the case of limit pricing and maintain a simple model, we adopt the following stage-game assumption.

ASSUMPTION 1 (Monopoly pricing). In a given product line j , the current incumbent and any former incumbents in the same line (with lower quality than the current incumbent) enter a two-stage price-bidding game. In the first stage, each firm pays a fee of ϵ that is arbitrarily close to zero. In the second stage, all firms that paid the fee announce their prices.

Under assumption 1, only the leader pays the fee and enters the second stage since other firms can never recover their fee in the second stage. Since the leader is the only firm bidding a price, the leader will always operate with monopoly pricing, as in Aghion and Howitt (1992).

The maximization in the final goods sector, together with (13), implies a wage rate

$$w = \tilde{\beta}\bar{q}, \quad (14)$$

where $\tilde{\beta} \equiv \beta^\beta [1 - \beta]^{1-2\beta}$. Incorporating the equilibrium wage rate, the constant part of the equilibrium profit simplifies to

$$\pi = L(1 - \beta)\tilde{\beta}. \quad (15)$$

Note that, using the equilibrium quantity (13) and the wage rate (14), aggregate output can now be expressed as a linear function of production workers L and the average quality \bar{q} such that

$$Y = \frac{[1 - \beta]^{1-2\beta}}{\beta^{1-\beta}} \bar{q}L. \quad (16)$$

Equations (4), (5), (13), and (14) determine the final good workers as a fraction of the aggregate unit measure of workers,

$$L = \frac{\beta}{(1 - \beta)^2 + \beta}. \quad (17)$$

2. Invariant Step Size Distribution and Expected External Step Size

We next compute the invariant step size distribution $\Psi(s)$ that determines the expected innovation size from external innovations \bar{s} . Let Ψ_k denote the equilibrium share of product lines with $k \in \mathbb{N}_0$ subsequent follow-up innovations such that $s_j = \eta\alpha^k$. A steady-state equilibrium requires a stable innovation size distribution. Thus, while the stochastic nature of innovation moves individual products up and down the k distribution, the overall share of products at each level k is stable. This stability requires equal inflows and outflows of products from each size level, resulting in the flow equations:

State:	Inflow	Outflow	
$k = 0:$	$(1 - \Psi_0)\tau\theta$	$= \Psi_0\tau(1 - \theta)$	(18)
$k \geq 1:$	$\Psi_{k-1}\tau(1 - \theta) = \Psi_k\tau.$		

The first line governs inflows and outflows among product lines where major innovations have just occurred. Outflows happen because of

follow-up innovations at the rate $\tau(1 - \theta)$, while inflows happen because of new leading innovations being realized at rate $\theta\tau$ throughout the innovation size distribution. Internal R&D within firms does not influence these k distributions. A similar reasoning governs the share of product lines with $k \geq 1$ consecutive follow-up innovations. As a result, flow equations (18) generate the invariant distribution

$$\Psi_k = \theta(1 - \theta)^k \quad \text{for } k \geq 0, \tag{19}$$

which yields the expected innovation size from external R&D:

$$\bar{s} = \mathbb{E}(s_j) = \sum_{k=0}^{\infty} \Psi_k \eta \alpha^k = \frac{\theta \eta}{1 - (1 - \theta)\alpha}. \tag{20}$$

This expected size is naturally increasing in the probability of a major innovation θ , the realized size of major innovations η , and for lower decay rates in innovation quality within a technology cluster (i.e., higher α).

3. Research and Development by Incumbents

The value functions of firms determine R&D choices. For simplicity we drop the firm subscript f from the firm variables when it causes no confusion. Consider a firm with a product portfolio \mathbf{q} that serves as the state variable in the firm’s problem. The firm takes the values of (r, τ, g) as given and chooses the optimal R&D efforts x and z_j for every $j \in \mathcal{J}$ to maximize the following value function:²¹

$$rV(\mathbf{q}) - \dot{V}(\mathbf{q}) = \max_{\substack{x \in [0, \bar{x}], \\ \{z_j \in [0, \bar{z}]\}_{\mathcal{J}}}} \left\{ \begin{array}{l} \pi q_j - \hat{\chi} z_j^{\hat{\psi}} q_j \\ \sum_{q_j \in \mathbf{q}} \left[+z_j [V(\mathbf{q} \setminus \{q_j\} \cup_+ \{q_j(1 + \lambda)\}) - V(\mathbf{q})] \right. \\ \qquad \qquad \qquad \left. + \tau [V(\mathbf{q} \setminus \{q_j\}) - V(\mathbf{q})] \right] \\ + x [\mathbb{E}_j V(\mathbf{q} \cup_+ \{q_j + \bar{q}s_j\}) - V(\mathbf{q})] \\ \qquad \qquad \qquad - \tilde{\chi} x^{\tilde{\psi}} \bar{q} - \Phi \bar{q} \end{array} \right\}. \tag{21}$$

The first line on the right hand side represents operating profits over currently held product lines minus internal R&D costs. The second line is the change in firm value after internal improvements to currently held products. The term $V(\mathbf{q} \setminus \{q_j\} \cup_+ \{q_j(1 + \lambda)\})$ denotes the firm value after improving one of the firm’s existing products by size λ . These terms are multiplied by the Poisson arrival rate z_j as the success of internal R&D

²¹ We do not index the portfolio or R&D efforts by f as $\mathbf{q}_f, x_f,$ and z_{fj} to simplify notation. The term \cup_+ indicates the multiset union operator such that $\{a, b\} \cup_+ \{b\} = \{a, b, b\}$. Similarly, \setminus indicates the multiset difference operator such that $\{a, b, b\} \setminus \{b\} = \{a, b\}$.

is stochastic. Firms choose innovation effort for each product line separately. The third line shows the change in firm value due to losing its product lines through creative destruction τ . The term $V(\mathbf{q} \setminus \{q_j\})$ denotes firm value after losing a product that had quality q_j .

The fourth line is the change in firm value after a successful external innovation that garners a new product line. The term $V(\mathbf{q} \cup_+ \{q_j + \bar{q}s_j\})$ denotes equilibrium firm value after a successful external innovation of size s_j that adds a new product into the firm's portfolio. This addition is multiplied by the Poisson arrival rate x as the success of external R&D is stochastic too. The final line represents external R&D costs and fixed costs. The $-\dot{V}(\mathbf{q})$ term on the left-hand side of equation (21) represents change in firm value without any material events for the focal firm due to economywide growth (i.e., \bar{q} changes).

The aggregate creative destruction rate is the sum of average external innovation effort by each incumbent, Fx , and the realized entry rate x_e ,

$$\tau = Fx + x_e. \quad (22)$$

The aggregate growth rate is determined by the frequency of innovations coming from creative destruction τ , consisting of new entry and external innovations by incumbents; the frequency of internal innovations z ; and their relevant innovation sizes as described in the following lemma.

LEMMA 1. Let the equilibrium R&D efforts be given by (τ, z) . The steady-state growth rate of the aggregate variables in the economy is

$$g = \tau\bar{s} + z\lambda. \quad (23)$$

Now we are ready to solve for the equilibrium value function. One technical detail needs particular attention. Our goal in this benchmark model is to generate new intuitions while preserving tractability. The Klette and Kortum (2004) model is very tractable since everything scales perfectly in the number of product lines of the firms; this includes the profits collected by the firm and the franchise value, which is an option value for external innovation arising because more product lines make the firm more innovative via the Cobb-Douglas R&D technology. In our baseline model, profits also scale perfectly, yet the franchise value is constant across all firms since the R&D technology depends on having positive product lines only at the extensive margin but not on the intensive margin. This introduces a nonlinearity to the firm value function. To generate a value function that scales perfectly with the number of product lines as in the Klette and Kortum (2004) model, we assume that the fixed cost of operation is equal to the franchise value as follows.²²

²² The equality simplifies the math for the rest of the baseline model. These technical conditions related to fixed costs are not important for our general framework, and thus fixed costs are set equal to zero in later sections.

ASSUMPTION 2 (Perfectly scaling value function). The value of fixed cost of operation satisfies

$$\Phi = \left[\frac{\nu}{\hat{\psi}\hat{\chi}} \right]^{\frac{\psi}{\psi-1}} \tilde{\chi}(\tilde{\psi} - 1).$$

The next proposition shows that the value function (21) and its components can be expressed in a very tractable form. We assume for now that there is positive entry and later impose a parameter restriction that is sufficient to verify this condition.

PROPOSITION 1. Under assumptions 1 and 2 and when there is positive entry $x_e > 0$, the value function (21) of a firm with a set of product lines \mathbf{q} can be expressed as $V(\mathbf{q}) = A \sum_{q_j \in \mathbf{q}} q_j$, where A (the value of holding a product line) is

$$A = \frac{\nu}{1 + \bar{s}}. \tag{24}$$

Moreover, the optimal R&D decisions are given by

$$z = \left[\frac{\lambda\nu}{(1 + \bar{s})\hat{\psi}\hat{\chi}} \right]^{\frac{1}{\psi-1}} \quad \text{and} \quad x = \left[\frac{\nu}{\hat{\psi}\hat{\chi}} \right]^{\frac{1}{\psi-1}}, \tag{25}$$

and the aggregate creative destruction rate is

$$\tau = \frac{1}{1 + \bar{s}} \left[\frac{\pi}{A} - \left[\frac{\lambda}{\hat{\psi}\hat{\chi}} \right]^{\frac{\psi}{\psi-1}} A^{\frac{1}{\psi-1}} \hat{\chi} - \rho \right]. \tag{26}$$

This proposition shows that the innovation efforts of incumbents, both internal and external, are positively related to the entry cost. Higher entry costs lower entry rates and thus provide longer expected durations and profits from owning product lines. Moreover, both internal and external R&D efforts decline in their own cost scale parameters.

Importantly, internal innovation is increasing in its own step size λ because of the higher marginal return to successful internal improvements, but internal investments are decreasing in the average step size of external innovation \bar{s} , since larger \bar{s} encourages more creative destruction that lowers the expected duration of monopoly power the firm has on the product line. By contrast, step sizes do not show up in the equilibrium external innovation rate since a bigger step size \bar{s} both encourages effort (due to higher return) and discourages it (due to higher entry); these two opposing effects cancel out.

To pin down the entry rate, we solve for the equilibrium measure of firms F . To achieve this, we first characterize the invariant distribution of the number of products. This distribution is the main proxy for the

firm size distribution in Klette and Kortum (2004). Let μ_n denote the equilibrium share of the incumbent firms that own n product lines such that $\sum_{n=1}^{\infty} \mu_n = 1$. The invariant distribution again depends on the following flow equations:

State:	Inflow	Outflow	
$n = 0$:		$F\mu_1\tau = x_e$	(27)
$n = 1$:		$F\mu_2 2\tau + x_e = F\mu_1(x + \tau)$	
$n \geq 2$:		$F\mu_{n+1}(n + 1)\tau + F\mu_{n-1}x = F\mu_n(x + n\tau)$.	

The first line characterizes outside entrepreneurs ($n = 0$). Inflows to outside entrepreneurs happen when firms with one product are destroyed, and outflows occur when outside entrepreneurs successfully develop a new product at rate x_e . Similarly, the second line considers inflows and outflows of firms with one product, and the third line considers n -product firms. The next proposition provides the explicit form solution of the invariant product number distribution.

PROPOSITION 2. The invariant distribution μ_n is equal to

$$\mu_n = \frac{x_e}{Fx} \left(\frac{x}{\tau}\right)^n \frac{1}{n!} \quad \text{for } n \geq 1. \tag{28}$$

Since (28) is a probability distribution, it must be that $\sum_{n=1}^{\infty} \mu_n = 1$, which implies $Fx/x_e = e^{x/\tau} - 1$. This condition and (22) deliver the entry rate as

$$x_e = \tau e^{-x/\tau} \quad \text{and} \quad F = \frac{\tau}{x} (1 - e^{-x/\tau}). \tag{29}$$

The entry rate is a fraction of the aggregate creative destruction rate. In order to ensure an equilibrium with positive aggregate creative destruction and entry, we make the following assumption.

ASSUMPTION 3 (Positive entry). The parameters of the model are such that

$$\pi > \left[\frac{\lambda}{\hat{\psi}\hat{\chi}} \right]^{\frac{\psi}{\psi-1}} \left[\frac{\nu}{1+\bar{s}} \right]^{\frac{\psi}{\psi-1}} \hat{\chi} + \frac{\nu\rho}{1+\bar{s}}.$$

This assumption is very easy to satisfy. For any given positive profit, there is always a low enough entry cost ν such that an equilibrium with positive entry exists.

The total R&D effort of the economy is

$$R = \hat{\chi} \left[\frac{\lambda\nu}{(1+\bar{s})\hat{\psi}\hat{\chi}} \right]^{\frac{\psi}{\psi-1}} \bar{q} + F\tilde{\chi} \left[\frac{\nu}{\tilde{\psi}\tilde{\chi}} \right]^{\frac{\psi}{\psi-1}} \bar{q} + \nu\tau e^{-\frac{x}{\tau}} \bar{q}, \tag{30}$$

and the total fixed cost is

$$K = F\Phi\bar{q}. \tag{31}$$

Combining (16) and (17) delivers the equilibrium output level,

$$Y = \frac{[1 - \beta]^{1-2\beta} \beta^\beta}{(1 - \beta)^2 + \beta} \bar{q}. \tag{32}$$

From this, consumption is determined through the resource constraint as

$$C = Y - K - R. \tag{33}$$

We end this section by summarizing the equilibrium.

DEFINITION 1 (Balanced growth path equilibrium). A balanced growth path equilibrium of this economy consists of the following tuple for every $t, j \in [0, 1], \bar{q}$, and q_j : $k_j^*, p_j^*, w^*, L^*, \tilde{L}^*, x^*, z_j^*, \tau^*, x_e^*, F^*, R^*, K^*, Y^*, C^*, g^*, \Psi_n^*, \mu_n^*, r^*$, such that (i) k_j^* and p_j^* satisfy (13); (ii) wage rate w^* satisfies (14); (iii) measure of final good production workers L^* satisfies (17) and \tilde{L}^* is simply $1 - L^*$; (iv) external (x^*) and internal (z_j^*) innovation flows are equal to (25); (v) aggregate creative destruction τ^* satisfies (26); (vi) entry flow x_e^* and measure of incumbent firms F^* satisfy (29); (vii) total R&D spending R^* satisfies (30); (viii) total amount of fixed cost expenses K^* satisfies (31); (ix) aggregate output Y^* satisfies (32); (x) aggregate consumption C^* satisfies (33); (xi) steady-state growth rate g^* satisfies (23); (xii) the invariant distribution of innovation sizes Ψ_n^* satisfies (19); (xiii) the invariant distribution of number of products μ_n^* satisfies (28); and (xiv) the interest rate satisfies the Euler equation (11).

E. Central Theoretical Results

The following propositions characterize the firm growth, R&D, and innovation dynamics of the model. These closely correspond to the empirical regularities described in the prior section. In our model, the ideal proxy for firm size is the total quality $Q = \sum_{q_j \in \mathbf{q}_j} q_j$ because firm sales, profits, and production workers are all proportional to Q .²³ Firm size also closely relates to the number of product lines, which we discuss in Sec-

²³
$$\text{Sales} = \sum_{q_j \in \mathbf{q}_j} p(q_j)k(q_j) = [(1 - \beta)/w]^{1-\beta} LQ_j,$$

$$\text{Profits} = \sum_{q_j \in \mathbf{q}_j} \pi q_j = \pi Q_j,$$

and

$$\text{Production workers} = \sum_{q_j \in \mathbf{q}_j} l_j = [(1 - \beta)/w]^{1/\beta} LQ_j.$$

tion VI.B. Therefore, we also use n_j to proxy for firm size in propositions when convenient.

PROPOSITION 3. Let $\mathcal{G}(Q) \equiv \mathbb{E}(\dot{Q}/Q)$ be the average growth rate of a firm with total quality Q . Then $\mathcal{G}(Q)$, in equilibrium, is given by

$$\mathcal{G}(Q) = \frac{x(1 + \bar{s})\bar{q}}{Q} + z\lambda - \tau,$$

where $\mathcal{G}(Q)$ is a strictly decreasing function.

This result suggests that small firms grow faster than large firms. This microfounded departure from Gibrat's law of proportionate growth occurs because of the lack of scaling of external innovation efforts. As a result, the growth coming from internal innovation is the same, on average, across different firm sizes ($z\lambda$), whereas the contribution of external R&D to firm growth gets smaller as firm size increases (the first ratio in $\mathcal{G}(Q)$). Combining these effects, overall firm growth declines with firm size.

PROPOSITION 4. Let $\mathcal{R}(Q) \equiv \text{R\&D/Sales}$ be the firm R&D intensity of a firm with total quality Q . Then $\mathcal{R}(Q)$, in equilibrium, is given by

$$\mathcal{R}(Q) = \frac{\beta c_x(x)\bar{q}}{\pi Q} + \frac{\beta c_z(z)}{\pi},$$

where $\mathcal{R}(Q)$ is a strictly decreasing function.

This result suggests that small innovative firms have a greater R&D intensity than large firms. Similarly to the previous proposition, the intuition is that total internal R&D effort is proportionate to the number of product lines of the firm. On the other hand, external R&D efforts do not scale with the number of product lines, which results in a declining R&D intensity for larger firms. In other words, adding additional product lines continually adds more R&D effort but further dilutes the external R&D effects with respect to intensity measures.

As noted earlier, our model does not require taking a stance on the relative sizes of internal versus external innovations. With some structure added that is consistent with our earlier empirical results, the model also makes predictions about the innovation size distribution and the relative frequency of firms by innovation size.

PROPOSITION 5. Let a major innovation be defined as an innovation with a step size larger than a certain threshold $s_k \geq s_{\hat{k}}$ for some $\hat{k} \in \mathbb{Z}_+$ and $s_{\hat{k}} > \lambda$. Moreover, let $\mathcal{M}(n)$ be the probability of making a major innovation conditional on having a successful innovation for a firm with n product lines. Then $\mathcal{M}(n)$ can be expressed by

$$\mathcal{M}(n) \equiv \frac{x \sum_{k=0}^{\hat{k}} \theta(1 - \theta)^k}{x + nz} = \frac{x[1 - (1 - \theta)^{\hat{k}+1}]}{x + nz},$$

where $\mathcal{M}(n)$ is a strictly decreasing function.

This result suggests that small firms and new entrants have a comparative advantage for achieving major advances. Large incumbents endogenously spend effort on maintaining and expanding existing products. Thus, while firms of all sizes obtain major advances, these major advances account for a smaller share of achieved innovations among larger firms.²⁴ An important distributional implication of proposition 5 is that these differences weaken when considering progressively larger thresholds s_i . The comparative advantage is weakest at the most extreme values (i.e., $s_{i=0} = \eta$).

We empirically estimated these predictions in Section II.B, and we use these results in our quantitative analysis. The baseline model makes many more predictions that we catalogue in appendix B and investigate further in our NBER working paper (Akcigit and Kerr 2010).

IV. Generalized Model

This section generalizes the innovation production function of the benchmark model. In particular, we assume that the production function for external innovations takes the form

$$X_n = \chi [R_x / \bar{q}]^\psi n^\sigma. \quad (34)$$

This production function nests two special forms. First, when $\sigma = 1 - \psi$, the model becomes the extended Klette and Kortum (2004) framework in which both internal and external investments scale up with firm size on a one-for-one basis with added product lines. Second, when $\sigma = 0$, we are back to the benchmark model of Section III. We describe here the solution of the model under this generalized production function, and Section VI quantifies this model and the σ parameter.

The static equilibrium of this generalized model follows exactly as the benchmark model; therefore, we skip it (eqq. [13]–[17] hold identically). Moreover, when $\sigma > 0$, a firm that loses all of its product lines exits the economy. As we are not seeking analytical results, but instead preparing the general model for quantification, we eliminate the fixed cost and set $\Phi = 0$.

Research and development by incumbents.—The production function in (34) delivers the R&D function

$$R_x = \bar{q} \tilde{\chi} n^{\tilde{\sigma}} x_n^{\tilde{\psi}},$$

where $x_n \equiv X_n/n$ is the innovation intensity per product line and

$$\tilde{\sigma} \equiv \frac{1 - \sigma}{\psi}, \quad \tilde{\chi} \equiv \chi^{-1/\psi}, \quad \text{and} \quad \tilde{\psi} \equiv \frac{1}{\psi}.$$

²⁴ The aggregate quantity of major innovations by small and large firms depends on these propensities and the firm size distribution.

In this case, the value function can be expressed as follows.

PROPOSITION 6. For a firm that has a quality portfolio \mathbf{q} , the value function has the following form:

$$V(\mathbf{q}, \bar{q}) = A \sum_{q_j \in \mathbf{q}} q_j + B_n \bar{q},$$

where

$$(r + \tau)A = \pi + A^{\frac{\psi}{\psi-1}} \left[\frac{\lambda}{\hat{\psi}} \right]^{\frac{\psi}{\psi-1}} (\hat{\psi} - 1) \hat{\chi}^{\frac{1}{1-\psi}}, \tag{35}$$

and

$$B_{n+1} = \left[\frac{(\rho + n\tau)B_n - n\tau B_{n-1}}{\hat{\psi} - 1} \right]^{\frac{\psi-1}{\psi}} \tilde{\psi} \tilde{\chi}^{\frac{1}{\psi}} n^{\frac{\sigma-\psi}{\psi}} + B_n - A[1 + \bar{s}]. \tag{36}$$

Moreover, the optimal innovation efforts are defined as

$$z_j = \left[\frac{A\lambda}{\hat{\psi}\hat{\chi}} \right]^{\frac{1}{\psi-1}} \quad \text{and} \quad x_n = \left[\frac{A[1 + \bar{s}] + B_{n+1} - B_n}{\tilde{\psi} n^{\sigma-1} \tilde{\chi}} \right]^{\frac{1}{\psi-1}}. \tag{37}$$

In this generalized model, the value function consists of two parts. The first part, which is denoted by A , is related to the discounted sum of future profits and internal innovations. By owning the product line, the firm will collect flow profits of πq_j until it is replaced at the rate τ . In addition, the firm can improve its quality q_j through internal innovations at the rate z_j , which also provides value to the firm. The second part, which is denoted by B_n , relates to the firm’s external innovation capacity. By owning a product line, the firm has a franchise value of extending into new product lines through external innovations, which happens at the rate x_n . Since the production function is dependent on the number of product lines, this franchise value now is a function of n as well. The Klette and Kortum (2004) model corresponds to $B_n = nB$, while the baseline model of Section III corresponds to $B_n = B$.

Accordingly, the new flow equations for the fraction of firms with n product lines are

State:	Inflow	Outflow
$n = 0 :$		$F\mu_1\tau = x_e$
$n = 1 :$		$F\mu_2 2\tau + x_e = F\mu_1(2x_2 + \tau)$
$n \geq 2 :$	$F\mu_{n+1}(n + 1)\tau + F\mu_{n-1}(n - 1)x_{n-1}$	$= F\mu_n(nx_n + n\tau).$

This summarizes the generalized model, and a final remark is in order.

REMARK 1. Proposition 6 shows that innovation intensity x_n can be expressed as $x_n = n^\xi f(n)$, where

$$\xi \equiv \frac{\psi + \sigma - 1}{1 - \psi} \tag{38}$$

and n^ξ captures the direct effect of n on x_n . Note that

$$f(n) = \left[\frac{A[1 + \bar{s}] + B_{n+1} - B_n}{\tilde{\psi}\tilde{\chi}} \right]^{\frac{1}{\bar{s}-1}}$$

captures the indirect effect of number of product lines on x_n through its impact on the franchise value B_n . When $\psi + \sigma = 1$, our model mirrors Klette and Kortum (2004) with $f(n)$ equal to some constant, whereas innovation intensity will be decreasing in firm size when $\psi + \sigma < 1$. Therefore, $\psi + \sigma$ dictates the amount of decreasing innovation intensity in firm size.

V. Patent Citation Behavior and Innovation Spillover Sizes

We now incorporate patent citation behavior across innovations into our benchmark model. As we have already defined the economy’s equilibrium, our specified citation behavior does not affect real outcomes. We undertake this extension, however, to derive the economic meaning behind patent citations. This in turn allows us to quantify the model using richer data. Second, this addition demonstrates how this class of endogenous growth models captures many important features uncovered in the empirical literature on patent counts and citations.²⁵

Forward patent citations.—Innovations are clustered in terms of their technological relevances. Major innovations generate new technology clusters that last until they are overtaken by a subsequent major innovation. An example of the sequential innovation process was illustrated in example 1 in Section III.B.

Let $m(j, t)$ be the number of patents in the active technology cluster in product line j . For instance, if t is between the innovation times of P_3 and P_4 in example 1, then $m(j, t) = 3$; or if t is between P_2 and P_3 , then $m(j, t) = 2$. Therefore, the number of citable patents in active technology clusters at time t is $M(t) = \int_0^1 m(j, t) dj$.

We next describe the citation distribution of patents by specifying citation behavior with rules that are consistent with the patent literature. Pat-

²⁵ Hall et al. (2001) provide a comprehensive introduction to patent citations. See also Trajtenberg (1990), Jaffe, Trajtenberg, and Henderson (1993), Jaffe, Trajtenberg, and Foray (2000), Hall, Jaffe, and Trajtenberg (2005), and Thompson and Fox-Kean (2005).

ents cite previous patents within the same technology cluster to specify how they build on the prior work and the boundaries of the innovations. Each new patent, by definition, improves the previous technologically relevant innovations on some dimensions. However, not all subsequent innovations improve an existing technology in the same direction. Therefore, major patents with broader scope are more likely to be cited by subsequent follow-on patents (e.g., Lerner 1994). We proxy this patent scope by the step size $s \in \{\lambda, \eta\alpha^k | k \in \mathbb{N}_0\}$ in our model. We assume that an innovation with size s will receive a citation from a subsequent patent within the same technology cluster with probability $s\gamma$, where $\gamma \in (0, 1/\eta)$. Finally, a major innovation replaces the previous cluster. Thereafter, future citations begin with the new major innovation. Empirically, Hall et al. (2001) and Mehta, Rysman, and Simcoe (2010) quantify the decline in relative citation rates over patent age that this model structure provides. The citation behavior of example 1 is illustrated in table 2.

With these simple modeling assumptions, we can characterize the flow properties of citation behavior. These traits depend on the real side of the economy and provide a richer description of it. Our upcoming quantitative analysis uses the citation distribution to inform the traits of internal and external innovation. Similarly to our earlier expressions, the economy's equilibrium requires an invariant citation distribution. Let $\Upsilon_{s_i,n}$ and $\Upsilon_{\lambda,n}$ denote the share of patents that are of size $\eta\alpha^k$ and λ , respectively, and receive n citations. These shares include all patents and naturally sum to one when aggregating over all levels of citation counts, $\sum_{n=0}^{\infty} \Upsilon_{\lambda,n} + \sum_{k=0}^{\infty} \sum_{n=0}^{\infty} \Upsilon_{s_i,n} = 1$. The next proposition provides the explicit form solutions for these distributions.

PROPOSITION 7. The invariant distribution of the total number of forward citations (n) given to a patent of size $s \in \{\lambda, s_k | k \in \mathbb{N}_0\}$ can be expressed as

$$\Upsilon_{s,n} = \Upsilon_{s,0} \Omega_s^n \quad \text{for } n \in \mathbb{N}_0,$$

TABLE 2
CITATION PATTERNS IN EXAMPLE 1

Cited	Probability	Citing
P_1	$\gamma\eta$	P_2-P_6
P_2	$\gamma\eta\alpha$	P_3-P_6
P_3	$\gamma\eta\alpha^2$	P_4-P_6
P_4	$\gamma\lambda$	P_5, P_6
P_5	$\gamma\lambda$	P_6
P_6	$\gamma\eta\alpha^3$	None
P_7	$\gamma\eta$	P_8, P_9
P_8	$\gamma\lambda$	P_9
P_9	$\gamma\eta\alpha$	None
P_{10}	$\gamma\eta$	P_{11}, P_{12}, \dots

where $M = (x + z)/x\theta$,

$$\Upsilon_{s,0} = \frac{\theta(1 - \theta)^k \tau}{M[\tau\theta + \gamma s_k(\tau(1 - \theta) + z)]},$$

$$\Upsilon_{\lambda,0} = \frac{z}{M[\tau\theta + \gamma\lambda(\tau(1 - \theta) + z)]},$$

and

$$\Omega_s \equiv \frac{\gamma s(\tau(1 - \theta) + z)}{\tau\theta + \gamma s(\tau(1 - \theta) + z)}.$$

Similarly, the invariant distribution of the total number of external forward citations is

$$\tilde{\Upsilon}_{s,n} = \tilde{\Upsilon}_{s,0} \tilde{\Omega}_s^n \quad \text{for } n \in N_0,$$

where

$$\tilde{\Upsilon}_{s,0} = \frac{\theta(1 - \theta)^k \tau}{M[\tau\theta + \gamma s_k \tau(1 - \theta)]},$$

$$\tilde{\Upsilon}_{\lambda,0} = \frac{z}{M[\tau\theta + \gamma\lambda\tau(1 - \theta)]},$$

and

$$\tilde{\Omega}_s \equiv \frac{\gamma s \tau(1 - \theta)}{\tau\theta + \gamma s \tau(1 - \theta)}.$$

This proposition shows the information available from citation distributions. As $\tau\theta$ gets smaller in the denominator, $\Upsilon_{s,n}$ generates a more highly skewed distribution of citations. This is intuitive as a slower arrival of new technology clusters tilts innovation toward follow-on inventions that cite prior inventions and thereby generate more extreme citation counts. Patent citation distributions can thus be used to discipline the traits of innovation in the economy that would otherwise be unobservable.

VI. Quantitative Analysis

We estimate our model using microdata described in Section II.B. Section VI.A describes our identification strategy. Section VI.B provides the main estimation results, and Section VI.C provides robustness checks.

Appendix D outlines the computational solution of the generalized model.

A. Identification

Our model has 13 structural parameters as listed in table 3. We identify these parameters in three ways. First, we fix three parameters (ρ , $\hat{\psi}$, $\tilde{\psi}$) using values developed in Section VI.A.1 from the literature and R&D-based regressions. Second, we use the observed distribution of patent citations to pin down three elements of the step size distribution (θ , α , $\eta\gamma$) in Section VI.A.2. Finally, for the remaining parameters and to parse $\eta\gamma$, we target the relevant firm moments in the data. One critical part of this third step is to identify the key decreasing returns parameter σ using an indirect inference approach, where we replicate the regressions of Sections II.C and II.D using data simulated from the model.

1. Externally Calibrated Parameters

We set the discount rate equal to $\rho = 2$ percent, which roughly corresponds to an annual discount factor of 97 percent.

We rely on prior literature for estimates of the curvature of the R&D cost function, which we will set equal across internal and external innovation $\hat{\psi} = \tilde{\psi}$ (the model retains shifters in these cost functions). One line of studies quantifies the elasticity of patents to R&D expenditures (e.g., Griliches 1990; Hall and Ziedonis 2001; Blundell, Griffith, and

TABLE 3
PARAMETERS OF THE MODEL

Parameter	Description	Equation	Identification
ρ	Discount rate	(2)	External calibration
$\hat{\psi}$	Curvature of internal R&D	(6)	External calibration
$\tilde{\psi}$	Curvature of external R&D	(7)	External calibration
θ	Probability of major advance	(8)	Match citation distribution
α	Multiplier of declining follow-up improvements	(8)	Match citation distribution
λ	Quality multiplier of internal innovation	(8)	Indirect inference
β	Quality share in final goods production	(3)	Indirect inference
ν	Entry cost	(9)	Indirect inference
$\hat{\chi}$	Scale of internal R&D	(6)	Indirect inference
$\tilde{\chi}$	Scale of external R&D	(7)	Indirect inference
σ	Product line share in external R&D	(34)	Indirect inference
η	Quality multiplier of major advance	(8)	Cite distribution + indirect inference
γ	Citation probability multiplier	(43)	Cite distribution + indirect inference

Windmeijer 2002). This literature often concludes that this elasticity is around 0.5, which implies a quadratic curvature. Acemoglu et al. (forthcoming) reach a similar estimate using the Census Bureau data as well when focusing on firms in the R&D Survey. The second set of papers examines the impact of R&D tax credits on the R&D expenditure of firms (e.g., Hall 1993; Bloom, Griffith, and Van Reenen 2002; Wilson 2009). In a survey of this literature, Hall and Van Reenen (2000) conclude that a tax price elasticity of around unity is typically found, which again corresponds to a quadratic cost function.²⁶ Given this common finding, we set $\hat{\psi} = \tilde{\psi} = 2$. Section VI.C.4 will study the robustness of the results with alternative elasticities of 0.4 and 0.6.

2. Citation Distribution

Our model yields an analytical solution for the patent citation distribution that is dictated by the innovation step size parameters. In particular, when we focus only on external citations ($z_j = 0$), the distribution of patents that are of quality s_k and receive n citations is simply

$$\Upsilon_{s_k,n} = \Upsilon_{s_k,0} \Omega_{s_k}^n \quad \text{for } n \in N_0,$$

where

$$\Upsilon_{s_k,0} = \frac{\theta^2(1 - \theta)^k}{\theta + \gamma\eta\alpha^k(1 - \theta)}$$

and

$$\Omega_{s_k} \equiv \frac{\gamma s_k(1 - \theta)}{\theta + \gamma s_k(1 - \theta)}.$$

The term $\Upsilon_{s_k,n}$ gives us the joint distribution of patents that are k -times incremented and have received n citations. Our model provides the analytical distribution of k -times incremented patents from (19) as $\Psi_k = \theta(1 - \theta)^k$ for $k \geq 0$. Hence, we can find the marginal distribution of n -times cited patents as

$$\mathcal{F}_n(\theta, \gamma, \eta, \alpha) = \sum_{k=0}^{\infty} \Psi_k \Upsilon_{s_k,n}.$$

The empirical tractability comes from the fact that the distribution of n -times cited patents depends only on four structural parameters: $\theta, \gamma, \eta,$

²⁶ The mapping to our setting is straightforward. To simplify the notation, let us denote a single R&D spending relationship $R = P x_n^\psi F_n$, where P is the price of R&D and F_n is a multiplicative term that can potentially depend on firm size. If the return to innovation is Π , the generic maximization problem can be written as $\max_{x_n} \{x_n \Pi - P x_n^\psi F_n\}$. Solving for the first-order condition,

$$R = P^{-\frac{1}{\psi-1}} F_n^{-\frac{1}{\psi-1}} [\Pi/\psi]^{\frac{\psi}{\psi-1}}.$$

Hence the price elasticity of R&D spending in our model corresponds to $d \ln R / d \ln P = -1/(\psi - 1)$. A unitary estimate corresponds to $\psi = 2$.

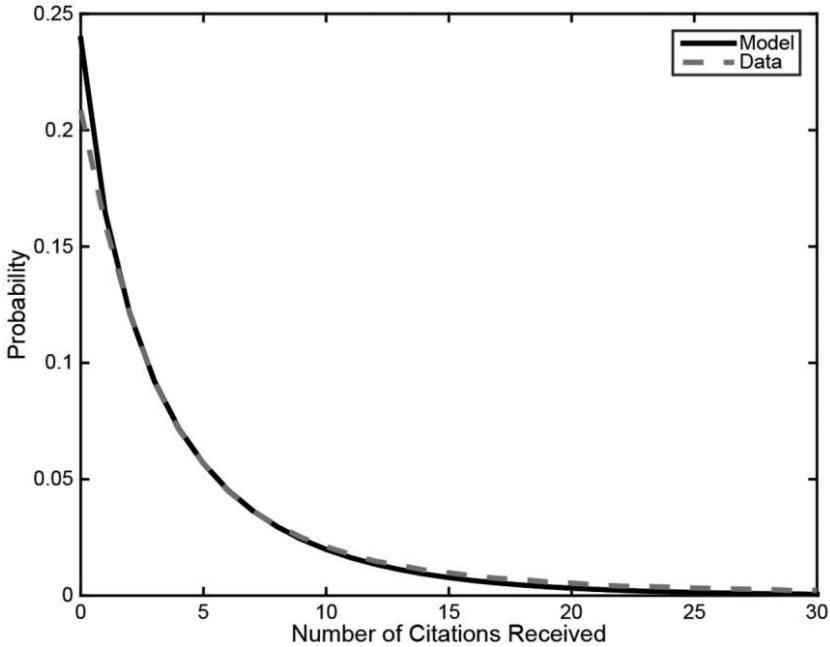


FIG. 7.—Citation distribution. Color version available as an online enhancement.

and α . Citation distributions do not allow one to distinguish between the overall quality level of external inventions (η) and factors that govern the general tendency of patents to cite each other (γ). Since γ and η always appear multiplicatively in the shape of the citation distribution, we can use these data to identify the three parameters θ , α , and the combined η , γ .

Figure 7 plots the empirical distribution together with the model-generated citation distribution. The model does a very good job of replicating the data.

Table 4 lists the resulting parameter estimates. Roughly 10 percent of external innovations are found to be significant enough to open new technology clusters, and the decay rate α for the quality of external inventions is fairly modest. The $\gamma\eta$ estimate suggests that patents that open a new technology cluster have a 75 percent probability of being cited by later patents in the cluster.

TABLE 4
CITATION DISTRIBUTION PARAMETERS

θ	$\gamma\eta$	α
.103	.750	.929

3. Indirect Inference

There are seven remaining parameters to be estimated: σ , $\tilde{\chi}$, $\hat{\chi}$, η , λ , β , and ν , which will also identify γ on the basis of the estimate in table 4. We identify these parameters using an indirect inference approach in the spirit of Lentz and Mortensen (2008). We compute various model-implied moments from the simulation strategy described above and compare them to the data-generated moments to minimize

$$\min \sum_{i=1}^7 \frac{|\text{model}(i) - \text{data}(i)|}{\frac{1}{2}|\text{model}(i)| + \frac{1}{2}|\text{data}(i)|},$$

where we index each moment by i . Our indirect inference procedure targets seven moments that we describe next. The generalized model does not yield an analytical solution, and thus we cannot express the targeted moments in this form. However, we build intuition by using the analytical solutions to Section II's benchmark model to guide us in choosing the right moments for identification. For ease of these depictions, we abstract from quality levels by setting $q_j = 1$ for all j , although innovation qualities are clearly included in the simulation of the general model.

Average profitability.—For both the benchmark and generalized models, the profit-to-sales ratio is equal to $\mathbb{E}(\text{profit}_f/\text{sales}_f) = (1 - \beta)^{(2\beta-1)/\beta} \tilde{\beta}^{1/\beta}$, where $\tilde{\beta} \equiv \beta^\beta [1 - \beta]^{1-2\beta}$. We therefore target the average profitability in the economy to help identify β . The profit-to-sales ratio in the model includes R&D expenditures, and thus we combine annual published Bureau of Economic Analysis (BEA) pretax profit rates with industrial R&D expenditure rates to determine an estimate of 10.9 percent for the 1982–97 period.

R&D intensity and internal-to-external citations ratio.—We discipline the R&D scale parameters $\hat{\chi}$ and $\tilde{\chi}$ through measures of R&D intensity and the citation ratio of internal versus external innovations. Aggregating across firms and using proposition 4, the baseline model shows the economywide R&D-sales ratio to be a linear combination of $\hat{\chi}$ and $\tilde{\chi}$. This ratio is 4.1 percent in our sample.²⁷ In addition, the citation ratio of internal versus external innovations informs the R&D scale parameters as

²⁷ For this purpose, we need to make use of the R&D Survey, which samples with certainty firms that conduct more than \$1 million of R&D and subsamples firms beneath this threshold. Our first step builds a sample of firm-period observations for which we observe reported R&D, sales, and employment. The 5-year periods match those of our core sample. We then merge in patents, including zero-valued outcomes. From this, we obtain an average conversion factor for relating R&D/sales to patents/employee. The second step applies this conversion factor to our full sample, where our aggregate patent/employee statistic includes firms that did not patent. This procedure gives us an aggregated value that closely aligns with other estimates of R&D/sales ratios. These values are determined through aggregates over the whole sample, not firm-level imputations. As the largest companies account for the substantial majority of these variables and will be surveyed directly by the R&D Survey, the procedures used here are quite robust.

$$\frac{\lambda \bar{\chi}}{(1 + \bar{s}) \hat{\chi}}$$

We define internal patents as those with 50 percent or more of citations being given to assignees of the same firm. This approach is similar to that in figure 3, with the explicit 10-year window from application date ensuring that the procedure is consistent across the sample period. We estimate this ratio using external citations to be 0.774 (= 5.023/6.488). These data inputs will inform the R&D scale parameters.

Fraction of internal patents and aggregate growth rate.—Our model has four parameters that govern the step size dynamics: θ , α , η , and λ . We previously identified θ and α through the citation distribution. The remaining two parameters are the step size for internal innovations λ and the step size of radical innovations η . Step sizes determine both the innovation incentives and the aggregate growth rate:

$$z_j = \left[\frac{\lambda \nu}{(1 + \bar{s}) \hat{\psi} \hat{\chi}} \right]^{\frac{1}{\psi-1}} \quad \text{and} \quad g = \tau^* \bar{s} + z^* \lambda.$$

We can therefore discipline η and λ by targeting the fraction of internal patents, $z/(z + \tau)$, and the growth rate. The internal patent share is 21.5 percent. The aggregate growth rate is calculated in deflated terms and on a per-employee basis to match the model and the BEA profit estimates. This ranges from 0.91 percent to 1.03 percent depending on details of the calculation, and we assign a value of 1.0 percent.

Entry rate.—The entry rate in the benchmark model is $x_e = \tau \exp(-x/\tau)$. Equations (24) and (26) show that the creative destruction rate is decreasing in the entry cost parameter ν , $d\tau/d\nu < 0$, and equation (25) shows that incumbent efforts are increasing in entrant costs, $dx/d\nu > 0$. Therefore, the impact of entry cost on the flow of entry is strictly negative, $dx_e/d\nu < 0$, and thus targeting the entry rate can help inform the entry parameter. The entry rate in our data is 5.82 percent, measured over 5-year intervals through employments among patenting entrants.

Firm growth versus firm size regression from Section II.C.—The extended Klette and Kortum (2004) approach, where $\sigma = 1 - \psi$, predicts that the unconditional firm growth would be independent of firm size, whereas the benchmark model with $\sigma = 0$ goes to the other extreme and predicts that firm growth is decreasing in firm size. In order to identify the actual value of σ , we mirror the same growth-size regressions with data generated from the simulated model. Every firm in the model has an innovation, as they would otherwise not exist, and we treat sample preparation and estimation exactly as we do in the data sample. The empirical coefficient of interest from the earlier analysis is -0.035 .

TABLE 5
MOMENTS

Moment	Data	Model
Profitability	.109	.106
R&D intensity	.041	.042
Internal/external cite	.774	.732
Fraction of internal patents	.215	.250
Entry rate	.058	.066
Average growth rate	.010	.010
Growth vs. size (fact 1)	-.035	-.035

B. Benchmark Estimation Results

Table 5 reports the empirical and simulated moments using the generalized model. Overall, the model matches closely the targeted moments. The resulting parameter estimates are reported in table 6.

Our estimates find that there are some decreasing returns in firm size for external innovation as captured by the value of $\sigma \approx 0.4$. Among the other results, the ratio of $\tilde{\chi}$ to $\hat{\chi}$ suggests that the R&D cost parameter for external innovations is about 12 times larger than for internal innovations. External innovations that open up a new technology cluster are estimated to have more than twice the potency of internal innovations. With the decay rate of $\alpha = 0.929$, roughly 10 follow-on external innovations occur before external innovations are less valuable than internal innovations.

1. Characterization of the Economy

To provide further intuition on how σ plays a role in generating size-dependent firm moments, figure 8 plots the franchise value function of a firm B_n as a function of the number of product lines n when $\sigma \in \{0, 0.2, 0.4, 0.5\}$. Figure 9 similarly plots the resulting external innovation intensity X_n . The franchise value function B_n for the baseline model in figure 8 is flat because external innovation does not scale, while it grows linearly in the Klette and Kortum (2004) scenario. The small dashed line shows that the franchise value with $\sigma = 0.4$ grows similarly to that in the Klette and Kortum framework among smaller firms, with more modest departures after that. Figure 9 likewise illustrates that external innovation intensity de-

TABLE 6
ESTIMATED MODEL PARAMETERS

σ	$\tilde{\chi}$	$\hat{\chi}$	η	λ	β	ν
.395	4.066	.346	.112	.051	.106	.830

NOTE.—Implied $\sigma + \psi = .895$.




FIG. 8.—Franchise value B_n . Color version available as an online enhancement.

clines with firm size but stabilizes in a way that limits the full dilution in the baseline model.

In our model, firm size is determined by the combination of the number of product lines and their quality distributions. Figure 10 illustrates the very tight correspondence of product lines to firm size in our model, with the latter normalized to the average quality level in the economy, which builds additional connections and intuitions to the frameworks of Klette and Kortum (2004) and Lentz and Mortensen (2008).

Figure 11 demonstrates that our framework generates an invariant product line distribution at the firm level that resembles an exponential distribution. Combined with the quality margin, the invariant firm size distribution is illustrated in figure 12. Similarly to prior papers, the tails of the sales distribution in our model are not as fat as in the data.²⁸

2. Growth Decomposition

We now use the structure of our model to document the sources of growth. In our model, growth is driven by (i) new entrants, (ii) incumbents doing

²⁸ See Gabaix (2009) for an excellent review of the literature on firm size distribution.




FIG. 9.—Innovation intensity x_n . Color version available as an online enhancement.

internal innovations on their existing lines, and (iii) incumbents expanding into other lines through external innovations:

$$g = \underbrace{x_e \bar{s}}_{\text{entry}} + \underbrace{\sum_{n=0}^{\infty} F \mu_n X_n \bar{s}}_{\text{incumbent external}} + \underbrace{z \lambda}_{\text{incumbent internal}} .$$

Table 7 reports the magnitudes of each of these components in our model.

Our model estimates that 26 percent of aggregate productivity growth is driven by new entry. Of the three-quarters of productivity growth that comes from the action of incumbent firms, the majority of it depends on external innovation efforts of firms. These figures are consistent with the empirical findings surveyed by Foster, Haltiwanger, and Krizan (2001), recognizing that some of our external innovation effect would be viewed as entry/exit in prior empirical calculations.

Another important distinction between external innovation and internal innovation is the differential impacts on qualities. The average step size associated with external innovations is $\bar{s} = 0.069$, whereas the step size of internal innovation is $\lambda = 0.051$, which implies that an average external innovation has 35 percent ($= 0.069/0.051 - 1$) higher impact than internal innovation.




FIG. 10.—Firm size versus number of product lines. Color version available as an online enhancement.

An interesting implication of the estimated model is that it costs more for large firms to produce major innovations. To see how big this additional cost is, let us define a cost multiplier $\mathcal{K}(n)$:

$$\mathcal{K}(n) \equiv \frac{R_x(x_n|n)/n}{R_x(x_n|1)},$$

where $R_x(x_n|n)$ is the cost of producing major innovations at the rate θx_n . Note that this cost multiplier captures the additional percentage cost of producing the same amount of major innovations per product line. Simple algebra shows that the cost multiplier can be expressed as

$$\mathcal{K}(n) = n^{\bar{\sigma}-1} = n^{(1-\sigma-\psi)/\psi}.$$

Figure 13 plots the cost multiplier according to our parameter estimates. This figure, together with figure 11, indicates that a firm at the 90th percentile pays 25 percent more compared to a one-product firm. Likewise, a firm that is at the 99th percentile pays 45 percent more, on average. Therefore, an important takeaway from our estimates is that even a small




FIG. 11.—Product line distribution. Color version available as an online enhancement.

departure from constant returns to scale ($\sigma + \psi = 0.9$) could result in a sizable increase in innovation cost with firm size.

Finally, figure 14 plots the fraction of major advances in a firm's innovation portfolio $\theta x_n / (x_n + z)$ against firm size. Moving from a median-sized firm to a 90th percentile firm reduces the fraction of major advances by around 10 percent; the decline is 16 percent when we move to a 99th percentile firm.




FIG. 12.—Firm size distribution. Color version available as an online enhancement.

TABLE 7
GROWTH DECOMPOSITION

ACTUAL VALUES			IN PERCENTAGE TERMS		
Internal	External	New Entry	Internal	External	New Entry
.0020	.0055	.0026	19.8%	54.5%	25.7%

3. Comparison of Untargeted Moments

We next compare our quantified model against untargeted features of the data. We do this through nonparametric regressions that compare variables across the firm size distribution. We include indicator variables by firm size quintile, with the smallest firm size category serving as the reference group. Our model estimation targets only the annual linear relationship for firm size and growth, and so the degree to which we observe comparable patterns for other variables across the firm size distribution provides confidence in the model's performance. For the exercises, we use the continuous innovation sample in both data sets so that all variables are defined and the samples remain consistent over tests. We struc-




FIG. 13.—Cost multiplier for major innovation, $\mathcal{K}(n)$. Color version available as an online enhancement.

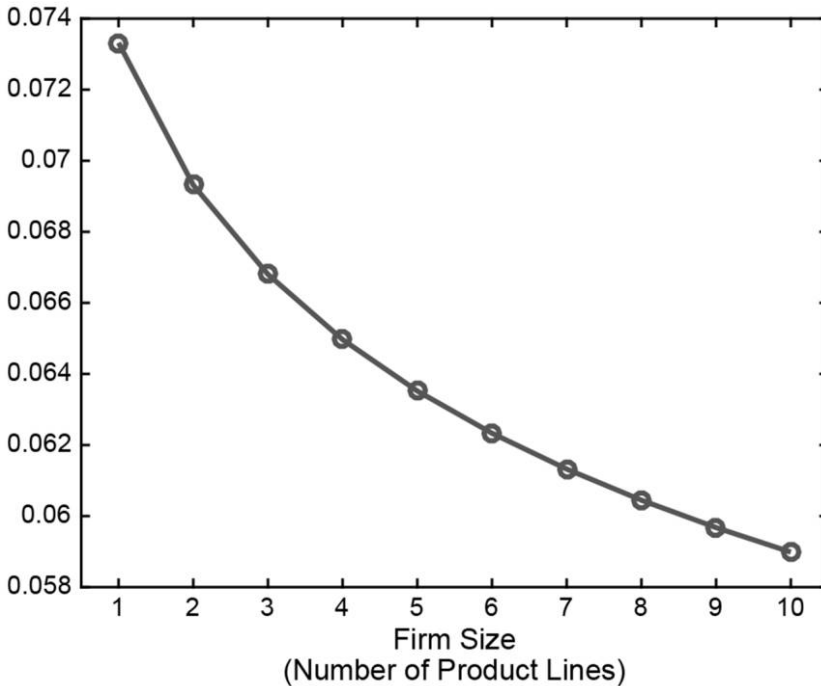


FIG. 14.—Share of major advances in firm's innovation portfolio. Color version available as an online enhancement.

ture our model simulation such that the model-developed data have statistical properties comparable to our Census Bureau data ($n = 32,768$).²⁹

Table 8 considers four main variables for which we have provided initial empirical evidence thus far. On all four dimensions, the model closely matches the data in terms of the direction of differences across the firm size distribution: slower growth, lower patents per employee, higher share of patents being internal, and a lower share of patents being in the top 10 percent in terms of external impact. The model predicts a larger 5-year growth differential between the smallest quintile and the second quintile than present in the data, but the differences for larger quintiles are quite similar. Patents per employee are very similar in levels and direction. The

²⁹ We continue to organize our sample around 5-year blocks. The three periods included in the regressions are 1978–82, 1983–87, and 1988–92, and we use earlier and later data to calculate variables as required. In estimations with Census Bureau data, we include $\eta_{i,t}$ fixed effects for the industry i and year t of the firm. Industries are assigned to firms at the two-digit level of the Standard Industrial Classification system using industries in which firms employ the most workers. All estimations cluster standard errors at the firm level and are unweighted.

TABLE 8
FIRM SIZE DISTRIBUTION AND DATA-MODEL COMPARISON

	Growth Rate to Next Period	Normalized Patent per Employee	Internal Patent Share	Top 10% Patent Share
A. Model, Effects Relative to Smallest-Size Quintile				
2nd quintile	-.1284 (.0210)	-.8194 (.0392)	-.0134 (.0114)	-.0032 (.0063)
3rd quintile	-.2159 (.0199)	-1.1065 (.0379)	-.0116 (.0111)	-.0055 (.0060)
4th quintile	-.3202 (.0191)	-1.3404 (.0372)	.0256 (.0105)	-.0059 (.0056)
Largest quintile	-.3866 (.0188)	-1.5507 (.0368)	.0538 (.0099)	-.0065 (.0053)
B. Data, Effects Relative to Smallest-Size Quintile				
2nd quintile	-.0133 (.0502)	-.9067 (.0336)	.0190 (.0044)	-.0030 (.0078)
3rd quintile	-.2790 (.0464)	-1.0780 (.0320)	.0356 (.0048)	-.0211 (.0076)
4th quintile	-.2865 (.0462)	-1.1166 (.0322)	.0413 (.0047)	-.0296 (.0072)
Largest quintile	-.4052 (.0448)	-1.1351 (.0323)	.0471 (.0045)	-.0211 (.0072)

NOTE.—Estimates are unweighted and cluster standard errors by firm.

model underpredicts the initial rise in internal patent shares present in the data, but the effects for the largest quintiles are very close. Finally, the model underpredicts the steepness of the decline in top/radical patents but otherwise shows a very similar coefficient pattern.³⁰ Overall, these results are very encouraging given that the model has not been targeting these firm size distribution components or time dimension.

Table 9 continues with this approach and considers the patent quality distribution more broadly. We calculate the share of patents for each firm-period that fall within the indicated quartile of the quality distribution. In the data, these quality distributions are measured through external citations relative to the application year and technology of the patent. The model again performs quite well in this untargeted test. Perhaps most striking, the model correctly predicts the disproportionate mass of patents for the largest firms falling within the second-quality quartile, and it gets the relative size of this effect very close to the data. This part of the distribution is where internal patents sit and is a very distinctive piece of the framework developed in this paper. The model also correctly predicts that most of this extra mass is being shifted from the top quartile of external impact.³¹

³⁰ The model coefficients are not statistically different from zero for the last column. In unreported estimations, we develop a larger model sample of 152,089 data points, where we find a largest-quintile impact of -0.0071 (0.0017). Thus, our attention focuses mainly on the coefficient magnitudes between the model and data vs. statistical precision. The complete results for tables 8–10 with the larger sample are available on request and are very similar to those reported.

³¹ The largest firms in panel B also show some modest mass at the lowest quartile. In the model, the constant internal step size λ concentrates the internal effect into a single quartile. The fact that overall we match the quality distribution so well indicates that this simplifying structure is a reasonable approximation.

TABLE 9
FIRM SIZE DISTRIBUTION AND PATENT QUALITY DISTRIBUTION COMPARISON

	SHARE OF FIRM PATENTS IN QUALITY DISTRIBUTION RANGE			
	[0, 25)	[25, 50)	[50, 75)	[75, 100]
A. Model, Effects Relative to Smallest-Size Quintile				
2nd quintile	.0039 (.0091)	-.0081 (.0122)	-.0133 (.0094)	-.0091 (.0097)
3rd quintile	.0111 (.0090)	-.0055 (.0119)	-.0051 (.0090)	-.0107 (.0094)
4th quintile	.0012 (.0082)	.0153 (.0112)	-.0028 (.0083)	-.0137 (.0088)
Largest quintile	-.0108 (.0077)	.0386 (.0104)	-.0045 (.0078)	-.0232 (.0082)
B. Data, Effects Relative to Smallest-Size Quintile				
2nd quintile	-.0079 (.0079)	.0054 (.0090)	.0074 (.0095)	-.0049 (.0106)
3rd quintile	.0039 (.0081)	.0317 (.0093)	-.0008 (.0094)	-.0349 (.0105)
4th quintile	.0122 (.0078)	.0405 (.0090)	.0025 (.0092)	-.0552 (.0102)
Largest quintile	.0140 (.0074)	.0327 (.0080)	.0037 (.0086)	-.0503 (.0099)

NOTE.—See table 8.

Table 10 finally compares firm-level growth regressions in the model and data. These tests evaluate whether the microdynamics of firms behave similarly as we consider all elements together. We use the continuous innovator samples and 5-year periods. The central regressors to explain employment growth to the next period are the firm's current employment, the firm's total patenting in the period, the quality distribution of the firm's own patents in this period (Patent Quality Share_{f,q}), and the share of a firm's patents that are internal in nature (Internal Share_{f,q}). Specifications take the form

$$\begin{aligned} \text{EmpGr}_{f,t} = & \eta_{i,t} + \gamma_E \ln(\text{Emp}_{f,t}) + \gamma_P \ln(\text{Patents}_{f,t}) \\ & + \sum_{q \in Q_p} (\beta_q \cdot \text{Patent Quality Share}_{f,q}) \\ & + \sum_{q \in Q_i} (\theta_q \cdot \text{Internal Share}_{f,q}) + \epsilon_{f,t}, \end{aligned}$$

where f and t index firms and 5-year periods. The set of patent quality quartiles Q_p are indexed by q , and we measure effects relative to the low-

TABLE 10
FIRM-LEVEL REGRESSION COMPARISON

	DEPENDENT VARIABLE IS GROWTH TO NEXT PERIOD		
	Model	Data Using Citations for Quality	Data Using Claims for Quality
Log employment _t	-.0980 (.0032)	-.0983 (.0075)	-.1012 (.0076)
Log patents _t	.1091 (.0048)	.1310 (.0125)	.1330 (.0125)
Share patents [50, 75] _t	.0894 (.0150)	.1004 (.0379)	-.0015 (.0397)
Share patents [75, 100] _t	.0734 (.0135)	.3659 (.0399)	.1274 (.0382)
(0, 1) medium internal patents _t	-.0579 (.1105)	-.0473 (.0329)	-.0431 (.0323)
(0, 1) high internal patents _t	-.1056 (.0085)	-.1870 (.0321)	-.2036 (.0323)

NOTE.—See table 8.

est two quality quartiles. For internal patents, we define indicator variables for internal patents being a (0, 20 percent] share of the firm's total innovation during the period or greater than 20 percent.

On the whole, the model and data display very similar properties at the micro level. Firm growth is increasing in total patents, increasing in the share of these patents falling in the upper half of the distribution, and decreasing in the share of the patents that are internal in nature. The data tend to show greater growth effects with patent quality than the model for the very top quartile, but most of the coefficient magnitudes are quite comparable. In the last column, we use patent claims to measure quality and find comparable results.³²

Sections B and C of appendix C report additional data analyses that confirm features present in the model. Section B shows that the patents that firms develop in their first 2 years of existence have higher external impact than those subsequently developed by the same firm. Section C shows that the external innovation that builds on a particular invention tends to have greater forward impact than the internal innovation that also builds on the same invention. These two features are distinctive elements of our model structure that are important to confirm in the data. Our NBER working paper (Akcigit and Kerr 2010) also provides additional empirical elements that support the model's features. We show, for example, that the external citation distributions that exist for an external patent do not depend on the size of the firm making the patent. This invariance provides support for our model's structure that relates firm size to choices over types of innovations, rather than firms of different sizes having inherently different capacities for producing high-quality innovations.

C. *Robustness*

This section considers robustness checks that extend the moments used to estimate parameters. Across these upcoming variations, we continue to conclude that $\sigma + \psi = 0.9$ is a good estimate for the level of decreasing returns to external innovation in firm size.

1. Adding the Fraction of Top Innovations as a Target

Our model predicts that the fraction of major innovations in a firm's portfolio tends to be decreasing in firm size if external innovation does not scale one-for-one. This theoretical prediction was empirically verified in Section II.E, and we used this as an untargeted moment to assess

³² While citations are the more commonly used measure, there is some concern that firm growth or survival could influence future external citations (e.g., out of fear of litigation). We thank a referee for pointing out this feature, which is not directly testable as quality would be observationally similar. Claims provide a check against this concern.

TABLE 11
ROBUSTNESS WITH FACTS 1 AND 2

Moment	Data	Model
Profitability	.109	.106
R&D intensity	.041	.041
Internal/external cite	.774	.767
Fraction of internal patents	.215	.250
Entry rate	.058	.066
Average growth rate	.010	.010
Growth vs. size (fact 1)	-.035	-.038
Top innovation vs. size (fact 2)	-.0034	-.0034

NOTE.—Estimated $\sigma = .395$; implied $\sigma + \psi = .895$.

the model. As an alternative exercise, we introduce this empirical moment as an additional target. Table 11 reports the new moments and the new estimate of σ . To save space, the rest of the parameter estimates are not reported.

The model replicates both facts very closely while also preserving the goodness of fit with the rest of the moments. The resulting estimated σ value is very similar at 0.395.

2. Adding Patent per Employment as a Target

Table 12 further incorporates the normalized patents per employment regression coefficient as an additional target. While the fit of the first two facts declines with this augmented model, all three relationships are still captured. Most important, the scaling estimate $\sigma = 0.407$ remains robustly identified.

3. Alternative Growth Cap

The major moment influencing σ in the benchmark estimation in table 5 is the empirical relationship between firm size and growth. To confirm that

TABLE 12
ROBUSTNESS WITH FACTS 1, 2, AND 3

Moment	Data	Model
Profitability	.109	.113
R&D intensity	.041	.049
Internal/external cite	.774	.806
Fraction of internal patents	.215	.272
Entry rate	.058	.059
Average growth rate	.010	.009
Growth vs. size (fact 1)	-.035	-.057
Top innovation vs. size (fact 2)	-.0034	-.0061
Patent per employment vs. size (fact 3)	-.182	-.120

NOTE.—Estimated $\sigma = .407$; implied $\sigma + \psi = .907$.

TABLE 13
ROBUSTNESS WITH GROWTH RATE MAXIMUM

Moment	Data	Model
Profitability	.109	.106
R&D intensity	.041	.041
Internal/external cite	.774	.732
Fraction of internal patents	.215	.252
Entry rate	.058	.066
Average growth rate	.010	.010
Growth vs. size (fact 1)	-.048	-.046

NOTE.—Estimated $\sigma = .384$; implied $\sigma + \psi = .884$.

these results are not sensitive to the winsorization imposed, in table 13 we keep all parameters at their baseline levels and reestimate σ with the maximum growth rate of 3,000 percent, versus 1,000 percent in our baseline.

This adjustment lowers σ to 0.384, which is intuitive given that the weaker winsorization allows us to pick up even more abnormal growth for smaller firms, but the influence on our results is overall quite modest.

4. Alternative R&D Elasticities

Table 14 studies the robustness of our results to alternative estimates of the R&D elasticity, centered on the $\psi = 0.5$ elasticity from the micro studies (see the discussion in Sec. VI.A.1). Panel A considers a lower value of $\psi = 0.4$, whereas panel B considers a larger value $\psi = 0.6$.

TABLE 14
ROBUSTNESS WITH DIFFERENT R&D ELASTICITIES

Moment	Data	Model
A. $\psi = .4$		
Profitability	.109	.097
R&D intensity	.041	.041
Internal/external cite	.774	.773
Fraction of internal patents	.215	.252
Entry rate	.058	.067
Average growth rate	.010	.009
Growth vs. size (fact 1)	-.035	-.036
B. $\psi = .6$		
Profitability	.109	.107
R&D intensity	.041	.039
Internal/external cite	.774	.798
Fraction of internal patents	.215	.228
Entry rate	.058	.068
Average growth rate	.010	.010
Growth vs. size (fact 1)	-.035	-.036

NOTE.—In panel A, estimated $\sigma = .497$; implied $\sigma + \psi = .897$. In panel B, estimated $\sigma = .283$; implied $\sigma + \psi = .883$.

The model continues to replicate the targeted moments well. The remarkable result is the robustness of the sum of the elasticity parameters $\sigma + \psi \approx 0.9$, which conforms to benchmark estimates.

VII. Conclusion

Firms come in many shapes and sizes, as do their innovations. An important step for research on the origins of innovation and endogenous growth is to build an apparatus that can handle more of this firm-level heterogeneity; it is equally important to discern when this apparatus adds value commensurate with its extra complexity. This paper takes a step forward on both of these dimensions. First, our model allows for internal and external innovations, links firm innovation choices to firm size, and traces out consequences of these differences for firm-level dynamics and aggregate growth rates. The model remains tractable with these added ingredients, laying bare some economic factors that can lie behind empirical regularities such as deviations from Gibrat's law or the disproportionate representation of small firms and start-ups among the producers of major innovations. We also quantified a generalized form of our model using US data from the Census Bureau for 1982–97, finding that decreasing returns to external innovation in larger firms are an important but not a radical departure from the perfect scaling of the Klette and Kortum (2004) framework.

Among these contributions, our paper is also quite novel in how it layers on patents and citations across patents to inform the model behavior, building on prior work such as Caballero and Jaffe (1993) and Eeckhout and Jovanovic (2002). Indeed, estimations of our model and the scaling parameters would not have been possible otherwise. This work also allows us to conclude that growth impacts of external innovation have exceeded internal innovation for the recent US economy, which in turn helps identify some of the special role that small, innovative firms and new entrants can play in economic growth. There is great potential for further developing this link of patents and patent citations and the information they contain into growth models. Our framework is a natural launching point for estimating the role of intellectual property protections for the incentives to innovate and the subsequent trade-offs that come with monopoly rights. As a second example, one could follow inventors out of large incumbent firms and into the formation of new companies to study the role of spawning new firms in economic growth and the implications of regulations like noncompete clauses. Growth models can garner greater insights and realism by layering information similar to patents and citations that can be studied in both the model and data.

Appendix A

Proofs of Propositions

Proof of Lemma 1

Note that

$$Y^* = (1 - \beta)^{(1-2\beta)/\beta} \tilde{\beta}^{(\beta-1)/\beta} L^* \bar{q}.$$

Therefore, the growth rate of aggregate output is equivalent to the growth rate of the average quality of product lines. We can express the level of $\bar{q}(t)$ after an instant Δt as

$$\begin{aligned} \bar{q}(t + \Delta t) &= \{\bar{q}(t)[\tau^* \Delta t(1 + \bar{s}) + z^* \Delta t(1 + \lambda)] \\ &\quad + \bar{q}(t)[1 - \tau^* \Delta t - z^* \Delta t]\}. \end{aligned}$$

Now subtract $\bar{q}(t)$ from both sides and divide by Δt and take the limit as $\Delta t \rightarrow 0$:

$$g = \frac{\dot{\bar{q}}(t)}{\bar{q}(t)} = \lim_{\Delta t \rightarrow 0} \frac{\bar{q}(t + \Delta t) - \bar{q}(t)}{\Delta t} \frac{1}{\bar{q}(t)} = \tau^* \bar{s} + z^* \lambda.$$

QED

Proof of Proposition 1

Conjecture that

$$V(\mathbf{q}) = A \sum_{q_j \in \mathbf{q}} q_j. \tag{A1}$$

Substituting this expression into the original value function,

$$r^* A \sum_{q_j \in \mathbf{q}} q_j = \max_{x, \{\tilde{s}\}_j} \left\{ \begin{aligned} &\sum_{q_j \in \mathbf{q}} \pi^* q_j - \sum_{q_j \in \mathbf{q}} \hat{\chi} z_j^\psi q_j - \Phi \bar{q} \\ &- \tilde{\chi} x^\psi \bar{q} + xA\bar{q}(1 + \bar{s}) \\ &+ \sum_{q_j \in \mathbf{q}} z_j A q_j \lambda - \sum_{q_j \in \mathbf{q}} \tau^* A q_j \end{aligned} \right\}.$$

This expression holds if and only if

$$r^* A = \max_z \left\{ \pi^* - \hat{\chi} z^\psi + zA\lambda - \tau^* A \right\} \tag{A2}$$

and

$$\max_x \left\{ xA(1 + \bar{s}) - \tilde{\chi} x^\psi \right\} - \Phi = 0. \tag{A3}$$

Assume for now that there is positive entry (we will verify this later in the proof). Then from the free-entry condition (10) we have

$$A = \frac{\nu}{1 + \bar{s}}. \tag{A4}$$

The maximization in (A2) implies $z = [A\lambda/\hat{\psi}\hat{\chi}]^{1/(\hat{\psi}-1)}$ or

$$z_j = \left[\frac{\lambda\nu}{(1 + \bar{s})\hat{\psi}\hat{\chi}} \right]^{\frac{1}{\hat{\psi}-1}}$$

and

$$\tau = \frac{\pi}{A} + \hat{\chi} \left[\frac{\lambda}{\hat{\psi}\hat{\chi}} \right]^{\frac{\hat{\psi}}{\hat{\psi}-1}} A^{\frac{1}{\hat{\psi}-1}} (\hat{\psi} - 1) - g - \rho,$$

where the last line used the fact that $r = g + \rho$. Since the growth rate is $g = \tau\bar{s} + z\lambda$, the above expression can be further refined as

$$\tau = \frac{1}{1 + \bar{s}} \left[\frac{\pi}{A} - \left[\frac{\lambda}{\hat{\psi}\hat{\chi}} \right]^{\frac{\hat{\psi}}{\hat{\psi}-1}} A^{\frac{1}{\hat{\psi}-1}} \hat{\chi} - \rho \right].$$

Now we turn to the maximization problem in (A3), which delivers the optimal innovation effort (together with [A4]) as

$$x = \left[\frac{\nu}{\hat{\psi}\hat{\chi}} \right]^{\frac{1}{\hat{\psi}-1}}.$$

Hence the condition in (A3) is

$$\max_x \left\{ xA(1 + \bar{s}) - \tilde{\chi}x^{\hat{\psi}} \right\} = \left[\frac{\nu}{\hat{\psi}\hat{\chi}} \right]^{\frac{\hat{\psi}}{\hat{\psi}-1}} \tilde{\chi}(\hat{\psi} - 1).$$

Hence assumption 2 guarantees (A3). QED

Proof of Proposition 2

Conjecture the form $\mu_n^* = \tilde{A}\tilde{B}^n(1/n!)$. Then the flow equations in (27) imply

$$F\tilde{A}\tilde{B}^2\tau + x_e = F\tilde{A}\tilde{B}(x + \tau)$$

and

$$\tilde{B}^2\tau = \tilde{B}(x^* + n\tau^*) - nx.$$

Combining these two equations implies

$$F\tilde{A}\tilde{B}n\tau^* - F\tilde{A}nx + x_e = F\tilde{A}\tilde{B}\tau.$$

This equation can hold for all $n \geq 2$ if and only if $\tilde{B} = x/\tau$ and $\tilde{A} = x_e/Fx$. QED

Proof of Proposition 3

Firm growth is equivalent to the growth of Q_f . After a small time interval, the quality index will be, on average,

$$Q_f(t + \Delta t) = \left\{ \begin{array}{l} x\Delta t [Q_f(t) + \bar{q}(1 + \bar{s})] + \sum_{q_j} z\Delta t [Q_f(t) + \lambda q_j] \\ + (1 - x\Delta t - n_f z\Delta t - n_f \tau \Delta t) Q_f(t) \\ + \sum_{q_j} \tau \Delta t (Q_f - q_j) \end{array} \right\}.$$

Then after some algebra the expected growth rate of a firm is

$$\mathcal{G}(Q_f) = \lim_{\Delta t \rightarrow 0} \frac{Q_f(t + \Delta t) - Q_f(t)}{\Delta t Q_f} = \frac{x\bar{q}(1 + \bar{s})}{Q_f} + z\lambda - \tau,$$

which is decreasing in Q_f . QED

Proof of Proposition 4

Immediate from the text. QED

Proof of Proposition 5

The total probability of having an innovation during Δt is $x\Delta t + n_f z\Delta t$. The probability of having a major innovation with $s_k \geq s_k > \lambda$ is $[1 - (1 - \theta)^{k+1}]x\Delta t$. Then the probability of having a major innovation conditional on a successful innovation is the ratio $[1 - (1 - \theta)^{k+1}]x\Delta t / (x\Delta t + n_f z\Delta t)$. QED

Proof of Proposition 6

Note that the new value function in general form is

$$rV(\mathbf{q}) - \dot{V}(\mathbf{q}) = \max_{\substack{x_n \in [0, \bar{x}], \\ \{z_j \in [0, \bar{z}]\}_{\mathcal{J}_j}}} \left\{ \begin{array}{l} \sum_{q_j \in \mathbf{q}} [\pi q_j - \hat{\chi} z_j^\psi q_j] - \bar{q} \bar{\chi} n^{\bar{\sigma}} x_n^{\bar{\psi}} \\ + n x_n [\mathbb{E}_j V(\mathbf{q} \cup_+ \{q_j + \bar{q}s_j\}) - V(\mathbf{q})] \\ + \sum_{q_j \in \mathbf{q}} z_j [V(\mathbf{q} \setminus_- \{q_j\} \cup_+ \{q_j(1 + \lambda)\}) - V(\mathbf{q})] \\ + \sum_{q_j \in \mathbf{q}} \tau [V(\mathbf{q} \setminus_- \{q_j\}) - V(\mathbf{q})] \end{array} \right\}.$$

Substituting the conjecture $V(\mathbf{q}, \bar{q}) = A \sum_{q_j \in \mathbf{q}} q_j + B_n \bar{q}$ into the above value function we get

$$r \sum_{q_j \in \mathbf{q}} Aq_j + rB_n \bar{q} - B_n \bar{q}g = \max_{\substack{x_n \in [0, \bar{x}], \\ \{z_j \in [0, \bar{z}]\}_{\mathcal{J}_j}}} \left\{ \begin{aligned} & \sum_{q_j \in \mathbf{q}} [\pi q_j - \hat{\chi} z_j^{\bar{\psi}} q_j] - \bar{q} \tilde{\chi} n^{\bar{\sigma}} x_n^{\bar{\psi}} \\ & + nx_n \left[A\bar{q}[1 + \mathbb{E}_j s_j] \right. \\ & \quad \left. + B_{n+1} \bar{q} - B_n \bar{q} \right] \\ & + \sum_{q_j \in \mathbf{q}} z_j Aq_j \lambda \\ & + \sum_{q_j \in \mathbf{q}} \tau [-Aq_j + B_{n-1} \bar{q} - B_n \bar{q}] \end{aligned} \right\}.$$

Now equating the terms with q_j and \bar{q} , we get

$$rA = \max_{z_j} \left\{ \pi - \hat{\chi} z_j^{\bar{\psi}} + z_j A \lambda - \tau A \right\}$$

and

$$rB_n - B_n g = \max_{x_n} \left\{ \begin{aligned} & -n^{\bar{\sigma}} \tilde{\chi} x_n^{\bar{\psi}} \\ & + nx_n [A[1 + \mathbb{E}_j s_j] + B_{n+1} - B_n] \\ & + n\tau [B_{n-1} - B_n] \end{aligned} \right\}.$$

Note that from log utility we have $\rho = r - g$. Hence the two value functions become

$$rA = \pi - \tau A + \max_{z_j} \left\{ z_j A \lambda - \hat{\chi} z_j^{\bar{\psi}} \right\},$$

$$\rho B_n = \max_{x_n} \left\{ \begin{aligned} & -n^{\bar{\sigma}} \tilde{\chi} x_n^{\bar{\psi}} \\ & + nx_n [A[1 + \bar{s}] + B_{n+1} - B_n] \\ & + n\tau [B_{n-1} - B_n] \end{aligned} \right\}.$$

Now we can take the first-order conditions

$$z_j = \left[\frac{A\lambda}{\hat{\psi} \hat{\chi}} \right]^{\frac{1}{\bar{\psi}-1}} \quad \text{and} \quad x_n = \left[\frac{A[1 + \bar{s}] + B_{n+1} - B_n}{\hat{\psi} n^{\bar{\sigma}-1} \tilde{\chi}} \right]^{\frac{1}{\bar{\psi}-1}}.$$

Hence A is defined by the following equation,

$$(r + \tau)A = \pi + A^{\frac{\bar{\psi}}{\bar{\psi}-1}} \left[\frac{\lambda}{\hat{\psi}} \right]^{\frac{\bar{\psi}}{\bar{\psi}-1}} (\hat{\psi} - 1) \hat{\chi}^{\frac{1}{1-\bar{\psi}}},$$

and B_n :

$$B_{n+1} = \left[\frac{(\rho + n\tau)B_n - n\tau B_{n-1}}{\hat{\psi} - 1} \right]^{\frac{\bar{\psi}-1}{\bar{\psi}}} \hat{\psi} \tilde{\chi}^{\frac{1}{\bar{\psi}}} n^{\frac{\bar{\sigma}-\bar{\psi}}{\bar{\psi}}} + B_n - A[1 + \bar{s}].$$

QED

Proof of Proposition 7

First we compute the number of citable patents M . The measure of citable patents after Δt is simply

$$M(t + \Delta t) = [M(t) + 1][x\Delta t(1 - \theta) + z\Delta t] + 1 \times x\Delta t\theta + (1 - x\Delta t - z\Delta t)M(t).$$

Imposing the steady-state condition $M(t + \Delta t) = M(t)$, we find $M = (1/\theta) + (z/x\theta)$.

For any given innovation size $s_k = \eta\alpha^k$, the flow equations for external patents with n citations take the following form:

State:	Inflow	Outflow
$n = 0$:	$\Psi_{k-1}\tau(1 - \theta)$	$= M\Upsilon_{s,0}\tau\theta + M\Upsilon_{s,0}\gamma\eta\alpha^k(\tau(1 - \theta) + z)$ (A5)
$n \geq 1$:	$M\Upsilon_{s,n-1}\gamma\eta\alpha^k(\tau(1 - \theta) + z)$	$= M\Upsilon_{s,n}\tau\theta + M\Upsilon_{s,n}\gamma\eta\alpha^k(\tau(1 - \theta) + z)$.

The first line represents size s_k innovations with no citations ($n = 0$). Inflows come from Ψ_{k-1} product lines where the latest follow-up innovation was of size $\eta\alpha^{k-1}$ and a new follow-up innovation brings the product line into the Ψ_k group. This occurs at rate $\tau(1 - \theta)$. This inflow is not dependent on the number of citable patents M , as it depends only on the rate of external advancement across product lines. All patents initially have zero citations, and only a single patent can arrive per product line at any instant. The inflow thus depends only on the rate of affected product lines.

The outflow of this $n = 0$ group depends on $M\Upsilon_{s,0}$, the number of patents for each innovation size s_k . The first part of the outflow occurs when the technology cluster is replaced through a new major innovation at the rate $\tau\theta$, as the affected patents become defunct and are no longer considered for citation. The second part of the outflow occurs when patents receive a new citation from subsequent innovations at the rate $\gamma\eta\alpha^k(\tau(1 - \theta) + z)$. This latter expression is the probability of citation based on step size ($\gamma\eta\alpha^k$) multiplied by the arrival rate of subsequent patents. In this case, patents remain active but move up the citation distribution.

Similar reasoning applies to the second row, where citations $n \geq 1$, except that the inflow occurs only from the $(k, n - 1)$ group. These innovations arrive at rate $\tau(1 - \theta) + z$, now also depending on incumbent advances, and they cite the specific patent at rate $\gamma\eta\alpha^k$.

Next we characterize the citation distribution of internal patents with flow equations:

State:	Inflow	Outflow
$n = 0$:	z	$= M\Upsilon_{\lambda,0}\tau\theta + M\Upsilon_{\lambda,0}\gamma\lambda(\tau(1 - \theta) + z)$ (A6)
$n \geq 1$:	$M\Upsilon_{\lambda,n-1}\gamma\lambda(\tau(1 - \theta) + z)$	$= M\Upsilon_{\lambda,n}\tau\theta + M\Upsilon_{\lambda,n}\gamma\lambda(\tau(1 - \theta) + z)$.

These flows have a similar interpretation. The substantive difference is that the inflow of zero-cited patents occurs at rate z for internal improvements, accumulating across realized success from the internal R&D efforts of the incumbent firm in each product line.

The equations (19) and (A5) imply

$$\Upsilon_{s,0} = \frac{\tau(1-\theta)^k \theta}{M[\tau\theta + \gamma s_k(\tau(1-\theta) + z)]},$$

and we can rewrite the second line of (A5) in a recursive form as

$$\Upsilon_{s,n} = \Upsilon_{s,n-1} \frac{\gamma s_k(\tau(1-\theta) + z)}{\tau\theta + \gamma s_k(\tau(1-\theta) + z)},$$

which implies

$$\Upsilon_{s,n} = \Upsilon_{s,0} \left[\frac{\gamma s_k(\tau(1-\theta) + z)}{\tau\theta + \gamma s_k(\tau(1-\theta) + z)} \right]^n.$$

Similar reasoning applies to $\Upsilon_{\lambda,n}$ and to the flow equations (A6).

For the second part of the theorem, we just rewrite the same flow equations without the internal citations z . Then the expressions follow. QED

Appendix B

Full Predictions of the Baseline Model

This appendix outlines the full set of predictions for the baseline theoretical model without scaling. Most predictions are general and do not depend on whether internal or external innovation has a larger average step size. Predictions C3, D5, and D6 are specific to the case of external innovation having the larger step size, which we find empirically to be true. Our NBER working paper (Akcigit and Kerr 2010) provides the proofs of these predictions.

A. Firm Size Distribution and Firm Growth Rates

- A1. The size distribution of firms is highly skewed.
- A2. The probability of a firm's survival is negatively related to its size.
- A3. Small firms that survive tend to grow faster than larger firms. Among larger firms, this negative relationship weakens.
- A4. The variance of growth rates is higher for smaller firms.
- A5. Younger firms have a higher probability of exiting, but those that survive tend to grow faster than older firms.

B. Firm Size Distribution and Innovation Intensity

- B1. R&D expenditures increase with firm size among innovative firms, but the intensity of R&D decreases with firm size.
- B2. Similarly, patent counts increase with firm size among innovative firms, but the intensity of patenting decreases with firm size.
- B3. Younger firms are more R&D and patent intensive than older firms.

C. *Patent Citation Behavior and Innovation Spillover Size*

- C1. A large fraction of patents receive zero external citations.
- C2. The distribution of citations is highly skewed.
- C3. An average external patent receives more external citations than an internal patent.
- C4. The distribution of patent citation life is highly skewed.

D. *Innovation Type and Firm Size Distribution*

- D1. The proportion of a firm's patents that receive zero future external citations rises with firm size.
- D2. The proportion of a firm's given citations that are self-citations rises with contemporaneous firm size.
- D3. Average future external citations per patent is decreasing in firm size.
- D4. The relative rate of major innovations (highly cited patents) is higher for small firms. This higher relative rate weakens with more stringent citation quality thresholds.
- D5. The average citations (received) of patents by entrants is higher than the average citations of patents by incumbents. Similarly, the average citations of patents by young firms is higher than the average citations of patents by older firms.
- D6. The patents made by firms at their entry, on average, receive more external citations than later patents of the same firm.

E. *Innovation Type and Firm Growth Rates*

- E1. More cited patents lead to higher growth for a firm. This effect is larger for small firms.
- E2. An external patent leads to higher growth than an internal patent on average.
- E3. More R&D- and patent-intensive firms grow faster.
- E4. Everything else equal, firms that obtain more external patents are more likely to survive. Firms that receive more external citations are more likely to exit the economy.

Appendix C

Additional Empirical Results

We include here some selected empirical results that provide special details relevant to our model. Our working paper (Akcigit and Kerr 2010) contains additional results.

A. *Monte Carlo Simulations of Internal Patent Citations*

Table C1 considers in greater detail the observation made in Section II that self-citation behavior rises with firm size. We study this issue using patent data and

assignees, which allows us to undertake the simulations outside of the Census Bureau. We consider patterns for patents filed in 1995 and their citations over the previous 5 years. This short period lowers the computation demands of the simulations, and this snapshot is very representative of the general behavior across the full sample. In 1995, the self-citation share grows from 9 percent for firms filing just one patent to 17 percent for firms filing two to five patents. The share further increases to 31 percent for firms filing over 100 patents.

The last three columns of table C1 evaluate these observed self-citation shares against counterfactuals. Large patenting firms are more likely to cite themselves because of the greater likelihood that they draw on their past inventions. This is true even if citations are random. If IBM and a small firm in 1995 draw a random citation for the computer industry from 1990–95, the likelihood that IBM draws itself is much greater. The likelihood of self-citing for a new entrant is naturally zero. This bias to firm size is particularly true where large firms dominate narrow technology fields.

To confirm that this mechanical effect is not driving the observed relationship in column 2, we undertake Monte Carlo simulations in which we replace observed patents with random counterfactuals. For each observed citation, we draw a counterfactual that matches the technology and application year of the cited patent. We include the original citation among the possible pool of patents, and we draw with replacement. We measure from the simulation a counterfactual self-citation share to assignee size relationship. As this relationship depends on the randomness of the simulation draws, we repeat the procedure 1,000 times.

We use these 1,000 simulations to generate 95 percent confidence bands for the self-citation ratio of each assignee. These confidence bands are specific to assignees on the basis of their size and underlying technologies. These confidence bands more rigorously test whether the observed self-citation relationships are a systematic departure from the null hypothesis of being randomly determined. As anticipated, column 3 shows that the mean value of the test statistic is rising in firm size.

Columns 4 and 5 confirm that the observed self-citation behavior is a significant departure among large assignees. Column 4 examines the prevalence of departures. For assignees with one patent during 1995, only 13 percent display self-citation behavior that we can reject as being random at a 95 percent confidence level. This nonrandom share grows to 97 percent for assignees with more than 100 patents in 1995. Column 5 also shows that average deviation of self-citation shares from the random baseline is growing in firm size. These departures indicate that our results are due to firm behavior rather than the mechanics of firm size. These self-citation findings hold in within-firm panel analyses, too.³³

³³ This analysis closely relates to the patent localization work of Jaffe et al. (1993) and Thompson and Fox-Kean (2005). Similar procedures are used in agglomeration calculations such as Duranton and Overman (2005) and Ellison, Glaeser, and Kerr (2010). Agrawal et al. (2010) discuss related issues with respect to large patenting firms in “company towns” and their self-citation behavior (e.g., Eastman Kodak in Rochester, NY).

TABLE C1
 CROSS-SECTIONAL RELATIONSHIP OF ASSIGNEE SIZE AND SELF-CITATION BEHAVIOR
 COMPARISON OF OBSERVED SELF-CITATION BEHAVIOR AGAINST
 1,000 MONTE CARLO SIMULATIONS REPLICATING TECHNOLOGIES
 AND CITATION YEARS

	COUNT OF ASSIGNEES BY NUMBER OF 1995 PATENTS WITH CITATIONS FOR PATENTS OVER THE PRIOR 5 YEARS (1)	MEAN OBSERVED SELF- CITATION SHARE FOR PATENTS OVER THE PRIOR 5 YEARS (2)	Mean Test Statistic for 95% Confidence Level by Size Category (3)	Share of Firms Deviat- ing at 95% Confidence Level from Random Behavior (4)	Mean Deviations of Observed Citation Shares (Col. 2 Minus Col. 3) (5)
1 patent	8,044	9%	1%	13%	8%
2-5 patents	3,382	17%	3%	35%	14%
6-10 patents	595	22%	4%	64%	18%
11-20 patents	307	23%	4%	73%	19%
21-100 patents	288	27%	4%	89%	23%
100+ patents	65	31%	6%	97%	25%

NOTE.—The table reports the results of Monte Carlo simulations of self-citation behavior by firm size. The sample is restricted to US-based, industrial patents in 1995 and their citations to other US-based, industrial patents over the prior 5 years. Rows group assignees by their patent counts in 1995. Column 2 indicates the share of observed citations that are self-citations. For the Monte Carlo simulations, we draw counterfactuals that match the technologies and application years of cited patents. We include the original citation among the possible pool of patents, and we draw with replacement. We measure from the simulation a counterfactual self-citation share to assignee size relationship. We repeat the simulations 1,000 times to generate 95 percent confidence bands for the self-citation ratio of each assignee. These confidence bands are specific to assignees based on their size and underlying technologies. Column 3 provides the mean test statistic by firm size. This statistic rises with firm size because firms with larger patent portfolios are more likely to cite themselves even if citations are random. Column 4 indicates the share of assignees by size category that exhibit self-citation behavior that exceeds a random pattern at a 95 percent confidence level. These deviations are strongly increasing in firm size. Column 5 presents the mean deviation of observed self-citation behavior from the simulation baselines. These deviations are also increasing in firm size.

B. Panel Relationship between Entry and Patent Quality

Table C2 presents some simple panel evidence on patent quality within firms over time. We restrict the sample to new entrants during 1977–94. We regress traits of patents on an indicator variable for whether or not the patent is filed in the first 2 years that a firm is observed. We include firm fixed effects to compare early patents of the firm to later patents. We also include technology-year fixed effects. Column 1 shows that the average external citation count is higher at entry. Column 2 shows that patents also have larger numbers of claims at firm entry than in later years. Columns 3–6 show the distribution of external citations in quartiles. Column 3 is the lowest-quality quartile, and column 6 is the highest-quality quartile. Entrants have disproportionate representation in the highest-quality quartile compared to later years for the same firm. The results describe the time path of firms in terms of invention quality.

TABLE C2
PANEL RELATIONSHIP BETWEEN ENTRY AND PATENT QUALITY

	NUMBER OF EXTERNAL CITATIONS (1)	NUMBER OF CLAIMS ON PATENT (2)	PREVALENCE OF PATENTS BY EXTERNAL CITATION RANKS (Coefficients Sum to Zero across Columns)			
			0%– 24% (3)	25%– 49% (4)	50%– 74% (5)	75%– 100% (6)
First 2 years the firm is observed	1.1621 (.1557)	.6920 (.1811)	–.0148 (.0048)	–.0042 (.0059)	–.0048 (.0063)	.0239 (.0058)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Technology-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes

NOTE.—The table quantifies changes in average patent quality within firms over time. Columns 1 and 2 show that external citation rates and claims per patent are higher at firm entry. Columns 3–6 show the distribution of external citations in quartiles. Column 3 is the lowest-quality quartile, and col. 6 is the highest-quality quartile. The coefficients for a row sum to zero across these columns. Entrants have disproportionate representation in the highest-quality quartile compared to later years for the same firm. The sample includes 260,972 US industrial patents for firms first observed between 1977 and 1994. Estimations include firm fixed effects and technology-year fixed effects, cluster standard errors at the firm level, and weight patents such that each firm receives constant weight.

C. Dynamic Evidence on Quality within Firms

Table C3 provides evidence to verify our model's assumption that major external innovations are followed within firms by internal innovations and refinements. This process requires that an external innovation be made to dramatically push forward the technology of a product line that is dominated by internal inventions within the currently leading firm. We can further verify these features by demonstrating that the mean quality of citing patents outside of the original firm for a given invention is higher than the mean quality of citing patents within the firm.

We use a linear specification of the form

$$\text{Cite}_{p_2, p_1} = \phi_{p_1} + \eta_{i,t}^{p_2} + \beta \cdot \text{External}_{p_2, p_1} + \epsilon_{p_2, p_1},$$

where Cite_{p_2, p_1} models traits of patents p_2 that cite patents p_1 . We include citations for

US industrial patents filed during 1975–84. We restrict the citations to be US industrial patents filed within a 10-year window of the original patent. We find similar patterns when using all citations, but the consistent window is more appropriate.

The primary regressor is the indicator variable $\text{External}_{p_2, p_1}$ that takes unit value if the assignee of citing patent p_2 differs from the assignee of cited patent p_1 . Three-quarters of citations are external. We include ϕ_{p_1} fixed effects for cited patents. We thus compare differences between internal and external citations on the same patent. We also include $\eta_{i,t}^{p_2}$ fixed effects for the technology i and year t of the citing patent p_2 ; the patent fixed effects naturally control for these traits for cited patents p_1 . We define $\eta_{i,t}^{p_2}$ through USPTO subcategories and 5-year time periods. We cluster standard errors by cited patents.

The first column of table C3 models the number of external citations on citing patents p_2 as the outcome variable. The second column alternatively tests the number of claims on the citing patent as a measure of quality. Columns 3–6 then test the quality distribution of citing patents in a format similar to table C2. Quality distributions are determined through ranks of external citations by technology and period. Coefficients across the final four columns for a row approximately sum to zero, but the relationship does not hold exactly given that quality distributions are calculated over a larger group than the regression sample.

Column 1 finds that the mean number of future citations for external innovations that builds on a given invention is 0.8 citations higher than the internal innovations that also build on the focal invention. This effect is large relative to the sample mean of 8.2. There is also a substantial external premium of 1.2 claims relative to the sample mean of 15.4. Columns 3–6 show that this effect mainly comes from a greater prevalence of upper-quartile patents among the external citing patents, with mass moved from the lowest two quartiles of the distribution. These patterns suggest that external innovation that builds on a given invention is stronger than the internal innovation that follows.

TABLE C3
ASSIGNEE SIZE AND BUILDING ON TECHNOLOGIES

	NUMBER OF EXTERNAL CITATIONS CITING PATENT (1)	NUMBER OF CLAIMS ON PATENT CIT- ING PATENT (2)	PREVALENCE OF PATENTS BY EXTERNAL CITATION RANKS AMONG CITING PATENTS (Coefficients Sum to Zero across Columns)			
			0%– 24% (3)	25%– 49% (4)	50%– 74% (5)	75%– 100% (6)
External citation	.849 (.053)	1.236 (.073)	-.015 (.002)	-.009 (.002)	-.005 (.002)	.029 (.002)
Cited patent fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Citing technology-year effects	Yes	Yes	Yes	Yes	Yes	Yes

NOTE.—The table characterizes differences in patent quality for internal vs. external patents that cite a particular invention. Columns 1 and 2 show that external citation rates and claims are higher. Columns 3–6 show the quality distribution of the citations by quartiles. Column 3 is the lowest-quality quartile, and col. 6 is the highest-quality quartile. External citations are consistently of higher quality. The sample includes 761,940 citations of US industrial patents from 1975–84 applied for within 10 years after the original patent. Estimations include cited patent fixed effects and technology-period fixed effects for citing patents. Estimations cluster standard errors by cited patent.

D. *Additional Empirical Figures*

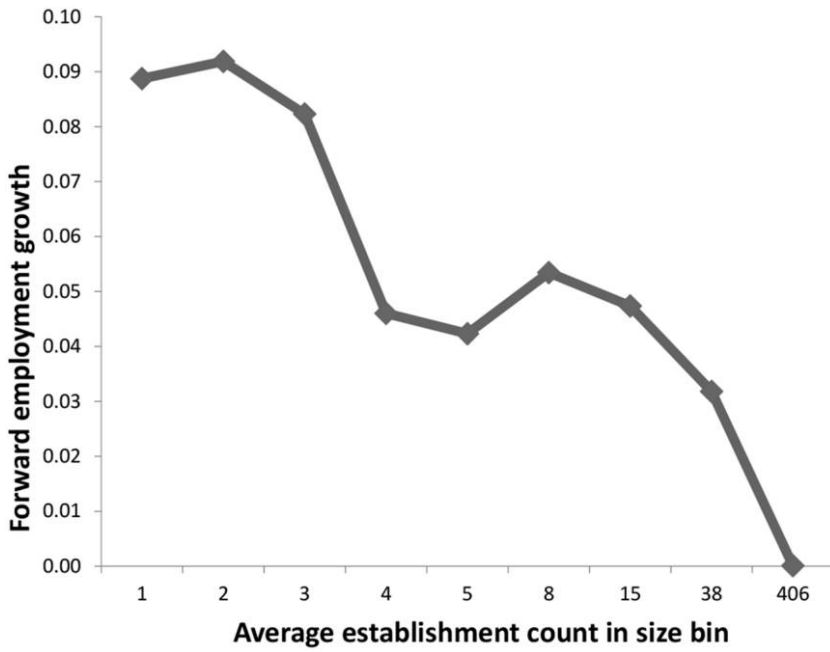


FIG. C1.—Firm growth by firm size Color version available as an online enhancement.




FIG. C2.—Innovation intensity by firm size. Color version available as an online enhancement.

Appendix D

Computer Algorithm

We solve the generalized model as a fixed point over the growth rate g . Our algorithm employs a computational loop with the following steps:

1. Guess a growth rate g .
 - a. Guess a creative destruction rate τ .
 - i. Solve for A in (35), the sequence $\{B_n\}$ in (36), and z_j and $\{x_n\}$ in (37).
 - ii. Verify the free-entry condition as a function of τ : $A[1 + \bar{s}] + B_1 = \nu$.
 - iii. If not converged, update τ and go to step 1(a)i.
 - b. Calculate the growth rate $g = \tau\bar{s} + z\lambda$.
 - c. Update the growth rate. If not converged, go to step 1(a).
2. End the equilibrium solver.
3. Simulate a sample of firms and compute the moments of interest.

The sequence of firm value functions in step 1(a)i is solved using the uniformization method (see Acemoglu and Akcigit [2012] for details). In step 3, we simulate a sample of 2^{13} firms (8,192), split the time into discrete intervals (e.g.,

months), and iterate for 500 years until we obtain convergence. At each iteration, firms gain and lose products according to the flow probabilities specified in the model.

References

- Acemoglu, D. 2008. *Introduction to Modern Economic Growth*. Princeton, NJ: Princeton Univ. Press.
- Acemoglu, D., and U. Akcigit. 2012. "Intellectual Property Rights Policy, Competition and Innovation." *J. European Econ. Assoc.* 10 (1): 1–42.
- Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr. Forthcoming. "Innovation, Reallocation, and Growth." *A.E.R.*
- Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr. 2016. "Transition to Clean Technology." *J.P.E.* 124 (1): 52–104.
- Acemoglu, D., and D. Cao. 2015. "Innovation by Entrants and Incumbents." *J. Econ. Theory* 157:255–94.
- Acs, Z. J., and D. B. Audretsch. 1987. "Innovation, Market Structure, and Firm Size." *Rev. Econ. and Statis.* 69 (4): 567–74.
- . 1988. "Innovation in Large and Small Firms: An Empirical Analysis." *A.E.R.* 78 (4): 678–90.
- . 1991. "Innovation and Size at the Firm Level." *Southern Econ. J.* 57 (3): 739–44.
- Aghion, P., U. Akcigit, and P. Howitt. 2014. "What Do We Learn from Schumpeterian Growth Theory?" In *Handbook of Economic Growth*, vol. 2, edited by P. Aghion and S. N. Durlauf, 515–63. Amsterdam: Elsevier.
- Aghion, P., C. Harris, P. Howitt, and J. Vickers. 2001. "Competition, Imitation and Growth with Step-by-Step Innovation." *Rev. Econ. Studies* 68 (3): 467–92.
- Aghion, P., and P. Howitt. 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60 (2): 323–51.
- Aghion, P., P. Howitt, and J. Vickers. 1997. "Competition and Growth with Step-by-Step Innovation: An Example." *European Econ. Rev.* 41 (3): 771–82.
- Agrawal, A., I. Cockburn, and C. Rosell. 2010. "Not Invented Here? Innovation in Company Towns." *J. Urban Econ.* 67 (1): 78–89.
- Akcigit, U. 2010. "Firm Size, Innovation Dynamics and Growth." Working paper, Univ. Pennsylvania.
- . 2017. "Economic Growth: The Past, the Present, and the Future." *J.P.E.* 125 (6): 1736–47.
- Akcigit, U., H. Alp, and M. Peters. 2015. "Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries." Working Paper no. 21905, NBER, Cambridge, MA.
- Akcigit, U., D. Hanley, and N. Serrano-Velarde. 2016. "Back to Basics: Basic Research Spillovers, Innovation Policy and Growth." Discussion Paper no. 11707, Centre Econ. Policy Res., London.
- Akcigit, U., and W. R. Kerr. 2010. "Growth through Heterogeneous Innovations." Working Paper no. 16443, NBER, Cambridge, MA.
- Arkolakis, C. 2011. "A Unified Theory of Firm Selection and Growth." Working Paper no. 17553, NBER, Cambridge, MA.
- Balasubramanian, N., and J. Sivadasan. 2011. "What Happens When Firms Patent? New Evidence from US Economic Census Data." *Rev. Econ. and Statis.* 93 (1): 126–46.
- Barro, R., and X. Sala-i-Martin. 1995. *Economic Growth*. New York: McGraw-Hill.

- Baumol, W. J. 2002. "Entrepreneurship, Innovation and Growth: The David-Goliath Symbiosis." *J. Entrepreneurial Finance* 7 (2): 1–10.
- Bernstein, S. 2015. "Does Going Public Affect Innovation?" *J. Finance* 70 (4): 1365–1403.
- Bloom, N., R. Griffith, and J. Van Reenen. 2002. "Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997." *J. Public Econ.* 85 (1): 1–31.
- Bloom, N., M. Schankerman, and J. Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4): 1347–93.
- Blundell, R., R. Griffith, and F. Windmeijer. 2002. "Individual Effects and Dynamics in Count Data Models." *J. Econometrics* 108 (1): 113–31.
- Caballero, R. J., and A. B. Jaffe. 1993. "How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth." In *NBER Macroeconomics Annual 1993*, vol. 8, edited by O. J. Blanchard and S. Fischer, 15–86. Cambridge, MA: MIT Press.
- Cabral, L., and J. Mata. 2003. "On the Evolution of the Firm Size Distribution: Facts and Theory." *A.E.R.* 93 (4): 1075–90.
- Cai, A. J. 2010. "Knowledge Spillovers and Firm Size Heterogeneity." Working paper, Univ. New South Wales.
- Caves, R. E. 1998. "Industrial Organization and New Findings on the Turnover and Mobility of Firms." *J. Econ. Literature* 36 (4): 1947–82.
- Christensen, C. 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Boston: Harvard Bus. School Press.
- Cohen, W. 1995. "Empirical Studies of Innovative Activity." In *Handbook of the Economics of Innovations and Technological Change*, edited by P. Stoneman. Oxford: Blackwell.
- Cohen, W. M., and S. Klepper. 1996. "Firm Size and the Nature of Innovation within Industries: The Case of Process and Product R&D." *Rev. Econ. and Statis.* 78 (2): 232–43.
- Davis, S., J. Haltiwanger, and S. Schuh. 1996. *Job Creation and Destruction*. Cambridge, MA: MIT Press.
- Dunne, T., M. J. Roberts, and L. Samuelson. 1988. "Patterns of Firm Entry and Exit in US Manufacturing Industries." *RAND J. Econ.* 19 (4): 495–515.
- Duranton, G. 2007. "Urban Evolutions: The Fast, the Slow, and the Still." *A.E.R.* 97 (1): 197–221.
- Duranton, G., and H. G. Overman. 2005. "Testing for Localization Using Micro-Geographic Data." *Rev. Econ. Studies* 72 (4): 1077–1106.
- Eeckhout, J., and B. Jovanovic. 2002. "Knowledge Spillovers and Inequality." *A.E.R.* 92 (5): 1290–1307.
- Ellison, G., E. L. Glaeser, and W. R. Kerr. 2010. "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns." *A.E.R.* 100 (3): 1195–1213.
- Foster, L., J. Haltiwanger, and C. J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*, edited by C. R. Hulten, E. R. Dean, and M. J. Harper. Chicago: Univ. Chicago Press (for NBER).
- Gabaix, X. 2009. "Power Laws in Economics and Finance." *Ann. Rev. Econ.* 1:255–94.
- Galasso, A., and T. S. Simcoe. 2011. "CEO Overconfidence and Innovation." *Management Sci.* 57 (8): 1469–84.
- Gans, J. S., D. H. Hsu, and S. Stern. 2002. "When Does Start-Up Innovation Spur the Gale of Creative Destruction?" *RAND J. Econ.* 33 (4): 571–86.
- Garcia-Macia, D., C.-T. Hsieh, and P. Klenow. 2016. "How Destructive Is Innovation?" Working paper, Stanford Univ.

- Geroski, P. A. 1998. "An Applied Econometrician's View of Large Company Performance." *Rev. Indus. Org.* 13 (3): 271–94.
- Gilbert, R. J., and D. M. G. Newbery. 1982. "Preemptive Patenting and the Persistence of Monopoly." *A.E.R.* 72 (3): 514–26.
- Griliches, Z. 1990. "Patent Statistics as Economic Indicators: A Survey." *J. Econ. Literature* 28 (4): 1661–1707.
- . 1992. "The Search for R&D Spillovers." *Scandinavian J. Econ.* 94:S29–S47.
- Gromb, D., and D. S. Scharfstein. 2002. "Entrepreneurship in Equilibrium." Working Paper no. 9001, NBER, Cambridge, MA.
- Grossman, G. M., and E. Helpman. 1991. "Quality Ladders in the Theory of Growth." *Rev. Econ. Studies* 58 (1): 43–61.
- Hall, B. H. 1993. "R&D Tax Policy during the 1980s: Success or Failure?" In *Tax Policy and the Economy*, vol. 7, edited by J. S. Poterba, 1–36. Cambridge, MA: MIT Press (for NBER).
- Hall, B. H., A. B. Jaffe, and M. Trajtenberg. 2001. "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." Working Paper no. 8498, NBER, Cambridge, MA.
- . 2005. "Market Value and Patent Citations." *RAND J. Econ.* 36 (1): 16–38.
- Hall, B. H., and J. Van Reenen. 2000. "How Effective Are Fiscal Incentives for R&D? A Review of the Evidence." *Research Policy* 29 (4): 449–69.
- Hall, B. H., and R. H. Ziedonis. 2001. "The Patent Paradox Revisited: An Empirical Study of Patenting in the US Semiconductor Industry, 1979–1995." *RAND J. Econ.* 32 (1): 101–28.
- Hausman, J., B. H. Hall, and Z. Griliches. 1984. "Econometric Models for Count Data with an Application to the Patents-R&D Relationship." *Econometrica* 52 (4): 909–38.
- Hellmann, T., and E. Perotti. 2011. "The Circulation of Ideas in Firms and Markets." *Management Sci.* 57 (10): 1813–26.
- Henderson, R. M. 1993. "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Equipment Industry." *RAND J. Econ.* 24 (2): 248–70.
- Henderson, R. M., and K. B. Clark. 1990. "Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms." *Administrative Sci. Q.* 35 (1): 9–30.
- Hopenhayn, H. A. 1992. "Entry, Exit, and Firm Dynamics in Long Run Equilibrium." *Econometrica* 60 (5): 1127–50.
- Hopenhayn, H. A., G. Llobet, and M. Mitchell. 2006. "Rewarding Sequential Innovators: Prizes, Patents, and Buyouts." *J.P.E.* 114 (6): 1041–68.
- Howitt, P. 1999. "Steady Endogenous Growth with Population and R&D Inputs Growing." *J.P.E.* 107 (4): 715–30.
- Hsieh, C.-T., and P. Klenow. 2014. "The Life-Cycle of Manufacturing Plants in India and Mexico." *Q.J.E.* 129 (3): 1035–84.
- Hurst, E., and B. W. Pugsley. 2011. "What Do Small Businesses Do?" *Brookings Papers Econ. Activity* (Fall): 73–142.
- Jaffe, A. B., M. Trajtenberg, and M. S. Fogarty. 2000. "Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors." *A.E.R. Papers and Proc.* 90 (2): 215–18.
- Jaffe, A. B., M. Trajtenberg, and R. Henderson. 1993. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations." *Q.J.E.* 108 (3): 577–98.
- Jarmin, R. S., and J. Miranda. 2002. "The Longitudinal Business Database." Working paper, Center Econ. Studies, US Census Bureau, Washington, DC.

- Jones, C. I. 1995. "R&D-Based Models of Economic Growth." *J.P.E.* 103 (4): 759–84.
- Jovanovic, B. 1982. "Selection and the Evolution of Industry." *Econometrica* 50 (3): 649–70.
- Jovanovic, B., and G. M. MacDonald. 1994. "The Life Cycle of a Competitive Industry." *J.P.E.* 102 (2): 322–47.
- Kerr, W. R. 2010. "Breakthrough Inventions and Migrating Clusters of Innovation." *J. Urban Econ.* 67 (1): 46–60.
- Kerr, W. R., and S. Fu. 2008. "The Survey of Industrial R&D Patent Database Link Project." *J. Technology Transfer* 33 (2): 173–86.
- Kerr, W. R., R. Nanda, and M. Rhodes-Kropf. 2014. "Entrepreneurship as Experimentation." *J. Econ. Perspectives* 28 (3): 25–48.
- Klepper, S. 1996. "Entry, Exit, Growth, and Innovation over the Product Life Cycle." *A.E.R.* 86 (3): 562–83.
- Klepper, S., and E. Graddy. 1990. "The Evolution of New Industries and the Determinants of Market Structure." *RAND J. Econ.* 21 (1): 27–44.
- Klette, T. J., and S. Kortum. 2004. "Innovating Firms and Aggregate Innovation." *J.P.E.* 112 (5): 986–1018.
- Kortum, S. 1997. "Research, Patenting, and Technological Change." *Econometrica* 65 (6): 1389–1420.
- Kortum, S., and J. Lerner. 2000. "Assessing the Contribution of Venture Capital to Innovation." *RAND J. Econ.* 31 (4): 674–92.
- Kueng, L., M.-J. Yang, and B. Hong. 2014. "Sources of Firm Life-Cycle Dynamics: Differentiating Size vs. Age Effects." Working Paper no. 20261, NBER, Cambridge, MA.
- Lamoreaux, N. R., K. L. Sokoloff, and D. Sutthiphisal. 2011. "The Reorganization of Inventive Activity in the United States in the Early Twentieth Century." In *Understanding Long-Run Economic Growth: Geography, Institutions, and the Knowledge Economy*, edited by N. R. Lamoreaux and D. Costa, 235–74. Chicago: Univ. Chicago Press.
- Lentz, R., and D. Mortensen. 2008. "An Empirical Model of Growth through Product Innovation." *Econometrica* 76 (6): 1317–73.
- . 2014. "Optimal Growth through Product Innovation." Working paper, Univ. Wisconsin.
- Lerner, J. 1994. "The Importance of Patent Scope: An Empirical Analysis." *RAND J. Econ.* 25 (2): 319–33.
- . 1997. "An Empirical Exploration of a Technology Race." *RAND J. Econ.* 28 (2): 228–47.
- . 2012. *The Architecture of Innovation: The Economics of Creative Organizations*. Boston: Harvard Bus. School Press.
- Lerner, J., M. Sorensen, and P. Strömberg. 2011. "Private Equity and Long-Run Investment: The Case of Innovation." *J. Finance* 66 (2): 445–77.
- Lucas, R. E. 1978. "On the Size Distribution of Business Firms." *Bell J. Econ.* 9 (2): 508–23.
- Luttmer, E. G. 2007. "Selection, Growth, and the Size Distribution of Firms." *Q.J.E.* 122 (3): 1103–44.
- . 2011. "On the Mechanics of Firm Growth." *Rev. Econ. Studies* 78 (3): 1042–68.
- March, J. G. 1991. "Exploration and Exploitation in Organizational Learning." *Org. Sci.* 2 (1): 71–87.
- Mehta, A., M. Rysman, and T. Simcoe. 2010. "Identifying the Age Profile of Patent Citations: New Estimates of Knowledge Diffusion." *J. Appl. Econometrics* 25 (7): 1179–1204.

- Nelson, R., and S. Winter. 1982. *An Evolutionary Theory of Economic Change*. Cambridge, MA: Harvard Univ. Press.
- Nicholas, T. 2014. "Scale and Innovation during Two U.S. Breakthrough Eras." Working paper, Harvard Bus. School.
- Peretto, P. F. 1998. "Technological Change, Market Rivalry, and the Evolution of the Capitalist Engine of Growth." *J. Econ. Growth* 3 (1): 53–80.
- Rausch, L. 2010. "Indicators of U.S. Small Business's Role in R&D." Info Brief no. 10-304, Nat. Sci. Found., Washington, DC.
- Reinganum, J. F. 1983. "Uncertain Innovation and the Persistence of Monopoly." *A.E.R.* 73 (4): 741–48.
- Romer, P. M. 1986. "Increasing Returns and Long-Run Growth." *J.P.E.* 94 (5): 1002–37.
- . 1990. "Endogenous Technological Change." *J.P.E.* 98, no. 5, pt. 2 (October): S71–S102.
- Rosen, R. J. 1991. "Research and Development with Asymmetric Firm Sizes." *RAND J. Econ.* 22 (3): 411–29.
- Samila, S., and O. Sorenson. 2011. "Venture Capital, Entrepreneurship and Economic Growth." *Rev. Econ. and Statis.* 93 (1): 338–49.
- Spence, M. 1984. "Cost Reduction, Competition, and Industry Performance." *Econometrica* 52 (1): 101–21.
- Sutton, J. 1997. "Gibrat's Legacy." *J. Econ. Literature* 35 (1): 40–59.
- Thomke, S. 2003. *Experimentation Matters: Unlocking the Potential of New Technologies for Innovation*. Boston: Harvard Bus. School Press.
- Thompson, P., and M. Fox-Kean. 2005. "Patent Citations and the Geography of Knowledge Spillovers: A Reassessment." *A.E.R.* 95 (1): 450–60.
- Trajtenberg, M. 1990. "A Penny for Your Quotes: Patent Citations and the Value of Innovations." *RAND J. Econ.* 21 (1): 172–87.
- Wilson, D. J. 2009. "Beggars Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits." *Rev. Econ. and Statis.* 91 (2): 431–36.
- Zucker, L., M. Darby, and M. Brewer. 1998. "Intellectual Human Capital and the Birth of U.S. Biotechnology Enterprises." *A.E.R.* 88:290–306.