The Quality of Innovation "Booms" During "Busts"

Christos A. Makridis and Erin McGuire *

July 30, 2020

Abstract

Models of creative destruction posit that recessions are periods of reallocation and disruption, generating new ideas that catapult new firms to the frontier. However, empirical evidence suggests that research and development (R&D) expenditures and patenting is pro-cyclical, not counter-cyclical. Using panel data on the quantity and quality of patents for nearly two decades, we provide a resolution to the tension by documenting that the quality of innovation is counter-cyclical: innovations produced during busts have a larger effect on the path of future research than those developed during booms. These results are a function of financial constraints that affect private sector firms most heavily, shifting the concentration of resources and time among R&D workers towards longer-term and basic science research in the public sector during busts. Our results suggest that the ongoing coronavirus pandemic could lead to large innovations in the future.

Keywords: Business Cycles, Financial Constraints, Innovation, Patenting, Research and Development

JEL Codes: L1, O32, O33

*Makridis: Massachusetts Institute of Technology, email: *makridis@mit.edu*, Mcguire: Scheller College of Business, Georgia Institute of Technology. telephone: +1(720) 402-7162, e-mail: *erin.mcguire@scheller.gatech.edu*.
These views are our own and do not reflect those of any affiliated institutions.

1 Introduction

Research and development (R&D) is a fundamental driver of innovation and productivity across firms and countries (Aghion and Howitt, 1992; Jones, 2002). While a common prediction from these models of endogenous growth is that innovation is countercyclical, many have suggested that R&D is actually procyclical (Griliches, 1990; Comin and Gertler, 2006; Barlevy, 2007). Recent theoretical and empirical contributions have reconciled these differences by focusing on strategic reasons for aligning innovation with booms (Shleifer, 1986) and exploring the impact of liquidity constraints during busts (Ouyang, 2011; Aghion et al., 2012; Fabrizio and Tsolmon, 2014). Understanding how large-scale shocks affect the quantity and quality of innovation comes at an especially moment in history as countries press through the coronavirus and research begins to unfold about how the pandemic affects innovation outcomes.

Our paper highlights a new dimension of the cyclicality in innovation, namely the quality of innovation and its potential sources. The quality, specifically the development of general purpose technologies (GPTs), may matter significantly for economic growth and social welfare (Bresnahan and Trajtenberg, 1995). Moreover, one explanation behind the recent productivity paradox (Syverson, 2017) is that lags between the development and diffusion of machine learning and artificial intelligence technologies could account for the sluggish productivity growth in recent years (Brynjolfsson et al., 2019). If the quality of innovation is fairly countercyclical, then the slump of in productivity growth over the past decade could be in part a function of a lag in the development of genuine GPTs. We show that the quality of innovation is indeed countercyclical and results in part through the presence of financial constraints and composition effects in the allocation of time towards R&D activities between public versus private institutions over the business cycle.

The first part of the paper begins by replicating the conventional results about the procyclicality of the *quantity* of innovation before documenting the countercyclicality of the *quality* of innovation. Our identifying variation exploits fluctuations in employment, real GDP, and financial constraints across industries and time, allowing us to compare firms that vary in their exposure throughout multiple business cycles between 1980 and 2015. Our results are robust to a standard Bartik-like instrument that controls for supply-side shocks, focusing instead on industries that are more exposed to certain occupations over others. Moreover, we control for national-level trends in grants and subsidies, like those from the National Science Foundation, that are important for explaining fluctuations in funding towards basic and applied science.

By examining industry trends in patenting and citations in the United States over the period 1980-2018, we find that patents produced during an industry downturn generate more citations in aggregate and have a greater influence on future technologies. Patents produced in a year when industry relative valuation is one percentage point lower have a one percentage point greater textual similarity to future patents and receive 0.34% more citations on average. Although patent production has been previously shown to be procyclical, patents produced during busts provide more inspiration and are more instrumental in the trajectory of future technology. Our results are consistent with Schumpeter (1939)'s view of recessions as times of creative destruction and higher levels of innovation, as well as more recent theory from Manso et al. (2019) that firms pursue greater experimentation (exploitation) during busts (booms), making GPTs countercyclical.

We also compare the similarity of patents produced in booms and busts to those that they cite. If patents produced during booms have less shared text with those that they cite, this may imply that boom patents are larger departures from what has previously been done in the industry. A similar argument would hold if boom patents cite less past work. These greater differences could drive the differences in similarity and citation rates of future work. However, the estimates imply that patents produced during booms have less citations and there are no business cycle-related differences in textual similarity to cited patents. The differences between patents produced in busts and booms is not driven by greater similarity or additional references to past patents.

The second part of our analysis provides evidence for potential mechanisms behind our results, namely the presence of financial constraints among firms in the private sector and composition effects that shift the concentration of R&D activities from applied to basic research during busts. First, using the set of publicly-traded firms between 1997 and 2015, we show that a percentage point (pp) increase in industry employment growth is associated with roughly a 0.91pp increase in the growth rate of R&D expenditures among firms that rank above the median in financial constraints, but only a 0.35pp increase for those that rank below. Furthermore, the cyclicality of R&D expenditures is concentrated exclusively in industries that have higher citation-adjusted patenting rates. Second, using a longitudinal sample of skilled workers between 2000 and 2013, we find that a 1pp increase in occupational employment growth is associated with a null response in the time that non-private skilled workers allocate to basic research and a 0.39pp decline in the probability that they allocate over 10% of their time towards applied research. However, the opposite is true for those in the private sector who exhibit a highly procyclical allocation of time towards applied and basic research. Cumulatively, we interpret these results as suggestive evidence that public R&D generates higher quality innovation during busts at least in part because of the contraction of expenditures and time allocated to R&D in the private sector.

Starting with Schumpeter (1939) who described recessions as periods of "creative destruction" where radical new innovations are created and the technologies they replace become obsolete, there has been a general recognition that innovation *can* be counter-cyclical (Caballero and Hammour, 1994; Aghion and Saint-Paul, 1998; Canton and Uhlig, 1999). In these models, periods of low aggregate demand reduce the opportunity cost of experimentation. An alternative strand of theoretical work argues that innovation does not necessarily have to be counter-cyclical. For example, firms may decide to delay implementation of innovation to periods of high demand (Shleifer, 1986; Barlevy, 2007); firms might be credit constrained during recessions (Aghion et al., 2012); industries might differ in their obsolescence and patent protection (Fabrizio and Tsolmon, 2014); and, individual inventor's might also experience financing constraints as a function of housing wealth shocks (Bernstein et al., 2020). Moreover, many empirical studies contradict the predictions of creative destruction models, finding that R&D spending and patenting is concentrated during booms (Griliches, 1990; Geroski and Walters, 1995; Comin and Gertler, 2006; Koyptov et al., 2018).

Our paper fills an important gap in the literature by focusing on the role of quality in the innovation process and investigating the production of these ideas. First, many studies use patent counts as the main measure of innovation, but that approach may confound differences in the contributions to technological development. Patents can differ greatly in their contributions to technological progress and their significance to development (Kelly et al., 2018; **?**). Moreover, many of these studies use national business cycles as measures of aggregate innovation. The use of national business cycles as a measure of economic conditions makes it impossible to disentangle business cycle effects from time trends and national policies designed to encourage research and development. It is possible that R&D subsidies concentrate in booms, leading to the pro-cyclicality of patenting and R&D. Third, data on research and development is often limited to publicly traded firms. The use of public firm data makes it impossible to determine whether there are different trends in innovation in private firms. It is possible that public firms must emphasize economic conditions more in their decisions regarding research and development strategy, as the value of their firm is tied to the stock market.

This study makes several contributions to the literature on innovation and business cycles. The main contribution is the additional dimension of analysis of the value of innovations at different points in the business cycle. We use a measure of technological similarity developed in **?** to quantify the influence that a patent has on future patents. Thus I am able to show that patents produced during busts are more technologically similar to future patents in the industry. This is the first study we're aware of to analyze the long-run influence of patents on future innovations. Third, we use the universe of patents produced between 1980-2019 and aggregate them to the NAICS-3 industry-year level, which means that we include patents produced by private firms, research organizations, and government agencies. The use of this sample accounts for the possibility that public firms react differently to business cycles, as their profits are more directly tied to the stock market. Finally, our estimation framework relies on industry booms and busts, disentangling the effect of trends in national innovation policies from business cycles.

This study also has important implications for firms and policymakers. If innovation during recessions is more influential than innovation produced at other points in the business cycle, it may be important to use policy to encourage R&D at times when the economy is in a downturn. Specifically, policymakers may want to increase R&D grants and subsidies at periods of low aggregate demand. Then firms will have a lower resource cost of experimentation and the opportunity effect will dominate. Firms also benefit from these results since we show that experimentation during recessions is present and pays off in the long run, which is especially important for publicly traded firms that often face pressure to meet quarterly earnings targets (Terry, 2017).

2 Data and Measurement

2.1 Sources

Industry-level cycles—We proxy for industry level business cycles using aggregate industry-level real output (in 2012 prices) and employment data obtained from the Bureau of Economic Analysis (BEA) for the years 1998-2015. Employment and real output growth are frequently used proxies for business cycle fluctuations, particularly at more disaggregated levels.[1] To mitigate concerns that employment growth is capturing supply-side shocks, we also use a standard Bartik shift-share instrument that exploits the exposure of an industry to employment growth across different six-digit occupations. To measure business cycle conditions at the firm-level, we also follow the approach from Hoberg and Phillips (2010) and aggregate up for our industry-level exercises.[2]

Patent data—We obtain patent data for 1980-2015 from the PatentsView database, which contains data on the universe of patents filed in the United States. Each patent is matched with a measure of the similarity to the previous patents it cited and the patents that build off of work in the patent. We obtain this measure from **?**. They use a vector-space model based on earlier work in Younge and Kuhn (2015) to estimate the technological distance between each citing/cited pair of patents based on the technical description of each patent. Finally, the patents are matched to industries using a concordance developed in Lybbert and Zolas (2014).

Occupation-level Panel (2003-2017)—The Bureau of Labor Statistics releases annual data on employment and wages (across different parts of the distribution) through the Occupation

[1]See, for example, Blanchard and Katz (1992) for one of the first illustrations and Hershbein and Kahn (2018) for a more recent example.

[2]See Appendix A for more details on the estimation of industry dynamics.

Employment Statistics (OES) data tables. While employment is measured at a six-digit level, and there are many new occupations introduced following 2012, we create an employment-weighted average at a five-digit level that matches the identification of occupations in the Census. The key feature of the data is the measurement of employment and income in every occupation.

Individual-level Panel of Educated Workers (2000-2013).—My primary dataset is the U.S. IPUMS Higher Ed data, which is based on a series of surveys of science and engineering graduates by the National Science Foundation since the 1990s, including the Survey of Doctorate Recipients, National Survey of College Graduates, and the Survey of Recent College Graduates. These surveys are combined into a single, integrated set of files called Scientists and Engineers Statistical Data System (SESTAT). The IPUMS Higher Ed dataset made significant strides to create comparable variables and definitions across the different versions of the NSF surveys, which has been used in several recent contributions in the research policy literature (Jung et al., 2017).

Among the variables used from the micro-data are measures of time use in different work activities. While they are not continuous measurements, respondents report a value equal to one if they allocated at least 10% of their time towards the activity in the typical work week. Research activities include either applied or basic research whereby applied research is defined as âstudy directed toward gaining scientific knowledge to meet a recognized needâ and basic research is defined as "study directed toward gaining scientific knowledge for its own sake." Other activities include development ("using knowledge gained from research for the production of materials or devices"), design ("designing equipment, processes, structures, and/or models"), employee relations ("recruiting, personnel development, and training"), management ("managing and supervising"), production ("production, operations, and maintenance"), quality management ("quality or productivity management"), sales and marketing ("sales, purchasing, marketing, customer service, or public relations"), professional services ("health care, counseling, legal services, etc"), teaching, and finance ("accounting, finance, or contracts").

Occupational Task Concentrations (O*NET).—To measure R&D tasks at a detailed six-digit occupational level, we use O*NET, which is the new companion to the well-known Dictionary of Occupational Titles (DOT) used in prior work (Autor et al., 2003). O*NET is a survey that the U.S. Department of Labor administers to a random sample of U.S. workers within detailed occupations. Respondents answer questions on an ordinal scale that measures both the importance of a task and the frequency at which different tasks occur on the job. The data has typically been used to measure the intensity of cognitive, non-routine, non-cognitive, routine, and manual skill intensities across occupations (Autor et al., 2003; Autor and Dorn, 2013; Autor et al., 2008).

Using additional information on the universe of over 10,000 task statements in O*NET, I create an occupational index of R&D intensity by extracting all task statements containing either the words "research" or "develop." Of course, there are examples of task statements that include the word "innovation," but these examples are few and far between.[3] Following prior work by (Autor and Handel, 2013), we take the product of the importance and frequency weights (when available) to generate an overall intensity for each sub-index task and construct a standardized z-score of their sum for each occupation $\times$ year.[4] Since there are many tasks that meet these requirements, the footnote documents several as examples and the full list will be posted on my website for downloading and experimenting.[5] While the advantage of this measure is that it

[3] In fact, only five task statements contain the prefix "innovat-" in the 2004 file. To illustrate that even inclusion of the prefix would not necessarily produce a more reliable overall R&D index, consider one such task statement: "Monitor the field of dance to remain aware of current trends and innovations."

[4] Hagerty and Land (2007) find that using equal weights over subset variables for these types of indices provides the greatest robustness and accuracy. The issue of potentially different weights is relatively minor when all are measures are positively correlated.

[5] A couple examples include: "Direct and conduct studies and research on issues affecting areas of responsibility"; "Conduct research to develop methodologies, instrumentation and procedures for medical application, analyzing

represents the task concentration of RD activities across the full set of occupations over time, its primary limitation is the lack of within-occupation variation. For example, pooling together 2004 to 2016 data, a regression of R&D intensity on occupation fixed effects produces an R-squared of 0.92, suggesting that the bulk of the variation is in the cross-section. Moreover, not all occupations are updated each year, which means that attenuation bias could result from measurement error in the the within-occupation task concentration variation. That leads us to use the measure primarily as a diagnostic and for classifying individuals in the Census as high or low R&D workers.

As an illustration of the variation contained in these data, Figure 1 plots the distribution of R&D task intensity for 2005-06 and 2017-18. For example, the mean z-score in 2015-06 is -0.14, whereas it is 0.11 by 2017-18. Moreover, R&D task intensities are more skewed to the right in 2017-18. These differences in the distribution suggest that R&D tasks have, on average, become more intense over time, which is consistent with evidence from Bloom et al. (2020) about the increasing intensity of R&D activities. Unlike other attempts to measure R&D activities, this approach focuses on the allocation of time among the workers in each occupation conducting the actual R&D tasks. These occupational task intensities allow us to classify occupations as either high or low R&D by partitioning off of the median z-score. That is, all occupations with R&D task intensities above the median are classified as R&D-intensive, zero otherwise.

[INSERT FIGURE 1 HERE]

data and presenting findings"; "Plan, supervise, and conduct psychological research, and write papers describing research results"; "Collaborate with colleagues to address teaching and research issues"; "Research and report on specialized fields such as medicine, science and technology, politics, foreign affairs, sports, arts, consumer affairs, business, religion, crime, or education." Although these examples point towards a broader definition of R&D that is not restricted purely towards traditional science and engineering occupations, they still capture the fundamental feature of R&D as an activity devoted towards invention.

2.2 Annual Industry Patent Summary Statistics

The first part of the analysis makes use of a panel at the NAICS-3/year level. The key variables are the measures of industry economic conditions, the total number of patents filed in that year (a weighted sum from the concordances), the total number of past patents cited, the total number of forward citations generated, the textual similarity of filed patents to their citations, and the textual similarity of filed patents to the citations they generate.

[INSERT TABLE **??** HERE]

Table **??** contains the summary statistics of the final industry-year panel. The average industry relative investment and industry relative valuation are similar to estimates in Hoberg and Phillips (2010). We estimate that the average industry has a 0.3% greater valuation than predicted in any year.[6] Investment is on average 2.97% greater than predicted, which is similar to the average 2.2% in Hoberg and Phillips (2010). Our minimums and maximums are greater, but this is because the years in this sample include more of the Great Recession and include the growth in venture capital investment between 2010 and the present.[7]

On average, there are 1831 patents filed in each year in a three digit NAICS category, but the distribution of patents over industries is highly skewed to the right. The median number of patents filed in a given industry-year is 196. Each patent cites 15 previous patents and generates 17.55 citations on average. Both citing and cited patents have a measure of textual similarity of 0.26 to the relevant patent.

[6]The mean relative valuation is -0.3% in their paper, but is not statistically significantly different from these results. Hoberg and Phillips also use the years 1980-2010 for their sample, while we have the additional years 2010-2015.

[7]https://news.crunchbase.com/news/there-are-more-vc-funds-than-ever-but-capital-concentrates-at-the-top

2.3 Annual Research and Development Worker Summary Statistics

Table 2 documents descriptive statistics about the individuals in the NSF sample, distinguishing between public and private sector workers. Beginning with demographic characteristics, males are more likely to sort into private sector jobs (60 versus 47 percent), but there are no statistically significant differences in the average age, share of whites, and number of children in the family. There are, however, substantial differences in educational attainment. For example, public sector workers are more likely to have a doctorate (9 versus 3 percent) or masters degree (38 versus 25 percent), but less likely to have a professional degree (6 versus 10 percent). These differences in the share of workers with different educational degrees reflect, at least in part, differences in the taste for science (Roach and Sauermann, 2010, 2014).

Turning towards the workplace characteristics, private sector workers earn considerably more than their public sector counterparts with an annual salary of $72,119 (versus $56,214 in the public sector). Moreover, the standard deviation in earnings is larger for the private sector, reflecting the fact that there is greater heterogeneity and flexibility in compensation schemes. Private sector workers also work longer hours than their counterparts and are more likely to reside within a "career paths jobs." This is consistent with past evidence that private sector workers are more likely to value financial incentives and career mobility over other non-pecuniary job characteristics (Bozeman and Kingsley, 1998; Bozeman, 2000).

To better understand these other job-specific features, the survey also contains questions asking respondents to rank their satisfaction with job characteristics, such as career opportunities and benefits, together with their overall perception of the job, all on a scale of one to four. Table 2 suggests that, whereas private sector workers report overwhelmingly greater career opportunities and financial (salary) compensation with averages of 0.02sd and 0.06, public sector workers report averages of -0.07sd and -0.14. However, private sector workers rank much lower in their reported benefits, job security, and perceived social impact. The fact that public workers earn less, but experience these other positive job-specific amenities, reflects the presence of compensating differentials in the market for scientists (Stern, 2004).

Finally, turning towards measurement of the job activities where students allocate over 10% of their time per week, there are also a number of stark differences. For example, development, design, and sales are also systematically higher among private sector workersâin fact, 40% of private sector workers report allocating over 10% of their time to sales activities in a week, relative to only 19% among public sector workers. These differences in time use reflect differences in the underlying tasks and output among workers in these sectors (Autor and Handel, 2013).

[INSERT TABLE 2 HERE]

3 Business Cycles and the Quality of Innovation

3.1 Background

There are two opposing effects that recessions can have on research and development activity. The first effect is the "supply of resources effect." When an economy experiences a recession, capital markets tighten and aggregate demand drops. It is more difficult for firms to obtain external financing and firm profits often drop. This means that firms have fewer financial resources to devote to research and development. Due to these changes in the market, research and development will increase in booms when finance is abundant and decrease in recessions when markets contract.

On the other hand, the "opportunity cost effect" posits that the cost of R&D also decline in recessions. First, the cost of inputs to research decreases. Labor market wages, interest rates, and the cost of materials all decrease during a recession. Second, the opportunity cost of research decreases. As profits often decline during recessions, the cost of diverting effort from production to research is lower in recessions. When aggregate demand increases in a boom, the benefit of devoting more time and resources to production is higher, so firms do less research and development in booms. Under the reduction in costs effect, innovation will be counter-cyclical.

The opposing effects imply the correlation between business cycles and the intensity of innovation is ambiguous. If the resources effect dominates the opportunity cost effect, R&D will concentrate in booms and decline in recessions. If the reduction in opportunity cost effect dominates the resources effect, innovative activities will increase in recessions and decline in booms.

While there is a large literature about the cyclicality of innovation and its various causes, there is much less research on the quality of innovation over the business cycle. Building on Schumpeter (1939) who argued that novel innovations concentrate at the bottom of a cycle, Manso et al. (2019) develop a two-armed bandit model that predicts firms are more likely to pursue novel exploration strategies in recessions. Their main theoretical result stems from the fact that firms find it less costly to pursue exploratory research during busts, whereas the marginal opportunity cost of a dollar of innovation on exploitation is lower during booms. We develop this further by exploring the cyclicality in the quality of innovation using similar data on patents across industries and time.

This trade off between experimentation and exploitation stems from the diffusion of product innovation, detailed in a five-step process by Gort and Klepper (1982). The first stage is initiated by the introduction of a new product by a producer and ends when new competitors begin to rapidly enter the market. The number of competitors continues to increase in stage two. Stage three occurs when the number of entrants is roughly equal to the number of firm exits. In stage four, the number of exits dominates the number of entrants. The final stage is defined as a second period of zero net entry and the obsolescence of the product that could be driven by, for example, a large in advance in a new technology.

Dating back to Burns (1934), a key assumption is the degree of product innovations decreases over the product life cycle. In the beginning of stage two, a larger fraction of the innovations by competitors have major consequences for product quality and the production process. As the product cycle continues, innovations become increasingly minor since opportunities for dramatic improvements decrease. Minor innovations, such as quality control or marketing, could be implemented during later stages of the cycle, decreasing the overall rate of technological advance.

[INSERT FIGURE 2 HERE]

When the implications of product life-cycle theory are combined with the insight that more novel innovations are produced during recessions, we can predict their influence relative to inventions at other points in the business cycle. The rate of technological advance should increase in recessions and decrease in booms: firms choose to research more novel technologies during recessions. These novel technological developments spark new product life cycles, leading to increases in market entry and imitation. In booms, firms shift resources to production and make small iterations on product technology because the opportunity cost effect dominates the resource effect. These less-novel innovations are consistent with the later stages of the product life cycle, meaning that the rate of technological advancement will decline. Thus innovations in recessions will have a greater influence on future technologies than innovations in booms.

3.2 Empirical Specification

We estimate fixed effects regressions that vary the quantity and quality of patents with a proxy for business cycle fluctuations, conditional on industry and year fixed effects. We measure innovation outcomes in three ways: the logged number of patents, logged number of patents cited, the logged number of future citations generated per patent, and the textual similarity of past patents cited and future citations to the patents filed in a given industry-year.[8] We also proxy for business cycle fluctuations in four ways: year-to-year growth in real GDP (2012 prices), year-to-year growth in employment, and both industry relative valuation and industry relative investment as defined in Hoberg and Phillips (2010). The relative valuation measure is the average difference in predicted firm valuation from actual valuation at the industry level. For example, if relative valuation for manufacturing in 1998 is 0.05, this implies that manufacturing firm valuations were on average 5% above predicted values in 1998. The interpretation of relative investment is similar.

Our inclusion of industry and year fixed effects purges heterogeneity in patenting that is time-invariant, capturing, for example, the fact that some industries are more innovative than others. The year fixed effects enable us to control for aggregate trends in patenting and national level business cycles. While we have also experimented with demographic controls, like the share of college graduates, to address concerns about changes in the composition of workers, we recognize the potential that changes in patenting might also reflect supply-side shifts. We, therefore, construct a Bartik-like measure of employment growth using three-digit industry $\times$ six-digit occupation em-

[8]The textual similarity measures come from Kuhn et al. (2020). They use a machine learning algorithm to create a measure of similarity across patent text. We use the measures of similarity for each pair of patents linked together by a citation.

ployment counts between 2002 and 2018 from the Occupational Employment Statistics. Following Goldsmith-Pinkham et al. (2019), we fix our base year to 2004 prior to the Great Recession and exploit the exposure of industries to different occupational concentrations.

3.3 Main Results

Table 3 documents our main results. Start with the first four columns, we see that there is some negative association between fluctuations in economic activity and the overall number of patents when we use the relative valuation or investment as a proxy. Here, a 10% rise in valuation and investment is associated with a 3.4% and 1.1% decline in the number of patents, suggesting that the quantity of innovation is counter-cyclical. However, when we use industry employment or real GDP growth as a proxy, we find less economically and statistically significant evidence. Each specification includes industry and year fixed effects.

[INSERT TABLE 3 HERE]

One limitation with these results, however, is that they fail to control for heterogeneity in the quality of innovation. That is, simply showing that there is some evidence of declines in patenting during economic booms says nothing about the impact of the innovation. In order to measure the level of impact that technologies produced have on future innovation flows, we construct measures of the number of times a patent is cited in the years after it is published and its textual similarity to future patents. We then aggregate these measures up to the industry-year level to measure the average impact of innovations produced during industry booms and busts. We now explore the contribution of patents to the future flow of innovation in an industry in columns 5-8. Here, we find even stronger evidence of counter-cyclicality, although only when we use employment and real GDP growth as our proxy for economic conditions. Under these preferred specifications in columns 7 and 8, a 1pp rise in employment and real GDP growth is associated with a 0.35% and 0.47% decline in the flow of future patent citations. We find similar results when we use patents cited as the outcome variable, but the estimates are less statistically significant.

Are these patents similar to those that have already been produced? In particular, how are business cycles associated with the novelty of an innovation, independent of citation counts? Columns 13 to 16 investigate the effects of business cycle conditions on the similarity of patents to those that have been cited before. We find that a 1pp rise in employment and real GDP growth is associated with a 0.44% and 0.39% decline in patent similarity. This augments our earlier result that the quality of innovation is increasing during recessions, pointing out that these innovations are also unique and different than prior patents.

[INSERT TABLE 3 HERE]

4 Understanding the Mechanisms

4.1 Exploring the Possibilities

We now turn towards the underlying mechanisms behind the countercyclical evidence of quality in innovation. We focus on three potential explanations.

First, if the firms that are more likely to experience liquidity constraints during busts are also less productive, then the quality of their R&D and other innovation efforts might also be lower for other unobserved reasons. For example, motivated by evidence from Terry (2017) that many publicly traded firms underinvest in R&D to meet quarterly earnings targets, these companies might also display lower levels of productivity overall.

Second, if R&D workers are less likely to experience job loss over a business cycle, then they might be less risk averse and more likely to invest in longer-run projects that ultimately pay large dividends during booms. For example, Acemoglu et al. (2018) develop a model of endogenous entry and exit, hiring skilled workers to perform R&D. To the extent these skilled workers are less likely to experience layoffs over the business cycle, they might be more suited to produce general purpose ideas that benefit their firms during booms.

Third, there could be a selection effect in the class of organizations that participate in innovation. For example, if universities process funding more slowly and/or obtain longer-run grants (Lane et al., 2015), then they might have a comparative advantage for innovating during recessions. Moreover, there is suggestive evidence that government funding does not crowd out private investment for R&D (Becker, 2015), meaning that the complementarity could enable public and non-profit organizations to innovate more during recessions.

4.2 Examining the Evidence

We begin by investigating the potential role of financial constraints. While others have historically recognized the role of financial constraints (Aghion et al., 2012), they have done so in the context of explaining the procyclicality of R&D, rather than the counter-cyclicality of innovation quality. Using data from Compustat, we regress year-to-year growth in R&D expenditures on year-to-year three-digit industry employment growth, allowing for separate treatment effects among firms that are more versus less financially constrained, conditional on firm and year fixed effects. This allows us to exploit within-firm variation in R&D intensity over the business cycle, tracing out the response of R&D to different industry shocks from over a decade. We measure financial constraints using the firm-year index introduced by Hoberg and Maksimovic (2015) who apply textual analysis techniques to financial statements among publicly traded firms.

Table 4 documents these results. Not surprisingly, we find that R&D expenditure growth is highly procyclical: a 1pp rise in industry employment growth is associated with a 0.74pp increase in R&D growth (column 1). Once we add firm controls, such as logged revenue, employment, and capital, the coefficient declines slightly to 0.69, but is statistically indistinguishable (column 2). However, once we allow for heterogeneity among firms that are more financially constrained, we find an even larger associated on the interaction: a 1pp rise in industry employment growth is associated with a roughly 0.90pp rise in R&D expenditure growth among the financially constrained firms, but only a 0.35pp rise among those that are not constrained.

While these results are consistent with the presence of financial constraints affecting the quantity of innovation, they are not fully connected with our focus on the quality on innovation, particularly if there are composition effects in the set of firms performing R&D over the business cycle. To make a link, we classify industries as having high versus low levels of citation-weighted patenting by creating an indicator for whether an industry ranks above the median in its 1998 to 2018 average future citation counts to total patents, zero otherwise. Table 4 subsequently partitions our estimates in columns 4 and 5 based on these two group. Importantly, all of the effect is coming from industries with higher quality patents; the elasticity is not statistically significant for other industries ranked below the median in patent quality. This suggests that financial constraints matter for not only the quantity of innovation, but also the quality.

[INSERT FIGURE 4 HERE]

We now explore whether reallocation between R&D workers into and out of the labor market could explain differences in the cyclicality of innovation activities. While we would ideally like to track the same worker over time and observe their output over the business cycle, these data are not available. Instead, we exploit variation in employment growth over the business cycle, differentiating between researchers and non-researchers at a six-digit occupation $\times$ CBSA level between 2005 to 2018. We measure researchers in two ways: using (i) an indicator for whether the occupation ranks above the median in the earlier index of R&D intensity constructed from O*NET, and (ii) an indicator for whether the occupation is classified as a science, technology, engineering, and mathematics (STEM) job from the Bureau of Labor Statistics. Specifically, we consider regressions of occupation $\times$ CBSA year-to-year employment growth on an indicator for whether the occupation is classified as research, CBSA year-to-year employment growth, and their interaction, conditional on at least CBSA, occupation, and year fixed effects.

Table 6 documents these results. While we find that increases in CBSA employment growth are associated with robust increases in occupational employment growth in the same CBSA area, there is no evidence of an asymmetry in the cyclicality for either high R&D intensity or STEM jobs. Columns 1 and 3 present the results using fixed effects on occupation, CBSA, and year, whereas columns 2 and 4 also control for occupation $\times$ year fixed effects, thereby exploiting the exposure of a given occupation to different regional shocks. The latter approach would control with time-varying shocks to skill prices that could be unique to occupation, but idiosyncratic across locations. Given that we find no evidence of asymmetric cyclicality, we conclude that it is unlikely that higher skilled workers would invest more in longer-term projects simply due to differences in their sensitivity to the business cycle.

[INSERT TABLE 6 HERE]

Finally, we turn towards evidence on the potential for selection effects between public and private institutions. Drawing on the NSF data covering high skilled workers, we are reminded of the stark differences between public and private institutions in the allocation of time towards research displayed in Table 2. For example, whereas 27% and 17% of workers in the private sector allocate over 10% of their time per week towards applied and basic research, these proportions are 33% and 25% in the public sector. In this sense, although the allocation of time towards applied research only varies by a few percentage points, public sector workers are over twice as likely to allocate more than 10% of their time towards basic research.

Motivated by these differences in time use, we investigate whether they also vary over the business cycle by estimating similar regressions of the form:

$$l_{iot}^{k} = \gamma z_{ot} + \phi PRIV_{iot} + \delta(z_{ot} \times PRIV_{iot}) + \omega w_{it} + \beta X_{iot} + g(t) + \epsilon_{it} \tag{1}$$

where l^k denotes whether individual i in five-digit occupation o in year t allocated over 10% of their time towards activity k, z denotes a productivity shock (i.e., five-digit occupation employment growth), $PRIV$ denotes whether the individual works in the private sector (versus a two or four year university or government agency), w denotes the individual's logged salary, X denotes a vector of individual covariates, and $g(t)$ denotes a polynomial time trend. Included within the vector of individual covariates, X, are the following variables: a quadratic in age, an indicator for being male, an indicator for being white, and education fixed effects (bachelors, masters, professional, normalized to having a PhD/MD). These covariates help mitigate concerns about self-selection into different types of jobs (e.g., risk preferences or career aspirations Bozeman and Kingsley (1998)) that also vary in their sensitivity to business cycle shocks. We measure z using year-to-year growth in five-digit occupational employment, capturing the heterogeneity in returns to different tasks over the cycle.

While not included in the baseline regression results, there are also other controls that are noteworthy to mention for robustness, especially hours worked and firm size. For example, one concern is that employees simply work longer hours during booms, giving the impression of reallocation when really the intensive margin of labor services is expanding. However, controlling for fixed effects on weekly hours worked and number of weeks worked (four bins each) produces quantitatively indistinguishable coefficient estimates. Similarly, another concern is that larger organizations are less sensitive to business cycle shocks, but individuals move between smaller and larger organizations over the cycle. While job mobility is generally common, it is less common among high-tech workers who often have considerable firm-specific investments. These considerations are examined later when studying the moderating effects of financial constraints. Since our equation is estimated using a weighted probit in the baseline specification, occupational fixed effects are not included (Lancaster, 2000). While the private and public sectors produce different kinds of output, we assume that time use at a task-level is comparable between the two sectors.

One of the potential concerns with our strategy is that it might fail to fully control for composition effects. Given that organizations tend to upskill during recessions (Hershbein and Kahn, 2018), the return to a marginal hour of work among private sector organizations may be higher if these organizations are also more likely to terminate less productive workers. Then, greater procyclicality in the time allocated towards R&D activities may simply reflect a shifting composition towards more productive workers. While this is possible, the fact that an individual's salary is included as a control helps mitigate concerns about composition effects. Nonetheless, one way of addressing this identification problem involves the inclusion of person and year fixed effects, exploiting variation in a fixed individual exposed to different labor market conditions.[9]

Table 5 documents the main results. Consistent with the first proposition, private sector (business) workers allocate 3.4-4.5% less time in applied research and 7.9-8.1% less time in basic research activities. Both basic and applied research respond similarly to cyclical shocks, consistent with macroeconomic models with endogenous growth (Akcigit et al., 2019). In contrast, private sector workers, not surprisingly, allocate much more time towards standard business activities, including 5.8-7.6% more time in development activities, 10.1-13.9% in design, 6.1% in production, and 14% in sales. These results are consistent with research on sorting into public versus private sector jobs based on a willingness to pay for autonomy and conducting research (Stern, 2004; Roach and Sauermann, 2014).[10]

Time allocated towards research in the private sector is overwhelmingly procyclical, whereas it is countercyclical or possibly acyclical in the public sector. For example, a one percentage point (pp) rise in occupational employment growth is associated with a 0.392pp decline in time allocated towards basic research in the public sector. The coefficient only declines slightly in magnitude to 0.33pp after introducing salary as a control. Time allocated towards basic research is also relatively countercyclical, but I fail to reject the null that it is acyclical in the public sector. However, a one pp rise in employment growth is associated with a 0.255pp (=0.647-0.392) to 0.308pp (=0.638-

[9]Estimating a non-linear model with individual fixed effects introduces the incidental parameters problem (Lancaster, 2000). However, when estimating using a linear probability model, we obtain qualitatively similar results. For example, time allocated towards research activities is still procyclical among private sector workers, but countercyclical among public sector workers. However, these marginal effects are not statistically significant.

[10]One concern with these result is that employees do not have discretion over their allocation of time. While high skilled workers are more likely to have discretion since they cluster in occupations that exhibit greater autonomy and non-cognitive tasks, whether the employee or employer decides the set of tasks to prioritize is not germane to the empirical resultâall that matters is that there is some cyclicality.

0.33) increase in time allocated towards applied research without and with salary as a control. Time allocated towards basic research in the private sector is also highly statistically significant with a precise marginal effect of 0.31 for a 1pp rise in employment growth.

Other business-based activities exhibit countercyclical time use among R&D workers. For example, a 1pp rise in employment growth is associated with a 0.248pp to 0.340pp decline in time allocated towards development activities in the public sector. Similarly, a 1pp increase in employment growth is associated with a 0.689pp to 0.806pp decline in time allocated towards design activities in the public sector, but only a 0.203pp (=-0.689+0.486) to 0.32pp (=-0.806+0.486) decline in the private sector (when salary is not included as a control). Moreover, production activities are actually procyclical in the private sector with a net gradient of 0.067pp (0.449-0.382) and 0.078pp (0.453-0.375) increase for a one pp rise in employment growth.

How should we interpret these results? On one hand, these heterogeneous cyclicalities of time use might reflect reallocation of tasks within organizations over the business cycle. On the other hand, they might reflect differences in the composition of workers in the firm. While the inclusion of salary and individual demographic controls help rule out concerns about the strength of incentives over the cycle, it also behaves as an implicit control for ability, on top of the inclusion of detailed demographic controls. Nonetheless, there are two points worth noting. First, while there is empirical evidence that organizations upskill over the business cycle (Hershbein and Kahn, 2018), unobserved shocks to the demand for higher skilled workers would, if anything, bias us against finding these results since there is, at the margin, slight negative selection into the private sector among knowledge workers (Stern, 2004). Second, one way to gauge the role of composition effects is by comparing the marginal effect of the interaction with and without controls. Compared to the baseline specification, which suggests that a 1pp rise in employment growth is associated with 0.255, without the controls the corresponding increase is 0.104pp (=0.485-0.381), which is consistent with the aforementioned direction of potential bias.

[INSERT TABLE 5 HERE]

Put together, we interpret these results as providing evidence of a composition effect in the inputs to the production of knowledge over the business cycle, which results at least in part due to financial constrained faced more acutely in the private sector. For example, our first piece of evidence suggests that financially constrained firms contract their R&D more than their counterparts, particularly in response to employment declines. These effects are coming almost exclusively from industries with higher quality innovation. Our second piece of evidence suggests that private sector workers allocate their time more elastically over the business cycle, particularly in business activities, whereas public sector workers exhibit more acyclical responses. The combination of a reallocation of tasks among individuals in public and private firms has the potential to explain the counter-cyclicality of innovation. We find no evidence of the third potential mechanism, which speculates that employment fluctuations might affect researchers more (the extensive margin of innovation) and make them even more risk averse (the intensive margin of innovation).

5 Conclusion

Though researchers have found that patent counts or R&D expenditures are procyclical, this prediction is at odds with the theory of creative destruction where firms invest more resources in research during periods of lower demand. It is possible that these results are in contrast because resources are also constrained in recessions, which limits the amount of capital the firm can devote to research. If the “resources effect” dominates the “opportunity cost effect,” the empirical results that innovation is procyclical are logical. The number of patents may decrease during recessions due to resource constraints, but the type of research that firms do may change.

We investigate the relative importance of innovations created during recessions and booms, exploiting within-industry variation in business cycles, rather than national business cycles, enabling us to disentangle the effect of aggregate changes in funding for innovation (resources) from effects of shifts in aggregate demand. We find that the number of patents is not correlated with industry dynamics after controlling for changes in innovation funding at the national level. However, the value of patents in future innovation does change depending on industry conditions: more influential innovations are produced during industry busts. We find that patents produced during busts share a greater similarity to future patents produced in the industry and that patents produced during busts generate a greater amount of citations than those produced during booms. This suggests that although resources are constrained, firms may work on riskier projects with potentially higher returns because the opportunity cost of failure has decreased. Furthermore, these results are driven by a combination of financial constraints and composition effects where private (public) sector R&D inputs are more (less) cyclical and more (less) focused on applied research.

Tables and Figures

Figure 1: Distribution of R&D Task Intensity, 2004-5 and 2015-16

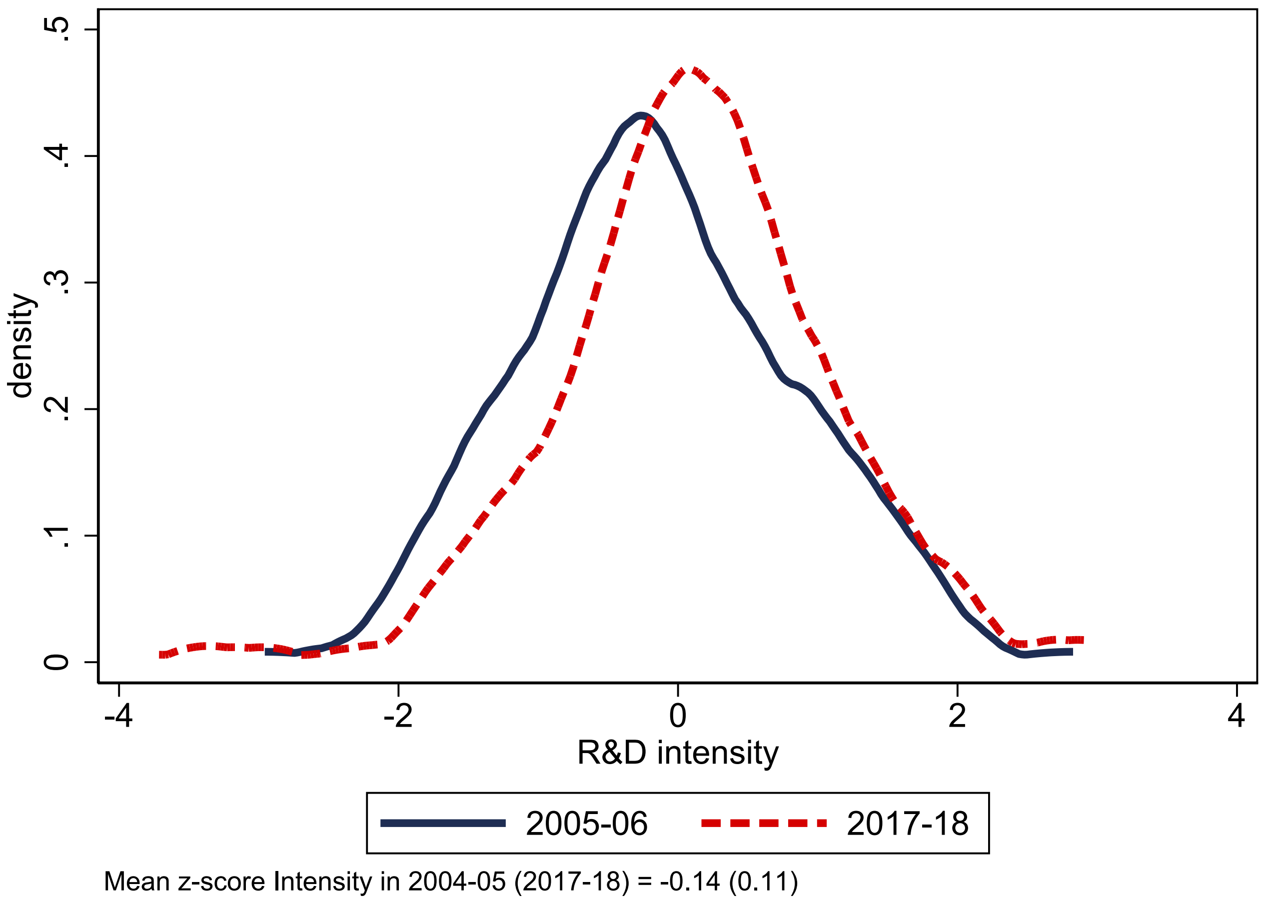

Source: Notes.—Sources: O*NET (2004-2016). The figure plots the distribution of research and development (R&D) task intensity in a standardized z-score format. The index is based off of the "importance" (IM) weight in O*NET collected for all task statements that have either the words "research" or "development" in their description.

Figure 2: Conceptual Map of the Product Cycle over the Business Cycle

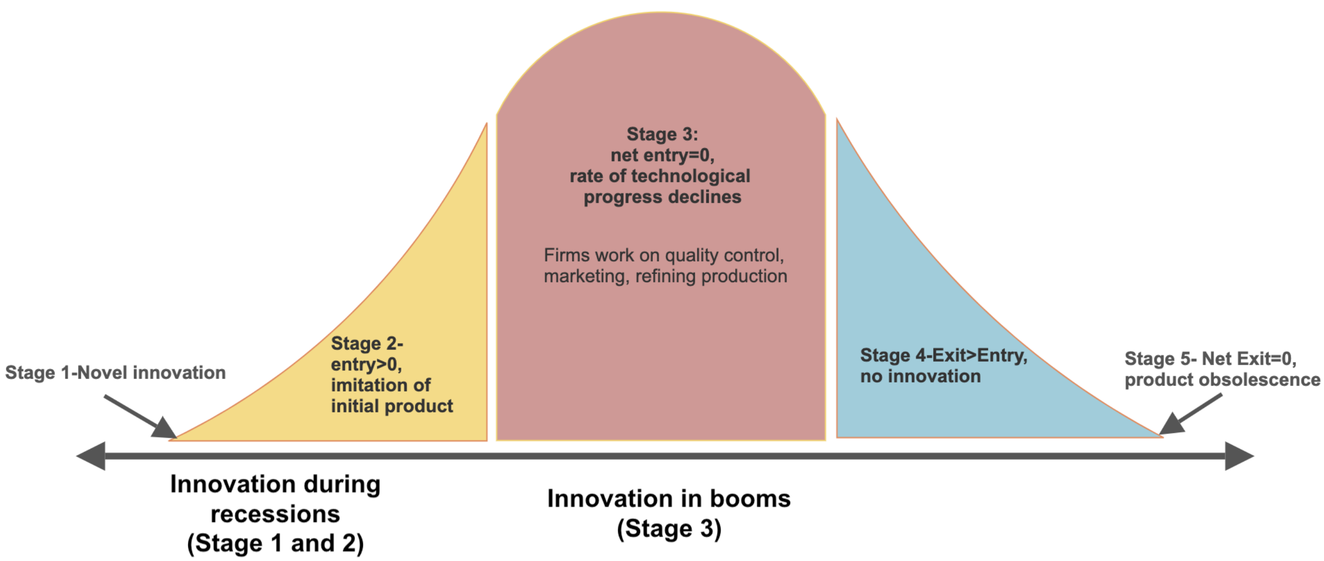

Source: Authors' calculations. The figure displays the five stages of a product life cycle as a function of business cycle conditions.

Table 1: Industry Characteristic Descriptive Statistics

	Sample Size	Mean	Median	SD	Min	Max
Year	2210	1997	1997	9	1980	2015
Industry Relative Valuation	2210	0.004	0.002	0.06	-0.76	2.80
Industry Relative Investment	2210	0.031	0.02	0.71	-7.11	11.64
Real GDP Growth	822	0.67	-0.25	2.10	-0.89	9.44
Employment Growth	790	0.003	0.009	0.05	-0.5	0.5
Total Patents	2210	1831	196	5503	0	63017
Average Citations	2210	17.56	13.07	19.52	0	174
Average Patents Cited	2210	15.33	15.04	3.89	1	47
Average Similarity to Cited Patents	2210	0.27	0.28	0.06	0	0.56
Average Similarity to Citing Patents	2210	0.26	0.26	0.06	0	0.94

Notes.—Source: Compustat and Patentsview. Observations are recorded at the three digit NAICS/year level and are recorded between 1980 and 2015. The calculations relating to patents created in each industry-year are described in Section 3 and in the Appendix.

Table 2: NSF Data Descriptive Statistics: Comparison of Private v. Public Sector

	Sample		Private		Public	
	Mean	SD	Mean	SD	Mean	SD
Demographics						
Male	0.56	0.50	0.60	0.49	0.47	0.50
Age	43.8	10.6	43.7	10.5	44.0	10.8
White	0.76	0.43	0.76	0.43	0.74	0.44
Children, #	2.27	0.92	2.28	0.92	2.24	0.93
Bachelor's Degree	0.57	0.50	0.61	0.49	0.47	0.50
Masters Degree	0.29	0.45	0.25	0.43	0.38	0.48
Professional Degree	0.09	0.29	0.10	0.30	0.06	0.25
PhD Degree	0.05	0.22	0.03	0.17	0.09	0.29
Years of Schooling	17.15	1.56	17.00	1.43	17.50	1.78
Work						
Salary	67275	40349	72119	42345	56214	32788
Hours Worked	1807	524	1838	518	1735	530
Job Satisfaction						
Overall Job Satisfaction	0.01	0.72	-0.00	0.73	0.05	0.69
Career Opportunities	-0.01	0.92	0.02	0.92	-0.07	0.90
Benefits	-0.06	0.93	-0.12	0.94	0.07	0.87
Independence	-0.00	0.72	0.02	0.72	-0.04	0.73
Responsibility	0.02	0.75	0.01	0.75	0.04	0.74
Salary	-0.00	0.85	0.06	0.83	-0.14	0.87
Job security	0.02	0.85	-0.02	0.87	0.11	0.82
Social impact	0.01	0.80	-0.09	0.84	0.24	0.65
Job Activities						
Development	0.27	0.44	0.29	0.45	0.23	0.42
Design	0.24	0.43	0.27	0.45	0.16	0.37
Employee Relations	0.29	0.46	0.30	0.46	0.28	0.45
Management	0.59	0.49	0.60	0.49	0.56	0.50
Maintenance	0.14	0.35	0.17	0.37	0.09	0.29
Quality Management	0.28	0.45	0.31	0.46	0.20	0.40
Sales	0.33	0.47	0.40	0.49	0.19	0.39
Professional Services	0.39	0.49	0.40	0.49	0.35	0.48
Teaching	0.32	0.47	0.22	0.41	0.56	0.50
Computer Services	0.26	0.44	0.28	0.45	0.20	0.40
Supervising	0.44	0.50	0.47	0.50	0.37	0.48
Applied Research	0.29	0.45	0.27	0.44	0.33	0.47
Basic Research	0.20	0.40	0.17	0.38	0.25	0.43
Observations	492730		298373		194357	

Notes.—Source: National Science Foundation Survey of College Graduates (1997-2013). The table reports the means and standard deviations associated with various demographic, work, job satisfaction, and job activity characteristics. Hours worked is made into a continuous variable by taking the product of different weekly hours worked and annual weeks worked groups. Hours worked is partitioned into the following four groups: 20 or less, 21-35, 36-40, and over 40; weeks worked is partitioned into the following four groups: 1-10, 11-20, 21-39, and 40-52. Job satisfaction indices are reported as z-scores, which are created by standardizing the indices, which range on a scale of one to four. Job activities are indicator variables that denote whether the worker allocates at least 10% of their time in the corresponding activity. The sample is restricted to individuals between the age of 25 and 65 years old and observations are weighted by the survey sample weights.

Table 3: Estimating the Cyclicality in the Quality of Innovation

Dep. var. =	log(Average Citations)				log(Similarity to Future Citations)				log(Patents Cited)				log(Similarity to Cited)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Relative Valuation	-0.340***				-0.008*				-0.107				-0.065			
	[0.118]				[0.005]				[0.100]				[0.105]			
Relative Investment		-0.113**				-0.012				-0.045*				-0.039		
		[0.060]				[0.104]				[0.027]				[0.028]		
Real GDP Growth			-0.030*				-0.352**				-0.076				-0.437*	
			[0.017]				[0.157]				[0.228]				[0.238]	
Employment Growth				-0.157				-0.467**				-0.384				-0.390
				[0.347]				[0.214]				[0.290]				[0.262]
R-squared	.932	0.930	.991	.992	.924	0.463	.991	.991	.989	0.928	.99	.99	.924	0.921	.992	.99
Sample Size	2197	2215	822	790	2196	2215	822	790	2197	2215	822	790	2197	2215	822	790
Time Dummies	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
Industry Fixed Effects	x	x	x	x	x	x	x	x x	x	x	x	x	x	x	x	x

Notes.—Source: Compustat and Patentsview (1997-2018). This table reports the coefficients associated with regressions of innovation outcomes (i.e., the flow of innovation) on industry economic conditions, conditional on various layers of fixed effects. The first two columns in the table estimate the effect of economic fluctuations on the impact a patent has on future patents. Namely, we examine the number of citations that a patent receives and the textual similarity of the patent to future patents. The two rightmost columns report the dynamics in similarity to previous patents. This estimates the level of departure of innovations from previous work. These measures are aggregated up to the industry-year level to measure the average impact of innovations produced during industry booms and busts. The key independent variables, four measures of economic conditions, are discussed in Section 3. Standard errors are clustered at the three-digit NAICS level.

Table 4: Examining the Presence of Financial Constraints Over the Business Cycle

Dep. var. =	Research and Development Growth				
	(1)	(2)	(3)	(4)	(5)
Employment Growth	.738***	.689***	.356***	.196	.056
	[.207]	[.206]	[.125]	[.374]	[.443]
× Financial Constraint			.563**	.773***	.022
			[.260]	[.265]	[.511]
Sample Size	23923	23828	23828	14575	9253
R-squared	.20	.20	.20	.22	.19
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Firm Controls	No	Yes	Yes	Yes	Yes
Sample	All	All	All	High Patent	Low Patent

Notes.—Source: Hoberg and Maksimovic (2015), Compustat (1997-2015), Bureau of Economic Analysis (BEA). The table reports the coefficients associated with regressions of year-to-year research and development (R&D) growth (trimmed at the 1st and 99th percentiles) on year-to-year growth in three-digit NAICS industry employment, an indicator for ranking above the median in financial constraints, and their interaction, conditional on firm controls (logged revenue, employment, and property, plant, and equipment). The measure of financial constraints comes from Hoberg and Maksimovic (2015) for each firm-year: "firms with higher values are more similar to a set of firms known to be at risk of delaying their investments due to issues with liquidity." We define industries as having high citation-weighted patenting if their average ratio of citation counts to total patents ranks above the median between 1998 and 2018. Standard errors are clustered at the three-digit NAICS level and observations are unweighted.

Table 5: Examining the Allocation of Time Towards Business and Research Activities Over the Business Cycle

Dep. var. =	Applied Research		Basic Research		Development		Design		Production		Sales	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Private Sector	-.034***	-.045***	-.081***	-.079***	.076***	.058***	.139***	.101***	.062***	.061***	.144***	.140***
	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]	[.004]	[.003]	[.004]	[.005]	[.005]
Employment Growth	-.393***	-.332***	-.051	-.065	-.340***	-.249***	-.806***	-.689***	-.382***	-.376***	-.679***	-.657***
	[.080]	[.081]	[.063]	[.063]	[.085]	[.086]	[.079]	[.076]	[.077]	[.077]	[.086]	[.087]
× Private Sector	.647***	.639***	.313***	.314***	.125	.123	.486***	.520***	.449***	.453***	.080	.081
	[.093]	[.093]	[.075]	[.076]	[.094]	[.094]	[.087]	[.084]	[.083]	[.083]	[.093]	[.093]
Sample Size	344263	343093	344263	343093	344263	343093	344263	343093	344263	343093	344263	343093
Interaction Effect	.25	.31	.26	.25	-.22	-.13	-.32	-.17	.07	.08	-.60	-.58
Interaction S.E.	.05	.05	.04	.04	.04	.04	.04	.04	.03	.03	.04	.04
Demographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Salary Control	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes

Notes.—Source: Notes.—Bureau of Labor Statistics (BLS) Occupation Employment Statistics, National Science Foundation Survey of College Graduates (2000-2013). The table reports the marginal effects on the coefficients (transformed using the margins command) associated with logit regressions of an indicator for whether time allocated in activity k on five-digit occupation employment growth, an indicator for the individual being in the private sector, an interaction, and controls, which include: logged annual salary, race, quadratic in age, education fixed effects, gender, and a quadratic time (year) trend. These k indicators are equal to one if the individual reports allocating over 10% of their time in the activity for a given week, including: applied research, basic research, development, design of equipment / processes / structures / models, productions / operations / maintenance, and sales / purchasing / marketing. Standard errors are clustered at the person-level and observations are weighted by the survey sample weights.

Table 6: Examining the Potential for Reallocation Among Researchers Over the Business Cycle

Dep. var. =	Occupation x Metro Employment Growth			
	(1)	(2)	(3)	(4)
CBSA Employment Growth	.701***	.705***	.706***	.710***
	[.054]	[.055]	[.050]	[.050]
× R&D Occupation	-.016	-.014		
	[.078]	[.078]		
× STEM Occupation			-.112	-.108
			[.132]	[.132]
Sample Size	149834	149833	149834	149833
R-squared	.02	.02	.02	.02
CBSA FE	Yes	Yes	Yes	Yes
SOC FE	Yes	Yes	Yes	Yes
SOC x Year FE	Yes	Yes	Yes	Yes
hassocYear	No	Yes	No	Yes

Notes.—Source: Bureau of Labor Statistics (BLS) Occupation Employment Statistics, BLS classification for science, technology, engineering, and mathematics (STEM) employment, and O*NET index of R&D intensity (authors' calculations). The table reports the coefficients associated with regressions of year-to-year six-digit occupation × CBSA employment growth on an indicator for whether the occupation is classified as research (either above the median R&D intensity or STEM), and their interaction, conditional on occupation, CBSA, and year fixed effects. Standard errors are clustered at the CBSA-level.

Online Appendix

A Constructing Measures of Industry Booms and Busts

Following methods used by Hoberg and Phillips (2010), we identify industry booms and busts using three proxies. These proxies are generated using firm level data from Compustat. The three proxies used are: (1) industry valuation relative to predicted valuation, (2) new investment to the industry relative to predicted investment, and (3) financing at the industry level. These variables are influenced both by beliefs about the future profitability of an industry (relative valuation) as well as current activity (new financing and investment).

The relative valuation and investment in the industry is generated through the same valuation model used in Hoberg and Phillips (2010), updated from Pastor and Veronesi (2003). We use a ten year lagged sample to estimate different coefficients for each industry and year, subsequently generating the estimates of valuation and investment using the current year's data in Compustat. Specifically, the procedure has three steps:

First, we estimate the PV valuation model using data from year t-10 to t-1 for all firms in industry j. We the following for each industry in an unbalanced panel with random effects:

$$\log(M/B)_{i\tau} = \alpha + \beta AGE_{i\tau} + \gamma DD_{i\tau} + \delta LEV_{i\tau} + \phi \log(SIZE)_{i\tau} + \psi VOLP_{i\tau} + \xi ROE_{i\tau} \qquad (2)$$

where $log(\frac{M}{B})$ is the log of the market to book ratio, Age is the negative reciprocal of one plus firm age, DD is a dividend dummy, LEV is the firm's leverage, $SIZE$ is the log of total assets, current firm return on equity (ROE), and $VOLP$ is the volatility of the firm's profitability. These variables are calculated from Compustat items following Pastor and Veronesi (2003) and Fama and French (1993). Following both Hoberg and Phillips (2010), we drop firm observations whose market equity, book equity, and total assets are smaller than one million, market to book ratios outside of the range (0.01,100), and we winsorize VOLP and ROE at the 1/99% level in each year.

Second, from the estimation in the first step, we derive a different set of coefficients for the valuation equation for each industry and year.[11] These coefficients are then used to compute predicted values for the firm market-to-book ratio in each year. This imposes the assumptions that the firm's predicted valuation is a function of its current characteristics and the industry-specific prices of these characteristics from previous years.

Third, we calculate the relative valuation by subtracting the predicted firm valuation in year t from the second step from the log of the firm's actual market to book value in year t. We then winsorize at the 1/99% level for each year and take the average relative valuation over all firms in the industry. Relative investment at the firm and industry level is calculated as follows:

$$\begin{aligned} \log(INVEST/PPE)_{i\tau-1} = \alpha + \beta TOBINQ_{i\tau} + \gamma ROE_{i\tau} + \delta DD_{i\tau} + \\ \phi AGE_{i\tau} + \psi LEV_{i\tau} + \xi VOLP_{i\tau} + \zeta \log(SIZE)_{i\tau} \end{aligned} \tag{3}$$

We define ROE, DD, AGE, LEV, $SIZE$, and $VOLP$ as in the valuation model. Lagged Tobin's Q is also included here but not in the valuation model. After running the time series estimations and getting the estimates of the coefficients for each industry-year, this model generates

[11] The set of coefficients for each year t and industry j is a function of the characteristics and valuations of firms in that industry from years t-10 to t-1.

predictions of investment for each firm in each year based on the firm's characteristics and the historic returns to these characteristics in investment. We define relative firm investment as the actual investment subtracted from the predicted investment in each year. Then, we aggregate to the industry level by averaging over all firms in each year.

Finally, total new financing to the firm in a year is defined as the sum of net equity issuing activity and the sum of net debt issuing activity divided by the firm's assets. This is averaged over the industry level to get new financing to the industry in each year.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, reallocation and growth. *American Economic Review*, 108(11):3450–3491.

Aghion, P., Askenazy, P., Berman, N., and Cette, G. (2012). Credit constraints and the cyclicality of r&d investment: Evidence from france. *Journal of the European Economic Association*, 10(5):1001–1024.

Aghion, P. and Howitt, P. (1992). A model of growth through creative destruction. *Econometrica*, 60(2):323–251.

Aghion, P. and Saint-Paul, G. (1998). Virtues of bad times: Interaction between productivity growth and economic fluctuations. *Macroeconomic Dynamics*, 2(3):322–344.

Akcigit, U., Hanley, D., and Serrano-Velarde, N. (2019). Back to basics: Basic research spillovers, innovation policy and growth. *Review of Economic Studies, R&R*.

Autor, D. and Dorn, D. (2013). The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H. and Handel, M. J. (2013). Putting tasks to the test: Human capital, job tasks and wages. *Journal of Labor Economics*, 31(2):S59–S96.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). Trends in U.S. wage inequality: Revising the revisionists. *Review of Economics and Statistics*, 90(2):300–323.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4):1279–1333.

Barlevy, G. (2007). On the cyclicality of research and development. *American Economic Review*, 97(4):1131–1164.

Becker, B. (2015). Public R&D policies and private R&D investment: A survey of the empirical evidence. *Journal of Economic Surveys*, 29(5):917–942.

Bernstein, S., McQuade, T., and Townsend, R. (2020). Do household wealth shocks affect productivity? evidence from innovative workers during the great recession. *Journal of Finance*.

Blanchard, O. J. and Katz, L. (1992). Regional evolutions. *Brookings Papers on Economic Activity*, 1:1–75.

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. (2020). Are ideas getting harder to find? *American Economic Review*, 110(4):1104–1144.

Bozeman, B. (2000). Technology transfer and public policy: A review of research and theory. *Research Policy*, 29:627–655.

Bozeman, B. and Kingsley, G. (1998). Risk culture in public and private organizations. *Public Administration Review*, 58(2):109–118.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies "engines of growth"? *Journal of Econometrics*, 65(1):83–108.

Brynjolfsson, E., Rock, D., and Syverson, C. (2019). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. *National Bureau of Economic Research, University of Chicago Press*, 1:23–57.

Burns, A. E. (1934). Some theoretic bases of the economic state. *The ANNALS of the American Academy of Political and Social Science*, 174(1):173–178.

Caballero, R. and Hammour, M. (1994). The cleansing effect of recessions. *American Economic Review*, 84(5):1350–1368.

Canton, E. and Uhlig, H. (1999). Growth and the cycle: Creative destruction versus entrenchment. *Journal of Economics*, 69(3):239–266.

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

Fabrizio, K. and Tsolmon, U. (2014). An empirical examination of the procyclicality of r&d investment and innovation. *Review of Economics and Statistics*, 96(4):662–675.

Fama, E. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33:3–56.

Geroski, P. A. and Walters, C. F. (1995). Innovative activity over the business cycle. *Economic Journal*, 105(431):916–28.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2019). Bartik instruments: What, when, why and how. *American Economic Review, conditionally accepted.*

Gort, M. and Klepper, S. (1982). Time paths in the diffusion of product innovations. *Economic Journal*, 92(367):630–653.

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

Hagerty, M. R. and Land, K. C. (2007). Constructing summary indices of quality of life: A model for the effect of heterogeneous importance weights. *Sociological Methods Research*, 35:455–496.

Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? Evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772.

Hoberg, G. and Maksimovic, V. (2015). Redefining financial constraints: A text-based analysis. *Review of Financial Studies*, 28(5):1312–1352.

Hoberg, G. and Phillips, G. (2010). Real and financial industry booms and busts. *Journal of Finance*, 65(1):45–86.

Jones, C. (2002). Sources of U.S. economic growth in a world of ideas. *American Economic Review*, 92(1):220–239.

Jung, J., Bozeman, B., and Gaughan, M. (2017). Impact of research collaboration cosmopolitanism on job satisfaction. *Research Policy*, 46(10):1863–1872.

Kelly, B., Papanikolaou, D., Seru, A., and Taddy, M. (2018). Measuring technological innovation over the long run. *NBER Working Paper No. 25266*.

Koyptov, A., Roussanov, N. L., and Taschereau-Dumouchel, M. (2018). Short-run pain, long-run gain? recessions and technological transformation. *NBER Working Paper No. w24373*.

Kuhn, J. M., Younge, K. A., and Marco, A. C. (2020). Patent citations reexamined. *RAND Journal of Economics*, 51(1):109–132.

Lancaster, T. (2000). The incidental parameter problem since 1948. *Journal of Econometrics*, 95:391–413.

Lane, J. I., Owen-Smith, J., Rosen, R., and Weinberg, B. A. (2015). New linked data on research investments: Scientific workforce, productivity, and public value. *Research Policy*, 44(9):1659–1671.

Lybbert, T. and Zolas, N. (2014). Getting patents and economic data to speak to each other: An 'algorithmic links with probabilities' approach for joint analyses of patenting and economic activity. *Research Policy*, 43(3):530–542.

Manso, G., Balsmeier, B., and Fleming, L. (2019). Heterogeneous innovation and the antifragile economy.

Ouyang, M. (2011). On the cyclicality of R&D. *Review of Economics and Statistics*, 93(2):542–553.

Pastor, L. and Veronesi, P. (2003). Stock valuation and learning about profitability. *Journal of Finance*, 58(5):1749–1789.

Roach, M. and Sauermann, H. (2010). A taste for science? PhD scientists' academic orientation and self-selection into research careers in industry. *Research Policy*, 39(3):422–434.

Roach, M. and Sauermann, H. (2014). Not all scientists pay to be scientists: PhDs' preference for publishing in industrial employment. *Research Policy*, 43(1):32–47.

Schumpeter, J. (1939). *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process*. McGraw-Hill Book Company, Inc., New York and London.

Shleifer, A. (1986). Implementation cycles. *Journal of Political Economy*, 94(6):1163–1190.

Stern, S. (2004). Do scientists pay to be scientists? *Management Science*, 50(6):835–853.

Syverson, C. (2017). Challenges to mismeasurement explanations for the US productivity slowdown. *Journal of Economic Perspectives*, 31(2):165–186.

Terry, S. J. (2017). The macro impact of short-termism. *Econometrica, R&R.*