

NBER WORKING PAPER SERIES

WHAT IS A PATENT WORTH? EVIDENCE FROM THE U.S. PATENT “LOTTERY”

Joan Farre-Mensa
Deepak Hegde
Alexander Ljungqvist

Working Paper 23268
<http://www.nber.org/papers/w23268>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue
Cambridge, MA 02138

March 2017, Revised December 2018

A previous version of this paper was circulated under the title “The Bright Side of Patents.” We are grateful to Emily Blanchard, Lauren Cohen, Wesley Cohen, Michael Ewens, Matt Fiedler, Lee Fleming, Teresa Fort, Ken French, Alberto Galasso, Bronwyn Hall, Zorina Khan, Josh Lerner, Ross Levine, Alan Marco, Ramana Nanda, Bhaven Sampat, Robert Seamans, Amit Seru (the Editor), Ted Sichelman, Scott Stern, Rick Townsend, Heidi Williams, Rosemarie Ziedonis, an anonymous associate editor, two anonymous reviewers, and audiences at the NBER Productivity Lunch Seminar, the NBER Summer Institute, the SFS Cavalcade, the WFA meetings, and various universities for helpful comments. We also thank the Institute for Exceptional Growth Companies for granting access to the NETS database. Hegde gratefully acknowledges the support of the United States Patent and Trademark Office’s Thomas Alva Edison Visiting Scholars program and the Kauffman Junior Faculty Fellowship. The views and comments expressed herein are solely the opinion of the authors, do not reflect the performance of duties in the authors’ official capacities, and are not endorsed by, nor should be construed as, any viewpoint official or unofficial of the United States Patent and Trademark Office. The authors confirm to the best of their knowledge that no information contained herein is privileged, confidential, or classified. The authors confirm to the best of their knowledge that no information contained herein is privileged, confidential, or classified.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w23268.ack>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Joan Farre-Mensa, Deepak Hegde, and Alexander Ljungqvist. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

What is a Patent Worth? Evidence from the U.S. Patent “Lottery”
Joan Farre-Mensa, Deepak Hegde, and Alexander Ljungqvist
NBER Working Paper No. 23268
March 2017, Revised December 2018
JEL No. D23,G24,L26,O34

ABSTRACT

We provide evidence on the value of patents to startups by leveraging the quasi-random assignment of applications to examiners with different propensities to grant patents. Using unique data on all first-time applications filed at the U.S. Patent Office since 2001, we find that startups that win the patent “lottery” by drawing lenient examiners have, on average, 55% higher employment growth and 80% higher sales growth five years later. Patent winners also pursue more, and higher quality, follow-on innovation. Winning a first patent boosts a startup’s subsequent growth and innovation by facilitating access to funding from VCs, banks, and public investors.

Joan Farre-Mensa
Cornerstone Research
699 Boylston Street, 5th Floor
Boston, MA 02116
j.farre-mensa@northeastern.edu

Alexander Ljungqvist
Stockholm School of Economics
Box 6501
SE-113 83 Stockholm, Sweden
aljungqv@stern.nyu.edu

Deepak Hegde
Stern School of Business
New York University
40 West Street
Suite 722
New York, NY 10012
dhegde@stern.nyu.edu

Patents award temporary monopoly rights over inventions to their inventors. In theory, patents can benefit their holders by deterring copycats, as defensive shields in litigation suits, as bargaining chips in licensing negotiations, and as signaling devices to attract investors and customers (Williams 2017). Policymakers and scholars argue that these benefits offset the social cost of patent rights (Nordhaus 1969, Scherer 1972). Startups may derive particularly large private benefits from patents, given that they typically have few assets early on in their lives, other than their ideas, and lack alternative mechanisms to protect their inventions (Lamoreaux and Sokoloff 1999, Federal Trade Commission 2011). At the same, there is mounting evidence that patents are abused, costly to enforce, and tax follow-on inventors (Heller and Eisenberg 1998, Becker and Posner 2013, Jaffe and Lerner 2004, Lemley and Shapiro 2005, Galasso and Schankerman 2015). How valuable, then, are patents to startups?

The value of a patent is the incremental economic benefit accruing to its holder from the legal right to exclude others from exploiting the invention, beyond what would be earned if the invention were not granted a patent. Separating this incremental benefit from the economic value of the underlying invention itself is challenging. Previous studies document that firm value increases in the number of patents a firm holds or when a firm is awarded a patent, but these correlations conflate the value of the underlying inventions and the incremental benefits of the exclusion rights conferred by the patents.¹ Our study aims to provide an estimate of these incremental economic benefits, focusing on how patents affect the growth and success of startups. To do so, we employ novel data and an identification strategy that disentangles the causal effects of patent rights from the economic value of the underlying inventions.

To isolate the incremental benefit of patents over and above the value of the underlying

¹ For example, Hall, Jaffe, and Trajtenberg (2005), Balasubramanian and Sivadasan (2011), and Kogan et al. (2017) report a positive correlation between patents and measures of firm value.

inventions, we assemble novel data on patent rejections. Until recently, researchers have lacked access to data on patent applications that were denied and so have instead compared firms with and without patents. Of course, firms could be without patents for many reasons: they may not have inventions to patent in the first place, may have had their patent applications denied, or may have chosen other ways to protect their inventions (such as trade secrets). Comparing firms with and without patents thus cannot identify the value of a patent right, as distinct from the value of the underlying invention. We overcome this challenge thanks to special access to the internal databases of the U.S. Patent and Trademark Office (USPTO), which cover the population of not only granted but also rejected applications.²

We combine the applications data with data on the performance of startups assembled from a variety of sources. We focus on five types of outcome variables: (i) growth in sales and employment (from Dun & Bradstreet's National Establishment Time Series or NETS database); (ii) follow-on patenting and patent citations (from the USPTO's patent database); (iii) the pledging of patent applications as collateral to raise debt (from the USPTO's patent assignment database); (iv) venture funding (from VentureXpert); and (v) fundraising by startups through initial public offerings (IPOs) (from VentureXpert and Thomson-Reuters' SDC database).

To disentangle the value of a patent from that of the underlying invention, we exploit plausibly exogenous variation in the patent approval process through an instrumental-variables approach pioneered by Sampat and Williams (2018).³ The validity of the approach rests on two features of the USPTO's patent examination process. First, the USPTO assigns applications in each technology field (or "art unit") to examiners randomly with respect to the characteristics of

² Access to the USPTO's internal databases was granted to one of the authors through the agency's Thomas Alva Edison Scholars program.

³ Sampat and Williams (2018) use the approach to measure the effect of gene patents on follow-on innovation in the human genome field.

the underlying invention (Lemley and Sampat 2012, Sampat and Williams 2018). Second, examiners vary in their propensity to approve applications: some are more lenient while others are stricter, leading Cockburn, Kortum, and Stern (2003) to observe that “there may be as many patent offices as patent examiners.” The assignment of applications to examiners who quasi-randomly differ in their leniency imparts a lottery-like element to the patent review process: among applicants with patentable inventions, some win patent rights while others do not, simply because the former were lucky enough to draw more lenient examiners.⁴ Thus, we are able to isolate the economic benefits of patent rights from those of the underlying invention.

Our sample covers all 34,215 first-time patent applications filed by U.S. startups at the USPTO since 2001 that received a final decision by December 31, 2013. Our estimates suggest that winning the patent lottery, by randomly drawing an examiner more likely to approve applications, increases the average startup’s employment growth over the next five years by 54.5 percentage points, compared to startups with only randomly different underlying inventions that fail to obtain a patent. The effect on sales growth—a 79.5 percentage-point increase over five years—is even larger. For the average startup in our sample, these estimates imply that receiving a patent leads to 16 additional employees after five years, and \$10.6 million in additional sales cumulated over five years after winning the patent lottery. A first patent grant also increases both the number of subsequent patents the firm is granted (by 56.5%) and their quality (with the average number of citations per subsequent patent increasing by 33%).

In contrast to the large benefits of receiving a first patent, subsequent patents are much less

⁴ Gaule (2018) uses a similar instrument to study the effect of patents on the likelihood of going public or being acquired using a sample of 2,191 VC-backed startups. Galasso and Schankerman (2015) exploit variation in the leniency of judges at the U.S. Court of Appeals of the Federal Circuit to study the effect of patent invalidation on follow-on innovation in related technologies. Of course, not all patent applications are equally affected by the leniency of their examiner: applications of marginal quality are likely to be the most affected. In the language of the treatment effects literature, these marginal applications comprise our instrument’s compliant subpopulation.

obviously beneficial to the startups in our sample. We find little evidence that approval of a second patent application boosts growth in employment or sales significantly (even for startups whose first application was rejected), though it does increase the number of subsequent patents (by 49.8%). Approval of a third patent shows similar patterns. These results suggest that the value of patents to startups changes over time. Securing a first patent, when they typically have few other assets, allows startups to innovate and embark on a high-growth trajectory. Subsequent patents encourage further innovation but play less of a direct role in boosting growth, which instead presumably depends on other resources the startups have acquired along the way.

What explains the large economic benefits startups derive from obtaining a patent, especially for their first application, over and above the economic benefits arising from the underlying invention? To make headway on this question, we investigate a funding channel. Empirically, we find that the first patent increases a startup's chances of securing funding from VCs over the next three years by 47%, and of securing a loan by pledging the patent as collateral by 76%. A patent also more than doubles the odds of the startup raising funding from public investors through an IPO. The VC funding effect is strongest for startups founded by inexperienced entrepreneurs and located in areas where attracting investors' attention is harder, and weakest for biochemistry startups. Mirroring the ambiguous effects of subsequent patents on the performance of startups, we find that the approval of a startup's second or third patent application appears to have no statistically significant impact on the startup's ability to raise VC funding.

Collectively, our findings suggest that patents facilitate startups' access to external finance in contexts where information frictions, and thus contractual hazards, are especially high. A patent grant sets a startup on a growth path through funding that helps transform its ideas into products and services that generate jobs, revenues, and follow-on inventions.

Our study makes the following contributions to the literatures on intellectual property, innovation, and entrepreneurial finance. First, cash-constrained startups account for thousands of patent applications every year, even though patents are expensive to obtain and maintain⁵ and there is significant uncertainty surrounding their enforceability and thus their expected value.⁶ Our study helps explain the apparent attraction of patents to startups by providing the first causal estimates of the incremental economic benefits U.S. startups derive from patent grants.⁷ Measuring patent value as the incremental growth in employment and sales, over and above that stemming from the underlying invention, as well as the incremental increase in follow-on innovation activity, we show that patents confer economically large benefits on startups. Interestingly, not all patents are created equal: second and subsequent patents, after purging out the benefits of the underlying inventions, are less effective at stimulating growth and success.

Second, we provide causal evidence that patents facilitate startup growth by enabling access to external funding,⁸ particularly in settings associated with high information frictions. These findings are consistent with the view that patents, through their property rights function, can help startups raise funding by providing some level of legal recourse against expropriation (Johnson et al. 2002) or by mitigating information asymmetry problems that hamper contracting, particularly between first-time inventors and investors (Arrow 1962, Long 2002).

Third, our identification strategy, based on quasi-random assignment of applications to patent

⁵ The average U.S. patent costs an estimated \$20,000 to obtain and thousands more to keep from expiring before its 20-year term (Lemley 2001).

⁶ Unlike traditional property rights, such as title to a piece of land, patents are awarded for technological inventions, the boundaries of which tend to be uncertain and so costly to enforce (Lemley and Shapiro 2005). The USPTO is blamed for aggravating this uncertainty by granting patents to worthless inventions (Heller and Eisenberg 1998, Jaffe and Lerner 2004, Becker and Posner 2013). For example, U.S. Patent #4,344,424 for an [anti-eating mouth cage](#), U.S. Patent #5,934,226 for a [bird diaper](#), and U.S. Patent #6,293,874 for a [buttock-kicking apparatus](#).

⁷ Schankerman (1998) provides the first, and only other, systematic evidence we are aware of on the private value of patents. His estimates are derived from patent renewal data for French patents granted between 1969 and 1982.

⁸ Our evidence on this point echoes prior arguments made in correlation studies (Conti, Thursby, and Thursby 2013, Hsu and Ziedonis 2013), survey studies (Graham and Sichelman 2008), and law reviews (Long 2002).

examiners of varying leniency, highlights the profound impact the luck of the draw can have in determining startups' fortunes. The median examiner in our sample grants 61.5% of applications. An applicant "lucky" enough to have drawn an examiner in the 75th percentile of leniency enjoys an 11.8 percentage-point higher probability of winning the patent lottery, and thus achieving the sales and employment growth associated with a patent grant, than an applicant who has drawn an examiner in the 25th percentile of leniency. We are agnostic about whether a system that results in such large variation in fortunes is "good" or "bad" for efficiency, except to note that variation among evaluators is to be expected in systems that depend on subjective human judgment.⁹

Fourth, patents being privately valuable is clearly not a sufficient condition for the patent system to improve welfare (we cannot ignore negative externalities imposed on other inventors or deadweight monopoly costs), but it *is* a necessary condition: patents are supposed to promote innovation through the promise of private value to inventors, and we establish that they do, particularly for entrants into patenting.¹⁰ The large economic magnitude of the private returns we document suggests that the patent system is at least partially performing its intended function by promoting the growth of entrepreneurs who are least likely to have access to alternative commercialization channels, at a time when they face the greatest information-related hurdles. Thus, we contribute to the growing debate about the welfare consequences of the patent system.

Finally, previous empirical tests of the Schumpeterian hypothesis linking innovation to economic growth report large positive effects of patented inventions on firm growth and productivity (Kogan et al. 2017). Our findings raise the possibility that at least some of these positive effects are due to the legal right to exclude conveyed by patent grants.

⁹ For example, teachers may be more or less lenient when grading examinations, evaluators at funding bodies vary in their stringency, and court judges may be more or less sympathetic towards the accused.

¹⁰ Of course, the fact that the market values of startups are not readily observable means that we cannot estimate the private value of patents in dollar terms.

1. Institutional setting and data

1.1 The patent examination process

The USPTO sends incoming applications to the appropriate “art unit” for review. Each art unit consists of a group of patent examiners who specialize in a narrowly defined technology field.¹¹ During our sample period, the USPTO employed some 13,000 examiners in over 900 art units. The median art unit has 13 examiners; the largest, more than 100.

Applications in each art unit’s holding queue are assigned to one of the unit’s examiners, who is responsible for assessing if the claims in the application meet the legal thresholds of novelty, usefulness, and non-obviousness. As we argue more fully in Section 2.3, the matching of application to examiner is orthogonal to the quality of the application or of the applicant. This quasi-random assignment is central to our identification strategy.

After receiving an assignment, the examiner evaluates the application and makes a preliminary ruling on its validity. This ruling, called the “first-action decision,” is communicated to the applicant via an official letter, and it is from this letter that the applicant first learns the examiner’s identity. On average, sample applications take 0.7 years to be assigned to an examiner, who then takes an additional year to make a first-action decision. The final accept/reject decision is on average made 1.5 years later (i.e., 3.2 years after the application date).¹²

1.2 Timing considerations

Firm outcomes could, in principle, be measured from three alternative starting points: the filing date, the first-action date, and the final-decision date. Our choice of starting point is guided

¹¹ To illustrate, the examiners in art unit 1641 are in charge of examining patent applications related to “peptide or protein sequence,” examiners in art unit 2831 are in charge of applications related to “electrical connectors,” examiners in art unit 3676 are in charge of applications related to “wells and earth boring,” and so on.

¹² Strictly speaking, patent applications are never irrevocably rejected by the USPTO; they are abandoned by applicants following what technically are appealable rejections issued by examiners (Lemley and Sampat 2008). For expositional clarity, we follow Sampat and Williams (2018) and refer to abandoned applications (i.e., the complement of those applications that are approved) as “rejected.”

by two considerations: how uncertainty about the patentability of a startup's invention evolves over time, and from what point onwards the startup's behavior could affect the timing of the USPTO's decision. Resolution of uncertainty is necessary (but not sufficient) for a patent application to affect firm outcomes. Endogenous timing of the final approval decision may contaminate our causal estimates.

Because the first-action letter is the first communication about the merits of its application that an applicant receives from the USPTO, there is no resolution of uncertainty before the first-action date. This rules out using the filing date as the starting point. The timing of the final decision is likely endogenous: the delay between first action and final decision is determined, in large part, by applicants' actions (in particular, how long they take to respond to any concerns the examiner may have raised in the first-action letter). This rules out using the final-decision date as the starting point. The appropriate choice of starting point in our setting is hence the first-action date. As Carley, Hegde, and Marco (2015) note, first-action letters contain a detailed summary of the merits the examiner sees in the application, so the first-action decision resolves a substantial amount of uncertainty about the application's ultimate fate, as required.

1.3 Patent data and sample selection

Establishing the private value of a patent requires data on both approved and rejected patent applications. Until recently, publicly available datasets, such as those maintained by the NBER or Harvard Business School, only covered approved patents (Lerner and Seru 2015), and even now, public disclosure of rejected applications is incomplete.¹³ Instead, we obtain patent data directly from the USPTO's internal databases, which contain a complete record of both approved

¹³ The USPTO's Patent Application Information Retrieval (PAIR) system, which is publicly accessible, provides no data on applications that are abandoned prior to public disclosure (around 15% of all unsuccessful applications) and no data on rejected applications filed before 2001.

and rejected patent applications going back to 1976.¹⁴ For applications filed since 2001, when the American Inventors Protection Act came into force, we know the applicant's identity, regardless of the USPTO's final decision. For earlier applications, the applicant's name is not recorded for rejected applications, though all other relevant aspects of the application are available. Our sample hence starts in 2001 (though we use the earlier data in the coding of our instrument).

Our goal is to identify the benefits of early patent grants to startups. The USPTO does not tag whether an applicant is a startup, so to code startups we proceed as follows. First, we restrict the sample to incorporated applicants based in the U.S. Second, we screen out not-for-profit entities such as charities, universities, or government research labs. Third, using data from a variety of sources, we manually screen out applicants that are or have been listed on a stock market or are or have been a subsidiary of another firm (whether listed or privately held, and whether domestic or foreign), at the time of the application.¹⁵ These three steps leave a set of patent applications filed by stand-alone for-profit U.S. firms. Not all are startups. To zero in on applicants likely to be startups, we apply two further filters. First, we focus on filers that qualify for reduced fees at the USPTO by virtue of satisfying the criteria defining a "small business entity" under Section 3 of the Small Business Act. Second, we focus on applicants without a history of patenting, by screening for firms that have filed at least one application on or after January 1, 2001 and no applications in the previous 25 years. This step requires standardizing applicant names and checking for name changes. Since many patents are subsequently reassigned, this step also requires identifying each patent's original applicant.

Our analysis focuses on how the outcome of a firm's early patent application affects its

¹⁴ Carley, Hegde, and Marco (2015) provide a comprehensive description of these data.

¹⁵ We also screen out a small number of applicants that are acquired between filing and first-action, on the ground that we cannot disentangle the effects of the patent decision from the effects of the acquisition.

ability to grow, continue innovating, and raise funding. To ensure we have sufficient time to study the long-term effects of early patent grants, we require firms to receive a first-action decision on their first application by the end of 2009 and a final decision by the end of 2013.¹⁶

Our final sample consists of 34,215 first-time patent applicants (called startups from here on). Of these, 31.6% operate in the electronics, computers, and communications industries (henceforth, IT); 17.8% are active in the pharmaceutical and biochemical sectors (henceforth, biochemistry); and the remaining 50.6% operate in “other” industries including transportation, construction, mechanical engineering, and manufacturing. Just under two-thirds (64.5%) of first-time patent applications in our sample are granted over our sample period.

1.4 Data on firm outcomes

Being privately held, the startups in our sample are not covered in standard financial databases such as Compustat, so we assemble data on firm outcomes from five other sources.

- *Dun and Bradstreet’s National Establishment Time Series (NETS) database.* NETS is similar to the U.S. Census Bureau’s Longitudinal Business Database (LBD) in that it aims to cover the universe of business establishments in the U.S., but offers the advantage of not requiring special permission for access. Matching patent assignees to NETS (and to other databases) requires matching on firm names and locations. We use a “fuzzy” matching algorithm (with each potential match manually verified), supplemented with information on name changes and location moves obtained from Capital IQ and the USPTO’s firm name and address register. We are able to match 80.01% of sample startups to firms in NETS – a higher match

¹⁶ The firm’s “first application” is the first application the USPTO rules on. (In 8% of cases, the first ruling a firm receives is not for its first-ever application but for a later application.)

rate than that achieved by studies using the Census Bureau's data.¹⁷

- *The USPTO's patent database.* This database allows us to track each sample company's subsequent patent applications as well as citations to each sample company's patents.
- *The USPTO's patent assignment database.* This dataset records transactions between patent owners and entities to whom owners transfer patent rights. We use it to track instances of startups pledging their patents or patent applications as collateral for loans.
- *VentureXpert.* This database tracks VC funding events. We use it to identify which of the 34,215 sample firms go on to raise VC funding at some point after the first-action date.
- *The Thomson Reuter's Securities Data Company (SDC) database.* We use data from SDC (and VentureXpert) to identify firms that raise capital from public investors via an IPO.

Table 1 provides summary statistics and compares startups whose first patent application is approved or rejected. Panel A shows that at the time of application, the median startup is 2 years old, has 8 employees, and \$800,000 in sales. Unsuccessful applicants have lower pre-filing growth in employment and sales than successful ones, suggesting that endogeneity concerns are rife in patent application data. After the USPTO's first-action decision, successful applicants grow employment and sales at substantially higher rates than unsuccessful applicants (Panel B), produce more and higher-quality follow-on inventions (Panel C), are more likely to raise funding from VCs or the IPO market (Panel D), and are more likely to raise debt by pledging their patent applications as collateral, even when not funded by VCs (Panel E). These simple correlations suggest that startups whose first patent application is approved experience superior outcomes.

¹⁷ Balasubramanian and Sivadasan (2011) are able to match 63.7% of patent assignees to firm names in the Census Bureau's Business Register, often considered the "gold standard" for its coverage of the entire population of U.S. business establishments with paid employees filing taxes with the Internal Revenue Service. Kerr and Fu (2008) report a match rate of about 70%.

2. The real effects of patent grants

Our focus in this section is on how obtaining a legal property right over an invention (as distinct from the invention itself) affects a startup's subsequent growth in employment and sales and its follow-on innovation. We first outline our empirical setup, review the main challenge our identification strategy needs to overcome, and discuss our identifying assumptions and associated threats to identification. We then report our findings on the real effects of a startup obtaining an early patent. In the next section, we investigate financing channels that link these outcomes to legal property rights.

2.1 Empirical setup and identification challenge

We estimate panel regressions of the following general form:

$$Firm\ outcome_{ijat+k} = \beta Patent\ application\ approved_{ijat} + \Phi X_{ija\tau} + \nu_{a\tau} + \varepsilon_{ijat+k} \quad (1)$$

where i indexes startups, j examiners, a art units, t first-action years, and τ application years.

$X_{ija\tau}$ includes HQ-state fixed effects to control for geographical differences in outcomes (Lerner and Seru 2015). Art-unit-by-application-year fixed effects, $\nu_{a\tau}$, are included to control for unobserved demand or technology shocks that could affect both applications and outcomes; for example, a breakthrough in a technology field may lead to an increase in both the number of patentable inventions and the growth rate of firms operating in that field.¹⁸ Standard errors are clustered at the art unit level to allow for arbitrary correlation of the errors within each art unit.

Assuming it can be identified, the coefficient of interest, β , captures the average treatment effect of being granted a patent on an invention on firm outcomes, that is, the conditional average difference in outcomes between a startup that has both an invention and is granted a patent over

¹⁸ Since art units are quite narrowly defined (they span 495 different technology fields in our sample), including these fixed effects allows us to control for time-varying demand and technological conditions at a very fine level and so ensures that our findings are not confounded by unobserved industry-level shocks. Applications belonging to art-unit-by-year singletons do not contribute to identification and are excluded.

it and a startup subject to the same demand and technology conditions that has not been granted a patent over its invention. The main identification challenge is to ensure that differences in outcomes reflect the effects of being granted a patent and not differences in the quality of the underlying invention (or other factors that correlate with its quality, such as unobserved characteristics of the startup).

2.2 Identification strategy and identifying assumptions

The ideal experiment to identify the causal contribution of a patent over an invention to a firm's success would randomize patent approvals, thus ensuring that successful applicants and their inventions do not differ systematically from unsuccessful ones *ex ante*.

We exploit a lottery-like feature of the USPTO's review process to get as close as possible to this ideal experiment. As prior work has already argued,¹⁹ the lottery-like feature in question has two components: patent applications are assigned to examiners in a quasi-random fashion; and patent examiners differ systematically in their propensity to approve patents, such that more lenient examiners are more likely to grant a patent than are stricter examiners, holding the quality of the invention constant. Under these assumptions, threats to which we discuss in the next subsection, examiner leniency is a valid instrument for whether or not a startup's patent application is approved so that β in equation (1) can be consistently estimated using 2SLS.

We measure examiner leniency as the approval rate of examiner j belonging to art unit a assigned to review startup i 's patent application submitted at time τ :

$$\text{Examiner approval rate}_{ija\tau} = \frac{n_{\text{granted}_{ja\tau}}}{n_{\text{reviewed}_{ja\tau}}}, \quad (2)$$

where $n_{\text{reviewed}_{ja\tau}}$ and $n_{\text{granted}_{ja\tau}}$ are the numbers of patents examiner j has reviewed and

¹⁹ Notably Lichtman (2004), Sampat and Lemley (2010), Lemley and Sampat (2012), Sampat and Williams (2018), and Gaule (2018).

granted, respectively, prior to the application date τ .^{20,21,22} Figure 1 shows that examiner approval rates vary substantially, even within a given art unit and application year. To illustrate the variation, consider a startup that is assigned to an examiner with an approval rate at the 75th percentile rather than at the 25th percentile in that art unit and year. Figure 1 suggests that this assignment would put the startup’s application on the docket of an examiner whose approval rate is 17.6 percentage points higher—a lucky draw indeed.²³

Table 2, column 1 reports the first stage of our 2SLS specification, that is:

$$\text{Patent application approved}_{ijat} = \theta \text{Examiner approval rate}_{ija\tau} + \Pi X_{ija\tau} + v_{a\tau} + u_{ijat}. \quad (3)$$

The results confirm that examiner leniency is a strong predictor of whether an application is approved, supporting the relevance condition required for identification.²⁴ The coefficient estimate for θ in column 1 implies that each percentage-point increase in an examiner’s prior approval rate leads to a 0.67 percentage-point increase in the probability that a startup patent she reviews is approved ($p < 0.001$). Thus, moving from an examiner at the 25th percentile to one at the 75th percentile would increase a startup’s probability of being granted a patent by 11.8 percentage points ($= 0.67 \times 17.6$), all else equal.

²⁰ Neither the numerator nor the denominator in (2) includes patent application i , as it has not been reviewed prior to date τ . To ensure that we measure approval rates accurately, we exclude startups whose application is assigned to an examiner with fewer than 10 prior reviews. All results are robust to using alternative cutoffs.

²¹ Due to data limitations, prior work often uses a somewhat noisier instrument, sometimes called the “leave-one-out” estimator, defined using an examiner’s future decisions as well as her past ones. Our results are little different using this alternative way of coding the instrument.

²² Our results are robust to computing the instrument as the approval rate of all first-time patent applications assigned to a given examiner, or to including only the last three years of applications. We thank an anonymous referee for suggesting these robustness tests.

²³ Our approval rates are based on a large number of reviewed applications: the average (median) examiner has reviewed 770 (409) applications by the time we measure her approval rate (the 10th percentile is 51). This suggests that the variation shown in Figure 1 reflects persistent inherent differences in examiners’ propensity to approve applications and not small-sample random differences in the quality distribution of the applications they review.

²⁴ We estimate a linear probability model in the first stage instead of a probit model to avoid using the nonlinearity of the probit model as an additional source of identifying information (Angrist and Pischke 2009, p. 190). All our conclusions are robust to estimating our first-stage equation (3) using a probit model instead.

The effect of an examiner's past approval rate on the probability of receiving a patent is not only large economically, it is also strong statistically, with F statistics well above the rule-of-thumb value of 10. This ensures that our results are not subject to weak-instrument bias.

2.3 Threats to identification

Conditional on patent applications being quasi-randomly assigned to patent examiners, the familiar exclusion restriction required for instrument validity is likely plausibly satisfied in our setting: it is difficult to see how an examiner's past leniency would affect a startup's future success *directly*, other than through its effect on the likelihood that the startup is granted a patent. Thus, threats to identification in our setting likely involve the possibility that examiner leniency correlates with omitted variables that in turn drive a startup's future success. This amounts to the possibility that our instrument is not, after all, "as good as randomly assigned conditional on covariates" (Angrist and Pischke 2009, p. 117). This would be the case if the characteristics of the application, the applicant, or the examiner influenced assignment.

We investigate this possibility in three ways. First, we argue—on institutional grounds and based on prior literature—that patent applications are indeed assigned quasi-randomly at the USPTO. Second, we implement a validation test of quasi-random assignment proposed by Righi and Simcoe (2017). Third, we test whether examiner leniency correlates with (observed) drivers of a startup's future success.

We are far from the first to argue that the USPTO assigns applications to examiners quasi-randomly (see Lichtman 2004, Sampat and Lemley 2010, Lemley and Sampat 2012, Sampat and Williams 2018, and Gaule 2018, among others). The precise details of the assignment process vary from art unit to art unit. For example, Lemley and Sampat (2012) report that some art units assign patents based on the last digit of the (randomly assigned) application serial number. Other

art units use a “first-in-first-out” rule: the application with the earliest filing date is assigned to the first available examiner. Importantly, though, the various assignment mechanisms in use at the USPTO are consistent with our identifying assumption that applications are assigned to examiners quasi-randomly (that is, randomly with respect to application or applicant quality). Reassuringly, Righi and Simcoe (2017) report finding “no evidence that particularly important applications (with large families) or broad applications (with short first independent claims) are assigned to specific examiners” (p. 2) and “no evidence that certain examiners specialize in applications that have greater importance or broader claims” (abstract).

Next, we implement a validation test of quasi-random assignment proposed by Righi and Simcoe (2017). As Righi and Simcoe note, under the null of quasi-random assignment, the first-stage coefficient estimate of θ in equation (3) should be invariant to the inclusion of controls for application, applicant, or examiner characteristics. We test this in the remainder of Table 2. Columns 2 through 4 investigate the possibility that examiner assignment depends on applicant or application quality. We use size (the log number of employees and log sales at first-action) and growth in employment and sales to proxy for applicant quality, and the log number of independent claims in the application to proxy for application quality. Columns 5 and 6 investigate the possibility of assignment based on examiner characteristics, specifically specialization (in columns 5 and 6) and experience and seniority (in column 6). We follow Righi and Simcoe, who suggest including technology-subclass-by-year fixed effects to control for specialization.²⁵ Data on examiner experience and seniority were made available to us by the USPTO pursuant to a Freedom of Information Act request.

²⁵ Technology subclass is the finest observable unit of technology classification at the USPTO. We include subclass-by-year fixed effects alongside the art-unit-by-year fixed effects of our baseline specification, which we can do because subclasses can span multiple art units. (In our sample, we have a total of 495 unique art units, 3,871 unique subclasses, and 9 years in the data.)

The reassuring upshot of Table 2 is that θ , the coefficient on the instrument, does not change across the six specifications: the coefficient estimates vary between 0.67 and 0.69. In light of this, our instrument passes Righi and Simcoe’s (2017) validation test. Table 2 further shows that our proxies for the quality of the applicant do not predict whether the application is approved (columns 2 and 3), and neither do our proxies for examiner experience and seniority (column 6).²⁶ Not surprisingly, we do find that applications with a larger number of claims, which we expect to be of higher quality, are more likely to be granted (column 4); however, this does not undermine identification insofar as the number of claims is orthogonal to examiner leniency, as suggested by the fact that the first-stage coefficient on the instrument does not change when the number of claims is included in equation (3)—and as we directly test next.

Table 3 tests whether examiner leniency correlates with observed characteristics of the applicant, the application, or the examiner. Neither the characteristics of the applicant (in columns 1 and 2) nor the characteristics of the application (in column 3) predict the leniency of the examiner to whom the application was assigned, consistent with quasi-random assignment. We do find that more experienced and more senior examiners are more lenient, all else equal (column 4), consistent with previous evidence reported by Lemley and Sampat (2012) and Frakes and Wasserman (2017). But this correlation does not invalidate our instrument if more experienced and more senior examiners are assigned applications quasi-randomly, as our Table 2 findings suggest.²⁷

The final two columns of Table 3 report the results of a placebo test. Some startups file for

²⁶ Similarly, we find no evidence that including examiner workload at the time an application is assigned to her affects the coefficient on the instrument. We thank an anonymous reviewer for suggesting this robustness test.

²⁷ In particular, it is worth noting that Table 2 shows that neither experience nor seniority have any explanatory power over and above leniency to predict whether an application is approved, and their inclusion in the first stage has no effect on the coefficient estimated for leniency in the first stage. Nonetheless, it is instructive to consider whether the second-stage results reported in the remainder of the paper are robust to controlling for examiner characteristics. As Tables IA.1-IA.4 in the Internet Appendix show, this is indeed the case.

patent protection not just in the U.S., but also in the European Union and/or Japan. Using foreign patent *grants* as a measure of the quality of the applicant or the underlying invention, we validate the quasi-random-assignment assumption by testing whether applications that are granted by a foreign patent office are more likely to have been assigned to more lenient U.S. examiners. Consistent with our identification assumption that leniency is uncorrelated with the quality of the application or the applicant, we find no such evidence.

2.4 Empirical results

2.4.1 Employment growth, sales growth, and firm survival

Table 4 reports our baseline results for the effects of approval of a startup's first patent application on its subsequent growth in employment and sales and its survival. Panel A focuses on employment.²⁸ The employment effects of quasi-randomly being granted a patent take some time to get going. After one year, a patent grant raises a startup's employment growth rate by a marginally significant 6.1 percentage points ($p=0.066$), all else equal.²⁹ After two years, the employment growth rates of successful and unsuccessful patent applicants differ by 22.8 percentage points, increasing to differences of 33.3, 48.9, and 54.5 percentage points after three, four, and five years, respectively (all significant at $p<0.01$ or better). Figure 2, Panel A visualizes this steady increase over time in the differential employment growth of successful and unsuccessful first-time patent applicants.

To gauge the economic significance, consider the median startup, which has eight employees at first-action. All else equal, patent approval results in the startup having 4.4 ($= 8 \times 0.545$) more

²⁸ To streamline the discussion, we report only 2SLS estimates. The corresponding OLS estimates can be found in the Internet Appendix, starting with Table IA.5. We discuss differences between the 2SLS and OLS estimates in Section 2.5.

²⁹ In addition to including art-unit-by-year and headquarter-state fixed effects, we also control for the log number of employees that the firm has at first-action (not shown). As expected, larger firms tend to grow more slowly.

employees five years later than if the application had been rejected as a result of being assigned to a stricter examiner. Summed over five years, a patent grant supports an additional 13.2 man-years of employment at the median startup. The equivalent numbers for the average successful applicant are 16.1 more employees by year 5 and 49 additional man-years of employment.

Table 4, Panel B and Figure 2, Panel B show a similar pattern for sales growth. On average, successful applicants grow their sales by a cumulative 10.0, 27.6, 51.2, 79.6, and 79.5 percentage points more than do their unsuccessful counterparts over the one through five years following patent approval ($p < 0.01$ or better from year 2 onwards). For the median startup, which in our sample has sales of \$800,000 at first-action, patent approval leads to sales being \$636,000 ($= \$800,000 \times 0.795$) higher in year 5 than if the application had been rejected. Summed over the five years, a patent grant allows the median startup to make an additional \$2 million in sales, all else equal. The equivalent numbers for the average successful applicant are \$3.4 million higher sales by year 5 and a cumulative \$10.6 million more in sales over the five years.

Table 4, Panel C shows the extensive-margin effect of a patent grant on a startup's probability of survival over the next few years. We code a firm as being alive in year $t + k$ if it continues to be included in the NETS database that year.³⁰ Compared to the intensive-margin effects on growth, the effect of a first patent grant on survival is quite modest. The estimates are economically small (measuring less than five percentage points) and statistically significant only at the three-year horizon ($p = 0.045$).

³⁰ Dun & Bradstreet (D&B), the source of the NETS database, is careful about identifying firm deaths and distinguishing them from relocations. See Neumark, Zhang, and Wall (2005) for a comprehensive account of D&B's methodology. We spot-check firm deaths according to NETS with information obtained from Capital IQ, which appears in text form (e.g., "On April 22, 2009, Technologies to Be, Inc. filed a voluntary petition for liquidation under Chapter 7 in the U.S. Bankruptcy Court for the Middle District of Florida.").

2.4.2 Potential confound: review speed

A causal interpretation of the intensive-margin estimates in Table 4 requires that our leniency instrument is uncorrelated with omitted variables that affect firm outcomes. A possible violation of this condition could arise if lenient examiners were also faster and review speed had an independent effect on outcomes. If so, our patent approval estimates would confound the benefits of approval with the benefits of speedy resolution of uncertainty about the application's fate. We investigate this concern by adding review speed (measured as the time from patent filing to first-action decision) as a control variable to equation (1).

Given that review speed is plausibly endogenous, we follow Hegde and Ljungqvist (2017) and instrument it using a measure of administrative delays. The instrument builds on the notion that the time it takes to receive a first-action decision on a patent application can be decomposed into two parts: the time from filing to the application being assigned to the examiner's docket (which reflects quasi-random administrative delays at the USPTO that are unrelated to invention quality or application complexity), and the time the examiner takes from the docket to the first-action date (which is potentially endogenous). Thus, we instrument review speed using the sum of the time an application takes from filing to being assigned to an examiner's docket, and the average time that examiner has taken in the past from docket to first-action.³¹

As Table IA.6 in the Internet Appendix shows, controlling for instrumented review speed yields point estimates of the patent approval effect on employment and sales growth that are well within the 95% confidence intervals around the Table 4 coefficients.

³¹ The rationale for using the examiner's average past review speed to construct this review-speed instrument is similar to that for using examiner leniency to instrument approval: both instruments exploit inherent characteristics of the examiners which, given the quasi-random assignment of applications to examiners, should be orthogonal to application and applicant quality.

2.4.3 Follow-on innovation

How does the outcome of a startup's first patent application affect its ability to continue innovating? We capture a startup's follow-on innovation using the log number of patent applications it files after the first-action decision on the first application; the log number of such subsequent applications that are approved; the approval rate of subsequent applications; the log number of citations received by all subsequent applications combined; and the log average number of citations per subsequent approved patent. (See Table 1 for descriptive statistics.) Since these models do not require NETS data, we use the full sample of startups. As in our baseline specification (1), we include art-unit-by-year and headquarter-state fixed effects.

Table 5 reports the results. Columns 1 and 2 show that approval of its first patent application leads to a 56.5% ($=e^{0.448}-1$) increase in the number of patents the startup subsequently applies for and a 42.3% increase in the number of patents it subsequently obtains ($p<0.001$ in each case). This may not be surprising; after all, Table 4 shows that successful applicants enjoy faster growth, and that faster growth may, directly or indirectly, support investment in R&D. Column 3 shows that it is not just the volume of subsequent patent applications and grants that increases: success in the first application leads to a 24.4 percentage-point increase in the approval *rate* of a startup's subsequent applications ($p<0.001$), suggesting that later applications may be of higher quality. Consistent with this interpretation, we find that a first patent grant boosts the number of citations received by the patents the firm is subsequently granted, both overall (up by 60.3% in column 4) and per patent (up by 33% in column 5).

2.4.4 Subsequent patent applications

Next, we investigate whether the approval of subsequent patent applications delivers similarly large benefits to startups or whether a startup's first patent is in some way special.

Table 6, Panel A shows that obtaining a second patent has a much smaller effect on employment growth than did the first patent. After one year, a successful second application increases a startup's employment growth by a marginally significant 10.8 percentage points ($p=0.06$), all else equal. Over longer time horizons, the effect is neither economically nor statistically significant. Table 6, Panel B reports economically large effects on sales growth over four and five years (measuring 20.2 and 49.6 percentage points, respectively), but these effects are noisily estimated and not statistically significant at conventional levels.³²

It is only with regards to follow-on innovation that a startup's second patent has reliably beneficial effects. Table 6, Panel C shows positive and statistically significant effects on the number of future patent applications (up by 64.2%) and approved patents (up by 49.8%), on the success rate of future patent applications (up by 14.1 percentage points), and on citations to future patents (up by 75.6%) ($p<0.01$ in each case), though not on average citations per patent.

The effects of approval of a third patent application follow a similar pattern: no effect on growth in sales or employment and a strong positive effect on follow-on innovation (not shown).

Overall, these results suggest that the benefits startups derive from obtaining patent grants (over and above the benefits they derive from the underlying inventions) change over time. While a successful first application provides a substantial boost to both growth and follow-on innovation, later patent grants only boost follow-on innovation. One plausible interpretation is that startups, once they have embarked on a high-growth trajectory thanks to winning their first patent lottery, are able to find alternative ways to keep their growth rates up. In other words,

³² Not even when we restrict the sample to startups whose first patent application was rejected do we find evidence that a successful second patent application boosts growth (see Table IA.10 in the Internet Appendix). In that sense, at least, startups that "lose the patent lottery" in their first application do not recover even when they later manage to obtain a patent in their second application. Having the first patent application rejected can thus have long-lasting negative consequences for a startup.

success in a startup's first patent application is indeed special, at least as far as helping to launch accelerated growth is concerned.

2.4.5 IT startups

Two influential surveys report that established U.S. companies in the IT sector consider patents the least effective mechanism to protect their R&D investments (Levin et al. 1987, Cohen, Nelson, and Walsh 2000). Instead, respondents rely on trade secrets, lead time, and design and manufacturing capabilities to profit from their inventions. IT products also tend to be complex and involve multiple patentable elements, making it difficult for any single patent to effectively protect a product (Cohen, Nelson and Walsh 2000). Table IA.9 in the Internet Appendix revisits these survey findings using our data on startups, showing that IT startups derive sizable benefits in terms of sales growth and follow-on innovation, contrary to the survey evidence from established IT firms. A plausible explanation for the difference in importance is that IT startups use patents as bargaining chips in licensing negotiations or as defensive shields that preserve their independence to operate in product markets (Hall and Ziedonis 2001).

2.5 External validity

We report the OLS counterparts to our 2SLS estimates of the effects of patents on growth and follow-on innovation in the Internet Appendix. A comparison of Tables 4 and 5 to Tables IA.5 and IA.7 reveals that the OLS estimates are smaller than the 2SLS estimates. To understand why this is the case, we need to be clear about whose outcomes our instrument identifies.

When treatment effects are heterogeneous (i.e., they affect firms differently), an instrument identifies a local average treatment effect (LATE). This means that our 2SLS estimates identify the effects of patents only for the subpopulation of startups whose patent application is affected by their examiner's leniency. Applications that are obviously bad will be rejected even by lenient

examiners; in the language of the treatment effects literature, such applicants are “never-takers.” Applications that are obviously great will be approved even by harsh examiners; such applicants are hence “always-takers.” The compliant subpopulation that we expect to be responsive to our instrument is thus the group of startups with applications of middling quality.³³ Accordingly, the leniency instrument identifies the effect of patent approval on firms with middling inventions.

An important consequence is that our 2SLS estimates, being local average treatment effects, cannot be generalized to the average startup. In fact, it is likely that our LATE estimates overstate the value of patents for the average startup. This is because we expect smaller treatment effects among never-takers and among always-takers than among the compliant subpopulation of startups with middling inventions. Startups with clearly bad inventions are unlikely to prosper even if they were somehow granted a patent; startups with clearly great inventions are likely to prosper regardless of whether they receive a patent. Average treatment effects (or ATE) are hence likely smaller than our LATE estimates.

This discussion helps explain why our 2SLS estimates in Tables 4 and 5 are larger than their OLS counterparts. By definition, OLS estimates equal ATE plus the selection bias. We expect the selection bias to be positive and thus our OLS estimates to be upwardly biased estimates of the ATE of patent grants: applicants of higher unobserved quality are both more likely to have produced a “novel, useful, and non-obvious” invention worthy of a patent and to perform better going forward, regardless of the patent. Despite this positive selection bias, 2SLS estimates can be larger than their OLS counterparts if the difference between our estimated LATE and the ATE of a patent grant is larger than the selection bias, that is, if $LATE > (ATE + \text{selection bias})$. This

³³ There is no reason to expect the presence of “defiers” in our setting: no startup should be *less* likely to receive a patent as a result of being randomly assigned to a more lenient examiner. The monotonicity assumption that is necessary for our IV estimates to be LATE is hence likely satisfied in our setting.

appears to be the case in our setting, which is consistent with our expectation that compliant startups benefit more from a patent grant than the average startup.³⁴

3. What drives the real effects of patents?

Our findings so far establish that winning the patent lottery (first time around) helps startups grow and innovate. As Figure 2 shows, the gains in employment and sales growth are not instantaneous. The same is true for follow-on innovation: the average sample startup submits its next patent application 1.5 years after first-action on its first application. These time lags suggest that the mechanism linking patent approval to positive firm-level outcomes is one that requires time to take effect. A prime candidate mechanism is access to external capital. Transforming a startup's patented ideas into new products and processes that eventually support jobs and yield revenues typically requires investments in operations and marketing, and these investments in turn require upfront funding.

In this section, we investigate whether patents facilitate access to three sources of capital: venture capital (VC), public investors in the IPO market, and specialized lenders. A significant body of research discusses the importance of these funding sources for startups (Gorman and Sahlman 1989, Hellmann and Puri 2000, Gompers and Lerner 2001, Hochberg, Ljungqvist, and Lu 2007, Bernstein, Giroud, and Townsend 2016, and Hochberg, Serrano, and Ziedonis 2018). However, establishing a causal link between a patent grant and funding is subject to the same identification challenges researchers face when measuring the effect of patents on growth and follow-on innovation. Our identification strategy helps overcome these challenges.

Specifically, we estimate 2SLS linear-probability models of the following general form:

$$\text{Capital access}_{ijat+k} = \beta \text{Patent application approved}_{ijat} + \Phi X_{ija\tau} + v_{a\tau} + \varepsilon_{ijat+k} \quad (4)$$

³⁴ Card (2001) uses the same argument to explain why IV estimates of the returns to schooling tend to be larger than OLS estimates, even though OLS likely produces upward-biased estimates of the ATE of schooling on earnings.

separately for each of our three capital sources, (i) VC funding, (ii) an initial public offering of equity, or (iii) a loan obtained by pledging a patent or patent application as collateral. We again use the examiner's prior approval rate to instrument for the likelihood that the application is approved and report OLS results in the Internet Appendix for completeness. As before, we include art-unit-by-year and headquarter-state fixed effects. In addition, we control for the log number of prior VC rounds the startup has raised when we estimate the probability of receiving funding from a VC or through an IPO.³⁵

3.1 Venture capital funding

Table 7 reports the results for VC funding. Approval of a firm's first patent application causes a startup's chances of obtaining VC funding in the following year to increase by 1.7 percentage points ($p=0.053$ in column 1). Extending the window increases the effect to 2.7, 3.0, 3.4, and 3.6 percentage points over two, three, four, and five years, respectively ($p<0.01$). These effects are economically large. To illustrate, the 3 percentage-point increase in column 3 represents a 46.9% increase relative to the 6.4% unconditional probability of a sample firm raising VC funding in the three years following the first-action decision.

These estimates point to a steep change in a startup's ability to raise VC funding around one to two years after the first-action date. The raw data support the inference that startups that win the patent lottery tend to raise VC funding quite quickly after the first-action decision: the median successful applicant that raises VC funding during our five-year window does so a mere 10 months after the first-action date. Figure 3 illustrates the extent to which fundraising events

³⁵ Of the 34,215 startups in our sample, 92.5% had raised no VC funding before the first-action date. For these, equation (4) identifies the effect of patent approval on their ability to raise their first VC round. For firms with at least one prior VC round, equation (4) identifies the effect on their ability to raise a follow-on round. Specifically, 2.4% of our sample firms had raised one VC round before their first-action; 1.9% had raised two prior VC rounds; 1.4% had raised three prior rounds; and the remaining 1.8% had raised four or more prior rounds. Pledging a patent application as collateral before the first-action date is virtually unheard of, and none of the sample startups had filed for an IPO before first-action.

are bunched shortly after first-action. This timing fits our conjectured mechanism well: startups use patent grants to first raise external capital; they then use the capital to fund investments in operations and marketing to turn their patented ideas into new products and processes which subsequently yield increases in sales.

Of course, this is not to say that VCs only help startups by supplying capital. In addition to capital, VCs have been shown to provide monitoring and advice (Hellmann and Puri 2002, Bernstein, Giroud, and Townsend 2016),³⁶ access to networks of potential customers, suppliers, and strategic partners (Hochberg, Ljungqvist, and Lu 2007), and help recruiting talented staff (Gorman and Sahlman 1989).

3.2 Fundraising in the IPO market

While the IPO market is sometimes viewed narrowly as a venue in which VCs “exit” their investments in startups, its dominant function in the U.S. has in fact traditionally been to help firms raise relatively large amounts of new capital at relatively low cost.³⁷ We thus investigate the extent to which patents facilitate startups’ access to the IPO market. 267 (or 0.78%) of the 34,215 startups in the sample go public between their (first) first-action date and November 2018, with the median startup taking 5.5 years to do so.³⁸ Column 6 of Table 7 shows that a successful first patent application boosts a startup’s chances of raising capital through an IPO by one percentage point ($p=0.048$), a 128% increase over the unconditional probability of 0.78%.

³⁶ For example, VCs could help a startup by sharing their experience on how to commercialize patents.

³⁷ Ljungqvist and Wilhelm (2003) report that 85.6% of VC-backed firms that went public between 1996 and 2000 saw the VCs sell no shares at all in the IPO. Including non-VC-backed firms, 92.5% of shares sold to IPO investors were “primary,” representing fundraising from IPO investors, rather than “secondary” sales by existing investors.

³⁸ Firms that remain private by November 2018 may yet go public in the future. The traditional way to deal with right-censoring of this kind is to estimate a duration model. However, duration models cannot accommodate instrumental-variable techniques, so we estimate a simple linear-probability model instead. The model includes art-unit-by-application-year fixed effects, which control for the fact that startups that applied for their first patent in the later years of our sample have had less time to go public than earlier applicants.

3.3 Loans from banks and specialized lenders

As Hochberg, Serrano, and Ziedonis (2018) document, firms frequently pledge their patent rights as collateral for loans obtained from banks or specialized patent lenders. What is less well known is that firms can also pledge pending patent *applications* (after first-action but before final approval) and even *rejected* patent applications. Typically, firms pledge not single patent rights (or applications) but a bundle of their intangible assets, which can include either accepted or rejected applications alongside trademarks, copyrights, etc. A rejected patent application may not have passed the standards of patentability but is still considered prior art and may serve as a description of the holder's intellectual assets, particularly when offered as part of a bundle.

This observation is important, because it allows us to estimate whether startups have an easier time obtaining a loan collateralized by their patent application if the patent application was approved rather than denied: even those with denied patent applications can, in principle, obtain such a loan, but presumably less easily. As expected, then, Table 8, Panel A shows that startups whose first patent application is approved are 8.6 percentage points ($p < 0.001$) more likely to pledge that application as collateral than are startups whose first application is rejected, a 119% increase over the unconditional sample probability of 7.2%. We find similar results in Panel B, which excludes sample startups that receive VC funding. For this restricted sample of startups, a successful patent application increases the probability of securing a collateralized patent loan by 5.9 percentage points ($p < 0.001$), a 93% increase over the unconditional probability of 5.5%.

4. How do patents facilitate access to capital?

Why do suppliers of external funding favor startups that have secured patents? Prior work suggests several factors may be at play. First, to the extent that patents are monopoly rights, the market power they confer may increase profits, thus making startups with patents more attractive

to investors. Second, patents can facilitate transactions such as licensing agreements and alleviate investors' concerns regarding the firm's ability to monetize an invention (Arora, Fosfuri, and Gambardella 2001). Third, entrepreneurs with patents may be more willing to share the details of their invention with investors without fear of expropriation (Arrow 1962, Anton and Yao 1994, Biais and Perotti 2008). Fourth, the patent application itself can help to credibly communicate the technical details of the invention (Hegde and Luo 2018). Finally, since patents are certified to have cleared the USPTO's requirements of novelty, non-obviousness, and utility, they can serve as a quality signal to investors in a market characterized by information asymmetries that are often severe (Long 2002, Hsu and Ziedonis 2013).³⁹

This section explores whether financiers respond to patents differently under different circumstances, which in turn may shed light on the reasons behind their preference for startups that have been granted a patent. We readily acknowledge that these tests are suggestive in nature: discriminating among the five or more factors at play is obviously challenging. To keep the amount of hand-collection of the required data manageable, we focus on estimating heterogeneous effects in the context of a startup's ability to raise VC funding (rather than an IPO or a collateralized patent loan). To conserve space, we report results for a three-year window following first-action; our conclusions are robust to using alternative time windows.

4.1 Variation in VC funding round

The first three columns of Table 9 split startups by the number of VC rounds raised before (first) first-action. Though somewhat noisily estimated, this exercise reveals that VCs appear to pay attention to a first patent grant for both new prospects and existing portfolio companies.

³⁹ Melero, Palomeras, and Wehrheim (2017) find that patent grants reduce the mobility of employee inventors, which they argue reflects the fact that patents make inventors' innovation implementation skills more specific to their current employer. Our findings in Section 3 suggest another complementary explanation: by facilitating startups' access to capital, patents may make it easier for them to retain their inventor talent.

Specifically, for startups without prior VC funding, approval increases the likelihood of raising VC funding by 1.3 percentage points ($p=0.08$ in column 1), or 76.5% from the unconditional probability of 1.7%. For startups with one prior VC round under their belt, approval increases the likelihood of raising a second VC round by 45.5 percentage points ($p=0.048$ in column 2), or 97.2% from the conditional probability of 46.8%. And for startups with two prior VC rounds, approval increases the likelihood of raising further VC funding by 26.9 percentage points ($p=0.059$ in column 3), or 43.7% from the unconditional probability of 61.5%.

The fact that the patent grant effect is largest (in relative terms) for firms raising their first or second VC round is consistent with the notion that informational asymmetries between VCs and startups are highest when raising early rounds, at which point patents serve as a screening device and/or a way to mitigate Arrow's (1962) disclosure paradox.⁴⁰ At the same time, the continued large effect of a first patent grant on the likelihood of raising funding for firms that have already raised two or more prior funding rounds suggests that at later stages, patents are valued for the monopoly rights they entail.

4.2 Variation in prior entrepreneurial experience

A potentially useful proxy for the uncertainty and information frictions surrounding a startup is the experience of its founders: all else equal, startups run by experienced entrepreneurs are less risky and thus easier for investors to finance (Hsu and Ziedonis 2013). To code prior founder experience, we use hand-collected data from Capital IQ for startups that raise VC funding at some point in their lives. This restricts the sample to firms with at least one prior VC round

⁴⁰ The large approval effect we observe for firms trying to raise a second VC round is consistent with the notion that the first (or seed) round is typically small, and so VCs may not do a lot of due diligence before their initial investment. Indeed, Ewens, Nanda, and Rhodes-Kropf (2018) show that VCs increasingly follow a “spray and pray” strategy, making small seed investments in a large number of startups. Our results suggest that VCs may use patent outcomes when deciding which startups they support beyond the seed round.

before first-action. Of these firms, 57% have a founding team with at least one experienced founder, while the rest are run by teams made up exclusively of first-time entrepreneurs.

Column 4 of Table 9 allows the effect of a patent grant on the likelihood of VC funding to vary with prior founder experience. The results confirm that patent approval facilitates access to capital the most for inexperienced founders. Patent approval increases a startup's likelihood of raising VC funding in the next three years by nearly 33 percentage points among inexperienced founders ($p=0.053$); for experienced founder teams, the effect is virtually zero ($p=0.989$).

4.3 Variation in startup agglomeration across U.S. states

Two facts combine to suggest that the value of a patent grant in obtaining VC funding varies geographically. First, VCs have a well-known preference for investing locally (Lerner 1995, Sorenson and Stuart 2001). Second, startup activity varies considerably across the country, with hotspots like California, Massachusetts, and New York being particularly popular places to start an innovative business. Combined, these two facts imply that VCs operating in areas with larger startup populations have more potential investments to choose among than those operating in areas with fewer startups. To deal with the larger number of investments to screen, VCs may rely more on easily observable signals such as patent grants in areas with high startup activity.⁴¹

Column 5 of Table 9 allows the effect of a patent grant on the likelihood of raising VC funding to differ between startups headquartered in a state with an above- or below-median startup agglomeration in the year of a startup's first patent application.⁴² This yields results consistent with the idea that patents play a key role in helping startups located in hubs of

⁴¹ This is likely to be the case even if such areas also concentrate a large number of VCs, given that even in high-VC-density areas like California, most VC firms tend to be small (Ewens and Rhodes-Kropf 2015) and competing VC firms do not coordinate their screening efforts.

⁴² We measure startup agglomeration using the number of first-time patent applicants in the state. We obtain similar results if we simply code California, Massachusetts, and New York (which have consistently been the three states with the most startup activity according to the 2016 NVCA Yearbook) as states with high startup agglomeration.

innovative activity to stand out from the crowd. In states with below median startup activity, the patent approval effect on the likelihood of raising VC funding is small (1.6 percentage points) and insignificant ($p=0.157$). In startup hubs like California, Massachusetts, or New York, the patent approval effect is 2.7 percentage points larger ($p=0.012$), a sizeable difference relative to the unconditional probability of 6.4%. Interestingly, startups whose first application is rejected are 1.6 percentage points *less* likely to raise VC funding in a startup hub than in a non-hub state ($p=0.058$), which suggests that it is indeed *a priori* harder to stand out in such hubs.

4.4 Variation across industries

IT and biochemistry have, for a long time, been the main focus of VCs in the U.S. (Gompers and Lerner 2001, Graham et al. 2009). There are reasons to expect the information value of a patent to be different in these two industries (Cohen, Nelson, and Walsh 2000). IT startups tend to be founded by younger entrepreneurs (Ewens, Nanda, and Rhodes-Kropf 2018) and their inventions often face substantial demand uncertainty and imitation risk. Thus, a favorable decision on an IT startup's first patent application can provide a particularly valuable early signal about the quality of its technology and its founders, while also allowing the founders to more freely discuss their idea with VCs without the fear of expropriation. Evidence from interviews at semiconductor firms suggests that the primary function of a patent in that industry is “securing capital from private investors [for firms] in the startup phase” (Hall and Ziedonis 2001).⁴³

Biochemistry startups, in contrast, tend to be founded by experienced scientists, the quality of whose research can more readily be evaluated using a variety of sources such as academic publications or National Institutes of Health grants (Li and Agha 2015). Biochemistry startups face relatively little demand uncertainty or risk of imitation, with the greatest uncertainty instead

⁴³ Hochberg, Serrano, and Ziedonis (2018) and Mann (2018) document the existence of a well-developed secondary market for IT patents, which alleviates investors' downside risk if the firm ends up being unviable.

coming from the probability of technical success and the regulatory process (DiMasi, Hansen, and Grabowski 2003). As a result, early patent decisions reveal relatively little information about the quality of the founders or the potential commercial success of their inventions.

Consistent with this discussion, column 6 of Table 9 shows that the approval of an IT firm's first patent application increases its probability of raising VC funds in the next three years by 6.3 percentage points ($p=0.003$). In biochemistry, on the other hand, patent approval has no statistically significant effect on VC funding.⁴⁴

4.5 Subsequent patent applications

Column 7 of Table 9 reports the effects of the second patent application on the probability of the startup raising any VC funding in the next three years. Although the coefficient estimate for the second patent grant is comparable in magnitude to the corresponding estimate for the first patent grant in Table 7 (2.8 versus 3.0 percentage points), it is not statistically significant ($p=0.139$). This pattern mirrors the finding reported in Section 2.4.4 that a startup's first patent is special: later patents do not boost growth. However, when we restrict the sample to startups whose first patent application was rejected, we find economically large (albeit statistically insignificant) evidence that a successful second patent application – i.e., the firm's first patent grant – may facilitate access to VC funding ($p=0.113$ in column 8).

5. Conclusions

We estimate the magnitude of the economic benefits (in terms of growth and follow-on innovation) that U.S. startups derive from obtaining an early patent grant, over and above the

⁴⁴ In fact, the sum of *First patent application approved* and *First patent application approved* \times *biochemistry* in column 6 is negative, though not significantly so ($p=0.109$). “Other industries” behave like IT firms, consistent with the fact that many of the technological groups that fall in the “other” category (in particular, electronic commerce, construction, mechanical engineering, transportation, and manufacturing) tend to be closer to IT than biochemistry in that their founders are usually not well-known, experienced scientists.

economic benefits startups derive from their inventions. To disentangle the value of patents as distinct from the value of the underlying inventions and to identify causal effects, we use plausibly exogenous variation in patent approvals generated by the random allocation of patent applications to examiners with varying propensity to approve applications at the USPTO. Our analysis shows that patent approvals have a substantial and long-lasting impact on startups, especially the first time startups apply for patent protection: startups whose first patent application is approved create more jobs, enjoy faster sales growth, and are more innovative than startups with only randomly different inventions that fail to win patent protection. Successful second and third applications only significantly boost follow-on innovation but not growth.

We relate the positive effects of patents to their role in facilitating startups' access to capital, which helps startups turn ideas into products and products into revenues. Patents are particularly beneficial for startups founded by inexperienced entrepreneurs, for those located in states with many startups competing for funding, and for IT startups.

While our results by no means rule out the existence of negative effects of patents, they show that patents convey substantial economic benefits on startups, not least by facilitating contracting between them and their investors. We emphasize that our findings do not imply that the patent system is optimal, or even net welfare-increasing. Instead, taken together with empirical analyses of the patent system's spillover effects on other firms (Griliches 1984, Jaffe 1986, Heller and Eisenberg 1998, Williams 2013, Galasso and Schankerman 2015, Sampat and Williams 2018), our findings on the beneficial effects of patents on their owners contribute to a more complete understanding of the patent system. Specifically, we disentangle the benefits of patent rights from those of the underlying inventions they protect identify an important channel through which how these benefits operates. Reforms of the patent system that do not take this channel into

account run the risk of stifling growth by negatively impacting the availability of capital for innovative startups.

References

- Angrist, Joshua D., and Jörn-Steffen Pischke, 2009, *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton, NJ: Princeton University Press).
- Anton, James J., and Dennis A. Yao, 1994, "Expropriation and Inventions: Appropriable Rents in the Absence of Property Rights," *American Economic Review* 84, 190–209.
- Arora, Ashish, Andrea Fosfuri, and Alfonso Gambardella, 2001, *Markets for Technology: The Economics of Innovation and Corporate Strategy* (Cambridge, MA: MIT Press).
- Arrow, Kenneth, 1962, "Economic Welfare and the Allocation of Resources for Invention," in: Harold M. Groves, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors* (Princeton, NJ: Princeton University Press).
- Balasubramanian, Natarajan, and Jagadeesh Sivadasan, 2011, "What Happens When Firms Patent? New Evidence from U.S. Economic Census Data," *Review of Economics and Statistics* 93, 126–146.
- Becker, Gary, and Richard A. Posner, 2013, "On Reforming the Patent System," available at <http://www.becker-posner-blog.com/2013/07/on-reforming-the-patent-system-becker.html>.
- Bernstein, Shai, Xavier Giroud, and Richard Townsend, 2016, "The Impact of Venture Capital Monitoring," *Journal of Finance* 71, 1591–1622.
- Biais, Bruno, and Enrico Perotti, 2008, "Entrepreneurs and New Ideas," *RAND Journal of Economics* 39, 1105–1125.
- Card, David, 2001, "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems," *Econometrica* 69, 1127–1160.
- Carley, Michael, Deepak Hegde, and Alan Marco, 2015, "What is the Probability of Receiving a U.S. Patent?," *Yale Journal of Law and Technology* 17, 204–223.
- Cockburn, Iain M., Samuel Kortum, and Scott Stern, 2003, "Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes," in: Wesley M. Cohen and Stephen A. Merrill, ed., *Patents in the Knowledge-Based Economy* (Washington, D.C.: National Academies Press).
- Cohen, Wesley, Richard Nelson, and John Walsh, 2000, "Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not)," NBER Working Paper No. 7552.
- Conti, Annamaria, Jerry Thursby, and Marie Thursby, 2013, "Patents as Signals for Startup Financing," *Journal of Industrial Economics* 61, 592–622.
- DiMasi, Joseph A., Ronald W. Hansen, and Henry G. Grabowski, 2003, "The Price of Innovation: New Estimates of Drug Development Costs," *Journal of Health Economics* 22, 151–185.
- Ewens, Michael, and Matthew Rhodes-Kropf, 2015, "Is a VC Partnership Greater Than the Sum of Its Partners?," *Journal of Finance* 70, 1081–1113.
- Ewens, Michael, Ramana Nanda, and Matthew Rhodes-Kropf, 2018, "Cost of Experimentation and the Evolution of Venture Capital," *Journal of Financial Economics* 128, 422–442.
- Federal Trade Commission, 2018, "The Evolving IP Marketplace: Aligning Patent Notice and

- Remedies with Competition,” (2011), retrieved Dec 12, 2018 from <https://www.ftc.gov/sites/default/files/documents/reports/evolving-ip-marketplace-aligning-patent-notice-and-remedies-competition-report-federal-trade/110307patentreport.pdf>.
- Frakes, Michael, and Melissa F. Wasserman, 2017, “Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Microlevel Application Data,” *Review of Economics and Statistics* 99, 550–563.
- Galasso, Alberto, and Mark Schankerman, 2015, “Patents and Cumulative Innovation: Causal Evidence from the Courts,” *Quarterly Journal of Economics* 130, 317–369.
- Gaule, Patrick, 2018, “Patents and the Success of Venture-Capital Backed Startups: Using Examiner Assignment to Estimate Causal Effects,” *Journal of Industrial Economics* 66, 350–376.
- Gompers, Paul, and Josh Lerner, 2001, “The Venture Capital Revolution,” *Journal of Economic Perspectives* 15, 145–168.
- Gorman, Michael, and William A. Sahlman, 1989, “What Do Venture Capitalists Do?,” *Journal of Business Venturing* 4, 231–248.
- Graham, Stuart J.H., and Ted M. Sichelman, 2008, “Why Do Start-Ups Patent?” *Berkeley Technology Law Journal* 23, 1064–1097.
- Graham, Stuart J.H., Robert P. Merges, Pamela Samuelson, and Ted M. Sichelman, 2009, “High Technology Entrepreneurs and the Patent System: Results of the 2008 Berkeley Patent Survey,” *Berkeley Technology Law Journal* 24, 255–327.
- Griliches, Zvi, 1984, *R&D, Patents and Productivity* (Chicago, IL: University of Chicago Press).
- Hall, Bronwyn H., Adam Jaffe, and Manuel Trajtenberg, 2005, “Market Value and Patent Citations,” *RAND Journal of Economics* 36, 16–38.
- Hall, Bronwyn H., and Rosemarie H. Ziedonis, 2001, “The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995,” *RAND Journal of Economics* 32, 101–128.
- Hegde, Deepak, and Alexander Ljungqvist, 2017, “Patent Scope and Speed,” Working Paper.
- Hegde, Deepak and Hong Luo, 2018, “Patent Publication and the Market for Ideas,” *Management Science* 64, 495–981.
- Heller, Michael A., and Rebecca S. Eisenberg, 1998, “Can Patents Deter Innovation? The Anticommons in Biomedical Research,” *Science* 280, 698–701.
- Hellmann, Thomas, and Manju Puri, 2000, “The Interaction Between Product Market and Financing Strategy: The Role of Venture Capital,” *Review of Financial Studies* 134, 959–984.
- Hochberg, Yael V., Alexander Ljungqvist, and Yang Lu, 2007, “Whom You Know Matters: Venture Capital Networks and Investment Performance,” *Journal of Finance* 62, 251–301.
- Hochberg, Yael V., Carlos J. Serrano, and Rosemarie H. Ziedonis, 2018, “Patent Collateral, Investor Commitment, and the Market for Venture Lending,” *Journal of Financial Economics* 130, 74–94.
- Hsu, David H., and Rosemarie H. Ziedonis, 2013, “Resources as Dual Sources of Advantage:

- Implications for Valuing Entrepreneurial Firm Patents,” *Strategic Management Journal* 34, 761–781.
- Jaffe, Adam, 1986, “Technological Opportunity and Spillovers of R&D: Evidence from Firms’ Patents, Profits and Market Value,” *American Economic Review* 76, 984–1001.
- Jaffe, Adam B., and Josh Lerner, 2004, *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do About It* (Princeton, NJ: Princeton University Press).
- Johnson, Simon, John McMillan, and Christopher Woodruff, 2002, “Courts and Relational Contracts,” *Journal of Law, Economics, and Organization* 18, 221–77.
- Kerr, William, and Shihe Fu, 2008, “The Survey of Industrial R&D-Patent Database Link Project,” *Journal of Technology Transfer* 33, 173–186.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, “Technological Innovation, Resource Allocation, and Growth,” *Quarterly Journal of Economics* 132, 665–712.
- Lamoreaux, Naomi R., and Kenneth L. Sokoloff, 1999, “Inventors, Firms, and the Market for Technology in the Late Nineteenth and Early Twentieth Centuries,” in: N.R. Lamoreaux, D.M.G. Raff, and P. Temins, eds., *Learning by Doing in Markets, Firms, and Countries* (Chicago IL: Chicago University Press).
- Lemley, Mark A., 2001, “Rational Ignorance at the Patent Office,” *Northwestern University Law Review* 95, 1497–532.
- Lemley, Mark A., and Bhaven N. Sampat, 2008, “Is the Patent Office a Rubber Stamp?,” *Emory Law Journal* 58, 181–209.
- Lemley, Mark A., and Bhaven N. Sampat, 2012, “Examiner Characteristics and Patent Office Outcomes,” *Review of Economics and Statistics* 94, 817–827.
- Lemley, Mark A., and Carl Shapiro, 2005, “Probabilistic Patents,” *Journal of Economic Perspectives* 19, 75–98.
- Lerner, Josh, 1995, “Venture Capitalists and the Oversight of Private Firms,” *Journal of Finance* 50, 301–318.
- Lerner, Josh, and Amit Seru, 2015, “The Use and Misuse of Patent Data: Issues for Corporate Finance and Beyond,” NBER Working Paper No. 24053.
- Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter, 1987, “Appropriating the Returns from Industrial Research and Development,” *Brookings Papers on Economic Activity* 3, 783–831.
- Li, Danielle, and Leila Agha, 2015, “Big Names or Big ideas: Do Peer-Review Panels Select the Best Science Proposals?,” *Science* 348, 434–438.
- Lichtman, Douglas, 2004, “Rethinking Prosecution History Estoppel,” *University of Chicago Law Review* 71, 151–182.
- Ljungqvist, Alexander, and William Wilhelm, 2003, “IPO Pricing in the Dot-Com Bubble,” *Journal of Finance* 63, 723–752.
- Long, Clarisa, 2002, “Patent Signals,” *University of Chicago Law Review* 69, 625–679.

- Mann, William, 2018, "Creditor Rights and Innovation: Evidence from Patent Collateral," *Journal of Financial Economics* 130, 25–47.
- Melero, Eduardo, Neus Palomeras, and David Wehrheim, 2017, "The Effect of Patent Protection on Inventor Mobility," *Academy of Management Proceedings* 1.
- Neumark, David, Junfu Zhang, and Brandon Wall, 2005, "Employment Dynamics and Business Relocation: New Evidence from the National Establishment Time Series," NBER Working Paper No. 11647.
- Nordhaus, William, 1969, *Invention, Growth, and Welfare: A Theoretical Treatment of Technological Change* (Cambridge, MA: MIT Press).
- Righi, Cesare, and Timothy Simcoe, 2017, "Patent Examiner Specialization," NBER Working Paper No. 23913.
- Sampat, Bhaven N., and Mark A. Lemley, 2010, "Examining Patent Examination," *Stanford Technology Law Review*.
- Sampat, Bhaven N., and Heidi L. Williams, 2018, "How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome," *American Economic Review*, forthcoming.
- Schankerman, Mark, 1998, "How Valuable is Patent Protection? Estimates by Technology Field," *RAND Journal of Economics* 29, 77–107.
- Scherer, Frederic, 1972, "Nordhaus' Theory of Optimal Patent Life: A Geometric Reinterpretation." *American Economic Review* 62, 422–427.
- Sorenson, Olav, and Toby Stuart, 2001, "Syndication Networks and the Spatial Distribution of Venture Capital Investments," *American Journal of Sociology* 106, 1546–1588.
- Williams, Heidi, 2013, "Intellectual Property Rights and Innovation: Evidence from the Human Genome," *Journal of Political Economy* 121, 1–27.
- Williams, Heidi, 2017, "Patents and Research Investments: Assessing the Empirical Evidence," NBER Working Paper No. w21889.

Appendix A. Variable definitions.

Firm survival during year t after the first-action decision on a firm's first patent application is set to 1 if the firm is matched with the NETS sample and employment (or sales) data are available either for the year t or for any year after t . The variable is set to zero if the firm is matched with the NETS sample and employment (or sales) data are not available for the year t or for any year after t .

Employment growth after the first-action decision on a firm's first (or second) patent application is $\text{employment}_{t+k} / \text{employment}_t - 1$, where t is the first-action year and $k = 1 \dots 5$. If a firm dies and thus does not appear in NETS in year $t+k$, we set $\text{employment}_{t+k} = 0$.

Sales growth after the first-action decision on a firm's first (or second) patent application is $\text{sales}_{t+k} / \text{sales}_t - 1$, where t is the first-action year and $k = 1 \dots 5$. If a firm dies and thus does not appear in NETS in year $t+k$, we set $\text{sales}_{t+k} = 0$.

Pre-patent-filing employment growth is $\text{employment}_t / \text{employment}_{t-1} - 1$, where t is the year that a firm's first patent application is filed.

Pre-patent-filing sales growth is $\text{sales}_t / \text{sales}_{t-1} - 1$, where t is the year that a firm's first patent application is filed.

No. subsequent patent applications is the number of applications with a filing date greater than the first-action date of a firm's first (or second) application.

No. subsequent approved patents is the number of approved applications with a filing date greater than the first-action date of a firm's first (or second) application.

Approval rate of subsequent patent applications is defined as (no. subsequent approved patents) / (no. subsequent patent applications). It is only defined for firms with at least one subsequent patent application.

Total citations to all subsequent patent applications is the number of citations received by all subsequent patent applications combined. (This number is zero for firms with no subsequent applications.) We measure citations over the five years following each patent application's public disclosure date, which is typically 18 months after the application's filing date.

Average citations-per-patent to subsequent approved patents is the average number of citations received by those subsequent patent applications that are approved. It is only defined for firms with at least one subsequent approved patent.

Examiner experience is the number of years since the examiner joined the USPTO.

Examiner grade is the examiner's grade according to the government's General Schedule. Most examiners start at grade GS-7 or GS-9. Examiners at grades GS-7 through GS-11 need senior examiners to sign off on their decisions. GS-13 examiners undergo a period when they have partial signatory authority (during which time their work is subject to random checks). Examiners at levels GS-14 and above have full signatory authority.

Experienced founder is an indicator set equal to one if at least one of the up to five key executives of the startup listed in Standard & Poor's Capital IQ database previously founded a different firm, according to the professional background provided by Capital IQ.

High startup agglomeration state is an indicator set equal to one if the startup is headquartered in a state with above median startup agglomeration in the year of the startup's first patent application. We measure startup agglomeration using the number of first-time patent applicants in the state.

Industry classification. IT startups are those whose first patent application is reviewed by an examiner belonging to an art unit in one of the following USPTO technology centers: 21 (computer architecture, software, and information security); 24 (computer networks, multiplex communication, video distribution, and security); 26 (communications);

or 28 (semiconductors, electrical and optical systems and components). Biochemistry startups are those whose first patent application is reviewed by one of the following technology centers: 16 (biotechnology and organic chemistry); or 17 (chemical and materials engineering). Startups belonging to other industries are those whose first patent application is reviewed by one of the following technology centers: 36 (transportation, construction, electronic commerce, agriculture, national security and license & review); or 37 (mechanical engineering, manufacturing, products).

Figure 1. Distribution of Patent Examiners' Approval Rates.

The figure shows the sample distribution of patent examiner approval rates, defined as in equation (2), estimated within an art unit and year using a regression of approval rates on a full set of art-unit-by-application-year fixed effects.





Figure 2. The Effect of Patent Grants on Startup Growth.

The figure plots the estimated patent approval effect on employment growth (Panel A) and sales growth (Panel B) over the five years following the first-action decision on a startup's first patent application. Specifically, the solid line shows the estimated patent approval effect obtained by estimating equation (1) by 2SLS separately over horizons from one to five years after the first-action date. We use the approval rate of the examiner reviewing each patent application as an instrument for the likelihood that the application is approved. The dashed lines show 95% confidence intervals.

Panel A. Employment growth.



Panel B. Sales growth.




Figure 3. Time Lag between Patent Decision and VC Funding Round.

The figure shows the distribution of the time lag (in months) between the first-action date and the VC investment date for successful first-time patent applicants that go on to raise funding from a VC. VC funding events that take place more than five years after the first-action decision are not shown.




Table 1. Summary Statistics.

The table reports summary statistics for the firms in our sample of first-time patent applicants (or “startups”), broken down by whether their first application is approved or rejected. Data on age, employment, and sales are only available for those startups that can be matched to the National Establishment Times Series (NETS) database. For variable definitions and details of their construction see Appendix A.

		Firms whose first patent application is ...	
		approved	rejected
No. firms		22,084	12,131
% of firms		0.6454	0.3546
Scope: number of allowed independent claims	mean	3.2	0
	median	3.0	0
	<i>st.dev.</i>	2.6	0
Panel A. Pre-filing characteristics			
Age at first patent filing (years)	median	2.0	2.0
Employees at filing date	mean	29.6	29.0
	median	8.0	8.0
	<i>st.dev.</i>	61.9	61.2
Sales at filing date (\$ million)	mean	4.3	4.2
	median	0.8	0.7
	<i>st.dev.</i>	9.9	10.0
Pre-patent-filing employment growth (%)	mean	16.2	16.0
	<i>st.dev.</i>	68.7	68.1
Pre-patent-filing sales growth (%)	mean	20.1	18.4
	<i>st.dev.</i>	87.9	83.8
Panel B. Subsequent growth in employment and sales			
Employment growth after first-action decision on the firm's first patent application, measured over the following ...			
... 1 year	mean	6.6	-0.1
	<i>st.dev.</i>	50.1	47.8
... 3 years	mean	19.3	2.3
	<i>st.dev.</i>	122.2	111.0
... 5 years	mean	24.7	-2.1
	<i>st.dev.</i>	159.3	124.7
Sales growth after first-action decision on the firm's first patent application, measured over the following ...			
... 1 year	mean	11.1	2.5
	<i>st.dev.</i>	73.7	66.1
... 3 years	mean	34.3	11.7
	<i>st.dev.</i>	184.0	161.2
... 5 years	mean	50.2	16.1
	<i>st.dev.</i>	255.9	211.8

Table 1. Continued.

		Firms whose first patent application is ...	
		approved	rejected
Panel C. Subsequent patenting: patent applications filed after first-action decision on firm's first application			
No. subsequent patent applications	mean	3.1	1.2
	<i>st.dev.</i>	11.7	5.7
No. subsequent approved patents	mean	1.8	0.5
	<i>st.dev.</i>	7.4	2.7
Approval rate of subsequent patent applications (%)		70.5	47.9
Total citations to all subsequent patent applications	mean	8.3	2.2
	<i>st.dev.</i>	77.8	26.6
Average citations-per-patent to subsequent approved patents	mean	2.0	1.5
	<i>st.dev.</i>	3.8	3.0
Panel D. Subsequent VC funding and IPOs			
% of startups that raise VC funding after first-action		8.0	5.6
% of startups that go public after first-action		0.9	0.6
Panel E. Subsequent pledges of patents as collateral			
% of startups that pledged their first patent application as collateral after first-action decision, measured after the following ...			
... 1 year	mean	1.3	0.9
... 3 years	mean	4.0	2.1
... 5 years	mean	6.6	2.6
% of startups without VC funding that pledged their first patent application as collateral after first-action decision, measured after the following ...			
... 1 year	mean	0.9	0.7
... 3 years	mean	3.1	1.6
... 5 years	mean	5.1	2.0

Table 2. Examiner Leniency: First-stage Results.

The table reports the results of estimating various versions of the first-stage equation (3) of our 2SLS analysis. The first stage uses the past approval rate of the patent examiner in charge of reviewing a startup's first patent application to predict whether the application will be approved. Identification assumes that applications are assigned to examiners quasi-randomly within an art unit and year. Accordingly, our baseline specification shown in column 1 includes art-unit-by-year fixed effects. Columns 2 through 6 consider threats to identification arising from potential violations of quasi-random assignment. Columns 2 through 4 investigate the possibility of quality-based assignment, using characteristics of the applicant or the application to proxy for quality. Columns 5 and 6 investigate the possibility of assignment based on examiner characteristics, controlling for examiner specialization by including technology-subclass-by-year fixed effects (in columns 5 and 6) and proxies for examiner experience and seniority (in column 6). The number of observations varies depending on data availability (e.g., sales and employment data are only available for startups that can be matched to NETS) and due to a varying number of singletons. All specifications are estimated using least squares. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	First patent application approved?					
	(1)	(2)	(3)	(4)	(5)	(6)
IV: patent examiner approval rate	0.670*** <i>0.017</i>	0.669*** <i>0.021</i>	0.674*** <i>0.022</i>	0.682*** <i>0.019</i>	0.692*** <i>0.022</i>	0.689*** <i>0.033</i>
Applicant characteristics						
ln(employees at first-action)		-0.003 <i>0.004</i>				
ln(1 + sales at first-action)		0.003 <i>0.004</i>				
employment growth at first action			0.010 <i>0.009</i>			
sales growth at first action			0.003 <i>0.007</i>			
Application characteristics						
ln(# independent claims in application)				0.032*** <i>0.004</i>		
Examiner characteristics						
ln(examiner experience)						-0.010 <i>0.007</i>
examiner grade GS-9						-0.020 <i>0.016</i>
examiner grade GS-11						-0.029* <i>0.017</i>
examiner grade GS-12						-0.020 <i>0.021</i>
examiner grade GS-13						-0.014 <i>0.021</i>
examiner grade GS-14						0.018 <i>0.023</i>
examiner grade GS-15						-0.051 <i>0.034</i>
Fixed effects						
art unit × year	Yes	Yes	Yes	Yes	Yes	Yes
HQ state	Yes	Yes	Yes	Yes	Yes	Yes
tech subclass × year	No	No	No	No	Yes	Yes
Diagnostics						
R ²	25.7%	27.8%	28.2%	25.7%	42.1%	42.2%
F-test: IV = 0	1,504.1***	1,062.6***	967.7***	1,316.0***	966.7***	440.7***
No. of observations (firms)	34,215	21,564	20,207	29,001	28,299	28,294

Table 3. Examiner Leniency: Instrument Validity.

The table reports the results of regressing the approval rate of the examiner reviewing each firm's first patent application on the characteristics of the applicant, the application, or the examiner. The number of observations varies depending on data availability (e.g., sales and employment data are only available for startups that can be matched to NETS) and due to a varying number of singletons. In columns 5 and 6, the sample is restricted to patent applications that are also filed with the European and Japanese patent offices, respectively. All specifications are estimated using least squares. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	IV: Patent examiner approval rate					
	(1)	(2)	(3)	(4)	(5)	(6)
Applicant characteristics						
ln(employees at filing date)	0.001					
	<i>0.002</i>					
ln(1 + sales at filing date)	-0.000					
	<i>0.001</i>					
employment growth during year prior to filing date		-0.002				
		<i>0.004</i>				
sales growth during year prior to filing date		0.000				
		<i>0.003</i>				
Application characteristics						
ln(# independent claims in application)			0.002			
			<i>0.002</i>			
Examiner characteristics						
ln(examiner experience)				0.039***		
				<i>0.004</i>		
examiner grade GS-9				0.038***		
				<i>0.007</i>		
examiner grade GS-11				0.080***		
				<i>0.009</i>		
examiner grade GS-12				0.131***		
				<i>0.010</i>		
examiner grade GS-13				0.191***		
				<i>0.011</i>		
examiner grade GS-14				0.226***		
				<i>0.012</i>		
examiner grade GS-15				0.197***		
				<i>0.020</i>		
Approval by foreign patent office						
European Patent Office					-0.005	
					<i>0.013</i>	
Japanese Patent Office						0.014
						<i>0.029</i>
Fixed effects						
art unit × year	Yes	Yes	Yes	Yes	Yes	Yes
HQ state	Yes	Yes	Yes	Yes	Yes	Yes
tech subclass × year	No	No	No	Yes	Yes	Yes
Diagnostics						
R ²	58.6%	59.1%	57.1%	80.4%	85.4%	92.4%
No. of observations (firms)	20,151	17,546	29,001	28,294	2,745	568

Table 4. How Does a Startup’s First Patent Application Affect Employment Growth, Sales Growth, and Survival?

Panels A and B report the results of estimating equation (1) to examine how the approval of a startup’s first patent application affects the startup’s subsequent growth in employment and sales, respectively, over the one to five years following the first-action date. For startups that die, we set the growth rate to –100% in the year of exit. Panel C reports the results of linear probability models of firm survival. We code a startup as being alive in year t if it continues to be included in the NETS database that year. The variable of interest in each panel is approval of a startup’s first patent application. Panels A and C control for log employment at first-action, while Panel B controls for log sales at first-action (not shown). All columns report 2SLS results using the approval rate of the examiner reviewing the patent application as an instrument for the likelihood that the application is approved. The weak-instrument test uses the Kleibergen-Paap rk Wald F statistic. Employment and sales data come from NETS; thus, startups that cannot be matched to NETS are excluded. NETS data are available through 2011, resulting in reduced sample sizes as we widen the window from one to five years. The sample is restricted to firms for which NETS reports non-zero sales and employment for the year of the first-action decision. For variable definitions and details of their construction see Appendix A. All specifications include art-unit-by-year and headquarter-state fixed effects. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	1 year (1)	2 years (2)	3 years (3)	4 years (4)	5 years (5)
<u>Panel A. Employment growth</u>					
First patent application approved	0.061* <i>0.033</i>	0.228*** <i>0.061</i>	0.333*** <i>0.093</i>	0.489*** <i>0.132</i>	0.545*** <i>0.149</i>
Diagnostics					
Weak-instrument test	1,063.6***	1,063.6***	879.3***	621.6***	490.2***
Unconditional mean of dep. variable	4.3%	10.2%	13.8%	17.2%	17.4%
No. of observations (firms)	21,564	21,564	18,745	15,417	12,655
<u>Panel B. Sales growth</u>					
First patent application approved	0.010 <i>0.013</i>	0.276*** <i>0.082</i>	0.512*** <i>0.137</i>	0.796*** <i>0.208</i>	0.795*** <i>0.246</i>
Diagnostics					
Weak-instrument test	1,065.0***	1,064.6***	880.1***	622.1***	493.2***
Unconditional mean of dep. variable	8.2%	18.3%	27.0%	36.5%	40.9%
No. of observations (firms)	21,530	21,537	18,729	15,410	12,651
<u>Panel C. Survival</u>					
First patent application approved	0.010 <i>0.013</i>	0.031* <i>0.018</i>	0.046** <i>0.023</i>	0.024 <i>0.031</i>	0.032 <i>0.039</i>
Diagnostics					
Weak-instrument test	1,063.6***	1,063.6***	879.3***	621.6***	490.2***
Unconditional mean of dep. variable	95.8%	91.2%	86.8%	83.8%	79.7%
No. of observations (firms)	21,564	21,564	18,745	15,417	12,655

Table 5. How Does a Startup’s First Patent Application Affect Subsequent Innovation?

The table reports the results of estimating equation (1) to examine how the approval of a startup’s first patent application affects the startup’s follow-on innovation. Data on subsequent applications come from the USPTO internal databases and include all applications that receive a final decision through December 31, 2013. Column 3 includes only startups filing at least one patent application after the first-action decision on the startup’s first patent application and for which we can measure the approval rate of subsequent applications. Column 5 includes only those startups with at least one subsequent patent approval and for which we can measure the average number of citations-per-patent to subsequently approved patents. We measure citations over the five years following each patent application’s public disclosure date, which is typically 18 months after the application’s filing date. In untabulated results, we find that the effects in columns 4 and 5 are even stronger when we measure citations over seven or ten years. For variable definitions and further details of their construction see Appendix A. All specifications are estimated by 2SLS and include art-unit-by-year and headquarter-state fixed effects. We use the approval rate of the examiner reviewing each patent application as an instrument for the likelihood that the application is approved. The weak-instrument test uses the Kleibergen-Paap rk Wald F statistic. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Follow-on innovation				
	Log (1 + subsequent patent applications) (1)	Log (1 + subsequent approved patents) (2)	Approval rate of subsequent patent applications (3)	Log (1 + total citations to subsequent patent applications) (4)	Log (1 + avg. citations-per-patent to subsequent approved patents) (5)
First patent application approved	0.448*** <i>0.039</i>	0.353*** <i>0.030</i>	0.244*** <i>0.043</i>	0.472*** <i>0.049</i>	0.285*** <i>0.094</i>
Diagnostics					
Weak-instrument test	1,504.1***	1,504.1***	570.5***	1,504.0***	305.0***
Uncond. mean of non-logged dep. var.	2.4	1.3	65.8%	6.1	1.9
No. of observations (firms)	34,215	34,215	12,595	34,214	9,793

Table 6. The Effect of a Startup’s Second Patent Application

This table repeats the analysis shown in Panels A and B of Table 4 and in Table 5, except that it models the effect of the approval of a startup’s *second* patent application. All columns report 2SLS results using the approval rate of the examiner reviewing the second patent application as an instrument for the likelihood that the application is approved. The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. For variable definitions and details of their construction see Appendix A. All specifications include art-unit-by-year and headquarter-state fixed effects. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	(1)	(2)	(3)	(4)	(5)
Panel A. Employment growth					
	<u>1 year</u>	<u>2 years</u>	<u>3 years</u>	<u>4 years</u>	<u>5 years</u>
Second patent application approved	0.108* <i>0.058</i>	0.000 <i>0.112</i>	-0.019 <i>0.156</i>	0.008 <i>0.220</i>	-0.091 <i>0.286</i>
Diagnostics					
Weak-instrument test	547.3***	547.3***	491.7***	308.6***	213.4***
Unconditional mean of dep. variable	6.2%	12.1%	16.4%	20.6%	19.8%
No. of observations (firms)	8,798	8,798	7,284	5,660	4,279
Panel B. Sales growth					
	<u>1 year</u>	<u>2 years</u>	<u>3 years</u>	<u>4 years</u>	<u>5 years</u>
Second patent application approved	0.148 <i>0.097</i>	0.024 <i>0.171</i>	0.092 <i>0.282</i>	0.202 <i>0.396</i>	0.496 <i>0.560</i>
Diagnostics					
Weak-instrument test	551.7***	551.9***	495.3***	309.1***	213.3***
Unconditional mean of dep. variable	13.1%	28.6%	41.8%	51.8%	58.9%
No. of observations (firms)	8,781	8,784	7,279	5,658	4,278
	<u>Log (1 + subsequent patent applications)</u>	<u>Log (1 + subsequent approved patents)</u>	<u>Approval rate of subsequent patent applications</u>	<u>Log (1 + total citations to subsequent patent applications)</u>	<u>Log (1 + avg. citations-per-patent to subsequent approved patents)</u>
Panel C. Follow-on innovation					
Second patent application approved	0.496*** <i>0.080</i>	0.404*** <i>0.062</i>	0.141*** <i>0.053</i>	0.563*** <i>0.094</i>	0.052 <i>0.096</i>
Diagnostics					
Weak-instrument test	768.9***	768.9***	276.0***	769.0***	211.2***
Unconditional mean of non-logged dep. var.	4.5	2.4	66.4%	12.4	2.1
No. of observations (firms)	12,653	12,653	6,103	12,652	4,883

Table 7. Does a Startup’s First Patent Application Affect Access to VC Funding and the IPO Market?

The table reports the results of estimating equation (1) to examine how the approval of a startup’s first patent application affects the startup’s ability to raise funding from a VC or in the IPO market. The dependent variable in columns 1 through 5 is an indicator set equal to one if the startup raises VC funding at some point in the 1...5 years following the first-action decision, respectively. The dependent variable in column 6 is an indicator set equal to one if the startup goes public after the first-action decision on its first patent application, and zero otherwise. All specifications are estimated by 2SLS and include art-unit-by-year and headquarter-state fixed effects. We use the past approval rate of the examiner reviewing each patent application as an instrument for the likelihood that the application is approved. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively. The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic.

	Following the first-action decision on its first patent application, does the startup raise VC funding ...					Does the startup raise capital in the IPO market? (6)
	in the next 1 year? (1)	in the next 2 years? (2)	in the next 3 years? (3)	in the next 4 years? (4)	in the next 5 years? (5)	
First patent application approved	0.017* <i>0.009</i>	0.027*** <i>0.010</i>	0.030*** <i>0.010</i>	0.034*** <i>0.011</i>	0.036*** <i>0.011</i>	0.010** <i>0.005</i>
Log (1 + no. prior VC rounds)	0.272*** <i>0.009</i>	0.380*** <i>0.009</i>	0.416*** <i>0.009</i>	0.425*** <i>0.009</i>	0.429*** <i>0.009</i>	0.044*** <i>0.004</i>
Diagnostics						
Weak-instrument test	1,497.9***	1,483.6***	1,483.6***	1,484.0***	1,486.0***	1,503.1***
Mean of dep. variable	3.9%	5.7%	6.4%	6.8%	7.0%	0.78%
Median no. months from first-action to VC round or IPO for successful applicants	5.2	8.1	9.2	10.0	10.3	65.7
No. of observations (firms)	34,167	34,111	34,060	34,013	33,981	34,215

Table 8. Does a Startup’s First Patent Application Affect Access to Debt?

The table reports the results of estimating equation (1) to examine how the approval of a startup’s first patent application affects the startup’s ability to raise debt. The dependent variable in columns 1 through 5 is an indicator set equal to one if the startup pledges its patent application as collateral in a security agreement (recorded as a patent reassignment by the USPTO) at some point in the 1...5 years following the first-action decision, respectively; the dependent variable in column 6 is an indicator set equal to one if the startup pledges its patent application as collateral in a security agreement at any point after first-action. All specifications are estimated by 2SLS and include art-unit-by-year and headquarter-state fixed effects. We use the past approval rate of the examiner reviewing each patent application as an instrument for the likelihood that the application is approved. The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Following the first-action decision on its first patent application, does the startup pledge the application as collateral ...					
	in the next 1 year? (1)	in the next 2 years? (2)	in the next 3 years? (3)	in the next 4 years? (4)	in the next 5 years? (5)	at any point in the future? (6)
Panel A: All firms						
First patent application approved	0.004 <i>0.006</i>	0.018** <i>0.008</i>	0.029*** <i>0.009</i>	0.039*** <i>0.010</i>	0.058*** <i>0.010</i>	0.086*** <i>0.012</i>
Diagnostics						
Weak-instrument test	1,530.0***	1,530.0***	1,530.0***	1,530.0***	1,530.0***	1,530.0***
Mean of dep. variable	1.1%	2.3%	3.3%	4.3%	5.2%	7.2%
Median no. months from first-action to application pledge	6.0	11.8	17.8	23.0	27.5	38.9
No. of observations (firms)	33,520	33,520	33,520	33,520	33,520	33,520
Panel B. Firms without VC funding						
First patent application approved	-0.002 <i>0.005</i>	0.011 <i>0.007</i>	0.019** <i>0.008</i>	0.026*** <i>0.009</i>	0.044*** <i>0.010</i>	0.059*** <i>0.011</i>
Diagnostics						
Weak-instrument test	1,370.7***	1,370.7***	1,370.7***	1,370.7***	1,370.7***	1,370.7***
Mean of dep. variable	0.8%	1.7%	2.5%	3.2%	3.9%	5.5%
Median no. months from first-action to application pledge	5.9	12.2	18.6	23.0	27.5	39.6
No. of observations (firms)	31,161	31,161	31,161	31,161	31,161	31,161

Table 9. How Do Patents Affect a Startup’s Access to VC Funding? Subsample Analyses.

The table examines how the effect of a patent on facilitating access to VC funding varies across different types of startups. The dependent variable equals one if the startup raises VC funding at some point in the three years following the first-action decision on its patent application. Columns 1 through 3 split startups by the number of VC rounds raised before first-action. Column 4 captures variation in founders’ prior entrepreneurial experience. Data on founder experience come from Capital IQ. Capital IQ’s coverage of founders’ backgrounds is most complete for firms that have raised VC funding, and so column 4 restricts the sample to firms that have raised at least one VC round before first-action. Column 5 splits the approval effect according to whether startups are headquartered in a state with above-median startup agglomeration in the year of their first patent application. (This indicator is time-varying and so is identified in the presence of headquarter-state fixed effects.) Column 6 splits the approval effect by industry. Columns 7 and 8 focus on approval of the second patent application, with column 8 restricting the sample to startups whose first application was rejected. For variable definitions and details of their construction see Appendix A. All specifications are estimated by 2SLS using the approval rate of the examiner reviewing each patent application as an instrument for the likelihood that the application is approved. Columns 3-8 include the log number of prior rounds; columns 4-6 include the interaction of the examiner approval rate and the splitting variable as instrument for the interacted patent approval indicator(s). The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. All specifications include art-unit-by-year fixed effects and headquarter-state fixed effects. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	In the three years following the first-action decision on its first patent application, does the startup raise ...							
	its first VC	its second	a higher VC	any VC funding?				
	round?	VC round?	round?	(4)	(5)	(6)	(7)	(8)
	(1)	(2)	(3)					
First patent application approved	0.013*	0.455**	0.269*	0.002	0.016	0.063***		
	<i>0.007</i>	<i>0.227</i>	<i>0.142</i>	<i>0.178</i>	<i>0.011</i>	<i>0.021</i>		
... × inexperienced founder				0.326*				
				<i>0.167</i>				
... × high startup agglomeration state					0.027**			
					<i>0.011</i>			
... × biochemistry						-0.107***		
						<i>0.035</i>		
... × other industries						-0.028		
						<i>0.024</i>		
Second application approved							0.028	0.091
							<i>0.019</i>	<i>0.057</i>
Inexperienced founder				-0.262**				
				<i>0.110</i>				
High startup agglomeration state					-0.016*			
					<i>0.009</i>			
Diagnostics								
Mean of dep. variable	1.7%	46.8%	61.5%	61.3%	6.4%	6.4%	11.3%	10.2%
Weak-instrument test	1,260.6***	20.8***	43.24***	17.4***	744.8***	212.9***	751.1***	139.3***
No. of observations (startups)	31,057	406	1,306	1,086	34,060	34,060	12,455	2,782

INTERNET APPENDIX

(NOT INTENDED FOR PUBLICATION)

Table IA.1. How Does a Startup's First Patent Application Affect Employment Growth, Sales Growth, and Survival?

The table repeats the analysis in Table 4 in the paper, except that it uses as the first stage of the 2SLS analysis the specification shown in column 6 of Table 2. These same additional controls are then also included in the second stage (not shown). The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	1 year (1)	2 years (2)	3 years (3)	4 years (4)	5 years (5)
<u>Panel A. Employment growth</u>					
First patent application approved	0.101** <i>0.051</i>	0.257** <i>0.101</i>	0.299* <i>0.156</i>	0.500** <i>0.204</i>	0.668*** <i>0.223</i>
Diagnostics					
Weak-instrument test	344.6***	344.6***	294.9***	229.2***	178.8***
Unconditional mean of dep. variable	4.3%	10.1%	13.5%	16.8%	16.4%
No. of observations (firms)	16,522	16,522	14,299	11,602	9,433
<u>Panel B. Sales growth</u>					
First patent application approved	0.080 <i>0.073</i>	0.212 <i>0.139</i>	0.433* <i>0.222</i>	0.672** <i>0.313</i>	0.893** <i>0.353</i>
Diagnostics					
Weak-instrument test	349.1***	349.5***	298.5***	228.6***	179.0***
Unconditional mean of dep. variable	8.1%	18.1%	26.5%	36.2%	40.3%
No. of observations (firms)	16,487	16,494	14,276	11,594	9,428
<u>Panel C. Survival</u>					
First patent application approved	-0.002 <i>0.025</i>	0.014 <i>0.032</i>	0.023 <i>0.045</i>	0.041 <i>0.059</i>	0.035 <i>0.062</i>
Diagnostics					
Weak-instrument test	344.6***	344.6***	294.9***	229.2***	178.8***
Unconditional mean of dep. variable	95.7%	91.0%	86.7%	83.5%	79.3%
No. of observations (firms)	16,522	16,522	14,299	11,602	9,433

Table IA.2. How Does a Startup's First Patent Application Affect Subsequent Innovation?

The table repeats the analysis in Table 5 in the paper, except that it uses as the first stage of the 2SLS analysis the specification shown in column 6 of Table 2. These same additional controls are then also included in the second stage (not shown). The weak-instrument test uses the Kleibergen-Paap rk Wald F statistic. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Follow-on innovation				
	Log (1 + subsequent patent applications) (1)	Log (1 + subsequent approved patents) (2)	Approval rate of subsequent patent applications (3)	Log (1 + total citations to subsequent patent applications) (4)	Log (1 + avg. citations-per- patent to subsequent approved patents) (5)
First patent application approved	0.489*** <i>0.063</i>	0.386*** <i>0.051</i>	0.286*** <i>0.084</i>	0.506*** <i>0.080</i>	0.250 <i>0.268</i>
Diagnostics					
Weak-instrument test	454.8***	454.8***	129.2***	454.9***	62.3***
Uncond. mean of non-logged dep. var.	2.4	1.3	65.8%	6.2	2.1
No. of observations (firms)	28,192	28,192	8,285	28,191	5,815

Table IA.3. Does a Startup's First Patent Application Affect Access to VC Funding and the IPO Market?

The table repeats the analysis in Table 7 in the paper, except that it uses as the first stage of the 2SLS analysis the specification shown in column 6 of Table 2. These same additional controls are then also included in the second stage (not shown). The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Following the first-action decision on its first patent application, does the startup raise VC funding ...					Does the startup raise capital in the IPO market? (6)
	in the next 1 year? (1)	in the next 2 years? (2)	in the next 3 years? (3)	in the next 4 years? (4)	in the next 5 years? (5)	
First patent application approved	0.012 <i>0.015</i>	0.023 <i>0.016</i>	0.039** <i>0.017</i>	0.041** <i>0.017</i>	0.046*** <i>0.017</i>	0.014* <i>0.007</i>
Log (1 + no. prior VC rounds)	0.264*** <i>0.009</i>	0.370*** <i>0.010</i>	0.405*** <i>0.010</i>	0.415*** <i>0.010</i>	0.419*** <i>0.010</i>	0.041*** <i>0.004</i>
Diagnostics						
Weak-instrument test	450.4***	452.6***	444.6***	445.7***	449.1***	452.9***
Mean of dep. variable	4.2%	6.1%	6.8%	7.2%	7.4%	0.82%
Median no. months from first-action to VC round or IPO for successful applicants	5.3	8.0	9.1	9.9	10.2	62.1
No. of observations (firms)	28,147	28,090	28,042	27,997	27,965	28,192

Table IA.4. Does a Startup's First Patent Application Affect Access to Debt?

The table repeats the analysis in Table 8 in the paper, except that it uses as the first stage of the 2SLS analysis the specification shown in column 6 of Table 2. These same additional controls are then also included in the second stage (not shown). The weak-instrument test uses the Kleibergen-Paap rk Wald F statistic. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Following the first-action decision on its first patent application, does the startup pledge the application as collateral ...					
	in the next 1 year? (1)	in the next 2 years? (2)	in the next 3 years? (3)	in the next 4 years? (4)	in the next 5 years? (5)	at any point in the future? (6)
Panel A: All firms						
First patent application approved	0.008 <i>0.009</i>	0.013 <i>0.013</i>	0.041*** <i>0.015</i>	0.053*** <i>0.017</i>	0.085*** <i>0.018</i>	0.109*** <i>0.021</i>
Diagnostics						
Weak-instrument test	456.9***	456.9***	456.9***	456.9***	456.9***	456.9***
Mean of dep. variable	1.1%	2.3%	3.4%	4.4%	5.3%	7.3%
Median no. months from first-action to application pledge	6.1	12.0	18.0	23.1	27.4	39.2
No. of observations (firms)	27,542	27,542	27,542	27,542	27,542	27,542
Panel B. Firms without VC funding						
First patent application approved	0.005 <i>0.008</i>	0.008 <i>0.011</i>	0.031** <i>0.013</i>	0.038*** <i>0.014</i>	0.066*** <i>0.015</i>	0.073*** <i>0.018</i>
Diagnostics						
Weak-instrument test	451.9***	451.9***	451.9***	451.9***	451.9***	451.9***
Mean of dep. variable	0.8%	1.7%	2.5%	3.2%	4.0%	5.5%
Median no. months from first-action to application pledge	5.9	12.4	18.7	23.0	27.4	40.1
No. of observations (firms)	25,293	25,293	25,293	25,293	25,293	25,293

Table IA.5. How Does a Startup’s First Patent Application Affect Employment Growth, Sales Growth, and Survival? OLS Results.

The analysis is analogous to Table 4, with the only difference being that we use OLS instead of 2SLS. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	1 year (1)	2 years (2)	3 years (3)	4 years (4)	5 years (5)
<u>Panel A. Employment growth</u>					
First patent application approved	0.065*** <i>0.008</i>	0.126*** <i>0.014</i>	0.152*** <i>0.019</i>	0.212*** <i>0.026</i>	0.247*** <i>0.032</i>
Diagnostics					
R^2	13.0%	13.6%	13.8%	14.9%	14.5%
Unconditional mean of dep. variable	4.3%	10.2%	13.8%	17.2%	17.4%
No. of observations (firms)	21,564	21,564	18,745	15,417	12,655
<u>Panel B. Sales growth</u>					
First patent application approved	0.083*** <i>0.048</i>	0.153*** <i>0.082</i>	0.214*** <i>0.137</i>	0.297*** <i>0.208</i>	0.333*** <i>0.246</i>
Diagnostics					
R^2	12.3%	13.1%	14.0%	13.7%	13.8%
Unconditional mean of dep. variable	8.2%	18.3%	27.0%	36.5%	40.9%
No. of observations (firms)	21,530	21,537	18,729	15,410	12,651
<u>Panel C. Survival</u>					
First patent application approved	0.029*** <i>0.003</i>	0.051*** <i>0.005</i>	0.060*** <i>0.006</i>	0.069*** <i>0.007</i>	0.078*** <i>0.008</i>
Diagnostics					
R^2	12.7%	13.6%	14.3%	15.2%	15.6%
Unconditional mean of dep. variable	95.8%	91.2%	86.8%	83.8%	79.7%
No. of observations (firms)	21,564	21,564	18,745	15,417	12,655

Table IA.6. Effects of the First Patent Application on Employment and Sales Growth Controlling for Review Speed.

The table reports the results of estimating the effects of the first patent application on employment and sales growth while controlling for patent review speed (not shown). Doing so removes a possible confound that could arise if lenient examiners were also faster and review speed had an independent effect on outcomes. Given that review speed is plausibly endogenous, we instrument it using a measure of administrative delays described in Hegde and Ljungqvist (2017). The instrument builds on the notion that the time it takes to receive a first-action decision on a patent application can be decomposed into two parts: the time from filing to the application being assigned to the examiner’s docket (which reflects quasi-random administrative delays at the USPTO that are unrelated to invention quality or application complexity) and the time the examiner takes from the docket to the first-action date (which is potentially endogenous). Thus, we instrument review speed using the sum of the time an application takes from filing to being assigned to an examiner’s docket and the average time that examiner has taken in the past from docket to first-action. Panel A models the effect of patent approval of a startup’s first patent application on the startup’s subsequent growth in employment over the next 1 to 5 years. Panel B models the effect on subsequent sales growth. Panel A controls for log employment at first-action, while Panel B controls for log sales at first-action (not shown). All columns report 2SLS results. The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. Employment and sales data come from NETS; thus, startups that cannot be matched to NETS are excluded. NETS data are available through 2011, resulting in reduced sample sizes as we widen the window from 1 to 5 years. For variable definitions and details of their construction see Appendix A. All specifications include art-unit-by-year and headquarter-state fixed effects. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Employment or sales growth after USPTO first-action on the firm’s first patent application, measured over the following ...				
	1 year (1)	2 years (2)	3 years (3)	4 years (4)	5 years (5)
<u>Panel A. Employment growth</u>					
First patent application approved	0.032 <i>0.036</i>	0.178*** <i>0.066</i>	0.285*** <i>0.101</i>	0.418*** <i>0.148</i>	0.465*** <i>0.161</i>
Diagnostics					
Weak-instrument test	447.5***	447.5***	359.3***	262.0***	208.3***
Unconditional mean of dep. variable	4.3%	10.3%	13.8%	17.2%	17.3%
No. of observations (firms)	21,531	21,531	18,711	15,390	12,632
<u>Panel B. Sales growth</u>					
First patent application approved	0.060 <i>0.051</i>	0.225** <i>0.090</i>	0.461*** <i>0.148</i>	0.664*** <i>0.230</i>	0.652** <i>0.263</i>
Diagnostics					
Weak-instrument test	447.7***	447.9***	361.1***	263.6***	210.5***
Unconditional mean of dep. variable	8.2%	18.3%	27.0%	36.5%	40.8%
No. of observations (firms)	21,498	21,504	18,695	15,383	12,628

Table IA.7. How Does a Startup's First Patent Application Affect Subsequent Innovation? OLS Results.

The table reports the results of estimating equation (1) to examine how the approval of a startup's first patent application affects the startup's follow-on innovation. The analysis here is analogous to Table 5, with the only difference being that we use OLS instead of 2SLS. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Follow-on innovation				
	Log (1 + subsequent patent applications) (1)	Log (1 + subsequent approved patents) (2)	Approval rate of subsequent patent applications (3)	Log (1 + total citations to all subsequent patent applications) (4)	Log (1 + average citations-per-patent to subsequent approved patents) (5)
First patent application approved	0.361*** <i>0.013</i>	0.294*** <i>0.010</i>	0.199*** <i>0.012</i>	0.344*** <i>0.017</i>	0.079*** <i>0.022</i>
Diagnostics					
R^2	21.5%	20.6%	21.0%	21.0%	29.6%
Mean of non-logged dep. var.	2.4	1.3	65.8%	6.1	1.9
No. of observations (startups)	34,215	34,215	12,595	34,214	9,793

Table IA.8. The Effect of a Startup’s Second Patent Application. OLS Results.

This table repeats the analysis shown in Panels A and B of Table 4 and in Table 5, except that it models the effect of the approval of a startup’s *second* patent application on the growth (Panels A and B) or follow-on innovation (Panel C) following the first-action decision on the second patent application. The analysis here is analogous to Table 6, with the only difference being that we use OLS instead of 2SLS. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	(1)	(2)	(3)	(4)	(5)
Panel A. Employment growth					
Second patent application approved	<u>1 year</u> 0.045*** <i>0.015</i>	<u>2 years</u> 0.062** <i>0.026</i>	<u>3 years</u> 0.097** <i>0.038</i>	<u>4 years</u> 0.099** <i>0.050</i>	<u>5 years</u> 0.142** <i>0.063</i>
Diagnostics					
R^2	21.0%	22.0%	22.9%	24.6%	25.2%
Unconditional mean of dep. variable	6.2%	12.1%	16.4%	20.6%	19.8%
No. of observations (firms)	8,798	8,798	7,284	5,660	4,279
Panel B. Sales growth					
Second patent application approved	<u>1 year</u> 0.054** <i>0.024</i>	<u>2 years</u> 0.089** <i>0.044</i>	<u>3 years</u> 0.175** <i>0.073</i>	<u>4 years</u> 0.223** <i>0.096</i>	<u>5 years</u> 0.405*** <i>0.131</i>
Diagnostics					
R^2	21.2%	22.5%	23.8%	24.5%	25.4%
Unconditional mean of dep. variable	13.1%	28.6%	41.8%	51.8%	58.9%
No. of observations (firms)	8,781	8,784	7,279	5,658	4,278
Panel C. Follow-on innovation					
Second patent application approved	<u>Log (1 + subsequent patent applications)</u> 0.448*** <i>0.023</i>	<u>Log (1 + subsequent approved patents)</u> 0.385*** <i>0.017</i>	<u>Approval rate of subsequent patent applications</u> 0.165*** <i>0.013</i>	<u>Log (1 + total citations to subsequent patent applications)</u> 0.463*** <i>0.030</i>	<u>Log (1 + avg. citations-per-patent to subsequent approved patents)</u> 0.026 <i>0.030</i>
Diagnostics					
R^2	29.3%	29.9%	30.6%	30.7%	38.0%
Unconditional mean of non-logged dep. var.	4.5	2.4	66.4%	12.4	2.1
No. of observations (firms)	12,653	12,653	6,103	12,652	4,883

Table IA.9. Cross-Industry Differences in the Effects of a First Patent Application.

The table reports the results of estimating equation (1) to examine how the approval of a startup's first patent application affects the startup's various outcomes separately in the IT sector. Each row represents one regression. Column 1 reports estimates obtained by running regressions in our full sample of startups, as reported in Tables 4 and 5. Column 2 reports estimates obtained by running regressions in the subsample of startups in the IT sector. For variable definitions and further details of their construction see Appendix A. All specifications are estimated by 2SLS using the approval rate of the examiner reviewing each patent application as an instrument for the likelihood that the application is approved. The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

Effect of patent grants for ...	All startups (1)	IT startups (2)
Employment growth		
1-year employment growth	0.061* <i>0.033</i>	-0.022 <i>0.058</i>
2-year employment growth	0.228*** <i>0.061</i>	0.190* <i>0.109</i>
3-year employment growth	0.333*** <i>0.093</i>	0.271* <i>0.150</i>
4-year employment growth	0.489*** <i>0.132</i>	0.283 <i>0.185</i>
5-year employment growth	0.545*** <i>0.149</i>	0.396* <i>0.210</i>
Sales growth		
1-year sales growth	0.096** <i>0.048</i>	-0.006 <i>0.081</i>
2-year sales growth	0.276*** <i>0.082</i>	0.386*** <i>0.134</i>
3-year sales growth	0.512*** <i>0.137</i>	0.630*** <i>0.205</i>
4-year sales growth	0.796*** <i>0.208</i>	0.602** <i>0.293</i>
5-year sales growth	0.795*** <i>0.246</i>	0.885*** <i>0.333</i>
Follow-on innovation		
No. subsequent patent applications	0.448*** <i>0.039</i>	0.569*** <i>0.068</i>
No. subsequent approved patents	0.353*** <i>0.030</i>	0.445*** <i>0.054</i>
Approval rate of subsequent patent applications	0.244*** <i>0.043</i>	0.262*** <i>0.091</i>
Total citations to all subsequent patent applications	0.472*** <i>0.049</i>	0.725*** <i>0.093</i>
Average citations-per-patent to subsequent approved patents	0.285*** <i>0.094</i>	0.441** <i>0.201</i>

Table IA.10. The Effect of a Startup's Second Patent Application Split By the First Application's Outcome.

This table restricts the sample used in Table 6, Panels A and B, to startups whose first patent application was either approved (Panel A) or rejected (Panel B). All columns report 2SLS results using the approval rate of the examiner reviewing the second patent application as an instrument for the likelihood that the application is approved. The weak-instrument test uses the Kleibergen-Paap *rk* Wald *F* statistic. For variable definitions and details of their construction see Appendix A. All specifications include art-unit-by-year and headquarter-state fixed effects. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	First application is approved			First application is rejected		
	Employment or sales growth after USPTO first-action on the firm's second patent application, measured over the following ...					
	1 year (1)	3 years (2)	5 years (3)	1 year (4)	3 years (5)	5 years (6)
<u>Panel A. Employment growth</u>						
Second patent application approved	0.083 <i>0.080</i>	-0.043 <i>0.234</i>	-0.333 <i>0.421</i>	0.109 <i>0.140</i>	-0.042 <i>0.387</i>	-0.382 <i>0.653</i>
Diagnostics						
Weak-instrument test	224.0***	222.5***	103.2***	95.4***	75.6***	24.7***
Unconditional mean of dep. variable	6.7%	19.6%	23.4%	5.0%	9.4%	10.3%
No. of observations (firms)	6,264	5,313	3,252	1,711	1,305	557
<u>Panel B. Sales growth</u>						
Second patent application approved	0.124 <i>0.124</i>	-0.038 <i>0.410</i>	-0.201 <i>0.849</i>	0.005 <i>0.262</i>	0.403 <i>0.762</i>	1.259 <i>1.251</i>
Diagnostics						
Weak-instrument test	227.8***	226.1***	103.4***	93.7***	75.0***	24.8***
Unconditional mean of dep. variable	13.4%	45.4%	64.7%	12.3%	28.6%	39.0%
No. of observations (firms)	6,251	5,309	3,252	1,704	1,303	557

Table IA.11. Does a Startup’s First Patent Application Affect Access to VC Funding and the IPO Market? OLS Results.

The table reports the results of estimating equation (4) to examine how the approval of a startup’s first patent application affects the startup’s ability to raise funding from a VC or in the IPO market. The analysis here is analogous to Table 7, with the only difference being that we use OLS instead of 2SLS. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Following the first-action decision on its first patent application, does the startup raise VC funding ...					Does the startup raise capital in the IPO market? (6)
	in the next 1 year? (1)	in the next 2 years? (2)	in the next 3 years? (3)	in the next 4 years? (4)	in the next 5 years? (5)	
First patent application approved	0.012*** <i>0.002</i>	0.020*** <i>0.003</i>	0.023*** <i>0.003</i>	0.024*** <i>0.003</i>	0.025*** <i>0.003</i>	0.004*** <i>0.001</i>
Log (1 + no. prior VC rounds)	0.012*** <i>0.002</i>	0.020*** <i>0.003</i>	0.023*** <i>0.003</i>	0.024*** <i>0.003</i>	0.025*** <i>0.003</i>	0.044*** <i>0.004</i>
Diagnostics						
Mean of dep. variable	3.9%	5.7%	6.4%	6.8%	7.0%	0.78%
Median no. months from first-action to VC round or IPO for successful applicants	5.2	8.1	9.2	10.0	10.3	65.7
R^2	35.3%	44.0%	46.3%	46.2%	46.0%	15.4%
No. of observations (firms)	34,167	34,111	34,060	34,013	33,981	34,215

Table IA.12. Does a Startup’s First Patent Application Affect Access to Debt? OLS Results.

The table reports the results of estimating equation (4) to examine how the approval of a startup’s first patent application affects the startup’s ability to raise debt. The analysis here is analogous to Table 8, with the only difference being that we use OLS instead of 2SLS. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	Following the first-action decision on its first patent application, does the startup pledge the application as collateral ...					
	in the next 1 year? (1)	in the next 2 years? (2)	in the next 3 years? (3)	in the next 4 years? (4)	in the next 5 years? (5)	at any point in the future? (6)
Panel A: All firms						
First patent application approved	0.004*** <i>0.001</i>	0.011*** <i>0.002</i>	0.022*** <i>0.002</i>	0.033*** <i>0.003</i>	0.044*** <i>0.003</i>	0.068*** <i>0.003</i>
Diagnostics						
Mean of dep. variable	1.1%	2.3%	3.3%	4.3%	5.2%	7.2%
Median no. of months from first-action to application pledge	6.0	11.8	17.8	23.0	27.5	38.9
R^2	10.4%	10.6%	10.6%	10.5%	10.6%	11.5%
No. of observations (firms)	33,520	33,520	33,520	33,520	33,520	33,520
Panel B. Firms without VC funding						
First patent application approved	0.003** <i>0.001</i>	0.009*** <i>0.002</i>	0.017*** <i>0.002</i>	0.025*** <i>0.002</i>	0.034*** <i>0.002</i>	0.054*** <i>0.003</i>
Diagnostics						
Mean of dep. variable	0.8%	1.7%	2.5%	3.2%	3.9%	5.5%
Median no. of months from first-action to application pledge	5.9	12.2	18.6	23.0	27.5	39.6
R^2	10.5%	10.4%	10.0%	9.8%	9.7%	10.3%
No. of observations (firms)	31,161	31,161	31,161	31,161	31,161	31,161

Table IA.13. How Do Patents Affect Access to VC Funding? OLS Subsample Analyses.

The table examines how the effect of patent grants on facilitating access to VC funding varies across different subsamples. The analysis here is analogous to Table 9, with the only difference being that we use OLS instead of 2SLS. Heteroskedasticity consistent standard errors clustered at the art unit level are shown in italics underneath the coefficient estimates. We use ***, **, and * to denote significance at the 1%, 5%, and 10% level (two-sided), respectively.

	In the three years following the first-action decision on its first patent application, does the startup raise ...							
	its first VC round?	its second VC round?	a higher VC round?	any VC funding?				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
First patent application approved	0.008*** <i>0.002</i>	0.199*** <i>0.067</i>	0.188*** <i>0.033</i>	0.167*** <i>0.054</i>	0.016*** <i>0.003</i>	0.044*** <i>0.006</i>		
... × inexperienced founder				-0.034 <i>0.081</i>				
... × high startup agglomeration state					0.013*** <i>0.004</i>			
... × biochemistry						-0.031*** <i>0.008</i>		
... × other industries						-0.030*** <i>0.006</i>		
Second application approved							0.033*** <i>0.006</i>	0.051*** <i>0.013</i>
Inexperienced founder				-0.045 <i>0.063</i>				
High startup agglomeration state					-0.007 <i>0.006</i>			
Diagnostics								
R^2	13.0%	59.1%	61.5%	50.4%	46.3%	46.4%	54.4%	57.6%
Mean of dep. variable	1.7%	46.8%	41.3%	61.3%	6.4%	6.4%	11.3%	10.2%
No. of observations (startups)	31,057	406	1,306	1,086	34,060	34,060	12,455	2,782