

Intellectual property protection lost and competition: An examination using large language models

Utku U. Acikalin^a, Tolga Caskurlu^b, Gerard Hoberg^c, Gordon M. Phillips^{d, 1}

^a Cornell University, Ithaca, NY, USA

^b University of Amsterdam, Amsterdam, the Netherlands

^c University of Southern California Marshall School of Business, Los Angeles, CA, USA

^d Tuck School of Business, Dartmouth College and NBER, Hanover, NH, USA

ARTICLE INFO

JEL classification:

O31
O34
D43
F13

Keywords:

Patents
Intellectual property protection
Innovation
Competition
Litigation
Alice

ABSTRACT

We examine the impact of lost intellectual property protection on innovation, competition, firm performance, and valuation. We consider firms whose ability to protect intellectual property (IP) using patents is weakened following a major Supreme Court decision. We use large language models (LLM) to identify firms' patent portfolios' exposure to this decision and find an unequal impact. Large firms gain and small firms lose. Large impacted firms benefit as their sales growth increases and their exposure to lawsuits decreases. Small impacted firms lose as they face increased competition, product-market encroachment, and lower profits and valuations. They increase R&D and nondisclosure agreements.

What though the field be lost? All is not lost

[Paradise Lost, John Milton, 1674.]

1. Introduction

Intellectual property protection is at the core of innovation and competition policy. Economic and legal scholars have debated extensively

whether intellectual property (IP) protection increases firm incentives to innovate and conduct R&D. Many articles argue that patents stifle innovation, as [Boldrin and Levine \(2013\)](#) summarize in their survey. Examining 60 countries over 150 years, [Lerner \(2009\)](#) also finds limited benefits of increasing patent protection for domestic firms. He finds decreased domestic patenting following increased IP protection but increases in foreign patenting, suggesting foreign competitors enter with the increased protection. Yet, not all studies agree that IP protection is harmful to innovation. [Budish et al. \(2015\)](#) models how the length

¹ Dimitris Papanikolaou was the editor for this article. We thank the editor, two anonymous referees, Raj Aggarwal, Janet Gao, Umit Gurun, Josh Lerner, Alexander Ljungqvist, Tim Loughran, N.R. Prabhala, Carl Shapiro, Andrei Shleifer, Richard Thakor, Vladimir Vladimirov and seminar participants at the AEA, Alberta, Amsterdam, Bogazici, Columbia, European Finance Association, Harvard, Johns Hopkins, Minnesota, Ohio State, NBER Innovation and NBER Big Data conferences, Notre Dame, Peking University, RCFS Winter Conference, Stanford, UCLA, Università Cattolica, University of Pennsylvania, and Yeshiva University for helpful comments and discussions. We gratefully acknowledge TÜBİTAK ULAKBİM, the High Performance and Grid Computing Center for providing access to TRUBA resources, and the Dutch National Computer Facilities Program for providing access to the Snellius supercomputers.

of patent protection should optimally increase for long-term innovation when commercialization occurs later in order to increase incentives for innovation.

A natural question is how strong to make IP protection? The theories behind optimal IP protection begin with Nordhaus (1969). In that study, the debate is about the trade-off between issuing patents to encourage innovation and the cost of reducing subsequent competition resulting from giving the patentee a local monopoly over the life of the patent. There are also issues with the scope of the patent. If patent protection is too broad, new entrants and new innovations may be discouraged as the protected scope of existing innovations might imply high entry barriers. Monopoly profits that arise from IP protection would also be high, harming consumers. If too weak, then firms would be discouraged from engaging in costly innovation, as the innovation would be potentially available to all to copy.

Our study examines the consequences of weakened IP protection across multiple industries in a setting that shocked both existing patents and also incentives for future innovation and patenting in the U.S. in multiple patent categories — thus changing the dynamics of innovation and competition in large technology areas. We examine entire firms' product spaces when multiple firms' existing patents in the business methods broad area are exogenously weakened by the Alice Corp v. CLS Bank International Supreme Court decision (Alice, henceforth). We examine the impact of lost intellectual property protection on a wide array of future firm decisions, including firm innovation, growth and valuation. We also examine competitive entry through venture capital financed entry and patent lawsuits.

The Alice decision revoked patent protection on business method patents whose fundamental idea is considered abstract, with a transformation that is not novel. As part of this decision, the Supreme Court also ruled that media and systems claims are similar to the business methods claims, and they are also patent ineligible. The outcome of this decision was in doubt, given prior court decisions, and we show that it had a large impact after the ruling. In the next section, we summarize the extensive lower court disagreements on this case preceding the final Supreme Court ruling.²

The Alice decision impacted multiple industries, including data processing, software, finance, games, and measuring or testing in microbiology and enzymology. We document that Alice has had a large impact on patent rejections, and it led to significant decreases in patenting in exposed areas after 2014. Even post-Alice, there was considerable uncertainty about whether a particular patent sufficiently transforms an abstract idea enough to make it patent-eligible. Rejections based on Alice represented approximately 20% of the patent subject matter rejections overall in 2015 and 2016. For example, in the commerce and data processing methods industry, 36.2% of patents filed in 2013 were rejected, citing Alice. We also conservatively estimate Alice's impact on future patents and find that Alice resulted in at least 6898 fewer patents being granted per year after Alice.³

While the decision had a large ex post impact on patenting, there was and still exists uncertainty about whether an existing or proposed patent transforms an idea sufficiently to be granted patent protection.

² Indeed, when the case was being considered at the Supreme Court, there were extensive Amicus briefs filed on both sides. Amicus briefs filed with the Supreme Court in support of CLS Bank included briefs filed by Google, Amazon, Dell, LinkedIn, Verizon, Microsoft, Checkpoint Software, the Software Freedom Law Center and Opensource Initiative, and prominent lawyers and economists. There were over 20 Amicus briefs in support of Alice Corp., including Advanced Biological Laboratories, IBM, Trading Technologies International, Inc., and prominent lawyers and economists. See http://www.alicecorp.com/fs_patents.html.

³ We also consider the potential effect other impactful Supreme Court cases that pre-date Alice: *Bilski v. Kappos* (561 U.S. 593, 2010), *Mayo v. Prometheus* (566 U.S. 66, 2012), and *Association for Molecular Pathology v. Myriad Genetics, Inc.* (569 U.S. 576, 2012). See Section 3.8.

Given the uncertain impact on each patent, we use the ModernBERT language model on patent textual corpora to assess whether a firm's patents would be overturned if reexamined post-Alice. Many legal scholars have written about the Alice decision and the difficulties of measuring and deciding whether there is a sufficient transformation of an abstract idea to warrant a patent and whether existing patents still provide intellectual property protection.⁴

We examine all patents in Alice-impacted areas that were granted by 2014 (the date of the Alice decision). Some of these patents are likely to fail to provide protection if challenged in a court in the post-Alice period. This is a challenging computational task as there are more than 3.8 million patents granted between 1994 and 2014. We thus focus on the patents with the same primary Cooperative Patent Classification (CPC) as those rejected by the United States Patent and Trademark Office (USPTO) per the Supreme Court's Alice criteria. Given the uncertainty about whether a given patent will be rejected, we use an LLM to gauge each patent's textual semantic similarity to patents previously rejected under Alice.

We use a deep learning-based LLM called ModernBERT to predict the likelihood that each of the pre-Alice granted patents in the sample may become patent-ineligible due to the Alice decision. ModernBERT is a recent transformer model that is an improvement over BERT as it allows long texts (i.e., 8192 tokens vs. 512 tokens).⁵ The breakthrough innovation of models like ModernBERT is that they process words in relation to all the other words in a sentence rather than one-by-one in order or in a fixed-sized sliding window approach. These models can examine the full context of a word by looking at the words that come before and after it.

We find a large impact of Alice on future patenting. Our ModernBERT framework indicates that 131,548 of the 642,678 patents (20.47% of our sample of impacted industries) are likely to lose protection if subjected to judicial scrutiny.⁶ We begin by verifying that ex post patenting declines significantly for affected firms overall. We then split our public firm sample by size as we hypothesize that smaller public firms may be hurt more as they have fewer resources (managerial, financial and organizational) to defend their product spaces.

Applying an entropy balancing technique to achieve covariate balance for the sample, we find that treated small firms face increased entry by VC-funded startups and complain more about competition. In response, small firms increase their R&D expenditures, a pattern consistent with efforts to replenish weakened innovative portfolios and to escape heightened competition through product differentiation. In contrast, we find more muted effects among large firms. We examine performance in the form of sales growth, profitability, and market valuations, and we continue to find an unequal impact. Large firms gain, and small firms lose. The myriad of treatment effects we find for small versus large firms are statistically significant in formal difference tests. We also examine pre-trends and formal structural break tests, which further illustrate the differential outcomes post Alice.

Our main contribution is to examine who gains and loses after a change in intellectual property protection that impacts whole areas of technology indefinitely for all public firms, both large and small. Thus our setting is about the broad inability to patent (present and future) in treated technology areas, which is different from previous studies using random assignments to examine the loss of protection for individual patents. Previous studies just show how firm's decisions vary when a

⁴ See Kesan and Wang (2020) and Lim (2020) for an extensive discussion of these debates and issues. These difficulties and the impact of Alice gave rise to U.S. Senate subcommittee hearings promoted in 2019 on potential revisions to strengthen intellectual property law in the "Stronger Patents Act".

⁵ BERT was released by Google in 2019 and achieves state-of-the-art performance on various Natural Language Processing (NLP) tasks (Devlin et al., 2019). The BERT model is also used in Google search queries.

⁶ Table IA4 provides detailed patent statistics on varying score thresholds and the frequencies of Cooperative Patent Classification (CPC) codes.

rival loses patent protection and do not address an overall loss of patent protection for everyone in the technological area.⁷ Understanding the consequences of invalidating protection for entire areas indefinitely is important as it indicates the likely consequences of future major policy changes that broadly weaken protection.

We show that differential losses for small firms are related to changes in competition relating to decreased IP protection. These small firms face increased competition on a number of different measures. Both large and small firms face increased venture capital-financed entry into their product space, but this entry is more severe for small firms. Small firms also complain more about increased competition. We also find, using textual analysis of 10-Ks, that small firms resort to non-disclosure agreements as an alternative way to protect IP post-Alice. Thus, small firms turn to increased secrecy to defend new IP in the face of lost patent protection. This highlights the importance of disclosure, as noted by [Sampat and Williams \(2019\)](#) for technologies that transition from patentable to unpatentable.

Further indicating non-uniform competitive effects, we find evidence that barriers to entry explain the superior outcomes for large firms. In subsample tests, we find that these better outcomes occur for firms with large stocks of tangible assets and for firms operating in areas where assets are difficult to redeploy. In contrast, we find that access to more liquidity is unlikely to be a factor. Overall, the results are consistent with large firms having better ex post outcomes, given that their tangible assets act as barriers to entry.

We next examine patent infringement litigation. [Lemley and Zyontz \(2021\)](#) descriptively illustrates using filed lawsuits that companies were more likely to lose their patent protection post-Alice than were non-producing entities (NPEs), also called patent trolls.⁸ Even though legal data is relatively noisy and we view these tests to be suggestive, our difference-in-difference tests suggest that large firms likely experience fewer costly infringement claims. This finding is intuitive as firms would be less likely to sue a deep-pocketed firm when the validity of the patents is questionable. We find mostly insignificant or weak results for small firms.

While we do not conduct a full welfare analysis, the results support the conclusion that patent protection benefits small firms. The withdrawal of patent protection hurts small firms while increasing the overall trend of production being concentrated in larger firms. Our results are consistent with large firms having more resources – technological, financial and managerial – which allows them to protect their market position. We conclude that patent protection is particularly important for small firms competing with larger firms.

In a review focusing on the role of patents on the incentives to innovate, [Williams \(2017\)](#) proposes three focal channels: disclosure, patent strength, and pre-existing patents. The Alice decision has a direct impact on the latter two: it reduces patent strength both for future and existing patents in entire technological areas that were previously protected. Consistent with the importance of protection for future innovation, we document a widespread reduction in patenting following the Alice decision. Yet we do not find corresponding reductions in R&D as smaller firms significantly increase R&D. Disclosure, the 1st channel, was also impacted as small firms shift from disclosed patents to increased use of non-disclosure agreements.⁹

⁷ [Galasso and Schankerman \(2015\)](#) showed that when larger firms' patents were invalidated, small firms increased innovation. [Farre-Mensa et al. \(2020\)](#) shows that small firms gain from patent protection beyond the value of the idea itself using an instrument of random assignment of patent examiners from [Sampat and Williams \(2019\)](#). They show that small firms gain access to increased funding post-patent.

⁸ One limitation of filed lawsuit analysis is that the primary impact of Alice might be on the number of patent disputes that were *not* filed, as lawsuits are not cost-effective under lost protection.

⁹ Although few existing articles focus on this topic, studies examining innovative secrecy include [Boot and Vladimirov \(2025\)](#) (theory and cooperation) and [Breuer et al. \(2019\)](#) (EU reporting regulations).

Our paper is the first paper in finance and economics to show that large language models can be used to study the economic impact of legal decisions, thus contributing to methodology in the law and finance literature. We test our Alice score's ability to predict Alice-related cases in USPTO post-grant reviews. Our findings demonstrate high predictive accuracy, making it a valuable tool for legal professionals in assessing patent validity.

2. Innovation and Alice v. CLS bank international

There is a substantial debate on how strong to make IP protection. The general academic consensus is summarized by [Boldrin and Levine \(2013\)](#), who state that there is no empirical evidence that patents increase innovation and productivity. They advocate for abolishing patents entirely and using other legislative instruments to increase innovation. [Galasso and Schankerman \(2015\)](#) document that patent invalidation of large patentees triggers follow on innovation by smaller firms. However, these were exogenous invalidations of particular existing patents and these tests are not about indefinite forward-looking changes to entire patent areas as is the case for Alice. [Lerner \(2002\)](#)'s comprehensive study of over 60 countries used patent law changes and showed some benefits of strengthening patent protection for countries with initially weaker patent protection. There is also evidence (see [Budish et al. \(2015\)](#) for example, using cancer clinical trials) that maintaining incentives to engage in innovation is important if the ideas take a long time and are costly to develop.

We examine firm R&D and firm outcomes, including changes in sales, profits, and market valuations, and the impact on competition between firms after the landmark Supreme Court case, *Alice Corp. v. CLS Bank International*, 573 U.S. 208 (2014). This decision has impacted large industry areas — and key for us, not just the finance area or a subset of an area. These areas previously had substantial patenting activity. We estimate that the areas impacted had over 600,000 existing patents. [Kesan and Wang \(2020\)](#) reviews the impact of Alice and documents large decreases in 11 patent categories, including bioinformatics, business methods, business methods of finance, business methods of e-commerce, software (in general), databases and file management, cryptography and security, telemetry and code generation, digital cameras, computer networks, and digital and optical communications. They showed significant rejections of patents under Alice based on whether the proposed invention sufficiently transforms an abstract idea or natural law.¹⁰

2.1. Legal background of the Alice case

In *Alice Corp. v. CLS Bank International*, 573 U.S. 208 (2014), the United States Supreme Court issued a landmark decision that reshaped the criteria for patent eligibility in broad business methods areas. The specific patent the court ruled on was using computer-implemented methods for settling financial transactions. This was merely a codification of an abstract idea, and thus fell outside the scope of patent protection. The Court issued a broader decision that has had an impact beyond patents on financial transactions, as it ruled that implementing a known abstract idea on a computer does not render it patentable.

The litigation involved considerable judicial disagreement. After a district court held the patents invalid, the case reached the Court of Appeals for the Federal Circuit (CAFC). In this court, a randomly assigned three-judge panel could not unanimously decide on the case, and the panel reversed the district court decision with a majority opinion.¹¹

¹⁰ Section 101 of the Patent Statute specifies four categories of invention that are patent-eligible: process, machine, manufacture, and composition of matter. However, there are three court-made exclusions to these categories that limit patent eligibility: laws of nature, natural phenomena, and abstract ideas.

¹¹ *CLS Bank Int'l v. Alice Corp. Pty. Ltd.*, 685 F.3d 1341, 1356 (Fed. Cir. 2012).

However, given the case's complexity and its importance for the whole industry, the CAFC vacated the panel's opinion and decided to hear the case in a full session of all ten judges.¹² The uncertainty in the en banc session was similar to that of the three-judge panel as there again was a split decision. Five of ten judges upheld the district court's decision that Alice's systems claims were not patent-eligible, and five disagreed. Seven of the ten judges upheld the district court's decision that Alice's method claims were not patent-eligible. However, these seven judges reached their opinions for different reasons. Overall, the judges could not agree on a single standard by which to determine whether a computer-implemented invention is a patent-ineligible abstract idea.

After the deep division in the CAFC, the Supreme Court of the US granted certiorari and affirmed the en banc decision of the Federal Circuit Court of Appeals. The Court held a two-step framework for determining the patent eligibility of applications that would be applied to claims of abstract ideas. The Court decided that the claims in Alice patents cover an abstract idea and the proposed method claims fail to transform the abstract idea into a patent-eligible invention. The Court also ruled that the media and systems claims are similar to the methods claim and that they are also patent ineligible.¹³

The Alice decision had a significant stock market reaction in many patenting areas. We compute abnormal returns using the market model and the Fama–French 3-factor model using the event study model provided on the WRDS platform. We then examine abnormal announcement returns for firms with high versus low Alice exposure surrounding three key Supreme Court dates: 12/6/2013 (date case granted), 3/31/2014 (date case argued), and 6/19/2014 (date case decided).¹⁴

We find that firms with one standard deviation higher exposure to Alice experience announcement returns of roughly -0.1% on the date the case was decided (at the 1% level) and similar but less significant returns on the date the case was granted. Although these average treatment effects are modest in magnitude, they increase to -0.4% for the top 5% of firms with the highest Alice score on the date the case was decided and -0.2% to -0.3% on the date the case was granted. These results indicate that the outcome of the case was unexpected and its impact was significantly negative for treated firms. Yet uncertainty was high and the total impact was much larger and realized over time as we show later.

2.2. Consequences of lost IP protection

The Alice case had a large impact on ex post patenting. In Table 1, we present statistics for the number of pre-Alice patent applications that were rejected by the USPTO, where the examiners cited Alice as the reason for rejection. Over 33,700 distinct patent applications made prior to Alice have been rejected in the 3 years post-Alice by examiners citing Alice. These rejected patents cover 5831 distinct CPC Subgroups (out of 126,540 total), 919 Groups, 283 Classes, and 8 CPC Sections.¹⁵ Rejection data from Lu et al. (2017) are available only through 2016; thus, rejection ratios for 2017 are coded as NA. The last column *Change* reports the percentage difference between the number of applications in 2013 and the average number filed during 2015–2017.

Kesan and Wang (2020) document that about 17.9% of office action final decisions in bioinformatics were rejected based on section 101 before Alice was decided. This rate increased to 72.4% of the rejections of applications filed before Alice but decided afterward, and 72.8% of applications filed after Alice. Substantial increases in §101 rejections also occurred in computer networks, graphical user interfaces,

document processing, and cryptography/security. Patent application volumes declined sharply in the aftermath of Alice, falling by 12%–31% across categories. For example, patent applications in the business methods area decreased 29.5%.¹⁶

While Alice had a large impact on patenting, the Supreme Court left substantial ambiguity about whether an individual patent transformed abstract ideas sufficiently to make them patent-eligible. The court did not define “abstract”, and the court did not define how to decide whether an abstract idea has been transformed sufficiently into an inventive concept by including additional limitations to the patent claim, which form the basis for rendering claims eligible for patent protection. Given this uncertainty about whether a given patent would be rejected because of Alice, we use a deep learning-based language model called ModernBERT to predict the likelihood that each of the pre-Alice granted patents in our sample might become ineligible due to the Alice decision. We then study ex post firm decisions and outcomes based on this predicted likelihood.

We split firms by large and small size as we hypothesize that smaller firms might be hurt more by lost patent protection, as they have fewer resources (managerial, financial, and organizational) to defend their product spaces, while large firms will be the leading firms with more resources who can defend their product area. The differential impact on innovation by leading vs. laggard firms has been modeled, for example, by Aghion et al. (2005). In our setting, given the large increase in competition post-Alice, we expect firms left behind will innovate more in order to escape competition from other firms. Larger firms are predicted to behave more like in Schumpeter (1912) and innovate less, as laggards will perform most innovation with low initial profits, or smaller firms will look to enter.

We thus examine firm R&D and performance outcomes, including changes in sales, profits, and market valuations, and the impact on competition between firms. While we hypothesize that the impact of the loss of IP protection may be negative for affected firms, such an unconditional prediction is not clear, given predicted differences in impact for firms with different abilities and innovative resources. We thus focus on testing predictions separately for small and large firms and also on their differences. Large firms might benefit from losses in IP protection in their sector, for example, as they may be able to adopt new ideas without paying the firms that originally created the ideas. These firms might see decreases in the competitive threats they face. Finally, we predict that firms might seek alternative ways to increase secrecy and protect IP after patent protection is lost.

3. Data and methods

We first train and validate a model of Alice's impact on lost IP protection at the patent level using the ModernBERT large language model and patent text. We then aggregate the patent-level impact to construct the firm-level Alice treatment variable central to our analysis. Finally, we describe the sample and present summary statistics.

3.1. Experimental challenges

Our experiment requires identifying patents granted in the pre-Alice period that would lose protection if tested in court post-Alice. Only a limited number of patents have been reexamined for post-Alice viability, as holders must sue for infringement and risk invalidation.¹⁷ Given that more than 3.8 million patents were granted between 06/19/1994

¹² *CLS Bank Int'l v. Alice Corp. Pty. Ltd.*, 484 Fed. Appx. 559 (Fed. Cir. 2012), *CLS Bank Int'l v. Alice Corp. Pty. Ltd.*, 717 F.3d 1269, 1273 (2013).

¹³ *Alice Corp. Pty. Ltd. v. CLS Bank Int'l*, 134 S. Ct. 2347, 2354 (2014).

¹⁴ Table IA2 reports the results.

¹⁵ Table IA5 provides CPC groups that are impacted by Alice (Panel A) and the industries they belong to (Panel B).

¹⁶ Kesan and Wang (2020) also shows using a difference-in-difference regression that section 101 Alice rejections increased significantly in 11 different patent categories.

¹⁷ For example, in *BlackBerry Limited v. Facebook Inc.* (C.D. Ca. Oct. 1, 2019), BlackBerry sued for infringement based on four patents. The court found three patents ineligible under Alice, noting the difficulty of discerning consistent boundaries in §101 jurisprudence.

Table 1

Annual patent applications and post-Alice rejections by industry.

This table reports annual statistics from USPTO patent applications and the corresponding percentage that were rejected in parentheses based on the Supreme Court's Alice decision for the top 12 industries with patent rejections. The rejection data provided by Lu et al. (2017) extends until 2016; therefore, the ratio of rejection is assigned NA for 2017. *Change* reports the percentage change from the number of patent applications in 2013 to the average number of patent applications for the 2015–2017 period. Corresponding CPCs for each industry are provided in Table IA5.

Number of Patent Applications & Rejection Percentage									
Industry	2008–2009	2010–2011	2012	2013	2014	2015	2016	2017	Change (2013 to 2015–2017)
Commerce (Data Processing Methods)	6582 (11.7%)	7675 (17.9%)	5032 (29.8%)	5564 (36.2%)	5221 (23.2%)	4246 (6.6%)	3406 (1.5%)	3253 (NA)	–34.7%
Administration (Data Processing Methods)	6681 (6.7%)	6250 (11.1%)	3658 (20.8%)	2958 (31.3%)	2969 (16.7%)	2500 (3.6%)	2531 (0.6%)	2574 (NA)	–14.3%
Finance (Data Processing Methods)	2297 (9.4%)	2662 (13.2%)	1545 (22.5%)	1752 (42.1%)	1512 (37.8%)	1038 (8.7%)	775 (1.9%)	715 (NA)	–51.9%
Payment Systems (Data Processing Methods)	1603 (9.9%)	2043 (12.9%)	1673 (26.6%)	1946 (36.7%)	2182 (24.4%)	2158 (5.8%)	2029 (1.9%)	1902 (NA)	4.3%
Coin-free Facilities or Services (Coin-free or Like Apparatus)	2385 (3.9%)	1665 (6.8%)	1221 (17.0%)	1407 (34.3%)	1134 (31.2%)	980 (14.9%)	939 (6.7%)	938 (NA)	–32.3%
Information Retrieval (Digital Data Processing)	7981 (0.5%)	8451 (1.2%)	5850 (2.4%)	6566 (4.1%)	6650 (5.1%)	6339 (2.0%)	6196 (1.0%)	5818 (NA)	–6.8%
Video Games (Games)	1414 (4.5%)	1504 (7.0%)	919 (12.5%)	1045 (27.4%)	1010 (19.6%)	781 (7.8%)	847 (3.4%)	931 (NA)	–18.4%
Specialized For Sectors (Data Processing Methods)	515 (4.9%)	918 (10.9%)	753 (15.5%)	845 (32.1%)	881 (19.3%)	669 (4.5%)	848 (0.6%)	809 (NA)	–8.2%
Computer Security (Digital Data Processing)	3886 (1.6%)	3926 (1.5%)	2617 (2.6%)	2684 (5.0%)	2641 (5.2%)	2604 (4.0%)	2675 (0.7%)	2881 (NA)	1.3%
Network Security (Transmission of Digital Information)	3522 (0.8%)	3208 (0.8%)	2206 (1.9%)	2864 (4.3%)	3433 (5.5%)	4042 (3.4%)	4126 (0.8%)	3822 (NA)	39.5%
Network Specific Applications (Transmission of Digital Information)	3389 (0.8%)	3441 (1.5%)	2282 (3.2%)	2891 (6.0%)	3174 (4.7%)	3172 (2.2%)	3099 (0.6%)	2420 (NA)	0.2%
Measuring or Testing Processes (Microbiology & Enzymology)	3759 (1.3%)	4311 (2.5%)	2237 (4.3%)	2356 (4.9%)	2336 (3.2%)	2105 (2.6%)	2101 (0.6%)	2087 (NA)	–11.0%

and 06/19/2014, manual classification is infeasible, necessitating the use of an LLM. To render the experiment tractable, we restrict attention to patents sharing the primary CPC codes of those rejected by the USPTO under Alice. This filtering yields 642,678 patents to be scored on their likelihood of losing protection. For training, we use patents applied for pre-Alice but granted post-Alice, where 33,734 patents were rejected (rejected = 1) with explicit reference to Alice. As negatives, we draw from the set of granted patents in the same CPC categories (rejected = 0). Together, these rejected and granted patents form the “ground truth” for model training and evaluation.

Basic text-based similarity measures such as term frequency-inverse document frequency (TF-IDF) suffer from two limitations. First, as technological vocabulary evolves, TF-IDF has limited ability to capture similarity across patents. Second, even patents with overlapping vocabulary might use it in different contexts and would be differentially affected by the Supreme Court's Alice decision. We therefore employ transformer models, which process words in context rather than independently or within a fixed sliding window.¹⁸ This mechanism, known as self-attention, captures semantic intent at the sentence level. For illustration, consider two sentences with similar meanings: (i) Symptoms of influenza include fever and nasal congestion; (ii) A stuffy nose and elevated temperature are signs you may have the flu. A TF-IDF model that filters stop words (e.g., “and”) assigns a similarity score of 0, whereas BERT and ModernBERT yield scores of 0.86 and 0.92, respectively.

¹⁸ Extensive empirical research demonstrates that LLMs outperform traditional NLP models, including Bag-of-Words (BOW), TF-IDF, word embedding models such as Word2Vec, FastText, and GloVe, as well as hybrid approaches combining embeddings with neural networks (Adhikari et al., 2020; Maltoudoglou et al., 2022; Esmaeilzadeh and Taghva, 2021; Minaee et al., 2021; Roman et al., 2021).


3.2. The ModernBERT model

Transformer-based large language models (LLMs) are neural network models that revolutionized the field of Natural Language Processing (NLP) (Kalyan et al., 2021). BERT is an early encoder-only model from Google that achieved state-of-the-art results on many NLP tasks (Devlin et al., 2019). Yet, a practical limitation is that it can only process 512 tokens, a limitation for long legal documents. We therefore adopt ModernBERT (Warner et al., 2025), a recent encoder-only Transformer-based LLM that extends the native context window to 8192 tokens and incorporates long-context improvements (e.g., rotary positional embeddings, alternating local-global attention, and FlashAttention). ModernBERT has been shown to provide a strong Pareto improvement over older encoder-only LLMs such as BERT and RoBERTa (Warner et al., 2025). Our results are consistent, as we find that ModernBERT also outperforms these older models. We thus use ModernBERT to predict the likelihood that the Alice decision leads pre-Alice granted patents to lose their IP protection.

ModernBERT is pre-trained on 2 trillion tokens of primarily English web documents, code, and scientific literature. We fine-tune the pre-trained model using labeled data.¹⁹ We use descriptions of patent

¹⁹ Transformer-based language models are fine-tuned for classification as follows. An additional linear layer is appended to the last layer. The parameters of this new layer, as well as the parameters of the language model, are then fine-tuned using the cross-entropy loss function. The weights of the model are not frozen. We fine-tune the ModernBERT model for 5 epochs with a batch size of 32 and a learning rate of $2e^{-5}$ using a linear learning rate scheduler and AdamW optimizer. We use the same fine-tuning parameters for each language model and use the Transformers Library version 4.52.3 (<https://huggingface.co/docs/transformers>).

Rejected Patent Applications because of Alice



pre-Alice Granted Patents





Fig. 1. Training and evaluation.

applications to train the model. However, the average patent has 13,168 tokens, and just 46.82% of our pre-Alice granted patents have descriptions shorter than 8192 tokens. We thus apply a TextRank-style extractive summarization method to reduce text size to ~ 4800 words. Specifically, we compute inter-sentence similarity using transformer-based embeddings from a pretrained language model, then rank sentences via the standard PageRank formulation (Mihalcea and Tarau (2004), Upasani et al. (2020)).

3.3. Rejected patent applications

We begin by collecting patents rejected under 35 U.S.C. §101 from the USPTO website.²⁰ We then identify Alice-related rejections following the method of Lu et al. (2017), yielding 56,709 rejected applications. Because some are reapplications with only minor textual changes (e.g., one or two sentences), we compute pairwise similarities using TF-IDF and flag applications with similarity scores ≥ 0.99 as duplicates, retaining only the most recent version. After deduplication, we obtain 33,734 unique rejected applications with both a document number and a Cooperative Patent Classification (CPC) code. These Alice-rejected patents cover 5831 unique CPCs.

3.3.1. Training the ModernBERT model

Our model operates in two phases: training and evaluation (see Fig. 1). For training, we use patent applications explicitly rejected under the Alice decision (rejected = 1) and those ultimately granted (rejected = 0). Model evaluation is assessed on a hold-out testing sample. We then apply the trained model to score the 642,678 pre-Alice granted patents.

We use the full set of post-Alice rejected patents from the three years following the decision for model training and testing. In this sample, 80.4% of training patents and 80.2% of testing patents were applied for prior to Alice. This design ensures both optimal sample size and balance, while also maintaining that the vast majority of patents were

submitted pre-Alice or shortly thereafter, when its full implications were not yet evident.

From the 33,734 Alice-rejected applications, we randomly select 10,000 for hold-out testing and use the remaining 23,734 as the positive training sample (rejected = yes). We then construct a negative sample (rejected = no) from patents successfully granted after 06/19/2014, the date of the Alice decision. To ensure robustness, we generate this negative sample using four levels of CPC granularity: (i) section, (ii) class, (iii) subclass, and (iv) group, as well as (v) main group or subgroup. For example, in CPC “B60K35/00” – B, 60, K, 35, and 00 correspond to the Section, Class, Subclass, Group, and Main Group, respectively.

In Experiment A, each of the 23,734 rejected = “yes” patents is matched to a rejected = “no” patent granted after 06/19/2014 within the same CPC Group. In Experiments B, C, and D, we sequentially expand the rejected = “no” pool by adding 23,734 patents at the CPC Subclass, Class, and Section levels, respectively. Thus, from A to D, each sample incorporates an additional 23,734 rejected = “no” patents drawn from progressively broader CPC categories.

3.3.2. Testing ModernBERT and other models

In this section, we compare our trained ModernBERT model with Longformer, SciBERT, BERT, RoBERTa, TF-IDF, and Word2Vec. For TF-IDF and Word2Vec, predictions are generated in combination with logistic regression, decision tree, and random forest classifiers. A detailed technical comparison of the transformer models is provided in IA Section 1.

For testing, we use the 10,000 hold-out rejected patents and match each with two accepted patents, yielding 20,000 accepted patents (rejected = 0). This 1:2 ratio reflects the lower than 50% expected rejection rate and mitigates inflated accuracy from models that trivially predict all patents as accepted.²¹ From the 708,115 post-Alice granted patents, we first select 50,000 accepted patents randomly, weighted by

²⁰ <https://developer.uspto.gov/product/patent-application-office-actions-data-stata-dta-and-ms-excel-csv>

²¹ For instance, with a 1:9 ratio of positives (rejected = 1) to negatives (rejected = 0), a naive model predicting only acceptance would achieve 90% accuracy.

Table 2

Comparison of predictions for ModernBERT vs. other models.

This table compares predictions of ModernBERT (Warner et al., 2025), Longformer (Beltagy et al., 2020), SciBERT (Beltagy et al., 2019), BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), TF-IDF (Robertson, 2004) and Word2Vec models (Mikolov et al., 2013) based on F_1 Score and Accuracy. F_1 Score is the harmonic mean of recall and precision, which are defined in Eq. (1). Accuracy is the ratio of correctly predicted observations to the total observations. For all models, we conduct four experiments in which the only difference is the way we create the training samples. In experiment A, for each of the 23,734 positives (rejected patent applications), we find a matching negative (a patent application that is granted) that is in the same CPC Main Group. In sample B, C, and D, without replacement, we keep adding 23,734 more matching *granted* patents to the negatives pool based on CPC Group, Subclass, and Class, respectively. Therefore, from A to D, each sample has 23,734 more negatives, i.e. granted patents, but the newly added ones are selected from a broader CPC. In the last column of the table, we use an ensemble of the two models that have the highest F_1 Score and Accuracy by taking the average of their prediction scores. The use of this ensemble is motivated by the fact that A typically has the highest F_1 Score and D has the highest Accuracy. For the testing, we only use applications and granted patents that were not used in the training. In the testing, we have 10,000 positives/rejected=1 in the sample of applications and for each positive, we choose two negatives/rejected=0. This 1:2 positives to negatives ratio is consistent with the expected rejection ratio having more negatives than positives, while not overestimating accuracy for models that do not learn but only predict negative outcomes. From the negatives pool, we thus sample 20,000 negatives 1000 times and boot-strap performance of the models. The table below then reports the average F_1 Score and Accuracy for each model.

Model name	A		B		C		D		$\frac{A+D}{2}$	
	F_1 Score	Accuracy	F_1 Score	Accuracy	F_1 Score	Accuracy	F_1 Score	Accuracy	F_1 Score	Accuracy
ModernBERT Finetune	0.677	0.766	0.653	0.797	0.655	0.811	0.662	0.815	0.700	0.825
Longformer Finetune	0.647	0.745	0.623	0.765	0.618	0.784	0.639	0.800	0.682	0.804
SciBERT Finetune	0.651	0.735	0.634	0.749	0.632	0.767	0.638	0.777	0.669	0.778
BERT Finetune	0.623	0.733	0.598	0.739	0.614	0.764	0.624	0.774	0.642	0.775
RoBERTa Finetune	0.600	0.716	0.555	0.740	0.540	0.756	0.515	0.758	0.592	0.765
TF-IDF + Logistic Regression	0.547	0.643	0.599	0.634	0.613	0.670	0.550	0.719	0.559	0.679
TF-IDF + Decision Tree	0.503	0.602	0.554	0.552	0.558	0.584	0.491	0.690	0.409	0.697
TF-IDF + Random Forest	0.628	0.743	0.368	0.717	0.263	0.696	0.209	0.689	0.387	0.723
Word2Vec + Logistic Regression	0.606	0.731	0.418	0.732	0.377	0.732	0.358	0.730	0.497	0.755
Word2Vec + Decision Tree	0.492	0.607	0.456	0.645	0.461	0.687	0.461	0.702	0.365	0.707
Word2Vec + Random Forest	0.619	0.747	0.439	0.746	0.387	0.739	0.365	0.735	0.500	0.766

the CPC frequency distribution of our prediction sample. We then draw 20,000 accepted patents from this pool 1000 times to bootstrap model performance. Evaluation is based on standard metrics – precision, recall, F_1 score, and accuracy –calculated from the confusion matrix: True Positives (TP), False Positives (FP), True Negatives (TN), and False Negatives (FN).

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN},$$

$$F_1 \text{ Score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}, \quad \text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}.$$

(1)

Table 2 reports the evaluation results. Since no single model achieves both the highest F_1 score and accuracy, we average the prediction scores of the two leading models (Model A with the highest F_1 score and Model D with the highest accuracy). This combined ModernBERT model outperforms all other algorithms, attaining the highest F_1 score (0.700) and accuracy (0.825).

Table IA4 reports additional statistics and compares actual Alice rejections with ModernBERT-predicted rejection statistics by CPC category. The results indicate that actual rejections closely align with ModernBERT predictions for the most affected CPC codes.

3.3.3. Further validation of Alice scores with post-grant reviews

In this section, we evaluate the predictive power of Alice scores in the context of USPTO post-grant reviews, trial proceedings before the Board that assess patentability under §282(b)(2) or (3). Post-grant reviews may be initiated by a third-party petition filed within nine months of the patent grant or reissue.²² We obtain data on post-grant reviews from Unified Patents, covering 391 unique patents reviewed between September 16, 2012 (the start date of post-grant reviews) and September 16, 2024. Of these, 75 petitions explicitly cited the Alice decision as grounds for invalidation.

Table 3

Post-grant review prediction.

The table displays the results from a logit regression in which the dependent variable, Alice Referenced Dummy, is a dummy equal to one if the given patent in the post-grant review references the Alice case and zero otherwise. The sample consists of 391 unique patents from September 16, 2012 (the start date of post-grant reviews) to September 16, 2024, that experienced post-grant reviews. We note that 75 of these reviews specifically referenced the Alice case when requesting invalidation. *Alice ModernBERT Score* is the predicted value of Alice rejection from the ModernBERT model. *Log(Number of Claims)* is the logged number of claims of the patent. *Petitioner Types* indicate the type of entity petitioning for post-grant review. These include defensive aggregators, government, large operating companies, small or medium enterprises, unified patents, or other. The *t*-statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

Dependent variable:	Alice referenced dummy			
	(1)	(2)	(3)	(4)
Alice ModernBERT Score	5.422*** (8.58)	5.399*** (8.51)	5.843*** (7.42)	7.576*** (4.38)
Log(Number of Claims)		−0.257 (−0.89)	−0.139 (−0.38)	−1.448** (−2.31)
Observations	391	391	339	223
Petitioner Type	NO	NO	YES	YES
Filing Year Fixed Effect	NO	NO	YES	YES
Grant Year Fixed Effect	NO	NO	YES	YES
CPC Fixed Effect	NO	NO	NO	YES
Pseudo R^2	0.2508	0.2528	0.3641	0.5553

Table 3 reports regression results where the dependent variable equals one if the post-grant review references Alice and zero otherwise. Controls include the log number of claims, petitioner type (defensive aggregator, government, large operating company, SME, Unified Patents, or other), filing and grant year fixed effects, and CPC fixed effects. Despite the small sample, the ModernBERT Alice Score significantly predicts Alice-referenced cases at the 1% level, with *t*-statistics ranging from 4.38 to 8.58 across specifications. These findings provide further evidence of the predictive accuracy of the ModernBERT Alice score.

²² <https://www.uspto.gov/patents/laws/america-invents-act-iaia/interpartes-disputes>

3.4. Validation of disappearing post-Alice patents

Online Appendix Table IA7 (and detailed discussion in Online Appendix 1.1) reports the distribution of Alice ModernBERT Scores for patents granted in the Top 20 technological areas most affected by Alice. Patents each year are sorted by Alice scores into ten equal-sized bins between 0 and 1, and density is computed as the share of patents in each bin. The table compares patenting rates for each bin in 2017 (post-Alice) to those in the Pre-Alice years 2011–2013. Ratios below one in the last column indicate post-Alice declines. We find sharp reductions in high-score bins, with decile 10 falling to 11% of pre-Alice levels. Focusing on the last two bins where ratios are below 50%, the estimated number of “likely lost patents” is an economically large 6898 per year. Fig. 2 graphically confirms these sharp drop-offs, and Table IA8 shows robustness for applications. These results add to our evidence of validation as they confirm our strong prediction that patents exposed to Alice concepts will drop off sharply post-Alice with economically large impact. These broad declines across areas also indicates that firms cannot easily bypass Alice by resubmitting altered filings.

3.5. Patent sample and treatment measure

We create the treatment measure for each firm i that we use in our regressions as follows:

$$Treatment_i = \frac{\sum_{j=1}^{N_i} PatentValue_{i,j} \times AliceScore_{i,j}}{Sales_i} \quad (2)$$

$Sales_i$ denotes firm i 's total sales in 2014. $PatentValue_{i,j}$ is the dollar value of patent j , drawn from the KPSS database (Kogan et al., 2017). The functional form is theoretically motivated as the fraction of a firm's patent portfolio value affected by Alice multiplied by the portfolio's importance to firm operations.²³ The treatment variable is calculated for each firm in 2014 using all valid patents granted prior to the Alice decision. Patent values are depreciated at an annual 20% rate relative to 2014, and adjusted for inflation.²⁴ $AliceScore_{i,j}$ is ModernBERT's predicted probability that patent j loses protection if challenged in court. N_i is firm i 's total number of patents.

As an alternative to KPSS valuations, we estimate $PatentValue_{i,j}$ in Eq. (2) using a citation-based metric. For each patent j , we count the number of granted patents citing j with application dates within five years of j 's grant date. Citation-based values are depreciated at an annual 20% rate relative to 2014.

We estimate all regressions using a difference-in-differences (DiD) approach, where *Post* indicates years after Alice and is interacted with *Treatment* to capture heterogeneous treatment effects.²⁵ To account for firm size heterogeneity, we interact *Treatment* with indicators for small and large firms. All regressions include firm and year fixed effects, and standard errors are clustered by firm. Given these fixed effects, the lower-order interactions of *Treatment*, size dummies, and *Post* are included but absorbed. We do note that the 2-way interactions of the size dummies**Post* are not absorbed, but we do not report them as these are the differences between firms in the post- and pre-periods at the mean treatment value, and are not our focus.

²³ Formally: (value of patents impacted by Alice/value of patent portfolio) × (value of patent portfolio/sales) = (value of patents impacted by Alice/sales).

²⁴ We follow Hall and Li (2020), who show that depreciation rates are likely higher than the 15% commonly assumed, particularly in high-technology sectors. Results are robust to a 15% rate.

²⁵ Because *Treatment* is continuous rather than binary, we also estimate a fuzzy DiD model following de Chaisemartin and D'Haultfoeuille (2018) and de Chaisemartin et al. (2019), with results reported in Tables IA21 and IA22 of the Internet Appendix.

Our focus on heterogeneous effects by firm size follows Aghion et al. (2005), who argue that larger firms are better able to defend product markets given greater resources and potential entry barriers (mechanisms examined in Section 4.4). We classify firms as large or small based on whether their 2013 assets are above or below the median within their TNIC-2 industry peers (Hoberg and Phillips, 2016).

In all subsequent regression tables, *Post* equals one for years after the Alice decision (2015–2017) and zero for years before (2011–2013). We exclude 2014 as partially treated, though Section 3 of the Internet Appendix reports results treating 2014 as a separate category. *Treatment* is a firm-level measure combining the importance of patents to the firm with the extent to which its portfolio is affected by Alice. Specifically, we weight each patent's Alice score by its importance to the firm, using two alternative measures: (i) KPSS-based patent values and (ii) citation-based values. Eq. (2) defines the measure, with citations substituted for patent values in the latter case. Results are broadly similar across the two weighting methods.

3.6. Sample and key variables

We include all public firms with at least one patent in a CPC category containing a rejected patent. Patents are linked to firms using the correspondence of Kogan et al. (2017), extended through 2020. Patent text data are obtained from the USPTO. We also add competitors using the TNIC-3 network of Hoberg and Phillips (2016). The final sample comprises 3646 unique firms: 1087 Alice-affected firms and 2559 competitors.

Table 4 reports summary statistics for the main sample of 19,808 firm-year observations from 2011–2017, excluding 2014, the treatment year.²⁶ Variables used in the analysis are described briefly below, with full definitions in Appendix A. The table presents statistics separately for the pre-Alice (2011–2013) and post-Alice (2015–2017) periods.

Panel A of Table 4 presents firm characteristics, including the size of firms measured by assets and sales, sales growth, age, profitability ((Sales-Cost of Goods Sold) / Sales), and valuation (market-to-book ratio = (market value of equity + book value of debt) / lagged book value of assets). The results indicate declines in both profit margins and sales growth.

Panel B presents the innovation, and legal variables used in our study. Note that all scaled variables are scaled by lagged sales or assets. *Treatment* measures the extent a firm's patent portfolio is impacted by Alice as measured using the ModernBERT model in Eq. (2). It captures how much a firm is dependent on patents and also the percentage of patents' value that are impacted by the Alice court decision. For ease of interpretation, we normalize the *Treatment* by subtracting the mean and dividing by the standard deviation, resulting in a z-score transformation. R&D/Sales is Compustat R&D divided by lagged sales of the firm and is set to zero if R&D is missing for our base tests. $Patent\ App_{i,t}/Sales_{i-1}$ ($Assets_{i-1}$) is the number of patent applications divided by lagged sales (assets).

The legal variables we examine are # *NPE Alleged*, # *OC Alleged*, and # *Alleged*. We compute the first three using information in the Public Access to Court Electronic Records (PACER) database, which provides public access to all cases litigated in the U.S. District Courts, and Stanford Non-Practicing Entity (NPE) Litigation Database. For the last one, we use textual queries of each firm's 10-K statement filed with the SEC. # *NPE Alleged* (# *OC Alleged*) refer to the number of lawsuits that the firm was alleged infringing on a patent lawsuit by a Non-Practicing Entity (Operating Company) in that year. # *Alleged* is the total number of lawsuits that the firm was alleged infringing on a patent. The table shows that patents, total lawsuits, and OC-based lawsuits decline post-Alice.

²⁶ Robustness results including 2014 are provided in Internet Appendix Section 3, Tables IA11–IA15.




Fig. 2. Ratio of post-Alice density to pre-Alice density.

The figure illustrates whether there is a decrease in applications and grants of patents with high Alice scores after the Supreme Court decision (the final column of Table IA7). We compute the density of the Alice score based on ten bins (increments of .1) from zero to unity both before Alice (2011 to 2013) and post Alice (2017). The figure reports the ratio of the density for each bin. The values below unity for the rightmost bins below indicate that many fewer patents with high Alice scores were applied for (Panel A) and granted (Panel B) after the Alice Supreme Court decision.

Panel C of Table 4 presents competition summary statistics. $VCF_t/Sales_{t-1}$ is a measure of VC entry and is total first-round dollars raised by the 25 startups from Venture Expert whose Venture Expert business description most closely matches the 10-K business description of the focal firm (using cosine similarities), scaled by lagged focal firm sales. The next three variables are constructed using the metaHeuristica software package to run high-speed queries on 10-Ks filed with the Securities and Exchange Commission. *Complaints* is the number of paragraphs in the firm's 10-K that complain about competition divided by the total number of paragraphs in the firm's 10-K. *Nondisclosure* is the total number of 10-K words mentioning "non-disclosure", "NDA agreements", "confidentiality agreement" or

"secrecy agreement" scaled by the total number of words in the 10-K. The table shows that VC-backed entries and nondisclosure agreements increase post-Alice.

Our treatment variable is not binary and represents the percentage of a firm's patent portfolio value that is exposed to Alice scaled by sales. Each patent's Alice exposure score is the probability from our ModernBERT model that the patent will be ruled ineligible if it is challenged in court. For roughly half of the regression sample, the treatment Alice score is close to 0. Prior to normalization, the treatment score has a median of zero and a mean of 0.008, while the 75th and 90th percentiles are 0.0004 and 0.0175, respectively. The standard deviation is 0.026. We show the full distribution of firm-level treatment

Table 4

Firm summary statistics.

This table provides summary statistics for our sample of public firms based on annual firm observations from 2011 to 2017 (excluding 2014, the treatment year). All variables are described in detail in the variable list in [Appendix A](#) and in Section 3 of the paper. *, **, and *** denote significant difference of the mean post-Alice vs. pre-Alice at the 10%, 5% and 1% level.

Variable	N	# Firms	Pre-Alice			Post-Alice			Diff
			Median	Mean	Std. Err.	Median	Mean	Std. Err.	
Panel A: Firm Characteristics									
Assets (in mil.)	19808	3646	830.133	9139.905	1226.523	1251.855	11 393.070	1395.515	
Sales (in mil.)	19808	3646	384.894	3509.421	253.765	551.892	4031.148	273.913	
Sales Growth	19808	3646	0.064	0.088	0.003	0.037	0.033	0.003	***
Age	19808	3646	18.000	21.347	0.265	22.000	25.510	0.280	***
Gross Profit Margin	19808	3646	0.434	0.446	0.007	0.403	0.381	0.011	***
M/B Ratio	19309	3603	1.133	1.906	0.040	1.187	1.756	0.031	***
Panel B: Innovation & Lawsuit Characteristics									
Treatment	19808	3646	-0.306	0.001	0.017	-0.306	0.001	0.017	
R&D _t /Sales _{t-1}	19808	3646	0.000	0.143	0.008	0.000	0.158	0.011	***
Patent App./Sales _{t-1}	19808	3646	0.000	0.015	0.001	0.000	0.008	0.000	***
Patent App./Assets _{t-1}	19808	3646	0.000	0.008	0.000	0.000	0.004	0.000	***
# NPE Alleged	19808	3646	0.000	0.163	0.008	0.000	0.154	0.007	
# OC Alleged	19808	3646	0.000	0.101	0.005	0.000	0.048	0.003	***
# Alleged	19808	3646	0.000	0.280	0.012	0.000	0.227	0.010	***
Panel C: Competition Measures									
VCF/Sales _{t-1}	19725	3646	0.014	0.122	0.005	0.012	0.195	0.012	***
Complaints	19727	3646	13.484	14.586	0.113	13.611	14.709	0.116	
Nondisclose	18486	3509	0.000	0.012	0.001	0.000	0.016	0.001	***

Alice scores for patenting firms in [Fig. 3](#). Panel A shows the histogram for our KPSS valuation-based treatment variable and Panel B for our citation-based treatment variable.

Both panels show that roughly 25% to 30% of our firms have zero Alice scores. An additional roughly 40% have positive scores that are close to zero. Thus, roughly 65% have Alice scores equal or slightly greater than zero (from 0.0 to .005). About 6% to 7% of firms have very high exposure to Alice with Alice scores in the rightmost bin. Using a continuous treatment score allows us to show how the ex post outcomes vary with the intensity of treatment. Henceforth, we refer to the Alice ModernBERT Score just as “Alice Score”.

[Table 5](#) displays eight examples of large and small firms with high Alice-scores. It presents their most impacted patent and also firm-level characteristics including the final firm-level treatment value. All dollar values are depreciated by our patent depreciation rate and also by the CPI so that values are in 2014 dollars. The table shows that both large and small firms have patents impacted by Alice. It also shows that firms that have highly Alice-impacted patents frequently also have many other patents that are not impacted. The observation that larger firms often have many non-Alice impacted patents along with Alice-impacted patents illustrates the depth of knowledge accumulation these firms have, which likely helps to explain why these firms fare better after the Alice shock than do smaller firms.

3.7. Entropy balancing approach

As treatment effects may be concentrated in specific industries, and firms exposed may have different characteristics, we use entropy balancing in our main regressions to ensure covariate balance and clean identification of Alice effects. To implement entropy balancing, we use [Tübbicke \(2022\)](#), which is the continuous treatment extension of [Hainmueller \(2012\)](#).

Entropy balancing is most intuitively understood as a more accurate and generalized version of nearest neighbor matching (NNM). Entropy balancing derives observation weights that fully balance industry and key control variables across treated and untreated groups. NNM cannot fully balance these characteristics due to imperfections and granularity in sorting methods. Notwithstanding its disadvantages, we also report results using NNM in the Internet Appendix Table IA17 to Table IA20 and find that our results are mostly robust.

Our implementation of entropy balancing imposes covariate balancing for: (i) Two-digit SIC codes (ii) Sales (iii) Assets (iv) Log(Age) and (v) the Number of Valid Patents as of 2013. [Table 6](#) tests the accuracy of entropy weights and presents the results from regressing the key variables on our treatment variable using equal weights versus entropy weights. The results in the odd-numbered columns show that equally weighted (non-entropy balanced) treated firms have higher sales, assets, and patent counts. The entropy-balanced even-numbered columns show that these estimates are insignificant using entropy weights. We conclude that covariate balance is well-achieved.

3.8. *Bilski, Mayo, and Myriad Supreme Court cases*

Although the Alice Supreme Court case has been more impactful than other recent cases that weaken patent protection for methods-based inventions, out of an abundance of caution, we also consider the influence of three other impactful Supreme Court cases that pre-date Alice: *Bilski v. Kappos* (561 U.S. 593, 2010), *Mayo v. Prometheus* (566 U.S. 66, 2012), and *Association for Molecular Pathology v. Myriad Genetics, Inc.* (569 U.S. 576, 2012). Using the USPTO actions data from [Lu et al. \(2017\)](#), we divide the number of rejections for each case by the number of patent applications filed between the date of the respective Supreme Court decision and the end of 2016. These actual rejections indicate that *Bilski* had 22.4%, *Mayo* had 8.2%, and *Myriad* had 4.6% the overall impact of Alice.

Given *Bilski* was impactful enough to estimate a machine learning model, we apply the same machine learning analyses that we have performed for Alice to separately estimate a *Bilski* score for each patent in our database. We find that 61,373 of our 642,678 patents would be predicted to lose their protection following the *Bilski* decision. Yet the impact was in different technological areas as the correlation of the *Bilski* score and our Alice score in a sample of Alice or *Bilski* affected CPCs is just 2.4%. This lack of correlation between these scores indicates that the impact of *Bilski*, while less than Alice in economic terms, is unlikely to contaminate our inferences as uncorrelated influences are unlikely to bias coefficients.

As the total number of rejections for the *Mayo* and *Myriad* cases is very low (a full order of magnitude smaller than Alice), we are unable to apply the same ML techniques to these two cases due to data inadequacy for training. For these more narrow cases, which focus on specific issues in biotechnology, we instead identify the top-5 CPC




Fig. 3. Histogram for treatment.

This figure displays the histogram of the treatment variables for firms with patents before Z-scoring. In Panel A, the treatment is based on KPSS, and it is based on citation in Panel B. The bin width is 0.01 and y-axis is the percentage of treatment falls into the bin.

codes for each to assess impacted patents. These top-5 CPC codes are indeed in the biology domain, and each constitutes less than 1% of the Alice-impacted patents.

To be conservative, for all subsequent regressions examining the impact of Alice, we drop all Bilski-related patents (61,373) noted above, and also all patents in the top-5 CPCs from Mayo and Myriad (2353 for Mayo and 3803 for Myriad) from our sample of 642,678 patents. This reduced sample forms the sample used in our baseline tables in the paper. Our economic inferences are robust to using the full sample of 642,678 patents. We also conduct additional robustness in Section 4.6.

4. The impact and outcomes of Alice

We now analyze the impact of Alice on innovation, firm performance, valuation, competition, and lawsuits, and then further explore

mechanisms such as barriers to entry. Throughout this section, we present results for large and small firms by separating them into yearly above and below median groups based on assets. Our tests are based on a continuous-treatment-intensity difference in differences (DiD) framework based on the extent of firms' exposure to Alice. We present formal DiD regression analysis and focus on testing for differential treatment effects for small versus large firms. We reinforce these DiD tests with pre-trend analysis and formal structural break tests that also focus on differential treatment effects.²⁷ Both the DiD and structural break

²⁷ We separately tabulate individual small and large firm treatment effects in Online Appendix Section 7.

Table 5

Sample of alicé-impacted patents and owner characteristics.

This table provides a representative selection of patents exhibiting a high Alice ModernBERT score, accompanied by firm characteristics pertaining to the patent owners. Panel A presents data concerning large firms within our sample classification, while Panel B offers corresponding information for small firms. After the first column (company name), the next four columns contain information about the company's most-exposed patent to Alice (as indicated by the high Patent Alice ModernBERT Score). The final three columns display firm characteristics summarizing patenting intensity and the value of discounted total patent portfolio.

Company name	Patent number	Patent description	Patent Alice-score	KPSS disc. patent value (2014 \$M)	Firm's total number of patents	Firm's # of alicé-impacted patents	Firm's total patent value (2014 \$M)
Panel A: Large Companies							
JPMORGAN CHASE & CO	6 502 080	A method of predicting a net reserve for a vehicle leased by a lessee from a lessor in accordance with a lease.	0.99	44.06	1,702	380	123,594
ORACLE CORP	8 738 489	A computer-implemented method that manages financial close schedules by creating, validating, duplicating, and setting parameters.	0.95	46.91	6,359	596	102,589
GENERAL ELECTRIC CO	8 751 423	A turbine performance diagnostic system that creates a performance report for one or more turbines.	0.93	22.17	6,359	908	44,037
WALT DISNEY CO	8 745 165	A system and method for managing distribution of rich media content.	0.93	69.19	1,640	114	23,960
Panel B: Small Companies							
ASPIRA WOMEN'S HEALTH INC	8 664 358	Methods to predict ovarian cancer presence, subtype, stage, treatment efficacy, and cancer development using diagnostic systems.	0.97	1.677	201	15	13.49
(a.k.a. VERMILLION INC) MULTIMEDIA GAMES HOLDING CO	8 708 806	A method that controls free plays in wagering games through a repeatable secondary game awarding additional plays.	0.90	9.44	308	68	206.67
GUIDEWIRE SOFTWARE INC	8 566 131	A system that manages insurance policy processing by enforcing holds, blocking operations, and releasing holds when criteria are met.	0.99	17.71	14	10	166.96
STAMPS.COM INC	8 612 361	System and method for handling payment errors with respect to delivery services	0.97	5.73	403	45	111.01

frameworks indicate significant treatment effects without pretrends for nearly all variables examined.

4.1. Alice and innovation

We first examine the direct impact of Alice on ex post patents using the number of patents scaled by assets and also by sales. Compared to scaling by assets, scaling by sales has the advantage that it is less subject to accounting differences and, in particular, differences from accounting for acquisition goodwill. In all tables, the *Treatment* variable is normalized by subtracting the mean and dividing by the standard deviation, resulting in a z-score transformation for ease of interpretation.

The results in columns (1)–(4) of Table 7 show that both large and small firms reduce patenting in the years after Alice. The results for small firms are highly significant at the 1% level, and the results for large firms are significant at the 5% level. The third row shows that the impact on smaller firms is significantly larger using our most well-motivated KPSS treatment variable, although it is insignificant using our citations measure. These findings confirm the importance of Alice in reducing patents for both small and large firms due to weaker IP protection in the post-Alice environment. The economic effects are large. For patents relative to sales, using the coefficients in column (3) for our interaction term with small firms (−0.015) for KPSS-scaled values, we calculate that small firm patenting decreases

by 31% relative to the pre-Alice mean for small patenting firms (0.048) with one standard deviation shift in the treatment variable. In the same column, the interaction for large firms is −0.002, which implies that large firm patenting decreases by 14% relative to the pre-Alice mean for large patenting firms (0.015).

We present pre-trend analysis and formal structural break tests for the difference between small and large firms graphically in Fig. 4. Our central results in Fig. 4 show that, regardless of scaling by sales or assets, (A) there is no significant pre-trend (the x-axis lies within the error bands) and (B) there is a very significant structural break after Alice as the error bands lie significantly below the projected pre-trend dotted line. Our results overall indicate significant decreases in patenting, and that decreases are more severe for small firms.

We also note that our regressions include firm fixed effects, and thus, we do not report the lower interactions, including the individual variables (Large, Small, and Treat), as these are absorbed by firm fixed effects given that they are defined in the treatment year. We do note that the 2-way interactions of the size dummies*Post are not absorbed, but we do not report them as these are the differences between firms in the post- and pre-periods at the mean treatment value, and are not our focus. Although we cluster standard errors by firm, our results are robust to more stringent clusters by industry, illustrating that industry differences in Alice-affected patents do not drive our results.

Table 6**Treatment and control group balance.**

The table displays the results from regressions where key characteristics, both equal-weighted and entropy-weighted, are regressed on our Alice treatment variable. The goal is to illustrate the degree of imbalance prior to entropy weighting (expect significant results) and the efficacy of entropy balancing after weighting (expect insignificant results). In the odd columns, dependent variables are equal-weighted sales, assets, logged age, and the total number of valid patents until that year. In the even-numbered columns, the dependent variables are entropy-weighted using weights from the entropy balancing approach explained in Section 3.7. In columns (1)–(2) and (5)–(6), the data only includes observations from 2013 (i.e., the year before the Alice decision). All years are included in the other columns. All variables are described in Appendix A. Even-numbered regressions include year fixed effects, and standard errors are clustered at the firm level. *t*-statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

PANEL A:	(1) Sales	(2) Weighted sales	(3) Sales	(4) Weighted sales	(5) Assets	(6) Weighted assets	(7) Assets	(8) Weighted assets
Treatment	−307.760*** (−2.66)	0.030 (0.25)	−312.214*** (−4.59)	0.041 (0.33)	−631.121*** (−2.99)	0.052 (0.22)	−674.716*** (−4.66)	0.115 (0.40)
Sample	Only 2013	Only 2013	All Years	All Years	Only 2013	Only 2013	All Years	All Years
Year Fixed Effects	NO	YES	NO	YES	NO	YES	NO	YES
Observations	3646	3646	19808	19808	3646	3646	19808	19808
Adj. R ²	0.002	−0.000	0.003	−0.000	0.002	−0.000	0.007	0.000
PANEL B:	(1) Log(Age)	(2) Weighted log(Age)	(3) Log(Age)	(4) Weighted log(Age)	(5) Patent count	(6) Weighted patent count	(7) Patent Count	(8) Weighted patent count
Treatment	−0.047*** (−3.58)	0.001 (0.01)	−0.050*** (−5.32)	0.000 (0.05)	212.629*** (12.33)	0.003 (0.31)	230.970*** (7.63)	0.006 (0.63)
Sample	Only 2013	Only 2013	All Years	All Years	Only 2013	Only 2013	All Years	All Years
Year Fixed Effects	NO	YES	NO	YES	NO	YES	NO	YES
Observations	3646	3646	19808	19808	3646	3646	19808	19808
Adj. R ²	0.003	0.000	0.030	0.001	0.040	−0.000	0.044	0.000

4.2. Alice and competition

Unlike some existing studies, which focus on the impact of individual patent invalidations, our study examines the impact of a technology-area-wide loss in IP protection. Such a market-wide shock impacts both existing patents and also the incentives to patent more in the future. These shifts, furthermore, affect the incentives of potential competitors, and, thus, it is important to examine the impact of Alice on competition coming from either new VC-funded entrants as well as from existing public firms.

Table 8 examines two measures of changes in firm-level competition and firm R&D. We also examine pairwise competitive encroachment later in this section. Columns (1) and (2) of Table 8 examine venture capital entry into a firm's local product market. The dependent variable, *VCF/Sales*, is the total first-round dollars raised by the 25 startups from Venture Expert whose Venture Expert business description most closely matches the 10-K business description of the focal firm (using cosine similarities), scaled by focal firm sales (see Hoberg et al. (2014)). We examine broad competition *Complaints* in columns (3) and (4). *Complaints* is the number of paragraphs in the firm's 10-K that complain about competition divided by the total number of paragraphs in the firm's 10-K. In columns (5) and (6), we examine investment in R&D.

Table 8 indicates that entry by VC-backed firms into small firm markets increases by 138% (51%) with a one standard deviation increase in our KPSS-treatment (citation-based treatment) relative to the average entry pre-Alice for small firms (0.22). The third row shows that these results are significantly higher for small than for large firms. Competition complaints also increase significantly more using our primary KPSS based measure although the positive difference is not significant for the cites-based measure. Complaints by small firms increase by 1.85%.

The results in columns (5)–(6) show that small firms increase R&D after Alice, while there has been a decrease for large firms. Using the coefficient from column (5), we calculate that small firms' R&D/sales increases by 82% relative to the mean of small firms pre-Alice (0.242) with a one standard deviation increase in the KPSS-based measure of

treatment (28% using the citation-based measure of treatment). Large firms decrease R&D by 29% using KPSS-based treatment value, and insignificant using the citation-based metric. The third rows shows that the difference between small and large firms are significant at the 1% level for both of the treatment calculation metrics.

The pre-trend analysis and structural break results for VC entry, competition complaints, and R&D are shown graphically in Fig. 5. As before, the tests examine the difference between small and large firm treatment effects. We observe a slightly significant but economically small pre-trend for VC entry and no significant pre-trend for competition complaints. The structural break analysis also confirms significant changes post-Alice as confidence intervals for both variables are quite far from the dotted line pre-trend projection.

However, there is a significant pre-trend for R&D. Moreover, the structural break test around 2014 for R&D suggests that its changes are not significant. These results are consistent with the view that firms hedge the potential impact of Alice by preemptively increasing R&D even before the final Alice decision. R&D is somewhat unique in this way as it is easy to increase pre-emptively. Overall, our R&D results are consistent with small firms trying to increase R&D to make up for lost intellectual property (or expectations of future losses), as reinforced by the higher VC-funded entry and higher competition we document above (structural breaks around Alice are unambiguous for these variables). In contrast, large firms insignificantly decrease R&D, indicating they were impacted less, making R&D increases unnecessary. This muted response by larger firms echoes our consistent finding that larger firms (presumably due to deeper knowledge capital and barriers to entry) were less (or even marginally positively) impacted by Alice than smaller firms.

We now examine local pairwise product market encroachment post-Alice in Table 9. *Delta TNIC Score* is computed as the change in the TNIC similarity of the pair of firms from year $t-1$ to year t . A high value indicates encroachment (increased similarity and thus higher competition within the pair) and a decrease indicates the two firms are becoming more differentiated. Our panel for this test is a large firm-pair-year panel.




Fig. 4. Difference in patent applications between small and large firms.

This figure presents the yearly point estimates for the difference between small and large firms by utilizing $\text{Small} \times \text{Treatment}$ interacted with year-specific dummies (without $\text{Large} \times \text{Treatment}$). The dependent variables are defined as $\text{Patents/Assets}_{t-1}$ and $\text{Patents/Sales}_{t-1}$. The year 2013, which immediately precedes the treatment, is designated as the reference year in the regression analysis. The treatment variable is constructed using KPSS values. The black dotted line represents a linear projection based on a regression fitted to the 2011, 2012, and 2013 data points, and extrapolated to cover the years 2015, 2016, and 2017. The shaded gray region corresponds to the 99% confidence interval.

The variable “ $\text{Treat} \times \text{Post}$ ” is the sum of the Alice treatment variable of the two firms in the given pair multiplied by the post-Alice dummy. The results in column (1) of Table 9 intuitively show that firm pairs jointly experiencing a large Alice treatment, on average, experience increased product market encroachment at the pair level. This is consistent with weaker IP protection resulting in rivals adopting the patented technologies of rivals and products of the pair becoming more similar. These results are highly significant despite the inclusion of firm-pair fixed effects and clustering standard errors by firm-pair.

Column (2) of Table 9 examines if treatment effects are different for small and large firms. We find that small firms are particularly

prone to encroachment after Alice. This is consistent with these firms holding narrower advantages in the product market, and any losses in protection might be catastrophic as rivals would have free access to their technologies. In contrast, larger firms experience modest increases in product differentiation (lower pairwise similarity) post-Alice. This is consistent with these firms having broad patent portfolios that span technology areas, making it harder for rivals to enter.

The final column (3) in Table 9 examines four-way interactions of size for the two firms in the pair. The results indicate that positive encroachment only occurs when there is a small firm in the pair that is treated by Alice. Indeed, $\text{Small} \times \text{Large} \times \text{Treat} \times \text{Post}$ has a




Fig. 5. Difference in competition between small and large firms.

This figure presents the yearly point estimates for the difference between small and large firms by utilizing Small $\times$ Treatment interacted with year-specific dummies (without Large $\times$ Treatment). The dependent variables are VCF/Sales, Complaints, and R&D/Sales. The year 2013, which immediately precedes the treatment, is designated as the reference year in the regression analysis. The treatment variable is constructed using KPSS values. The black dotted line represents a linear projection based on a regression fitted to the 2011, 2012, and 2013 data points, and extrapolated to cover the years 2015, 2016, and 2017. The shaded gray region corresponds to the 99% confidence interval.

positive coefficient as does Small $\times$ Small $\times$ Treat Small $\times$ Post. However, once the treated firm is a large firm, the coefficient magnitude shrinks by roughly 70% and flips to negative, indicating that larger firms experience much different outcomes. Indeed many scholars argue that patent protection could either be harmful or helpful, and our findings illustrate that firm size is a significant mediator of this outcome.

4.3. Alice and firm performance

We now examine the profitability of firms post-Alice. Table 10 displays panel data regressions that examine whether the sales, profitability and market value of large vs. small firms were differently affected by the Alice decision. In columns (1)–(2), the dependent variable is *Sales Growth*, calculated as the natural logarithm of total sales in the current year t divided by total sales in the previous year $t - 1$. In columns (3) and (4), the dependent variable is *Gross Profit Margins*, which is Sales minus cost of goods sold scaled by sales. In columns (5)–(6), the dependent variable is the *M/B Ratio* (market value of equity plus book debt and preferred stock, all divided by book value of assets).

Table 10 shows that large firms whose patent portfolios are exposed to Alice experience an increase in sales growth and profitability based on the KPSS treatment but results are insignificant for M/B Ratio post-Alice. With one-standard deviation increase in treatment, large firms' sales growth (profitability) increases by 1.5 (1.9) percentage points which is 16.8% (3.8%) of 2013 sales growth (profitability), for the KPSS-based treatment measure. Thus, large firms appear to benefit

some when they are operating in technology markets that experience market-wide losses in patent protection. Our earlier results suggest that these gains at least partially come at the expense of small firms, as large firms face less competition when their smaller rivals scale back.

Consistent with this view, Table 10 shows that small firms experience losses after Alice in the form of decreased profits and lower market valuations. These results are significant at the 1% level. Column (3) shows that a one-sigma shift in the KPSS-based treatment reduces small firm profit margins by 43.7% percent relative to the small-firm pre-Alice mean (0.393). Column (5) shows that small firm M/B ratios decline by 12.5% relative to the analogous pre-Alice mean (2.432). Row three shows that the difference between small and large firm profitability and market value (but not sales growth) are statistically significant.

Fig. 6 examines pre-trends and structural break tests regarding the difference between small and large firms as before. The tests for M/B produce strong findings of (A) no pre-trend and (B) a highly significant structural break with small firms losing value relative to large firms. This result is among the most central findings in our study as firm valuations impound the full future value of the impact of Alice, and impacts relating to innovation are likely to be long-lived (especially because the Supreme Court decision was not expected to become reversed). Our thesis would thus predict stronger results for valuation and these findings support that prediction. As above, we do not find strong results for sales growth, and the results for profitability indicate a structural break that is almost but not quite statistically significant (consistent with profitability effects taking longer to fully manifest).




Fig. 6. Difference in performance between small and large firms.

This figure presents the yearly point estimates for the difference between small and large firms by utilizing $\text{Small} \times \text{Treatment}$ interacted with each year dummies (without $\text{Large} \times \text{Treatment}$). The dependent variables are Sales Growth, Profitability, and M/B. The year 2013, which immediately precedes the treatment, is designated as the reference year in the regression analysis. The treatment variable is constructed using KPSS values. The black dotted line represents a linear projection based on a regression fitted to the 2011, 2012, and 2013 data points, and extrapolated to cover the years 2015, 2016, and 2017. The shaded gray region corresponds to the 99% confidence interval.

4.4. Barriers to entry and financial liquidity

We run subsample tests to further understand why large and small firms experience different outcomes following Alice. Our central hypothesis rests on two economic rationales. (1.) Large firms do better because their pre-existing stocks of tangible capital and higher market shares provide strong barriers to entry that protect their IP. These barriers make it harder for entrants to displace large incumbents even when overall IP protection weakens. (2.) Large firms may do better due to their superior liquidity (higher cash balances). This would increase flexibility and allow larger firms to deploy resources to introduce new technologies.

We test these channels by re-running our analysis using several specific subsamples. For barriers to entry (BTEs), we consider subsamples based on asset non-redeployability (from Kim and Kung (2017), high indicates BTEs),²⁸ high asset tangibility (measured using PPEGT/assets, high indicates BTEs), and TNIC market shares (high indicates BTEs). For financial liquidity, we consider cash balances measured as cash/assets.

Table 11 displays the results. We find that firms with low BTEs are most negatively affected by Alice relative to those with high BTEs. In contrast, we find no differential results for financial liquidity (cash/assets). For each variable in the third row, we report significance levels for differences across the two respective groups. The tests for the three BTE variables indicate many significant differences that match those for our large versus small tests. These results illustrate that large

firms fared better likely because their pre-existing barriers to entry (especially large stocks of tangible capital and barriers to redeployment) allowed them to maintain their pre-existing rents following the loss of IP protection following Alice. In contrast, smaller firms with weaker barriers to entry experienced negative outcomes as they had few alternatives for protecting pre-Alice rents.

4.5. Suggestive evidence on legal impacts

In Table 12, we examine whether patent lawsuits involving small and large firms were differentially affected by the Alice decision. In columns (1)–(4), as described in Section 3.6, we use the Stanford Non-Practicing Entity (NPE) Litigation Database to find NPE and operating company (OC) initiated lawsuits. In columns (5) and (6), we examine whether Alice had an effect on the total number of alleged infringement suits. We find a positive (at the 5% level for KPSS) increase in NPE lawsuits for small firms. This result for small firms, especially when viewed alongside the increased use of non-disclosure agreements post-Alice we report below for these firms, is consistent with the broader “losing-ground scenario” we document. These firms are likely forced to test IP boundaries more (thus increasing lawsuit exposure), and they also attempt to replace lost patent protection using other contracts such as NDAs.

The results are different for large firms, whose exposure significantly *decreases* for lawsuits sourced from both NPEs and OCs after Alice. These results are intuitively interpreted through two impacts. (1.) Alice reduced IP protection, resulting in lawsuits becoming less viable as a means to extract wealth from another party (one needs strong IP to make a claim of infringement successfully). (2.) The gains associated

²⁸ We use *non-redeployability* rather than just *redeployability* to keep the sign consistent with higher values indicating more BTEs.

Table 7
Effects of Alice on patent applications (ModernBERT).

The table displays the results from entropy-balanced panel data regressions in which patent applications are scaled by assets and sales as dependent variables. In columns (1)–(2), the dependent variable is the number of patent applications in year t divided by assets in year $t-1$; and in columns (3) and (4), the dependent variable is the number of patent applications in year t divided by sales in year $t-1$. *Treatment* is the relative impact of the Alice decision on the firm's patent portfolio measured using KPSS dollar values or citations scaled by sales (see Eq. (2) for the formula). The *Treatment* variable was normalized by subtracting the mean and dividing by the standard deviation, resulting in a z-score transformation. In the odd- and even-numbered columns, respectively, we use the KPSS and the number of citations approach to compute the *Patent Value* for the treatment variable. *Small* is a binary variable equal to one if a firm's total assets are less than the median total asset of its TNIC peers in 2013 and is zero otherwise. *Large* is 1-*Small*. *Post* is a dummy variable that equals one if the year is after the Alice decision and zero otherwise. *Small-Large Difference* is the difference between the estimates of $\text{Small} \times \text{Post} \times \text{Treatment}$ and $\text{Large} \times \text{Post} \times \text{Treatment}$. All variables are described in detail in the variable list in Appendix A. All regressions include firm and year fixed effects. Standard errors are clustered at the firm level. t -statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

Dependent variable:	<i>#of Patent App.</i> <i>Assets_{t-1}</i>		<i>#of Patent App.</i> <i>Sales_{t-1}</i>	
	(1)	(2)	(3)	(4)
Small $\times$ Post $\times$ Treatment	−0.006*** (−7.32)	−0.004*** (−4.91)	−0.015*** (−9.79)	−0.008*** (−4.69)
Large $\times$ Post $\times$ Treatment	−0.002*** (−2.71)	−0.003** (−1.99)	−0.002** (−2.07)	−0.004** (−2.26)
Difference	−0.005*** (−4.54)	−0.001 (−0.84)	−0.013*** (−6.54)	−0.004 (−1.49)
Observations	19808	19808	19808	19808
Firm Fixed Effects	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES
Treatment Calculation	KPSS	Citation	KPSS	Citation
Adj. R ²	0.119	0.088	0.107	0.071

with having fewer lawsuits accrued mostly to larger firms who are sued less. Smaller firms, whose ability to defend IP might be limited, were less able to achieve this outcome. Row three shows that, for all three lawsuit variables, the differences between large and small firms are highly significant.

In the final two columns (7) and (8) of Table 12, we consider *Nondisclosure*, which is the total number of 10-K words mentioning “non-disclosure”, “NDA agreements”, “confidentiality agreement” or “secrecy agreement” scaled by the total number of words in the 10-K. The table also shows that small firms use more non-disclosure agreements post-Alice — economically increasing by 32.2% percent of their pre-Alice means. We find no significant changes for large firms, and we find that differences between small and large firms are highly significant. Overall, this evidence suggests that small firms use more nondisclosure agreements to protect their IP after Alice weakened protection from patents most for small firms. This also suggests that NDAs are a potential “second best” contract that can improve IP protection after protection through patents is lost.

Overall, we note that these tests are only suggestive given data limitations. Yet our evidence suggests that large firms appear to benefit, and small firms lose ground, regarding lawsuits and legal protection following the Alice ruling. These results are consistent with large firms having more resources — technological, financial and managerial — which allows them to more effectively protect their product market position post-Alice. In contrast, small firms are less likely to sue, as they face a high risk of losing patent challenges.

4.6. Robustness and placebo tests

Entropy balancing (used in our baseline regressions) is a generalized form of nearest neighbor matching (NNM) that precisely balances characteristics across treated and untreated observations. We also report robustness to using the less sophisticated NNM approach. We match firms on sales and require them to be in the same 4-digit SIC. If there are fewer than 3 matches, we consecutively require that firms are in the same 3-digit SIC and 2-digit SIC industries. We then create seven sales categories: sales < \$20mil, \$20mil ≤ sales < \$100mil, \$100mil ≤ sales < \$500mil, \$500mil ≤ sales < \$2.5bil., \$2.5bil. ≤ sales < \$15bil., \$15bil. ≤ sales < \$100bil., and sales ≥ \$100bil. We require that matches are in the same sales category. We report results in Online Appendix Tables IA17 to Table IA20 and conclude that our findings are robust.

Our Alice Score is estimated using machine learning and has some noise. Although noise generally results in understated results, recent work (de Chaisemartin and D'Haultfoeulle, 2018) suggests a fuzzy difference-in-differences model. In Online Appendix Tables IA21 and Table IA22, we report the local average treatment effect (LATE) from the fuzzydid model (see de Chaisemartin et al. (2019) for Stata implementation) and generally find robust results. For small firms, all main results maintain sign and significance except M/B. M/B also becomes significant and positive for large firms, and some large firm results become insignificant.

We also consider a number of placebo tests that we report in Internet Appendix Tables IA23, IA24, IA25 and IA26. In IA23 and IA24, we randomly swap the firm-level treatment variable among all firms in our sample. In IA25 and IA26, we replace the firm-level treatment variable with random draws from a standard uniform distribution in odd-numbered columns and we replace the firm treatment variable with random draws from a uniform distribution from 0 to the max treatment value in the sample in the even-numbered columns. Results for all of these tests show insignificant “placebo” treatment effects.

Finally, we provide sensitivity analysis for our differences-in-differences analysis using the technique developed in Rambachan and Roth (2023). Instead of requiring that parallel trends hold exactly, the authors impose restrictions on how different post-treatment violations of parallel trends can be from the pre-treatment differences (“pre-trends”). To conduct this analysis, we need to apply a binary treatment. We thus transform our continuous treatment variable into a binary variable by defining treated=0 when the treatment is zero, and treated=1 otherwise. Figures IA7, IA8, and IA9 of the Internet Appendix present results for patent applications and competition measures. These analyses impose that the post-treatment violation of parallel trends is not more than a constant ($\bar{M}$) larger than the maximum violation of parallel trends in the pre-treatment period. The figures show that post-treatment violations are robust to a maximum violation in the pre-treatment.

5. Conclusions

We examine the impact of lost intellectual property protection on firm innovation, performance, market valuations, and competition. We examine firms whose existing granted patents would likely lose protection if challenged following the Alice Corp v. CLS Bank International, 573 U.S. 208 (2014) Supreme Court decision. This decision revoked protection for patents whose fundamental idea is considered abstract with a non-novel transformation. It impacted multiple areas, including business methods, software, and bioinformatics. The outcome of this decision was highly uncertain and unanticipated.

Given that the decision still had ex post uncertainty about whether existing patents would provide sufficient IP protection, we use LLMs to analyze patent textual corpora and assess how much each firm's patent portfolio is exposed to Alice. We document that ex post patenting by firms whose patent stock portfolio is identified as being exposed to

Table 8
Competition and patent protection (ModernBERT).

The table displays the results from entropy-balanced panel data regressions in which competition variables are the dependent variables. In columns (1)–(2), the dependent variable, $VCF/Sales_t$, is the measure of VC entry into a given firm's product market and is the total first-round dollars raised by the 25 startups from Venture Expert whose Venture Expert business description most closely matches the 10-K business description of the focal firm (using cosine similarities) in year t , scaled by focal firm sales in year $t-1$. Complaints is the number of paragraphs in the firm's 10-K that complain about competition divided by the total number of paragraphs in the firm's 10-K. $R\&D/Sales_t$ is R&D expenses in year t scaled by sales in year $t-1$. *Treatment* is the relative impact of the Alice decision on the firm's patent portfolio measured using KPSS dollar values or citations scaled by sales (see Eq. (2) for the formula). The *Treatment* was normalized by subtracting the mean and dividing by the standard deviation, resulting in a z-score transformation. In the odd and even-numbered columns, respectively, we use the KPSS and the number of citations approach to compute the *Patent Value* for the treatment variable. *Small* is a binary variable equal to one if a firm's total assets are less than the median total asset of its TNIC peers in 2013 and is zero otherwise. *Large* is $1 - Small$. *Post* is a dummy variable that equals one if the year is after the Alice decision and zero otherwise. *Small-Large Difference* is the difference between the estimates of $Small \times Post \times Treatment$ and $Large \times Post \times Treatment$. All variables are described in detail in the variable list in [Appendix A](#). All regressions include firm and year fixed effects. Standard errors are clustered at the firm level. *t*-statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

Dependent variable:	$\frac{VCF}{Sales_{t-1}}$		Complaints		$\frac{R\&D}{Sales_{t-1}}$	
	(1)	(2)	(3)	(4)	(5)	(6)
Small $\times$ Post $\times$ Treatment	0.305*** (5.43)	0.113*** (2.94)	0.305* (1.94)	0.175 (1.66)	0.200*** (2.94)	0.075** (2.16)
Large $\times$ Post $\times$ Treatment	0.002*** (2.91)	0.003*** (3.20)	−0.191 (−1.58)	−3.775 (−1.41)	−0.012** (−2.22)	−0.018 (−1.40)
Small-Large Difference	0.304*** (5.40)	0.110*** (2.85)	0.496** (2.50)	3.949 (1.47)	0.211*** (3.10)	0.093** (2.51)
Observations	19 725	19 725	19 727	19 727	19 808	19 808
Firm Fixed Effects	YES	YES	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES	YES	YES
Treatment Calculation	KPSS	Citation	KPSS	Citation	KPSS	Citation
Adj. R^2	0.153	0.061	0.008	0.043	0.091	0.025

Alice significantly decreases for both large and small firms. We find a significant increase in R&D for small firms. These results are consistent with small firms' attempting to replenish their innovative portfolio as predicted by [Aghion et al. \(2005\)](#). Examining ex-post changes in sales growth and profitability along with firm value, we find an asymmetric impact of Alice. Large impacted firms gain, while small firms lose. Exposed large firms experience higher sales growth. Exposed small firms experience decreased market valuations and increased competition through increased venture capital entry.

We show that these differential losses experienced by small firms can be attributed to changes in competition, barriers to entry, and limited legal options for small firms to replace lost IP protection. Small firms face increased competition, but large firms are only minimally impacted. In the post-Alice period, small firms face increased venture capital financed entry into their product space, they increasingly complain about competition, and they face more product market encroachment from rivals. Consistent with the effort to protect IP that has lost protection, small firms also tend to resort more to non-disclosure agreements post-Alice. In contrast, large firms face fewer lawsuits and less product market encroachment. Finally, our subsample analysis shows that firms with strong barriers to entry (high stocks of tangible assets or low asset redeployability) are significantly less impacted by Alice.

Overall, our study provides new insights regarding patent protection when one firm loses a patent, as we examine the impact on both large and small firms when *all* firms lose patent protection in broad areas. In our setting, firms are randomly treated, and participants update their expectations as the court ruling diminishes patent rights for everyone. Our results illustrate an uneven impact of lost IP protection. While we do not conduct a full welfare analysis, the results suggest that patent protection benefits small firms and its withdrawal increases the trend of production being concentrated in larger firms. Our results

are consistent with large firms having more resources – technological, financial and managerial – which allows them to protect their market position. We conclude that patent protection is most important for small firms competing with larger firms.

CRedit authorship contribution statement

Utku U. Acikalin: Writing – original draft, Data Curation, Software, Resources, Methodology, Validation. **Tolga Caskurlu:** Writing – original draft, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Gerard Hoberg:** Writing – original draft, Project administration, Methodology, Investigation, Formal analysis, Conceptualization. **Gordon M. Phillips:** Writing – original draft, Project administration, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Variable definitions

See [Table A.13](#).

Table 9**Firm level competition and encroachment.**

The table displays firm-pair-year panel data regressions in which pairwise product market encroachment (Delta TNIC Score) is the dependent variable. Delta TNIC Score is computed as the change in pairwise TNIC similarity (see Hoberg and Phillips 2016) from year $t-1$ to year t . A large value indicates increased similarity and product market encroachment. To compute the RHS variables, we first sort firms into above and below median sales (relative to TNIC-2 peers) in 2013. We refer to the firms above and below the median as large and small, respectively. The variable $Treat \times Post$ is the sum of the Alice Scores for the two firms in the pair. Analogously, $Large \times Treat \times Post$ is the sum of $Treat \times Post$ when either firm in the pair is large and is zero otherwise. $Small \times Treat \times Post$ is analogously defined for when either firm in the pair is small. The last four variables focus on the properties of the pair. $Large \times Large \times Treat \times Post$ is the sum of $Treat \times Post$ for both firms in the pair when both are large (and is zero otherwise). $Small \times Small \times Treat \times Post$ is analogously defined when both are small. $Large \times Small \times Treat \times Post$ is $Treat \times Post$ for the large firm in the pair when the pair has one large firm and one small firm. $Small \times Large \times Treat \times Post$ is analogously defined using the small firm's treatment. We note that all level effects and Smaller-order interactions are subsumed by the fixed effects and, thus, are not reported. All regressions include firm-pair and year fixed effects and standard errors are clustered at the firm-pair level. t -statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

Dependent variable:	Delta TNIC score		
	(1)	(2)	(3)
$Treat \times Post$	0.0049*** (7.19)		
$Large \times Treat \times Post$		-0.0116*** (-13.17)	
$Small \times Treat \times Post$		0.0283*** (28.00)	
$Large \times Large \times Treat \times Post$			-0.0124*** (-10.17)
$Large \times Small \times Treat \times Post$			-0.0105*** (-8.59)
$Small \times Large \times Treat \times Post$			0.0254*** (19.34)
$Small \times Small \times Treat \times Post$			0.0313*** (20.51)
Observations	6,696,982	6,696,982	6,696,982
Pair Fixed Effects	YES	YES	YES
Year Fixed Effects	YES	YES	YES

Table 10**Profitability (ModernBERT).**

The table displays the results from entropy-balanced panel data regressions that examine whether the profitability of large and small firms are differently affected by the Alice decision. In columns (1)–(2), the dependent variable is sales growth, calculated as the natural logarithm of total sales in the current year t divided by total sales in the previous year $t-1$; and in columns (3) and (4), Gross Profit Margin is sales minus cost of goods sold scaled by lagged sales. In columns (5)–(6), the dependent variable is the M/B Ratio, calculated as the market-to-book ratio (market value of equity plus book debt and preferred stock, all divided by book assets). *Treatment* is the relative impact of the Alice decision on the firm's patent portfolio measured using KPSS dollar values or citations scaled by sales (see Eq. (2) for the formula). The *Treatment* was normalized by subtracting the mean and dividing by the standard deviation, resulting in a z-score transformation. In the odd and even-numbered columns, respectively, we use the KPSS and the number of citations approach to compute the *Patent Value* for treatment. *Small* is a binary variable equal to one if a firm's total assets are less than the median total asset of its TNIC peers in 2013 and is zero otherwise. *Large* is 1-*Small*. *Post* is a dummy variable that equals one if the year is after the Alice decision and zero otherwise. *Small-Large Difference* is the difference between the estimates of $Small \times Post \times Treatment$ and $Large \times Post \times Treatment$. All variables are described in detail in the variable list in Appendix A. All regressions include firm and year fixed effects. Standard errors are clustered at the firm level. t -statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

Dependent variable:	Sales growth		Gross profit margin		M/B ratio	
	(1)	(2)	(3)	(4)	(5)	(6)
$Small \times Post \times Treatment$	0.013 (0.72)	0.002 (0.18)	-0.175*** (-3.12)	-0.082** (-2.48)	-0.305*** (-3.79)	-0.150* (-1.68)
$Large \times Post \times Treatment$	0.015*** (3.63)	0.015** (1.99)	0.019** (2.03)	0.026*** (2.60)	-0.053 (-0.80)	0.128 (0.88)
Small-Large Difference	-0.002 (-0.10)	-0.013 (-1.02)	-0.191*** (-3.38)	-0.106*** (-3.10)	-0.252** (-2.41)	-0.278* (-1.62)
Observations	19 808	19 808	19 808	19 808	19 309	19 309
Firm Fixed Effects	YES	YES	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES	YES	YES
Treatment Calculation	KPSS	Citation	KPSS	Citation	KPSS	Citation
Adj. R^2	0.032	0.035	0.059	0.028	0.050	0.043

Table 11
Barriers to entry and financial liquidity (ModernBERT).

The table displays the results from entropy-balanced panel data regressions where we explore subsample results relating to barriers to entry and financial liquidity using models analogous to our main regressions for *Small* and *Large*. We report results for multiple dependent variables specified in the first column. We only report results for our treatment variable interacted with post and a high and low dummy for each subsample (specified in the headers over each column). For each regression, we thus report these two primary treatment effect coefficients. In the first two columns of results, the low and high are indicator variables based on below and above median assets, respectively, as in our baseline models (these are re-displayed here for comparison). The remaining columns report results in similar pairs and define low and high based on below or above median values of barriers to entry measures, including Asset Redeployability from Kim and Kung (2017) and PPEGT/Assets for low and high levels of tangible assets. Market Share is based on sales and the TNIC-2 industry classification. Regarding liquidity, we define cash as Compustat cash and short-term investments (che) scaled by assets. Due to space restrictions, we only show the results for the KPSS treatment variable. All regressions include controls and firm and year fixed effects, and standard errors are clustered by firm. *t*-statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

Dependent variable:	Assets (Baseline)			Non-Redeployability High = BTEs Stronger			PPEGT/Assets High = BTEs Stronger			Market Share High = BTEs Stronger			Cash/Assets High = More Liquidity		
	Small	Large	S-L	Low	High	L-H	Low	High	L-H	Low	High	L-H	Low	High	L-H
Patent App./Assets	-0.006*** (-7.32)	-0.002*** (-2.71)	***	-0.005*** (-5.60)	-0.007** (-2.51)		-0.006*** (-6.10)	-0.004*** (-2.63)		-0.007*** (-6.79)	-0.001*** (-2.62)	***	-0.003 (-1.50)	-0.006*** (-6.00)	
Patent App./Sales	-0.015*** (-9.79)	-0.002** (-2.07)	***	-0.012*** (-8.12)	-0.013** (-2.33)		-0.014*** (-8.69)	-0.007*** (-2.72)	**	-0.015*** (-10.06)	-0.002*** (-2.80)	***	-0.006 (-1.43)	-0.013*** (-8.85)	
VC Investments	0.305*** (5.43)	0.002*** (2.91)	***	0.323*** (4.97)	0.039 (0.90)	***	0.316*** (4.86)	0.044 (1.37)	***	0.330*** (5.21)	-0.002* (-1.83)	***	0.085 (1.62)	0.296*** (4.53)	**
Complaints	0.295* (1.77)	-0.186 (-1.46)	**	0.154 (0.89)	0.076 (0.39)		0.101 (0.62)	0.055 (0.23)		0.147 (0.83)	-0.085 (-0.52)		0.101 (0.57)	0.088 (0.53)	
Nondisclosure	0.006*** (2.68)	0.000 (0.04)	***	0.004** (2.21)	0.000 (0.07)	**	0.003* (1.93)	0.003 (1.51)		0.004** (2.08)	0.000 (0.85)	*	0.001 (1.27)	0.003* (1.92)	
RD/Sales	0.200*** (2.94)	-0.012** (-2.22)	***	0.202** (2.46)	-0.003 (-0.18)	**	0.197** (2.35)	0.020 (0.60)	**	0.207** (2.43)	-0.003 (-1.01)	**	0.067* (1.69)	0.177** (2.20)	
Sales Growth	0.013 (0.72)	0.015*** (3.63)		0.013 (0.60)	0.019*** (2.62)		0.005 (0.22)	0.030*** (2.81)		0.014 (0.59)	0.015*** (2.80)		0.027*** (2.93)	0.013 (0.60)	
Profitability	-0.172*** (-3.08)	0.019** (2.02)	***	-0.185*** (-2.76)	0.046*** (4.03)	***	-0.185*** (-2.71)	0.005 (0.14)	**	-0.183*** (-2.65)	0.019* (1.92)	***	-0.011 (-0.29)	-0.162** (-2.48)	**
Tobins'Q	-0.305*** (-3.79)	-0.053 (-0.80)	**	-0.295*** (-3.38)	-0.155 (-1.57)		-0.298*** (-3.43)	-0.204** (-2.19)		-0.322*** (-3.80)	-0.047 (-0.48)	**	-0.017 (-0.30)	-0.313*** (-3.71)	***

Table 12
Lawsuits (ModernBERT).

The table displays the results from entropy-balanced panel data regressions examining whether lawsuit frequencies of large and small firms were differently affected by the Alice decision. *Num. of NPE Alleged* is the total number of lawsuits that the firm was alleged by a non-practicing entity in that year. *Num. of OC Alleged* is the number of lawsuits that the firm was alleged by an operating company in that year. *Number of Alleged* is the total number of lawsuits that the firm was alleged in that year. *Nondisclosure* is the total number of 10-K words mentioning “non-disclosure”, “NDA agreements”, “confidentiality agreement” or “secrecy agreement” scaled by the total number of words in the 10-K. *Treatment* is the relative impact of the Alice decision on the firm’s patent portfolio measured using KPSS dollar values or citations scaled by sales (see Eq. (2) for the formula). The *Treatment* was normalized by subtracting the mean and dividing by the standard deviation, resulting in a z-score transformation. In the odd and even-numbered columns, respectively, we use the KPSS and the number of citations approach to compute the *Patent Value* for treatment. *Small* is a binary variable equal to one if a firm’s total assets are smaller than the median total asset of its TNIC peers in 2013, and it is zero otherwise. *Large* is 1-*Small*. *Post* is a dummy variable that equals one if the year is after the Alice decision and zero otherwise. *Small-Large Difference* is the difference between the estimates of Small × Post × Treatment and Large × Post × Treatment. All variables are described in detail in the variable list in Appendix A. All regressions include firm and year fixed effects. Standard errors are clustered at the firm level. *t*-statistics are reported in parentheses; *, **, and *** denote the significance at the level of 10%, 5%, and 1%.

Dependent variable:	Num. of NPE Alleged		Num. OC Alleged		Number of Alleged		Nondisclosure	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Small × Post × Treatment	0.008** (2.14)	-0.002 (-0.33)	0.007 (1.27)	-0.002 (-0.28)	0.017** (2.18)	-0.003 (-0.37)	0.006*** (2.67)	0.002** (2.50)
Large × Post × Treatment	-0.063** (-2.25)	-0.182** (-2.13)	-0.053*** (-2.63)	-0.171*** (-3.41)	-0.124*** (-2.91)	-0.455*** (-5.07)	-0.000 (-0.04)	-0.001* (-1.77)
Small-Large Difference	0.072** (2.52)	0.180** (2.10)	0.060*** (2.87)	0.170*** (3.36)	0.142*** (3.27)	0.452*** (5.01)	0.006*** (2.58)	0.003*** (3.06)
Observations	19 808	19 808	19 808	19 808	19 808	19 808	18 486	18 486
Firm Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES
Year Fixed Effects	YES	YES	YES	YES	YES	YES	YES	YES
Treatment Calculation	KPSS	Citation	KPSS	Citation	KPSS	Citation	KPSS	Citation
Adj. R ²	0.012	0.021	0.018	0.028	0.015	0.034	0.045	0.021

Table A.13
Variable Definitions.

Variable	Definition
Panel A: Firm Characteristics	
Assets	Compustat item AT
Sales	Compustat item SALE
Cash	Compustat item CHE
Gross Profit Margin	Compustat items Sales minus COGS divided by lagged total sales.
M/B Ratio	Compustat sum of market equity (CSHO * PRCC _F), DLC, DLTT, PSTKL, all scaled by lagged book assets.
Sales Growth	Natural logarithm of total sales in the current year t divided by total sales in the previous year t-1.
Log(Age)	Natural logarithm of one plus the current year of observation minus the first year the firm appears in the Compustat database.
Asset Redeployability	Calculated using the procedure from Kim and Kung (2017).
PPEGT	Compustat item Property Plant and Equipment, Gross Total
Market Share	Firm sales/total industry sales using the TNIC-2 industry classification.
Panel B: Innovation, Acquisition & Lawsuit Characteristics	
Alice ModernBERT Score	Alice ModernBERT Score is the predicted probability that the patent would be rejected based on the Alice decision using the ModernBERT LLM.
Treatment	Treatment is a weighted average of a firm's patent values multiplied by the Alice ModernBERT Score scaled by sales. For a firm's patent portfolio, we gather all patents that are valid by the third quarter of 2014. Firm's patent value is calculated in two ways: i) dollar amount provided by KPSS; ii) citations that the patent received. The mathematical notation is provided in Eq. (2).
R&D/Sales _{t-1}	Compustat XRD divided by lagged total sales, set to zero if XRD is missing.
Num. Patent App.	The number of patent applications.
Num. Alleged	It is the number of lawsuits that the firm was alleged for infringing on a patent.
Num. NPE Alleged	It is the number of lawsuits that the firm was alleged by an NPE for infringing on a patent.
Num. OC Alleged	It is the number of lawsuits that the firm was alleged by an OC for infringing on a patent.
Nondisclose	Total number of 10-K words mentioning "non-disclosure", "NDA agreements", "confidentiality agreement" or "secrecy agreement" scaled by the total number of words in the 10-K.
Panel C: Competition Measures	
VCF/Sales _{t-1}	A measure of VC entry in a given firm's product market computed as the total first-round dollars raised by the 25 startups from Venture Expert whose Venture Expert business description is most similar to the 10-K business description of the focal firm (using cosine similarities), scaled by focal firm lagged sales.
Complaints	The number of paragraphs in the firm's 10-K that mention competition divided by the total number of paragraphs in the firm's 10-K.

References

- Adhikari, A., Ram, A., Tang, R., Hamilton, W.L., Lin, J., 2020. Exploring the limits of simple learners in knowledge distillation for document classification with DocBERT. In: *Proceedings of the 5th Workshop on Representation Learning for NLP*. pp. 72–77.
- Aghion, P., Bloom, N., Blundell, R., Griffith, R., Howitt, P., 2005. Competition and innovation: an inverted u relationship. *QJE* 120, 701–728.
- Beltagy, I., Lo, K., Cohan, A., 2019. SciBERT: A pretrained language model for scientific text. In: *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*. pp. 3615–3620.
- Beltagy, I., Peters, M.E., Cohan, A., 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Boldrin, M., Levine, D.K., 2013. The case against patents. *J. Econ. Perspect.* 27 (1), 3–22.
- Boot, A., Vladimirov, V., 2025. Disclosure, patenting, and trade secrecy. *J. Account. Res.* 63 (1), 5–65. <http://dx.doi.org/10.1111/1475-679X.12580>, URL: <https://onlinelibrary.wiley.com/doi/10.1111/1475-679X.12580>.
- Breuer, M., Leuz, C., Vanhaverbeke, S., 2019. Reporting Regulation and Corporate Innovation. Technical Report, National Bureau of Economic Research.
- Budish, E., Roin, B.N., Williams, H., 2015. Do firms underinvest in long-term research? Evidence from cancer clinical trials. *Am. Econ. Rev.* 105 (7), 2044–2085.
- de Chaisemartin, C., D'Haultfoeuille, X., 2018. Fuzzy Differences-in-Differences. *Rev. Econ. Stud.* 85 (2), 999–1028. <http://dx.doi.org/10.1093/restud/rdx049>, URL: <https://doi.org/10.1093/restud/rdx049>, <https://academic.oup.com/restud/article-pdf/85/2/999/24473453/rdx049.pdf>.
- de Chaisemartin, C., D'Haultfoeuille, X., Guyonvarch, Y., 2019. Fuzzy differences-in-differences with stata. *Stata J.* 19 (2), 435–458.
- Devlin, J., Chang, M., Lee, K., Toutanova, K., 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In: *Burstein, J., Doran, C., Solorio, T. (Eds.), Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, pp. 4171–4186. <http://dx.doi.org/10.18653/v1/n19-1423>.
- Esmailzadeh, A., Taghva, K., 2021. Text classification using neural network language model (NNLM) and BERT: An empirical comparison. In: *Intelligent Systems and Applications: Proceedings of the 2021 Intelligent Systems Conference (IntelliSys)*. vol. 296, Springer Nature, p. 175.
- Farre-Mensa, J., Hegde, D., Ljungqvist, A., 2020. What is a patent worth? Evidence from the US patent "lottery". *J. Financ.* 75 (2), 639–682.
- Galasso, A., Schankerman, M., 2015. Patents and cumulative innovation: Causal evidence from the courts. *Q. J. Econ.* 130 (1), 317–369.
- Hainmueller, J., 2012. Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies. *Political Anal.* 20 (1), 25–46. <http://dx.doi.org/10.1093/pan/mpi025>.
- Hall, B., Li, W., 2020. Depreciation of business R&D capital. *Rev. Income Wealth* 66 (1), 161–180.
- Hoberg, G., Phillips, G., 2016. Text-based network industries and endogenous product differentiation. *J. Political Econ.* 124 (5), 1423–1465.
- Hoberg, G., Phillips, G., Prabhala, N., 2014. Product market threats, payouts, and financial flexibility. *J. Financ.* 69 (1), 293–324. <http://dx.doi.org/10.1111/jofi.12050>, URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/jofi.12050>, <https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12050>.
- Kalyan, K.S., Rajasekharan, A., Sangeetha, S., 2021. AMMU: a survey of transformer-based biomedical pretrained language models. *J. Biomed. Informatics* 103982.
- Kesan, J., Wang, R., 2020. Eligible subject matter at the patent office: An empirical study of the influence of alice on patent examiners and patent applicants. *Minn. Law Rev.* 105, 527–617.
- Kim, H., Kung, H., 2017. The asset redeployability channel: How uncertainty affects corporate investment. *Rev. Financ. Stud.* 30 (1), 245–280, URL: <http://www.jstor.org/stable/26142292>.
- Kogan, L., Papanikolaou, D., Seru, A., Stoffman, N., 2017. Technological Innovation, Resource Allocation, and Growth. *Q. J. Econ.* 132 (2), 665–712. <http://dx.doi.org/10.1093/qje/qjw040>, <https://academic.oup.com/qje/article-pdf/132/2/665/30637466/qjw040.pdf>.
- Lemley, M., Zyontz, S., 2021. Does alice target patent trolls? *J. Empir. Leg. Stud.* 18

- (1), 47–89.
- Lerner, J., 2002. 150 years of patent protection. *Am. Econ. Rev.* 92 (2), 221–225.
- Lerner, J., 2009. The empirical impact of intellectual property rights on innovation: Puzzles and clues. *Am. Econ. Rev.* 99 (2), 343–348.
- Lim, D., 2020. The influence of Alice. *Minn. L. Rev. Headnotes* 105, 345.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V., 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lu, Q., Myers, A., Beliveau, S., 2017. Patent prosecution research data: Unlocking office action traits. USPTO Economic Working Paper No. 10, Available at SSRN: <https://ssrn.com/abstract=3024621> or <http://dx.doi.org/10.2139/ssrn.3024621>.
- Maltoudoglou, L., Paisios, A., Lenc, L., Martínek, J., Král, P., Papadopoulos, H., 2022. Well-calibrated confidence measures for multi-label text classification with a large number of labels. *Pattern Recognit.* 122, 108271.
- Mihalcea, R., Tarau, P., 2004. TextRank: Bringing order into text. In: *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*. pp. 404–411.
- Mikolov, T., Chen, K., Corrado, G., Dean, J., 2013. Efficient estimation of word representations in vector space. In: *Bengio, Y., LeCun, Y. (Eds.), 1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2–4, 2013, Workshop Track Proceedings*. URL: <http://arxiv.org/abs/1301.3781>.
- Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., Gao, J., 2021. Deep learning-based text classification: A comprehensive review. *ACM Comput. Surv.* 54 (3), 1–40.
- Nordhaus, W.D., 1969. An economic theory of technological change. *Am. Econ. Rev.* 59 (2), 18–28.
- Rambachan, A., Roth, J., 2023. A more credible approach to parallel trends. *Rev. Econ. Stud.* rdad018.
- Robertson, S., 2004. Understanding inverse document frequency: on theoretical arguments for IDF. *J. Doc.*.
- Roman, M., Shahid, A., Khan, S., Koubaa, A., Yu, L., 2021. Citation intent classification using word embedding. *IEEE Access* 9, 9982–9995.
- Sampat, B., Williams, H.L., 2019. How do patents affect follow-on innovation? Evidence from the human genome. *Am. Econ. Rev.* 109 (1), 203–236.
- Schumpeter, J., 1912. *The Theory of Economic Development*. Harvard University Press, Cambridge, Mass.
- Tübbicke, S., 2022. Entropy balancing for continuous treatments. *J. Econ. Methods* 11 (1), 71–89. <http://dx.doi.org/10.1515/jem-2021-0002>, URL: <https://doi.org/10.1515/jem-2021-0002>.
- Upasani, S., Amin, N., Damania, S., Jadhav, A., Jagtap, A.M., 2020. Automatic summary generation using TextRank based extractive text summarization technique. 07.
- Warner, B., Chaffin, A., Clavié, B., Weller, O., Hallström, O., Taghadouini, S., Gallagher, A., Biswas, R., Ladhak, F., Aarsen, T., Adams, G.T., Howard, J., Poli, I., 2025. Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference. In: *Che, W., Nabende, J., Shutova, E., Pilehvar, M.T. (Eds.), Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Vienna, Austria, pp. 2526–2547. <http://dx.doi.org/10.18653/v1/2025.acl-long.127>, URL: <https://aclanthology.org/2025.acl-long.127/>.
- Williams, H., 2017. Patents and research investments: What inventions are we missing? *Annu. Rev. Econ.* 9.