Paper to be presented at
the DRUID16 20th Anniversary Conference
Copenhagen, June 13-15, 2016

Independent Boards and Innovation

Benjamin Balsmeier
TU Freiberg
Int. Management and Corporate Strategy
benjamin.balsmeier@gmail.com

Abstract

Much research has suggested that independent boards of directors are more effective in reducing agency costs and improving firm governance. Less clear, however,is how they influence Innovation and Innovation search strategies.Relying on regulatory changes for identification, we Show that firms that Transition to independent boards focus on more crowded and familiar areas of technology. They patent and claim more and receive more total future citations to their patents, though the citation increase comes mainly from incremental patents in the middle of the citation distribution; the numbers of uncited and highly cited patents-arguably corresponding to riskier and completely failed or breakthrough inventions,respectively - do not change significantly. Overall the results indicate that strengthened governance improves innovation performance along existing trajectories, without changing the possibility of a breakthrough.

Jelcodes:G34,-

Independent Boards and Innovation*

February 29, 2016

Abstract

Much research has suggested that independent boards of directors are more effective in reducing agency costs and improving firm governance. Less clear, however, is how they influence innovation and innovation search strategies. Relying on regulatory changes for identification, we show that firms that transition to independent boards focus on more crowded and familiar areas of technology. They patent and claim more and receive more total future citations to their patents, though the citation increase comes mainly from incremental patents in the middle of the citation distribution; the numbers of uncited and highly cited patents - arguably corresponding to riskier and completely failed or breakthrough inventions, respectively - do not change significantly. Split samples indicate stronger effects for research-intensive firms and firms with more entrenched managers. Overall the results indicate that strengthened governance improves innovation performance along existing trajectories, without changing the possibility of a breakthrough.

Keywords: Corporate Governance, Innovation, Patents, Board Composition, Independent Directors

JEL Classification: G34, L14, L25, M21

*

1 Introduction

The board of directors has an important role in the governance of corporations. Charged with overseeing and advising managers, it can effectively reduce agency costs that arise from the separation of ownership and control.

Several authors have argued that independent directors, with no ties to the company other than their directorship, are better suited to perform this role as they can credibly limit managerial discretion by punishing managers after undesirable outcomes. Independent boards are thus more likely to produce decisions that are consistent with shareholder-wealth maximization (e.g. Fama and Jensen, 1983; Williamson, 1983).

Such limited managerial discretion, however, may have unintended effects on corporate innovation. Career concerns might discourage innovation (Aghion, Van Reenen, and Zingales, 2013); the fear of being punished (or dismissed) after poor performance could induce the managers to play it safe and avoid risky innovation strategies (Manso, 2011). A manager with limited discretion may be reluctant to engage in exploratory projects, since the value of those projects depends on the flexibility to adapt after observing outcomes (Manso, Balsmeier, Fleming, 2015).

We investigate the effect of board independence on search and innovation processes. Consistent with many classical models, we first propose that patent counts grow as boards become more independent. Independent boards are more likely to terminate the manager in case of poor performance (Weisbach, 1988) and this threat provides an incentive to the manager to work hard (Sitglitz and Weiss, 1983). Increased monitoring from independent boards may alleviate agency problems such as shirking or tunneling of corporate resources. Managers should also take actions that are - and appear to be - closer to the interests of shareholders (Harris and Raviv, 1978; Holmstrom, 1979; Holmstrom and Milgrom, 1991). When under increased scrutiny and demands for results, managers will focus on quantifiable results - such as a greater number of patents. They will adduce an increase in patent count to satisfy demands for performance.

On the other hand, by punishing managers after failure and limiting managerial discretion, board independence may also stifle - or at least redirect - innovation. A second and more subtle hypothesis is that as boards become more independent, patents filed by the firm will become less novel and less creative. Greater oversight will increase managerial focus upon immediate commercial gain; rather than embark on risky exploration of new technologies, managers will focus on harvesting currently successful approaches. Because they are more likely to get fired for poor performance, and are under pressure for immediate results, managers become risk adverse and invest less in potentially lengthy and fruitless searches for truly novel solutions. This occurs despite the possibility that exploration could increase the chance for a breakthrough, because it may also increase the chance for complete failure (the probability distribution flattens as failures and breakthroughs both become more likely). Even if it occurs, a breakthrough becomes less valuable to a manager, because an independent board is less likely to agree upon its subsequent commercialization strategy. Hence, a manager is less likely to search for that breakthrough in the first place. Rather than explain why they are moving the firm into an area in which it - and possibly no firm - has experience, managers will stick to proven competencies that the board and the market can understand. All these mechanisms cause a manager to exploit in order to maximize the mean outcome, rather than explore and increase the variance that could open up a breakthrough (March 1991).

These behaviors cohere with a number of theoretical models. Because board independence limits managerial discretion, it can affect the type of projects managers choose to engage in (Manso, Balsmeier, Fleming, 2015); the loss of flexibility produced by independent boards induces the manager to choose less exploratory projects, because they require less future adaptability. Independent boards are more likely to fire a manager after poor performance. As argued by Manso (2011), this lack of job security induces managers to pursue less exploratory projects. Managers will avoid new technologies that might be construed as empire building (Jensen, 1986). Boards may also directly resist exploration of new areas, if they fear that in the short-term the stock market fails to properly value investments in innovation (Stein, 1989; Cohen, Diether, Malloy, 2013). Due to potential conflict of interests between independent board members and the manager, or alternately, less familiarity with the firm's industry and technology, the quality of research advice given by friendly boards may be higher than by independent boards (Adams and Ferreira 2007).

These arguments imply observable outcomes. Firms whose boards become independent will patent and claim more but less novel technology. The patents will be in areas the firm has previously patented in. The patents will not be in new areas that other firms also avoid. Citations to the firm's patents will increase, though this increase will not result from investments in risky technologies that might provide a breakthrough or fail completely; instead, the citation increase will be generated by patents in the middle of the citation distribution. Furthermore, this increase in citations will be mediated by the movement of the firm into better known and more crowded areas of technical search; in other words, the firms' patents will be more highly cited simply as an artifact of their search strategy and the citation norms of more crowded fields. Rather than start new technological trajectories, managers will direct their efforts towards maximizing the return on previously proven trajectories. They will increase the first moment of innovative outcomes and decrease the second.

Evidence comes from observing search strategies for firms that were forced by regulatory changes to adopt more independent boards. Starting in 2002, stock exchanges and the Sarbanes-Oxley Act (SOX) required firms to have a majority of independent directors (for a similar approach, see Duchin et al. 2010). Comparing firms that changed from less to more independent boards against firms that already had independent boards, we find increased patent and claim output - but less creative and explorative inventions. Firms whose boards become more independent patent more and receive more citations to their patents, however, the effects are not significant for uncited and highly cited patents. Firms whose boards become more independent also work in more crowded and more familiar technologies; the rates of prior and self-citation increase. Moreover, results are more pronounced for firms with high R&D stock and a high entrenchment index. The evidence supports arguments for a more nuanced relationship between oversight and innovation; greater oversight appears to lead to greater focus and incremental output but have no impact on breakthrough creativity and exploration.

2 Literature review

A large literature studies the role and influence of board characteristics (for an overview see Adams, Hermalin, and Weisbach, 2010; for the economic relevance of boards see Ahern and Dittmar, 2012). Much of the literature focuses on the role of independent board members (most recently e.g. Masulis and Mobbs, 2014; Brochet and Srinivasan, 2014). Several studies have analyzed how independent directors influence CEO compensation (e.g. Faleye, Hoitash, and Hoitash 2011; Coles, Daniel, and Naveen, 2008; Denis and Sarin, 1999; Core, Holthausen and Larcker, 1999), CEO appointments and dismissals (Knyazeva, Knyazeva, and Masulis, 2013; Guo and Masulis, 2011; Borokhovich, Parrino, and Trapani, 1996; Weisbach, 1988), adoption of antitakeover defenses (Brickley, Coles, and Terry, 1994) or takeover premiums (Cotter, Shivdasani, and Zenner, 1997; Byrd and Hickman, 1992). From these studies the picture emerges that independent board members increase board oversight. Whether such intensified board monitoring is beneficial or detrimental to shareholder wealth is less clear and may depend on the complexity of a firm's operations (Faleye, Hoitash, and Hoitash, 2011; Duchin et al, 2010; Nguyen and Nielsen, 2010).

Several recent papers use patent data to empirically study how corporate governance affects innovation. Raw patent counts are usually supplemented by the number of citations that a patent receives, as this measure correlates with financial and technical value (Harhoff 1999; Hall et al., 2005); future cites are sometimes broken down by whether the firm cites its own work (Lerner, Sorensen, and Stromberg, 2011). Though less common, some papers have analyzed technology classes or the tails of the citation distribution (Gonzalez-Uribe and Xu, 2015; Byun, Oh, and Xia, 2015; Cerqueiro, Hegde, Penas, and Seamans, 2015). Measures of originality and generality (Hall, Jaffe, and Trajtenberg 2001) have also been used, though these measures depend on the US Patent and Trademark Office's changing and now discontinued classification of technologies (see Lerner, Sorensen, and Stromberg, 2011 and Hsu, Tian, and Xu, 2014). Lerner and Seru (2014) detail a number of problems with the use of patent measures in the finance literature, including failures to correct for differences in time periods and truncation (caused by the lag between application and patent grant, or delay in the accumulation of future prior art citations), economic value (typically proxied by future prior art citation), technology (typically measured by the United States Patent and Trademark Office classes), and disambiguation of assignees (determining which firms own which patents). One empirical contribution of this paper is to offer improved and easily calculated measures that can address some of these issues.

The results of the recent surge of empirical patent work on governance and innovation are decidedly mixed. Much of the contradictory work is well identified, so resolution will have to rely on sharper theory or more careful measurements of governance and innovation. A variety of papers find that stronger governance leads to greater innovation (alternately, weaker governance leads to decreased innovation). Aghion, van Reenen, and Zingales (2013) show that greater institutional ownership correlates with greater patenting and citations to patents. Bernstein (2014) finds that firms experience no change in the amount of patenting following an IPO (when they would assumedly transition from strong oversight by venture capitalists to weaker public oversight), however, they do experience a decrease in citations. Atanassov (2013) found that a strengthening of anti-takeover provisions in a state (assumedly implying weaker governance) led to fewer patents and citations, but that institutional shareholders decreased the effect. Supra, Subramanian, and Subramanian (2015) used a similar research context to Atanassov (2013) but found a non-monotonic effect, where innovation increased for firms that experienced very weak and very strong external takeover pressure.

In contrast, a variety of papers finds that weaker governance leads to increased innovation (alternately, stronger governance leads to decreased innovation). Atanassov (2016) finds that firms with a greater proportion of bank financing invented more and more highly cited patents (and that the volatility of citations was greater). In contrast to Atanassov (2013) and partial contrast to Supra, Subramanian, and Subramanian (2015), Chemmanur and Tian (2012) find that firms with greater anti-takeover provisions receive more and more highly cited patents.

Most similar to the current study, Faleye, Hoitash, and Hoitash (2011) find that monitoring intensity, as measured by the proportion of independent board directors on at least two monitoring committees, correlates negatively with research and development expenditure and future prior art citation counts. While they present well specified panel data models, their Sarbanes Oxley regressions (the main identification strategy used in the current paper) investigate the effect of SOX on firm value - but not, however, on R&D and patent data. Kang et al. (2014) find no correlation with social connections between the CEO and board members and research and development spending (arguably a social connection implies weaker governance); they find a positive correlation with patents and citations.

Using a differences in differences identification strategy based on the regulatory requirements of SOX, we find no effect of a firm's transition to an independent board upon R&D spending but a positive effect on total patenting and citations and a focusing of the firm's innovative search on known and previously successful areas; these results remain robust across a variety of matched, fixed effects, and trend control models.

Taking heed of the critiques of Lerner and Seru (2014), this paper assembles a suite of more detailed and nuanced measures of innovation. This battery of measures offers additional and consistent insights into the mechanisms of how board independence influences innovation, while retaining the advantages of the SOX identification strategy. Of more general interest, the battery of measures enables cleaner identification of a firm's search strategy; it illustrates how a firm can invent more highly cited and valuable patents by exploiting its current area of expertise. Such exploitation may be characterized as less novel or creative and is probably less risky, however, it is still innovation and an effective and valuable search strategy for the firm.

3 Identification strategy

Identification for our study relies upon regulatory changes that forced public firms to increase the presence of independent directors on their boards in the early 2000s. The effects of those regulatory changes on variables other than innovation have been analyzed elsewhere (see e.g. Duchin et al., 2010, for a setup that is most similar to ours). In this section, we briefly describe the regulatory framework that is relevant to our analysis.

Initiated by recommendations of the Blue Ribbon Committee (BRC) in 1999, stock market rules of the NYSE and Nasdaq have been built upon the assumption that independent board members are better able to monitor managers. Subsequent to the BRC recommendations, the Securities and Exchange Commission (SEC) approved new rules in December 1999, requiring public firms to move to a fully independent audit committee with the next re-election or replacement of audit committee members. Further motivated by prominent corporate scandals, e.g. Enron, this rule was written into U.S. law in 2002 as a part of the Sarbanes-Oxley Act (SOX). It was followed by subsequent NYSE and Nasdaq regulations in 2003 that imposed even stricter requirements on board composition. In addition to having an audit committee composed of exclusively independent directors, both stock exchanges forced firms to have a majority of independent directors as regular board members, and the compensation and nomination committees had to consist of 100% independent board members (>50% if firms are listed on Nasdaq only).

Definitions of director independence vary slightly across each rule. SOX states in section 301 that a given director is independent if the person does not "accept any consulting, advisory, or other compensatory fee from the issuer" (except for serving the board), and is not an "affiliated person of the issuer or any subsidiary" (NYSE speaks of "no material relationship"; and Nasdaq requires no relationship that would interfere with "independent judgment"). The NYSE and Nasdaq regulations are clear; the independence assumption is violated, for instance, if a director him- or herself or a direct family member was an employee of the firm during the previous three years, or a family member works for a third firm with which the given firm has a professional relationship, or a family member is connected to the firm's auditor.

These regulations made board changes necessary for a large group of firms. The number and fraction of independent board members was fairly stable until the year 2000. As the described board regulations came into effect, more and more independent directors were appointed to corporate boards. Figure 1 illustrates the changes in board composition for the sample of firms used in our study. It resembles a pattern that has been documented in other studies for differing sets of public firms (e.g. Linck et al, 2008; Duchin et al., 2010). Board composition data are taken from the Investor Responsibility Research Center (IRRC). From 1996 to 2006 the IRRC tracked individual board members of all major public U.S. firms and indicated in their database whether an individual board member is independent, an employee of the firm or otherwise affiliated (former employee, employee of an organization that receives charitable gifts from the company, employee of a customer or supplier to the company, relative of an executive director, etc.).

Reflecting the previously introduced regulatory changes, Figure 1 shows an increase of independent director presence on corporate boards from 2001 to 2006. Theoretical considerations about board control suggest that a crucial difference arises when a board switches from a minority to a majority of independent board members (Harris and Raviv, 2008).[1] It was further an explicit requirement of regulatory reforms. Thus, our

[1]The fraction of independent board members provides more variation but has two major disadvantages. First, considering board voting behavior, it is likely that the influence of independent directors on board oversight does not increase linearly with the number or fraction of independent members but exhibits a jump when independent directors gain or lose the majority of votes. Second, the switch from a minority to a majority of independent directors was an explicit requirement of regulation, such that it is more likely that observed changes in that regard happened involuntarily, which in turn improves the

Figure 1: Fractions of independent boards and directors

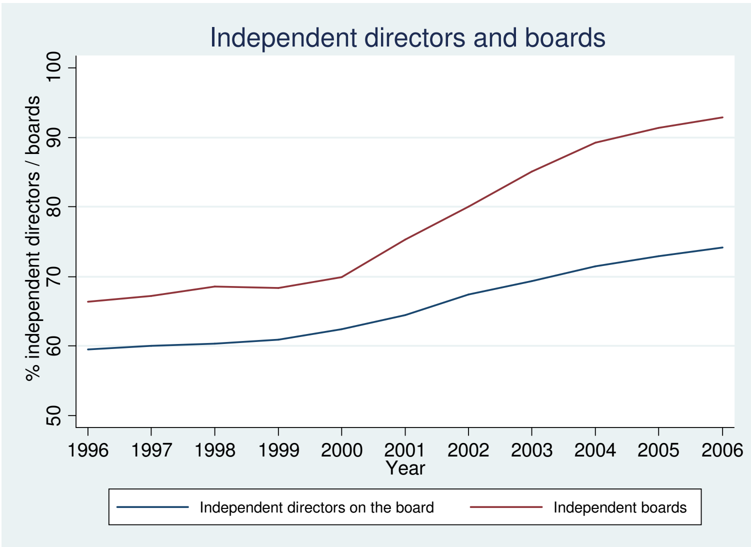

Notes: This graph illustrates the evolution of independent boards over the sampling period. A board is defined as independent in the empirical estimations if the majority of board members are classified as independent by the Investor Responsibility Research Center (IRRC). In this graph, independent directors represents the average fraction of independent board members of all firms in the study.

analysis focuses on this variable. Our data also show that the proportion of firms with a majority of independent board members stayed rather stable around 68% before 2000 and moved up to about 94% by 2006.

Our empirical identification of the relationship between board independence and innovation stems from the difference between firms who were already in compliance with the regulatory changes before 2001 and those firms who switch to a majority of independent directors after regulatory changes became effective. Hence, all firms that were not required to change their board serve as a control group. In line with Duchin et al. (2010), we define firms as treated when they switch to an independent board in 2001 or later and have an audit committee that contains 100% independent board members. The latter requirement helps to sort out potential voluntary switches, increasing the amount of truly exogenous increases of independent board members and making our main variable of interest less likely to be confounded by endogenous choice. The fraction of independent directors increased by 25% during 2001 to 2006 within noncompliant firms and by 9% within firms that had already fulfilled the regulatory requirements before 2001.

4 Sample selection

The dataset we built up for our study is determined by the joint availability of data on the composition of corporate boards and committees from the IRRC, information on basic firm characteristics from Compustat, and patent data from the NBER, the Fung Institute and the United States Patent and Trademark Office (USPTO). The IRRC provides data on corporate board members for 3000 major public U.S. based firms from 1996 to 2006. Compustat has further information on almost all of the firms covered by IRRC. A major challenge for the empirical researcher interested in those firms' innovative activities is the identification and compilation of the corresponding patent portfolios. Researchers involved in the NBER patent data project have spent significant amounts of resources to identify patents that have been granted to U.S. based firms. The NBER patent database contains, however, only those patents that have been granted through 2006. Due to the time lag with which inventions are granted property rights (1-5 years) and the publication of corresponding data by the USPTO, this results in significantly truncated data for patents filed after 2001. Researchers have found ways to use incomplete patent data for the years 2002 to 2006, exploiting the distribution of applications before 2002, but those approaches add noise to econometric analyses, and lead to significant estimation errors in our case, because our sample of board data covers 50%

identification of causal effects.

of years for which the NBER data is severely truncated. The issue becomes even more prevalent if researchers want to take citations to patents into account that often occur several years after a patent has been granted. In terms of patent applications, the NBER data misses 18 percent of patent applications of U.S. based assignees identified in 2002, rising to 99 percent by 2006.[2]

Newly available disambiguations (see Fiero et al., 2014) provide more recent data, avoid the truncation of the NBER patent database, and identify comprehensive patent portfolios of the firms in our sample up to the year 2007.[3] Following the literature (e.g. He and Tian 2013), we assign an eventually granted patent to the year it was applied for. Disambiguation of firm names presents a major challenge, since patent documents do not contain a unique identifier of assignees. Following disambiguation, patents are aggregated to the firm level and merged with other databases such Compustat and IRRC.

We extended the reach of the NBER patent database by combining it with USPTO and Fung Institute data, including patent citations and other detailed information within each patent document. We started with standardized assignee names provided by the USPTO for all patents granted through December 31, 2012. These standardized assignee names are largely free of misspellings but still contain many name abbreviations for individual firms. The standardized USPTO assignee names remain consistent throughout time and have been used by the NBER patent project team to disambiguate firm names. For almost all US firms that received at least one patent between 1975 and 2006, the NBER provides a unique time invariant assignee. We took all variations of standardized assignee names that belong to a given single firm as a training set, and gave all granted patents that appear with the same standardized assignee name the same unique NBER identifier.[4] These information enabled us to track firms' patenting activity over significantly longer time periods, overcoming truncation issues of patent applications and generally increasing the accuracy of available patent portfolios.

Finally, we merged unique time invariant Compustat identifiers to the patent assignee identifiers as they are provided by the NBER. It is worthwhile to note that in our analysis we take only those firms into account for which the NBER has identified Compustat matches, and we assigned zero patents only to those firms where the NBER team searched for but could not find matches with any patent. In this regard we devi-

[2]The numbers are derived by comparing all patent applications in the NBER database with all patents in the Fung Institutes database as published in April 2014.

[3]We gather patent data through 2007, because we will estimate regressions of firms' patenting activities in year t on board data and controls in t-1, reflecting that patenting activities need some time to be influenced by boards and simultaneous determination of variables may otherwise confound the estimation.

[4]Based on the first assignee that appears on the patent document. It allowed us to identify ~250k additional patents granted to U.S. based assignees after 2006.

ate from other studies that assign zero patents also to those firms that have not been tested to appear as a patent assignee or not. Thus, we avoid measurement errors at the expense of a smaller but more accurate dataset.

In order to circumvent potential selection effects to confound our estimation of the relationship between board independence and innovation, we further removed all firms that appear only before the year 2000 or entered the sample in the year 2000 or later, such that the remaining firms can be observed over a timespan where the previously described regulatory changes took place. Finally, we arrive at a sample of 6107 observations on 713 firms observed during the period from 1996 to 2006 for which we could gather all information of interest. All firms in the sample combined have applied for and been granted 328,463 patents during the sample period.

4.1 Measuring innovative search

Much recent empirical work on corporate governance and innovation has relied on patent data (e.g. Atanassov, 2013; He and Tian, 2013). Raw patent counts are used as well as the number of future prior art citations that a patent receives, as the number of future cites correlates with financial and technical value; highly cited patents are much more valuable commercially and the relationship is highly skewed in favor of very highly cited patents (Harhoff 1999; Hall et al., 2005). To be comparable with the extant literature we will show how board independence influences patent counts and citations. Our results go on to illustrate, however, that raw patent counts and total citation counts are of limited use in identifying differences in innovative search strategies, specifically towards more or less exploration. Therefore, we introduce a suite of measures, consistent with the arguments of Lanjouw and Schankerman (2004) for the use of multiple indicators of patent quality. These serve as additional dependent variables beside raw patent counts and citations, thus enabling the illustration of a richer and more robust picture of how board independence affects not only the rate but also the type and direction of innovation.

First, we calculate the number of citations that each patent makes to prior patents (Lanjouw and Schankerman 2004). An increase in the number of backward citations reflects direct relations to more prior art that must be specified in the patent application (required by law). This correlates with innovative search in relatively more crowded, better-known, and typically more mature technological areas.

Second, we take the number of times a given patent cites other patents owned by the same company (Sorenson and Stuart, 2000; similar measures are used in Faleye, Hoitash, and Hoitash, 2011). More self-cites indicate search within previously known areas of expertise while fewer self-citations indicate a broadening of innovative search or efforts to explore areas that are new to the firm.

Third, we calculate the number of patents that are filed in technology classes previously unknown to the firm. Unknown patent classes are defined as those in which a given firm has not applied for any patent beforehand (starting in 1976). The complement is the number of patents applied for in known classes. Addressing one concern of Lerner and Seru (2014), we consistently use the original patent class at time of patent grant; hence, if the USPTO defines a brand new class and issues a new patent, it will be observed, but if the USPTO redefines an old patent into a new class, it will not change the measure.

A continuous measure of whether firms stay or deviate from known research areas is the technological proximity between the patents filed in year t and the existing patent portfolio held by the same firm up to year t-1 (Jaffe 1989):

$$P_{it} = \sum_{k=1}^{K} f_{ikt} f_{ikt_{-1}} \Big/ \left(\sum_{k=1}^{K} f_{ikt}^2 \cdot \sum_{k=1}^{K} f_{ikt_{-1}}^2 \right)^{\frac{1}{2}}$$

where f_{ikt} is the fraction of firm i's patents that belong to patent class k at time t, and $f_{ikt_{-1}}$ is the fraction of firm i's patent portfolio up to t-1 that belongs to patent class k. P_{it} ranges between 0 and 1. The highest possible value indicates that the patents filed in year t are distributed across patent classes in the exact same way as the portfolio of all patents of the same firm up to the previous year.[5] Positive coefficients in a regression would thus indicate a more narrow innovation trajectory within known areas.

Fourth, we categorize patents according to how many citations they have received relative to other granted patents that have been applied for in the same technology class and year (Azoulay, Graff Zivin and Manso, 2011). In addition to limiting comparison of similar patents, we exclusively and exhaustively bin all patents according to their location in the distribution of citations. This is intended to clearly separate different types and degrees of innovative outcomes, ranging from highly successful breakthroughs (highly cited) to completely failed inventions (not cited at all) and moderately successful outcomes that lie between. We estimate separate models for each of the four non-overlapping categories: top 1%, 2nd-10th%, not in the top 10% but cited at least once, and never cited at all. We count a patent as a top 1% (2-10%) patent if the patent falls into the highest percentile (centile) of the citation distribution in the same technology class and application year. We also separately count all patents that received no citation at all and those that have received at least one citation but do not fall in the top 10% category.

[5] Reflecting that a value of one indicates no change, the measure takes value one if no patent was applied for in a given year. All results presented below are robust to excluding non-patenting firms.

Fifth and finally, we calculate the total number of claims made by a firm's patent portfolio each year (Lanjouw and Schankerman 2004). It is difficult to algorithmically interpret ex ante the innovative value of any particular claim, however, as claims can be added as scope conditions which typically act as limitations on the basic invention. An increase in the total number of claims should correlate, however, with the effort a firm puts into the patenting process, and this effort should increase in response to pressures for immediate and quantifiable results. Balsmeier, Manso and Fleming (2015) present a principal components analysis and more detailed characterizations of these measures.

We do not use measures of originality and generality because their correspondence to exploration and exploitation remains unclear. The measures calculate the spread of classes covered by forward and backward citations, however, they do not take history into account; the spread may be novel and unique, or it may be old and common. For example, a patent may be measured as original because it cites other patents across a wide variety of classes, yet that citation pattern may have already appeared on any number of patents. Additional pragmatic issues make the measure unattractive: 1) it is only calculated for the NBER sample, 2) any calculation relies upon the concordance of classes which changes as each new class is defined, and 3) the USPTO recently stopped using the US class system, hence it will be impossible to update the measure going forward. Unreported regressions available from the first author show no significant effect of board transition on the average of a firm's patent scores of originality and generality, for the patents in the NBER sub-sample. Individual level patent regressions similarly show no significant relationship.

4.2 Control variables

Following the extant literature, we control for a vector of firm characteristics that could confound the relation between board independence and a firm's innovative search and success. We compute all variables for firm i over its fiscal year t. *Board size* measures the number of board members as we want to insulate the effect of board independence from contemporary changes in the number of directors. Further, we found that the firms in our sample differ significantly in terms of *R&D* spending over total assets and firm size as measured by total assets - two variables that are naturally positively related to firms' innovation activities. In order to reduce the skewness in total assets we take the logarithm of total assets in all multivariate econometric analyses. In addition, we control for firm age (the number of years since the initial public offering date), as older firms may search in older technological areas. Moreover, *leverage* (long term debt over total assets) and *capital expenditures* (scaled by total assets) account for financial constraints that are known to influence corporate innovation. Finally, *Tobin's Q* enters the regression to control for differences in growth opportunities.

4.3 Methodological remarks

In order to analyze how a transition to an independent board affects innovative search we follow the literature on corporate governance and innovation (e.g. Atanassov, 2013; He and Tian, 2013; Kortum and Lerner, 2000) and estimate the baseline model in OLS:

$$log(1 + patents_{i,t+1}) = \beta_0 + \beta_1 \cdot independentboard_{it} + \gamma \cdot Z_{it} + \theta_t + \alpha_i + \epsilon_{it}$$

where $patents_{i,t+1}$ is the number of eventually granted patents of firm i applied for in year *t+1*. In alternative regressions we will exchange the number of patents with our previously introduced measures of innovation that allow us to assess the firms' innovative search strategy in more detail.[6] Our main explanatory variable of interest, *independent board*$_{it}$, is a dummy that indicates firms that have transitioned from a minority to a majority of independent board members in the year 2001 or later when regulatory changes became effective.[7] Under the assumption that changes in patenting by firms that transitioned would have been comparable to changes in patenting by other firms in the absence of a transition, β_1 captures the effect of board independence on innovation by the affected firms.[8] Z_{it} is a vector of the previously introduced firm characteristics, and year fixed effects θ_t control for changes in the macroeconomic environment and systematic changes in patenting activities over time. Our preferred specifications include firm fixed effects α_i that control for any unobserved firm heterogeneity that is time invariant. Hence, we basically estimate a DiD model, where those firms that switch from a minority to a majority of independent directors on the board in 2001 or later are the 'treated firms', and all others are 'non-treated firms'. In order to unravel the influence of firm fixed effects in our regressions we also show alternative models with

[6] In case the dependent variable is a count, all results are robust to alternatively estimating Poisson models (not shown).

[7] All results presented below are robust to alternatively taking the years 2000 or 2002 as the threshold value.

[8] As can be seen in Figure 2, not all firms transitioned from a friendly to an independent board at the same time, because directors were allowed to fulfill their contracts that were signed before the law change. In principle, this gives firms room for strategic choice that could confound our identification. Therefore, we checked whether the time between the law change and compliance is correlated with pre-SOX innovative activity of the firms in our sample. In order to test this, we first defined a variable that measures the years until the board actually changed from friendly to independent although SOX and other regulations were already active (2003). We found 17 firms with a one year lag, 14 with a two year lag and 8 with a three year lag. Then, we regressed time lag until compliance on firms' average amount of R&D, patents and cites before 2001 (results are robust to taking 2000 or 2002 instead). The lack of significant correlation between compliance lags and pre-treatment innovative activity increases confidence that the estimation is not biased by systematic choice of more or less innovative firms to transition later or earlier.

industry fixed effects, based on 3-digit SIC industry dummies, instead of firm fixed effects. To stay within the DiD framework, we include a dummy variable that marks all treated firms in those regressions without firm fixed effects.

Identification hinges in all models upon the parallel trend assumption; treated and non-treated firms show similar trends in the dependent variable of interest in the absence of treatment. To increase our confidence in this assumption, we estimate the dynamics of the treatment effect, which provides evidence that the DiD estimator is not significantly different from zero in the absence of treatment.

Our estimation might still be biased, however, if other remaining cross-sectional heterogeneity of the firms in our sample change systematically with the transition to an independent board and our measures of innovative search. In order to minimize concerns in this regard, we further re-estimate all our models based on a balanced sample, where treated and non-treated firms are comparable in terms of key observable characteristics before 2002. To achieve a balanced sample we use Coarsened Exact Matching (CEM) (Blackwell et al., 2009).[9] CEM has several features that bound the degree of model dependence, reduce causal estimation error, bias, and inefficiency (Iacus, King, and Porro, 2009a, 2009b, 2011, for a similar application, see Azoulay, Zivin, and Wang 2010). Based on CEM's coarsening function we match treated and non-treated firms on the joint distribution of firms' R&D spending over total assets, firm size as measured by the natural logarithm of total assets, the natural logarithm of Tobin's Q, boardsize and 26 two-digit SIC industry code dummies. We took the average values of those variables over the years 2000 and 2001 as matching criteria to ensure highest comparability before treatment.[10] Table 1 presents the differences in mean values of all control variables before and after the matching procedure.

Panel A and B of Table 1 show that treated firms in the full sample are on average a little smaller, invest less in R&D and have a smaller board. Except with regard to R&D spending, the relative differences of the two firm groups appear small in magnitude. Both groups are not statistically significant with regard to the mean values of the other control variables that have not explicitly been included in the matching. In order to eliminate any statistically significant differences of observable firm characteristics, while keeping as many treated firms as possible in the sample, we ran CEM with the side condition to differentiate firms according to ten categories of R&D spending and three categories of firm size, board size and Tobin's Q. Based on this procedure, 4 out of the 125 treated firms remain unmatched. For the remaining 121 treated firms, CEM selected 430 comparable firms, i.e. 158 incomparable firms are subsequently discarded

[9]In alternative models we balanced the sample based on propensity score matching, taking only the nearest neighbor of each treated firm as a control, and find qualitatively the same results.

[10]The results are robust to taking all available observations before 2001 into account.

Table 1: CEM matching of treated and non-treated firms

Variable	**number of firms**	**mean**
Panel A: Treated firms before matching		
log(total assets)	125	7.02
R&D /assets	125	0.04
Age	125	2.45
Leverage	125	0.18
Cap. exp.	125	0.06
log(Q)	125	1.34
Board size	125	8.45
Panel B: Non-treated firms before matching		
log(total assets)	588	7.33**
R&D /assets	588	0.05*
Age	588	2.43
Leverage	588	0.20
Cap. exp.	588	0.05
log(Q)	588	1.25
Board size	588	8.99**
Panel C: Non-treated firms after matching		
log(total assets)	430	6.99
R&D /assets	430	0.04
Age	430	2.37
Leverage	430	0.20
Cap. exp.	430	0.05
log(Q)	430	1.21
Board size	430	8.56

Notes: This table reports mean values of treated and non-treated observable firm characteristics, averaged over the years 2000 and 2001, before and after matching, based on the joint distribution of firms' R&D spending over total assets, firm size as measured by the natural logarithm of total assets, the natural logarithm of Tobin's Q, and board size. ***, **, * denote significance levels of 1%, 5%, and 10% of two sided *t*-tests on the difference between mean values of Panel A and B, and Panel A and C, respectively.

from the analysis. Panel B of Table 1 shows that, after matching, there are no statistically significant differences between the treated and non-treated firms according to two sided t-tests. Although not necessary for a consistent DiD estimation, it is worthwhile to mention that both groups do not differ in terms of the average amount of applied patents after matching.

While balancing the sample should improve identification (at least for firms that are similar to the treated firms), potential remaining differences in innovation trends might still have an influence on the estimation. Therefore, we also estimate models that allow for separate firm specific linear trends in innovation before 2002, using the following specification:

$$\begin{aligned} log(1 + patents_{i,t+1}) = \beta_0 + \beta_1 \cdot independent\ board_{it} + \gamma \cdot Z_{it} \\ + \delta \cdot firm_i \cdot pre2002_t \cdot t + \theta_t + \alpha_i + \epsilon_{it} \end{aligned}$$

where $pre2002_t$ equals one if the year of observation is 2001 or earlier.

Finally, in alternative specifications we further control for potential systematic changes in the influence of our controls on innovation after 2001, which may coincide with changes in board independence, by estimating:

$$\begin{aligned} log(1 + patents_{i,t+1}) = \beta_0 + \beta_1 \cdot independent\ board_{it} + \gamma \cdot Z_{it} \\ + \delta \cdot firm_i \cdot pre2002_t \cdot t + \zeta \cdot Z_{it} \cdot post_t + \theta_t + \alpha_i + \epsilon_{it} \end{aligned}$$

5 Results

We first present results on research and development spending, the number of patents, and the total number of citations to and claims within a firm's patent portfolio. We then present measures of innovative search strategy, including a breakdown of the citation distribution, backward and self-citations, and movement into new classes and across technological distance.

5.1 R&D, patents, citation-weighted patents, and claims

Tables 2 to 5 estimate regressions of firms' R&D investments, the number of eventually granted patents applied for, the total number of citations made to the firm's patents, and the total number of claims contained within a firm's patent portfolio. Each table contains 5 specifications of the same model. Specification (a) is a standard OLS model with industry fixed effects, (b) is a standard firm fixed effects model, (c) is the same as (b) but estimated on the previously described balanced CEM sample, (d) adds trend controls, and (e) adds interaction terms of all controls with a post SOX marker. For all models with firm fixed effects the R squared values refer to the explained within firm variance. The first model assesses potential changes in R&D investments after board independence changed, which might drive subsequent changes in patenting.[11] The next two models differentiate between a change in the number of patents and a change in citations to those patents. Cite-weighted patent counts (the total number of citations to a firm's patent portfolio) have been shown to correlate with financial value and patent renewals (Harhoff 1999; and Hall et al., 2005). The last model estimates effects on the total number of claims in a firm's portfolio.

Table 2 illustrates that a transition to an independent board appears unrelated to the level of firms' R&D investments. In contrast, Tables 3 and 4 illustrate how patenting and total citations both increase. Reading across the models, the effect on patenting ranges between a 31% to 20% increase in the number of patents, and a 59% to 41% increase in total citations. Figure 2 illustrates the dynamics of the latter two effects. For the graphs we defined dummy variables for the specific times before and after firms changed to an independent board. t_0 defines the year of the switch and serves as the baseline category, t_{n-1} defines the number of years before the switch, and t_{n+1} the years after the switch. Then, we ran regressions including these variables instead of the single dummy variable in the baseline model beforehand. As we still include year fixed effects, the coefficients represent the relative change in patenting per year that is attributable to the board change.

Table 5 illustrates that the number of claims in a firm's patents increase following a transition to an independent board. The effect on the number of claims ranges between a 50% to 36% increase in the number of claims. Figure 3 illustrates the dynamics.

The results are consistent with classic agency theory and our first hypothesis, suggesting that intensified monitoring leads to increased effort of the agent, which results in increased claims and patenting of inventions. That firms patent more, but do not spend significantly more on R&D, raises the question whether firms just work more efficiently or exploit extant knowledge at the expense of explorative innovation (models of patenting efficiency were not significant). The arguments for our second hypothesis argue that increased board independence leads to a shift from explorative to more exploitative innovative activities. The following models illustrate a consistent shift towards exploitation but no clear signal of the effect on exploration.

[11] Alternative regressions with R&D investments scaled by total assets reveal a significant positive effect only in specifications without firm fixed effects. Inclusion of controls for time invariant firm heterogeneity leads to statistically insignificant results.

Table 2: Independent boards and R&D

	(a)	(b)	(c)	(d)	(e)
	b/se	b/se	b/se	b/se	b/se
log(total assets)	0.822***	0.564***	0.601***	0.609***	0.602***
	(0.017)	(0.044)	(0.040)	(0.049)	(0.049)
log(age)	-0.153***	0.002	-0.006	-0.017	-0.013
	(0.021)	(0.029)	(0.038)	(0.056)	(0.054)
Leverage	-0.562***	0.040	-0.085	-0.211	-0.462**
	(0.113)	(0.107)	(0.124)	(0.152)	(0.212)
Cap. exp.	0.753	0.562	0.542	0.378	0.820
	(0.616)	(0.351)	(0.391)	(0.431)	(0.518)
log(Q)	0.366***	-0.016	-0.015	-0.014	0.022
	(0.025)	(0.024)	(0.029)	(0.032)	(0.035)
Boardsize	0.024**	0.007	0.004	0.006	-0.004
	(0.009)	(0.008)	(0.011)	(0.013)	(0.014)
Independent board	0.071	-0.052	-0.057	-0.059	-0.043
	(0.090)	(0.055)	(0.056)	(0.064)	(0.061)
Observations	6107	6107	4414	4414	4414
R^2	0.733	0.256	0.254	0.450	0.508
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is log(R&D). All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

Table 3: Independent boards and number of patents

	(a)	(b)	(c)	(d)	(e)
	b/se	b/se	b/se	b/se	b/se
log(total assets)	0.767***	0.273***	0.284***	0.369***	0.425***
	(0.017)	(0.060)	(0.064)	(0.067)	(0.079)
R&D	5.561***	0.941*	0.842	0.711	0.835
	(0.568)	(0.517)	(0.668)	(0.713)	(0.896)
log(age)	0.105***	0.068	0.000	0.004	-0.019
	(0.023)	(0.044)	(0.039)	(0.048)	(0.058)
Leverage	-0.468***	-0.112	-0.094	-0.253	-0.250
	(0.123)	(0.176)	(0.196)	(0.188)	(0.212)
Cap. exp.	1.635***	0.147	0.127	0.321	0.325
	(0.490)	(0.484)	(0.518)	(0.522)	(0.561)
log(Q)	0.199***	0.057*	0.057	0.081**	0.066
	(0.027)	(0.034)	(0.037)	(0.040)	(0.041)
Boardsize	0.015	0.017	-0.003	-0.016	-0.012
	(0.010)	(0.014)	(0.016)	(0.015)	(0.017)
Independent board	0.308***	0.272***	0.215***	0.208**	0.198**
	(0.083)	(0.079)	(0.080)	(0.087)	(0.087)
Observations	6107	6107	4414	4414	4414
R^2	0.571	0.207	0.176	0.410	0.414
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is the logarithm of one plus the number of eventually granted patents. All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

Table 4: Independent boards and number of cite-weighted patents

	(a) b/se	(b) b/se	(c) b/se	(d) b/se	(e) b/se
log(total assets)	0.919***	0.329***	0.287***	0.326***	0.525***
	(0.027)	(0.090)	(0.099)	(0.115)	(0.127)
R&D	7.750***	2.451***	2.671**	3.257***	4.779***
	(0.870)	(0.835)	(1.036)	(1.162)	(1.413)
log(age)	0.145***	0.072	0.013	0.029	-0.052
	(0.038)	(0.056)	(0.059)	(0.080)	(0.090)
Leverage	-0.374*	0.110	0.323	0.191	0.162
	(0.200)	(0.261)	(0.301)	(0.304)	(0.388)
Cap. exp.	2.636***	0.163	0.233	0.610	0.567
	(0.803)	(0.821)	(0.856)	(0.972)	(1.118)
log(Q)	0.351***	0.220***	0.237***	0.240***	0.272***
	(0.046)	(0.056)	(0.062)	(0.077)	(0.086)
Boardsize	-0.002	-0.005	-0.031	-0.048*	-0.048
	(0.015)	(0.021)	(0.026)	(0.027)	(0.032)
Independent board	0.472***	0.738***	0.599***	0.498**	0.499**
	(0.146)	(0.161)	(0.169)	(0.214)	(0.212)
Observations	6107	6107	4414	4414	4414
R^2	0.505	0.316	0.284	0.445	0.454
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is the logarithm of one plus the number of citation-weighted patents. All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

Figure 2: Dynamics of independent board effect on patents and citations

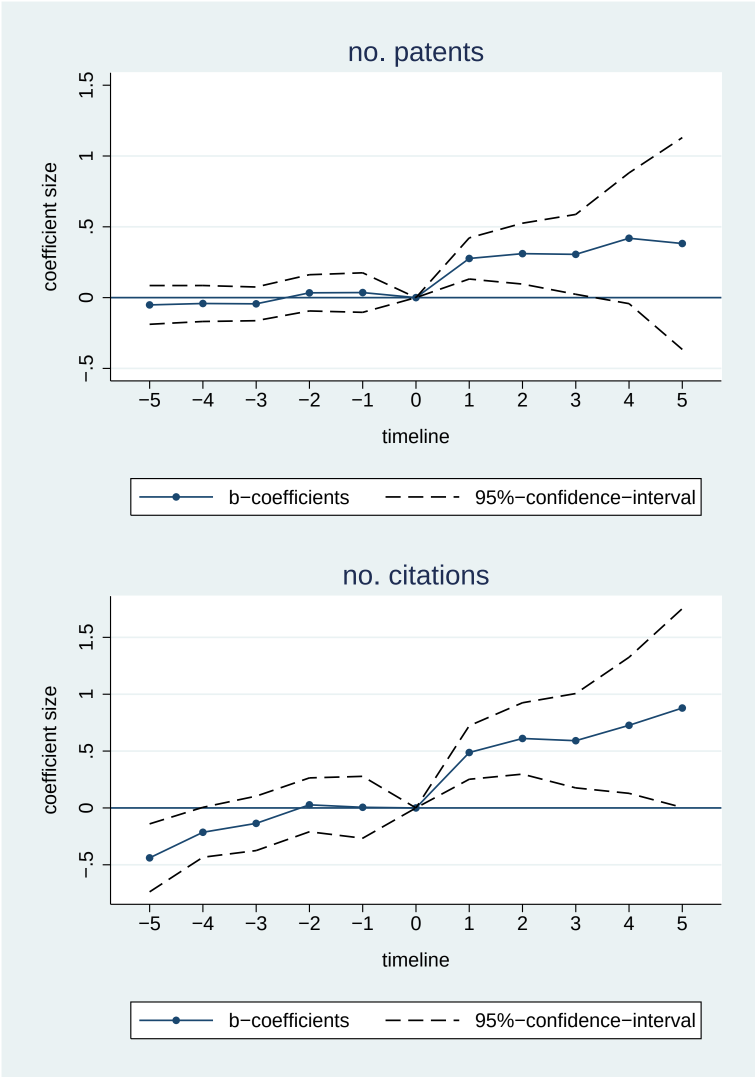

Notes: These figures illustrate the effect of a change in board independence on patenting and citations over time. For the graphs we defined dummy variables for the time firms changed from a minority of independent board members to an independent board. t_0 indicates the year of the switch and serves as the reference category. t_{n-1} indicate the years before the switch, and t_{n+1} the corresponding years after the switch. Coefficients are taken from the last regression model of section 4.3, but with the t_n dummies instead of the one dummy variable indicating a majority of independent board members.

Table 5: Independent boards and number of claims

	(a) b/se	(b) b/se	(c) b/se	(d) b/se	(e) b/se
log(total assets)	1.013***	0.380***	0.362***	0.512***	0.477***
	(0.030)	(0.099)	(0.114)	(0.132)	(0.146)
R&D	8.685***	1.326	1.362	1.486	0.580
	(0.981)	(0.950)	(1.182)	(1.300)	(1.368)
log(age)	0.146***	0.030	-0.002	-0.001	-0.097
	(0.043)	(0.059)	(0.067)	(0.083)	(0.094)
Leverage	-0.291	0.188	0.230	-0.099	0.084
	(0.228)	(0.282)	(0.322)	(0.330)	(0.414)
Cap. exp.	1.448	-0.023	0.012	0.447	0.338
	(0.881)	(0.934)	(0.989)	(1.059)	(1.183)
log(Q)	0.280***	0.110*	0.152**	0.184**	0.172**
	(0.051)	(0.062)	(0.068)	(0.078)	(0.087)
Boardsize	0.000	0.002	-0.014	-0.037	-0.036
	(0.017)	(0.022)	(0.027)	(0.028)	(0.033)
Independent board	0.501***	0.488***	0.476***	0.365**	0.359**
	(0.153)	(0.137)	(0.142)	(0.178)	(0.178)
Observations	6107	6107	4414	4414	4414
R^2	0.466	0.133	0.119	0.304	0.307
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is the logarithm of one plus the total number of claims of a patent portfolio. All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

Figure 3: Independent boards and number of claims

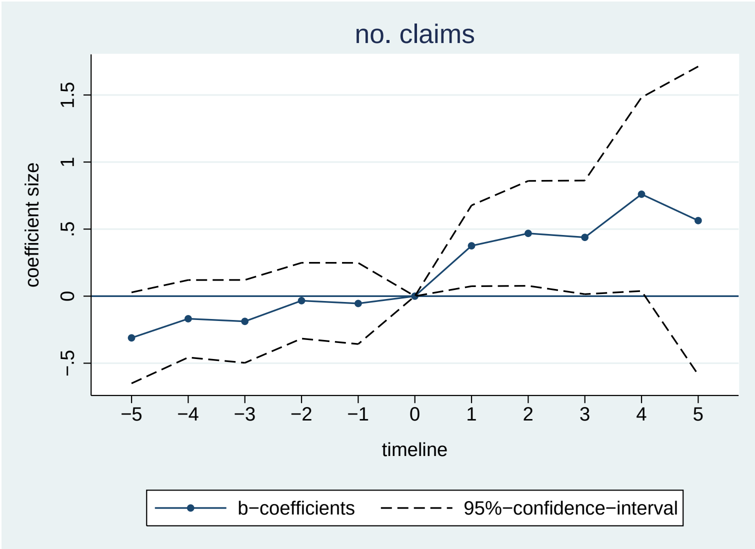

Notes: These figures illustrate the effect of a change in board independence on claims over time. For the graphs we defined dummy variables for the time firms changed from a minority of independent board members to an independent board. t_0 indicates the year of the switch and serves as the reference category. t_{n-1} indicate the years before the switch, and t_{n+1} the corresponding years after the switch. Coefficients are taken from the last regression model of section 4.3, but with the t_n dummies instead of the one dummy variable indicating a majority of independent board members.

5.2 The distribution of citations

Most recent research that uses patent data considers raw counts and total citations to the raw count patents, but less research considers the distribution of citations in careful detail. In this section we model the number of breakthrough, important, incremental, and failed inventions that a firm makes. These estimations are motivated by the argument that responding to increased oversight will increase tangible and countable but incremental patents at the expense of risky patents; such risky patents are more likely to fail completely or provide a breakthrough.

To model each of these four possible outcomes, we split the distribution into subcategories: (1) the number of patents that the firm invents that received cites within the highest percentile (top 1%) among all patents in the same 3-digit patent class and application year, (2) the number of patents that received cites within the highest centile (10%) among all patents in the same 3-digit patent class and application year but not including the top 1%, (3) the number of patents that received at least one citation (the median of the entire distribution is 0) but not including the top 10%, and (4) the number of patents that received no citation. Hence, the measures should be interpreted as 1) the number of breakthroughs, 2) the number of important patents, 3) the number of incremental patents that have small value, and 4) the number of patents that have little or no value to the firm. As an example, in the year 2000, IBM invented 4367 patents, of which 24 were in the top 1% of their field, 360 in the top 10 but not including the top 1%, 3374 with at least one cite but not in the top 1 or 10%, and 609 of which received no citations. Tables 6 to 9 present the corresponding results for each of the bins.

Consistent with the models in Tables 3 and 4 we see a positive effect of board transitions on patenting and citation rates. The estimated effect is by far the most significant and largest - from 35% to 22% - for incremental patents that received at least one citation (but not in the top 10% of the distribution), while the estimated effect on particularly successful patents (top 1% or top 10%) is very small in magnitude and significant at $p<.10$ for only two out of 10 regressions. Taking also into account that the effect on the number of unsuccessful patents (no cites) is most often statistically insignificant, the evidence is consistent with the argument that firms focus on less risky opportunities when the board becomes independent. Inclusion of a measure of backward citations weakens these effects further, implying that the increase in citations is mediated by movement of the firm into more crowded areas of technological search (models not shown but available from first author). In other words, the increase in citations may not correspond to an increase in patent value, rather, it may be an artifact of the exploitation strategy.

5.3 Discussion

The results consistently describe a shift towards innovative exploitation for firms that transition to an independent board. They do not provide consistent evidence for any influence on exploration; there are positive but weakly significant increases in the tail of the citation distribution and no impact on patenting in new classes. Here we summarize robustness checks reported in the Appendix, explore potential mechanisms that could accomplish the shift towards exploitation, and discuss why independent boards (and managers) might have less influence on exploration.

Firms which transition to independent boards patent more; this raises the concern that the increased backward and self citation results might simply be artifacts of the increased patenting. To rule out this possibility, we estimate regressions of backward and self-citations per patent. As can be seen in Tables **??** and **??** in the Appendix, the proportion of backward and self-citations also increases for firms which transition to independent boards. Effects sizes range from 22-18% for backward citations and 16-14% for

Table 6: The number of breakthrough inventions: independent boards and top 1% patents

	(a) b/se	(b) b/se	(c) b/se	(d) b/se	(e) b/se
log(total assets)	0.166***	0.037**	0.056***	0.054***	0.037*
	(0.008)	(0.016)	(0.015)	(0.019)	(0.020)
R&D	0.724***	-0.092	-0.060	-0.045	-0.102
	(0.097)	(0.136)	(0.223)	(0.290)	(0.364)
log(age)	0.036***	0.013	0.004	-0.002	-0.008
	(0.007)	(0.010)	(0.008)	(0.011)	(0.013)
Leverage	-0.198***	-0.049	-0.113***	-0.145**	-0.110
	(0.035)	(0.042)	(0.043)	(0.058)	(0.070)
Cap. exp.	0.489***	-0.109	-0.094	-0.110	-0.049
	(0.150)	(0.118)	(0.107)	(0.130)	(0.152)
log(Q)	0.031***	-0.000	-0.015	-0.021	-0.023
	(0.009)	(0.011)	(0.012)	(0.015)	(0.015)
Boardsize	0.000	0.004	0.000	-0.002	0.002
	(0.003)	(0.003)	(0.004)	(0.004)	(0.005)
Independent board	0.027	0.043*	0.030	0.045*	0.041
	(0.027)	(0.024)	(0.025)	(0.027)	(0.026)
Observations	6107	6107	4414	4414	4414
R^2	0.312	0.009	0.014	0.179	0.182
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is the logarithm of one plus the number of patents that fall in the top 1% percentile of the citation distribution within patent class and application year. All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

Table 7: The number of important inventions: independent boards and top 2% to 10% patents

	(a) b/se	(b) b/se	(c) b/se	(d) b/se	(e) b/se
log(total assets)	0.389***	0.103***	0.113***	0.109***	0.058
	(0.014)	(0.030)	(0.032)	(0.035)	(0.038)
R&D	2.283***	0.218	-0.072	-0.120	-0.232
	(0.265)	(0.222)	(0.344)	(0.420)	(0.553)
log(age)	0.072***	0.040**	0.027*	0.030	0.035
	(0.015)	(0.017)	(0.014)	(0.019)	(0.022)
Leverage	-0.300***	0.049	-0.046	-0.079	-0.064
	(0.076)	(0.072)	(0.083)	(0.095)	(0.110)
Cap. exp.	0.997***	-0.236	-0.196	-0.228	-0.068
	(0.330)	(0.207)	(0.219)	(0.228)	(0.277)
log(Q)	0.101***	0.030	0.026	0.034	0.028
	(0.018)	(0.019)	(0.022)	(0.030)	(0.028)
Boardsize	0.003	0.004	-0.004	-0.007	-0.001
	(0.007)	(0.006)	(0.007)	(0.008)	(0.009)
Independent board	0.069	0.064*	0.051	0.062	0.061
	(0.054)	(0.039)	(0.040)	(0.055)	(0.054)
Observations	6107	6107	4414	4414	4414
R^2	0.407	0.017	0.021	0.208	0.214
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is the logarithm of one plus the number of patents that fall in the top 10% centile of the citation distribution within patent class and application year (excluding the top 1%). All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

Table 8: The number of incremental inventions: independent boards and cited patents, not in top 10%

	(a)	(b)	(c)	(d)	(e)
	b/se	b/se	b/se	b/se	b/se
log(total assets)	0.678***	0.268***	0.227***	0.251***	0.316***
	(0.018)	(0.055)	(0.057)	(0.062)	(0.069)
R&D	4.879***	1.123**	0.820	0.857	1.210
	(0.497)	(0.459)	(0.566)	(0.615)	(0.755)
log(age)	0.097***	0.045	0.001	0.004	-0.024
	(0.023)	(0.034)	(0.032)	(0.041)	(0.048)
Leverage	-0.433***	-0.045	-0.031	-0.103	-0.064
	(0.116)	(0.148)	(0.157)	(0.162)	(0.189)
Cap. exp.	2.093***	0.284	0.407	0.553	0.544
	(0.481)	(0.401)	(0.419)	(0.455)	(0.519)
log(Q)	0.183***	0.091***	0.103***	0.097***	0.090**
	(0.027)	(0.031)	(0.032)	(0.037)	(0.039)
Boardsize	0.004	0.009	-0.003	-0.016	-0.014
	(0.009)	(0.012)	(0.014)	(0.014)	(0.016)
Independent board	0.348***	0.339***	0.260***	0.229***	0.220***
	(0.076)	(0.067)	(0.067)	(0.073)	(0.074)
Observations	6107	6107	4414	4414	4414
R^2	0.536	0.248	0.207	0.416	0.421
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is the logarithm of one plus the number of patents that are cited but do not fall in the top 10% of the citation distribution. All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

Table 9: The number of failed inventions: independent boards and patents without citations

	(a)	(b)	(c)	(d)	(e)
	b/se	b/se	b/se	b/se	b/se
log(total assets)	0.635***	0.223***	0.278***	0.390***	0.299***
	(0.015)	(0.068)	(0.075)	(0.081)	(0.095)
R&D	3.953***	0.184	-0.206	-0.677	-1.452
	(0.433)	(0.557)	(0.868)	(0.990)	(1.322)
log(age)	0.085***	0.071	-0.011	-0.004	-0.025
	(0.019)	(0.047)	(0.039)	(0.051)	(0.066)
Leverage	-0.418***	-0.233	-0.299	-0.492**	-0.273
	(0.103)	(0.175)	(0.209)	(0.217)	(0.244)
Cap. exp.	1.043***	-0.264	-0.304	-0.338	-0.343
	(0.400)	(0.472)	(0.513)	(0.532)	(0.568)
log(Q)	0.114***	0.003	-0.012	0.006	-0.025
	(0.023)	(0.035)	(0.040)	(0.044)	(0.043)
Boardsize	0.019**	0.028*	0.010	-0.000	0.003
	(0.008)	(0.015)	(0.015)	(0.015)	(0.018)
Independent board	0.167**	0.106	0.077	0.099	0.098
	(0.071)	(0.089)	(0.090)	(0.094)	(0.091)
Observations	6107	6107	4414	4414	4414
R^2	0.510	0.045	0.040	0.323	0.332
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Firm fixed effects	No	Yes	Yes	Yes	Yes
Trend control	No	No	No	Yes	Yes
Controls * post-SOX	No	No	No	No	Yes

Notes: The dependent variable is the logarithm of one plus the number of patents that are not cited. All explanatory variables are lagged by one period. Specification (a) includes untabulated 3-digit SIC industry dummies and a dummy that marks all treated firms. Independent board is a dummy that indicates firms after they switched from a minority of independent board members to a majority of independent board members in 2001 or later. Control variables are defined in section 4.2. Heteroskedasticity-robust standard errors that account for autocorrelation at the firm level are reported in parentheses. Coefficients: *** Significant at 1%, ** Significant at 5% level, * Significant at 10% level.

self-citations. We also investigated the coefficient of variance of citations to firms that undergo the transition to independent boards. While the results were not significant on a yearly basis, an aggregation of the four years following the transition demonstrated a significant decrease; consistent with a shift to exploitation and thinning of the tails, citations to firms with independent boards become less variable after the transition. We found no evidence that the transition to independent boards influences innovative efficiency (Cohen, Diether, and Malloy, 2013); there is no statistically significant result for the regression of patents per R&D investment and citations per R&D are positive but lose significance in the trend models.

Table **??** shows that our results are more pronounced for firms with high research and development spending and stock. Board independence appears to have a stronger impact on firms for which innovation is more important (high R&D), probably because for those firms the tension between exploration/exploitation is more significant and boards need to be more concerned about their innovation strategies. If innovation is less important (low R&D), there is probably less board involvement in innovation and thus our results are less pronounced.

As proposed in the introduction, multiple mechanisms could cause a firm whose board becomes independent to shift towards exploitation. For example, managers may shirk less and work harder in response to greater oversight, take less risk out of career concerns, respond to advice, or search less because they fear an independent board will constrain future flexibility. Ruling out one versus another mechanism empirically remains difficult, as they imply similar predictions and probably co-exist in practice. Nonetheless, the split sample tests described below remain consistent with models where greater oversight results in increased effort and risk aversion.

We split the sample into firms with high and low managerial entrenchment, using the index of Bebchuk, Cohen, and Ferrell (2009). This entrenchment or "e-index" indicates how many corporate governance provisions are in place that shield a manager from getting fired, e.g. poison pills, golden parachutes, and supermajority requirements for mergers and charter amendments. [12] Table **??** shows that the effects of independent boards are consistently stronger for firms with high managerial entrenchment. Managers in firms with low entrenchment index are already subject to career concerns and takeover pressures, even before the board becomes independent. Therefore, the transition to an independent board does not have much impact on these managers. It is for entrenched managers that transitioning into independent boards can trigger career concerns and greater exploitation.

[12]The E-Index is given by Bebchuk, Cohen and Ferrell (2009) for all equal years and is fairly stable over time. In order to keep the sample size as large as possible we imputed with the lagged value where the E-Index was missing; if the lagged value was missing we took the forward value.

The claims results provide additional corroboration for an oversight mechanism and potentially explain the quick response illustrated in the graphs following the year of transition. Research takes time: programs must be funded, staffed, and executed, after which any successful results must be patented. Even though the measures use the year of application (rather than the year of grant, which can be 1-4 years later typically), it is still surprising that most of the graphs show an immediate response following the year of transition. It is possible though less likely that a new R&D strategy could effect such immediate impact; a more plausible explanation is that the patenting processes changed, and in particular, that the firm's engineers and lawyers looked more carefully for patentable technology within the firm's extant portfolio. This would be reflected in the immediate increase of patents, claims, and even citations. Indeed, when claims are included as a mediating variable, most of the results weaken (in particular, the occasionally significant results in the tails of the citation distribution become consistently insignifcant).

Most likely many mechanisms play overlapping and possibly complementary roles in explaining the shift towards exploitation. The aggregate evidence is most consistent, however, with increased oversight and career concerns mechanisms; the shift towards exploitation when a firm's board becomes independent is most likely due to a combination of greater managerial effort and an increased aversion to innovative risk.

We do not see an increase in exploitation at the expense of exploration, indeed, we see almost no impact of board transition on exploration at all. This might imply that exploration and exploitation are not ends of a continuum; rather, they could be orthogonal and possibly complementary strategies, especially for firms with large research portfolios and relatively independent research teams. This raises the possibility that the contradictory results in the literature on governance and innovation are caused by the conflation of two successful innovation strategies. For example, this research and that of Lerner, Stromberg, and Sorenson (2011) illustrate an exploitation strategy; in contrast, Chemmanur and Tian (2015) illustrate an exploration startegy (as evidenced by the increase in citation variance). Both probably lead to an increase in firm value, though the mechanisms, risk, and payoff vary greatly.

It is likely, however, that exploitation crowds out exploration in general (March 1991). To begin with, the first order goal of boards and managers is value creation for shareholders, not scientific or technical discovery. It is also likely that boards and managers have less fundamental ability to influence exploration; it is simply easier to manage and organize extant and tangible possibilities than to encourage breakthrough creativity. By its nature, exploration also takes longer to measure and appropriate. To explore this hypothesis with the current dataset, we ran lagged outcomes of later years, but found no significant differences. Such delayed and uncertain payback is less attractive to managers and boards. It may also be more difficult to appropriate explorative innovation, as it requires a stable workforce and patient investment; this hypothesis could be explored with patent citation models.

Some of the conflicting results in the governance and innovation literature might be profitably revisited with more nuanced measures and an effort to better identify the particular mechanisms and strategies that result in increased patenting and citations. For example, are the non-monotonic results of Supra, Subramanian, and Subramanian due to differing search strategies? Does higher citation come from focus and exploitation on the strong governance end and search and exploration on the weak governance end? Similarly, does the weak governance result come from the tails of the distribution and greater volatility in patent citations, consistent with an exploration strategy?

6 Conclusion

We proposed that firms which undergo a transition to more independent boards increase exploitation of previously successful areas of expertise. We argued that the shift towards exploitation results from stronger board oversight which increases both managerial effort and risk aversion. Supporting evidence came from the regulatory changes of Sarbanes-Oxley; firms that transition to more independent boards invent more but less creative patents. On average the patents receive more citations, however, those citations occur to incremental patents in the middle of the distribution, and not to breakthrough or completely failed patents. Furthermore, the increase in cites is due partly to an increase in claims within each patent as well as movement into more crowded areas of technology. This implies that the increase in citations is due to a more thorough patenting of extant portfolios and an artifact of the citation norms in more crowded fields. Firms that transition also patent more heavily in technology classes of their current portfolio; they do not patent more on new classes. The effects are more pronounced for research intensive firms and those that score high on measures of managerial entrenchment. Speaking to the larger literature on governance and innovation, our results indicate that strengthened governance improves innovation performance along existing trajectories, without harming the probability of a breakthrough.

We offered more nuanced and easily calculated patent measures that enable greater insight into the search and innovation process. These measures highlight the importance of differentiating between the greater incremental output of exploitation and the riskier breakthrough output of exploration. The results indicate that firms can increase their patent counts - and even future citations to those patents - through exploitation of their existing portfolios. Further work should differentiate, both theoretically and empirically, between greater and focused effort and riskier search; it should not assume that an increase in patent counts or citations implies an increase in risk-taking and creativity.

Independent boards appear to move firms towards innovative exploitation and have little impact on exploration, but what is best for performance? Other research has found mixed evidence for the impact of independent boards on overall performance (see e.g. Duchin, Matsusaka, and Oguzhan, 2010; Nguyen and Nielsen, 2010; Adams, Hermalin, and Weisbach, 2010; Duchin, Matsusaka, and Oguzhan, 2010). Lack of exploration may cause long term obsolescence and competency traps, but where is the optimal tradeoff? Can large and diverse firms avoid the stark tradeoff, by developing portfolios that simultaneously explore and exploit? These are topics for future research.