How can Innovation Screening be Improved? A Machine Learning Analysis with Economic Consequences for Firm Performance

Xiang Zheng*

December 28, 2021

Abstract

Using USPTO patent application data, I apply a machine-learning algorithm to analyze how the current patent examination process in the U.S. can be improved in terms of granting higher quality patents. Surprisingly and unsurprisingly, combining human examiners' expertise with machine learning predictions would yield a 15.5% gain in patent generality and a 35.6% gain in patent citations for granted patents. My economic analysis in the second part of the paper shows that getting a patent on which humans and machines disagree is a winner's curse. First, these patents are more likely to expire early. Second, patents granted by examiners with higher false acceptance rates cause lower announcement returns around patent grant news for public firms. Third, public firms who get these patents are more likely to get sued in patent litigation cases. Consequently, these firms cut R&D investments and have worse operating performance. Fourth, private firms whose patents are granted by such examiners are less likely to exit successfully by an IPO or an M&A. Overall, this study suggests that the social and economic cost of the current patent screening system is large and could potentially be mitigated with the help of a machine learning algorithm.

Keywords: Machine Learning, Patent Screening, Economic Consequences, Firm Performance
JEL classification: C55, G32, O31

*Assistant Professor of Finance, School of Business, University of Connecticut, Storrs, CT 06269, Email: xiang.zheng@uconn.edu.

I am grateful to my dissertation advisors, Thomas Chemmanur, Francesco D'Acunto, Jonathan Reuter, and Hassan Tehranian for their invaluable guidance. For helpful comments and discussions, I thank Rui Albuquerque, Ran Duchin, Slava Fos, Edith Hotchkiss, Seung Jung Lee (discussant), Svetlana Petrova (discussant), Alberto Rossi, Ronnie Sadka, Philip Strahan, Wanli Zhao (discussant), conference participants at the Doctoral Consortium of the Financial Management Association Annual Meeting 2020, Southern Finance Association Annual Meeting 2020, Eastern Finance Association Annual Meeting 2021, Financial Management Association Annual Meeting 2021, and seminar participants at Boston College, California State University at Fullerton, Chinese University of Hong Kong (Hong Kong & Shenzhen), City University of Hong Kong, Durham University, Fudan University (Fanhai), Georgetown University Global Virtual Seminar Series on Fintech, Luohan Academy, Oregon State University, Peking University (Guanghua), San Diego State University, Shanghai Jiaotong University (Antai), Shanghai University of Finance and Economics, University of Connecticut, and University of Florida. Any errors and omissions remain my own responsibility.

How can Innovation Screening be Improved? A Machine Learning Analysis with Economic Consequences for Firm Performance

Abstract

Using USPTO patent application data, I apply a machine-learning algorithm to analyze how the current patent examination process in the U.S. can be improved in terms of granting higher quality patents. Surprisingly and unsurprisingly, combining human examiners' expertise with machine learning predictions would yield a 15.5% gain in patent generality and a 35.6% gain in patent citations for granted patents. My economic analysis in the second part of the paper shows that getting a patent on which humans and machines disagree is a winner's curse. First, these patents are more likely to expire early. Second, patents granted by examiners with higher false acceptance rates cause lower announcement returns around patent grant news for public firms. Third, public firms who get these patents are more likely to get sued in patent litigation cases. Consequently, these firms cut R&D investments and have worse operating performance. Fourth, private firms whose patents are granted by such examiners are less likely to exit successfully by an IPO or an M&A. Overall, this study suggests that the social and economic cost of the current patent screening system is large and could potentially be mitigated with the help of a machine learning algorithm.

Keywords: Machine Learning, Patent Screening, Economic Consequences, Firm Performance
JEL classification: C55, G32, O31

1. Introduction

"The strength and vitality of the U.S. economy depends directly on effective mechanisms that protect new ideas and investments in innovation and creativity."

– The U.S. Patent and Trademark Office

The patent system is viewed as one of the most important institutions that provide firms with innovation incentives by granting them temporary monopoly rights over their inventions. This, in turn, contributes to the technological growth in the economy (see, e.g., Nordhaus, 1969; Arrow, 1972; Mansfield, 1986). However, there has been considerable criticism of the patent system for granting too many low-quality patents through an inefficient screening process (see, e.g., Heller and Eisenberg, 1998; Jaffe and Lerner, 2011; Feng and Jaravel, 2020). Critics argue that inefficient screening of patent applications reduces, instead of increases, firms' incentives to innovate (see, e.g., Cornelli and Schankerman, 1999; Lemley and Shapiro, 2005; Schankerman and Schuett, 2016; Bessen and Maskin, 2009). Many factors may have contributed to this issue. First of all, patent examiners have faced increasing time constraints over time. On the one hand, the number of patent applications filed at the U.S. Patent and Trademark Office (USPTO) has skyrocketed over the last two decades. For example, Figure 1 shows that the number of patent applications filed at the USPTO from 2001 to 2018 has increased from 345,732 in 2001 to 643,303 in 2018.[1] On the other hand, Figure 1 also shows that the number of patent examiners working in the USPTO can not kept up the same pace as the number of newly filed application. As a result, patent examiners spend only 19 hours, on average, reviewing an application, but it takes around 25 months for an application to get its screening result (Frakes and Wasserman, 2017). Second, the USPTO also faces human capital constraints: it constantly fails to recruit and retain the best examiners due to fierce competition from the booming private sectors (Jaffe and Lerner, 2011). Last, the incentive structure in the USPTO favors acceptances over rejections (Merges, 1999; Frakes and Wasserman, 2015). For example, an examiner's compensation directly ties to the number of patent applications that he/she has finished reviewing. However, it usually takes less efforts to accept an application than to reject one. Consequently, inefficient patent screening not only reduce firms' incentives to

[1] Data source: U.S. Patent Statistics Chart and Patent Examination Data from the USPTO website.

innovate but also mislead investors, and hurt firm performance.

Motivated by the above facts and critics of the current patent screening system in the U.S., this paper explores how to improve the patent screening process with the help of machine learning algorithms. On the one hand, unlike humans, machine learning algorithms are able to process large datasets quickly, which could potentially relax the time constraints faced by patent examiners and resource constraints faced by the USPTO. On the other hand, machine learning algorithms neither have career concerns nor do they have (compensation) incentive issues. Therefore, they could potential reduce the agency frictions between the principle (the USPTO) and the agent (the patent examiner). In fact, this idea is partly supported by the patent office itself according to a recent news article published in the Wall Street Journal. The patent office is currently seeking help from artificial intelligence (including machine learning) to drive efficiencies in the patent examination process.[2] The director of the USPTO, Andrei Iancu, said in the news that "our need is high and technology has advanced, so this is a good time to take advantage of these new tools to help our examiners."

The key idea here is that the patent screening process can be viewed as a prediction process. To fulfill the mandate of the Patent and Copyright Clause of the Constitution, the U.S. Patent Act (35 U.S. Code §101 – §103) requires a granted patent to be "new," "useful," and "non-obvious" with its purpose of making new discoveries public knowledge in the future by rewarding inventors with a limited exclusive right on their invention.[3] Therefore, the grant of a patent based on the U.S. patent law hinges on the prediction of its future social value to the society by a patent examiner. More specifically, I argue that the patent office's objective is to grant higher quality patents (i.e., patents with higher social values) while rejecting lower quality patent applications. The USPTO itself also supports this argument in its 2018–2022 strategic plan, that is, the most important goal for the office is to continue optimizing patent quality. [4] However, as discussed earlier on, the objective of patent examiners may not be closely aligned with the objective of the patent office due to various problems: i.e., time constraints, talent constraints, career concerns, and compensation incentives.

[2]For the full news story, please see: https://www.wsj.com/articles/patent-office-seeks-help-from-ai-11572297295.

[3]See the "Patent and Copyright Clause" of the U.S. Constitution. To quote: [Congress shall have the power] "to promote the progress of science and useful arts, by securing for limited times to authors and inventors the exclusive right to their respective writings and discoveries."

[4]For the full USPTO 2018-2022 strategic plan, please see:
https://www.uspto.gov/sites/default/files/documents/USPTO_2018-2022_Strategic_Plan.pdf.

Therefore, it is possible that a machine learning algorithm can better execute the task and mitigate the misalignment problem.

Using detailed data on both granted and rejected applications that are recently available from the USPTO website, I train a supervised machine learning algorithm that maps patent application characteristics to patent quality using earlier patent applications (i.e., with standard quality measures that capture patent's social value innovativeness).[5] I then use this trained algorithm to predict the quality of more recent patent applications out-of-sample. My out-of-sample prediction results show that the current patent examination system grants many low-quality patents while rejecting many high-quality patent applications. I also show that the above machine learning algorithm performs significantly better than an OLS regression function out-of-sample, in terms of predicting standard quality measures of the patents such as "citation" and "generality" measures that capture the social value of patents that have been used extensively in the literature (see, e.g., Trajtenberg, Henderson, and Jaffe, 1997; Hall, Jaffe, and Trajtenberg, 2005; Chemmanur, Loutskina, and Tian, 2014).[6]

Next, I want to test whether human examiners can do a better job with the help of such a machine learning algorithm. Ideally, one wants to compare the average quality of patent applications granted by human examiners with the help of an algorithm to the average quality of the patents granted by human examiners along. However, the main challenge here is the missing counterfactuals: we do not observe actual quality information for applications rejected by humans but accepted by the algorithm.[7] To address this selection issue, I make use of the quasi-random assignment of patent applications to examiners who have different levels of leniency (or, in other words, different grant rates).[8] Following the methodology first introduced by Kleinberg et al. (2017), I divide ex-

[5]The machine learning algorithm used in this paper falls into the category of supervised learning, namely, training a prediction function that maps inputs (X) to an output (y) based on training input-output pairs. The inputs (X) used in my setting include numerical statistics of claims text, the text-based numerical vector of claims that capture the text similarity across contemporaneous patent applications, backward citations from prior patents, patent applications, foreign patents, and scientific literature, the total number of novel words, filing year dummies, inventor nationality dummy, small entity dummy, NBER classes dummies, and art unit dummies. The output (y) used in my setting includes the generality index of patents, and forward citation counts of patents.

[6]Machine learning generally makes much more accurate (out-of-sample) predictions by imposing fewer restrictions on the prediction function form compared to traditional statistical tools.

[7]This challenge is not unique in this setting. It shows up in most machine learning applications trying to improve screening efficiency (e.g., recruiting decision, admission decision, and bail decision). See Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan (2017) for a detailed discussion.

[8]Because of the quasi-random assignment, I argue that the average quality of patent applications reviewed by examiners with different levels of leniency is similar. Many recent studies exploit this feature to make causal inferences in their studies (see, e.g., Maestas, Mullen, and Strand, 2013; Farre-Mensa, Hegde, and Ljungqvist, 2020; Sampat and

aminers into two groups based on their leniency (grant rate): i.e., more lenient examiners approve 78% of applications in my sample while less lenient examiners accept only 50% of patent applications. We can view these two group of examiners as two independent patent screening systems. I focus only on patents which are already granted by more lenient examiners: (1) I first rank them based on the above machine learning algorithm; (2) I then reject an additional 28% of patents from the lowest predicted quality to match the grant rate of less lenient examiners (i.e., reduce the grant rate of more lenient examiners from 78% to 50%). This identification strategy enables me to compare the observable quality of patents granted by more lenient examiners with the help of an machine learning algorithm to those granted by less lenient examiners. Normally, we would expect that less lenient examiners are the better examiners, given that they set a higher bar to approve patent applications. However, we find that more lenient examiners, with the help of a machine learning algorithm, are able to do a significantly better job than less lenient examiners in terms of granting higher-quality patents.

To explain the above identification more clearly, Figure 2 provides an illustrative example of the above exercise. More lenient examiners approve 700 applications, while less lenient examiners approve 500 applications. I first rank the 700 applications granted by more lenient examiners based on their predicted quality and then reject additional 200 patents from the lowest predicted quality. After that, I can quantify the quality gain made by an algorithm by comparing observable quality measures (such as the patent generality index or the number of patent citations) of the 500 patents granted by the more lenient examiners with the help of an algorithm and the 500 patents granted by less lenient patent examiners. I find such comparison yields to economically significant improvements in patent quality: an algorithm trained against patent generality results in about a 15.5% gain in patent generality and a 35.6% gain in the number of patent citations compared to decisions made by less lenient examiners. In other words, the more lenient examiners (or worse-performed examiners) can significantly out perform less lenient examiners (or better-performed examiners) with the help of a machine learning algorithm. These results also demonstrate that a machine learning algorithm not only results in significant improvements in an objective that is targeted by the algorithm (i.e., generality) but also results in significant improvements in an alternative measure

Williams, 2019).

of patent quality (i.e., the number of citations) that is not targeted.[9] In addition, the above analysis also provides some suggestive evidence on why patent examiners fall short in making screening decisions. For example, the algorithm suggests that important factors (predicting patent generality) include a numerical vector that captures the text-similarity across different patent applications and a measure of patent application originality that captures the extent of combining knowledge from different technological fields by a given application. These factors require patent examiners either to spend a significant amount of time or to do a more careful job. I also find that busy examiners, more experienced examiners, and male examiners make more false acceptance mistakes. All these findings echo back to the various constraints and issues faced by the patent office.

So far, the findings suggest that examiners could do a better job with the help of a machine learning algorithm in terms of granting higher quality patents. Does such an algorithm also improve the economic outcomes of firms? To answer this question, the second part of the paper examines the economic consequence of the current patent screening process on firm performance. To do so, I first label patents that would be rejected by the algorithm to be "falsely accepted". I then construct an *ex-ante* screening efficiency measure for each examiner by computing his/her past false acceptance rate. I find that firms who get patents from examiners with higher false acceptance rates suffer from the winner's curse. In particular, I find that patents granted by such examiners have, on average lower announcement returns around their grant news. These patents are also more likely to expire early. In other words, these patents are winner's curses for their owners. Further, I find that public firms whose patents are granted by such examiners are more likely to get sued subsequently in both the short-term and long-term future. Consequently, they cut R&D expenditures and have worse operating performance (measured by either ROA or Cash Flow). Additionally, such a negative impact is more significant for firms in the high-tech or health industries. In the cases of private firms, they are less likely to exit successfully by an IPO or an M&A in the short-term and long-term future. The above effects are economically significant. For example, the annual ROA for public firms would increase by 1.3 percentage points, and the probability of private firms going public or getting acquired in three years would increase by 3.6 percentage points if the above machine

[9] Although this study finds that machine learning algorithms can make better screening decisions in terms of granting higher quality patents, replacing human examiners with machine learning algorithms may incur unintended consequences. Instead, such am algorithm can serve as an auditing process, in which examiners are responsible to reexamine those patent applications identified as questionable screenings. Combining the expertise of human examiners and the strength of machine learning mitigates unintended consequences while achieving better screening outcomes.

learning algorithm screened all patent applications. The above results can be viewed as causal evidence since patent applications are randomly assigned to patent examiners whose characteristics are unlikely to be correlated with firm characteristics.

The rest of the paper is organized as follows. Section 2 discusses the relation of my paper to the existing literature. Section 3 discusses the institutional background of the patent examination process. Section 4 describes the patent application data and sample statistics. Section 5 discusses the empirical design and results of the machine learning analysis. Section 6 describes the firm-level data and discusses the empirical analysis of firm performance. Section 7 concludes.

2. Relation to the existing literature

My paper is related to four different strands in the literature. The first strand is the theoretical and legal literature that explores the question of improving the patent screening process by reforming the patent system itself. For example, Dreyfuss (2008) argues that the patent system systematically creates type II errors (i.e., erroneous grants) due to resource constraints faced by patent examiners and the incentive structure at the USPTO. Dreyfuss (2008) proposes to increase the nonobviousness threshold in order to reduce the number of type II errors (see also, e.g., Duffy, 2008; Eisenberg, 2008; Mandel, 2008). On the other hand, Scherer (1972) and several other theoretical papers focus on reforming the optimal patent right (i.e., patent length and breadth) to improve the innovation incentive and quality (see, e.g., Gilbert and Shapiro, 1990; Matutes, Regibeau, and Rockett, 1996). Finally, a set of related papers also study the cost and benefit of the patent litigation system in affecting patent validity and scope (see, e.g., Meurer, 1989; Choi, 1998; Lanjouw and Schankerman, 2001; Bessen and Meurer, 2006). In this paper, I depart from the above literature and analyze how machine learning techniques are able to improve the effectiveness of the patent screening process without changing the current patent system itself. Additionally, I also document the economic consequence of the current patent screening system on firms owning these patents.

The second strand is the literature that applies machine learning techniques to economics and finance research. For example, Athey and Imbens (2017) argue that supervised machine learning has great potential for prediction problems but has not been widely utilized in social science research. Several studies apply machine learning to issues in finance: e.g., measuring asset risk premia (Gu, Kelly, and Xiu, 2020), predicting stock returns (Rossi, 2018), classifying fund types (Abis, 2017), and selecting the boards of directors (Erel, Stern, Tan, and Weisbach, 2018). However, there also exist challenges to apply machine learning in social science research. Kleinberg et al. (2017) use New York judges' decisions over bail cases as a setting to discuss unique potential endogeneity problems when applying machine learning to social science and provide methodologies to address these problems using econometric identifications.[10] However, using an algorithm to screen human activities may result in unintended algorithm biases. For example, training an algorithm using income or education may implicitly include hidden information such as race and gender, thereby enhancing existing racial and gender biases (inequality). In contrast, applying machine learning in evaluating the technological advance is less subject to such biases resulting from hidden information. Overall, mine is the first paper to apply machine learning algorithms to evaluate patent examiners' innovation screening efficiency and make use of the quasi-random assignment of patent applications to patent examiners to address potential selection issues.[11]

Third, my paper also contributes to the empirical literature that studies the relationship of patent quality and firm performance. For example, Hall et al. (2005) empirically document that a larger number of citations per patent leads to higher market values for firms holding these patents (see also, e.g., Zucker, Darby, and Armstrong, 2002). Chemmanur, Gupta, and Simonyan (2017) also show that private firms with a large number of patents and citations per patent have higher IPO valuations and future operating performance. However, the innovation measures used in these studies are only *ex-post* available. Alternatively, Bowen III, Frésard, and Hoberg (2020) measure the disruptive technological potential of startups using textual analyses and show that those firms with higher disruptive technological potentials are more likely to go public and are less likely to be sold. Kelly, Papanikolaou, Seru, and Taddy (2018) also use the textual analysis method to create indicators of technological innovation for each patent based on its textual similarity to earlier and

[10] See also Kleinberg, Ludwig, Mullainathan, and Obermeyer (2015), and Mullainathan and Spiess (2017) for detailed discussions on how to use machine learning as an applied econometrics tool.

[11] In addition, my paper also discusses and addresses the forward-looking issue that existed in both the training data and the training algorithm itself. For example, when training the algorithm using future information, it will result in unfair comparisons between humans and machines in the test set since humans in the test set are not able to access such future information. The existing literature largely overlooks this forward-looking issue in training algorithms (see, e.g., Erel, Stern, Tan, and Weisbach, 2018; Kleinberg, Ludwig, Mullainathan, and Obermeyer, 2015).

later patents. Kogan, Papanikolaou, Seru, and Stoffman (2017) measure the economic value of a patent as the stock price announcement effect of the patent grant and study its relationship with aggregate economic growth and TFP. Kline, Petkova, Williams, and Zidar (2019) follow a similar approach to estimate the *ex-ante* value of accepted and rejected patent applications and study the relationship between patent-induced shocks and labor productivity. Unlike the measures used in these above papers, the screening efficiency measure of patent examiners constructed in my paper can be viewed as an *ex-ante* measure of patent application quality. It is also less likely to be related to firm characteristics, given that patent applications are randomly assigned to each patent examiner within each art unit.

A set of recent papers also exploit the quasi-random assignment of applications to examiners with different leniency to make causal inferences on the relationship between current innovation and follow-on innovation. For example, Farre-Mensa et al. (2020) find that obtaining its first patent causally increases a startup's subsequent growth, follow-on innovation, and VC funding. On the other hand, Sampat and Williams (2019) examine whether patents in the field of human genes affect follow-on innovation and find that gene patents, on average, have no quantitatively important effects on follow-on innovation. Unlike these papers, my paper focuses on the economic consequences of weak screening by patent examiners and studies its impact on the future performance of both public and private firms. Overall, my paper complements the above literature by documenting causal evidence of the importance of corporate innovation on subsequent performance and investment of both public and private firms.

Finally, my paper is also related to the strand of literature that analyzes the value of innovations by examining stock market reactions to innovation-related announcements. For example, Eberhart, Maxwell, and Siddique (2004) examine the market valuation of firms' innovation inputs (R&D expenditures) and show that the market consistently underreacts firms' unexpected increases in R&D expenditures. Cohen, Diether, and Malloy (2013) also show that the stock market does not take firms' past successes in innovation into considerations when valuing their future innovation. Fitzgerald, Balsmeier, Fleming, and Manso (2021) show that firms with exploitation innovation strategies are undervalued relative to firms with exploration innovation strategies. On the other hand, Hirshleifer, Hsu, and Li (2013) explore the market valuation of firms' output and show that firms' innovation efficiency (measured as patents scaled by R&D expenditures) can predict firms' future stock returns. Shu, Tian, and Zhan (2019) test whether the workload of each patent examiner can predict firms' future stock market returns and show that investors underreact to the negative effect of examiner's workload on patent quality. My paper complements the above literature by providing additional evidence that the stock market incorporates (at least partially) the quality of firms' new patents from the past performance of the patent examiners examining these patent applications prior to patent grants.

3. Patent examination process and patentability

3.1. Patent examination process

The patent examination process starts with filing a patent application to the USPTO, where the USPTO will forward this newly filed application to a relevant art unit for examination.[12] Next, that patent application will be assigned to a patent examiner, a specialized technology employee with training and experience pertinent to the invention, for examination. Though there are no explicit policies regarding how patent applications are assigned to examiners within each art unit, many recent studies show that patent applications are randomly assigned to examiners within each art unit: an application that has filed in the earliest date is assigned to the first available examiner (see, e.g., Maestas, Mullen, and Strand, 2013; Farre-Mensa, Hegde, and Ljungqvist, 2020; Sampat and Williams, 2019).

After receiving a patent application, examiners first compare the claimed invention to the existing state of knowledge in the "prior art," consisting of patent documents as well as the scientific and commercial literature to determine whether the invention satisfies legal requirements for patentability. If an invention fails the patentable requirement, the examiner will issue an office action rejecting that application as not patentable and explain the reasons for the rejection. Following such a rejection, the inventor may revise the application and submit it again or withdraw it. My paper only focuses on the earliest application of all regular non-provisional utility applications to mitigate the concern that these subsequent applications may not be randomly assigned (Righi and

[12]There are nine patent examining group centers where each of them consists of several art units examining patents in the relevant field.

Simcoe, 2018).

3.2. *The legal requirements for patentability*

Patent and Copyright Clause of the Constitution (Article I, Section 8, Clause 8, of the Constitution) grants Congress the power "to promote the progress of science and useful arts, by securing for limited times to authors and inventors the exclusive right to their respective writings and discoveries." To fulfill its mandate, the U.S. Patent Act (35 U.S. Code §101) sets the requirements for patent protection as follows:

"Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvements thereof, may obtain a patent, subject to the conditions and requirements of this title."

Under the U.S. Patent Act, an invention is patentable after satisfying the following three criteria: new, useful, and non-obvious. Specifically, the novelty requirement (35 U.S. Code §102) states that an invention cannot be patented if the invention has been publicly disclosed before the applicant filed for patent protection and the usefulness requirement states that the subject matter must be useful. Usually, a patent application can easily pass both the novelty and usefulness requirements. However, the non-obvious requirement (35 U.S. Code §103), which requires the invention to be a non-obvious improvement over the prior art, is an ambiguous threshold that attracts many criticisms from the law literature for approving many low-quality patents (see, e.g., Duffy, 2008; Dreyfuss, 2008; Eisenberg, 2008; Mandel, 2008).

Since the goal of U.S. Patent Act is to reward patent applicants with a limited exclusive right on their invention for providing new discoveries to the public, I argue that the main objective for patent examiners is to identify and grant patents of higher quality (or higher social value) while rejecting those of lower quality.

3.3. *Measuring patent quality*

Recent papers start to use excess stock market returns to measure firms' private value of a patent (Kogan, Papanikolaou, Seru, and Stoffman, 2017). However, the private value of a patent is unlikely to capture the objective of patent examiners for the following reasons. First, the private value of a given patent depends not only on its own quality but also on whom the patent belongs to: i.e., a patent may have different private values to different owners, while examiners make grant decisions based on the characteristics of a patent application itself. Second, the private value of a patent can only be measured if it is filed by a public firm, while examiners also need to evaluate patent applications filed not only by public firms but also by private firms, governments, universities, and individual inventors to make grant decisions. On the other hand, citation-based measures, which have been used extensively in existing literature (see, e.g., Trajtenberg, 1990; Trajtenberg, Henderson, and Jaffe, 1997; Hall, Jaffe, and Trajtenberg, 2005), not only are available for any patent granted by the USPTO regardless of whom filed the patent application but also, more importantly, capture the social value (or social spillovers) of a patent (Bloom, Schankerman, and Van Reenen, 2013).

In this paper, I use patent generality as my primary measure of patent quality: the generality index of a patent captures the industry dispersion of citing patents in the following four years after being granted.[13] Explicitly, I compute the generality index following the existing literature (see, e.g., Trajtenberg, Henderson, and Jaffe, 1997; Hall, Jaffe, and Trajtenberg, 2005): $G_i = 1 - \sum_j^{n_i} s_{ij}^2$, where s_{ij} denotes the fraction of forward citations received by patent i in patent class j from the total number of patent classes n_i and $\sum_j^{n_i} s_{ij}^2$ is the Herfindahl-Hirschman index (Hirschman, 1980). By definition, if a patent is cited by later patents that belong to more fields, the generality of this patent will be higher. For example, if subsequent patents cite a patent in the field of biology in social science, medical science, and engineering, we would expect this patent to have a higher degree of generality than a similar patent that received the same number of citations but all from patents in the same field. With regard to patent classes, the USPTO has developed its own U.S. Patent Classification (USPC) system that consists of more than 450 unique classes and 150,000 subclasses. However, USPC classes provide no straightforward link to the established product and industry classifications (Marco, Carley, Jackson, and Myers, 2015a). Hall, Jaffe, and Trajtenberg (2001) developed a hierarchical classification (NBER classification) by aggregating USPC classes into 37 (two-digit) sub-categories.[14] Therefore, I construct two generality measures based on either

[13] I have also used citation counts (the number of citing patents in the following four years after a patent gets granted) as an alternative measure of patent quality. The results using citation counts are reported in Section IA.2.2 in the Internet Appendix and are robust to the findings in the main paper.

[14] The NBER classification comes from the NBER Patent Data Project: https://sites.google.com/site/patentdataproject.

the USPC classification or the NBER classification. All results presented in empirical sections are using the generality measure based on the NBER classification.[15]

4. Patent application data and sample selection

4.1. *Patent application data*

I collect data on patent applications from the USPTO website that provides various research datasets.[16] In particular, I collect patent application examination data from Patent Examination Research Dataset (Graham, Marco, and Miller, 2018; Marco, Toole, Miller, and Frumkin, 2017), patent application claims data from Patent Claims Research Dataset (Marco, Sarnoff, and Charles, 2019), patent application citation data from Office Action Research Dataset for Patents (Lu, Myers, and Beliveau, 2017) and PatentsView, and patent assignment data from Patent Assignment Dataset (Marco, Myers, Graham, D'Agostino, and Apple, 2015b).[17]

4.1.1. *Turning patent claims text into numerical variables*

The claim section in each patent application defines the extent of the protection sought in a patent application. A typical patent contains several claims, where each claim represents an original contribution and thereby being viewed as a good measure of the real invention in a patent (Tong and Frame, 1994). If claims in a patent application are very similar or closed to claims in other patent applications, we would expect that this patent application's quality (innovativeness) to be low. To capture the similarity of each patent application filed in a given year compared to all patent applications filed in that year, I take all claims text in each patent application to produce

[15] The results are quantitatively similar using the generality measure based on the USPC classification and are reported in Section IA.2.1 in the Internet Appendix.

[16] For a complete list of research datasets provided by the USPTO please see: https://www.uspto.gov/ip-policy/economic-research/research-datasets.

[17] Public PAIR data have been recently available from the USPTO website. Though not all patent applications received by the USPTO are included in Public PAIR, more than 83% of all patent applications are available after the implementation of The American Inventors Protection Act (AIPA) in late 2000. For regular utility patent applications that this paper focuses on, inclusion in Public PAIR increases to 95% since 2001 as a consequence of AIPA according to Graham et al. (2018).

a vector of 50 dimensions from claims text using the *Word2vec* algorithm.[18] I use this vector of 50 numerical variables as well as numerical statistics of claims and other patent application characteristics discussed in the later section as input variables in my machine learning prediction, where I find the prediction accuracy of the machine learning algorithm is improved with this set of text-based variables.

4.2. *Summary statistics*

Table 1 reports summary statistics of numerical variables for all patent applications used in my machine learning prediction. Out of 637,305 applications, 434,496 (68.2%) are approved, 236,643 of which have non-zero 4-year forward citations: the average 4-year forward citations and the generality index per patent among patents with non-zero citations are 3.886 and 0.072, respectively. In terms of numerical statistics of claims, each patent application on average has 2.791 independent claims and 15.528 dependent claims, where the average length of an independent claim (around 138 words) tends to be longer than that of a dependent claim (around 42 words). The average number of novel words per patent is 0.309.[19] I also compute the originality index for each patent application, which is defined similarly as generality except that it is based on backward citations each application has made. The average backward patent citations and the originality index are 8.511 and 0.166, respectively. In addition to citing prior patents, a patent application may also cite previous applications, scientific literature, and foreign patents. The average backward citations from patent applications, scientific literature, and foreign patent citations are 2.755, 3.837, and 2.905, respectively.[20] Besides patent application characteristics, 26.9% of patent applications are submitted by small entities, and 43.9% of primary inventors are from the U.S.

[18]The *Word2vec* algorithm learns vector representations of words from the input text corpus and places words that share similar context in the corpus in close proximity to one another in the vector space, where the vector space is set to 50 dimensions (see, e.g., Mikolov, Chen, Corrado, and Dean, 2013a, Mikolov, Le, and Sutskever, 2013b, and Mikolov, Yih, and Zweig, 2013c for details).

[19]The number of novel words for each patent is produced by Balsmeier, Assaf, Chesebro, Fierro, Johnson, Johnson, Li, Lück, O'Reagan, Yeh, et al. (2018), which I used as an input variable when I train the algorithm. My results remain quantitatively similar without including the number of novel words.

[20]I exclude citations made by examiners when counting backward citations for each patent application.

5. Machine learning prediction design and results

The empirical design to analyze the efficiency of the patent screening process follows three steps (Kleinberg, Lakkaraju, Leskovec, Ludwig, and Mullainathan, 2017). First, I partition my sample into a training set and a test set, as described in Subsection 5.1. Second, I train an algorithm using the training set by mapping the characteristics a patent application to its quality and present results in Subsection 5.2. Third, I evaluate the predicting accuracy of this algorithm using patent applications in the out-of-sample test set and present results in Subsection 5.3. Last, I test whether my prediction function can improve screening decisions of actual patent examiners by comparing the decision of this algorithm to that of patent examiners and present relevant results in Subsection 5.4.

5.1. Sample partition

I use the unique application number to merge across different data sets and obtain an initial sample of 3,473,251 patent applications with screening outcomes available (i.e., either granted or rejected) filed at the USPTO from 2001 to 2014.[21] When we train a machine learning algorithm to compare its prediction with human decisions, we have to make sure the data used to train the algorithm is *ex-ante* available for actual examiners in the test set in order to make fair comparisons. In my setting, I use patent application characteristics and patent outcomes of earlier applications in the training set. Since my outcome variable for training the algorithm, the generality index of patent applications, is constructed based on 4-year forward citations and is only available four years after each application being granted, I set a 4-year gap between the training sample and the test sample. In particular, I use applications filed from 2001 to 2005, which have their screening status available before 2006 for the training sample to train the machine learning algorithm, and use applications filed from 2010 to 2013 with their status available before 2014 for the test sample to evaluate the algorithm.[22]

[21]The patent application claims data from the "Patent and Patent Application Claims Research Dataset for Academia and Researchers" section is available until the end of 2014.

[22]I have partitioned my sample using alternative ways: partition the whole sample randomly to a training sample and test sample; partition the whole sample along the time but without a 4-year gap. Though these alternative ways of sample partition are subject to concerns raised in this section, results using these alternative ways of sample partition are similar to the main findings in this paper.

Partitioning my sample in this way, both my trained machine learning algorithm and the quality measure of patent applications in the training sample is available at the beginning of 2010. In other words, whatever information is needed to train the algorithm is also available for patent examiners in the test sample. Such a sample partition allows me to make a fair comparison between the algorithm and actual examiners in terms of screening any patent application in the test set. Figure 3 presents the sample partition along the timeline. The final sample used in my machine learning prediction consists of 280,243 patent applications in the training set and 357,101 patent applications in the test set.

5.2. *Training a machine learning algorithm*

To train a supervised machine learning algorithm, I need both input variables of patent application characteristics and an output variable of patent application quality from applications in the training data: the output variable y is the generality index of each patent as described in Subsection 3.3; and input variables, X, include numerical statistics of claims text as described in Subsection 4.2, the text-based numerical vector of claims, backward citations from prior patents, patent applications, foreign patents, and scientific literature, the total number of novel words, filing year dummies, inventor nationality dummy, small entity dummy, NBER classes dummies, and art unit dummies. As I mentioned earlier in Subsection 5.1, my training set consists of 280,243 patent applications, including 81,352 rejected applications, 83,558 accepted applications with zero 4-year forward citations, and 115,333 accepted applications with the number of 4-year forward citations larger than zero. Since the number of the 4-year forward citation to construct the generality index of a patent (an accepted application) needs to be larger than zero, 115,333 accepted applications with their generality index available are used for training the machine learning algorithm.

I train the prediction function called "Extreme Gradient Boosting," an ensemble method of decision trees based on tree boosting.[23] A decision tree is a tree-like prediction function that can be trained by splitting the data set into subsets based on particular values of input variables, where the process is repeated until splitting no longer adds value to predictions (see, e.g., Rokach and Maimon, 2008). Since a single decision tree may produce a weak learning function subject to

[23] Section IA.1 of the Internet Appendix provides for a detailed discussion about the supervised machine learning problem and the Extreme Gradient Boosting algorithm.

noise, gradient boosting algorithms optimize a cost function by iteratively choosing a weak learning function that follows the negative gradient direction to produce a strong learning function (see, e.g., Friedman, 2001; Chen and Guestrin, 2016). The strength of an Extreme Gradient Boosting algorithm is finding the best feature across different subsamples. In addition, I implement 5-fold cross-validation when training the algorithm to alleviate the in-sample over-fitting problem.

Figure 6 shows 10 important features identified by the machine learning algorithm in terms of predicting patent generality. The most important feature is the numerical vector that captures the text-similarity across different patent applications filed in the same year. The 50 variables in this feature collectively explain 43.8% of the total predictive power in the trained algorithm. The second most important feature is the originality measure of patent applications, which captures the dispersed knowledge cited by each patent application. Together, these two features explain 70.9% of the total predictive power in the trained algorithm, suggesting that patent applications with original ideas are more likely to be high-quality patents. Other important features include the number of cited scientific literature, cited patents, claims, and words in claims. Interestingly, the inventor's nationality also explains 1.5% of the total predictive power in the trained algorithm.

5.3. Evaluating the out-of-sample predicting performance of machine learning and OLS

In this subsection, I compare the out-of-sample predicting performance between machine learning and an OLS function. In an OLS regression, I regress patent generality on all input variables used in the machine learning prediction with patent applications in the training sample. I then use the fitted model to predict the generality of patent applications in the test set. Figure 4 presents the correlation between predicted generality and actual generality using patent applications in the out-of-sample test set. The left panel of Figure 4 plots the predicted generality based on the machine learning algorithm against the actual generality of granted patents, where I find that most of the data is centered around the 45-degree line, suggesting that the accuracy of the out-of-sample prediction is high. Yet the right panel of Figure 4 plots the predicted generality based on an OLS regression against the actual generality of granted patents, where the out-of-sample fitting is much less close to the 45-degree line.[24]

[24]Formally, the out-of-sample mean square error (MSE) of the algorithm is 0.032. I also separately regress the actual generality on predicted generality by this machine learning algorithm (XGBOOST) and the OLS function. I find that the

Next, I test whether this algorithm can identify patent applications with the highest quality (i.e., so-called "tail innovation"). In particular, I compare the predicted generality distribution from a machine learning algorithm to that from an OLS function and present the results in Table 2. The second column of Table 2 shows that only 20.5% of patent applications identified as the top 1 % highest quality in the predicted generality by an ML algorithm are also identified as the top 1 % highest quality by an OLS function. The actual generality of patent applications identified as the top 1% generality by the machine learning algorithm, as reported in the third column of Table 2, is 0.171, which is significantly higher than that of patent applications identified as the top 1% generality by OLS reported in the fourth column of Table 2: 0.136. The difference between the machine learning algorithm and OLS is persistent and significant when we compare the results of the machine learning algorithm and OLS in terms of top 5%, 10%, and 25% of the quality distribution as reported in the second, third, and fourth rows of Table 2.

5.4. *Improve screening decisions with the help of a machine learning algorithm*

5.4.1. *Do examiners reject high-quality patents?*

To answer this question, I examine the grant rate of actual examiners across patent applications with different predicted generality. To visualize the results, I divide patent applications in the test set equally into 1,000 bins based on their predicted generality and compute the grant rate of patent applications made by actual examiners in each of these 1,000 bins. Figure 5 plots the correlation between the grant rate of actual examiners and the average predicted generality of patent applications in each bin. I find that the grant rate of examiners indeed increases with the predicted generality of patent applications. However, I also notice a significant portion of patent applications with very high predicted quality (i.e., patent applications in the rightmost bins) being rejected by actual examiners.[25]

coefficient of predicted generality by XGBOOST is 0.838, while that from OLS is 0.374.

[25]Figures IA.1 and IA.3 in the Internet Appendix show similar results using the generality measure based on the USPC classification and the number of citations.

5.4.2. *Using variation in the leniency of examiners to quantify the improvement of screening decisions by a machine learning algorithm*

One way to quantify the potential quality gain achieved by the algorithm is to rank all patent applications based on my predicted generality and then set the grant rate of the algorithm to be the same as that of examiners. I can then compare the average generality of all patent applications granted by the algorithm to the average generality of the actually granted patents. However, measuring the improvement in this way may be misleading since I do not have information on the actual generality of those patent applications rejected by examiners but approved by the algorithm. To address this issue, I make use of the fact that patent applications are randomly assigned to examiners who have different grant rates: more lenient examiners (i.e., with an above-median grant rate) accept around 77.6% of patent applications and less lenient examiners accept 49.5% of patent applications. Thus, given all patents granted by more lenient examiners, I can reject additional applications based on predicted generality to match the grant rate of less lenient examiners (i.e., examiners with a below-median grant rate). Now, I can compare the average actual generality of applications granted by the algorithm to that of applications granted by less lenient examiners.

More importantly, comparing across examiners with different leniency allows me to track the quality (generality) of marginal applications that get rejected. Figure 7 shows the results of such comparisons. I sort patent applications by predicted generality and divided them equally into 20 bins. At the bottom of a given bin, the black bar shows the fraction of patent applications being rejected by more lenient examiners. The red bar on the top of the black bar in a given bin shows the fraction of additional applications being rejected by less lenient examiners, while the blue bar on the top of the black bar in a given bin shows the share of additional applications that would be by the algorithm. The top panel of Figure 7 shows that less lenient examiners would reject additional applications from patent applications in both the low- and high-quality bins. However, the bottom panel of Figure 7 shows that the machine learning algorithm would reject additional applications starting from the lowest quality of predicted generality, suggesting that examiners do not screen out the low-quality applications identified by the algorithm.

Next, I quantify the quality gain for the above exercise by comparing the actual outcome resulting from examiners to that from the algorithm. I find that the magnitude of improvement in generality (by comparing the actual generality of patents granted by the algorithm to that granted by less lenient examiners) is 15.5%. Moreover, I find that training the algorithm using generality also significantly improves citations of granted patents. In particular, I find that the magnitude of improvement in citations (by comparing the actual 4-year forward citations of patents granted by the algorithm to that granted by less lenient examiners) is 35.6%.[26]

5.4.3. *Why do examiners underperform?*

To answer this question, I link examiners' characteristics with their screening performance. I measure examiners' screening performance based on the disagreement between the machine learning predictions and actual screening decisions of patent examiners. To do so, I first compute the number of applications granted by actual examiners within each art unit in any given year. Then I rank all patent applications filed within each art unit in a given year based on their predicted generality by the algorithm and hypothetically grant the same number of patent applications as examiners within each art unit in a given year. So far, each patent application has an actual grant decision made by examiners and a hypothetical grant decision made by the machine learning algorithm. Finally, I label a patent to be "falsely accepted" if it is accepted by an actual examiner but rejected by the algorithm and label a patent application to be " falsely rejected" if it is rejected by an actual examiner but accepted by the algorithm. I then construct the following four measures of examiners' screening performance in a given year: the number of falsely rejected cases, the false rejection rate, the number of falsely accepted cases, and the false acceptance rate. I also construct the following three measures of examiner's characteristics: work experience in a given year, the workload in a given year, and examiner gender. To identify the gender of each examiner, I make use of the following Social Security Administration's data set: National Data on the relative frequency of given names in the population of U.S. births where the individual has a Social Security Number.[27] This data set contains all given names and their associated genders with a population greater than 5. I match examiners' first names with this data set to obtain examiners' genders.[28]

[26] I have trained a similar algorithm using the number of 4-year forward citations to proxy patent quality, where I find that the magnitude of improvement in the number of citations reaches 28.7%. All results based on the number of citations are reported in Figures IA.3 and IA.4 in the internet appendix.

[27] To access this data set, please see: https://www.ssa.gov/oact/babynames/limits.html.

[28] When a given name is associated with both genders, I first calculate its probability of being a specific gender based on the gender-specific population and assign its gender with the probability $> 90\%$.

I test the relationship between examiners' characteristics with their screening performance with the following regression:

$$\begin{aligned} y_{i,t} \;\; = \;\; & \alpha + \beta_1 \textit{WorkExperience}_{i,t} + \beta_2 \textit{WorkLoad}_{i,t} + \beta_3 \textit{MaleExaminer}_i \\ & + \textit{Art Unit}_a + \textit{Status Year}_t + \epsilon_{i,t}, \end{aligned} \tag{1}$$

where i indexes patent examiner; a indexes art unit; and t indexes the issue year of a patent. y includes # False Rejection, False Rejection Rate, # False Acceptance, and False Acceptance Rate. $\textit{WorkExperience}_{i,t}$ measures the work experience of a given examiner in year t and is calculated as the natural logarithm of the number of years worked in the patent office for examiner i. $\textit{WorkLoad}_{i,t}$ measures the work load of examiner i in year t and is calculated as the natural logarithm of the number of patent applications reviewed by that examiner. $\textit{MaleExaminer}_i$: is a dummy variable that equals to one if the gender of examiner i is male and zero otherwise. $\textit{Art Unit}_a$ and $\textit{Status Year}_t$ indicate the art unit fixed effect and the status year fixed effect.

Table 3 presents the results of the above regression. First, the negative coefficients of *WorkExperience* in columns (1) and (2) of Table 3 suggest that more experienced patent examiners tend to make fewer false rejection mistakes. However, they tend to make more false acceptance mistakes as suggested by the positive coefficient of *WorkExperience* in columns (3) and (4) of Table 3. These findings are also consistent with the agency problem induced by the compensation structure in the patent office: examiners get higher compensation by accepting more patent applications. Second, busy examiners tend to make more mistakes in both ends, as suggested by the positive coefficient of *WorkLoad* in columns (1), (3), and (4) of Table 3. These findings are consistent with how increasing time constraints faced by patent examiners reduce their screening efficiency. Last, I also find that male examiners tend to make more cases of false rejections and false acceptances.

5.4.4. Robustness test: disagreement between humans and machine algorithms, and early patent expiration

In this subsection, I test whether these granted patents in the out-of-sample test set, which this algorithm would reject, should or should not be granted in the first place as another robustness test. To measure the disagreement between the machine learning predictions and actual screening decisions of patent examiners, I first compute the number of applications granted by actual examiners within each art unit in any given year. Then I rank all patent applications filed within each art unit in a given year based on their predicted generality by the algorithm and hypothetically grant the same number of patent applications as examiners within each art unit in a given year. So far, each patent application has an actual grant decision made by examiners and a hypothetical grant decision made by the machine learning algorithm. Finally, I label a patent to be "falsely accepted" if it is accepted by an actual examiner but rejected by the algorithm.

Section 154 of the U.S. Patent Law (35 U.S. Code §154 (a)) sets forth the term of a utility patent filed on or after June 8, 1995, in the U.S. to be 20 years from the earliest filing date of the application on which the patent was granted. Section 41 of the U.S. Patent Law (35 U.S. Code §41 (b) & (c)) states that maintenance fees are required to be paid in every certain period in order to maintain utility patents in force.[29] If these "falsely accepted" patents should not be granted in the first place, we would expect that these patents are more likely to get expired early as a result of delaying and defaulting in payment of maintenance fees. In particular, I test whether these "falsely accepted" patents are properly maintained with the following regression.

$$y_i = \alpha + \beta \mathit{FalseAccept}_i + \mathit{ArtUnit}_a + \mathit{IssueYear}_t + \mathit{Small\&MicroEntity}_s + \mathit{USPC}_j + \epsilon_i, \tag{2}$$

where i indexes patent; a indexes art unit; t indexes the issue year of a patent; s indexes the size of a patentee; and j indexes the USPC class. y represents the patent-maintenance-related dummies indicating the following four aspects: payment of maintenance fee in the 4th year, payment of maintenance fee in the 8th year, maintenance fee reminder mailed, patents expired for failure to pay maintenance fees. $\mathit{FalseAccept}_i$ is a dummy variable, equaling to one if a patent is accepted by actual examiners but would be rejected by the algorithm. $\mathit{ArtUnit}_a$, $\mathit{IssueYear}_t$, $\mathit{SmallEntity}_s$, and USPC_j represent art unit fixed effects, issue year fixed effects, small entity dummies, and USPC class fixed effects.[30]

Table 4 presents the results of regressing Equation (2). The negative coefficients of *FalseAccept* in columns (1) and (2) of Table 4 suggest that "falsely accepted" patents are less likely to be

[29] A patentee needs to pay maintenance fees before the 4th year, the 8th year, and the 12 years to keep its patent in force.

[30] A patentee only needs to pay 1/2 or 1/4 of maintenance fees paid by a large entity if it is a small entity or a micro entity.

maintained four years and eight years after being granted. The positive coefficient of *FalseAccept* in column (3) of Table 4 suggests that patentees who own these "falsely accepted" patents are more likely to receive maintenance fee reminders. Further, the positive coefficient of *FalseAccept* in column (4) of Table 4 indicates that these "falsely accepted" patents are more likely to expire for the failure of paying maintenance fees. These results collectively show that these "falsely accepted" patents turn out to be not very useful to their holders.

6. Innovation screening and firm performance

This section extends my empirical analysis to study the (potential) economic consequences of the current patent screenings on firm performance. First, I describe firm data as well as an *ex-ante* measure of innovation screening efficiency of patent examiners in Subsection 6.1. Second, I present empirical findings on the relationship between innovation screening efficiency and stock market returns of public firms in Subsection 6.2. Third, I discuss empirical results on the effect of innovation screening on the subsequent operating performance of public firms in Subsection 6.3. Lastly, I also examine the effect of innovation screening on subsequent exits of private firms in Subsection 6.4.

6.1. Firm data and sample selection

I use all patent applications that have been filed since 2010 with their screening results available by 2018 in my analysis. In addition to the data on patent applications and patent examiners that I have used in the previous section, I have also collected data on patent assignees from the USPTO website, accounting and financial data for public firms from Compustat and CRSP, firm characteristics, and VC financing for private firms from VentureXpert. I match each dataset with firm names standardized by the NBER Patent Data Name Standardization Routine.[31] By construction, both public and private firms analyzed in this section should have at least one patent application filed since 2010.

[31]The name standardization routine comes from the NBER Patent Data Project: https://sites.google.com/site/patentdataproject.

6.1.1. *Measure innovation screening efficiency*

I construct an *ex-ante* measure of innovation screening efficiency based on the disagreement between the machine learning prediction and the actual screening decision of patent examiners. Based on the "falsely accepted" label I have assigned to each patent as described in Section 5.4.4, I compute the false acceptance rate of each examiner using all patent applications he/she has examined prior to any newly filed patent application. Specifically, I calculate the false acceptance rate of examiner e in art unit a who reviews patent application p at date t as follows:

$$\textit{ExaminerFalseAcceptRate}_{p,e,t,a} = \frac{\textit{\#FalseAccept}_{e,t,a}}{\textit{\#Reviewed}_{e,t,a}}, \tag{3}$$

where $\textit{\#Reviewed}_{e,t,a}$ and $\textit{\#FalseAccept}_{e,t,a}$ are the numbers of patents reviewed and falsely accepted by examiner j prior to date t, respectively.[32] A simple plot in Figure 8 shows that the false acceptance rate of patent examiners has been increasing since 2010, which is consistent with my findings in the previous section that patent examiners are less able to screen in high-quality patents over time.

To match the time horizon of financial and accounting data on firm performance, I further measure the patent screening of examiners associated with each firm in each quarter by averaging false acceptance rates of examiners who have examined that firm's patent applications in the past three years (i.e., a three-year rolling window).[33] For example, the false acceptance rate of firm i in quarter q is calculated as follows:

$$\textit{AvgExaminerFalseAcceptRate}_{i,q} = \frac{1}{N}\sum_{a=1}^{N}\left(\sum_{t=q-13}^{q-1}\textit{ExaminerFalseAcceptRate}_{p,e,t,a}\right), \tag{4}$$

where $\textit{ExaminerFalseAcceptRate}_{p,e,t,a}$ is the false acceptance rate of examiner e who reviews firm i's patent application p in the past three years, and N is the total number of patent applications filed by firm i with screening results available in the past three years.

By construction, the false acceptance rate of an individual examiner is *ex-ante* available for any

[32] I exclude the patent application p in both the numerator and the denominator. I also exclude firms whose patent application is assigned to a patent examiner who has reviewed less than 10 patent applications prior to the patent application p. All results in this section are robust to removing the above exclusions.

[33] I have used different time windows to measure firm-level innovation screening efficiency (i.e., a 1-quarter, 1-year, and 2-year window), and all my empirical results in this section remain qualitatively similar.

newly filed patent application in my sample. More importantly, it is also unlikely to be correlated with firm characteristics due to the quasi-random assignment of patent applications to patent examiners within each art unit (see, e.g., Maestas, Mullen, and Strand, 2013; Farre-Mensa, Hegde, and Ljungqvist, 2020; Sampat and Williams, 2019).

6.1.2. Summary statistics

Table 5 reports summary statistics for my measures of innovation screening efficiency as well as firm characteristics. Panel A of Table 5 presents summary statistics of stock returns for the sample of public firms at the firm-event level. For example, the average and the median false acceptance rates of examiners who screen a firm's patents are 16.6% and 16.0%. I estimate abnormal returns using the market model with CRSP value-weighted index return as the market return, where market model variables (alphas and betas) are estimated over 150 days ending 50 days before the screening decision date of each patent application.[34] The average cumulative abnormal returns over a 3-trading-day window around patent grant news are 3.2 basis points (bps), while the average 1-quarter, 2-quarter, 3-quarter, and 4-quarter buy-and-hold abnormal returns are 0%, -0.6%, -1.8%, and -3.8%, respectively.[35]

Panel B of Table 5 presents the summary statistics of firm performance and firm characteristics for public firms at the firm-quarter level. The average false acceptance rate for public firms is 16.7%; the median number of patent applications being reviewed and granted for public firms in a given three-year window are 15 and 12; the median quarterly *ROA* and *Cash Flow*, which are defined as net income and cash flow divided by total assets, are 0.6% and 1.6%. Public firms, on average, have the logarithm of book assets of 7.2, a leverage ratio of 0.2, the logarithm of the market to book ratio of 1.1, and R&D expenditures of 3.5%.[36] Most of the public firms in my sample are not involved in any patent litigation as defendants after their patents were granted: the average quarterly number of patent litigation for public firms is 0.1.

[34]I also estimate abnormal returns using alternative models such as Fama-French three-factor model, and Carhart four-factor model(see, e.g., Fama and French, 1993; Carhart, 1997). My results remain qualitatively similar using these alternative estimation models.

[35]The negative long-run stock return after patent being granted is somewhat surprising. However, my results are consistent with that in Cao, Jiang, and Ritter (2013) and they show that firms with patent filings (regardless of application outcomes) have negative profitability on average in the five years after IPOs.

[36]All accounting variables (i.e., *ROA*, *Cash Flow*, *R&D Expenditures*, *Leverage*, *Ln(M/B)*) are winsorized at 0.1% and 99.9%. All regression results are qualitatively similar before winsorizing and are robust to different winsorizing thresholds.

Panel C of Table 5 presents the summary statistics of firm performance and firm characteristics for private firms at the firm-quarter level. The average false acceptance rate for private firms is 16.6%; the median number of patent applications being reviewed and granted for private firms in a given three-year window are 4 and 3, which is much lower compared to those for public firms. In terms of firm characteristics, the average age of private firms is 10.6; private firms have the logarithm of quarterly VC financing amount of 0.2, and the quarter number of VC funds of 0.3. Finally, the average rate of successful exits through IPOs or M&As is 21.2%.

6.2. *Innovation screening and stock market returns of public firms*

In this subsection, I test whether my measure of innovation screening efficiency can explain stock market reactions to patent grant news and predict post-granting long-run stock returns. I measure stock market reactions to patent grant news using the cumulative abnormal return on a firm's equity over a 3-trading-day window (from day -1 to day 1) around the patent grant date (*CAR [-1:1]*) and long-run stock returns using the 1-quarter, 2-quarter, 3-quarter, and 4-quarter buy-and-hold abnormal returns on a firm's equity after the patent grant date (*BHAR [1:63]*, *BHAR [1:125]*, *BHAR [1:188]*, and *BHAR [1:250]*).[37]

I separately regress each stock return measure on the average false acceptance rates of examiners who examine firms' patent applications and report regression results in Table 6. Panel A of Table 6 reports regression results testing the effect of innovation screening efficiency on stock market reactions to patent grant news. The coefficient of the constant term in column (1) shows that the announcement returns (*CAR [-1:1]*) on average are positive and significant, which is consistent with the findings in Kogan et al. (2017). However, the coefficient of *ExaminerFalseAcceptRate* is negative and statistically significant in column (2), suggesting that the false acceptance rates of examiners are able to explain some variation in stock market reactions to those patents they have granted. Economically, a one-standard-deviation increase in *ExaminerFalseAcceptRate* decreases the 3-day announcement return by 2 bps. In other words, if all patent applications were screened by the machine learning algorithm used in this paper (i.e., *ExaminerFalseAcceptRate* decreases from 0.166 to 0), the 3-day announcement return would increase by 4 bps.

Finally, panel B of Table 6 reports regression results testing the effect of innovation screening

[37] I restrict my sample in this subsection to those patents with applications publicly available before they are granted.

efficiency on post-granting long-run stock market returns. Again, the coefficients of *ExaminerFalseAcceptRate* are negative and statistically significant in all four regressions, suggesting that my *ex-ante* measure of innovation screening efficiency negatively predicts firms' long-run stock market returns out of sample, thus can be viewed as *ex-ante* measures of patent quality. Economically, a one-standard-deviation increase in *ExaminerFalseAcceptRate* decreases the following 1-quarter, 2-quarter, 3-quarter, and 4-quarter buy-and-hold abnormal return by 9 bps, 20 pbs, 37 pbs, and 64 bps, respectively. In other words, if all patent applications were screened by the machine learning algorithm (i.e., *ExaminerFalseAcceptRate* decreases from 0.166 to 0), the 1-year buy-and-hold abnormal return would increase by 1.2%.

6.3. *Innovation screening and subsequent operating performance of public firms*

As I have shown, the average quality of patents would be higher if the algorithm granted them. If this is indeed the case, we would expect that firms should have worse performance if examiners with higher false acceptance rates granted their patents. In this subsection, I empirically test the effect of innovation screening efficiency on the subsequent operating performance of public firms with my baseline regression as follows:

$$y_{i,q+n} = \alpha + \beta AvgExaminerFalseAcceptRate_{i,q} + \gamma X_{i,q} + Industry_j + Quarter_q + \epsilon_{i,q}, \tag{5}$$

where i indexes firm; j indexes industry; q indexes quarter; and n equals 1, 4, 8, or 12. *y* is the operating performance of each public firm, which is measured using either *ROA* or *Cash Flow*. For example, $ROA_{i,q+4}$ measures the subsequent 4-quarter (or 1-year) operating performance of each public firm. $AvgExaminerFalseAcceptRate_{i,q}$ is my screening efficiency measure of examiners who have examined firm i's patent applications in the past three years (or twelve quarters) as described in 6.1.1. X is a vector of control variables including the number of patents reviewed and granted in the past three years, firm size in quarter t, leverage in quarter t, market to book ratio in quarter t, and R&D expenditures in quarter t as described in 6.1.2. $Industry_j$ and $Quarter_q$ represent two-digit SIC industry fixed effects and quarter fixed effects. All standard errors in my baseline regressions are double clustered at the firm and quarter level.

The baseline results using *ROA* as the dependent variable are reported in Table 7. Table 7 shows that coefficient of *AvgExaminerFalseAcceptRate* is negative and statistically significant in all regressions, suggesting that public firms whose patent applications are granted by examiners with higher past false acceptance rates perform worse in both the short- and long-term. These results are also economically significant: a one-standard-deviation increase in *AvgExaminerFalseAcceptRate* decreases the following 1-quarter, 1-year, 2-year, and 3-year *ROA* by 32 bps, 83 bps, 157 bps, and 156 bps, respectively. In other words, if all patent applications were screened by the machine learning algorithm (i.e., *AvgExaminerFalseAcceptRate* decreases from 0.167 to 0), ROA would increase by 0.8% and 3.9% over the following 1-quarter and 3-year periods.[38] More importantly, since patent applications are randomly assigned to patent examiners, the effect of the current patent screening system on firm performance is likely to be causal due to the quasi-random assignment of patent applications to patent examiners.

6.3.1. *Potential channels: innovation screening, subsequent R&D expenditures, and subsequent patent litigation*

In this subsection, I test two potential channels behind the effect of innovation screening on firm performance. Specifically, I study the impact of innovation screening on subsequent R&D expenditures and the subsequent number of patent litigation using the same baseline specification as described in Equation (5).

Table 8 presents regression results with the subsequent R&D expenditures as the dependent variable and shows that the coefficient of *AvgExaminerFalseAcceptRate* is negative and statistically significant in all regressions, suggesting that firms lower their R&D expenditures after their patents reviewed by examiners with lower screening efficiency. These results are also economically significant. For example, a one-standard-deviation increase in *AvgExaminerFalseAcceptRate* decreases the following 1-quarter and 3-year R&D expenditures by 4 pbs and 45 pbs (i.e., a 1.1%, and 1.1% decrease compared to the median 1-quarter and 3-year R&D expenditures). All these results suggest that innovation screening has a causal and real effect on the innovation input of public firms that might hurt their short-term and long-term performance.

Table 9 presents regression results with the number of subsequent patent litigation as depen-

[38] Due to space limitation, the baseline results using *Cash Flow* as the dependent variable are reported in Table IA.1 in the Internet Appendix. The results on firm cash flows are both qualitatively and quantitatively similar to the results reported in Table 7.

dent variables and shows that coefficients of *AvgExaminerFalseAcceptRate* are positive and statistically significant in all regressions. Economically, a one-standard-deviation increase in *AvgExaminerFalseAcceptRate* increases the number of patent litigation in the next one quarter and three years by 0.012 and 0.134 (i.e., a 9.3% and 8.6% increase compared to the average number of patent litigation over the same period). These results suggest that firms whose patents are granted by examiners with higher past false acceptance rates are more likely to be involved in subsequent patent litigation, which in turn might harm their short- and long-term performance.[39]

6.3.2. A cross-industry analysis: innovation screening and subsequent operating performance

In this subsection, I empirically test whether the effect of innovation screening on firm performance is larger in innovation-intensive industries with the following specification:

$$\begin{aligned} y_{i,q+n} \quad = \quad & \alpha + \beta_1 \mathit{AvgExaminerFalseAcceptRate}_{i,q} + \beta_2 \mathit{HiTechAndHealth} \\ & + \beta_3 \mathit{AvgExaminerFalseAcceptRate}_{i,q} \times \mathit{HiTechAndHealth} \\ & + \gamma X_{i,q} + \mathit{Industry}_j + \mathit{Quarter}_q + \epsilon_{i,q}, \end{aligned} \tag{6}$$

where *HiTechAndHealth* is a dummy variable, which equals to one if a firm belongs to the High-Tech or Health industry and zero otherwise. The High-Tech and Health industry definition is based on the Fama and French 5 industry groups.[40] I add an industry dummy (*HiTechAndHealth*) and its interaction with $\mathit{AvgExaminerFalseAcceptRate}_{i,q}$ to Equation (5) as described in Equation (6).

Table 10 presents the regression results of the cross-industry analysis. Panels A and B of Table 10 show that β_3 is negative and statistically significant in all regressions. These results suggest that firms in industries that rely more heavily on technological innovation experience a significantly larger impact from the current patent screening system. However, β_3 is not statistically significant in Panel C of Table 10 but is significantly negative in Panel D of Table 10, suggesting that the larger impact experienced by firms in the High-Tech and Health industry is related to subsequent litigation costs.

[39] I have also run the same set of regressions with firm fixed effects and reported the results in Table IA.2 in the Internet Appendix. Most of the results remain statistically significant, suggesting that the effect of innovation screening on firms' outcome exists within each firm and persistent across different time horizons.

[40] For a complete list of four-digt SIC code in each industry provided by Kenneth R. French's data library please see: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/Data_Library/det_5_ind_port.html.

6.4. *Innovation screening and subsequent exits of private firms*

In this subsection, I study the relationship between innovation screening and subsequent exits of private firms with the following specification:

$$y_{i,q+n} = \alpha + \beta \mathit{AvgExaminerFalseAcceptRate}_{i,q} + \gamma Z_{i,q} + \mathit{State}_s + \mathit{Industry}_j + \mathit{Quarter}_q + \epsilon_{i,q}, \quad (7)$$

where y is *SuccessfulExit*, which measures the successful exit of each private firm. For example, $\mathit{SuccessfulExit}_{i,q+4}$ is a dummy variable that equals one if a firm successfully exits either by an IPO or M&A in the following 1-year (4-quarter) period, and zero otherwise; $\mathit{SuccessfulExit}_{i,q+12}$ is a dummy variable that equals one if a firm successfully exits either by an IPO or M&A in the following 3-year (12-quarter) period, and zero otherwise. $\mathit{AvgExaminerFalseAcceptRate}_{i,q}$ is my screening efficiency measure of examiners who have examined firm i's patent applications in the past three years (twelve quarters) as described in 6.1.1. Z is a vector of control variables including the number of patents reviewed and granted in the past three years, firm age in quarter t, total funding received ending in quarter $t-1$, VC funding received in quarter t, and the number of funds invested in quarter t as described in 6.1.2. State_s, $\mathit{Industry}_j$, and $\mathit{Quarter}_q$ represent the state of incorporation fixed effects, two-digit SIC industry fixed effects, and quarter fixed effects. Standard errors are clustered at the state level.

The regression results reported in Table 11 show that the coefficient of *AvgExaminerFalseAcceptRate* is negative and statistically significant in all regressions, suggesting that private firms whose patent applications are granted by examiners with higher past false acceptance rates are less likely to exit successfully either by an IPO or by an M&A in both the short- and long-term. These results are also economically significant: a one-standard-deviation increase in *AvgExaminerFalseAcceptRate* decreases the following 1-quarter, 1-year, 2-year, and 3-year probabilities of exiting successfully by an IPO or an M&A by 15 bps, 72 bps, 139 bps, and 165 bps, respectively. In other words, if all patent applications were screened by the machine learning algorithm (i.e., *AvgExaminerFalseAcceptRate* decreases from 0.166 to 0), the probability of exiting successfully by an IPO or M&A increases by 3.6% over the following three-year period. More importantly, these results suggest that weak innovation screenings in the current patent screening system causally reduce the probability of subsequent exits by IPOs or M&As for private firms due to the quasi-random assignment of patent applications to patent examiners.[41]

7. Conclusion

In this paper, I examine whether the patent screening process can be improved under the current patent system in terms of granting better quality patents. I argue that examiners may not process relevant information efficiently to screen out low-quality applications due to their increasing time constraints and their own incentives. However, machine learning algorithms have much larger capacities to process information efficiently and potentially reduce human biases. Using utility patent applications filed at the USPTO from 2001 to 2018, I train a machine learning algorithm using earlier patent applications and predict the quality of more recent patent applications out of sample. I show that the current patent system screens in many low-quality patents, which can be mitigated with the help of a machine learning algorithm. To compare the performance between humans and machine learning algorithms, I make use of the quasi-random assignment of patent applications to examiners who have different levels of leniency. I find that the improvement in quality is substantial and significant: training an algorithm targeting the generality of patents results in a 15.5% gain of patent generality and a 35.6% gain of the number of patent citations. Further, regression analyses show that these patents, which an algorithm would reject, are more likely to expire early, suggesting that these "falsely accepted" patents indeed turn out to be useless to their holders, or in other words, an winner's curse for their owners.

To examine the economic consequences of current patent screening, I study the impact of innovation screening efficiency on the future performance of firms who have at least one patent application filed at the USPTO since 2010. To do so, I construct an *ex-ante* efficiency measure of innovation screening by computing the false acceptance rate of examiners who examine firms' patent applications. I find that my measure of innovation screening efficiency is able to predict

[41] To make sure my empirical results are not primarily driven by the art-unit level of screening efficiency, I have also constructed a measure of art-unit adjusted innovation screening efficiency and rerun all the regressions in Sections 6.3 and 6.4 as a robustness test. Due to space limitation, the results using this alternative measure are reported in Tables IA.3 and IA.4 in the Internet Appendix and are consistent with my findings reported in Sections 6.3 and 6.4.

both the announcement return around patent grant news and the subsequent long-run stock return and thereby can be viewed as an *ex-ante* measure of patent quality. Next, I find that public firms whose patent applications are accepted by examiners with higher false acceptance rates are likely to have lower operating performance (measured by ROA and Cash Flow) and lower their R&D expenditures; also more likely to be involved in more patent litigation in both the short-term and long-term future. Such a negative impact is larger for firms in the High-Tech and Health industries. Lastly, I find that private firms whose patent applications are accepted by examiners with higher false acceptance rates are less likely to exit successfully by an IPO or an M&A in the short-term and long-term future. The above results are also economically significant. For example, the 3-year ROA for public firms increases by 3.9 percentage points, and the 3-year probability of exiting successfully by an IPO or an M&A for private firms increases by 3.6 percentage points if the machine learning algorithm screened all patent applications. More importantly, these findings can be interpreted as causal evidence for the economic consequences of current patent screening since patent applications are randomly assigned to patent examiners that are unlikely to be correlated with firm characteristics.

Overall, this study shows how new technologies such as machine learning algorithms can help improve human decisions and thereby generating policy implications for USPTO policymakers. Based on findings in this paper, machine learning algorithms could potentially serve as a supporting tool in assisting human examiners to make better decisions. For example, human examiners may use a machine learning algorithm as a tool to double-check those questionable screening decisions identified by the machine learning algorithm. While human examiners may or may not change their decisions after reexaminations of those patent applications, such a reexamination process may potentially reduce human bias from their behavioral issues or the increasing time constraint faced by them.

Although this study finds that machine learning algorithms can potentially make better screening decisions in terms of granting higher-quality patents, replacing human examiners with machine learning algorithms may incur unintended consequences: i.e., inventors may strategically file patent applications to respond to such replacements. Therefore, this study proposes that such a machine learning algorithm can serve at most as an auditing tool, which presumably could be implemented at a relatively low cost. Combining the expertise of human examiners and the strength of machine learning mitigates such unintended consequences while achieving better screening outcomes.

References

Abis, S., 2017. Man vs. machine: Quantitative and discretionary equity management. Unpublished working paper. Columbia Business School .

Arrow, K. J., 1972. Economic welfare and the allocation of resources for invention. In: *Readings in Industrial Economics*, Springer, pp. 219–236.

Athey, S., Imbens, G. W., 2017. The state of applied econometrics: Causality and policy evaluation. Journal of Economic Perspectives 31, 3–32.

Balsmeier, B., Assaf, M., Chesebro, T., Fierro, G., Johnson, K., Johnson, S., Li, G.-C., Lück, S., O'Reagan, D., Yeh, B., et al., 2018. Machine learning and natural language processing on the patent corpus: Data, tools, and new measures. Journal of Economics & Management Strategy 27, 535–553.

Bessen, J., Maskin, E., 2009. Sequential innovation, patents, and imitation. The RAND Journal of Economics 40, 611–635.

Bessen, J. E., Meurer, M. J., 2006. Patent litigation with endogenous disputes. American Economic Review 96, 77–81.

Bloom, N., Schankerman, M., Van Reenen, J., 2013. Identifying technology spillovers and product market rivalry. Econometrica 81, 1347–1393.

Bowen III, D. E., Frésard, L., Hoberg, G., 2020. Rapidly evolving technologies and startup exits. Swiss Finance Institute Research Paper .

Cao, J., Jiang, F., Ritter, J. R., 2013. Patent and Innovation-Driven Performance in Venture Capital-Backed IPOs. Available at SSRN 2364668.

Carhart, M. M., 1997. On persistence in mutual fund performance. The Journal of Finance 52, 57–82.

Chemmanur, T. J., Gupta, M., Simonyan, K., 2017. Management quality and innovation in private firms and the IPO market rewards to innovative activity .

Chemmanur, T. J., Loutskina, E., Tian, X., 2014. Corporate venture capital, value creation, and innovation. The Review of Financial Studies 27, 2434–2473.

Chen, T., Guestrin, C., 2016. Xgboost: A scalable tree boosting system. In: *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, ACM, pp. 785–794.

Choi, J. P., 1998. Patent Litigation as an Information-Transmission Mechanism. The American Economic Review 88, 1249–1263.

Cohen, L., Diether, K., Malloy, C., 2013. Misvaluing innovation. The Review of Financial Studies 26, 635–666.

Cornelli, F., Schankerman, M., 1999. Patent renewals and R&D incentives. The RAND Journal of Economics 30, 197–213.

Dreyfuss, R. C., 2008. Nonobviousness: a comment on three learned papers. Lewis & Clark Law Review 12, 431.

Duffy, J. F., 2008. A Timing Approach to Patentability. Lewis & Clark Law Review 12, 343.

Eberhart, A. C., Maxwell, W. F., Siddique, A. R., 2004. An examination of long-term abnormal stock returns and operating performance following R&D increases. The Journal of Finance 59, 623–650.

Eisenberg, R. S., 2008. Pharma's nonobvious problem. Lewis & Clark Law Review 12, 375.

Erel, I., Stern, L. H., Tan, C., Weisbach, M. S., 2018. Selecting Directors Using Machine Learning. Tech. rep., National Bureau of Economic Research.

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

Farre-Mensa, J., Hegde, D., Ljungqvist, A., 2020. What is a patent worth? Evidence from the US patent "lottery". The Journal of Finance 75, 639–682.

Feng, J., Jaravel, X., 2020. Crafting intellectual property rights: Implications for patent assertion entities, litigation, and innovation. American Economic Journal: Applied Economics 12, 140–81.

Fitzgerald, T., Balsmeier, B., Fleming, L., Manso, G., 2021. Innovation search strategy and predictable returns. Management Science 67, 1109–1137.

Frakes, M. D., Wasserman, M. F., 2015. Does the us patent and trademark office grant too many bad patents: evidence from a quasi-experiment. Stanford Law Review 67, 613.

Frakes, M. D., Wasserman, M. F., 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.

Friedman, J. H., 2001. Greedy function approximation: A gradient boosting machine. The Annals of Statistics 29, 1189 – 1232.

Gilbert, R., Shapiro, C., 1990. Optimal patent length and breadth. The RAND Journal of Economics 21, 106–112.

Graham, S. J., Marco, A. C., Miller, R., 2018. The USPTO patent examination research dataset: A window on patent processing. Journal of Economics & Management Strategy 27, 554–578.

Gu, S., Kelly, B., Xiu, D., 2020. Empirical asset pricing via machine learning. The Review of Financial Studies 33, 2223–2273.

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

Hall, B. H., Jaffe, A. B., Trajtenberg, M., 2001. The NBER patent citation data file: Lessons, insights and methodological tools. Tech. rep., National Bureau of Economic Research.

Heller, M. A., Eisenberg, R. S., 1998. Can patents deter innovation? The anticommons in biomedical research. Science 280, 698–701.

Hirschman, A. O., 1980. National power and the structure of foreign trade, vol. 105. University of California Press.

Hirshleifer, D., Hsu, P.-H., Li, D., 2013. Innovative efficiency and stock returns. Journal of Financial Economics 107, 632–654.

Jaffe, A. B., Lerner, J., 2011. Innovation and its discontents: How our broken patent system is endangering innovation and progress, and what to do about it. Princeton University Press.

Kelly, B., Papanikolaou, D., Seru, A., Taddy, M., 2018. Measuring technological innovation over the long run. Tech. rep., National Bureau of Economic Research.

Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., Mullainathan, S., 2017. Human decisions and machine predictions. The Quarterly Journal of Economics 133, 237–293.

Kleinberg, J., Ludwig, J., Mullainathan, S., Obermeyer, Z., 2015. Prediction policy problems. American Economic Review 105, 491–95.

Kline, P., Petkova, N., Williams, H., Zidar, O., 2019. Who profits from patents? rent-sharing at innovative firms. The Quarterly Journal of Economics 134, 1343–1404.

Kogan, L., Papanikolaou, D., Seru, A., Stoffman, N., 2017. Technological innovation, resource allocation, and growth. The Quarterly Journal of Economics 132, 665–712.

Lanjouw, J. O., Schankerman, M., 2001. Characteristics of patent litigation: A window on competition. The RAND Journal of Economics 32, 129–151.

Lemley, M. A., Shapiro, C., 2005. Probabilistic patents. Journal of Economic Perspectives 19, 75–98.

Lu, Q., Myers, A., Beliveau, S., 2017. USPTO patent prosecution research data: Unlocking office action traits. Unpublished, USPTO Economic Working Papter .

Maestas, N., Mullen, K. J., Strand, A., 2013. Does disability insurance receipt discourage work? Using examiner assignment to estimate causal effects of SSDI receipt. American Economic Review 103, 1797–1829.

Mandel, G. N., 2008. Another missed opportunity: the Supreme Court's failure to define nonobviousness or combat hindsight bias in KSR v. Teleflex. Lewis & Clark Law Review 12, 323.

Mansfield, E., 1986. Patents and innovation: an empirical study. Management Science 32, 173–181.

Marco, A., Carley, M., Jackson, S., Myers, A., 2015a. The USPTO historical patent data files: two centuries of invention. Unpublished, USPTO Economic Working Papter .

Marco, A. C., Myers, A., Graham, S. J., D'Agostino, P., Apple, K., 2015b. The USPTO patent assignment dataset: Descriptions and analysis. Unpublished, USPTO Economic Working Papter .

Marco, A. C., Sarnoff, J. D., Charles, A., 2019. Patent claims and patent scope. Research Policy 48, 103790.

Marco, A. C., Toole, A. A., Miller, R., Frumkin, J., 2017. USPTO Patent Prosecution and Examiner Performance Appraisal. Unpublished, USPTO Economic Working Papter .

Matutes, C., Regibeau, P., Rockett, K., 1996. Optimal Patent Design and the Diffusion of Innovations. The RAND Journal of Economics 27, 60–83.

Merges, R. P., 1999. As many as six impossible patent before breakfast: Property rights for business concepts and patent system reform. Berkeley Technology Law Journal 14, 577.

Meurer, M. J., 1989. The settlement of patent litigation. The RAND Journal of Economics 20, 77–91.

Mikolov, T., Chen, K., Corrado, G., Dean, J., 2013a. Efficient estimation of word representations in vector space. arXiv preprint arXiv:1301.3781 .

Mikolov, T., Le, Q. V., Sutskever, I., 2013b. Exploiting similarities among languages for machine translation. arXiv preprint arXiv:1309.4168 .

Mikolov, T., Yih, W.-t., Zweig, G., 2013c. Linguistic regularities in continuous space word representations. In: *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 746–751.

Mullainathan, S., Spiess, J., 2017. Machine learning: an applied econometric approach. Journal of Economic Perspectives 31, 87–106.

Nordhaus, W. D., 1969. An economic theory of technological change. American Economic Review 59, 18–28.

Righi, C., Simcoe, T., 2018. Patent examiner specialization. Research Policy .

Rokach, L., Maimon, O. Z., 2008. Data mining with decision trees: theory and applications, vol. 69. World scientific.

Rossi, A., 2018. Predicting Stock Market Returns with Machine Learning. Unpublished working paper. Georgetown University .

Sampat, B., Williams, H. L., 2019. How do patents affect follow-on innovation? Evidence from the human genome. American Economic Review 109, 203–36.

Schankerman, M. A., Schuett, F., 2016. Screening for patent quality: Examination, fees, and the courts. CentER Discussion Paper Series .

Scherer, F. M., 1972. Nordhaus' theory of optimal patent life: A geometric reinterpretation. American Economic Review 62, 422–427.

Shu, T., Tian, X., Zhan, X., 2019. Patent Quality and Firm Value: Evidence from Patent Examiners Busyness. Available at SSRN 3137894 .

Tong, X., Frame, J. D., 1994. Measuring national technological performance with patent claims data. Research Policy 23, 133–141.

Trajtenberg, M., 1990. A penny for your quotes: Patent citations and the value of innovations. The RAND Journal of Economics 21, 172–187.

Trajtenberg, M., Henderson, R., Jaffe, A., 1997. University versus corporate patents: A window on the basicness of invention. Economics of Innovation and New Technology 5, 19–50.

Zucker, L. G., Darby, M. R., Armstrong, J. S., 2002. Commercializing knowledge: University science, knowledge capture, and firm performance in biotechnology. Management Science 48, 138–153.

Figure 1: The number of patent applications and patent examiners at the USPTO from 2001 to 2018

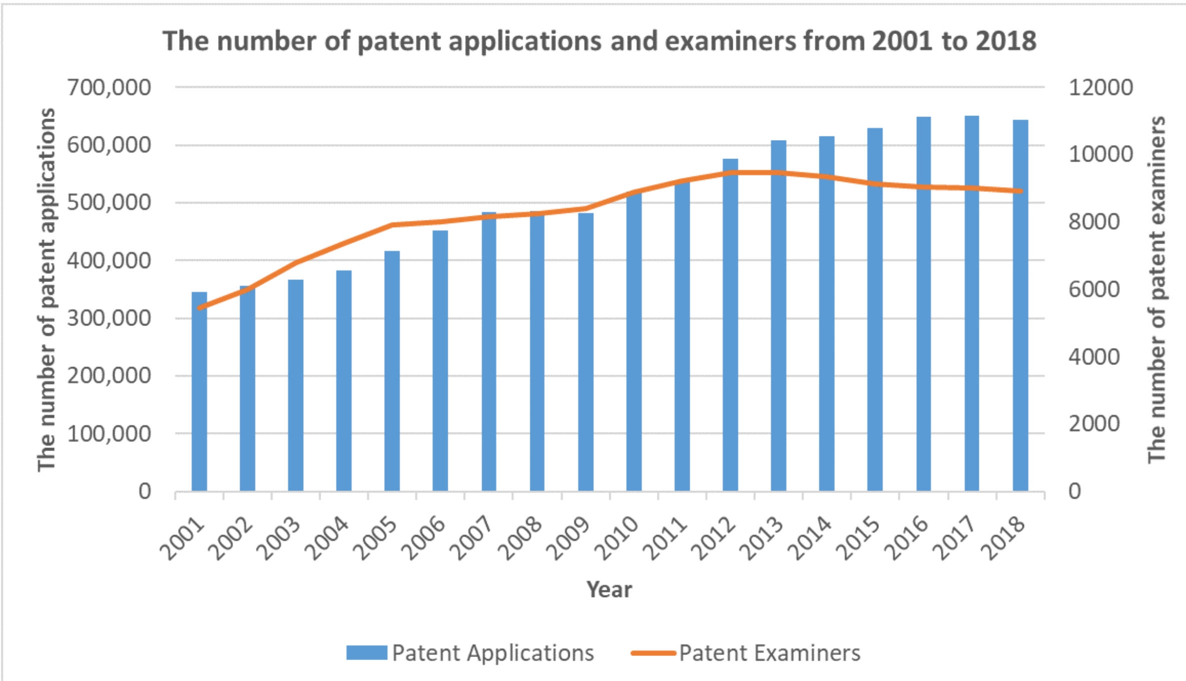

This figure shows the number of patent applications and patent examiners at the USPTO from 2001 to 2018. Each blue bin represents the number of patent applications and the yellow line represents the number of patent examiners. Data source: Patent Statistics Chart and Patent Examination Data from the USPTO website.

Figure 2: An illustrative example of using examiner leniency to evaluate the screening performance of a machine learning algorithm

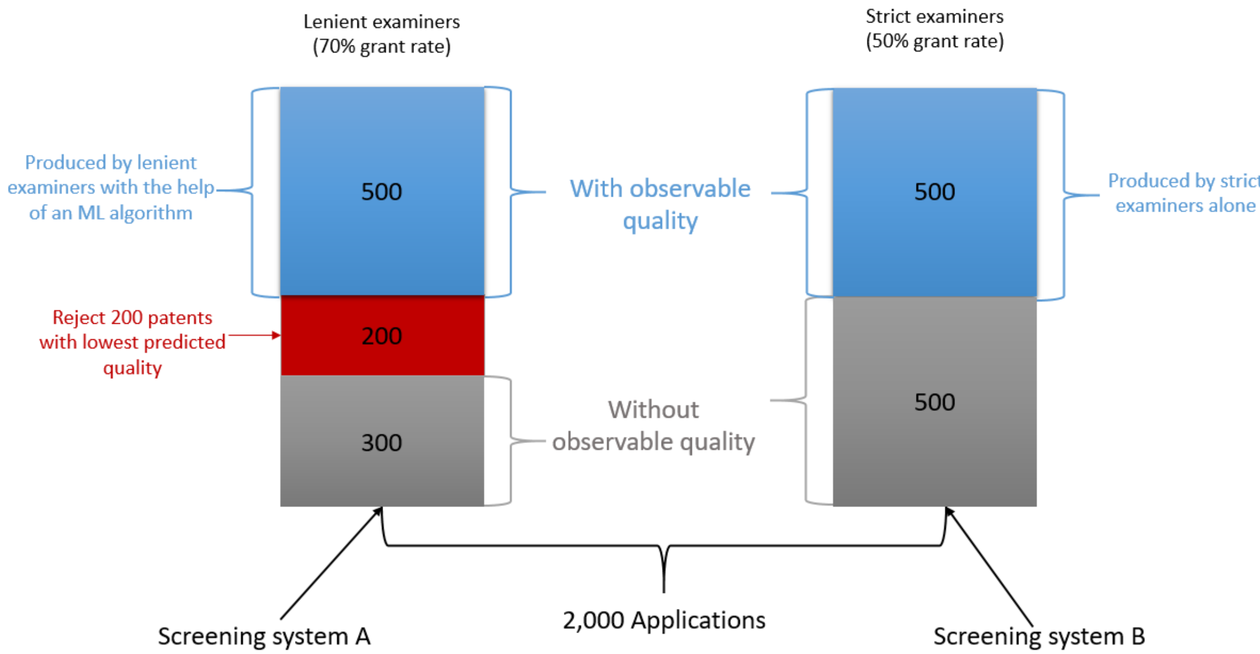

This figure provides an illustrative example of using examiner leniency to compare the performance of actual examiners and a machine learning algorithm.

Figure 3: Training and testing data used for my machine learning prediction

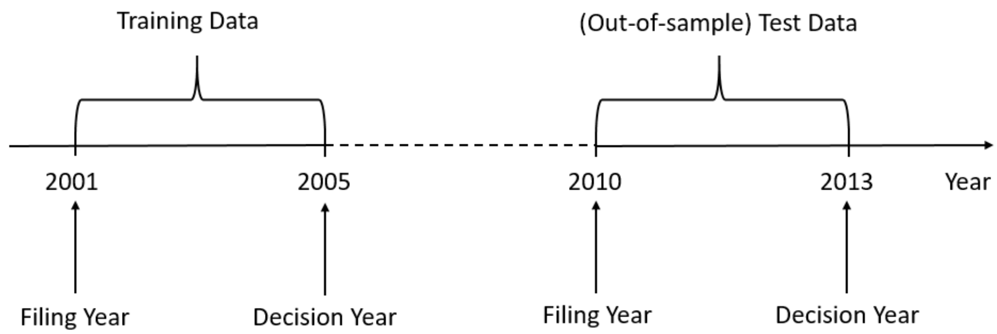

The figure shows the partition for the training and test data used for my machine learning prediction. I select applications filed from 2001 to 2005 with screening status available before the beginning of 2006 into the training set, and applications filed from 2010 to 2013 with screening status available before the beginning of 2014 into the test set. The training set is used to form the algorithm for my prediction and the test set is used to evaluate all of my results. The final sample used in my machine learning prediction consists of 280,243 patent applications in the training set and 357,101 patent applications in the test set.

Figure 4: The relation between predicted generality and actual generality in the test set

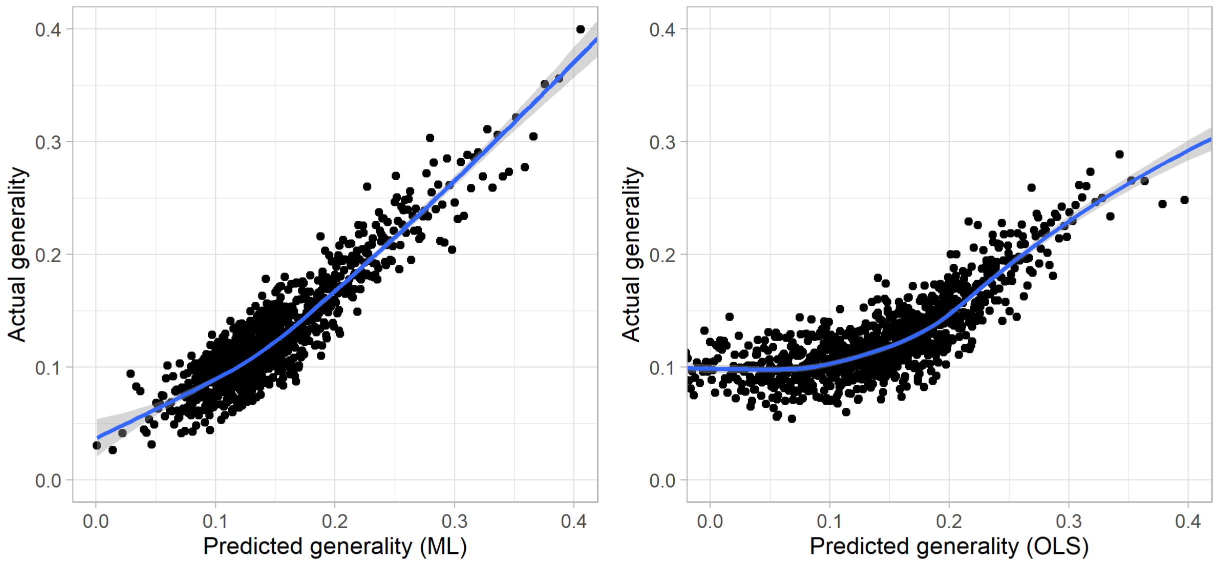

The figure shows the results of the machine learning algorithm (in the left panel) and OLS regressions (in the right panel) built using applications in the training set, applied to applications in the out-of-sample test set. The average predicted generality of patent application in each bin based on the machine learning algorithm and the OLS regression are on the x-axis of the left panel and the right panel. The actual generality is on the y-axis of both panels.

Figure 5: The relation between predicted generality by the machine learning algorithm and actual examiner grant decisions

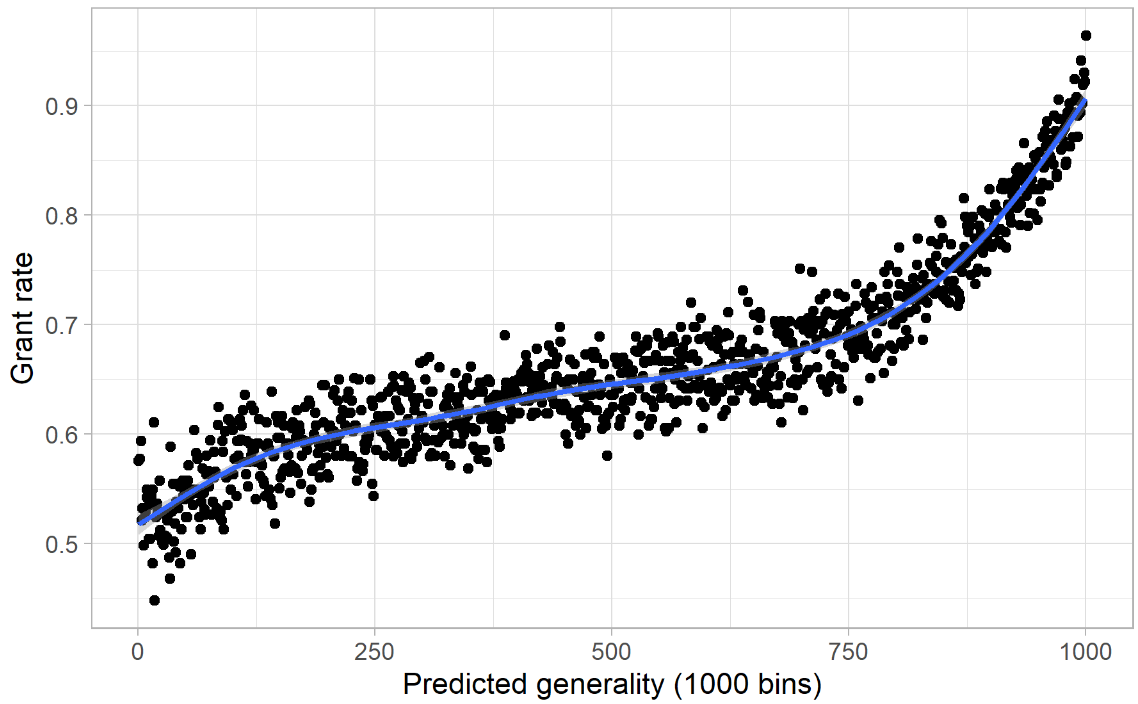

The figure shows the relation between predicted generality by the machine learning algorithm and actual examiner grant decisions. The rank of average predicted generality of all patent application in each pin is on the x-axis. The grant rate is on the y-axis.

Figure 6: Ten important features identified by the machine learning algorithm

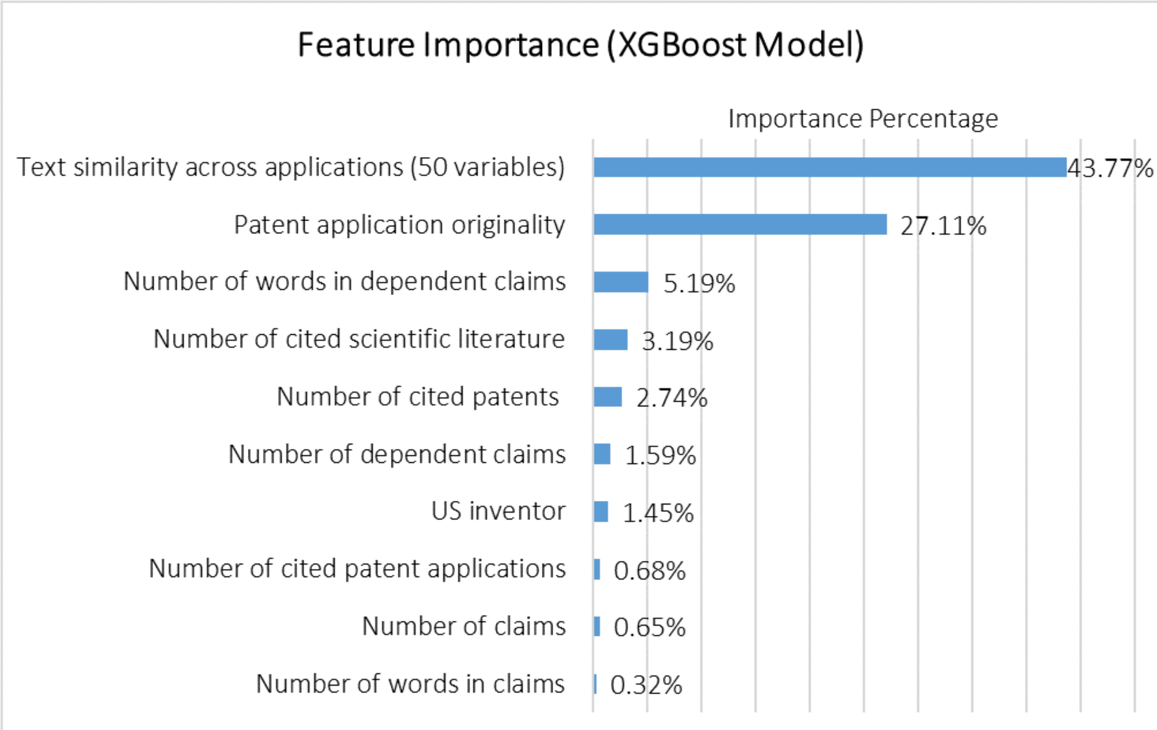

The figure shows ten important features identified by the machine learning algorithm. The predictive power by each feature measured as the percentage of total predictive power is on the x-axis. The name of each of the ten features is on the y-axis.

Figure 7: Comparison between applications rejected by stricter examiners and applications rejected by the algorithm

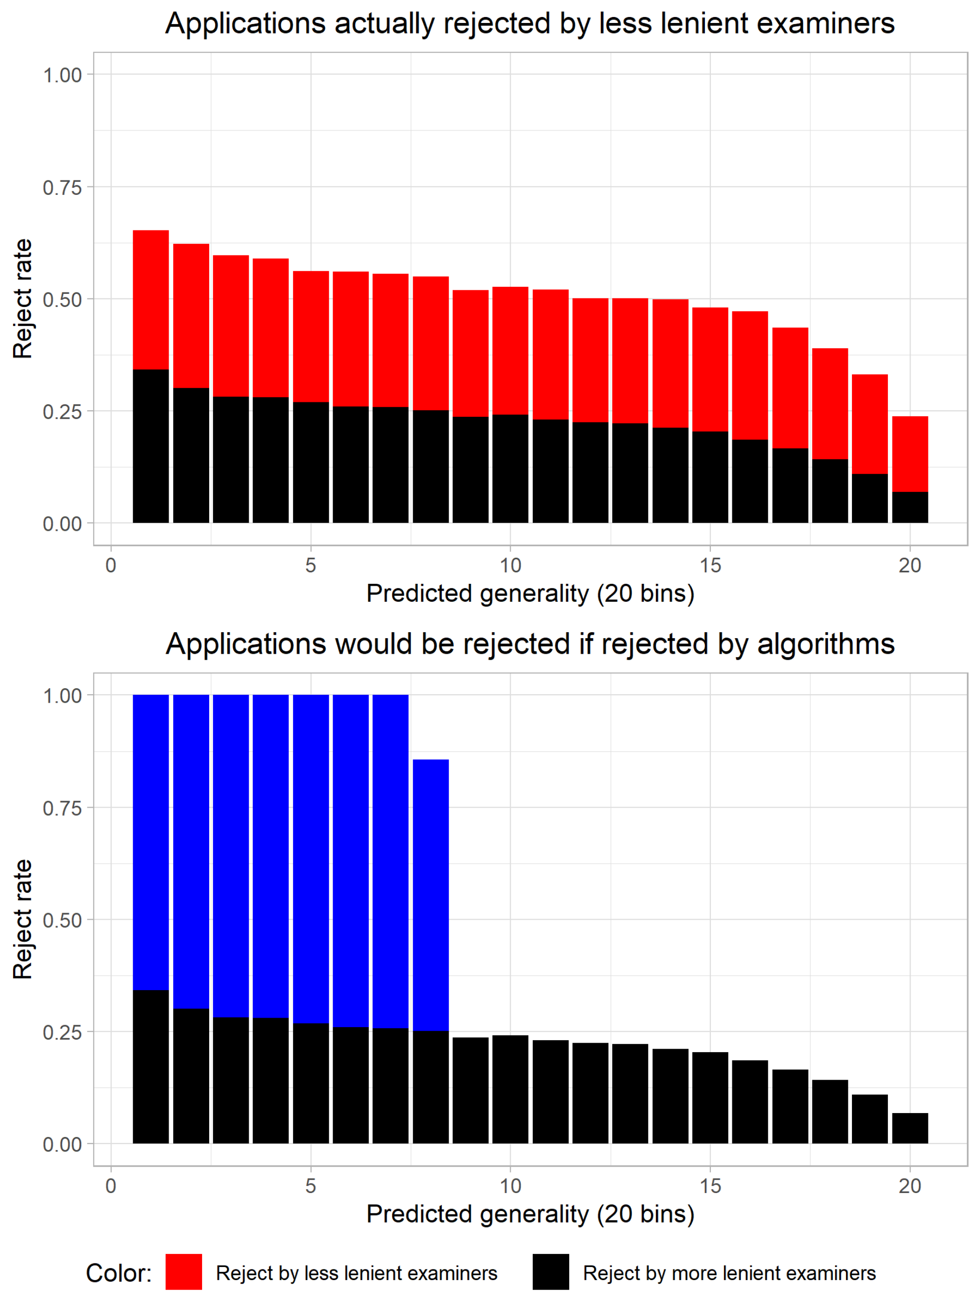

This figure shows comparison between applications rejected by stricter examiners and applications rejected by the algorithm. I divide patent applications in the test set into 20 bins by predicted generality (x-axis). In both panels, the black bar at the bottom of a given bin shows the fraction of patent applications being rejected by more lenient examiners. The red bar in the top panel shows which applications less lenient examiners actually reject. The blue bar in the below panel shows which applications the algorithm would reject to match the grant rate of less lenient examiners.

Figure 8: The screening efficiency of patent examiners since 2010

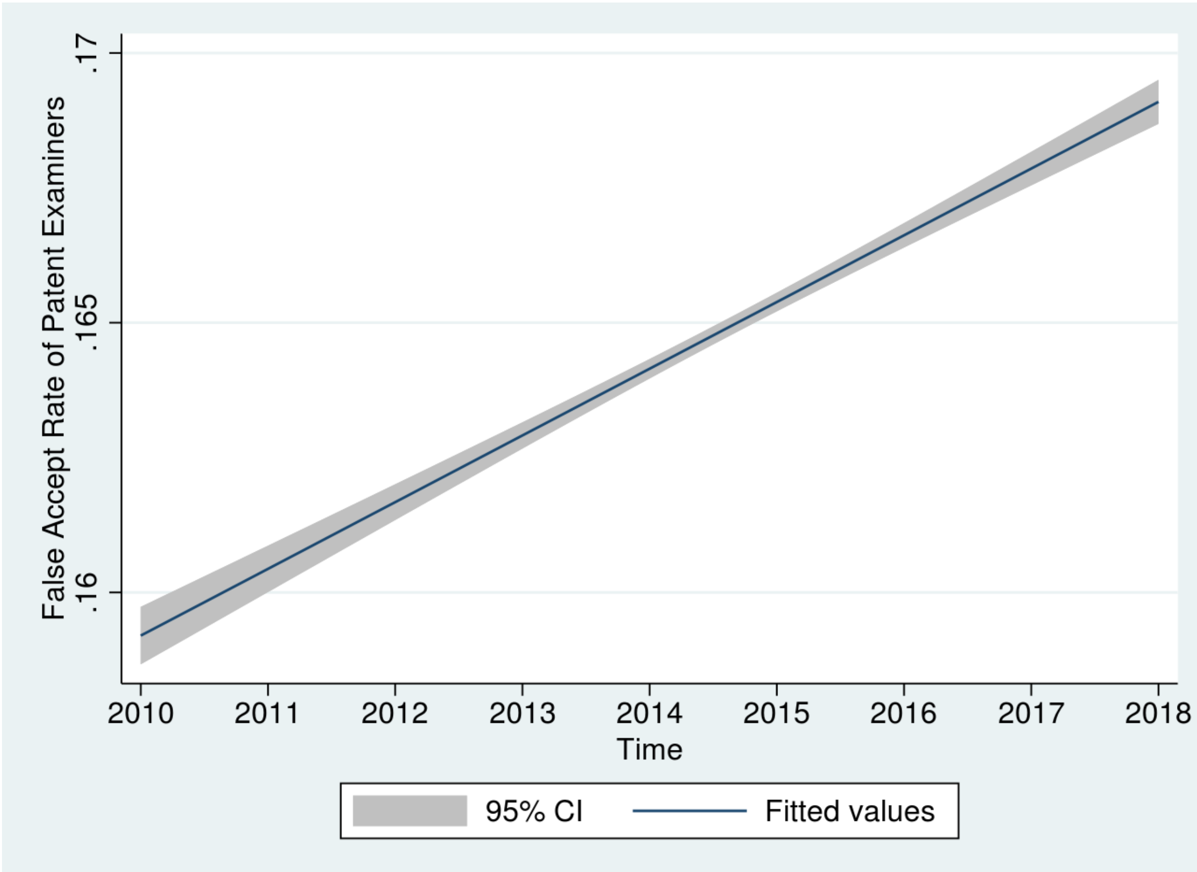

The figure shows the fitted false acceptance rate of patent examiners since 2010. The time variable is on the x-axis. The solid line is the fitted value from regressing the false acceptance rate of individual patent examiners on the time variable; the gray shade represents the 95% confidence interval of the fitted value.

Table 1: Summary statistics (patent applications)

This table shows descriptive statistics for the sample of patent applications from 2001 to 2013 used in my machine learning analysis. *ForwardCitations* counts the number of future citation that each patent has received over a 4-year period after it being granted. *Generality* captures the industry dispersion of 4-year forward citing patents, which equals to one minus Herfindahl-Hirschman index of industries that citing patents belong to. *NumberIndepClaims* and *NumberDepClaims* count the number of independent claims and dependent claims for each patent application. *NumberWordsIndepClaims* and *NumberWordsDepClaims* count the total number of words in independent claims and dependent claims for each patent application. *MinNumberWordsIndepClaims* and *MinNumberWordsDepClaims* count the minimum number of words in independent claims and dependent claims for each patent application. *AvgNumberWordsIndepClaims* and *AvgNumberWordsDepClaims* count the average number of words per independent claim and per dependent claim for each patent application. *NumberCitedForeignPatents* counts the number of foreign patents that each patent application has cited. *NumberCitedForeignPatents* counts the number of novel words that each patent application has. *NumberCitedLiterature* counts the number of scientific literature that each patent application has cited. *NumberCitedApplications* counts the number of patent applications that each patent application has cited. *OriginalityApplication* captures the industry dispersion of backward cited patent applications that each patent application has made, which equals to one minus Herfindahl-Hirschman index of industries that cited patent applications belong to. *NumberCitedPatents* counts the number of patents that each patent application has cited. *OriginalityPatent* captures the industry dispersion of backward cited patents that each patent application has made, which equals to one minus Herfindahl-Hirschman index of industries that cited patents belong to. *USInventorDummy* is a dummy variable indicating whether an investor is from U.S. or not. *SmallEntityDummy* is a dummy variable indicating whether a patent application is from a small entity or not.

Panel A: Patent application quality variables

	N	Mean	Median	p10	p90	S.D.
ForwardCitations	236,643	3.464	2	1	7	6.709
Generality	236,643	0.133	0	0	0.500	0.220

Panel B: Patent application characteristics

	N	Mean	Median	p10	p90	S.D.
NumberIndepClaims	637,344	2.791	2	1	5	2.545
NumberDepClaims	637,344	15.528	14	4	27	13.438
NumberWordsIndepClaims	637,344	361.497	258	85	695	499.292
NumberWordsDepClaims	637,344	601.290	475	135	1,134	879.788
MinNumberWordsIndepClaims	637,344	115.735	92	32	210	130.584
MinNumberWordsDepClaims	637,344	21.837	17	11	30	64.500
AvgNumberWordsIndepClaims	637,344	138.185	114	51.500	235.333	136.399
AvgNumberWordsDepClaims	637,344	42.348	34.125	20.875	64.500	69.755
NumberCitedForeignPatents	637,344	2.905	0	0	7	10.680
NumberNovelWords	637,344	0.309	0	0	1	5.103
NumberCitedLiterature	637,344	3.837	0	0	6	22.254
NumberCitedApplications	637,344	2.755	0	0	5	15.446
OriginalityApplication	637,344	0.155	0	0	0.776	0.309
NumberCitedPatents	637,344	8.511	0	0	18	35.589
OriginalityPatent	637,344	0.166	0	0	0.618	0.259
USInventorDummy	637,344	0.439	0	0	1	0.496
SmallEntityDummy	637,344	0.269	0	0	1	0.444

Table 2: Comparing OLS to machine learning prediction of high-quality patents

This table compare the performance of a machine learning algorithm and an OLS function in terms of identifying high-quality patents in the test set. The first column indicates the top 1%, 5%, 10%, and 25% of the predicted generality distribution and the second column shows the percentage of applications that identified by both ML and OLS as the top 1%, 5%, 10%, and 25% of predicted generality distribution. The third and fourth columns report the actual generality among the applications within each of the predicted generality distribution that are identified either by ML only, or by OLS only. The last column shows the statistical difference between results in the third and fourth columns.

Predicted generality	ML&OLS overlap	Average actual generality for applications identified as high predicted generality by:		
		ML Only	OLS Only	Difference (t-statistic)
Top 1%	20.5%	0.281	0.218	0.063*** (6.82)
Top 5%	33.3%	0.227	0.162	0.065*** (13.33)
Top 10%	38.4%	0.198	0.136	0.062*** (16.47)
Top 25%	42.7%	0.151	0.113	0.038*** (14.95)

Table 3: Relationship between patent examiner characteristics and screening performance

The sample consists of patent examiners in the out-of-sample test set. *# False Rejections* counts the number of false rejections made by a given examiner in each year, where *False Rejection* equals to one if a patent application is rejected by that examiner but accepted by the algorithm and zero otherwise. *False Rejection Rate* measures the false rejection rate by a given examiner in each year. *# False Acceptances* counts the number of false acceptances made by a given examiner in each year, where *False Acceptance* equals to one if a patent is accepted by that examiner but rejected by the algorithm and zero otherwise. *False Acceptance Rate* measures the false acceptance rate by a given examiner in each year. *WorkExperience* measures the work experience of a given examiner in a given year and is calculated as the natural logarithm of the number of years worked in the patent office for that examiner. *WorkLoad* measures the work load of a given patent examiner in a given year and is calculated as the natural logarithm of the number of patent applications reviewed by that examiner. *MaleExaminer* is a dummy variable that equals to one if the gender of a given examiner is male and zero otherwise. Art Unit fixed effects and issue year fixed effects are included in all regressions. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable Dependent Variable	# False Rejections	False Rejection Rate	# False Acceptances	False Acceptance Rate
	(1)	(2)	(3)	(4)
WorkExperience	-0.153***	-0.010***	0.165***	0.005***
	(-21.45)	(-16.46)	(17.89)	(9.26)
WorkLoad	1.267***	-0.047***	2.051***	0.018***
	(34.72)	(-15.65)	(43.71)	(6.92)
MaleExaminer	0.114**	0.006	0.181***	0.001
	(2.33)	(1.46)	(2.88)	(0.23)
Constant	-5.026***	0.481***	-11.044***	-0.020
	(-9.57)	(11.11)	(-16.35)	(-0.54)
Art Unit FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
R^2	0.463	0.134	0.411	0.081
Observations	18003	18003	18003	18003

Table 4: Relationship between weak patent screening and subsequent patent maintenance

The sample consists of granted patents in the out-of-sample test set. *FalseAccept* equals to one if a patent is accepted by an actual examiner but rejected by the algorithm and zero otherwise as described in Section 5.4.4. Small & Micro Entity Dummies, Art Unit fixed effects, issue year fixed effects, and patent USPC class fixed effects are included in all regressions. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable	Payment of Maintenance Fee in the 4th Year	Payment of Maintenance Fee in the 8th Year	Maintenance Fee Reminder Mailed	Patent Expired for Failure to Pay Maintenance Fees
	(1)	(2)	(3)	(4)
FalseAccept	-0.032***	-0.011***	0.026***	0.032***
	(-19.26)	(-7.75)	(13.22)	(18.14)
Small & Micro Entity Dummies	Yes	Yes	Yes	Yes
Art Unit FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Patent USPC Class FE	Yes	Yes	Yes	Yes
R^2	0.092	0.458	0.081	0.056
Observations	235552	235552	235552	235552

Table 5: Summary statistics (firms)

This table shows descriptive statistics for the sample of both public and private firms that have at least one patent application filed since 2010 and with status available before (and including) 2018. Panels A and B show summary statistics for the sample of public firms; Panel C shows summary statistics for the sample of private firms. *ExaminerFalseAcceptRate* is the false acceptance rate of an examiner associated with each patent application, which is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *CAR [-1:1]* is the cumulative abnormal return on a firm's equity over a 3-trading-day window (from day -1 to day 1) around each patent grant date. *BHAR [1:63]*, *BHAR [1:125]*, *BHAR [1:188]*, and *BHAR [1:250]* are the buy-and-hold abnormal returns on a firm's equity over a 63-trading-day, 125-trading-day, 188-trading-day, and 250-trading-day window after each patent grant date. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *ROA* is the ratio of quarterly net income over book assets. *Cash Flow* is the quarterly cash flow over book assets. *R&D Expenditures* are the quarterly R&D expenditures over book assets. *#PatentLitigation* counts the quarterly number of patent litigation that firms act as defendants. *FirmSize* is the natural logarithm of book assets. *Leverage* is the total debt (both current liability and long-term debt) over book assets. *Ln(M/B)* is the natural logarithm of the market to book ratio. *SuccessfulExit* is a dummy, which equals one if a given private firm has exited through an IPO or an M&A by the end of my sample period and zero otherwise. *LnVCFinancingAmount* is the natural logarithm of the quarterly investment amount for each firm. *LnNumberFundInvested* is the natural logarithm of the quarterly number of invested funds for each firm. *TotalFundingToDate* is the natural logarithm of total funding each firm has received prior to a given quarter. *LnFirmAge* is the natural logarithm of firm age, which equals the current year minus the firm founding year plus one. All accounting variables (i.e., *ROA*, *Cash Flow*, *R&D Expenditures*, *Leverage*, *Ln(M/B)*) are winsorized at 0.1% and 99.9%.

Panel A: public firm sample – stock returns (firm-event level)

	N	Mean	Median	p10	p90	S.D.
ExaminerFalseAcceptRate	115,673	0.166	0.160	0.062	0.273	0.091
CAR [-1:1]	115,664	0.032	0.016	-2.540	2.568	2.592
BHAR [1:63]	115,669	-0.004	-0.275	-16.943	17.371	15.408
BHAR [1:125]	115,353	-0.571	-0.301	-28.009	27.763	26.214
BHAR [1:188]	114,183	-1.834	-0.257	-40.514	36.723	38.290
BHAR [1:250]	112,607	-3.862	0.063	-55.678	45.625	52.162

Panel B: public firm sample – operating performance (firm-quarter level)

	N	Mean	Median	p10	p90	S.D.
AvgExaminerFalseAcceptRate	13,416	0.167	0.164	0.105	0.227	0.066
#ApplicatiosReviewe	13,416	160.624	15	2	200	839.933
#PatentsGranted	13,416	127.365	12	1	162	686.466
ROA	13,130	-0.022	0.006	-0.118	0.034	0.102
Cash Flow	12,768	-0.012	0.016	-0.111	0.043	0.101
R&D Expenditures	13,141	0.035	0.020	0	0.087	0.050
#PatentLitigation	13,416	0.131	0	0	0	0.622
FirmSize	13,354	7.196	6.985	4.214	10.560	2.429
Leverage	12,800	0.199	0.165	0	0.459	0.220
Ln(M/B)	12,718	1.130	1.051	0.075	2.255	0.909

Panel C: private firm sample (firm-quarter level)						
	N	Mean	Median	p10	p90	S.D.
SuccessExit	13,496	0.212	0	0	1	0.409
AvgExaminerFalseAcceptRate	13,496	0.166	0.164	0.080	0.248	0.077
#ApplicationsReviewed	13,496	9.772	4	1	20	23.827
#PatentsGranted	13,496	7.833	3	0	16	20.929
LnVCFinancingAmount	13,496	0.209	0	0	0	0.774
LnNumberFundInvested	13,496	0.319	0	0	1	1.203
TotalFundingToDate	13,496	0.461	0	0	2.401	1.382
FirmAge	13,496	10.623	10	6	17	4.683

Table 6: Relationship between screening efficiency of patent examiners and stock market returns of public firms

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *CAR [-1:1]* is the cumulative abnormal return on a firm's equity over a 3-trading-day window (from day -1 to day 1) around each patent grant date. *BHAR [1:63]*, *BHAR [1:125]*, *BHAR [1:188]*, and *BHAR [1:250]* are the buy-and-hold abnormal returns on a firm's equity over a 63-trading-day, 125-trading-day, 188-trading-day, and 250-trading-day window after each patent grant date. *ExaminerFalseAcceptRate* is the false acceptance rate of an examiner associated with each patent application, which is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Panel A: Relationship between screening efficiency of patent examiners and stock market reactions around each patent grant date

Dependent Variable	CAR [-1:1]	CAR [-1:1]
	(1)	(2)
ExaminerFalseAcceptRate		-0.233***
		(-2.78)
Constant	0.032***	0.071***
	(4.19)	(4.45)
Observations	115664	115664

Panel B: Relationship between screening efficiency of patent examiners and long-run stock market return

Dependent Variable	[BHAR [1:63]	BHAR [1:125]	BHAR [1:188]	BHAR [1:250]
	(1)	(2)	(3)	(4)
ExaminerFalseAcceptRates	-0.992**	-2.195***	-4.060***	-7.031***
	(-1.99)	(-2.59)	(-3.27)	(-4.14)
Constant	0.161*	-0.205	-1.158***	-2.691***
	(1.71)	(-1.28)	(-4.91)	(-8.34)
Observations	115669	115353	114184	112609

Table 7: Relationship between screening efficiency of patent examiners and subsequent operating performance of public firms

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *ROA* is the ratio of quarterly net income over book assets. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *FirmSize* is the natural logarithm of book assets. *Leverage* is the total debt (both current liability and long-term debt) over book assets. *Ln(M/B)* is the natural logarithm of the market to book ratio. *R&D Expenditures* are the quarterly R&D expenditures over book assets. All accounting variables (i.e., *ROA*, *Cash Flow*, *R&D Expenditures*, *Leverage*, *Ln(M/B)*) are winsorized at 0.1% and 99.9%. Quarter fixed effects and industry (two-digit SIC code) fixed effects are included in all regressions. All regressions are OLS regressions with standard errors double clustered at the firm and quarter level. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable	Subsequent ROA			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	-0.048***	-0.125***	-0.238***	-0.236***
	(-4.83)	(-3.93)	(-4.35)	(-2.96)
#ApplicationsReviewed	-0.018***	-0.070***	-0.142***	-0.196***
	(-5.01)	(-5.05)	(-4.80)	(-4.08)
#PatentsGranted	0.018***	0.066***	0.136***	0.181***
	(4.86)	(4.94)	(4.72)	(3.91)
FirmSize	0.011***	0.047***	0.090***	0.130***
	(14.26)	(18.30)	(18.14)	(15.41)
Leverage	-0.084***	-0.296***	-0.481***	-0.636***
	(-13.18)	(-14.22)	(-11.66)	(-8.87)
Ln(M/B)	0.013***	0.045***	0.084***	0.128***
	(11.22)	(12.35)	(11.27)	(9.78)
R&D Expenditures	-0.936***	-3.466***	-6.545***	-9.661***
	(-18.58)	(-23.38)	(-20.92)	(-18.97)
Constant	-0.143***	-0.748***	-1.589***	-2.577***
	(-5.15)	(-6.95)	(-8.55)	(-6.74)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.409	0.523	0.543	0.537
Observations	11954	10536	8170	5995

Table 8: Relationship between screening efficiency of patent examiners and subsequent R&D expenditures of public firms

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *R&D Expenditures* are the quarterly R&D expenditures over book assets. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *FirmSize* is the natural logarithm of book assets. *Leverage* is the total debt (both current liability and long-term debt) over book assets. *Ln(M/B)* is the natural logarithm of the market to book ratio. *R&D Expenditures* is the R&D expenditure over the book value total assets. All accounting variables (i.e., *R&D Expenditures*, *Leverage*, *Ln(M/B)*) are winsorized at 0.1% and 99.9%. Quarter fixed effects and industry (two-digit SIC code) fixed effects are included in all regressions. t-statistics are in parentheses. All regressions are OLS regressions with standard errors double clustered at the firm and quarter level. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable	Subsequent R&D Expenditures			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	-0.006*	-0.021**	-0.036*	-0.069**
	(-1.85)	(-2.00)	(-1.77)	(-2.15)
#ApplicationsReviewed	0.000	0.001	0.004	0.001
	(0.38)	(0.26)	(0.49)	(0.06)
#PatentsGranted	0.001	0.004	0.008	0.022*
	(0.56)	(1.11)	(1.05)	(1.76)
FirmSize	-0.002***	-0.011***	-0.025***	-0.042***
	(-7.50)	(-10.82)	(-11.43)	(-10.38)
Leverage	-0.001	-0.013**	-0.062***	-0.128***
	(-0.74)	(-2.10)	(-4.07)	(-5.02)
Ln(M/B)	0.002***	0.006***	0.012***	0.017***
	(3.62)	(4.04)	(3.55)	(2.87)
R&D Expenditures	0.808***	2.961***	5.465***	7.679***
	(25.61)	(27.49)	(23.01)	(17.34)
Constant	0.032***	0.197***	0.482***	0.938***
	(4.85)	(5.98)	(9.10)	(8.09)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.760	0.796	0.784	0.767
Observations	11965	10572	8215	6039

Table 9: Relationship between screening efficiency of patent examiners and the subsequent number of patent litigation of public firms

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *#PatentLitigation* counts the quarterly number of patent litigation that firms act as defendants. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *FirmSize* is the natural logarithm of book assets. *Leverage* is the total debt (both current liability and long-term debt) over book assets. *Ln(M/B)* is the natural logarithm of the market to book ratio. *R&D Expenditures* are the quarterly R&D expenditures over book assets. All accounting variables (i.e., *R&D Expenditures*, *Leverage*, *Ln(M/B)*) are winsorized at 0.1% and 99.9%. Quarter fixed effects and industry (two-digit SIC code) fixed effects are included in all regressions. t-statistics are in parentheses. All regressions are OLS regressions with standard errors double clustered at the firm and quarter level. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable	Subsequent #PatentLitigation			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	0.184*	0.749***	1.331**	2.043**
	(1.70)	(2.63)	(2.49)	(2.21)
#ApplicationsReviewed	0.029**	0.119**	0.260***	0.506***
	(2.02)	(2.57)	(2.73)	(3.25)
#PatentsGranted	0.004	0.017	0.031	-0.030
	(0.34)	(0.42)	(0.36)	(-0.21)
FirmSize	0.069***	0.270***	0.571***	0.924***
	(11.13)	(11.89)	(11.51)	(10.77)
Leverage	-0.284***	-1.229***	-2.699***	-4.457***
	(-5.63)	(-6.35)	(-6.24)	(-5.81)
Ln(M/B)	0.018**	0.066**	0.155**	0.347***
	(2.45)	(2.41)	(2.52)	(3.03)
R&D Expenditures	0.780***	2.868***	5.351***	7.308***
	(8.33)	(9.26)	(7.98)	(6.02)
Constant	-0.244	-0.787	-2.522**	-4.666**
	(-1.03)	(-1.28)	(-2.02)	(-2.53)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.121	0.159	0.179	0.197
Observations	12204	11197	9228	7182

Table 10: Relationship between screening efficiency of patent examiners and subsequent operating performance of public firms (A cross-industry analysis)

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *ROA* is the ratio of quarterly net income over book assets. *R&D Expenditure* is the R&D expenditure over the book value total assets. *#PatentLitigation* counts the quarterly number of patent litigation that firms act as defendants. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *HiTechAndHealth* is a dummy, which equals one if a firm belongs to the High-Tech industry or the Health industry based on Fama and French 5 industry groups. Control variables are defined as in Table 7. Quarter fixed effects and industry (two-digit SIC code) fixed effects are included in all regressions. All regressions are OLS regressions with standard errors double clustered at the firm and quarter level. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Panel A: Relationship between screening efficiency of patent examiners and subsequent ROA

Dependent Variable	Subsequent ROA			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
A: AvgExaminerFalseAcceptRate	-0.015	-0.036	-0.148**	-0.209**
	(-1.53)	(-1.15)	(-2.52)	(-2.40)
B: HiTechAndHealth	0.012***	0.030**	0.027	0.017
	(3.26)	(2.50)	(1.26)	(0.49)
A $\times$ B	-0.058***	-0.155**	-0.161	-0.044
	(-3.11)	(-2.56)	(-1.52)	(-0.28)
Controls	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.409	0.523	0.543	0.537
Observations	11954	10536	8170	5995

Panel B: Relationship between screening efficiency of patent examiners and subsequent R&D expenditures				
Dependent Variable	Subsequent R&D Expenditures			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
A: AvgExaminerFalseAcceptRate	-0.003	-0.017**	-0.036**	-0.073**
	(-1.18)	(-2.28)	(-2.11)	(-2.45)
B: HiTechAndHealth	0.006***	0.028***	0.064***	0.101***
	(4.77)	(6.72)	(7.49)	(6.58)
A $\times$ B	-0.001	0.016	0.054	0.098*
	(-0.20)	(0.91)	(1.51)	(1.67)
Controls	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.762	0.799	0.788	0.772
Observations	11965	10572	8215	6039

Panel C: Relationship between screening efficiency of patent examiners and subsequent patent litigation				
Dependent Variable	Subsequent #PatentLitigation			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
A: AvgExaminerFalseAcceptRate	-0.052	0.210	0.247	-0.076
	(-0.47)	(0.76)	(0.52)	(-0.10)
B: HiTechAndHealth	-0.015	0.060	0.092	-0.002
	(-0.39)	(0.55)	(0.44)	(-0.01)
A $\times$ B	0.467**	1.151**	2.265**	4.255**
	(2.35)	(2.22)	(2.35)	(2.55)
Controls	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.123	0.161	0.180	0.198
Observations	12204	11197	9228	7182

Table 11: Relationship between screening efficiency of patent examiners and subsequent exits of private firms

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *SuccessfulExit* is a dummy, which equals one if a given private firm has exited through an IPO or an M&A by the end of my sample period and zero otherwise. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted applications for each firm in the past three years as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *LnVCFinancingAmount* is the natural logarithm of the quarterly investment amount for each firm. *LnNumberFundInvested* is the natural logarithm of the quarterly number of invested funds for each firm. *TotalFundingToDate* is the natural logarithm of total funding each firm has received prior to a given quarter. *LnFirmAge* is the natural logarithm of firm age, which equals the current year minus the firm founding year plus one. Year fixed effects, industry (two-digit SIC code) fixed effects, and state fixed effects are included in all regressions. t-statistics are in parentheses. All regressions are OLS regressions with standard errors clustered at the state level. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable	Subsequent SuccessfulExit			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	-0.019***	-0.094***	-0.182***	-0.214***
	(-4.04)	(-4.74)	(-3.90)	(-4.75)
#PatentsGranted	0.006***	0.021***	0.038***	0.059***
	(3.23)	(4.55)	(3.46)	(3.42)
#ApplicationsReviewed	-0.003	-0.008	-0.013	-0.032
	(-1.28)	(-1.37)	(-1.04)	(-1.68)
InvestmentAmount	-0.005**	-0.005**	0.007*	0.004
	(-2.58)	(-2.28)	(1.83)	(0.56)
NumberFundInvested	0.005***	0.010***	0.010**	0.017***
	(4.01)	(3.77)	(2.49)	(3.42)
TotalFundingToDate	0.001	0.004***	0.008***	0.009***
	(1.27)	(4.27)	(3.08)	(3.46)
LnFirmAge	0.010***	0.020***	0.022	0.019
	(3.34)	(2.69)	(1.55)	(0.82)
Constant	-0.038***	-0.045	0.077	0.136
	(-6.93)	(-1.20)	(0.75)	(0.98)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
R^2	0.013	0.026	0.041	0.051
Observations	13478	12545	10022	7413

Internet Appendix to "How can Innovation Screening be Improved? A Machine Learning Analysis with Economic Consequences for Firm Performance"

IA.1. The supervised machine learning problem and the algorithm used in this paper

IA.1.1. The supervised machine learning problem

Supervised learning is a machine learning problem of learning a function that maps input variables to an output variable using the training data with both input and output variables available. The goal of supervised learning is to predict well with a new out-of-sample dataset (which we usually called it the test data).

In the context of this paper, I use the training data to construct $\hat{f}(X) = \hat{y}$ from input variables X about patent applications to predict an outcome variable y about the performance of patent applications such that $\hat{f}(X)$ predicts well out of sample. Specifically, I use the training data to train $f(X)$ as follows:

$$\hat{y} = \hat{f}(X) = \arg\min_{f \in \mathcal{F}} L(f(X), y) + R(f(X)), \tag{IA.1}$$

where $L(f(X), y)$ is the training loss function, $\mathcal{F}$ is the set of all possible functions f, and $R(f(X))$ is the regularization term.

The goal of minimizing the training loss function is to increase the in-sample prediction accuracy as much as possible, while adding the regularization term is to avoid in-sample over fitting by penalizing the algorithm for choosing more expressive functions.

IA.1.2. The "Extreme Gradient Boosting" algorithm

The "Extreme Gradient Boosting" algorithm (XGBoost) is an implementation of gradient boosting machines, which is used for the supervised machine learning prediction described above (see, e.g., Chen and Guestrin, 2016; Friedman, 2001). XGBoost is a decision tree ensemble based on tree boosting. A decision tree ensemble consists of a set of decision trees, where each tree i itself is a prediction function $f_i(X)$. Tree boosting is to train the each prediction function $f_i(X)$ using an additive strategy: add one new tree at a time from what we have learned. Specifically, we have

$$\hat{y}_0 = \hat{f}_0(X) = 0 \tag{IA.2}$$
$$\hat{y}_1 = \hat{f}_1(X) = \hat{f}_0(X) + f_1(X) = f_1(X) \tag{IA.3}$$
$$\hat{y}_2 = \hat{f}_2(X) = \hat{f}_1(X) + f_2(X) = f_1(X) + f_2(X) \tag{IA.4}$$
$$\ldots$$
$$\hat{y}_t = \hat{f}_t(X) = \hat{f}_{t-1}(X) + f_t(X) = \sum_{i=1}^{t} f_i(X), \tag{IA.5}$$

and the goal at step t is to find $f_t(X)$ that solves the following minimization problem:

$$\hat{y}_t = \hat{f}_t(X) = \arg\min_{f \in \mathcal{F}} L(f_t(X) + \hat{y}_{t-1}, y) + R(f_t(X)). \tag{IA.6}$$

Here, each prediction function $f_i(X)$ and the corresponding regularization term $R(f_i(X))$ are defined as

$$f_i(X) = \omega_{q(X)}, q : \mathbb{R}^m \rightarrow T, \omega \in \mathbb{R}^T, \tag{IA.7}$$

$$R(f_i(X)) = \gamma T + \frac{1}{2}\lambda \sum_{j=1}^{T} \omega_j^2 \tag{IA.8}$$

where ω are the leaf weights, q is a function mapping each data point to the corresponding leaf index, T is the total number of leaves in the tree, both γ and λ are parameters to weight each of these two complexity measures in order to avoid over-fitting (see Chen and Guestrin, 2016 for a detailed discussion).

IA.2. Additional Figures

IA.2.1. Results using the generality measure based on the USPC classification

Figure IA.1: The relation between predicted USPC-based generality, actual USPC-based generality, and grant rate of actual examiners in the test set

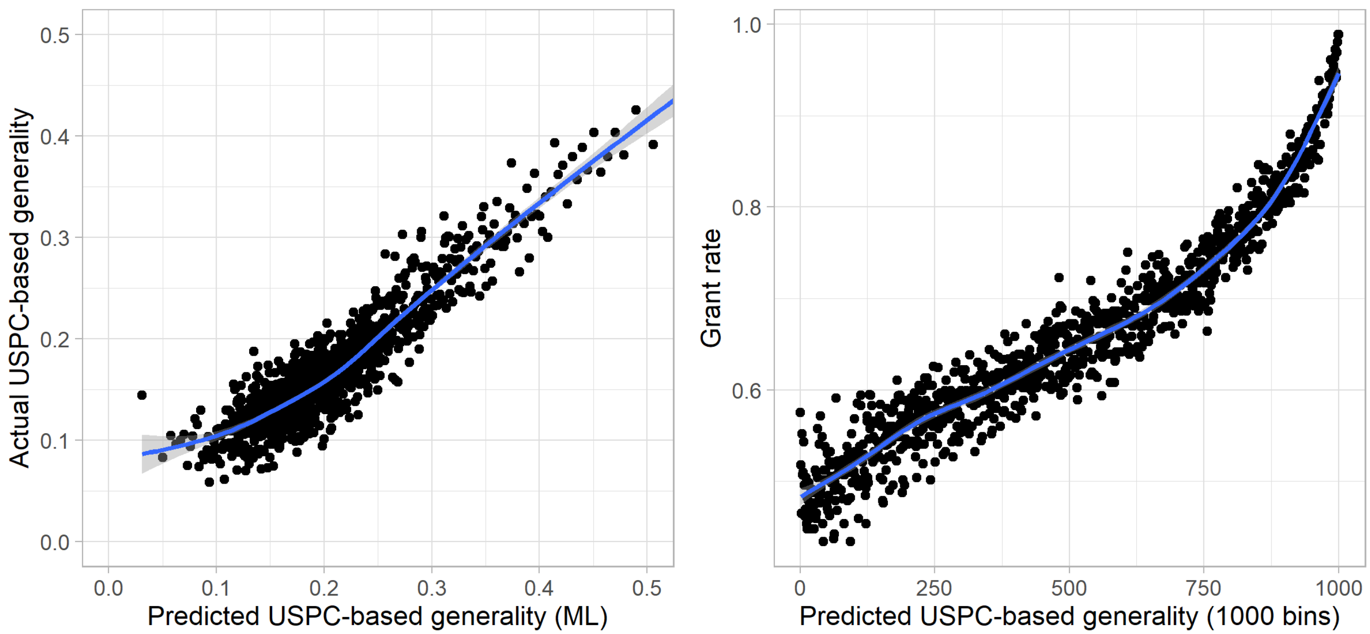

The figure shows the relation between predicted USPC-based generality, actual USPC-based generality, and grant rate of actual examiners in the test set. In the left panel, the average predicted USPC-based generality of patent applications in each bin based on the machine learning algorithm is on the x-axis and the actual USPC-based generality is on the y-axis. In the right panel, the rank of the average predicted USPC-based generality of patent applications in each bin based on the machine learning algorithm is on the x-axis and the grant rate is on the y-axis.

Figure IA.2: Comparison between applications rejected by stricter examiners and applications rejected by the algorithm

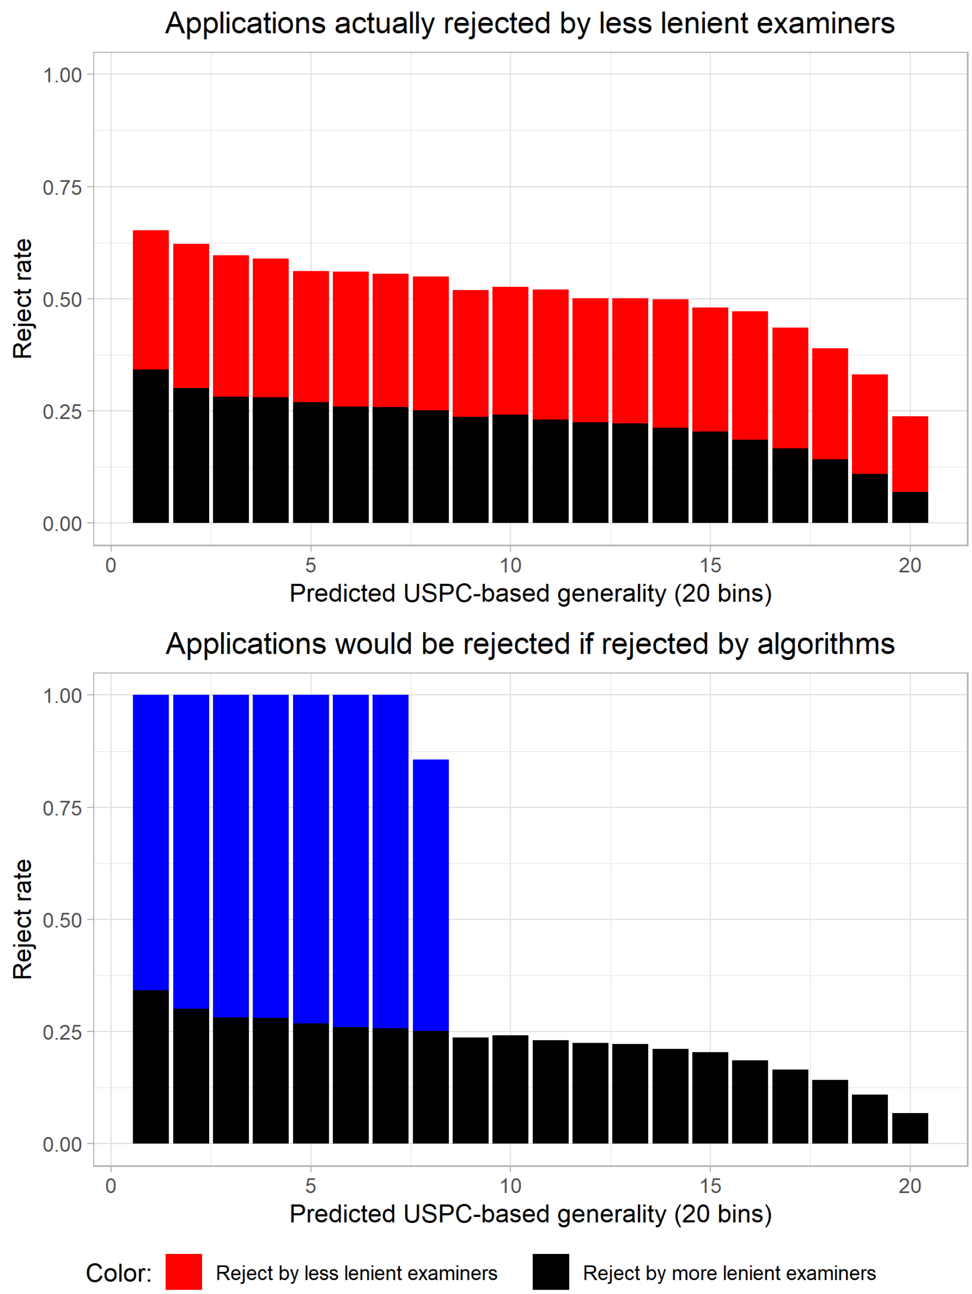

This figure shows comparison between applications rejected by stricter examiners and applications rejected by the algorithm. I divide the sample up equally into 20 bins by predicted USPC-based generality (x-axis). In both panels, the black bar at the bottom of each bin shows the fraction of patent applications rejected by more lenient examiners. The red bar in the top panel shows which applications less lenient examiners actually reject. The blue bar in the below panel shows which applications the algorithm would reject to match the grant rate of less lenient examiners.

IA.2.2. Results using the number of citations

Figure IA.3: The relation between predicted citations, actual citations, and grant rate of actual examiners in the test set

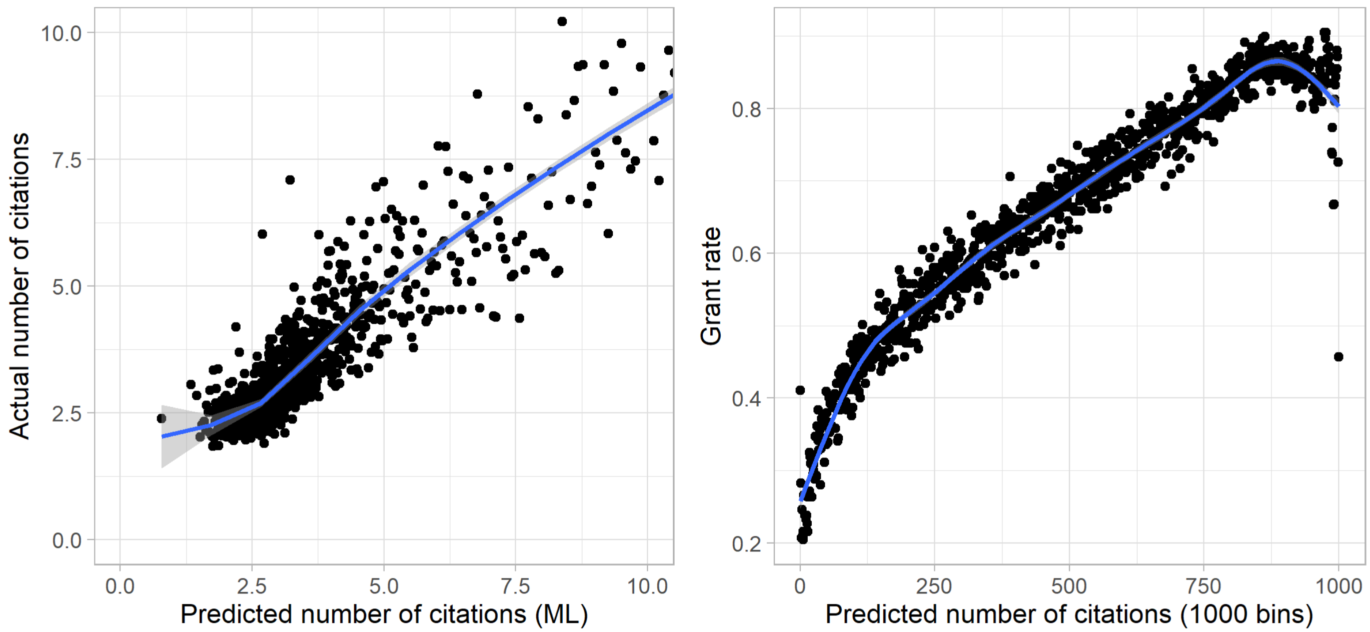

The figure shows the relation between predicted citations, actual citations, and grant rate of actual examiners in the test set. In the left panel, the average predicted number of citations of patent application in each bin based on the machine learning algorithm is on the x-axi and the actual citation is on the y-axis. In the right panel, the rank of the average predicted number of citations of patent application in each bin based on the algorithm is on the x-axis and the grant rate is on the y-axis.

Figure IA.4: Comparison between applications rejected by stricter examiners and applications rejected by the algorithm

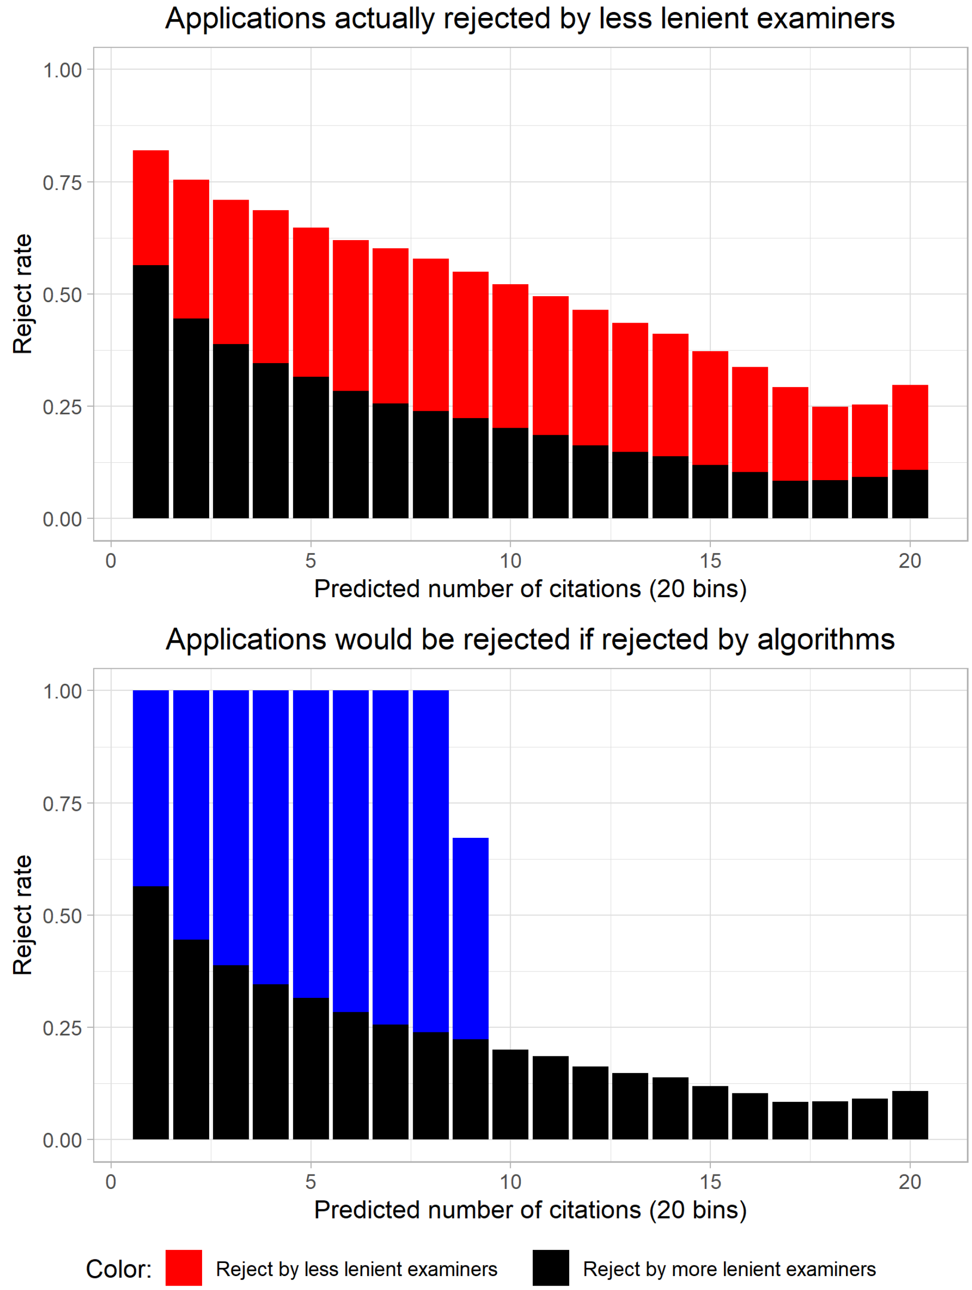

This figure shows comparison between applications rejected by stricter examiners and applications rejected by the algorithm. I divide patent applications in the test set into 20 bins by the predicted number of citations (x-axis). In both panels, the black bar at the bottom of each bin shows the fraction of patent applications rejected by more lenient examiners. The red bar in the top panel shows which applications less lenient examiners actually reject. The blue bar in the below panel shows which applications the algorithm would reject to match the grant rate of less lenient examiners.

IA.3. Additional Tables

Table IA.1: Relationship between screening efficiency of patent examiners and subsequent operating performance of public firms (Cash Flow)

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *Cash Flow* is the quarterly cash flow over book assets. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *FirmSize* is the natural logarithm of book assets. *Leverage* is the total debt (both current liability and long-term debt) over book assets. *Ln(M/B)* is the natural logarithm of the market to book ratio. *R&D Expenditures* are the quarterly R&D expenditures over book assets. All accounting variables (i.e., *Cash Flow*, *R&D Expenditures*, *Leverage*, *Ln(M/B)*) are winsorized at 0.1% and 99.9%. Quarter fixed effects and industry (two-digit SIC code) fixed effects are included in all regressions. All regressions are OLS regressions with standard errors double clustered at the firm and quarter level. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable	Subsequent Cash Flow			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	-0.045***	-0.119***	-0.240***	-0.265***
	(-5.09)	(-3.96)	(-4.35)	(-3.28)
#ApplicationsReviewed	-0.020***	-0.076***	-0.160***	-0.237***
	(-5.44)	(-5.38)	(-5.24)	(-4.71)
#PatentsGranted	0.019***	0.073***	0.155***	0.224***
	(5.33)	(5.36)	(5.24)	(4.63)
FirmSize	0.011***	0.046***	0.088***	0.124***
	(13.89)	(17.39)	(16.96)	(14.07)
Leverage	-0.074***	-0.254***	-0.390***	-0.470***
	(-11.65)	(-11.99)	(-8.95)	(-6.13)
Ln(M/B)	0.011***	0.037***	0.069***	0.101***
	(9.84)	(10.17)	(8.90)	(7.29)
R&D Expenditures	-0.924***	-3.381***	-6.361***	-9.494***
	(-18.61)	(-22.71)	(-20.03)	(-18.08)
Constant	-0.130***	-0.695***	-1.464***	-2.340***
	(-4.77)	(-6.58)	(-8.00)	(-6.28)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.432	0.536	0.546	0.542
Observations	11622	10125	7744	5591

Table IA.2: Relationship between screening efficiency of patent examiners and subsequent outcomes of public firms (A within-firm analysis)

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *ROA* is the ratio of quarterly net income over book assets. *R&D Expenditure* is the R&D expenditure over the book value total assets. *#PatentLitigation* counts the quarterly number of patent litigation that firms act as defendants. *AvgExaminerFalseAcceptRate* is defined as the average false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *FirmSize* is the natural logarithm of book assets. *Leverage* is the total debt (both current liability and long-term debt) over book assets. *Ln(M/B)* is the natural logarithm of the market to book ratio. Quarter fixed effects and firm fixed effects are included in all regressions. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Panel A: Relationship between screening efficiency of patent examiners and subsequent ROA

Dependent Variable	Subsequent ROA			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	-0.013	-0.045*	-0.091***	-0.146***
	(-1.23)	(-1.90)	(-2.62)	(-3.50)
#ApplicationsReviewed	-0.002	0.018**	0.036***	0.033**
	(-0.59)	(2.38)	(2.93)	(2.05)
#PatentsGranted	0.002	-0.014*	-0.017	-0.008
	(0.49)	(-1.93)	(-1.49)	(-0.55)
FirmSize	0.014***	0.025***	-0.007	-0.012
	(6.78)	(4.91)	(-0.85)	(-1.01)
Leverage	-0.087***	-0.218***	-0.174***	-0.012
	(-12.34)	(-12.30)	(-5.69)	(-0.28)
Ln(M/B)	0.023***	0.068***	0.079***	0.075***
	(18.13)	(22.16)	(15.36)	(10.40)
R&D Expenditures	-0.365***	-1.081***	-1.090***	-0.741***
	(-13.53)	(-16.98)	(-10.60)	(-5.03)
Constant	-0.117**	-0.361***	-0.087	-0.377*
	(-2.06)	(-2.85)	(-0.51)	(-1.95)
Firm FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.692	0.875	0.936	0.964
Observations	11954	10536	8170	5995

Panel B: Relationship between screening efficiency of patent examiners and subsequent R&D expenditures				
Dependent Variable	Subsequent R&D Expenditures			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	-0.007*	-0.019*	-0.036**	-0.034*
	(-1.96)	(-1.76)	(-2.16)	(-1.70)
#ApplicationsReviewed	0.000	-0.010***	-0.022***	-0.031***
	(0.25)	(-2.97)	(-3.81)	(-3.92)
#PatentsGranted	-0.000	0.008**	0.017***	0.029***
	(-0.27)	(2.48)	(3.07)	(3.85)
FirmSize	-0.007***	-0.017***	-0.013***	-0.021***
	(-9.24)	(-7.14)	(-3.12)	(-3.46)
Leverage	0.008***	0.018**	0.024	0.001
	(3.05)	(2.17)	(1.64)	(0.05)
Ln(M/B)	-0.001	-0.006***	-0.007***	-0.010***
	(-1.36)	(-4.20)	(-2.95)	(-2.95)
R&D Expenditures	0.351***	0.838***	0.611***	0.016
	(35.03)	(28.58)	(12.29)	(0.23)
Constant	0.056***	0.182***	0.219***	0.394***
	(2.65)	(3.12)	(2.62)	(4.21)
Yes	Yes			
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.844	0.917	0.956	0.975
Observations	11965	10572	8215	6039

Panel C: Relationship between screening efficiency of patent examiners and subsequent patent litigation

Dependent Variable	Subsequent #PatentLitigation			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate	0.140	0.511**	0.594	1.008*
	(1.49)	(2.01)	(1.37)	(1.82)
#ApplicationsReviewed	-0.122***	-0.474***	-0.949***	-1.248***
	(-4.21)	(-5.89)	(-6.26)	(-5.84)
#PatentsGranted	0.006	0.002	-0.006	-0.110
	(0.22)	(0.02)	(-0.04)	(-0.54)
FirmSize	-0.008	-0.049	-0.138	-0.201
	(-0.42)	(-0.92)	(-1.33)	(-1.30)
Leverage	-0.187***	-0.816***	-1.236***	-0.989*
	(-2.92)	(-4.38)	(-3.38)	(-1.84)
Ln(M/B)	-0.011	-0.066**	-0.169***	-0.204**
	(-0.98)	(-2.04)	(-2.68)	(-2.17)
R&D Expenditures	0.142	0.520	0.586	0.610
	(0.58)	(0.76)	(0.46)	(0.32)
Constant	0.024	0.261	2.836*	4.158
	(0.05)	(0.19)	(1.65)	(1.52)
Firm FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.499	0.691	0.788	0.871
Observations	12204	11197	9228	7182

Table IA.3: Relationship between screening efficiency of patent examiners and subsequent outcomes of public firms (Robustness tests)

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *ROA* is the ratio of quarterly net income over book assets. *R&D Expenditure* is the R&D expenditure over the book value total assets. *#PatentLitigation* counts the quarterly number of patent litigation that firms act as defendants. *AvgExaminerFalseAcceptRate*$_{Adj}$ is defined as the average (art-unit adjusted) false acceptance rates of examiners that are related to all granted and rejected applications for each firm in a given past three-year rolling window as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *FirmSize* is the natural logarithm of book assets. *Leverage* is the total debt (both current liability and long-term debt) over book assets. *Ln(M/B)* is the natural logarithm of the market to book ratio. All accounting variables (i.e., *ROA*, *R&D Expenditures*, *Leverage*, *Ln(M/B)*) are winsorized at 0.1% and 99.9%. Quarter fixed effects and industry (two-digit SIC code) fixed effects are included in all regressions. All regressions are OLS regressions with standard errors double clustered at the firm and quarter level. t-statistics are in parentheses. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Panel A: Relationship between screening efficiency of patent examiners and subsequent ROA

Dependent Variable	Subsequent ROA			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
AvgExaminerFalseAcceptRate$_{Adj}$	-0.041***	-0.093***	-0.199***	-0.231***
	(-3.84)	(-2.69)	(-3.33)	(-2.68)
#ApplicationsReviewed	-0.018***	-0.068***	-0.139***	-0.193***
	(-4.85)	(-4.92)	(-4.70)	(-4.02)
#PatentsGranted	0.017***	0.065***	0.133***	0.178***
	(4.72)	(4.82)	(4.63)	(3.86)
FirmSize	0.011***	0.047***	0.090***	0.130***
	(14.24)	(18.27)	(18.12)	(15.41)
Leverage	-0.085***	-0.297***	-0.483***	-0.638***
	(-13.22)	(-14.29)	(-11.73)	(-8.93)
Ln(M/B)	0.013***	0.045***	0.085***	0.129***
	(11.29)	(12.41)	(11.32)	(9.81)
R&D Expenditures	-0.935***	-3.463***	-6.539***	-9.653***
	(-18.56)	(-23.36)	(-20.91)	(-18.97)
Constant	-0.149***	-0.764***	-1.618***	-2.605***
	(-5.36)	(-7.07)	(-8.66)	(-6.80)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.408	0.522	0.542	0.537
Observations	11954	10536	8170	5995

Panel B: Relationship between screening efficiency of patent examiners and subsequent R&D expenditures

Dependent Variable	Subsequent R&D Expenditures			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
$\text{AvgExaminerFalseAcceptRate}_{\text{Adj}}$	-0.006*	-0.020	-0.042*	-0.079**
	(-1.76)	(-1.62)	(-1.86)	(-2.24)
#ApplicationsReviewed	0.000	0.001	0.004	0.001
	(0.42)	(0.33)	(0.52)	(0.11)
#PatentsGranted	0.001	0.004	0.008	0.022*
	(0.54)	(1.06)	(1.03)	(1.72)
FirmSize	-0.002***	-0.011***	-0.025***	-0.042***
	(-7.49)	(-10.81)	(-11.43)	(-10.38)
Leverage	-0.001	-0.014**	-0.062***	-0.129***
	(-0.76)	(-2.12)	(-4.08)	(-5.05)
Ln(M/B)	0.002***	0.006***	0.012***	0.017***
	(3.64)	(4.06)	(3.57)	(2.89)
R&D Expenditures	0.808***	2.961***	5.466***	7.680***
	(25.63)	(27.51)	(23.02)	(17.36)
Constant	0.031***	0.195***	0.479***	0.931***
	(4.83)	(5.95)	(9.12)	(8.09)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.760	0.796	0.784	0.767
Observations	11965	10572	8215	6039

Panel C: Relationship between screening efficiency of patent examiners and subsequent patent litigation

Dependent Variable	Subsequent #PatentLitigation			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
$\text{AvgExaminerFalseAcceptRate}_{\text{Adj}}$	0.229*	0.890**	1.659**	2.641**
	(1.70)	(2.57)	(2.55)	(2.35)
#ApplicationsReviewed	0.029**	0.116**	0.253***	0.497***
	(2.00)	(2.51)	(2.69)	(3.24)
#PatentsGranted	0.005	0.019	0.034	-0.025
	(0.36)	(0.48)	(0.40)	(-0.18)
FirmSize	0.069***	0.270***	0.572***	0.925***
	(11.14)	(11.91)	(11.53)	(10.78)
Leverage	-0.284***	-1.227***	-2.695***	-4.449***
	(-5.63)	(-6.35)	(-6.24)	(-5.81)
Ln(M/B)	0.018**	0.065**	0.153**	0.344***
	(2.42)	(2.37)	(2.49)	(3.01)
R&D Expenditures	0.777***	2.854***	5.322***	7.244***
	(8.31)	(9.22)	(7.94)	(5.99)
Constant	-0.226	-0.712	-2.396*	-4.485**
	(-0.96)	(-1.16)	(-1.93)	(-2.46)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
R^2	0.122	0.159	0.179	0.197
Observations	12204	11197	9228	7182

Table IA.4: Relationship between screening efficiency of patent examiners and subsequent exits of private firms (Robustness tests)

The sample consists of firms that have at least one patent application filed since 2010 and with application outcome available by 2018. *SuccessfulExit* is a dummy, which equals one if a given private firm has exited through an IPO or an M&A by the end of my sample period and zero otherwise. $AvgExaminerFalseAcceptRate_{Adj}$ is defined as the average (art-unit adjusted) false acceptance rates of examiners that are related to all granted applications for each firm in the past three years as described in Section 6.1.1, where the false acceptance rate of an examiner associated with each patent application is defined as the ratio of falsely accepted applications over all applications he/she has made decisions prior to that patent application. A patent application is falsely accepted if it is accepted by the actual examiner but rejected by the machine learning algorithm. *#ApplicationsReviewed* and *#PatentsGranted* count the number of patent applications being reviewed and accepted for each firm in a given past three-year rolling window. *LnVCFinancingAmount* is the natural logarithm of the quarterly investment amount for each firm. *LnNumberFundInvested* is the natural logarithm of the quarterly number of invested funds for each firm. *TotalFundingToDate* is the natural logarithm of total funding each firm has received prior to a given quarter. *LnFirmAge* is the natural logarithm of firm age, which equals the current year minus the firm founding year plus one. Year fixed effects, industry (two-digit SIC code) fixed effects, and state fixed effects are included in all regressions. t-statistics are in parentheses. All regressions are OLS regressions with standard errors clustered at the state level. ***, ** and * indicate significance at the 1%, 5%, and 10% level, respectively.

Dependent Variable	Subsequent SuccessfulExit			
	1 Quarter	1 Year	2 Years	3 Years
	(1)	(2)	(3)	(4)
$AvgExaminerFalseAcceptRate_{Adj}$	-0.012	-0.065**	-0.131**	-0.151***
	(-1.24)	(-2.37)	(-2.29)	(-2.81)
#PatentsGranted	0.006***	0.019***	0.035***	0.056***
	(2.85)	(4.07)	(3.23)	(3.20)
#ApplicationsReviewed	-0.002	-0.006	-0.010	-0.028
	(-1.04)	(-1.03)	(-0.79)	(-1.46)
InvestmentAmount	-0.005**	-0.005**	0.007*	0.004
	(-2.61)	(-2.34)	(1.80)	(0.53)
NumberFundInvested	0.005***	0.010***	0.010**	0.017***
	(4.02)	(3.79)	(2.55)	(3.52)
TotalFundingToDate	0.001	0.004***	0.008***	0.010***
	(1.28)	(4.33)	(3.11)	(3.51)
LnFirmAge	0.010***	0.020**	0.022	0.019
	(3.34)	(2.68)	(1.53)	(0.81)
Constant	-0.038***	-0.044	0.078	0.136
	(-7.08)	(-1.15)	(0.76)	(0.98)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes
R^2	0.012	0.026	0.040	0.050
Observations	13478	12545	10022	7413