Technological Innovation, Resource Allocation, and Growth*

Leonid Kogan, Dimitris Papanikolaou,
Amit Seru and Noah Stoffman

Abstract

We explore the role of technological innovation as a source of economic growth by constructing direct measures of innovation at the firm level. We combine patent data for US firms from 1926 to 2010 with the stock market response to news about patents to assess the economic importance of each innovation. Our innovation measure predicts productivity and output at the firm, industry and aggregate level. Furthermore, capital and labor flow away from non-innovating firms towards innovating firms within an industry. There exists a similar, though weaker, pattern across industries. Cross-industry differences in technological innovation are strongly related to subsequent differences in industry output growth.

JEL classifications: G14, E32, O3, O4

*We thank Hal Varian for helping us in extracting information on Patents from Google Patents database. We are grateful to John Cochrane, Josh Lerner, Petra Moser, and Tom Nicholas for detailed comments. We also thank Nick Bloom, Martin Eichenbaum, Carola Frydman, Jonathan Parker, Monika Piazzesi and seminar participants at NYU, Northwestern Brownbag lunch and SITE conference for helpful discussions. We are grateful to Tom Nicholas for sharing his patent citations data. The authors thank the Fama-Miller Center at University of Chicago, the Zell Center and the Jerome Kenney Fund for financial assistance. Leonid Kogan is from MIT Sloan (lkogan@mit.edu), Dimitris Papanikolaou is from Kellogg School of Management (d-papanikolaou@kellogg.northwestern.edu), Amit Seru is from University of Chicago (amit.seru@chicagobooth.edu) and Noah Stoffman is from Kelly School of Business (nstoffma@indiana.edu).

Introduction

Economists since Schumpeter have argued that technological innovation, combined with resource reallocation, is the engine that sustains long-term economic growth. However, the impact of technical change on economic growth and business cycle fluctuations remains difficult to quantify. Similarly, while technology shocks play a central role in macroeconomic real business cycle models, there is little consensus on whether these shocks represent actual technological improvements, or are reduced-form representations of other economic forces.[1] The primary reason for these ambiguities is the difficulty in measuring technological innovation in the data. This paper aims to fill this gap.

We construct a novel economic measure of innovation that combines information from a patent dataset with stock market data over the period 1926 to 2010.[2] Measuring technological innovation through patents offers important advantages. Patents are a direct measure of innovation that are available as far back as the eighteenth century. However, the use of patents as a measure of innovation has a significant shortcoming: not all patents are of equal economic value. Thus, an increase in the number of patents granted need not coincide with greater technological innovation. Our central idea is to use the stock market reaction around the day the patent is granted to appropriately weigh each patent. On the day that the patent is granted, market participants learn the full details of the patent. We use this stock market reaction as a measure of patent quality to construct measures of innovation at the firm, industry and economy level which allows us to evaluate the reallocation and growth dynamics within and across industries after bursts of innovative activity.

Our approach to measuring the quality of patents offers distinct advantages over the existing measures of patent quality. Patent citations contain valuable information that can be used to assess the quality of patents.[3] However, patent citations suffer from two major drawbacks. First, measuring the number of future citations each patent generates requires information over the entire sample. In many economic applications – such as when exploring the short- and medium-run response of investment or hiring decisions to innovation – it may be more desirable to use a measure that depends on the contemporary assessment of the value of a patent, as is the case with our measure. Second, the patent citation data is reliably available only in the later part of our sample.[4] This lack of information creates problems in assessing the quality of earlier patents, since patents often tend to cite only the most recent ones (Caballero and Jaffe, 1993).[5] In contrast, our measure is reliably available over a long time period allowing us to make meaningful comparisons. Despite these two drawbacks, patent citations provide a valuable independent measure of the *realized* value of a patent. Hence, we use patent citations as a validation of our procedure. We find that the firm's stock market reaction when the patent is granted is a strong predictor of the number of citations the patent receives in the future.

[1]See, for instance, Cochrane (1994).

[2]Several new studies exploit the same source of patent data (Google Patents) as we do in our paper. For instance, see Moser and Voena (2011), Moser, Voena, and Waldinger (2012) and Lampe and Moser (2011).

[3]See, for example, Harhoff, Narin, Scherer, and Vopel (1999), Hall, Jaffe, and Trajtenberg (2005) and Moser, Ohmstedt, and Rhode (2011).

[4]Moser and Nicholas (2004) and Nicholas (2008) discuss issues in extracting citations data from patent documents before 1975. In addition, even in the post-1975 period citation outcomes are affected by the identity of the patent examiner (Cockburn, Kortum, and Stern, 2002).

[5]For instance, the telephone patent by Alexander Graham Bell (patent number 174,465) has only one citation in the Google Patent database. The first year that patent citations are officially included on patent documents is 1947.

Our measure of technological innovation captures known periods of high technological progress as well as firms driving these waves (e.g., technologically progressive 1960s and early 1970s, see Laitner and Stolyarov (2003)). In addition, the empirical distribution of firm-level innovation measure is extremely fat-tailed, since a few large firms contribute disproportionately to the aggregate rate of innovation in the economy. The identity of these firms varies by decade. This finding is consistent with past research which describes the nature of radical innovations (Harhoff, Scherer, and Vopel (1997)). Furthermore, we find that characteristics of innovating firms using our measure match those of innovators as described by Baumol (2002), Griliches (1990) and Scherer (1983).

Armed with our measure, we examine the relation between innovation and economic growth. First, we explore the link between firm productivity and innovation. Our innovation measure is strongly linked to productivity of capital and labor, both at the firm and at the industry level. Firms and industries that innovate experience a surge in productivity and output. In addition, we find several patterns in the data that are consistent with Schumpeter's notion of "creative destruction". The innovation activity of competing firms has a negative effect on firm productivity in the short run. In addition, capital and labor are reallocated towards firms that innovate, away from firms that do not. We find similar patterns across industries. Furthermore, an increase in industry innovation is associated with an increase in the rate of firm exit, consistent with the view that innovation leads to industry shakeouts.

Next, we relate aggregate growth to innovation by estimating the impulse response of aggregate total factor productivity (TFP) and output to our aggregate innovation measures. Our innovation measure accounts for a substantial fraction of movements in aggregate TFP. An increase in innovation is associated with an increase in aggregate output, although with a lag of three to four years. We find similar patterns in the cross-section. Differences in innovation are strongly related to differences in subsequent growth both at the industry and firm level. These findings make a strong case for innovation as a source of long-run firm growth, consistent with the equilibrium model of Klette and Kortum (2004).

Our paper is connected to several strands of the literature. Our work is closely related to the literature in macroeconomics that aims to measure technological innovation. Broadly, there are three main approaches to identify technology shocks. First, researchers have measured technological change through Solow residuals, after accounting for non-technological effects such as imperfect competition and varying utilization (e.g., Basu, Fernald, and Kimball (2006)). Second, researchers have imposed long-run restrictions on vector auto-regressions (VARs) to identify technology shocks. Both of these approaches measure technology indirectly. The resulting technology series are highly model-dependent, as they depend on the identification assumptions.

Our approach falls into the third category, which constructs direct measures of technological innovation using micro data. Shea (1999) constructs direct measures of technology innovation using patents and R&D spending and finds a weak relationship between TFP and technology shocks. Our contrasting results suggest that this weak link is likely the result of assuming that all patents are of equal value. Indeed, Kortum and Lerner (1998) show that there is wide heterogeneity in the economic value of patents. Furthermore, fluctuations in the number of patents granted are often the result of changes in patent regulation, or the quantity of resources available to the US patent office (see e.g. Griliches (1990) and Hall and Ziedonis (2001)). As a result, a larger number of patents does not necessarily imply greater technological innovation. Using R&D spending to measure innovation overcomes some of these issues, but doing so measures innovation indirectly. The link between inputs and output may vary as the efficiency of the research sector varies over time or due to other economic forces.[6] The measure proposed by Alexopoulos (2011) based on books published in the field of technology overcomes many of these shortcomings. However, this measure is only available at the aggregate level, and does not directly capture the economic value of innovation. In contrast, our measure is available at the firm level, which allows us to evaluate reallocation and growth dynamics across firms and sectors.

Our paper is not the first to link firm patenting activity and stock market value (Pakes, 1985; Austin, 1993; Hall et al., 2005; Nicholas, 2008). In particular, Pakes (1985) examines the relation between patents and the stock market rate of return in a sample of 120 firms during 1968–1975. His estimates imply that, on average, an unexpected arrival of one patent is associated with an increase in the firm's market value of $810,000. The ultimate objective of these papers is to measure the economic value of patents; in contrast, we use the stock market reaction as a means to an end—to construct appropriate weights for an innovation measure which we employ to study reallocation and growth dynamics. Our paper is also related to work that examines whether technological innovation leads to positive knowledge spillovers or business stealing. Closest to our paper is the work of Bloom, Schankerman, and Reenen (2010), who disentangle the externalities generated by R&D expenditures on firms competing in the product and technology space. We contribute to this literature by proposing a measure of patent quality based on stock market reaction and assessing within- as well as between-industry reallocation and growth dynamics after bursts of innovative activity.

Our work is also related to literature on endogenous growth and creative destruction (see Acemoglu (2009) for a textbook treatment). Closest to our work are the papers that explore the impact of innovation on firm productivity and growth (Caballero and Jaffe, 1993; Akcigit and Kerr, 2010; Acemoglu, Akcigit, Bloom, and William, 2011). Finally, our paper is related to work that explores the micro-foundations of aggregate economic shocks. In particular, Gabaix (2011) proposes that if the distribution of firm size is sufficiently fat-tailed, as is the case in the US and in most of the world, firm-specific shocks can have substantial effects on aggregate quantities due to the failure of the law of large numbers. Consistent with this view, the empirical distribution of firm-level innovation measure is fat-tailed, suggesting that the innovative activity of a few large firms can have a large aggregate impact. However, we find evidence of comovement of our innovation measures across firms, suggesting that common shocks play an important role as well.

The remainder of the paper is organized as follows. In Section 2 we describe the constriction of our innovation measure. Section 3 studies the response of individual firms and industries on our innovation measure and documents patterns of reallocation. Section 4 explores the response of aggregate variables on our innovation measure. Section 5 discusses the connection of our findings with existing models and concludes.

[6] Kortum (1993) documents that the patent-to-R&D ratio has shown a secular decline in the US.

2 Measuring innovation

In this section we explain how we construct our firm, industry and aggregate level measures of innovation. In our analysis, we use patent data from Google Patents and CRSP. See the Online Appendix (Sections A-C) for details.

Our innovation in this paper is to identify the value of a patent from the stock price reaction around the days that the market learns that a firm has applied for a specific patent, or that a patent has been granted to the firm. In order to examine stock market reactions, we need to define what constitutes an information event. Prior to 2000, patent application filings were not publicized (see e.g. Austin (1993)). In contrast, information does become widely available when patents are granted. The USPTO's publication, *Official Gazette*, which is published every Tuesday, lists patents that are granted that day and reports details of the patent. Subsequent to the American Inventors Protection Act of 1999, the USPTO also began publishing applications 18 months after filing even if the patents had not yet been granted. Publication of these applications occurs on Thursday of each week. When application publication dates are available, we combine the stock market reaction around both information events to construct our innovation measure.

When constructing our innovation measure, we only use information on patents by publicly-traded firms. Hence, one worry is that we do not include private companies, several of which might be responsible for large and more important technology shocks.[7] This omission is likely to bias our findings toward zero. The magnitude of any bias, however, is likely to be small. First, Bloom et al. (2010) show that public firms in Compustat account for most of the R&D expenditures in the United States. Second, Baumol (2002) notes that while several independent and private firms might provide initial innovation, large publicly traded firms conduct most of the refinements that lead to large improvements in welfare.

We should stress that while our method identifies the value of a patent, relying on stock market reaction suffers from two limitations. First, market participants may have advance knowledge of the patent, either through information leakages, or because the firm has chosen to make its patent application public. If so, the stock market reaction on the patent grant day or publication date would underestimate the economic value of the patent. Second, our method only allows us to measure the private value of the patent. In contrast, the social value of a patent can be higher, or lower, depending on whether the patent generates research spillovers or steals business from existing firms. Notably, the challenge of accurately measuring the private and social value of an innovation is not unique to our paper, but confronts other measures, such as R&D or patent citation counts, as well.

2.1 Extracting patent value from stock price reaction

We extract information about the value of each patent from stock price reactions using two methods: a simple measure that ignores measurement error, and a more sophisticated measure that incorporates the error into the estimation procedure.

[7]Kortum and Lerner (2000) find that venture capital, which accounts for 3% of total R%D expenditures, is responsible of 15% of industrial innovations.

A simple measure

To isolate market movements we focus on the firm's idiosyncratic return, r_{ft}, defined as the firm's return minus the return on the market portfolio. By using this 'market-adjusted-return model' (Campbell, Lo, and MacKinlay, 1997), we avoid the need to estimate the firm's stock market beta, therefore removing one source of measurement error. As a robustness check, we construct the idiosyncratic return as the firm's stock return minus the return on the beta-matched portfolio (CRSP: bxret). This has the advantage that it relaxes the assumption that all firms have the same amount of systematic risk, but is only available for a smaller sample of firms. Unless noted otherwise, our results are quantitatively similar when using this alternative definition.

Given our measure of idiosyncratic firm return, we construct the idiosyncratic stock price reaction as the firm's idiosyncratic return during the announcement window, r^l_{jd}, times the market capitalization of the firm, S_{jd-1}, on the previous day:

$$A_j = r^l_{jd}\, S_{jd-1}, \qquad r^l_{fd} \equiv \sum_{t=0}^{l} r_{jd,d+t}. \tag{1}$$

The next step is to choose the length of the announcement window, l. As we show below, trading volume is higher on the two days following a patent being granted, suggesting that the stock price movements in days after the announcement are also informative. The downside is that increasing the announcement window can potentially add noise to our estimates. In the baseline case, we choose a three-day window ($l = 2$). As a robustness test, we extend the window to five days ($l = 4$).

The private value of a patent is generally nonnegative because a firm can always choose not to implement it. Therefore, when we construct our innovation measure, we restrict attention only to positive stock price responses:

$$A^+_j = \max[A_j, 0]. \tag{2}$$

Our first measure of innovation A^+_j is easy to construct, since it involves no estimation of parameters. The downside, however, is that it ignores the possibility of measurement errors. In particular, by truncating returns at zero we are introducing an upward bias in our estimate of the dollar value of innovation. The magnitude of this bias is increasing in the volatility of the firm's idiosyncratic return. To ensure that the variation in our measure A^+_j does not result from variation in the firm volatility, we control for idiosyncratic volatility σ_{ft} throughout.

Adjusting for measurement error

We construct an alternative measure of innovation to explicitly account for measurement error introduced while constructing the simple measure.[8] In other words, we account for that fact that the stock price of innovating firms may fluctuate for reasons unrelated to innovation during the announcement window. The idiosyncratic stock return during the announcement day window can be decomposed as:

$$r^l_{jd} = x_j + e_{jdl}, \tag{3}$$

[8]We are grateful to John Cochrane for this suggestion.

where x denotes the value of the patent (as percentage of market value) and e_{dl} the component of firm stock return that is unrelated to the patent. Under the assumption that $e_{jdl} \sim \mathcal{N}(0, \xi_j)$ and x is distributed according to a Gaussian $\mathcal{N}(0, v_j)$ truncated at zero, we can construct the conditional expectation of the value of the patent as a function of the firm's stock return:

$$E[x_j | r^l_{jd}] = \delta_j \, r^l_{jd} + \sqrt{\delta_j \, \xi_j} \, \frac{\phi(R_j)}{1 - \Phi(R_j)}, \tag{4}$$

where ϕ and Φ are the standard normal pdf and cdf, respectively, and

$$R_j = -\sqrt{\delta_j} \, \frac{r^l_{jd}}{\sqrt{\xi_j}}, \qquad \delta_j = \frac{v_j}{v_j + \xi_j}. \tag{5}$$

To implement our procedure, we need estimates of v_j and ξ_j, preferably at the firm level. To reduce the number of parameters, we assume that $\delta_j = \delta$, that is, the signal-to-noise ratio is constant across firms and time. To estimate δ, we regress log squared returns on a patent announcement-day dummy variable, I_{fd},

$$\ln \left(r^l_{fd} \right)^2 = a_0 + a_{ft} + b_d + \gamma I_{fd} + u_{fd}, \tag{6}$$

controlling for firm-year (a_{ft}) and day-of-week (b_d) fixed effects. The signal-to-noise estimate is then:

$$\hat{\delta} = 1 - \frac{var(r^l_{fd} | I_{fd} = 0)}{var(r^l_{fd} | I_{fd} = 1)} = 1 - e^{-\hat{\gamma}}. \tag{7}$$

We estimate (6) using a three-day ($l = 2$) and a five-day ($l = 4$) return announcement window. We obtain estimates $\hat{\delta} \approx 0.031$ in both cases, so we use this as our benchmark value.

Next, we estimate the measurement error ξ_j. There is strong evidence that firm-volatility varies both in the time-series and the cross-section, hence it is important to allow ξ_j to vary both across firms and time. For every firm f and year t we estimate its idiosyncratic variance, σ^2_{ft}, from daily returns. This variance is estimated over both announcement and non-announcement days, so it is a mongrel of both v and ξ. Given the estimate of the daily variance σ^2, the fraction of trading days that are announcement days, μ, and our estimate for the signal-to-noise ratio, $\hat{\delta}$, we recover the measurement error by $\xi_{ft} = \sigma^2_{ft} \, (1 + l) \left(1 + \mu_{ft} (1 + l) \frac{\hat{\delta}}{1 - \hat{\delta}} \right)^{-1}$.

We then construct our second innovation measure as:

$$\hat{A}_j = E[x_j | r^l_{jd}] \times S_{jd-1}. \tag{8}$$

Our second innovation measure (8) explicitly accounts for the fact that a firm's idiosyncratic return may contain information unrelated to the value of a patent. The conditional value of a patent in equation (4) is an increasing and convex function of the daily firm return, and thus has a similar shape as our simple innovation measure (2), up to a scale parameter that depends on the signal to noise ratio.

2.2 Information in stock-price responses

We now provide evidence that the stock market reaction contains valuable information about the value of a patent. First, we document that trading volume increases around the days that patents are granted (or their applications are published).[9] Second, we document that the stock market reaction of a patent is correlated with an independent measure of its *realized* value–the number of future citations the patent receives.

Trading volume

We regress a firm's share turnover x (trading volume divided by shares outstanding) on an announcement day dummy variable I_{fd},

$$x_{fd+k} = a_0 + a_{ft} + b_d + b(k)\, I_{fd} + u_{fd}, \tag{9}$$

controlling for firm-year ζ_{ft} and day-of-week b_d fixed effects. We vary k from -1 to 5. We find that there is a statistically significant increase in share turnover around the day that the firm is granted a patent or its application is publicized. Volume increases on the day of the announcement, and remains temporarily higher for the next two days. We find that the total turnover in the first three days after the announcement increases by 0.21-0.40%. Given that the daily median turnover rate is 1.29%, this is an economically significant increase in trading volume, consistent with the view that patent issuance conveys important information to the market. See the Online Appendix (Table 3) for the full set of results.

Patent citations

The next step is to explore whether the stock price reaction around the day of the announcement carries information about the likelihood of the patent receiving citations in the future. We look at patent citations because they represent an independent measure of the *realized* value of a patent.

We examine whether the firm's stock price reaction when granted patent j is correlated with the number of future citations, C_j, the patent subsequently receives:

$$C_j = a + b\, A_j + y_j + \gamma\, \log \sigma_j + e_j. \tag{10}$$

We include grant-year (or publication-year) fixed effects (y_j) in the regression because older patents have had more time to accumulate citations. We include the firm's idiosyncratic volatility σ to control for the truncation-induced bias in our simple measure A^+. We consider both three-day ($l = 2$) and five-day ($l = 4$) announcement day windows. We cluster the standard errors by year.

We show the estimation results in Table 1. Our truncated measure A^+ is informative about the number of future citations. As we see in Panel A, the coefficient of patent citations on our innovation measure is statistically significant across patent length windows. The economic magnitudes are moderately significant. The median number of citations a patent receives is 5. An increase from

[9] Though prices can adjust to new information absent any trading, the fact that stock turnover increases following a patent grant or publication is consistent with the view that some information is released to the market, and not all agents share the same beliefs.

the median to the 90th percentile in terms of stock price reaction around the day the patent is granted (approximately 24 million 1982 US dollars) is associated with 0.27 more citations. The corresponding numbers for the patent publication date is 31 million US dollars and 0.1 citations respectively. The previous numbers correspond to the three-day (l=2) window, but the results using the five-day (l=4) window are quantitatively similar. The magnitudes are substantially larger when we use our measure adjusted for measurement error, as we see in Panel B. An increase from the median to the 90th percentile in terms of our innovation measure $\hat{A}$, corresponding to approximately 13 to 27 million 1982 US dollars depending on the window, is associated with 1.7 to 1.8 (or 0.4 to 0.5) more citations, measured around the patent grant (or publication) day.[10]

Next, we repeat the exercise, replacing A_j with its logarithm, $\ln A_j$. This serves two purposes. First, it ameliorates the effects of outliers. Second, for our truncated measure A^+, it explores whether the positive effect on citations comes from the transition from zero to positive, or it also exist if we focus on the positive responses alone. As we see in Panels C and D, the semilog specification yields estimates that are economically more significant. An increase in A^+ and $\hat{A}$ from the median to the 90th percentile is associated with 1.8-2.1 (0.7) more patent citations, using data on patent grant (publication) day.

In addition, we perform a number of robustness tests. First, our findings are quantitatively similar if we estimate equation (10) with a Poisson or negative binomial regression. Second, the results using our second idiosyncratic return measure (the firm's return minus the beta-matched portfolio) are similar, though one-third smaller in magnitude. Third, we explore what happens if we do not truncate the idiosyncratic dollar return A_j. We find that the simple non-truncated dollar return A_j is essentially uncorrelated with future citations.

The results of this section suggest that the stock price reaction within a few days after the patent is granted contains important information about the value of the patent. We use this information to weigh the number of patents when we construct measures of innovation at the firm, industry or aggregate level. Since the point estimates are a bit higher when we use a three-day versus a five-day window, we focus on the former throughout the paper. Finally, the stock price reaction around both the grant as well as the publication date appear to be informative. Thus, in what follows, we measure the value of each patent as the sum of the values obtained using the grant-day and publication-day windows.

2.3 Some illustrative case-studies

Before turning to our main results, we provide some illustrative case studies to highlight the success of our method in identifying valuable patents. For these examples we performed an extensive search of online and print news sources to confirm that no other news events could account for the return around the patent dates.

The first example is patent 4,946,778, titled "Single Polypeptide Chain Binding Molecules", which was granted to Genex Corporation on August 7, 1990. As shown in Panel A of Figure 1, the stock price increased 67% (in excess of market returns) in the three days following the patent

[10] Note that small changes in citations generated by a patent (around the median number) can be associated with large value implications for the firm producing the patent (Hall et al. (2005)).

announcement. Investors clearly believed the patent was valuable, and news of the patent was reported in the media. For example, on August 8 *Business Wire* quoted the biotechnology head of a Washington-based patent law firm as saying "The claims issued to Genex will dominate the whole industry. Companies wishing to make, use or sell genetically engineered SCA proteins will have to negotiate with Genex for the rights to do so."

The patent has subsequently proved to be important on other dimensions as well. The research that developed the patent, Bird, Hardman, Jacobson, Johnson, Kaufman, Lee, Lee, Pope, Riordan, and Whitlow (1988), was published in *Science* and has since been cited over 1300 times,[11] while the patent itself has been subsequently cited by 775 patents. Genex was acquired in 1991 by another biotechnology firm, Enzon. News reports at the time indicate that the acquisition was made in particular to give Enzon access to Genex's protein technology.

Another example from the biotechnology industry is patent 5,585,089, granted to Protein Design Labs on December 17, 1996. The stock rose 22% in the next two days on especially high trading volume (Panel B of Figure 1). On December 20, the *New York Times* reported that the patent "could affect as much as a fourth of all biotechnology drugs currently in clinical trials."

Finally, consider the case of patent 6,317,722 granted to Amazon.com on November 13, 2001 for the "use of electronic shopping carts to generate personal recommendations". When Amazon filed this patent in September 1998, online commerce was in its infancy. Amazon alone has grown from a market capitalization of approximately $6 billion to over $100 billion today. The importance of a patent that staked out a claim on a key part of encouraging consumers to buy more – the now-pervasive "customers also bought suggestions" – was not missed by investors: The stock rose 34% in the two days after the announcement, adding $900 million in market capitalization (see Panel C of Figure 1).

Other patents associated with large returns include an ink jet technology granted to Canon in 1982 (Panel D of Figure 1), and a digital storage device granted to Sperry Rand in 1959. These examples, and a number of others we carefully investigated indicate that our method of identifying important patents by looking at stock returns appears to work well.

2.4 Construction of innovation measures

We now explain how we use the stock price reaction of innovating firms to construct measures of innovation at the firm, industry and aggregate level. In addition, we discuss various properties of our measures, which together strongly reaffirm that these measures are reasonable indicators of innovative activity.

Firm-level measures of innovation

In most of our analysis, the unit of observation is a year because of the availability of macroeconomic data. Hence, we need to construct measures of innovation at annual frequencies. We do so by summing over the stock price reaction across all patents granted, or its application published, to

[11]Google Scholar citation count.

firm f in year t,

$$A^v_{ft} = \sum_{j \in J^g_{ft}} A_j + \sum_{j \in J^p_{ft}} A_j, \qquad A_j \in [A^+_j, \hat{A}_j], \tag{11}$$

where J^g_{ft} and J^p_{ft} denote the sets of patents granted and published applications, respectively, to firm f in year t.

In our firm-level analysis, we scale the dollar value of innovation by the end-of-year firm market capitalization, S, in year t:

$$A_{ft} = \frac{A^v_{ft}}{S_{ft}}. \tag{12}$$

Hence, our firm-level innovation measures can be interpreted as the fraction of firm f's value that can be attributed to innovation in year t.

As we see in Table 2, the distribution of our firm-level measure is skewed to the right. In addition, as we show in the Online Appendix (Figure 2) the distribution of our firm-level measures of innovation A^+_f and $\hat{A}_f$ has fat tails. Restricting attention to the top 10 percent of the distribution, the relation between the log complementary empirical cdf, $\log(1 - F(A))$, and the log innovation measure, $\log A$ is close to linear, with a slope coefficient of approximately -1.9. Hence, the tail behavior of A can be well approximated by a power law. A simple estimator of a power law exponent (Newman, 2005) yields a point estimate of -2.75.[12] Our findings that the distribution of patent quality is fat-tailed is consistent with Harhoff et al. (1997), who show similar results for the distribution of patent citations.

In addition to some patents being very valuable, our results indicate that a few large firms are very important for the aggregate rate of innovation in the economy. The identity of these firms varies by decade. In the 1930s and 1940s, AT&T and GM are responsible for a large share of innovative activity. In the 1950s and 1960s, du Pont and Kodak take a leading role. In 1970s and 1980s, a large share of innovation takes place in Exxon, GE, 3M and IBM. Finally, in the 1990s and 2000s, "new economy" firms are responsible for a large share of innovation, namely Sun, Oracle, Microsoft, Intel, Cisco, Dell, and Apple.

Next, we explore how our firm-level measures of innovation are related to firm characteristics, in particular Tobin's Q, firm size, K, and R&D spending (normalized by assets):

$$A_{ft} = a_0 + a_1 \log Q_{t-1} + a_2 \log K_{t-1} + a_3 \log RD_{t-1} + \rho\, A_{ft-1} + u_{it}. \tag{13}$$

We estimate equation (13) using the entire sample of Compustat firms from 1950 to 2010 using a Tobit model.[13] We include industry dummies to account for industry-level time invariant characteristics; and time dummies to account for changing state of the business cycle as well as changes in patent law or changes in the efficiency and resources of the USPTO (see e.g. Griliches (1989)) during our sample period. We cluster the errors by firm.

We find that firms that are large, have higher Tobin's Q, and have higher R&D expenditures are more likely to innovate. These findings are similar to those discussed in Baumol (2002), Griliches

[12] This estimator assumes that A is i.i.d. across firms and across time. After removing firm and time fixed effects, the point estimate of the power law exponent is equal to -3.70 and -3.55 for A^+ and $\hat{A}$ respectively.

[13] Note that information on R&D expenditure is reliably reported in Compustat only from 1975 onwards. As a result our sample period for regressions that use R&D stock is restricted to 1975–2010.

(1990), Scherer (1965) and Scherer (1983) on the characteristics of firms that have conducted radical innovation and have been responsible for technical change in the U.S. See the Online Appendix (Table 4) for the full set of results.

Aggregate measures of innovation

We construct industry-level and economy-wide measures of innovation by aggregating our firm-level measures across firms. In particular, we construct dollar measures of innovation, by aggregating our firm-level measures across the set N_t of firms across the entire economy or at the industry level:

$$A_t^v = \sum_{f \in N_t} A_{ft}^v, \qquad A_{ft}^v \in [A_{ft}^{+v}, \hat{A}_{ft}^v]. \tag{14}$$

Our dollar measure A^v will be mechanically affected by economic forces that affect the level of stock prices but are likely to be unrelated to innovation, such as changes in discount rates. Hence, as before, we scale our dollar measure A^v by the total market capitalization

$$A_t = \frac{A_t^v}{S_t}, \tag{15}$$

where $S_t = \sum_{f \in N_t} S_{ft}$. Thus, our aggregate measure, A_t, is a value-weighted average of our firm-level innovation measure, A_{ft}.

We compare our two measures of aggregate innovation with three aggregate measures proposed in the literature: the log number of total patents granted; the log stock of R&D capital from the BEA; and the log number of technology books published from Alexopoulos (2011). Some of these measures show a secular time trend, so we remove a deterministic time-trend from all measures.

We plot these series in Figure 2. Our measures of innovative activity line up well with the three major waves of technological innovation in the U.S. First, our measures suggest high values of technological innovation in the 1930s, consistent with the views expressed in Field (2003). When we dissect our measures we find that firms that primarily contribute to technological developments during the thirties are in the automobiles (such as General Motors) and telecommunication (such as AT&T) sectors. This description fits well with studies that have examined what sectors and firms led to technological developments and progress in the 1930s (Smiley, 1994).

Second, our measures suggest higher innovative activity during 1960s and early 1970s – a period commonly recognized as a period of high innovation in the U.S (see Laitner and Stolyarov (2003)). As has been noted, this was a period that saw development in chemicals, oil and computing/electronics – the same sectors we find to be contributing the most to our measure with major innovators being firms such as IBM, GE, 3M, Exxon, Eastman Kodak, du Pont and Xerox.

Third, developments in computing and telecommunication have brought about the latest wave of technological progress in the 1990s and 2000s, which coincides with the high values of our measure. In particular, it is argued that this is a period when innovations in telecommunications and computer networking spawned a vast computer hardware and software industry and revolutionized the way many industries operate. We find that firms that are main contributors to our measure belong to these sectors with firms such as Sun Microsystems, Oracle, EMC, Dell, Intel, IBM, AT&T, Cisco, Microsoft and Apple being the leaders of the pack. We next turn to providing firm level evidence that lends additional support to validity of our measures.

Comparing our aggregate innovation series, we also note four important points. First, our two aggregate innovation series are very similar to each other, suggesting that the truncation bias in A^+ is diversified across firms. Second, our measure displays different behavior than the total number of patents, especially in the beginning and towards the end of the sample. The correlation between log A^+ ($\hat{A}$) and the log number of patents is equal to 0.36 (0.42) in levels and 0.16 (0.11) in first differences. Third, our two innovation measures capture similar low-frequency movements to R&D spending and the number of technology books published in the Library of Congress, in particular the rise in innovative activity during the 1960s and early 1970s. Finally, our innovation measures displays substantial high-frequency variability relative to either the stock of R&D or the number of technology books. Some of this variability comes from variation in the number of patents granted, but a significant part comes from changes in the average response of the stock market on these patent grant dates. In contrast, the stock of R&D capital and the number of technology books display mostly low-frequency variation.

We end this section by discussing the economic source of time-series variation in our aggregate innovation measure. One possibility, in the spirit of the 'granularity' hypothesis of Gabaix (2011), is that the observed time-series variation in aggregate innovation is the result of disproportionately large idiosyncratic shocks that fail to be diversified away. This view is consistent with our findings above that the right tail of our firm-level innovation measure follows a power law. The alternative hypothesis is that there is an underlying macroeconomic shock that affects the firm-level propensity to innovate and the distribution of patent outcomes.

To shed some light on this, we decompose the aggregate measure of innovation A^v_t into A^{v1}_t, the dollar value of innovation that is contributed by the top 1% of firms; and A^{v99}_t, the value that is contributed by the remaining firms. Indeed, our aggregate innovation measure is dominated by a few large firms. Focusing on the sample of firms with positive innovation A^v_{ft}, the top 1% of firms in terms of innovation account for an average of 32% of the total dollar value of innovation A^v_t. If aggregate innovation is determined by large idiosyncratic shocks, we would expect innovations in A^{v1}_t to be uncorrelated with A^{99}_t. Instead, we find that the sample correlation between $\Delta \ln A^{v1}_t$ and $\Delta \ln A^{v99}_t$ ranges from 75% to 78%, depending on our measure. Thus, the data suggests that a systematic shock affecting all firms is responsible for a large portion of the time-series variation in A^v.

3 Innovation, productivity and reallocation

To maximize the economy's overall level of production, its resources need to be allocated to the most productive firms and industries.[14] Here, we explore this mechanism in more detail on two broad fronts. First, we document the link between innovation and productivity. Second, we show that, consistent with economic optimization, productive resources flow into the innovating firm away

[14] There exists a large literature on the importance of resource allocation for economic growth (see, e.g. Restuccia and Rogerson (2008); Hsieh and Klenow (2009); Jones (2011); Acemoglu et al. (2011)).

from firms that do not innovate. In both of these cases we perform our analysis both within and across industries.

3.1 Firm-level evidence

We begin by exploring the productivity and reallocation dynamics subsequent to innovative activity within an industry. In particular, we examine the response of productivity, Tobin's Q and factor demand to a firm's own innovation activity, A_f, and also to the innovation output of its competitors. We construct our measure of innovation of a firm's competitors, A_{If} as the average innovative activity of all firms in the same industry excluding firm f, weighted by market capitalization S:

$$A_{Ift} = \sum_{h \neq f \in N_{It}} A^v_{ht} \Big/ \sum_{h \neq f \in N_{It}} S_{ht}. \tag{16}$$

We define industries using 3-digit SIC codes.[15] We explore the effect of innovation of a firm and its competitors on various firm outcome variables, x, by estimating the regression

$$x_{ft+1} = a_0 + a_1 \, A_{ft} + a_2 \, A_{Ift} + b \, Z_{ft} + \gamma_t + c_I + \rho \, x_{ft} + u_{ft+1}. \tag{17}$$

We include lags of the dependent variable, industry c_I and year γ_t dummies, and a vector of controls Z. We control for firm idiosyncratic volatility, σ_{ft}, when using our truncated measure A^+ because the magnitude of the truncation bias increases with volatility. We control for firm size, measured as either physical capital or number of employees, because large firms innovate more. In addition, one source of concern is that unobservable variables at the firm or industry level jointly drive innovation outcomes and the outcome variable x. Thus, depending on the specification, we control for firm productivity, profitability, Tobin's Q, and firm and industry stock returns. We present results with and without these controls, and cluster the standard errors by firm.

We are interested in the estimates of a_1 and a_2, which capture the impact of innovation by the firm and its competitors. A firm's innovative output, A_f, is highly skewed so we focus on inter-decile movements in firm-level innovation to explore the economic magnitude of a_1. In addition, the innovation of other firms can have a positive or a negative effect on a firm's outcome variables. An increase in the innovative output of competing firms can have a positive effect on the firm because of knowledge spill-overs. However, innovation of competitors can also have a negative effect due to business stealing or an increase in factor prices. We should note that the presence of unobserved variables that drive the common propensity of firms to innovate are likely to bias our estimate a_2 upwards. For instance, common productivity shocks could impact many firms in the same industry – thereby creating a positive correlation between innovative activity of a firm's competitors and a firm's own productivity.

Productivity

First, we examine whether firms that innovate have higher productivity subsequent to innovative activity. We consider both capital- and labor-productivity (mpk_{ft} and mpl_{ft}), defined as firm output—

[15]We obtain quantitatively similar results when we define industries according to their 4-digit SIC.

total sales plus change in inventories—divided by capital and number of employees, respectively. We evaluate the relation between subsequent productivity of capital and labor and innovation by a firm A_f or its competitors A_I by estimating (17) with $x_{ft} = [\log mpk_{ft}, \log mpl_{ft}]$. Depending on whether we focus on the productivity of capital or labor, we measure firm size as the stock of physical capital or number of employees respectively.

We report the results in Panel A of Table 3. We find a substantial increase in firm-level productivity subsequent to an innovation. Our estimates of a_1 imply that an increase in innovation by the firm from the 50th to the 90th percentile leads to an 0.6% to 1.5% increase in the productivity of capital and a 1.7% to 2.1% increase in the productivity of labor. Furthermore, we find some evidence that the business-stealing effect dominates, as the estimated coefficient a_2 is negative and statistically significant across specifications. In particular, a one-standard deviation increase in the amount of innovation by the firm's competitors is associated with a 1.5% to 1.9% decline in the productivity of capital and a 1.5-1.8% decline in the productivity of labor. Our finding that labor productivity increases following innovation, suggests that during our sample period, innovation is more likely to be labor augmenting than labor saving on average (see, e.g. Acemoglu (2010)).

Our estimates imply that the business-stealing effect is substantial. However, this finding may be an artifact of the short horizon considered in our analysis if the business-stealing effect and positive spillovers operate at different frequencies. In particular, positive spillovers may affect firms with a lag, so in the medium-run, the response of productivity may be different. To explore this possibility, we estimate a dynamic version of equation (17) with k-year ahead productivity, x_{ft+k}, as the regressand. We consider horizons of one to five years $k = [1..5]$. To conserve space, we present results with only size, lagged productivity and volatility controls. Including additional controls leads to quantitatively similar findings.

As we see in Figure 3, the negative effect of competitor innovation, A_{If}, on productivity is stronger in the short run. As we increase the horizon k, the estimated coefficients $a_2(k)$ increase, becoming zero or positive after 5 years. In contrast, the positive effect of firm innovation on productivity increases with the horizon k. After 5 years, the response of productivity of capital is between 37% to 50% higher than on impact. Labor productivity displays a similar, though quantitatively stronger response. The positive effect of firm innovation on labor productivity increases with the horizon by 65 percent.

In summary, our findings are consistent with the view that positive spillovers and business stealing operate at different horizons. In the short run, firms that do not innovate when their competitors do experience a decline in their productivity. However, in the medium run, the innovation of other competitors has either a zero or a positive effect. This positive effect can arise because competitor innovations affect the firm either directly, for instance through knowledge spillovers, or indirectly, by spurring future firm innovation. Last, another possibility, which we explore below, is that the firm scales down operations in response to innovation by competitors and therefore operates at a higher marginal product of capital and labor.

Tobin's Q

Next, we explore the effect of innovation on the market value of the firm. In particular, the firm's Tobin's Q should respond positively to a firm's innovation output. The response of Q to the innovation of the firm's competitors will depend on whether the business-stealing or positive spillover effects dominate in terms of market value. We estimate equation (17) with $x_{ft} = [\log Q_{ft}]$ and present the results in Panel B of Table 3.

We find that Tobin's Q responds positively to a firm's own innovation activity. Our estimates of a_1 imply that an increase in innovation by the firm from the 50th to the 90th percentile leads to an 1.5% to 1.6% increase in the firm's Tobin's Q. These magnitudes are in line with those reported in Hall et al. (2005). In addition, we find some evidence of positive spillovers. A one standard deviation increase in the innovation activity of other firms in the industry is associated with a 0.5% to 0.7% increase in Tobin's Q. However, we must be careful when interpreting this as evidence of positive spillovers, given the fact that the coefficient a_2 is likely to be biased upwards due to the presence of unobservables, as we discuss above.

Reallocation

In this section we explore the reallocation dynamics subsequent to innovation by a firm. In particular, we explore how our innovation measures are related to reallocation of physical capital and labor. We focus on the firm's investment and hiring rate. In addition, since adjusting a firm's capital and labor input often involves upfront costs, we explore the allocation of financial resources. We focus on the net financial inflows to the firm, defined as new issuance of equity and debt minus payouts to stock- and bond-holders.

We estimate equation (17), using firm investment, i, net hiring rate, h, and financial inflows, e, as outcome variables $x_{ft} = [i_{ft}, h_{ft}, e_{ft}]$. As before, our main estimates of interest are a_1 and a_2, which capture the change in factor inputs and financial inflows following innovation by the firm and its competitors, respectively.

We first examine how physical capital gets reallocated subsequent to innovation by a firm or by its competitors. Table 4 shows that subsequent to an innovation by a firm, there is a substantial increase in its investment rate. In particular, our estimates imply that an increase in innovation by the firm from the 50th to the 90th percentile leads to an increase in the firm's investment rate by 0.5% to 1%. This increase is statistically but also economically significant given that the median firm investment rate is 12% in our sample. Furthermore, we find evidence that physical capital flows from firms that do not innovate to firms that do. If the firm does not innovate but its competitors do, then its investment rate is substantially lower. A one-standard deviation increase in the level of innovation by the firm's competitors leads to a decline in the firm's investment rate of 0.6-1.6%.

Next, we examine reallocation of labor subsequent to innovation by a firm. Table 5 shows that subsequent to an innovation by a firm, there is a substantial increase in its employment using either innovation measure. As before, the economic magnitudes are significant. Our estimates imply that an increase in innovation by the firm from the 50th to the 90th percentile leads to increase in employment of the firm by 0.2% to 0.5%, compared to the median firm-level hiring rate of 2.7%. In addition, labor declines when a firm does not innovate but its competitors in the same industry do.

A one-standard deviation increase in the average innovation of the firm's competitors leads to a reduction of 0.7% to 1.4% in the firm's hiring rate. Our finding that innovating firms increase their labor demand suggests that innovation is more likely to be labor augmenting, and is consistent with the findings of Lentz and Mortensen (2008) for Danish firms.

Last, we examine the reallocation of financial capital subsequent to innovation by a firm and present the results in Table 6. Following an innovation by a firm, there is a substantial increase in its financial capital inflows. Our estimates imply that an increase in innovation by a firm from the 50th to the 90th percentile leads to an increase of capital inflows to book assets of 0.5% to 0.9%, compared to the median level of zero capital flows. We also find that a firm is more likely to increase payout and decrease new issuance when it does not innovate but its competitors do. In particular, a one-standard deviation increase in the average innovation of the firm's competitors leads to a reduction of 0.8% to 1.6% in net financial capital flows to the firm. This negative effect suggests that firms that fail to innovate in an industry where other firms do, have fewer investment opportunities and thus increase their net payout to investors.

In summary, our results in this section suggest that, consistent with economic optimization, resources are reallocated to innovating firms and away from firms that fail to innovate when their competitors do. In addition, we find that relative to their median value, new hiring exhibits a quantitatively stronger response than capital, both in terms of inflow and outflow. This increased reallocation response of labor relative to firm capital within industries is consistent with the view that capital is more firm-specific than labor.

3.2 Industry-level evidence

So far we have focused on the dynamics of productivity and reallocation within an industry. We now conduct a similar exercise examining the response of productivity and reallocation of inputs at the sector level. To do so, we use the KLEMS industry-level output data provided by Dale Jorgenson. First, we document the dynamic response of capital and labor productivity, defined as the ratio of the quantity of output to the quantity of capital and labor services, respectively. Second, we focus on the reallocation of inputs, namely the growth rate in the quantity of capital and labor services. Last, we explore the rate of establishment exit to our innovation measures, using information on establishment exit rates at the industry level from the US Census tables on Business Dynamics Statistics (BDS).

We estimate specifications similar to (17), but at the industry level:

$$x_{It+1} = a_0 + a_1\, A_{It} + a_2\, A_{MIt} + b\, Z_t + \gamma_t + \rho\, x_{It} + u_{It+1}. \tag{18}$$

Here, A_I is our measure of innovation at the firm-level, and A_{MI} is the average level of innovation in the economy, excluding industry I constructed in a manner similar to (16). Depending on the specification, we include a vector of controls, Z, which includes stock return and, in the case of our truncated measure A^+, volatility, as well as lagged values of the dependent variable and time effects γ_t. In the presence of time dummies γ_t, the interpretation of the coefficient a_2 is unclear, so we only include one of the two. We cluster the standard errors by industry.

Productivity

First, we explore the dynamic response of industry productivity to its own innovation A_I and the innovation of the other industries A_{MI}. We are interested in the coefficients a_1 and a_2, which measure the response of productivity to an industry and economy-wide (excluding the given industry) innovation shock respectively. The coefficient a_1 is informative as to whether innovation creates net value or is a zero-sum game that merely affects the distribution of rents within an industry. We estimate (18) with k-period ahead productivity as the regressand, $x_{ft+1} = [\log mpk_{ft+k}, \log mpl_{ft+k}]$. We consider horizons of one to five years $k = [1..5]$. We plot the results in Figure 4, controlling for lagged productivity and volatility. Controlling for firm- or industry-level stock returns leads to similar results.

We find that both labor and capital productivity increase in response to own industry innovation. A one-standard deviation A_I shock is associated with a 2.5% increase in the productivity of capital and labor, after a period of 5 years. By contrast, capital and labor productivity show no statistically significant response to the innovation activity of other industries.

Reallocation and Creative Destruction

Next, we examine the response of capital and labor to an industry innovation shock, as well as to the innovation of other industries. We estimate equation (18), using as the outcome variable the growth rate in the quantity of capital and labor services $x_{It} = [i_{It}, h_{It}]$. As before, the main estimates of interest in this specification are a_1 and a_2, which capture the change in the quantity of factor inputs in response to innovation in the industry and the rest of the economy respectively. We show our results in Table 7.

We find that an increase in the amount of industry innovation increases the quantity of capital and labor services in the industry, though in some specifications the effect is not statistically different from zero. As before, we find that the response of labor is greater than the response of capital. An increase in industry innovation is associated with a 0.2% to 0.4% increase in capital services and a 0.3% to 0.7% increase in labor services. These magnitudes are economically significant, given that the median growth in capital and labor services equals 3.1% and 0.7% respectively.

Our results suggest that increases in economy-wide innovation lead to cross-industry reallocation of labor and capital. In particular, a one standard deviation increase in the economy-wide innovation measure is associated with a 0.5% to 1.0% decline in the growth of capital services and a 1.4% to 2.2% decline in the growth of labor services.

We also examine patterns of firm exit at the industry level. If industry innovation spurs creative destruction, we expect to find a positive relation between the rate of firm exit and the level of industry innovation. In contrast, innovation in other industries should not affect the decision of firms to exit a particular industry. We estimate specifications similar to (18), but we replace the outcome variable with the rate of firm exit and examine the response of this variable to own industry innovation A_I and innovation of other industries A_{IM}. Table 8 presents the results.

Industry innovation is accompanied by an increase in creative destruction. The estimated coefficient a_1 is positive and statistically significant across specifications. Innovation accounts for an economically significant fraction of the variation in firm exit rates. A one standard deviation increase in industry innovation is associated with an increase in the firm exit rate by 0.2% to 0.4%, while the unconditional volatility of exit rates is equal to 2.1%. In contrast, economy-wide innovation A_{MI} has no statistically significant effect on firm exit.

3.3 Innovation and long-run growth

The results of the previous two sections imply that innovation is followed by increased productivity of capital and labor, as well as reallocation of resources towards innovating firms. These findings suggest that own innovation should be followed by increased output growth, both at the firm as well as at the industry level. The response of output to innovation by other firms or industries is more ambiguous. It depends on whether productivity increases in the long run, as well as whether the patterns of reallocation we document are reversed in the long run. To answer these questions, we estimate specifications at firm and industry level similar to (17) and (18)

$$\log y_{ft+k} = a_0 + a_1\, A_{ft} + a_2\, A_{Ift} + b\, Z_{ft} + \rho_1 \log y_{ft} + \rho_2 \log y_{ft-1} + e_{t+k}, \tag{19}$$

$$\log y_{It+k} = a_0 + a_1\, A_{It} + a_2\, A_{MIt} + b\, Z_{It} + \rho_1 \log y_{It} + \rho_2 \log y_{It-1} + e_{t+k}. \tag{20}$$

We again examine horizons of $k = 1$ to $k = 5$ years. We control for stock return volatility (in the case of our truncated measure A^+). In the firm-level regression (19), we also include controls for firm size, and industry and time fixed effects. We cluster the standard errors by firm or industry, respectively. We plot the estimated coefficients $a_1(k)$ and $a_2(k)$ in Figure 5, along with 90% confidence intervals.

We find that both firm and industry output displays a statistically significant response to an own-innovation shock. A firm that experiences an innovation shock from the median to the 90th percentile experiences a 1.5% increase in output over a period of 5 years. The response of output is quantitatively more significant at the industry level. A one standard deviation shock to industry innovation is associated with a 5.0% output growth over a period of 5 years. Furthermore, a positive innovation by other firms, or industries, is associated with a decline in output. Output falls by 2.5% to 3.5% at the firm level, and by 3.5% to 6.8% at the industry level following innovation by other firms or industries respectively. These findings suggest that part of the long-run increase in the average productivity of capital and labor (see Figure 3) in response to competitor innovation A_{If} may be the result of the firm scaling down operations.

In summary, innovation is associated with substantial subsequent increases in output. The results of this section can also be summarized by examining the relation between industry innovation in the first half of the sample (1960–1982) and subsequent output growth in the second half of the sample (1983–2006). In Figure 6 we plot the industry innovation measure A_I averaged over the first half of the sample (1960–1982) on the X axis and the corresponding output growth of the industry in the second half of the sample (1983–2006) on the Y axis. The correlation between the two series is 41% with a robust t-statistic of 2.6. Industries which experienced high technological innovation in the first half of the sample were also the ones whose growth rate was subsequently higher in the second half of the sample. For example, industries such as Electrical Machinery, Automotive and Communication, which are in the highest quartile of innovation during the first half of the sample, had an annualized growth rate of more than 4% over the second part of the sample. Similar correlation is found for low-innovative industries such as Textile and Utilities.[16]

4 Innovation and Aggregate Dynamics

Our results in the previous section suggest that innovation is an important determinant of industry-level productivity and growth, especially in the medium term. In this section, we analyze the effect of innovation at the level of the U.S. economy.

4.1 Total Factor Productivity and Output

Impulse Responses

In this section, we examine the extent to which our innovation measures account for medium-run fluctuations in aggregate productivity and output. We start by exploring the relation between measures of innovation and quantities of interest using VARs and VECMs. Then, we explore whether our results are sensitive to the details of the specification or the construction of our innovation measures. We focus on aggregate productivity and output, with productivity measured using utilization-adjusted TFP from Basu et al. (2006) and output measured as the real per capita gross domestic product. Our aggregate innovation measures A^+ and $\hat{A}$ are constructed according to (14) and (15).

We estimate bivariate VARs of the form $Z = [\log X, \log A]'$, where X is our variable of interest and A is our measure of innovation. We include a deterministic trend, following Alexopoulos (2011). When exploring the responses to our truncated innovation measure A^+, we also include the cross-sectional average of idiosyncratic volatility $\bar{\sigma}$ to ensure that our innovation measure does not pick up movements in firm-level volatility. In addition, we also compute responses using a vector-error-correction model (VECM). We select the number of cointegrating relations using the Johansen test, which suggests the presence of one cointegrating relation in all systems. We select the number of lags using the Akaike-Information Criterion, which advocates a lag length of one to two years for each of the systems. We compute standard errors by a bootstrap simulation of 500 samples. We plot the impulse-response functions in Figures 7 and 8, along with 90% confidence intervals. We compute impulse responses by ordering the innovation shock A last, so the technology shock affects the variables of interest only with a lag.

We find that TFP increases by 0.8% to 2% over 8 years following a one-standard deviation increase in innovation output. The forecast error variance attributed to our innovation measures ranges from 17% to 70% at the 8-year horizon, depending on the specification. Our findings are comparable to the results in Alexopoulos (2011), but in contrast to Shea (1999) who uses only information on patents and finds a negative relation.

[16] One source of concern with our analysis could be that the relation between innovation and output growth is driven by omitted variables. To alleviate these concerns we generate exogenous changes in R&D activity across industries by employing the Bloom et al. (2010) instrument for firm-level R&D activity. As discussed in Bloom et al. (2010), the firm-level tax price of R&D can be decomposed into a component that is relatively exogenous since it is based solely on federal rules. In unreported tests we use the Bloom et al. (2010) firm-level R&D instrument and construct its industry counterpart by taking the average of this tax price across firms in a given industry. We find qualitatively similar results to those reported in the table when we instrument the endogenous innovation variables (A).

Aggregate output displays a mild U-shaped response. In the first two years, the response of output to a one-standard deviation shock is negative at 0.5% to 0.8% and statistically significant. However, output increases in the long run by a substantial amount: a one-standard deviation innovation shock results in a net 1.5% to 4% increase in aggregate output after 8 years. The share of 8-year forecast-error variance attributed to our innovation measures ranges from 7% to 16%.

Last, we explore whether our measure of innovation contains incremental information to stock prices. Following Beaudry and Portier (2006), we include the level of the stock market in our VAR, scaled by the consumption deflator and population.[17] This helps us evaluate the extent to which our results are driven by variation in the denominator of A (the stock market capitalization). We order the level of stock prices second, so now $Z_t = [\log(X_t), \log(M_t), \log \sigma_t, \log(A_t)]'$. Our results are qualitatively similar in terms of statistical significance, but the economic magnitudes are smaller. Productivity and output increase by 0.8 and 2.6% respectively at the peak following a one-standard deviation innovation shock. In contrast, the response of output to a one-standard deviation shock in $\log M$ is not statistically significant beyond the one-year horizon. Furthermore, the innovation shock accounts for a comparable fraction of the variance of productivity (2.7-18.9%) and output (7.0-15.7%) relative to the stock market shock (7.9-8.3% and 4.7-5.1%, respectively). Hence, our innovation shock contains incremental information about future productivity and output to the level of the stock market. We report the full set of results in the Online Appendix (Figures 3 and 4).

Comparison to Patents or R&D capital

We explore whether our measure of innovation contributes information relative to other commonly employed measures of technological innovation: the stock of R&D capital, and the log number of patents. We estimate bivariate VARs for productivity and output, with the log number of patents or R&D capital series ordered last. The number of patents has some ability to predict TFP, but the results are quantitatively weaker. A one-standard deviation shock to the log number of patents is associated with a 0.4% increase in TFP, and the patent shock accounts for 13.1% of the forecast error variance. Output shows no statistically significant response to number of patents. In contrast, R&D Capital has some ability to predict output, but not productivity. Output drops in the short run by 0.4%. At the eight-year horizon output displays a statistically significant increase of 0.3%. See the Online Appendix (Figure 5) for the full set of results.

Granger causality

If our measure of innovation indeed represents a fundamental shock, then it should not be predictable by output or productivity. In addition, we explore whether our measure of innovation is predictable by other measures of technological growth in the literature, for instance the book-based measures of Alexopoulos (2011) and the stock of R&D capital.

The top panel of Table 9 shows that output and TFP do not Granger-cause either of our measures of innovation. In addition, the middle panel of Table 9 shows that our measures of innovation are distinct from the measures of Alexopoulos (2011), in that neither causes the other. Our measure is

[17] We depart from Beaudry and Portier (2006) in that we include the level of the CRSP value-weighted rather the level of the S&P 500 index, since the former includes all stocks traded on the three major exchanges.

somewhat correlated with the number of patents at 32-43%, but not with R&D spending (less than 10%). The bottom panel of Table 9 shows that our measure is not Granger-caused by either the number of patents or R&D spending.

4.2 Consumption

Next, we analyze the impulse response of aggregate consumption to our innovation measures. The response of consumption is informative about whether our innovation measure is an example of an embodied or disembodied shock. If technological innovation is a free factor of production, in that it affects all firms in the manner of a disembodied shock, we expect that consumption should increase immediately. Agents anticipating an increase in future consumption would like to increase their consumption today. In contrast, if innovation is not free, because for instance it is embodied in new vintages of capital or due to adoption costs, then consumption may only increase in the long run. In the short run, agents will divert resources away from consumption towards adopting new innovations.

We analyze the response of real per capita consumption of non-durables and services using VARs and VECMs, as in Section 4.1. We plot the impulse-response functions in Figure 9, along with 90% confidence intervals. We find that consumption displays a U-shaped response to innovation. In the first two years consumption displays a statistically significant drop of 0.5% to 0.7%. Subsequently, consumption increases, leading to a 0.2% to 0.5% net increase after 8 years. However, the increase in consumption is not consistently statistically significant across specifications. The innovation shock accounts for 6% to 8% of the forecast-error variance of consumption growth after 8 years.

The short-run decline in consumption is consistent with the delayed response of output in Section 4.1. Innovation affects output with a lag, so the positive response of consumption is necessarily delayed. However, our finding that consumption declines in the short run, whereas output does not, suggests presence of significant adoption costs.

4.3 Innovation and Tobin's Q

We conclude our analysis by examining the relationship between our measures of innovation and Tobin's Q at the aggregate level. The theoretical relation between innovation and Tobin's Q is ambiguous. If innovation represents an increase in TFP that costlessly affects all firms, then standard models will imply that average Tobin's Q should rise (see, e.g. Hayashi (1982)). However, it is also possible that innovation renders part of the capital stock obsolete (see e.g. Laitner and Stolyarov (2003)) or a reduction in profits for incumbent firms (e.g. Greenwood and Jovanovic (1999); Hobijn and Jovanovic (2001); Garleanu, Kogan, and Panageas (2012)). In these cases the relationship between innovation and average Tobin's Q is less clear.

We estimate the contemporaneous response of Tobin's Q to our innovation measures

$$\Delta \log Q_t = a + b\, \Delta \log A_t + c\, Z_t + e_t, \tag{21}$$

where the vector of controls includes lagged values of Q, our innovation measure A and in the case of our truncated measure A^+ changes in the cross-sectional average of idiosyncratic volatility, $\Delta \log \sigma$.

We show the results in Panel A of Table 10. We find that our innovation measure is negatively correlated with average Q. This negative correlation is statistically and economically significant. A one standard deviation increase in innovation is associated with a 8.1% to 12% contemporaneous drop in aggregate Tobin's Q. Our findings echo the stylized facts reported in Greenwood and Jovanovic (1999), Hobijn and Jovanovic (2001) and Laitner and Stolyarov (2003), who argue that Tobin's Q was too low in the 1960s and 1970s, despite the technological advances taking place.

One source of concern is that our aggregate innovation measure A may be mechanically negatively related to Q due to our choice of scaling by the market capitalization of all firms S. As a robustness test, we scale our aggregate innovation measure by the market capitalization *of innovating firms*, $S_{It} = \sum_{f \in N_t} S_{ft} \times 1_{A_{ft}>0}$. Thus, this alternative normalization ameliorates somewhat the concern that this finding is mechanical. However, as we show in Panel B of Table 10, we obtain similar results using this alternative normalization.

5 Conclusion

We explore the role of technological innovation as a source of economic growth by constructing direct measures of innovation at the firm level. We combine patent data for US firms from 1926 to 2010 with the stock market response to news about patents to identify the economic importance of each innovation. Our measures allow us to uniquely identify the reallocation and growth dynamics within- and across industry after bursts of innovative activity.

We document a strong link between innovation and productivity at the firm and industry level. Our evidence suggests that innovation is accompanied by "creative destruction" in the form of resource reallocation, both within and between sectors. Resources flow to innovating firms and sectors, away from firms and sectors that do not innovate. There are stronger patterns of reallocation for labor than for capital, consistent with the view that capital is more specific than labor (Ramey and Shapiro, 2001).

Technological innovation has a significant impact on aggregate variables in the medium run. Our innovation measure is strongly related to aggregate movements in TFP. In addition, aggregate output shows a delayed positive response, consistent with the presence of short-term adoption costs. A positive shock to innovation has a U-shaped effect on consumption growth: consumption is lower in the short run but increases in the long run. This is consistent with a reallocation of resources away from consumption in the short run towards the implementation of innovation. Finally, we find that an increase in innovative activity leads to a fall in aggregate Tobin's Q, consistent with the models of Greenwood and Jovanovic (1999), Laitner and Stolyarov (2003), and (Garleanu et al., 2012).

Our empirical findings link medium-run macroeconomic fluctuations to a direct measure technological innovation, consistent with the idea of medium-term cycles of Comin and Gertler (2006). Furthermore, our findings make a strong case for innovation as a source of long-run firm growth, consistent with the equilibrium model of Klette and Kortum (2004).

References

Acemoglu, D. (2009). *Introduction to modern economic growth.* Princeton University Press.

Acemoglu, D. (2010). When does labor scarcity encourage innovation? *Journal of Political Economy 118*(6), 1037–1078.

Acemoglu, D., U. Akcigit, N. Bloom, and K. William (2011). Innovation, reallocation and growth. Working paper, Stanford University.

Akcigit, U. and W. R. Kerr (2010). Growth through heterogeneous innovations. Technical report.

Alexopoulos, M. (2011). Read all about it!! What happens following a technology shock? *American Economic Review 101*(4), 1144–79.

Austin, D. H. (1993). An event-study approach to measuring innovative output: The case of biotechnology. *American Economic Review 83*(2), 253–58.

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review 96*(5), 1418–1448.

Baumol, W. (2002). *The free-market innovation machine: analyzing the growth miracle of capitalism.* Princeton University Press.

Beaudry, P. and F. Portier (2006). Stock prices, news, and economic fluctuations. *American Economic Review 96*(4), 1293–1307.

Bird, R., K. Hardman, J. Jacobson, S. Johnson, B. Kaufman, S. Lee, T. Lee, S. Pope, G. Riordan, and M. Whitlow (1988). Single-chain antigen-binding proteins. *Science 242*(4877), 423–426.

Bloom, N., M. Schankerman, and J. V. Reenen (2010). Identifying technology spillovers and product market rivalry. Working paper, Stanford University.

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, Volume 8.* National Bureau of Economic Research, Inc.

Campbell, J., A. Lo, and A. MacKinlay (1997). *The econometrics of financial markets.* Princeton University Press.

Cochrane, J. H. (1994). Shocks. *Carnegie-Rochester Conference Series on Public Policy 41*(1), 295–364.

Cockburn, I. M., S. Kortum, and S. Stern (2002). Are all patent examiners equal? The impact of examiner characteristics. Technical report.

Comin, D. and M. Gertler (2006). Medium-term business cycles. *American Economic Review 96*(3), 523–551.

Field, A. J. (2003). The most technologically progressive decade of the century. *American Economic Review 93*(4), 1399–1413.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*, 733–772.

Garleanu, N., L. Kogan, and S. Panageas (2012). Displacement risk and asset returns. *Journal of Financial Economics forthcoming.*

Greenwood, J. and B. Jovanovic (1999). The information-technology revolution and the stock market. *American Economic Review 89*(2).

Griliches, Z. (1989). Patents: Recent trends and puzzles. *Brookings Papers on Economic Activity. Microeconomics*, 291–330.

Griliches, Z. (1990). Patent statistics as economic indicators: A survey. *Journal of Economic Literature 28*(4), 1661–1707.

Hall, B. and R. Ziedonis (2001). The patent paradox revisited: An empirical study of patenting in the U.S. semiconductor industry, 1979-1995. *The RAND Journal of Economics 32*(1), 101–128.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2005). Market value and patent citations. *The RAND Journal of Economics 36*(1), pp. 16–38.

Harhoff, D., F. Narin, F. M. Scherer, and K. Vopel (1999). Citation frequency and the value of patented inventions. *The Review of Economics and Statistics 81*(3), 511–515.

Harhoff, D., F. M. Scherer, and K. Vopel (1997). Exploring the tail of patented invention value distributions. Technical report.

Hayashi, F. (1982). Tobin's marginal Q and average Q: A neoclassical interpretation. *Econometrica 50*(1), pp. 213–224.

Hobijn, B. and B. Jovanovic (2001). The information-technology revolution and the stock market: Evidence. *American Economic Review 91*(5), 1203–1220.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing tfp in China and India. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Jones, C. I. (2011). Misallocation, economic growth, and input-output economics. Working paper, Stanford University.

Jorgenson, D. W. and K. J. Stiroh (2000). U.S. economic growth at the industry level. *American Economic Review 90*(2), 161–167.

Klette, T. J. and S. Kortum (2004). Innovating firms and aggregate innovation. *Journal of Political Economy 112*(5), 986–1018.

Kortum, S. (1993). Equilibrium R&D and the patent-R&D ratio: U.S. evidence. *American Economic Review 83*(2), 450–57.

Kortum, S. and J. Lerner (1998). Stronger protection or technological revolution: what is behind the recent surge in patenting? *Carnegie-Rochester Conference Series on Public Policy 48*(1), 247–304.

Kortum, S. and J. Lerner (2000). Assessing the contribution of venture capital to innovation. *RAND Journal of Economics 31*(4), 674–692.

Laitner, J. and D. Stolyarov (2003). Technological change and the stock market. *The American Economic Review 93*(4), pp. 1240–1267.

Lampe, R. and P. Moser (2011). Do patent pools encourage innovation? Evidence from 20 industries in the 1930s.

Lentz, R. and D. T. Mortensen (2008). An empirical model of growth through product innovation. *Econometrica 76*(6), 1317–1373.

Moser, P. and T. Nicholas (2004). Was electricity a general purpose technology? Evidence from historical patent citations. *The American Economic Review, Papers and Proceedings 94*(2), 388–394.

Moser, P., J. Ohmstedt, and P. Rhode (2011). Patents, citations, and inventive output - evidence from hybrid corn.

Moser, P. and A. Voena (2011). Compulsory licensing: Evidence from the trading with the enemy act. *The American Economic Review forthcoming.*

Moser, P., A. Voena, and F. Waldinger (2012). German-jewish emigres and U.S. invention.

Newman, M. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics 46*(5), 323–351.

Nicholas, T. (2008). Does innovation cause stock market runups? Evidence from the great crash. *American Economic Review 98*(4), 1370–96.

Pakes, A. (1985). On patents, R&D, and the stock market rate of return. *Journal of Political Economy 93*(2), 390–409.

Ramey, V. A. and M. D. Shapiro (2001). Displaced capital: A study of aerospace plant closings. *Journal of Political Economy 109*(5), 958–992.

Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous plants. *Review of Economic Dynamics 11*(4), 707–720.

Scherer, F. M. (1965). Firm size, market structure, opportunity, and the output of patented inventions. *The American Economic Review 55*(5), pp. 1097–1125.

Scherer, F. M. (1983). The propensity to patent. *International Journal of Industrial Organization 1*(1), 107–128.

Shea, J. (1999). What do technology shocks do? In *NBER Macroeconomics Annual 1998, volume 13.* National Bureau of Economic Research, Inc.

Smiley, G. (1994). *The American economy in the twentieth century.* College Division, South-Western Pub. Co.

Tables

Table 1: Number of future citations and announcement day return

Number of cites	grant		publication		Number of cites	grant		publication	
	$l=2$	$l=4$	$l=2$	$l=4$		$l=2$	$l=4$	$l=2$	$l=4$
	(1)	(2)	(3)	(4)		(5)	(6)	(7)	(8)
	Panel A					Panel B			
A_i^+	0.112	0.087	0.038	0.027	$\hat{A}_i$	1.400	0.990	0.277	0.193
Change in firm value,	(6.53)	(6.28)	(2.35)	(2.42)	Change in firm value,	(6.74)	(6.77)	(1.68)	(1.66)
truncated at zero					adj. for meas. error				
Obs.	1842345	1842345	416282	416282	Obs.	1828616	1828616	415254	415254
R^2	0.096	0.096	0.079	0.079	R^2	0.100	0.099	0.079	0.079
	Panel C					Panel D			
$\log A_i$	0.704	0.695	0.252	0.256	$\log \hat{A}_i$	0.939	0.941	0.297	0.297
Change in firm value,	(7.11)	(6.99)	(3.01)	(2.93)	Change in firm value	(7.80)	(7.82)	(2.76)	(2.74)
log					adj. for meas. error, log				
Obs.	894697	894697	202194	202194	Obs.	1828616	1828616	415254	415254
R^2	0.099	0.101	0.083	0.082	R^2	0.102	0.102	0.084	0.084

Table shows output of a regressions of number of future citations N_i on the dollar return A_i following the day the patent is issued to the firm (columns 1-2, 5-6) or the details of the patent are disclosed by the USPTO (columns 3-4, 7-8). We construct the change in firm value A_i as the return of the firm minus the return of the market portfolio r_{fd}, times the firm's market capitalization on the day before the announcement in 1982 US dollars (billion) S_{fd-1}. We report results for three-day (l=2) and five-day (l=4) windows. Panel A shows results for change in firm value, truncated at zero; Panel B shows results for change in firm value, adjusted for measurement error; Panels C and D show corresponding results using log changes in firm value. We control for grant-year or announcement-year fixed effects and log firm idiosyncratic volatility ($\log \sigma_{ft}$). We cluster standard errors by announcement year and report t-statistics in parenthesis.

Table 2: Descriptive statistics on firm-level innovation variables

statistic	A_{ft}^+	$\hat{A}_{ft}$
Mean	0.055	0.044
St. Dev.	0.169	0.129
Percentiles		
50	0.000	0.000
75	0.014	0.024
90	0.172	0.129
95	0.339	0.250

Table presents descriptive statistics for our firm-level innovation measures A^+ and $\hat{A}$.

Table 3: Firm-level productivity and Tobin's Q

	A. Productivity					B. Tobin's Q	
mpk_{t+1}, mpl_{t+1}	Capital		Labor		q_{t+1}		
	(1)	(2)	(3)	(4)		(5)	(6)
A^+_{It}	-0.066	-0.049	-0.057	-0.046	A^+_{It}	0.031	0.030
	(-6.68)	(-5.03)	(-6.35)	(-5.12)		(2.69)	(2.62)
A^+_{ft}	0.060	0.091	0.108	0.128	A^+_{ft}	0.095	0.098
	(5.77)	(8.64)	(11.71)	(13.73)		(7.13)	(7.33)
R^2	0.844	0.847	0.847	0.850	R^2	0.684	0.686
$\hat{A}_{It}$	-0.087	-0.049	-0.084	-0.056	$\hat{A}_{It}$	0.039	0.048
	(-6.27)	(-3.56)	(-6.72)	(-6.47)		(2.62)	(3.16)
$\hat{A}_{ft}$	0.053	0.102	0.132	0.163	$\hat{A}_{ft}$	0.128	0.114
	(3.70)	(7.02)	(10.41)	(12.58)		(7.08)	(6.32)
R^2	0.844	0.847	0.847	0.850	R^2	0.684	0.686
Observations	125678	125678	120020	120020	Observations	123540	123540
Fixed Effects	I,T	I,T	I,T	I,T	Fixed Effects	I,T	I,T
Controls					Controls		
(Size, mpk or mpl)	Y	Y	Y	Y	(Size)	Y	Y
(R_f, R_I, q, σ)	-	Y	-	Y	(R, R_I, y/k, σ)	-	Y

Table shows output of the regression: $x_{ft+1} = a_0 + a_1\, A_{ft} + a_2\, A_{It} + b\, Z_{ft} + \gamma_t + c_I + \rho\, x_{ft} + u_{it}$, where $X_{ft} = [y^K_{ft}, y^L_{ft}, q_{ft}]$ is log productivity of capital, labor and Tobin's Q. Depending on the specification, the vector Z of controls includes lagged values of log Tobin's Q, firm stock return (R), firm volatility (σ, in the case of our truncated measure A^+ only), industry c_I or time γ_t fixed effects. We control for lagged firm size (log capital (Columns 1-2,5-6) or number of employees (Columns 3-4)) and productivity throughout. Standard errors are clustered by firm. All variables are winsorized by year at the 1% level.

Table 4: Firm-level reallocation: Investment

i_{t+1}	(1)	(2)	(3)	(4)	(5)
A^+_{It-1}	-0.046	-0.046	-0.040	-0.031	-0.027
	(-10.35)	(-10.45)	(-12.10)	(-9.78)	(-8.49)
A^+_{ft-1}	0.039	0.047	0.042	0.040	0.031
	(9.13)	(11.02)	(13.29)	(12.79)	(10.14)
R^2	0.085	0.093	0.221	0.260	0.273
$\hat{A}_{It}$	-0.070		-0.059	-0.042	-0.036
	(-12.27)		(-13.52)	(-9.92)	(-8.50)
$\hat{A}_{ft}$	0.040		0.034	0.041	0.039
	(7.10)		(8.23)	(9.96)	(9.66)
R^2	0.085		0.215	0.257	0.273
Observations	126727	126727	126727	126727	126727
Fixed Effects	I,T	I,T	I,T	I,T	I,T
(Size - K)	Y	Y	Y	Y	Y
(σ)	-	Y	Y	Y	Y
(i_{t-1})	-	-	Y	Y	Y
(R_f, R_I, Q)	-	-	-	Y	Y
(MPK, ROA)	-	-	-	-	Y

Table shows output of regressing firm investment i_t, defined as capital expenditures (Compustat capx) over lagged capital stock (Compustat ppegt), on our firm-level innovation measure A_f and innovation of other firms in the same industry A_I. Depending on the specification, we control for lagged values of log Tobin's Q, firm size (log capital), sales-to-capital (MPK), earnings to assets (ROA), firm stock return (R), firm volatility (σ, in the case of our truncated measure A^+ only), industry (I) or time (T) fixed effects, and lagged values of the dependent variable. Standard errors are clustered by firm. All variables are winsorized by year at the 1% level. In column (2), we only control for return volatility when we use our truncated measure. Controlling for return volatility using our measure adjusted for measurement error leads to quantitatively similar results.

Table 5: Firm-level reallocation: Labor hiring

Δn_{t+1}	(1)	(2)	(3)	(4)	(5)
A^{+}_{It-1}	-0.045	-0.046	-0.046	-0.033	-0.024
	(-7.35)	(-7.60)	(-7.93)	(-5.92)	(-4.30)
A^{+}_{ft-1}	0.008	0.016	0.024	0.020	0.017
	(1.37)	(2.80)	(4.41)	(3.61)	(3.05)
R^2	0.039	0.044	0.053	0.086	0.090
$\hat{A}_{It}$	-0.062		-0.063	-0.040	-0.029
	(-7.65)		(-8.18)	(-5.22)	(-3.74)
$\hat{A}_{ft}$	-0.008		0.014	0.020	0.017
	(-0.97)		(1.83)	(2.71)	(2.25)
R^2	0.039		0.053	0.086	0.090
Observations	119760	119760	119760	119760	119760
Fixed Effects	I,T	I,T	I,T	I,T	I,T
(Size - H)	Y	Y	Y	Y	Y
(σ)	-	Y	Y	Y	Y
(Δn_{t-1})	-	-	Y	Y	Y
(R_f, R_I, Q)	-	-	-	Y	Y
(MPL, ROA)	-	-	-	-	Y

Table shows output of regressing firm log employment growth Δn_{t+1} (Compustat item emp), on our firm-level innovation measure A_f and innovation of other firms in the same industry A_I. Depending on the specification, we control for lagged values of firm size (log no. of employees), sales-to-employees (MPL), earnings to assets (ROA), firm stock return (R), firm volatility (σ, in the case of our truncated measure A^+ only), industry (I) or time (T) fixed effects, and lagged values of the dependent variable. Standard errors are clustered by firm. All variables are winsorized by year at the 1% level. In column (2), we only control for return volatility when we use our truncated measure. Controlling for return volatility using our measure adjusted for measurement error leads to quantitatively similar results.

Table 6: Firm-level reallocation: Financial inflows

e_{ft+1}	(1)	(2)	(3)	(4)	(5)
A^{+}_{It-1}	-0.007	-0.006	-0.008	-0.001	-0.009
	(-1.28)	(-1.22)	(-1.84)	(-0.32)	(-2.08)
A^{+}_{ft-1}	0.043	0.037	0.034	0.025	0.014
	(6.80)	(5.90)	(6.30)	(4.56)	(2.81)
R^2	0.114	0.117	0.155	0.184	0.219
$\hat{A}_{It}$	-0.017		-0.017	-0.005	-0.018
	(-2.32)		(-2.68)	(-0.82)	(-2.86)
$\hat{A}_{ft}$	0.052		0.046	0.048	0.022
	(6.03)		(6.25)	(6.28)	(3.09)
R^2	0.114		0.154	0.182	0.219
Observations	126727	126727	126727	126727	126727
Fixed Effects	I,T	I,T	I,T	I,T	I,T
(Size - K)	Y	Y	Y	Y	Y
(σ)	-	Y	Y	Y	Y
(Δfin_{ft-1})	-	-	Y	Y	Y
(R_f, R_I, Q)	-	-	-	Y	Y
(ROA)	-	-	-	-	Y

Table shows output of regressing firm financial inflows, defined as debt issuance plus equity issuance minus payout (Compustat sstk + dltis - prstkc-dv-dltr) over book assets, on our firm-level innovation measure A_f and innovation of other firms in the same industry A_I. Depending on the specification, we control for lagged values of firm size (log capital), log Tobin's Q, earnings to assets (ROA), firm stock return (R), firm volatility (σ, in the case of our truncated measure A^+ only), industry (I) or time (T) fixed effects, and lagged values of the dependent variable. Standard errors are clustered by firm. All variables are winsorized by year at the 1% level. In column (2), we only control for return volatility when we use our truncated measure. Controlling for return volatility using our measure adjusted for measurement error leads to quantitatively similar results.

Table 7: Industry reallocation

x_{t+1}	Quantity of capital services, growth				Quantity of labor services			
A^+_{It}	0.018	0.015	0.026	0.022	0.024	0.032	0.031	0.039
	(2.02)	(1.78)	(3.11)	(2.49)	(1.42)	(1.88)	(1.82)	(2.35)
A^+_{MIt}	-0.112	-0.069			-0.274	-0.207		
	(-7.21)	(-3.72)			(-7.56)	(-4.60)		
R^2	0.037	0.098	0.165	0.184	0.051	0.080	0.164	0.186
$\hat{A}_{It}$	0.023	0.015	0.032	0.027	0.021	0.036	0.029	0.043
	(1.99)	(1.39)	(2.83)	(2.24)	(0.94)	(1.73)	(1.34)	(2.09)
$\hat{A}_{MIt}$	-0.153	-0.126			-0.263	-0.239		
	(-8.77)	(-4.31)			(-6.35)	(-5.49)		
R^2	0.048	0.094	0.164	0.184	0.033	0.072	0.163	0.183
Observations	1395	1395	1395	1395	1395	1395	1395	1395
Controls								
(R, σ)	-	Y	-	Y	-	Y	-	Y
Time Effects	-	-	Y	Y	-	-	Y	Y

Table reports results from a regression of the quantity $[k, n]$ of capital and labor services on the amount of innovation at the industry level $[A_I]$ and on the amount of innovation of all other industries $[A_{MI}]$. We control for time effects (T), industry stock return R^I, industry volatility σ^I (in the case of our truncated measure A^+ only) and one lag of the dependent variable. Data is from Dale Jorgenson's 35-sector KLEM, described in Jorgenson and Stiroh (2000). Sample is 1960-2005 and covers 31 industries after excluding the finance, utilities and government enterprises sector. We report t-statistics in parenthesis, with standard errors clustered by industry.

Table 8: Innovation and Firm Exit

x_{t+1}	Rate of establishment exit			
A^+_{It}	0.882	0.881	1.057	0.746
	(3.48)	(3.45)	(2.49)	(1.85)
A^+_{MIt}	-2.740	-2.778		
	(-1.52)	(-1.49)		
R^2	0.502	0.502	0.732	0.763
$\hat{A}_{It}$	1.313	1.322	2.227	2.171
	(2.14)	(2.20)	(3.31)	(2.94)
$\hat{A}_{MIt}$	-4.691	-4.745		
	(-1.32)	(-1.38)		
R^2	0.509	0.510	0.732	0.735
Observations	231	231	231	231
Controls				
(R, σ)	-	Y	-	Y
Time Effects	-	-	Y	Y

Table reports results from a regression of the rate of establishment exit on the amount of innovation at the industry level $[A_I]$ and on the amount of innovation of all other industries $[A_{MI}]$. We include industry fixed effects throughout. Depending on the specification, we include time effects (T), industry stock return R^I, and industry volatility σ^I (in the case of our truncated measure A^+ only). Data is from the tables of Business Dynamics Statistics at the US Census, and cover 7 industries, after dropping the finance sector and utilities, over the period 1977 to 2009. Industries correspond to the one-digit SIC code level. We report t-statistics in parenthesis, with standard errors clustered by year.

Table 9: Granger causality tests

Variable	Variable does not Granger cause A		A does not Granger cause variable	
	A	$\hat{A}$	A	$\hat{A}$
Output and productivity				
Productivity	0.246	0.634	0.006	0.001
Output	0.523	0.121	0.068	0.848
Technology measures of Alexopoulos (2011)				
Bowkers technology books	0.236	0.224	0.131	0.704
Library of Congress new technology books	0.111	0.005	0.615	0.647
Computer software and hardware books	0.386	0.477	0.579	0.245
Computer software, hardware, and network books	0.383	0.484	0.626	0.245
Telecommunications books	0.501	0.962	0.237	0.054
Other technology measures				
R&D Spending	0.672	0.988	0.448	0.456

Table features p-values of Granger causality tests, based on a 3-variable VAR $[X_t, \sigma_t, A_t]$ with a deterministic trend.

Table 10: Innovation and Tobin's Q

$\Delta \log Q_t$	A. Benchmark Measure			B. Alt. Normalization		
$\Delta \log A_t^+$	-0.182	-0.256	-0.218	-0.199	-0.194	-0.143
	(-4.88)	(-6.14)	(-3.91)	(-5.09)	(-4.89)	(-2.66)
R^2	0.253	0.373	0.387	0.266	0.297	0.305
$\Delta \log \hat{A}_t$	-0.185	-0.223		-0.196	-0.188	
	(-3.87)	(-4.17)		(-3.83)	(-3.67)	
R^2	0.188	0.266		0.186	0.223	
Observations	58	58	58	58	58	58
$(\log A_{t-1})$	-	Y	Y			
$(\log Q_{t-1})$	-	Y	Y			
$(\Delta \log \sigma_t)$	-	-	Y			

Table shows output of a regression of changes in log Tobin's Q at the firm (Panel A) or aggregate level (B) on our measure of innovation, controlling for changes in volatility (in the case of our truncated measure A^+ only) and lagged value of Q and innovation measures. Sample is 1952-2008. Results in Panel A use data from Compustat, include time-fixed effects and standard errors clustered at the firm level. Results in Panel B use data from the flow of funds, and standard errors are computed using the Newey-West estimator.

Figure 1: Some illustrative examples

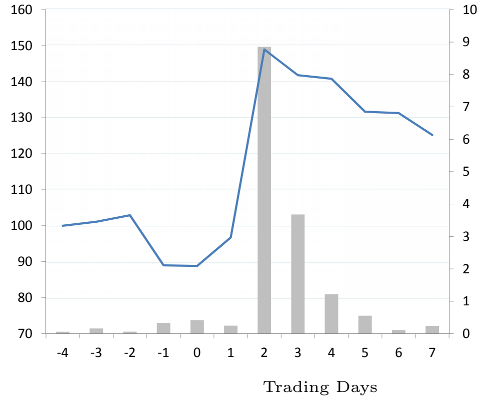

(a) Patent 4,946,778 granted to Genex on Aug, 7 1990, “Single Polypeptide Chain Binding Molecules.”

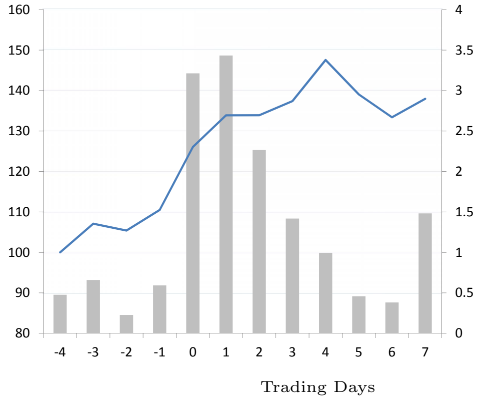

(b) Patent 5,585,089 granted to Protein Design on Dec 17, 1996, “Humanized Immunoglobulins.”

(c) Patent 6,317,722 granted to Amazon.com on Nov 13, 2001, “Use Of Electronic Shopping Carts To Generate Personal Recommendations.”

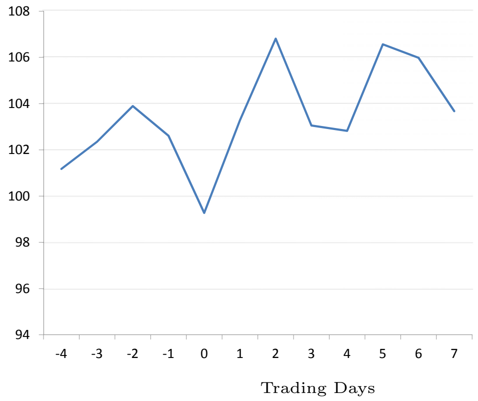

(d) Patent 4,345,262 granted to Canon on Aug 17, 1982, “Ink Jet Recording Method.”

Figure plots cumulative abnormal returns (left axis) and turnover (right axis) around the date the patent is granted for illustrative examples discussed in the text. Volume data is not available for Canon. Note that Canon reported a 6% fall in pre-tax profits on Aug 19 (two days subsequent to the patent grant).

Figure 2: Aggregate measures of innovation

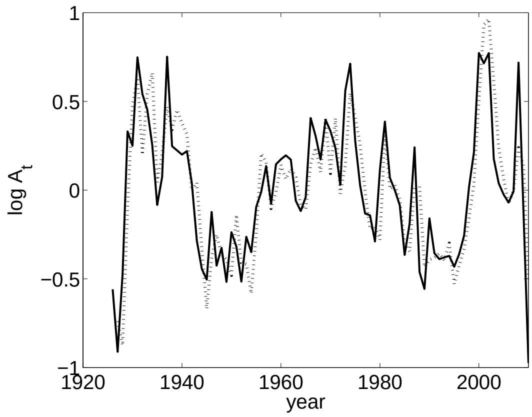

(a) Innovation A_t

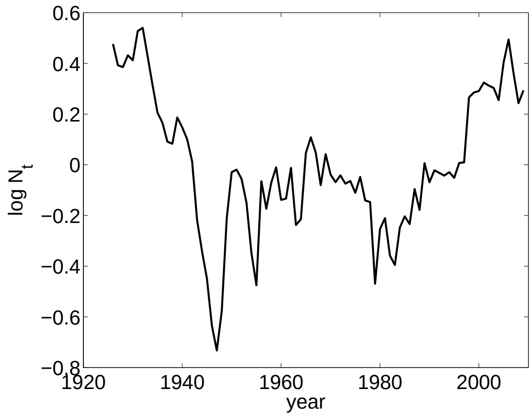

(b) Number of patents

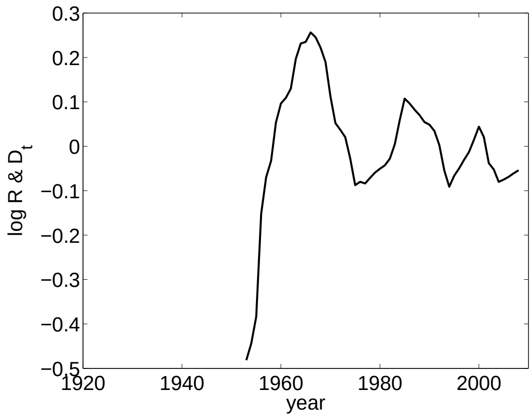

(c) R&D Stock

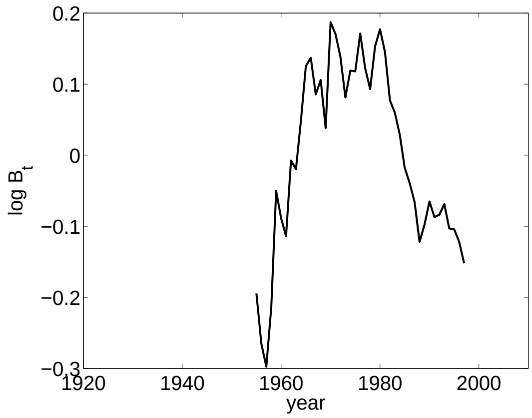

(d) New technology books (Library of Congress)

Figure plots log values of a) our measure of innovation A^+ (solid line) and $\hat{A}$ (dotted line) ; b) total number of patents granted; c) R&D Capital stock (from BEA); d) number of new technology books in the Library of Congress (from Alexopoulos (2011)).

Figure 3: Firm productivity – Dynamic response

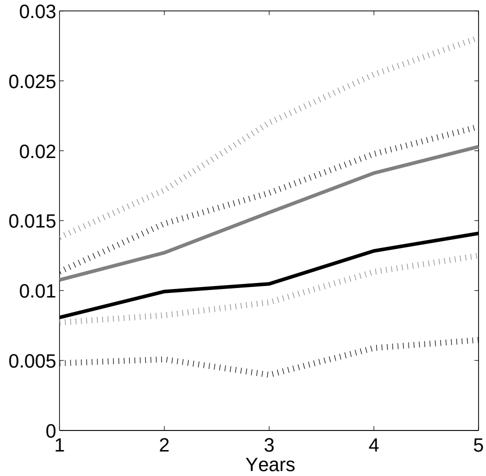

(a) Capital productivity to own innovation A_f

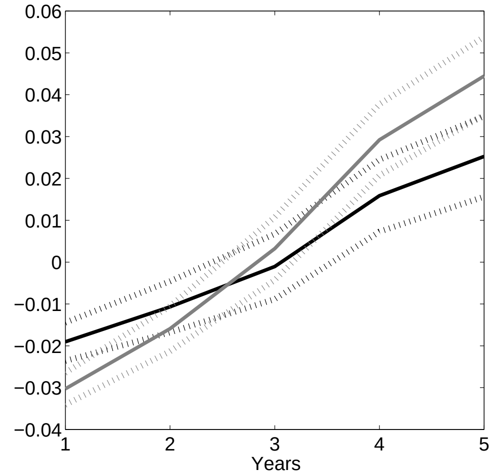

(b) Capital productivity to competitor innovation A_I

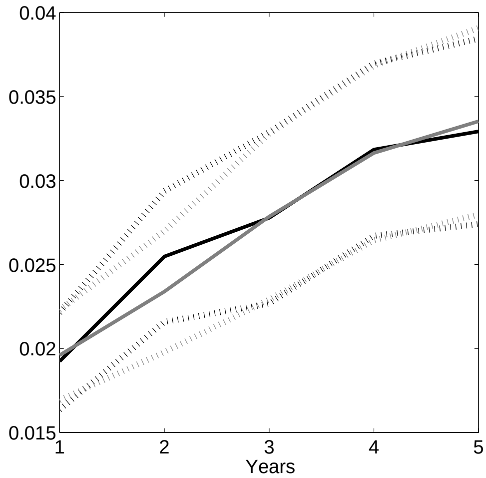

(c) Labor productivity to own innovation A_f

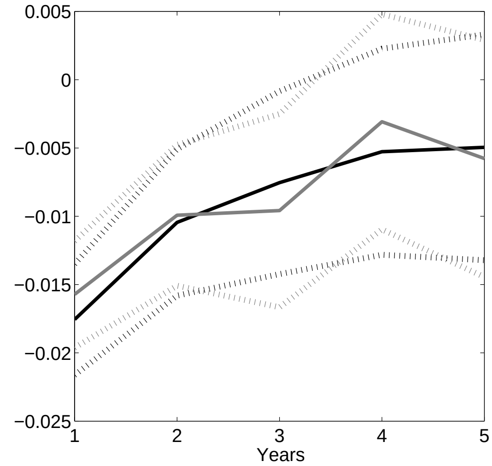

(d) Labor productivity to competitor innovation A_I

Figure plots coefficients $a_1(k)$ and $a_2(k)$ (and 90% confidence intervals) of a regression of long-run capital (left panel) or labor (right panel) firm-level productivity on innovation of firm (A_f) and competitors (A_I): $y_{ft+k} = a_0 + a_1(k)\, A_{ft-1} + a_2(k)\, A_{It-1} + \beta(k)\, Z_{ft-1} + \rho(k)\, y_{ft-1} + u_{it}$. The vector Z of controls includes lagged values of firm size (log capital or number of employees), log Tobin's Q, firm stock return (R), firm volatility (σ, in the case of our truncated measure A^+ only), industry (I) or time (T) fixed effects. More details on estimation are in the text. Grey lines correspond to results using our measure of innovation A^+ based on simple truncation; black lines correspond to results using our measure of innovation $\hat{A}$ that accounts for measurement error.

Figure 4: Industry productivity – Dynamic response

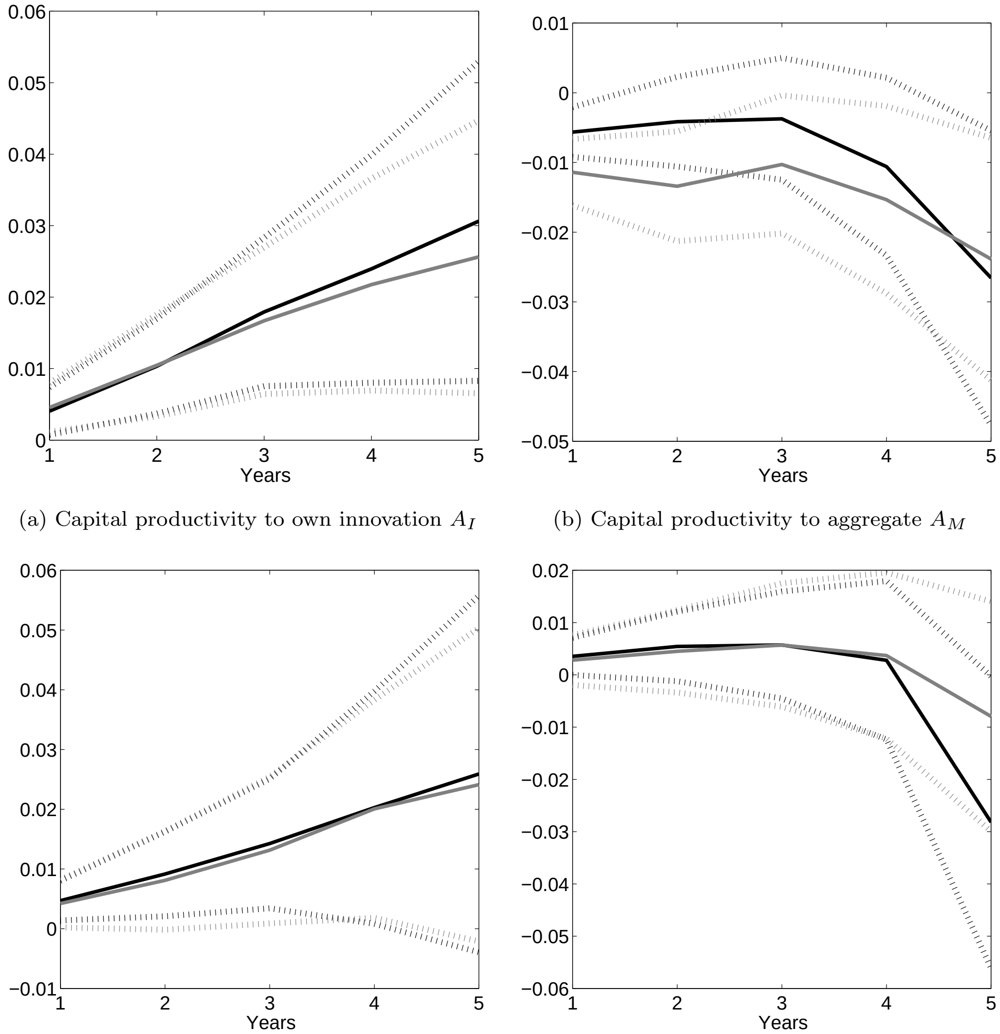

(a) Capital productivity to own innovation A_I

(b) Capital productivity to aggregate A_M

(c) Labor productivity to own innovation A_I

(d) Labor productivity to aggregate A_M

Table plots coefficients from a regression of log output y relative to aggregate output, capital mpk and labor productivity mpl on amount of innovation A_I at the industry level, and the aggregate amount of innovation excluding that industry A_M. Controls include industry stock return; industry volatility (in the case of our truncated measure A^+ only); lagged capital productivity (panels a and b); lagged labor productivity (panels c and d); and lagged output growth (panels e and f); Data is from Dale Jorgenson's 35-sector KLEM. We plot coefficient estimates and 90% confidence intervals using standard errors clustered by industry. Grey lines correspond to results using our measure of innovation A^+ based on simple truncation; black lines correspond to results using our measure of innovation $\hat{A}$ that accounts for measurement error. Right hand-side variables are standardized to unit standard deviation.

Figure 5: Firm and industry output – Dynamic response

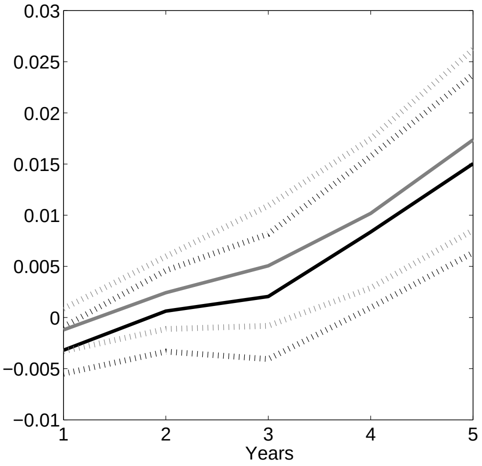

a) Firm output to own innovation A_f

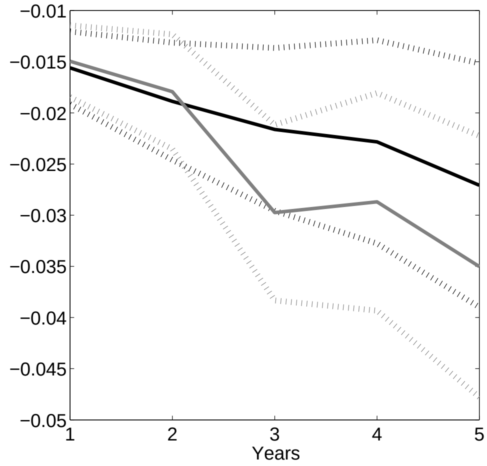

(b) Firm output to competitor A_{If}

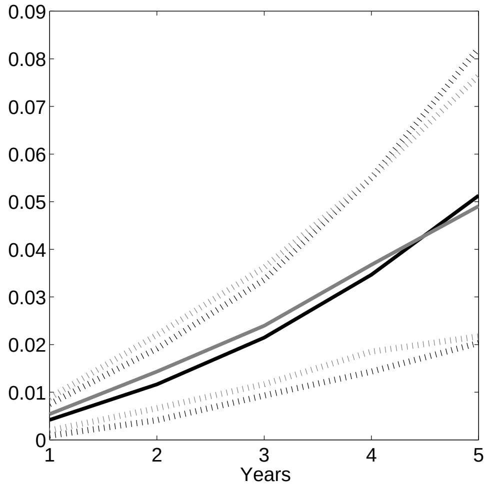

(c) Industry output to own innovation A_I

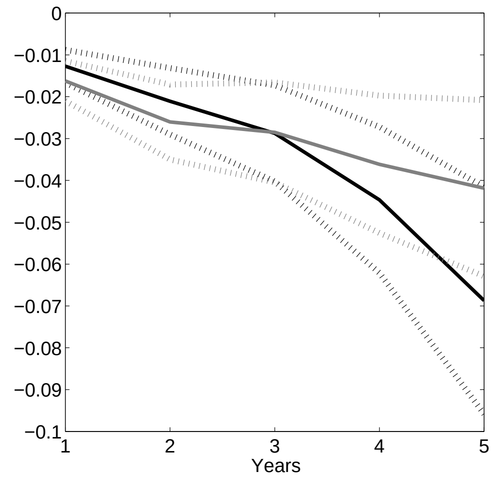

(d) Industry output to other industry innovation A_{IM}

Figure plots coefficients $a_1(k)$ and $a_2(k)$ (and 90% confidence intervals) of a regression of long-run capital (left panel) or labor (right panel) firm-level productivity on innovation of firm (A_f) and competitors (A_I): $y_{ft+k} = a_0 + a_1(k)\, A_{ft-1} + a_2(k)\, A_{It-1} + \beta(k)\, Z_{ft-1} + \rho(k)\, y_{ft-1} + u_{it}$. The vector Z of controls includes lagged values of firm size (log capital or number of employees), log Tobin's Q, firm stock return (R), firm volatility (σ, in the case of our truncated measure A^+ only), industry (I) or time (T) fixed effects. More details on estimation are in the text. Grey lines correspond to results using our measure of innovation A^+ based on simple truncation; black lines correspond to results using our measure of innovation $\hat{A}$ that accounts for measurement error.

Figure 6: Innovation and Industry Growth

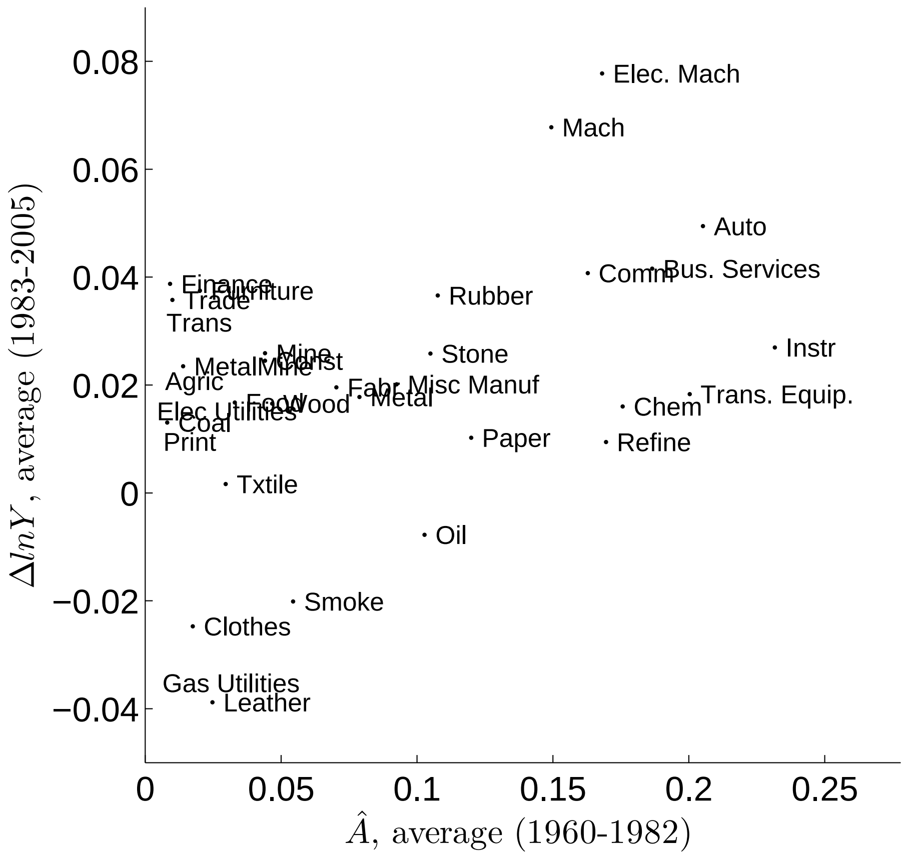

Figure plots the average output growth rate of 34 industries during the 1983-2006 period, versus the amount of innovation in 1960-1982. We use our innovation measure adjusted for measurement error $\hat{A}$. Data is from Dale Jorgenson's website. Output is measured as value added in constant prices.

Figure 7: Impulse responses, Productivity

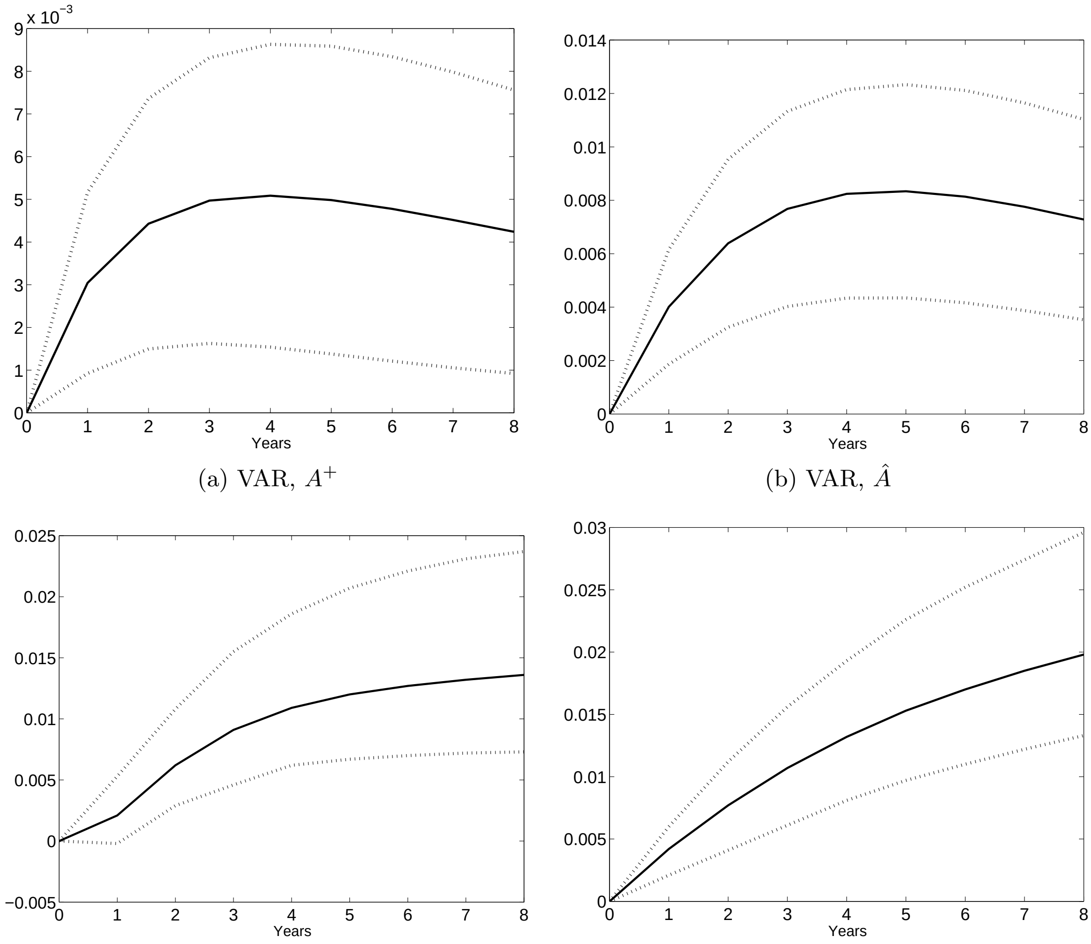

(a) VAR, A^{+}

(b) VAR, $\hat{A}$

(c) VECM, A^{+}

(d) VECM, $\hat{A}$

Figure shows impulse response of productivity from a VAR (top) and VECM (bottom). We obtain impulse responses by ordering our innovation measure last. Panels a and c present results using our measure of innovation A based on simple truncation. Panels b and d present results using our measure of innovation $\hat{A}$ that accounts for measurement error. We include a deterministic trend in all specifications. We select lag length based on the AIC criterion. In panels a and c we include the log cross-sectional average of idiosyncratic volatility. Dotted lines represent 90% confidence intervals using standard errors are computed using 500 bootstrap simulations.

Figure 8: Impulse responses, Output

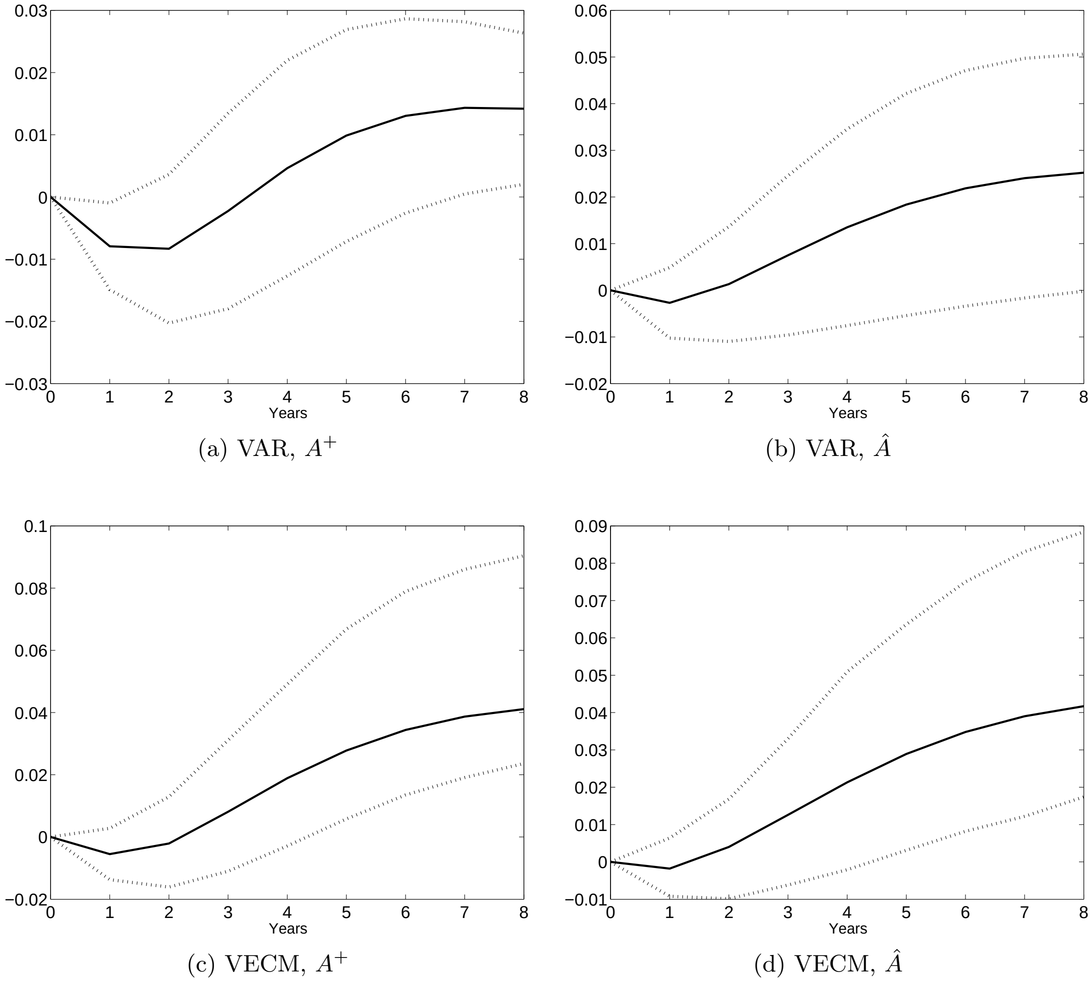

(a) VAR, A^{+}

(b) VAR, $\hat{A}$

(c) VECM, A^{+}

(d) VECM, $\hat{A}$

Figure shows impulse response of output from a VAR (top) and VECM (bottom). We obtain impulse responses by ordering our innovation measure last. Panels a and c present results using our measure of innovation A based on simple truncation. Panels b and d present results using our measure of innovation $\hat{A}$ that accounts for measurement error. We include a deterministic trend in all specifications. We select lag length based on the AIC criterion. In panels a and c we include the log cross-sectional average of idiosyncratic volatility. Dotted lines represent 90% confidence intervals using standard errors are computed using 500 bootstrap simulations.

Figure 9: Impulse responses, Consumption

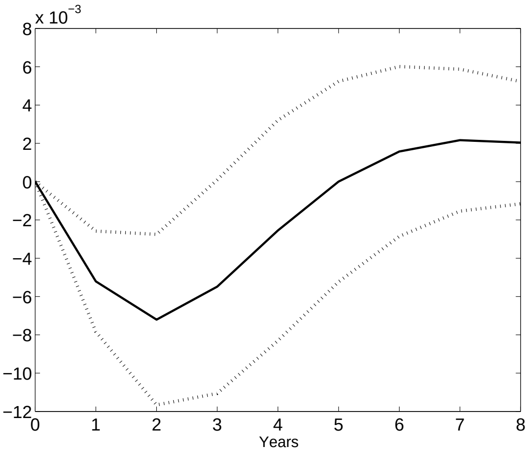

(a) VAR, A^+

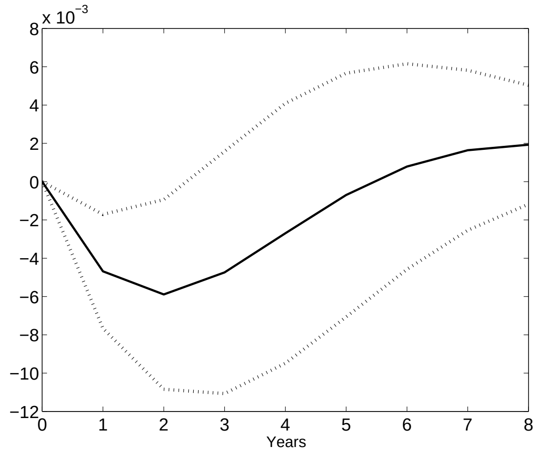

(b) VAR, $\hat{A}$

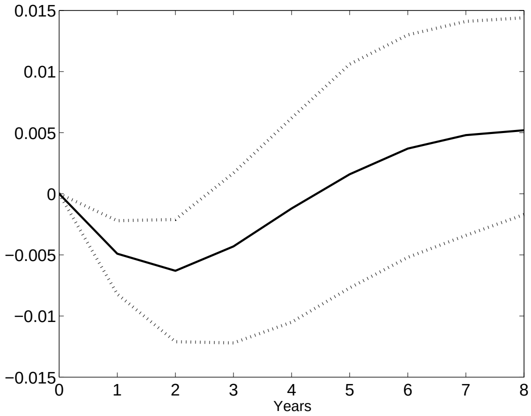

(c) VECM, A^+

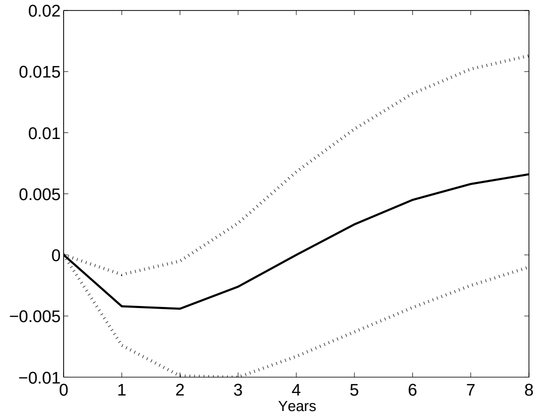

(d) VECM, $\hat{A}$

Figure shows impulse response of output from a VAR (top) and VECM (bottom). We obtain impulse responses by ordering our innovation measure last. Panels a and c present results using our measure of innovation A based on simple truncation. Panels b and d present results using our measure of innovation $\hat{A}$ that accounts for measurement error. We include a deterministic trend in all specifications. We select lag length based on the AIC criterion. In panels a and c we include the log cross-sectional average of idiosyncratic volatility. Dotted lines represent 90% confidence intervals using standard errors are computed using 500 bootstrap simulations.

Online Appendix to "Technological Innovation, Resource Allocation and Growth" –Not for Publication–

Leonid Kogan, Dimitris Papanikolaou,
Amit Seru and Noah Stoffman

A Patent Data

Our measure of innovation relies on using information on patents that a firm creates and the stock market response to news about these patents. We now discuss the data that we employ in our analysis.

Patents in the United States are granted by the United States Patent and Trademark Office (USPTO). We download the entire history of U.S. patent documents from Google Patents.[1] Each of about 7.8 million patent files was downloaded using an automation script.[2]

To construct our measure of innovation, we match all patents in the Google data to corporations whose returns are in the CRSP database. Patent regulations require that only an individual, not a corporation, can be an inventor. However, the inventor can assign the granted property rights to a corporation or another person. Therefore, when patents are granted they always have an inventor, and sometimes an "assignee", that is, one or more corporations or persons.

For most patents, Google provides a text version of the patent document, created using OCR software. We use this text version of the document to extract the names of corporations to which patents are assigned. However, OCR technology is imperfect, and many of the downloaded documents include a great deal of garbled text. We therefore make use of a number of text analysis algorithms to extract relevant information from the documents.

Our sample covers patents granted between 1926 and 2010 matched to firms with returns in CRSP database. Since we merge our patent data with data on stock returns, we are limited to the period after 1926, when the CRSP database begins.

---

[1] `http://www.google.com/patents`

[2] Google also makes available for downloading bulk patent data files from the USPTO. The bulk data does not have all of the additional "meta" information including classification codes and citation information that Google includes in the individual patent files. Moreover, the quality of the text generated from Optical Character Recognition (OCR) procedures implemented by Google is better in the individual files than in the bulk files provided by the USPTO. As explained below, this is crucial for identifying patent assignees.

Matching patents to firms

Here, we briefly discuss the steps our matching procedure followed, and provide extensive details Section C. We search the document for the words "assignee" or "assigned" and extract the text that immediately follows. This text is either a company name, or the name of an individual to whom the patent is assigned. We then count the number of times each assignee name appears across all patent documents. We compare each assignee name to more common names, and if a given name is "close", in the sense of the Levenshtein distance, to a much more common name, we substitute the common name for the uncommon name.[3] For example, one of the most common names is "General Electric Company", which is associated with over 43,000 patents. We substitute this name for the far less common, but quite similar, names "General Electbic Oohpany", "General Electbic Cqhpany", and "Genebal Electbic Compakt".

At this point, we have an assignee name for each patent. These names must be matched to a company identifier such as the CRSP permco. This is accomplished in two steps. We begin by looking only at patents that are also in the NBER database. For each assignee name identified in the steps above, we count how many different permcos are matched to patents in the NBER database. For example, all of the patents with an assignee name "General Electric Company" are matched to one permco in the NBER database. We can therefore safely assume that *all* of the patents assigned to the General Electric Company can be matched to that permco, *even for patents not included in the NBER data*. Remaining assignee names are matched to CRSP firm names using a name matching algorithm.[4] The algorithm uses a score based on the inverse word frequency to match assignee names to possible company names. For example, the word "American" is quite common in company names, and so contributes little to name matching; the word "Bausch" is quite uncommon, so it is given much more weight. Visual inspection of the matched names confirms very few mistakes in the matching.

Extracting patent citations

We extract patent citations from three sources. First, all citations for patents granted between 1976 and 2011 are contained in text files available for bulk downloading from Google. These citations are simple to extract and likely to be free of errors, as they are official USPTO data. Second, for patents granted before 1976, we extract citations from the OCR text generated from the patent files. We search the text of each patent for any 6- or 7-digit numbers, which could be patent numbers. We then check if these potential patent numbers are followed closely by the corresponding grant date for that patent; if the correct date appears, then we can be certain that we have identified a patent citation. Since we require the date to appear near any potential patent number, it is

[3]The Levenshtein distance is the number of edits required to make one string match another string, where an edit is inserting, deleting, or substituting one character.

[4]The algorithm is based on code written by Jim Bessen, available at `http://goo.gl/m4AdZ`.

unlikely that we would incorrectly record a patent citation – it is far more likely that we would fail to record a citation than record one that isn't there. Third, we complement our citation data with the hand-collected reference data of Nicholas (2008). See Section C of this Appendix for a detailed explanation of this process.

Summary statistics

We now provide some statistics that lend credence to our method for extracting patent information. Table 1 shows the number of patents we match to companies. Of the 6.2 million patents granted in or after 1926, we find the presence of an assignee in 4.4 million. The matching procedure provides us with a database of 1.9 million matched patents, of which 523,301 (27%) are not included in the NBER data. Figure 1 graphs the total number of patents matched by the year the patent was granted. Patents included in the NBER data, which is the most comprehensive database previously available, are shown in light shading. Patents unique to our database are presented in dark shading. Note that the two sets of data appear to fit together fairly smoothly, and that even during the period covered by the NBER data, our database adds an average of 2,187 patents to the NBER data.[5]

Table 2 provides additional summary statistics. Overall, our data provides a matched permco for 66% of all patents with an assignee, or 31% of all granted patents. By comparison, the NBER patent project provides a match for 32% of all patents from 1976–2006, so our matching technique works quite well, even using only data extracted from OCR documents for the period before the NBER data. Another point of comparison is Nicholas (2008), who uses hand-collected patent data covering 1910 to 1939. From 1926–1929, he matches 9,707 patents, while our database includes 8,858 patents; from 1930–1939 he has 32,778 patents while our database includes 47,036 matches during this period.

B Other Data

Stock Market and Financial Data

The return data used to assess the stock market response to news about patents are from CRSP over the period 1926–2010. In several of our analyses we use financial and accounting data that are from Compustat. The sample in these cases is determined by the availability of Compustat data (available from 1951 onwards). As is standard, we omit financial firms and utilities from our analysis.

[5]We use information on the patent-assignee match in the NBER data to assist with our matching, so the match during the overlapping period is mostly the same, by construction. An exception is for cases where there is apparently a mistake in the NBER match and our patent-assignee frequency-based matching system corrects an error.

Business-cycle data

Productivity is utilization-adjusted TFP from Basu, Fernald, and Kimball (2006). Populations is from the U.S. Census Bureau (http://www.census.gov/popest/national/national.html). Output and consumption are from the Bureau of Economic Analysis. Output is gross domestic product (NIPA Table 1.1.5) divided by the consumption price index (St Louis Fed, CPIAUCNS). Consumption is consumption of non-durables plus services, deflated by the price index of non-durables and services respectively (NIPA Tables 1.1.5, 2.3.4). To get hours worked, we merge series CEU0500000007 and EEU00500005 from the BLS, times total private employment(BLS, CEU0500000001) divided by population. The aggregate Tobin's Q is computed using NIPA and FRB Flow of Funds Data as in Laitner and Stolyarov (2003). Finally, the time series information on R&D expenditure spending in the US is obtained from the NSF website.

Industry Data

The industry-level data is from the KLEMS dataset of Dale Jorgenson. We use industry value added (constant prices) as measure of industry output.

Firm-level data

We define the investment rate as capital expenditures (Compustat: capx) divided by lagged gross property, plant and equipment (ppegt); labor hiring as the percentage change in the number of employees (emp); financial capital inflows as debt issuance plus equity issuance minus payout (Compustat sstk + dltis - prstkc-dv-dltr) normalized by assets (at); return on assets as operating income (ib) plus depreciation divided by lagged gross property, plant and equipment; Tobin's Q as the sum of the market value of common equity (CRSP December market capitalization), the book value of debt (dltt), the book value of preferred stock (pstkrv), minus the book value of inventories (invt) and deferred taxes (txdb), divided by gross property, plant and equipment (ppegt); productivity of capital as sales (sale) plus change in inventories (invt) over gross property, plant and equipment (ppegt); productivity of labor as sales (sale) plus change in inventories (invt) over number of employees (emp).

C Patent Data Construction – Details

In this section we explain in detail how we constructed our new patent data set. The raw data are very large and not very well structured, and thus required a great deal of effort to clean. We used a number of techniques to extract, clean, and match assignees from patents. As with any such project there is a trade-off between type-I and type-II errors (in this case, failing to match an assignee to CRSP or incorrectly matching an assignee to CRSP). Our approach was to be as conservative as possible, attempting to minimize mismatches while at the same time extracting as many correct matches as possible.

C.1 Data sources

We use three sources of data to construct the new patent database:

1. Details of patents granted from 1976–2010 is available in high-quality text files available for bulk downloading from Google, through a special data-hosting arrangement with the United States Patent and Trademark Office (USPTO). The text files use one of two data structures that allows relatively straightforward data extraction: files for 2001–present use XML, while files for 1976–2000 use a fixed-width data structure with labeled fields.

2. Patents granted prior to 1976 are also stored on Google, but only in individual web pages (one per patent). Information during this period is drawn from Optical Character Recognition (OCR) of original patent documents, and is of highly-variable quality. There is very limited, if any, structure to these files.

3. We use the NBER patent data (Hall and Trajtenberg, 2001), which covers the period 1976–2006, to help with the matching and to validate our other data extraction methods.

Due to varying data sources and quality over time, it worth stressing that from 1976–2010 we use the *official records* of the USPTO. As we discuss below, we are able to provide some additions and corrections to the NBER data during the period of overlap with our data. Prior to 1976 the data are more difficult to work with, but we have implemented a number of sophisticated text analysis algorithms to create a very high-quality database.

Downloading individual patent files

We downloaded individual patent data from Google. The URL for each patent's summary page is of the form `http://www.google.com/patents/?id=RD0yAAAAEBAJ`, where `RD0y` is a 4-character code used by Google to identify each patent. The IDs use any of the characters {a, …, z, A, …, Z, 0, …, 9, _, -}. There are $64^4 = 16.8$ million possible IDs, but only about 8 million patents. However, all 16.8 million URLs must be checked, because there is no publicly-available mapping of patent numbers to the Google ID.

A screen shot of the summary page for the patent with id RD0y, which is patent 4,345,262, is shown in Figure 1. The main page includes—when available—the title of the patent, the filing and grant dates, the abstract, inventor(s), original assignee(s), current classifications, and a record of citations (out-cites) and references (in-cites). The information reported on this page by Google was gleaned from the OCR analysis of the original patent document, and consequently less information is reported for older documents, especially patents granted before 1976.

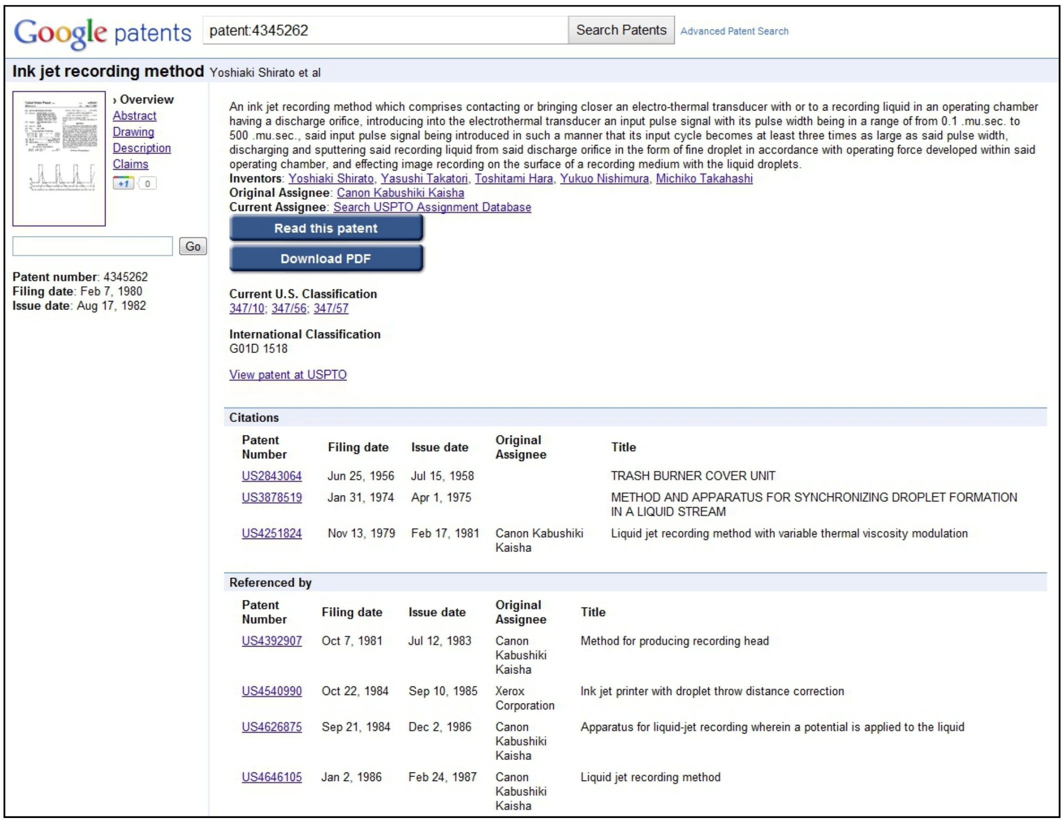

Figure 1: Google summary page for U.S. patent 4,345,262

Using a Perl automation script, we sequentially navigated to each of the 16.8 million patent summary pages.[6] From this page, we stored all available information. The script then loaded the "Read this patent" link, which loads a PDF version of the patent document. From here, we loaded the "plain text" version of the document, which is simply the text derived from OCR of the PDF document. Examples of these pages are shown in Figures 2 and 3. We saved the complete text of the plain text version of each patent. After compression, the complete archive of text requires

[6]Google generally blocks users from downloading so many web pages. We are grateful to Hal Varian for his assistance with arranging permission to access these pages.

approximately 56 gigabytes of disk space.

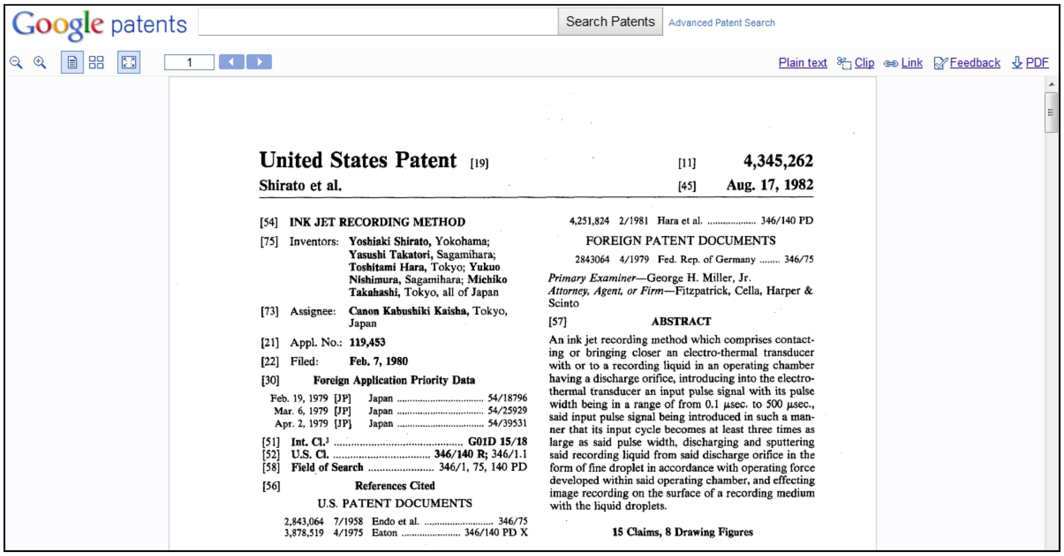

Figure 2: PDF view

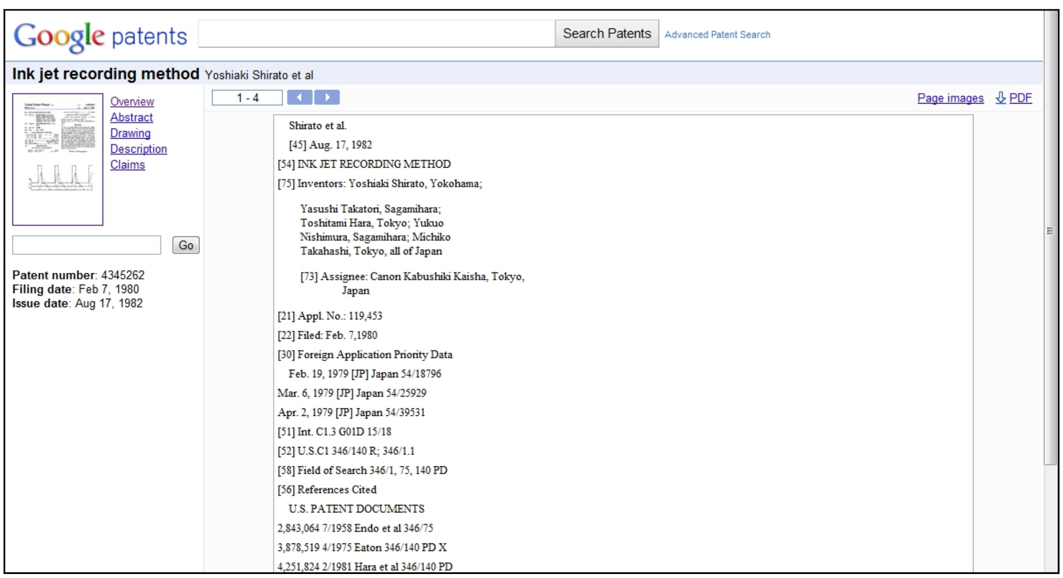

Figure 3: Plain text view

Download bulk patent files

As part of a special arrangement with the USPTO, Google also makes available for downloading bulk patent data files. The bulk data does not have all of the additional "meta" information including classification codes and citation information that Google includes in the individual patent files.

Moreover, the quality of the text generated from OCR procedures implemented by Google is better in the individual files than in the bulk files provided by the USPTO. We therefore do not use the bulk download files for data in the pre-NBER period.

For the post-NBER period, however, the bulk data files are of extremely high quality because they are based on digital patent records as opposed to OCR data drawn from images of patent documents. These data files are provided either in XML format or in a fixed-width record format. In both cases, all fields (inventor name, grant date, etc.) are clearly identified. We rely on these files to construct the database during the post-NBER period (2006–2009) and to make additions and corrections to the NBER data.

C.2 Identifying assignees

Extracting assignee names

For data during the post-1976 period, we can use the XML files available for bulk download to identify the assignee with virtually no errors.

During the pre-1976 period, we cannot rely solely on Google's extraction of the filing and grant dates or the assignee name because the OCR for patents frequently has errors. As an example, consider patent 1,131,249, shown in Figure 4.


UNITED STATES PATENT OFFICE.

EARLE R. KNIGHT, OF NORWOOD, OHIO, ASSIGNOR, BY MESNE ASSIGNMENTS, TO ALLIS-CHALMERS MANUFACTURING COMPANY, A CORPORATION OF DELAWARE.

RETAINING-RING FOR DYNAMO-ELECTRIC MACHINES.

1,131,249. Specification of Letters Patent. Patented Mar. 9, 1915.

Application filed May 1, 1909. Serial No. 492,399.

To all whom it may concern:

Be it known that I, EARLE R. KNIGHT, a citizen of the United States, residing at Norwood, in the county of Hamilton and State of Ohio, have invented certain new and useful Improvements in Retaining-Rings for

ring is shown in perspective in Fig. 3. It comprises two segmental duplicate interchangeable parts each part of which has two male portions 16 and two female portions 17. The male and female portions of one part of the split retaining rings are respectively

Figure 4: Title page of patent 1,131,249

It is clear to a human reader that this patent was assigned to the Allis-Chalmers Manufacturing Company, but the OCR for this patent reads

```
EASLS B. KNIGHT, OF NORWOOD, OHIO, ASSIGNOR,, BY MESH'S ASSIGN1IBNTS, TO ALUSCHALME&S
MANOTAC/rURING- COMPANY, A COBPOBAT'lOH OF DELAY/ABE.
```

Consequently, Google records the assignee as "BY MESH S ASSIGNIIBNTS", which is clearly not accurate.

We therefore rely on a number of textual analysis algorithms to extract the assignee name from the full text files we saved for each patent. In general, our approach to performing a "fuzzy" match on a text string is to use the maximum likelihood n-gram approach described by Norvig (2009).

We begin by identifying the text where the assignee, if there is one, will be named. We do this by searching for words that appear similar to "assign", "assignor", or "assignee". When found near the beginning of the patent document, this word is typically followed closely by the name of the assignee, so we extract a text string of 200 characters for further processing. The assignee may be a person, or a corporation, in which case the name will include a word like "company", "corporation" or "incorporated". If the word "assign" and its variants are not found, we assume the inventor did not assign the patent to another entity.

Cleaning assignee names

After extracting the string that is likely to contain the assignee name, additional cleaning is necessary. Because of OCR errors, company names may be garbled. For example, the General Electric Company, which has more than 43,000 patents in our data, appears as "General Electbic Oohpany", "General Electbic Cqhpany", and "Genebal Electbic Compakt", among hundreds of other misspellings. To fix these, we first count how many patents have been granted to each assignee name, regardless of how the assignee name is spelled. In this example, General Electric Company appears in 42,693 patents, while each of the misspelled variants appears fewer than 5 times.

We then calculate the Levenshtein edit distance[7] between each assignee name and all other names that have more patents. If any assignee name is close to another assignee name that is associated with many more patents, then the more common assignee name is substituted for the less common name. This algorithm correctly identifies all of the misspellings noted above as being General Electric.

After cleaning assignee names, we manually checked which misspelled names were matched to the 500 assignees with the most patents to confirm that no significant errors were introduced in this step.

[7] The Levenshtein distance is the number of edits required to transform one string into another string, where allowed edits are inserting, deleting, or substituting one character. For example, the Levenshtein distance between "patent" and "parent" is 1, while the distance between "patent" and "apparent" is 3.

Matching to CRSP

Having extract a list of assignee names, the next step is to match company names to the CRSP permco identifier. This is accomplished in three steps.

We begin by looking only at those patents that are included in the NBER patent database. For each assignee name identified in the steps above, we count how many *different* permcos are matched to patents in the NBER database. For example, all of the patents with an assignee name "General Electric Company" are matched to one permco in the NBER database. We can therefore safely assume that *all* of the patents assigned to the General Electric Company can be matched to that permco, *even for patents not included in the NBER data*. This step allows us to draw on the extensive data cleaning and matching project undertaken by Hall and Trajtenberg (2001) while at the same time identifying some errors in the NBER database. For example, patent 4,994,660 was assigned to General Electric but is identified in the NBER data as being assigned to Hitachi, Ltd. Because our algorithm relies on name matching, and the assignee name in that patent is General Electric, the patent is correctly identified in our data.

The first step only helps us match assignees with patenting activity during the period covered by the NBER database. We therefore proceed with a second step to match remaining assignee names. We do this with a name matching algorithm based on code written by Jim Bessen, available at `http://goo.gl/m4AdZ`. The algorithm uses a score based on the inverse word frequency to match assignee names to possible company names. For example, the word "American" is quite common in company names, and so contributes little to name matching; the word "Bausch" is quite uncommon, so it is given much more weight. Visual inspection of the matched names confirms very few mistakes in the matching.

Finally, we identify the top 250 assignees (by patents) with no CRSP matches. We manually matched these to CRSP whenever possible. Examples of firms requiring manual matching include research subsidiaries such as 3M Innovative Properties Company, which was not successfully matched to CRSP because its name differs substantially from its parent. Although we only checked 250 assignees, this manual check allowed us to match an additional 64,000 patents. Firms with high patenting activity but not matched to CRSP are either private companies or foreign firms that are not listed on U.S. exchanges, an example of which is Hoffmann-La Roche, the large Swiss drug company.

C.3 Correcting grant dates

The filing and grant dates of the patents are subject to the same sort of OCR errors as the assignee information. The grant dates are particularly important for our purposes because we use them to calculate the return around the grant date. Since patent numbers are sequential by grant dates, it is easy to infer missing or incorrect grant dates by comparing patent dates to the grant dates of adjacent patents. The same is not true of filing dates, but do not use filing dates in our current work.

To populate missing patent dates and correct mistakes we identify the 3 non-missing grant dates immediately preceding and following each patent. For example, if patent k's grant date is missing but patents $(k-3,\ldots,k-1,k+1,\ldots,k+3)$ have grant date D, then we set patent k's grant date to D. By applying this procedure iteratively we are able to correct most grant dates, with the exception of patents whose grant dates are missing and lie at a boundary between two grant dates. We fill in these missing boundary dates by manually checking their grant dates on the USPTO's web site.

While we don't rely on filing dates in the paper, it is possible to correct large errors in filing dates by identifying cases where filing dates occur after the grant date, or much earlier than the filing dates of adjacent patents. These errors often occur only in the year, so we can keep the recorded month and day the same while setting the year of the patent filing to the median filing year of a 20-patent window centered on a patent with an apparent error.

C.4 Extracting citations

Extracting patent citations from the patent text documents presents another challenge. The format of a patent document has changed several times, as has the location and formatting of citations within the document. For example, Figure 5 shows the references section of patent 2,423,030, granted in 1947. The format seen here is the first format used after patent citation began in February, 1947.

A human reader has no problem identifying the citations in this patent. But to understand the considerable challenge faced in automating this identification, consider the OCR for this part of the patent:

```
other side. 35 REFERENCES CITED
By this invention I am able satisfactorily and The following references are of record in the
conveniently to effect the drying of'Shaped pot- jjle of tllis patent:
tery or other ceramic articles either in their
```

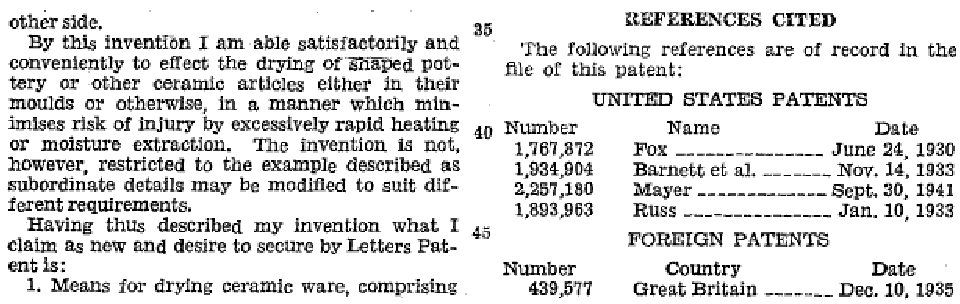

other side.

By this invention I am able satisfactorily and conveniently to effect the drying of shaped pottery or other ceramic articles either in their moulds or otherwise, in a manner which minimises risk of injury by excessively rapid heating or moisture extraction. The invention is not, however, restricted to the example described as subordinate details may be modified to suit different requirements.

Having thus described my invention what I claim as new and desire to secure by Letters Patent is:

1. Means for drying ceramic ware, comprising

REFERENCES CITED

The following references are of record in the file of this patent:

UNITED STATES PATENTS

Number	Name	Date
1,767,872	Fox	June 24, 1930
1,934,904	Barnett et al.	Nov. 14, 1933
2,257,180	Mayer	Sept. 30, 1941
1,893,963	Russ	Jan. 10, 1933

FOREIGN PATENTS

Number	Country	Date
439,577	Great Britain	Dec. 10, 1935

Figure 5: A patent citation section

```
moulds or otherwise, in a manner which min- UNITED STATES PATENTS
imises risk of injury by excessively rapid heating 40 Number Name Date
or moisture extraction. The invention is not, 1,767,872 Pox June 24, 1930
however, restricted to the example described as 1^934,904 Barnett et al Nov. 14', 1933
subordinate details may be modified to suit dif- 2,257,180 Mayer Sept. 30, 1941
ferent requirements. 1,893,963 Russ Jan. 10,1933
Having thus described my invention what I 45
claim as new and desire to secure by Letters Pat- * ("uu-^ f A 1 Jun 11>
entis: Number Country Date
1. Means for drying ceramic ware, comprising 439,577 Great Britain Dec. 10,1935
```

Our approach is to identify any text that could be a patent number (a 6- or 7-digit number, perhaps separated by commas, spaces, or other "noise" characters) and is closely followed by the correct grant date for the cited patent. In particular, for every potential patent number we identify, we determine its grant date and then search near the possible citation for that date. If the date appears, we can be very confident that we have correctly identified a citation. For example, for the patent shown in Figure 5 we extract the patent number 1,767,872 and then confirm that its grant date—June 24, 1930—appears somewhere nearby in the text. By using this two-step process to identify citations, our citation extraction is of very high quality—the probability that some random 7-digit number will be followed closely by the correct date is clearly extremely small.

Our citation extraction method provides more citations than what is available on the Google summary page. For example, the Google summary page for the patent shown in the previous example provides no citations at all, while our algorithm correctly extracted all four citations. (We exclude citations to foreign patents, as these patents are not in our database.) In general, Google does not currently report out-cites from patents granted before 1976, so we use this extraction method on all patents granted between 1926 and 1975.

C.5 Data validation

As previously mentioned, any data extraction project such as this can lead to two types of errors: matching a patent to a firm that is not the assignee, or failing to match a patent to a any firm when it does have an assignee. Our strategy makes the first error very unlikely, as a match occurs only when a name closely resembling a CRSP company name appears around the word "assignee" at the beginning of patent document. We cannot be sure how many errors of the second type we made, but we have taken care to ensure that our algorithms allow as flexible matching as possible.

We also did two final checks to check the quality of our matching strategy. First, we visually inspected a random sample of 500 patents granted between 1926 and 1975 and confirmed that assignees had been correctly extracted, and correctly matched if the assignee appeared in CRSP. This is obviously a very small sample of patents, but this careful check confirmed that no serious errors existed.

Second, we applied the extraction and matching algorithms we used in the pre-1976 period to a random sample of 25,000 patents granted between 1976 and 1999. We then compared our matches to the matches in the NBER data. None of our matches was incorrect, and only 3 patents were incorrectly not matched to an assignee. In other words, applying the techniques we used on pre-1976 data to data from the NBER period yields results that are virtually identical to those in the NBER database.

Table 1: Number of patents

Data step	Number of patents
Total downloaded patents	7,797,506
Granted in 1926 or later	6,272,428
Identified as having an assignee	4,374,524
Matched to CRSP	1,928,123
Of which:	
Present in NBER data	1,404,822
New to this paper	523,301

The table provides details on patents in our sample. We begin with all patents downloaded from Google Patents, and restrict the sample to post-1926. Not all patents have assignees, and among those that do, not all are companies in CRSP. We are able to match 1,928,123 patents to CRSP firms, of which 523,301 (27%) are new to this study. Further details are reported in Table 2 and Figure 1.

Table 2: Assignee matching by Decade

	Number of patents			Number of unique	
Years	Total	With assignee	Matched to CRSP	Matched firms	CRSP firms
1926–1929	174,022	48,433	8,858	182	786
1930–1939	442,700	172,925	47,029	355	951
1940–1949	307,499	141,345	60,616	451	1,042
1950–1959	425,953	171,157	82,255	587	1,246
1960–1969	567,599	265,524	165,409	1,175	3,177
1970–1979	690,459	393,661	247,102	2,086	7,204
1980–1989	708,735	579,518	235,525	2,756	11,715
1990–1999	1,109,398	933,705	352,005	3,664	14,882
2000–2010	1,846,063	1,668,256	729,324	4,415	11,900
All years	6,272,428	4,374,524	1,928,123	7,864	26,660

The shows summary statistics for patents in our sample by decade. Column 2 shows the total number of patents, and column 3 shows how many patents are identified as having an assignee. Column 4 shows how many of those patents with assignees are matched to a company in CRSP. (The remaining assignees are either individuals, private companies, or the matching process was unable to identify the correct company.) Columns 5 and 6 show how many unique firms there are matched to patents or in CRSP.

Table 3: Stock turnover around patent announcement days

Event	$l=-1$	$l=0$	$l=1$	$l=2$	$l=3$	$l=4$
A. Patent grant	-0.396	0.046	0.082	0.074	0.006	-0.377
	(-8.93)	(2.55)	(5.23)	(4.51)	(0.23)	(-9.12)
B. Patent publication	0.094	0.182	0.283	-0.136	0.015	0.147
	(1.65)	(3.72)	(4.33)	(-2.04)	(0.27)	(2.92)

Table shows the output of the regression of stock return turnover ($x_{t+l} = vol_t/shrout_t$) on a dummy variable taking the value 1 if a patent was granted to the firm on day t (Panel A), or the USPTO publicized the grant application of the firm on day t (Panel B). We include firm-year and day-of-week fixed effects. We cluster standard errors by year and report t-statistics in parenthesis. We restrict the sample to firms that have been granted at least one patent.

Table 4: Which firms innovate?

	A^{+}				$\hat{A}$			
A_f	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\ln K_{t-1}$	0.099	0.106	0.116	0.066	0.062	0.067	0.076	0.035
	(42.86)	(43.71)	(46.79)	(39.40)	(39.55)	(40.40)	(43.12)	(34.17)
$\ln Q_{t-1}$		0.052	0.042	0.029		0.035	0.028	0.020
		(19.80)	(15.04)	(17.79)		(19.82)	(14.96)	(20.11)
$\ln RD_{t-1}$			0.103	0.057			0.068	0.030
			(28.53)	(25.24)			(28.38)	(23.07)
$A_{f,t-1}$				0.665				0.742
				(68.24)				(88.33)
Observations	141695	141695	65234	65058	141695	141695	65234	65058
pseudo R^2	0.476	0.490	0.469	0.653	0.644	0.671	0.739	1.242

Table shows Tobit regressions of firm-level innovation A_{ft} on firm characteristics: log firm size (K_{ft}, gross PPE), log Tobin's Q and log R&D expenditures to book assets $\ln RD$. All specifications include year (T) and industry (I) fixed effects. Standard errors are clustered by firm.

Figure 1: Number of Patents with Matched Assignees

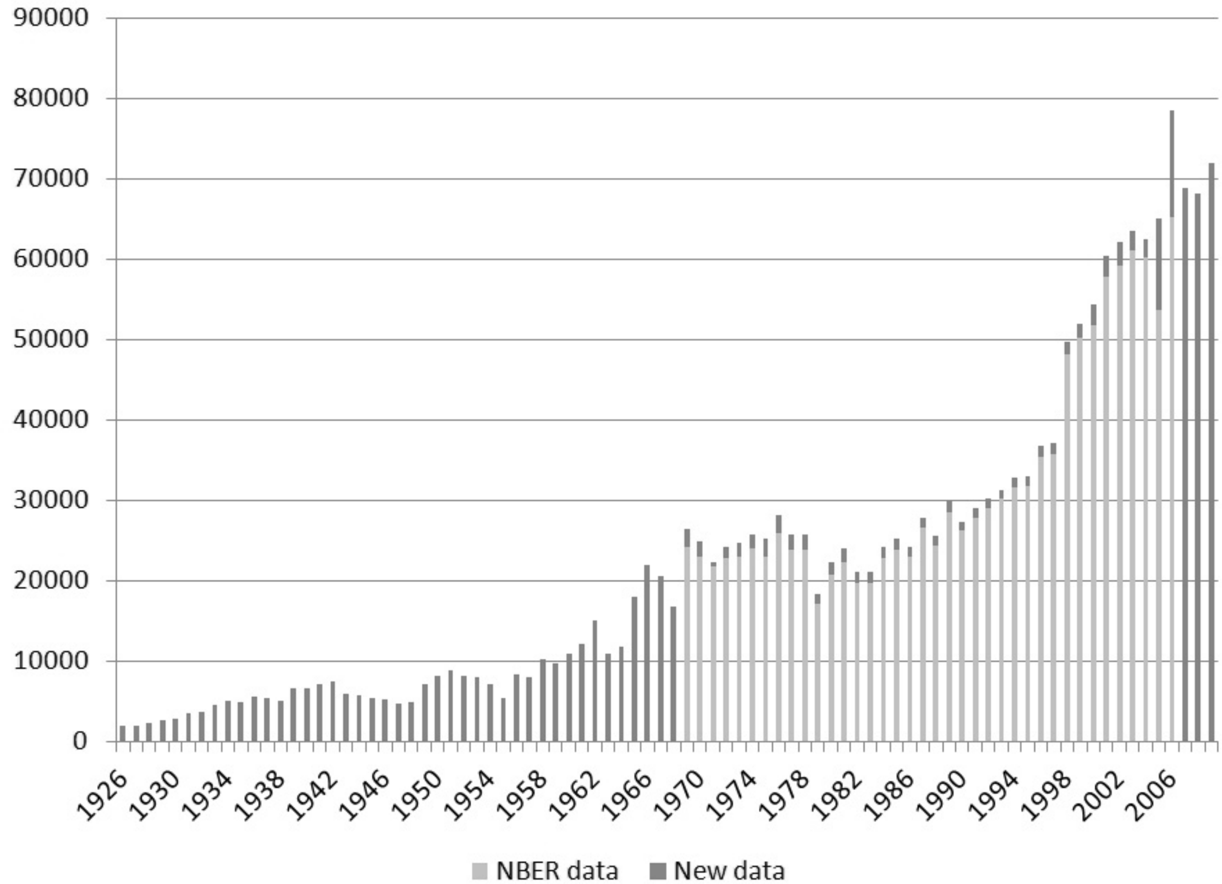

Figure shows the number of patents matched to CRSP firms by year of patent grant. Light shading denotes patents included in the NBER patent data set, while dark shading denotes patents that are new in our paper.

Figure 2: Firm innovation measures – Tail distribution

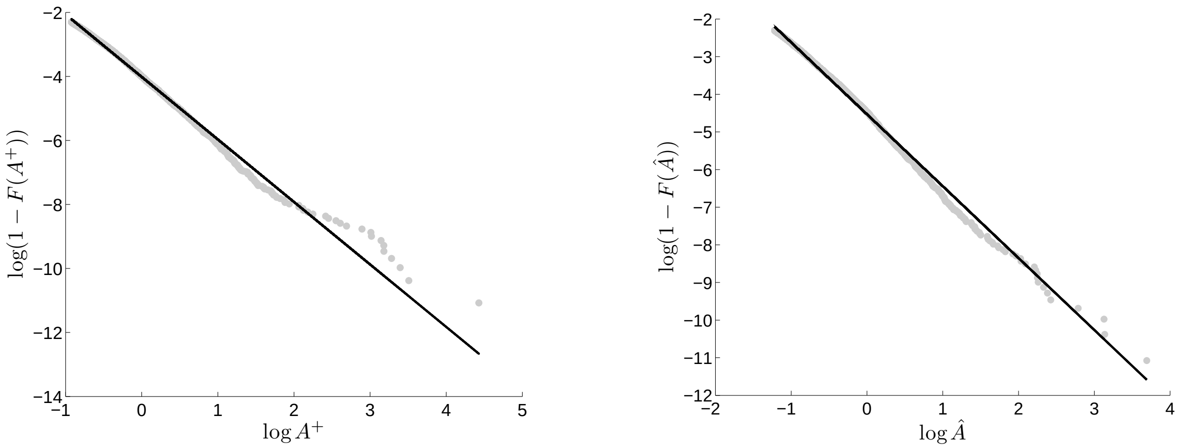

(a) Firm innovation measure A_f^{+}, truncated (b) Firm innovation measure $\hat{A}_f$, adjusted for measurement error

Figure plots the log complementary empirical cdf, $\log(1 - F(A))$, versus the log value of the firm-level innovation measure, $\log A$, for the top 10 percent of the distribution.

Figure 3: Impulse responses: Productivity, controlling for level of stock prices

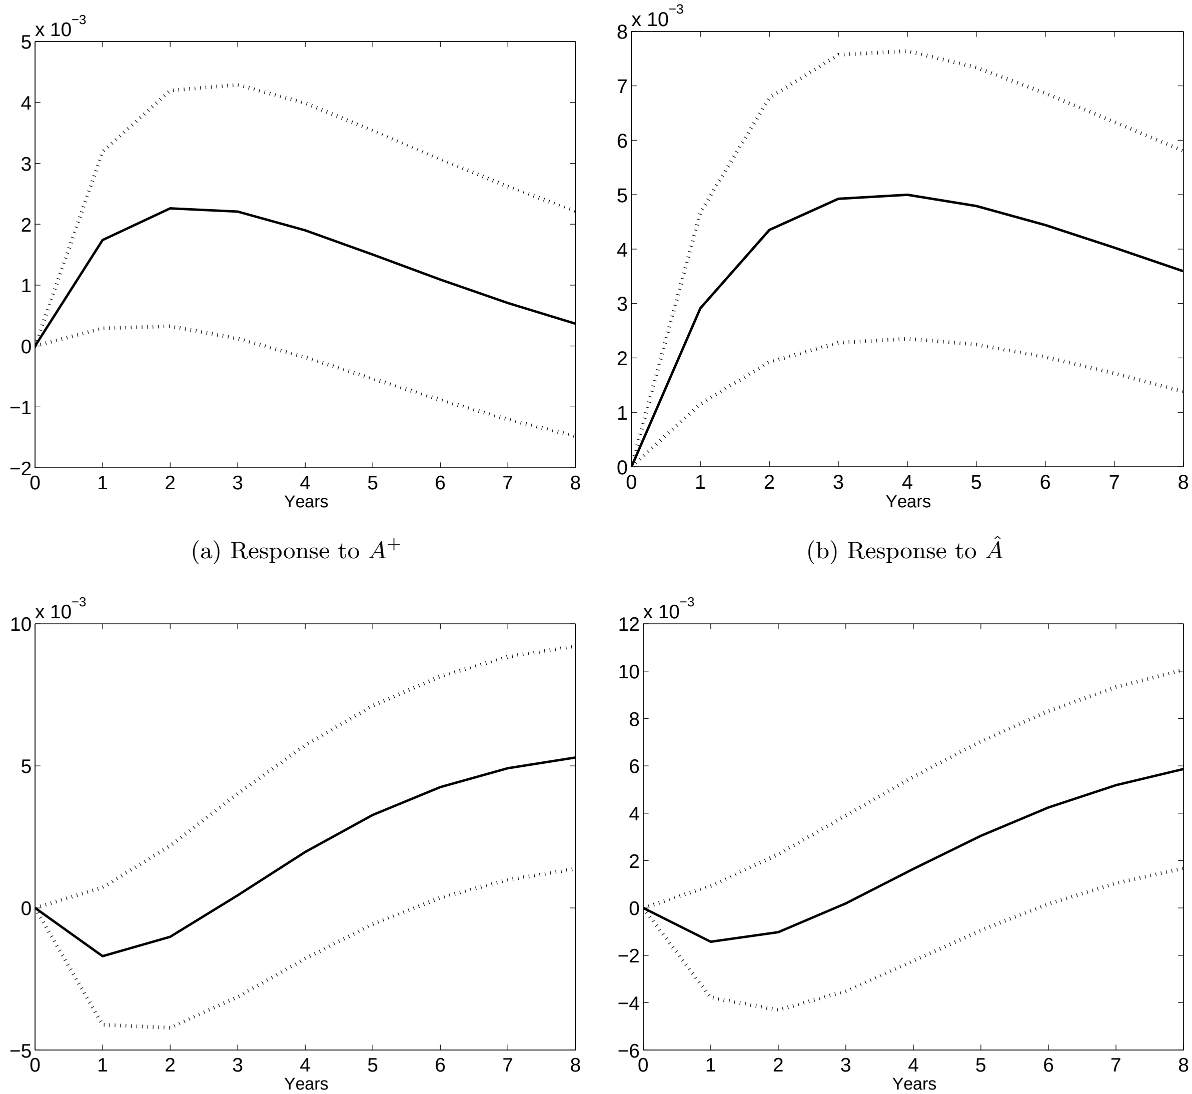

(a) Response to A^+

(b) Response to $\hat{A}$

(c) Response to M

(d) Response to M

Figure shows impulse response of total factor productivity productivity (TFP) from a VAR including $[\log TFP, \log M, \log v, \log A^+]$ (left) and $[\log TFP, \log M, \log \hat{A}]$ (right). Panels a and c present results using our measure of innovation A based on simple truncation A^+. Panels b and d present results using our measure of innovation $\hat{A}$ that accounts for measurement error. Top panel shows impulse response of productivity to our innovation measures. Bottom panel shows response to per-capita real stock market level. We include a deterministic trend in all specifications. We select lag length based on the AIC criterion. In panels a and c we include the log cross-sectional average of idiosyncratic volatility v. Dotted lines represent 90% confidence intervals using standard errors are computed using 500 bootstrap simulations.

Figure 4: Impulse responses, Output, controlling for level of stock prices

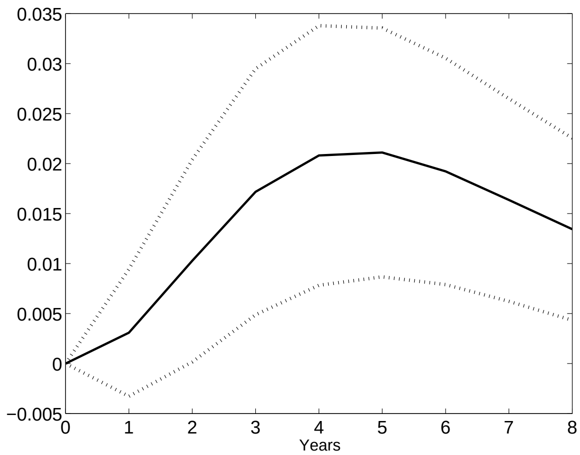

(a) Response to A^+

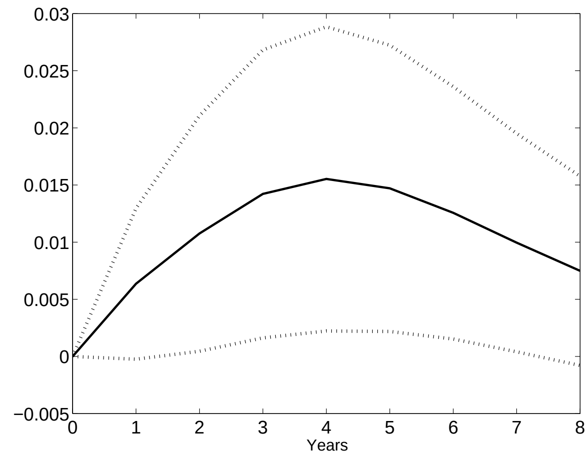

(b) Response to $\hat{A}$

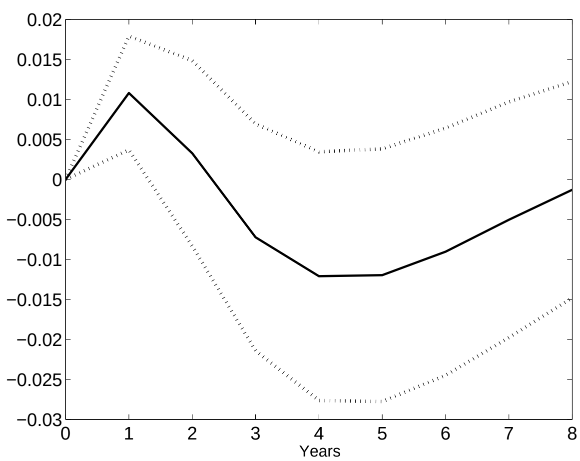

(c) Response to M

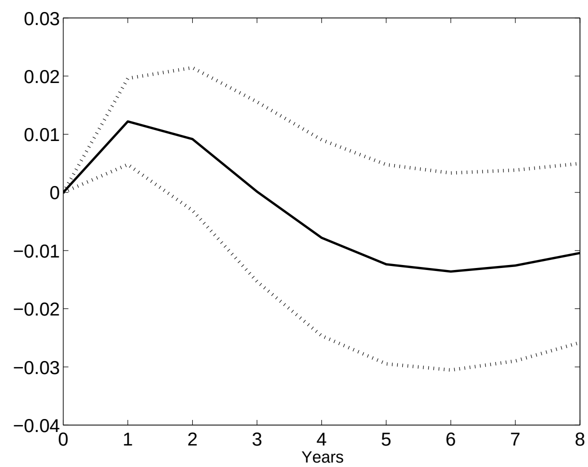

(d) Response to M

Figure shows impulse response of per capita real output Y from a VAR including $[\log Y, \log M, \log v, \log A^+]$ (left) and $[\log Y, \log M, \log \hat{A}]$ (right). Panels a and c present results using our measure of innovation A based on simple truncation A^+. Panels b and d present results using our measure of innovation $\hat{A}$ that accounts for measurement error. Top panel shows impulse response of productivity to our innovation measures. Bottom panel shows response to per-capita real stock market level. We include a deterministic trend in all specifications. We select lag length based on the AIC criterion. In panels a and c we include the log cross-sectional average of idiosyncratic volatility v. Dotted lines represent 90% confidence intervals using standard errors are computed using 500 bootstrap simulations.

Figure 5: Impulse responses using Number of Patents and R&D Stock

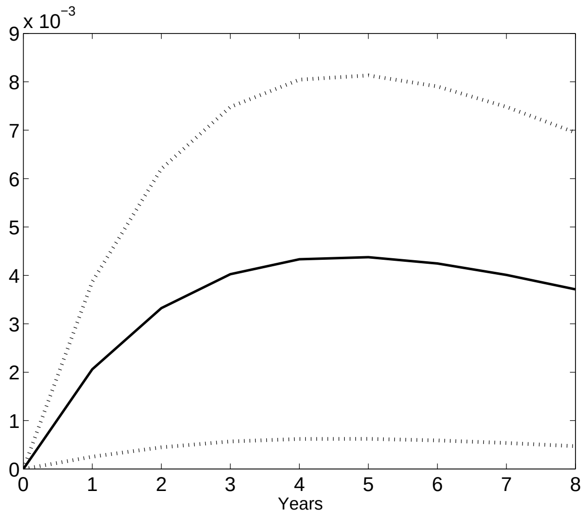

(a) No. patents - Productivity

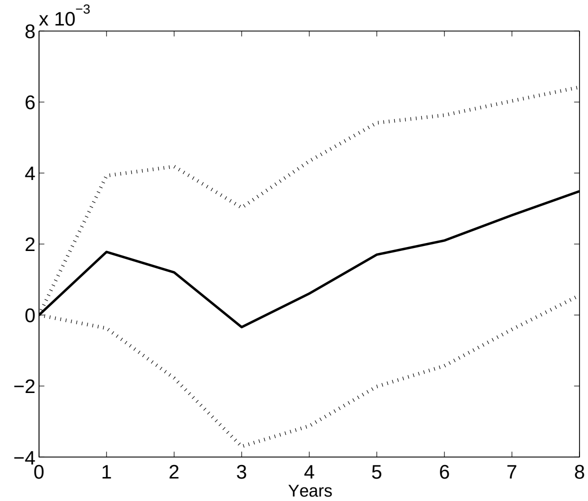

(b) R&D Stock - Productivity

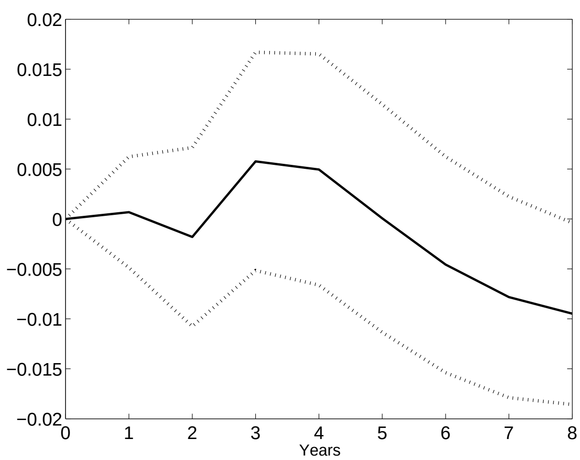

(c) No. patents - Output

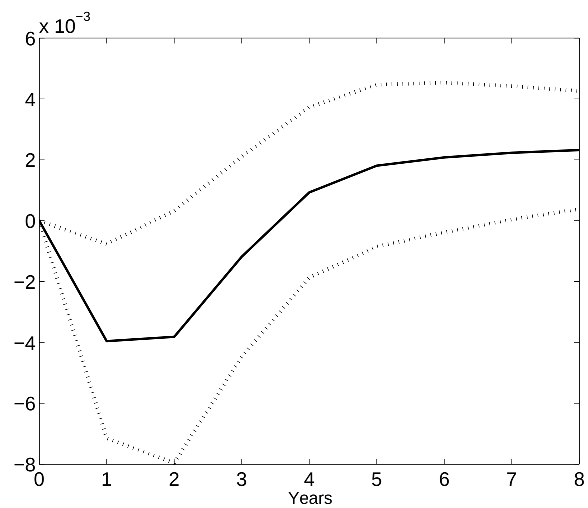

(d) R&D Stock - Output

Left panel shows impulse response of log productivity, output and consumption from a bi-variate VAR $[\log X_t, \log N_t]$ with a deterministic trend. Right panel shows impulse response of log productivity, output and consumption from a bi-variate VAR $[\log X_t, \log RD_t]$ with a deterministic trend. Dotted lines represent 90% confidence intervals.

References

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review 96*(5), 1418–1448.

Hall, B., A. J. and M. Trajtenberg (2001). The NBER patent citation data file: Lessons, insights and methodological tools. Technical report, NBER Working Paper 8498.

Laitner, J. and D. Stolyarov (2003). Technological change and the stock market. *The American Economic Review 93*(4), pp. 1240–1267.

Nicholas, T. (2008). Does innovation cause stock market runups? Evidence from the great crash. *American Economic Review 98*(4), 1370–96.

Norvig, P. (2009). Natural langugage corpus data. In T. Segaran and J. Hammerbacher (Eds.), *Beautiful Data*. O'Reilly Media.