The Impact of the U.S. Supreme Court ruling on Alice Corp v. CLS Bank on Patenting and Uncertainty in U.S. Innovation[1]

Jesse Frumkin, Nicholas A. Pairolero, Asrat Tesfayesus, and Andrew Toole[2]

This paper establishes a causal effect between the decision in Alice Corp v. CLS Bank and increased legal uncertainty in patent rights. Utilizing a new measure of abstract language in patent documents, and a quasi-natural experiment, we show that the Supreme Court decision in *Alice v CLS Bank (2014)* significantly and persistently increased examiner uncertainty surrounding patentability at the United States Patent and Trademark Office (USPTO). Further, we find that the decision increased the probability of receiving a first action rejection, suggesting that *Alice* changed the standard of patentability.

1. Introduction

Uncertainty negatively affects investment and therefore impacts the business cycle (Bernanke 1983; Dixit and Pindyck 1994; Abel, Dixit, Eberly and Pindyck 1996; Bloom 2009; Czarnitski et al. 2011; Bachman et al. 2013; Bloom 2014; Baker et al. 2016). Since research and development expenditure is likely irreversible, delaying investment until uncertainty is resolved may be beneficial. Using a measure of policy uncertainty based on textual analysis of news articles, Baker et al. 2016 finds that policy uncertainty reduces investment in industries most sensitive to the public sector. More broadly, increases in policy uncertainty anticipate reductions in investment, employment and output in the United States, and other major economies (Baker et al. 2016). Czarnitski et al. (2011) find that market uncertainty leads to lower research and development expenditure, although effective patent rights mitigate the relationship by securing the firm's competitive position.

Legal uncertainty in patent rights may impact investment by affecting the value of patent rights, and therefore the underlying value of research and development expenditures (Sherry and Teece 2004). For example, an innovation in patent rights uncertainty reduces the likelihood that a valid patent is enforceable, and reduces the value of transactions in the markets for

[1] Keywords: Intellectual Property, Patentability, Uncertainty, Abstractness, Examination.
[2] The views expressed are those of the individual authors and do not necessarily reflect official positions of the Office of the Chief Economist or the U.S. Patent and Trademark Office.

technology (Arora et al. x; Gans, Hsu and Stern 2008) and therefore negatively impacts ex-ante returns to investment. Alternatively, patent rights uncertainty increases the likelihood that an invalid patent is enforceable, increasing its value in the markets for technology, and may reduce entry and slow investment (Meurer 2003; Lemley and Shapiro 2005). Reductions in the value of patent rights through increased legal uncertainty reduces their benefit for mitigating the impact of market uncertainty, and therefore may further reduce investment spending (Czarnitski et al. 2011).

We analyze a recent Supreme Court decision modifying the subject matter eligibility of inventions containing abstract ideas, and the resulting variation in examination outcomes by United States Patent and Trademark Office (USPTO) patent examiners. Our main contribution is to go beyond the qualitative and theoretical nature of the literature to provide causal evidence suggesting that judicial decision-making may increase legal uncertainty in patent rights. We find that rather than clarifying the standard of patentability, the Supreme Court decision in *Alice v. CLS Bank (2014)* both substantially and persistently increased uncertainty surrounding the standard of patentability.

The U.S. Supreme Court has been active in recent patent litigation cases. In the *Alice* case, the Supreme Court invalidated several business method patents on the grounds of subject matter ineligibility under the United States code 35 USC § 101. In making this determination, the court applied a two-part test.[3] In the first step, the court analyzed whether the claims[4] were directed to laws of nature, natural phenomena, or abstract ideas; all of which are judicial exceptions to patentability. The court determined that the claims, which refer to computerized escrow arrangements, were directed to abstract ideas. In the second step, the court analyzed whether the claims contained "significantly more" than the abstract idea itself to overcome patent ineligibility. The court found that the claims did not.

[3] What we refer to as the first and second step are respectively referred to as steps 2A and 2B in the Manual of Patent Examination and Procedure, see MPEP § 2106 III https://mpep.uspto.gov/RDMS/MPEP/current#/current/d0e197244.html.
[4] By claims, we mean patent claims. The claims of a patent precisely define the invention.

The *Alice* decision provides specific criteria for developing examiner guidance at the USPTO.[5] Despite this, the decision may have introduced several uncertainties. In particular, while the USPTO provides context-specific examples set by precedential litigation cases, it remains unclear precisely what *abstract ideas* and *significantly more* mean. Qualitatively, this is reflected by sentiment that the Supreme Court "has failed to articulate objective, predictable criteria for allowing judges, patent examiners, or the public to determine whether a claim is drawn to eligible or ineligible subject matter."[6] Consequently, USPTO patent examiners may be inconsistently prosecuting patent applications in light of *Alice*.

Our identification strategy relies on a quasi-natural experiment, new application level data on USPTO patent prosecution, and difference-in-differences to study the impact of the *Alice* decision on patent prosecution in *Alice* at-risk technologies. The richness of recently released USPTO office action data (Lu, Myers and Beliveau 2017) allows us to identify specific reasons for rejection at the application level, and therefore identify the degree of variation in examiner decision making. We exploit internal USPTO data to control for fixed characteristics of examiners that affect prosecution decisions. Additionally, we exploit patent application classifications and abstract language use in patent claims to identify *Alice* at-risk patent applications.

Our results provide two insights on the effect of *Alice*. First, we find that variation in patent examiner rejection outcomes increases significantly and persistently post-*Alice*. This suggests that the *Alice* decision increased examiner uncertainty. Second, we find a positive and significant effect on the probability of receiving a § 101 rejection post-*Alice*. Importantly, our results identify a plausibly exogenous channel to identify the impact of uncertainty in patent

---

[5] The two-part analysis from Mayo Collaborative Services v. Prometheus Laboratories, Inc. applies to both product and process claims, including business methods.

[6] Patent Eligible Subject Matter: Report on Views and Recommendations from the Public (July 2017). Beginning in October of 2016, the USPTO solicited public commentary through a Federal Register Notice on the state of subject matter eligibility law and organized two public roundtables on the matter. These programs attracted large and diverse crowds from the patent community who also submitted a number of written comments. In its July 2017 report based on these public inputs, the USPTO explains that "[m]any members of the public argued that the two-part test provides an unworkable framework for the USPTO to make patent eligibility determinations with any reliability."

rights on economic outcomes. Survey evidence suggests that the recent Supreme Court decisions affecting subject matter eligibility have had a negative, yet modest impact on investment (Taylor 2019); however, the overall effect remains un-quantified, and the precise channels through which the impact occurs are currently unknown. Our measures of the two *Alice* induced channels (the *standard* of patentability, and *uncertainty* over the standard), along with the *Alice* decision as a quasi-natural experiment could be used to disentangle the effects of these channels on investment, and other economic outcomes.

The paper is organized as follows. In Section x, we provide a description of recent litigation decisions affecting the subject matter eligibility of abstract ideas. In Section x, we identify our data sources and section x describes our empirical methodology. In Section x, we provide descriptive results to illustrate changes in the patent prosecution process after *Alice.* In Section x, we report and describe the econometric results. We conclude in Section x.

2. Institutional Background

2.1 Patent Examination

In this section, we provide a brief review of USPTO patent examination since it's used to identify the impact of the *Alice* decision on uncertainty in patent rights. We only provide details on the aspects of examination used in our empirical strategy, although more thorough descriptions exist in the literature (Marco, Toole, Miller and Frumkin 2017; Lu, Myers and Beliveau 2017; Graham, Marco and Miller 2018).

Patent applications contain a variety of information on the invention; for example, a background description, technology classifications, inventors, and patent citations. The patent application text is separated into several parts. For our purposes, we are most interested in the patent claims, since they precisely define the underlying invention in the application. Additionally, each patent application is assigned at least one United States Patent Classification (USPC), which characterizes the underlying technology of the invention. Historically, the USPTO has used the USPC to route applications to examiners, although in 2019 the office started the transition to Cooperative Patent Classification (CPC) based routing. Our data pre dates 2019, so that the examiners in our data are organized within art units at the USPTO, each of which is responsible for certain categories of the USPC.

Patent examiners initially evaluate the validity of a patent application in the First Office Action. Generally, for new applications, the examiner will issue an allowance on the first action, or reject the claims based upon several criteria; including, but not limited to, novelty (35 USC § 102), non-obviousness (35 USC § 103), subject matter eligibility (35 USC § 101), and various clarity and enablement issues (35 USC § 112). For this paper, we are particularly interested in 35 USC § 101 since it defines issues related to subject matter eligibility, including the abstract ideas judicial exception. The next section carefully describes § 101, and the evolving interpretation of subject matter eligibility in recent litigation decisions.

2.2. 35 USC § 101: Subject Matter Eligibility

Patent litigation cases addressing subject matter eligibility under 35 USC § 101 have occurred across multiple centuries. These changes, including what constitutes non-statutory grounds for subject matter ineligibility (i.e. judicial exceptions), have been substantial. While judicial exceptions created in the courts include abstract ideas, laws of nature, and natural phenomena, we focus here on recent decisions affecting *abstract ideas* since it's the most relevant exception in the *Alice* case.

In *Bilski v. Kappos* (2010), a divided court invalidated a business method patent since it was directed to an abstract idea. In *Mayo Collaborative Services v. Prometheus Laboratories* (2012), the court introduced a two-part test when it unanimously decided the patent was not subject matter eligible. First, the court said that the patent, which claimed the administration and patient-specific appropriate dosage of a drug, was directed to a "law of nature" (which is a judicial exception). Second, the court devised a test for "whether the claims do significantly more than simply describe these natural relations," and found that the claims did not add enough to overcome the judicial exception.

While *Mayo* does not involve software, the test devised by the court has been applied in subsequent software-related cases. For example, *Alice* applied the *Mayo* analysis to invalidate patents directed to abstract ideas, indicating that the test extends to method claims. Thus, *Alice* expanded the application of the two-part test devised in *Mayo* and introduced information that changed the understanding of subject matter eligibility in "abstract" technologies.

3. Data

To identify the impact of *Alice* on examiner uncertainty, we use publically available USPTO patent application data. Our underlying sample consists of patent applications and examiner decisions on the first action, used to identify the impact of *Alice* on the probability of receiving a subject matter eligibility rejection at the USPTO. To identify examiner uncertainty over time, we aggregate application level rejection observations to the technology (USPC) time period level.

We use a variety of application characteristics. For each application, where available, we use the USPC at pre-grant publication since the classification at PGPUB is most likely the classification at the time of first action. For the patent applications without a PGPUB (since they were granted before publication), we use the USPC at grant. Finally, for a small number of applications (how many?) we use the latest USPC in Public Pair since the USPTO discontinued putting the USPC on the PGPUB in 2015. We utilize the Patent Claims Research Dataset (Marco, deGrazia and Sarnoff 2016) to control for application level patent scope. Specifically, we use independent claim count (ICC) and shortest independent claim from the claims data. Finally, we use a text-based measure of abstract language at the patent application level in our identification strategy. The source for patent application claims text is the USPTO bulk XML data, as recorded in the RedBook. We extract the full text of both published (i.e. PG-PUBS) and granted patent (i.e. Grants) claims with a publication date between January 6, 2005 through November 10, 2016 and January 4, 2005 through November 15, 2016, respectively. For each patent application, we use the PG-PUB text if available, and the grant text otherwise.

To identify the technologies affected by *Alice* (the treated technologies), we utilize patent applications assessed for validity under the "abstract idea" judicial exception in patent litigation cases post *Alice*. The USPTO Office of Patent Legal Administration (OPLA) maintains a public and comprehensive list of cases from the U.S. Court of Appeals for the Federal Circuit (CAFC) and the U.S. Supreme Court.[7] A team of lawyers monitors these cases and identifies patents litigated for subject matter eligibility. One of the fields recorded for each case is the "Exception Type." This refers to the relevant judicial exception type (i.e. laws of nature, natural phenomena, or abstract ideas). For the purpose of this paper, we consider only the cases directed toward abstract ideas. Patent applications from this set of USPCs form the treatment group.

Our outcome variables are based on patent examination decisions at the USPTO. We obtain examiner § 101 rejection data from recently released USPTO data (Lu, Myers and Beliveau 2017) that provides comprehensive information on examiner-issued rejections at any stage of prosecution. The data covers office actions mailed during the period 2008 to mid-2017.[8] This dataset, containing over 4.4 million office actions and over 10 million rejection observations, indicates whether the office action is a non-final or final rejection and provides a description of the rejection type. This identifies each of the sections from 35 USC that were applied as grounds for rejection in the Office action. The rejection type of interest is 35 USC § 101 since it is the section governing subject matter eligibility.

To identify the mailing date of the first action, we use the transaction history in the Patent Examination Research Dataset [PatEx] (Graham, Marco and Miller 2018). From PatEx, we extract all applications that contain instances of any of four office action types along with their mailing dates. The office action types are Non-Final Rejection, Final Rejection, Ex Parte Quayle, and Notice of Allowance. We merge the collected data with the Office action dataset using application number and the mailing date. We then map the PatEx first Office action to its corresponding § 101 rejection from the Office actions data, if one exists. This provides a dataset containing the first Office action date for each application, along with whether or not the application was rejected under § 101. Additionally, we use examiner fixed effects for our application level regressions. We identify the examiner issuing the first action using internal

[7] An updated list of these cases can be found at https://www.uspto.gov/sites/default/files/documents/ieg-mar-2017-sme_crt_dec.xlsx.

[8] Specifically, this dataset covers office actions for applications in the 12, 13, 14 and 15 series. Series identifiers for USPTO patent applications are the first two digits of the application number.

USPTO data. This patent application level dataset is used to estimate the impact of *Alice* on the probability of receiving a § 101 rejection on the first action.

To calculate the variance of examiner § 101 rejection rates (used to capture examiner uncertainty), first we compute the examiner-level rate of first action § 101 rejections for each half-year and USPC. We then compute the variance of the resulting examiner § 101 rejection rates within each USPC half year. We also record the variance of each examiner's § 101 rate in each USPC half year, and extract the average of these variances for each half-year and USPC. We use this variable to control for increases in § 101 rejection variance that stem directly from increased uncertainty over the examiner's § 101 rates themselves. Finally, we extract the number of examiners, and the number of applications per USPC half-year period as additional controls for the variance regressions.

4. Methodology

This paper uses a novel methodology to causally identify the impact of *Alice* on legal uncertainty. Specifically, we use appellate court case patent classifications and analyze the abstract content of patent claims to evaluate the treatment effect of the *Alice* decision. The next section describes the construction of the abstractness measure for patent claims text.

4.1 Abstractness Measure

Our measure of abstract language in patent claims is derived from WordNet, a lexical database developed in 1985 in the Cognitive Science Laboratory at Princeton University. The purpose of WordNet is to organize content based on the conceptual relationship of words. The words in WordNet are arranged in a *hypernym/hyponym* hierarchy. *Hyponym* words can be thought of as a subset of a given *hypernym* word in the sense that the former is broader than the latter. For example, a *hyponym* of the *synset* containing automobile is "ambulance", while a *hypernym* is "motor vehicle." In other words, *hyponyms* generally identify a narrower or more specific instance of a given *hypernym*. The broadest noun layer in WordNet is "entity", which is defined as "that which is perceived or inferred to have its own distinct existence (living or nonliving)." The next layer below "entity" or its direct *hyponyms* are "physical entity," "abstraction, abstract entity," and "thing." Of particular relevance for this paper is the *synset* "abstraction, abstract entity," defined as "a general concept formed by extracting common features from specific examples. All *hyponyms* of the *synset* "abstraction, abstract entity" comprise our primary set of our abstract words. The set contains 53,815 words. In the Appendix, we examine an example from the *Alice* claims that demonstrates the effectiveness of using WordNet for the measurement of abstract language use in patent documents.

At present, WordNet contains 155,287 unique words and 117,659 *synsets* for a total of 206,941 word-sense pairs. Despite this impressive coverage, an important feature of WordNet is that some of the words are from non-technical corpus. Thus, while our descriptive evidence in the Appendix indicates that the "abstraction, abstract entity" *synset* closely corresponds to abstract language in patent claims; its raw content may not systematically and comprehensively map to technical patent text. Further, WordNet might contain abstract language not yet deemed problematic by the courts, and therefore might be overly inclusive. For these reasons, we refine the abstract language from WordNet by intersecting the words derived from the WordNet *synset* "abstraction, abstract entity" and the words from the claims in the list of litigated patents for "abstract ideas" described above.

Our abstractness measure of patent claims relies on the natural language processing approach called term frequency inverse document frequency cosine similarity (or more concisely, *tfidf*). To compute this measure, first we concatenate each word (with a space between each word) from the abstract language set into one document. Second, we compare the patent claims from a given patent application to the abstract document using *tfidf.* [9] This procedure generates the cosine similarity that quantifies the similarity between the abstract words set and a given patent document.

To validate the abstractness measure, we compare it to text classification models that predict the probability of receiving a first action § 101 rejection during USPTO patent examination. A correlation between these two independently derived measures serves as a

[9] A detailed definition of *tfidf* cosine similarity is provided in the Appendix.

reliable validation of the abstractness measure; especially considering the fact that the prediction models do not distinguish between the different types of § 101 rejections. We find a strong positive correlation, suggesting that our WordNet based abstraction measure is an adequate proxy for language relevant to the "abstract ideas" judicial exception.[10]

4.2 Alice as a Quasi Natural Experiment

The Supreme Court decided *Alice Corp v. CLS Bank* on June 19, 2014; however, the case made its way through the lower courts in the preceding years. Despite this, the Supreme Court decision in *Alice* is plausibly exogenous because of significant disagreement in the lower courts. First, the district court found all of the *Alice* patents invalid. Then, the CAFC reversed the district court's decision, only to reverse course and affirm the district court in an En banc decision. Even though the CAFC eventually affirmed the district court, there was significant uncertainty in the affirmation[11]. The Supreme Court then affirmed the second CAFC decision. Disagreement in the lower courts, uncertainty involved in the CAFC decision, and the very fact that enough uncertainty existed in the CAFC decision to support an appeal to the Supreme Court implies that the Supreme Court decision is plausibly exogenous.

The Effect of Alice on Legal Uncertainty

[10] First, we estimate text classification models to obtain predicted § 101 rejection probabilities based on text contained in the patent documents. In this exercise, we use both logistic regression and multinomial naïve Bayes models for text classification of the § 101 first action rejections. These models serve as standard baselines in the machine learning literature on text classification (Wang et al. 2012). Second, the predictions from these models are compared to the abstractness measure for each patent application. We find a strong positive correlation. This is especially significant since the abstractness measure only covers one reason for issuing a § 101 rejection (i.e. abstract ideas), whereas the text classification algorithms are based upon all possible reasons for issuing a § 101 rejection. We find that the abstractness measure has a correlation of 0.3097 and 0.4101 with the output from the multinomial naïve bayes and the logistic regression predictions, respectively.

[11] In particular, as noted in CAFC decision, "While Chief Judge Rader is correct to note that no single opinion issued today commands a majority, seven of the ten members, a majority, of this en banc court have agreed that the method and computer-readable medium claims before us fail to recite patent-eligible subject matter. In addition, eight judges, a majority, have concluded that the particular method, medium, and system claims at issue in this case should rise or fall together in the § 101 analysis (717 F.3d 1269 (2013))."

To estimate the impact of *Alice* on examination uncertainty, we use a difference-in-differences model given by

$$Var_{kt} = \beta_0 + \beta_1(PatGrp_{kt} * Alice_t) + \delta_k + \gamma_t + \boldsymbol{X}_{kt}\alpha + \epsilon_{kt} \quad (1)$$

where k indicates the technology area; t is time in half-year intervals; δ_k is a technology (USPC class) fixed effect; γ_t is a half-year time fixed effect; $Alice_t$ is a dummy variable taking a value of one after the *Alice* decision; and $\boldsymbol{X}_{kt}$ is a vector of USPC-level controls. $PatGrp_{it}$ is a dummy variable indicating whether or not the application belongs to a USPC in the treatment group. Recall that the treatment group is composed of USPCs of litigated patents on the grounds of "abstract ideas". The dependent variable Var_{kt} is the variance of receiving a § 101 non-final rejection across examiners in technology k in period t. Specifically, we measure the variance in examiner § 101 rates within each USPC half-year. The variable measures the spread in examiner average § 101 rates. All else equal, greater variation is reflective of larger differences in examiner behavior.

In order to identify the uncertainty effect, we assume the assignment process of applications to examiners is constant over time. This assumption is significantly weaker than the 'random assignment' argument required in recent statistical work in the economics of intellectual property.[12] Our identification strategy would be invalid if application assignment changed because of *Alice*. The only nonrandom assignment noted in the literature is assignment based on technological specialization (Lemley and Sampat 2012; Righi and Simcoe 2018), and therefore we are comfortable assuming that assignment policies are constant over time. However, we add an additional control for robustness. In particular, we control for the variance of incoming application abstract language across examiners within each half-year USPC. The control serves two purposes. First, if applications with varying abstract language are differentially distributed after *Alice*, then we would expect § 101 variance across examiners to mechanistically increase, regardless of any change in uncertainty. Second, if examiners differentially receive applications with varying degrees of abstract language, and *Alice* changed how they treat these

[12] Lemley et al. 2012; Frakes and Wasserman 2017; Williams 2013; Farre-Mensa, Hegde and Ljungqvist 2017; Righi and Simcoe 2018; Sampat and Williams 2019.

applications, then we would expect § 101 variance to again mechanistically increase. Controlling for the variance in incoming abstract language alleviates these concerns.

We add several other controls since the examiner variation measure may be influenced in other ways unrelated to actual examiner uncertainty. We control for within examiner § 101 rejection rate variance (since the standard errors of each examiner § 101 rate affects the variance across examiners), and the number of examiners and applications within each USPC at time t (again, because the number of examiners and applications affects the variance across examiners).

A potential source of endogeneity for identifying examiner uncertainty post *Alice* is applicant behavior. In particular, applicants may react to the decision by sending in fundamentally different applications, and prosecute applications differently post *Alice*. To alleviate these endogeneity concerns, we restrict our sample to applications filed before the *Alice* decision. This restriction ensures that applicant drafting strategies post Alice do not influence the patent applications in our sample. Second, our primary dependent variables of interest are at the first office action. The first office action describes the initial decision by the examiner and is completed with minimal applicant interaction. This further reduces the influence of applicant behavior post *Alice*.

Even though the *Alice* decision is plausibly exogenous, assignment to the treatment group is not random since it relies on the technology of the patent application. To alleviate concerns of unobserved heterogeneity, we control for a variety of patent application characteristics $\boldsymbol{X}_{it}$ (including patent claim scope ICC and ICL). As an additional application level co-variate, for our second specification, we also include the average abstractness measure. The specification includes the additional variable Avg_Abst_i as follows:

$$Var_{kt} = \beta_0 + \beta_1(PatGrp_{kt} * Alice_t) + \beta_2 Avg\_Abst_i + \delta_k + \gamma_t + \boldsymbol{X}_{kt}\alpha + \epsilon_{kt} \quad (2)$$

The average abstractness variable is the average abstractness of the patent applications receiving a first action in the USPC at time t. Thus, β_1 is the effect of abstract language content on examiner uncertainty.

Finally, as a second complementary identification strategy, we add a term that interacts the USPC group dummy with the average abstractness measure and the *Alice* indicator. Using this specification, we quantify the extent to which the *Alice* decision increased examiner uncertainty, specifically through abstract language in more vulnerable USPC classes.

$$Var_{kt} = \beta_0 + \beta_1(PatGrp_{kt} * Alice_t) + \beta_2 Avg\_Abst_i + \beta_3(PatGrp_{it} * Avg\_Abst_i * Alice_t) + \delta_k + \gamma_t + \boldsymbol{X}_{kt}\alpha + \epsilon_{kt} \quad (3)$$

To assess the difference-in-differences pre-trend assumption, we allow the treatment group to affect variation in examiner rejection outcomes for several periods before Alice. We use a flexible model to separate treatment coefficients for each half-year t for the period 2008-2017. In these estimations, we use the most flexible specification and the coefficients of interest are the interactions between the treatment variable and time. The equation is given by

$$Var_{kt} = \beta_0 + \sum_t \beta_t(PatGrp_{it} * Alice_t) + \psi_1 Avg\_Abst_i + \psi_2(PatGrp_{it} * Avg\_Abst_i * Alice_t) + \delta_k + \theta_j + \gamma_t + \boldsymbol{X}_{it}\alpha + \epsilon_{it} \quad (4)$$

The Effect of *Alice* on First-Action § 101 Rejections

To estimate the effect of *Alice* on the probability of receiving a § 101 rejection, we employ a linear probability difference-in-differences fixed effect model. The baseline estimating equation is

$$pr(101)_{it} = \beta_0 + \beta_1(PatGrp_{it} * Alice_t) + \delta_k + \theta_j + \gamma_t + \boldsymbol{X}_{it}\alpha + \epsilon_{it} \quad (5)$$

where i indicates the patent document; j indicates the examiner assigned to the patent document; k indicates the technology area by USPC; t is time; $pr(101)_{it}$ is a binary variable that indicates whether the application receives a § 101 rejection on the first action; δ_k is a technology (USPC class) fixed effect; γ_t is a time fixed effect; θ_j is an examiner fixed effect; and $\boldsymbol{X}_{it}$ is a vector of patent characteristics. β_1 is the treatment effect for the at-risk *Alice* technology group on the probability of receiving a § 101 rejection. We use robust standard errors since the errors in linear probability models are heteroskedastic by construction. We also estimate equivalent models to

(2), (3) and (4) that add the abstractness measure, interact the abstractness measure with the post-*Alice* and treatment group indicators, and test for equivalent pre-trends.

5 Results

5.1 Descriptive Analysis

Before providing the formal regression results, we examine the impact of *Alice* on examination uncertainty and the probability of receiving a § 101 rejection descriptively. For examiner uncertainty, we compute the treatment and control average examiner variance measure for each half-year across the USPCs. Figure 1 displays the results. From the figure, the pre-trend assumption of difference and differences is satisfied since both the treatment and control show similar trends before the *Alice* decision. Notably, in the period after *Alice,* the examiner variance measure increases in the treatment relative to the control. Further, the change in trend continues throughout the end of the sample. This provides preliminary evidence that the *Alice* decision persistently increased examiner uncertainty in *Alice* at-risk technologies. Additionally, in the online appendix (figure 3) we provide the graph for the mean probability of receiving a § 101 rejection on the first action for both the treatment and control groups.[13] Similarly, we see a sharp increase in the probability of receiving a § 101 rejection post *Alice*.

[13] We use month years for the probability of receiving a § 101 rejection on the first action. The reason for this choice is described in the regression results section.

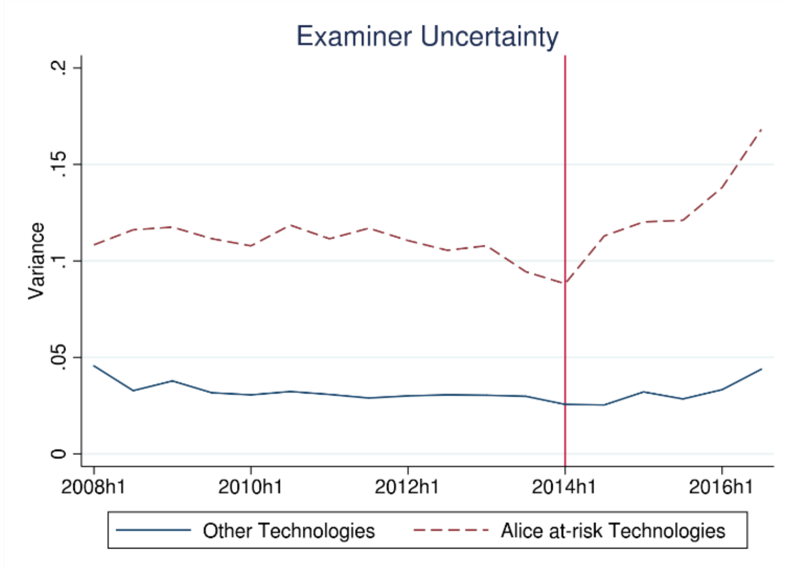

Figure 1: Examiner Uncertainty in *Alice* at-risk technologies relative to other technologies. Examiner uncertainty increases relative to the other technologies immediately after the *Alice* decision on June 19th 2014. Further, the increase in trend is persistent throughout the end of the sample.

5.2 The Effect of Alice on Legal Uncertainty

The first set of results examine the impact of *Alice* on variation in examination outcomes. Table 2 reports the USPC-level variance in § 101 rejection rates across examiners. Column 1 contains the results from equation (1), where the variable of interest is the interaction term between the *Alice*-sensitive USPC indicator and the Post-*Alice* decision indicator. The coefficient is positive and statistically significant. Specifically, the variance across examiner § 101 rates increases by approximately 1/3 of a pre-Alice standard deviation in at-risk *Alice* technologies (.0223/.064; See table 1 for pre-*Alice* standard deviations). The result shows that *Alice* substantially increased the variance across examiner § 101 rejection rates, suggesting that the decision increased legal uncertainty in patentability.

Column 2 reports the results from equation (2), with the USPC-level average abstractness variable. The coefficient is not statistically significant. Therefore, the average abstractness of the USPC does not influence the § 101 variance across examiners on average across the entire sample. The treatment effect (the at-risk group post *Alice*) is virtually unchanged from specification (1). Finally, in column 3, we report the results from equation (3) with the triple interaction term. We find no statistically significant effect on the triple interaction. These results suggest that examiner uncertainty is not sensitive to higher levels of abstractness in the patent documents. This is intuitive since extreme levels (either too low or too high) of abstractness in patent documents should make the determination of subject matter eligibility less ambiguous. On the other hand, patent documents in *Alice*-sensitive USPCs that have average levels of abstractness may be more likely to elicit disagreement from examiners. Despite this, the treatment effect (the at-risk group post *Alice*) is positive and statistically significant in all regressions.

Of additional interest are the variance controls in the regressions. As expected, the number of examiners, and the average number of applications per examiner reduce the variation in § 101 rejection rates across examiners. Also, as expected, higher standard errors for the within examiner 101 rates, and greater variation in incoming abstraction across examiners all increase the variation in § 101 rejection rates across examiners. As discussed earlier, these controls are important to eliminate heterogeneity that mechanistically increases variance across examiners regardless of changes in legal uncertainty. Additionally, the variation in examiner incoming abstraction is an important control in case the assumption of constant application assignment over time does not hold.

Figure 2 shows the coefficient plot from equation (4). Specifically, we report the interactions between the *Alice*-sensitive USPC group and half-year interval indicators. The base period is the half year directly preceding *Alice* (the first half of 2014). There is a substantial and persistent increase immediately after *Alice*. Further, the pre-trend assumption is satisfied throughout the entire pre-event period (none of the pre-*Alice* coefficients are statistically different from zero). We observe no change in variance, even for the period surrounding the *Bilski* decision in 2010 (which also directly affected Business Methods patents).

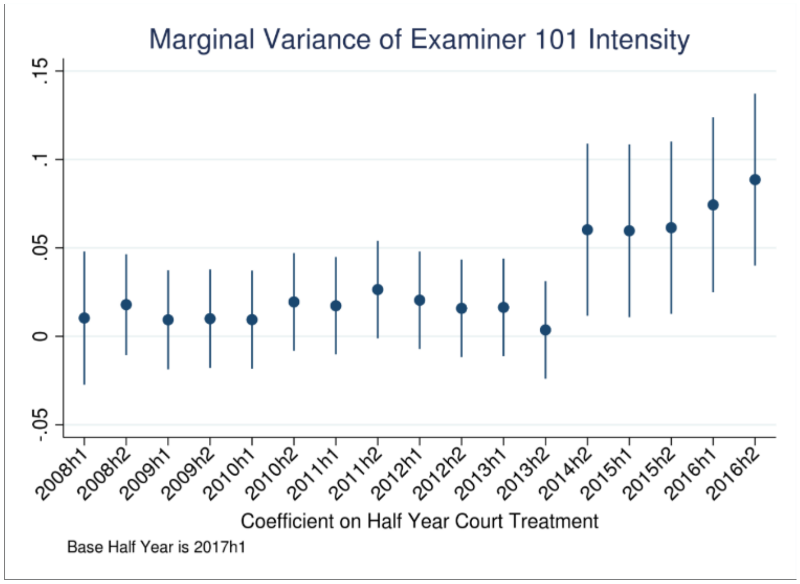

Figure 2: Examiner uncertainty in *Alice* at-risk technologies relative to other technologies. This figure shows the coefficients for PatGrp interacted with a dummy variable for each half year time period in model (3), rather than interacting PatGrp with the after *Alice* indicator aalice. Importantly, the pre-trend is statistically indistinguishable from zero, where the base period is the first half of 2014. Further, examiner uncertainty increases in the *Alice* at-risk technologies post *Alice*.

5.3 The Effect of *Alice* on First-Action § 101 Rejections

Next, we discuss the impact of *Alice* on the probability of receiving a § 101 rejection on the first action. Even though *Bilski* didn't affect examination uncertainty in at-risk *Alice* technologies, *Bilski* did impact the probability of receiving a § 101 rejection in *Alice* at-risk technologies.[14] This leads to a non-stable half-year pre-trend. Because of this, we restrict the sample to every observation with first action one year before *Alice*. This ensures that no other Supreme Court decision affecting subject matter eligibility occurs in the pre-*Alice* period, the most recent before *Alice* being *Myriad* (decided on June 13, 2013). Additionally, we use month year fixed effects.

[14] Figure 6 in the online appendix graphs the half-year interactions with the treatment indicator from regression model (8). Notice that the pre-trend in figure 6 looks very similar to the trend in examination citations to the *Bilski* case pre-*Alice,* as seen in figure 4, also in the online appendix. Although the focus of this paper is on the impact of *Alice* on legal uncertainty, it's interesting to note that *Bilski* had an impact on the probability of receiving a § 101 rejection, but not on examination uncertainty; especially since *Bilski* and *Alice* affected similar technologies (Business Methods).

Table 3 identifies the impact of *Alice* on the probability of receiving a § 101 rejection. Column 1 reports the baseline estimation from equation 5. The coefficient β_1 on the interaction term ($POST \times PatGrp_{it}$) is 0.0584, and is statistically significant. The result suggests that *Alice* increased the probability of receiving a § 101 rejection on the first action by 5.84 percentage points for a patent document in an *Alice*-sensitive USPC. Since the average probability of receiving a § 101 rejection pre-*Alice* of such documents was 27.4 percent, this implies that *Alice* increased the probability a § 101 rejection by 12.5 percent for patent documents in *Alice*-sensitive USPCs.

Column (2) contains the document-level abstractness measure as a separate variable. While the coefficient on the interaction term remains unchanged, the coefficient on the abstractness term is 0.600. This suggests that one standard deviation increase in the abstractness measure leads to a 0.6 percentage point increase in the probability of getting a § 101 rejection.

Column 3 reports the results from the third specification containing the triple interaction between the *Alice*-sensitive USPC dummy, the abstractness measure, and the post-*Alice* dummy. The triple interaction is statistically significant with a coefficient of 1.335. Since the pre-*Alice* average abstractness for applications in *Alice*-sensitive USPCs was 0.0255, the estimate suggests that the *Alice* decision increased the probability of receiving a § 101 rejection on the first action by 4.5 (i.e. 1.335*0.0255) percentage points for a document with average abstractness. This is an increase of 17.8 percent in the probability of receiving a § 101 rejection due to the *Alice* decision. Further, the coefficient on the post-*Alice* dummy and *Alice*-sensitive USPC interaction term is now .0252 and statistically significant. This suggests that the effect of the *Alice* decision on *Alice*-sensitive patent documents occurs through both the technology, as well as the abstractness content of the document itself. This contrasts with the uncertainty regressions, where a marginal increase in abstractness did not affect the increase in uncertainty post-*Alice*. More importantly, this suggests that the abstractness of patent applications is unable to capture the entire *Alice* effect. This is not surprising since the *Alice* decision utilized the two-step test. First, the invention must be directed to an 'abstract idea'. Second, the application should only be given a § 101 rejection if the invention is not directed to 'significantly more' than the 'abstract idea' itself. The abstractness measure does not capture the 'significantly more' aspect of the test.

Figure 5 in the online appendix reports the point estimates for each time interval (month-year) from the second model. Specifically, using 2014m6 as the base period, we report the coefficients for the *Alice*-sensitive USPC dummy variable interacted with each month-year dummy during the period 2013m6-2015m6. There is a sharp increase in the treatment affect in the period immediately following *Alice* (2014m7). The one-year pre-*Alice* is not statistically different from zero in any month.

5. Conclusion

Legal uncertainty in patent rights may affect investment, and other economic outcomes. Although scholars and policy makers have discussed the possibility for ambiguous judicial decisions to increase legal uncertainty in patentability, causal empirical evidence is missing from the literature. The primary contribution of this paper is identifying exogenous variation in patent rights uncertainty driven by judicial decision-making. The measures derived in this paper (for patent uncertainty and the standard of patentability), and *Alice* as a quasi-natural experiment could be used to test theoretical assertions in the literature on the impact of patent uncertainty on economic outcomes.

In June 2014, the Supreme Court decided *Alice v. CLS Bank*. The court applied the two-part test from *Mayo* to find the business methods patents subject matter ineligible. Beyond this, many in the patent community contend that the decision increased legal uncertainty about patent subject matter eligibility. Our empirical results show that *Alice* did increase examiner uncertainty surrounding subject matter eligibility. Further, the increase is persistent. Additionally, we showed that *Alice* increased the probability of receiving a rejection on the first action.

Our paper is first in quantifying an increase in legal uncertainty from ambiguous judicial decision-making. Further research is needed to identify the tradeoff between modifying the standard of patentability, and the resulting change in legal uncertainty from judicial decision-making. This requires a careful theoretical and empirical consideration of the economic impacts of both of these channels. The exogenous sources of variation identified in this paper should be useful for this purpose.

References

Abel A., Dixit A., Eberly J., Pindyck R., "Options, the Value of Capital, and Investment", Quarterly Journal of Economics, 1996

Acemoglu D., Finkelstein A., "Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector", Journal of Political Economy, 2008

Ai C., Norton E., "Interaction Terms in Logit and Probit Models", Economic Letters, 2003

Baker S., Bloom N., Davis S., "Measuring Economic Policy Uncertainty", The Quarterly Journal of Economics, 2016

Bar-Ilan A., Strange W., "Investment Lags", American Economic Review, 1996

Bessen J., Hunt RM., "An Empirical Look at Software Patents", Journal of Economics and Management Strategy, 2007

Bessen J., "A Generation of Software Patents", SSRN, 2011

Bloom N., "Fluctuations in Uncertainty", The Journal of Economic Perspectives, 2014

Bloom N., Bond S., Van Reenen J., "Uncertainty and Investment Dynamics", Review of Economic Studies, 2007

Bloom N., Floetotto M., Jaimovich N., Saporta-Eksten I., Terry S., "Really Uncertain Business Cycles", Econometrica, 2018

Bloom N., Van Reenen J., "Patents, Real Options, and Firm Performance", The Economic Journal

Brynjolfsson E., McAfee A., "The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies", W.W. Norton and Company, 2016

Choi J., "Patent Litigation as an Information – Transmission Mechanism", The American Economic Review, 1998

Cockburn I., Henderson R., Stern S., "The Impact of Artificial Intelligence on Innovation", NBER, 2018

Czarnitzki D., Toole A., "Patent Protection, Market Uncertainty, and R&D Investment", The Review of Economics and Statistics, 2011

Dixit A., Dixit R., Pindyck R., "Investment Under Uncertainty", Princeton University Press, 1994

Farre-Mensa J., Hegde D., Ljungqvist A., "What is a Patent Worth? Evidence from the U.S. Patent 'Lottery'", SSRN, 2017

Farrell J., Merges R., "Incentives to Challenge and Defend Patents: Why Litigation Won't Reliably Fix Patent Office Errors and Why Administrative Patent Review Might Help", Berkeley Technology Law Journal, 2004

Farrell J., Shapiro C., "How Strong are Weak Patents?, American Economic Review, 2008

Frakes M, Wasserman M., "Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents? Evidence from Microlevel Application Data", Review of Economics and Statistics, 2017

Gans J., Hsu D., Stern S., "The Impact of Uncertain Intellectual Property Rights on the Market for Ideas: Evidence from Patent Grant Delays", Management Science, 2008

Graham S., Marco A., Miller R., "The USPTO Patent Examination Research Dataset: A Window on Patent Processing", Journal of Economics and Management Strategy, 2018

Gilbert R., "Antitrust for Patent Pools: A Century of Policy Evolution." Stanford Technology Law Review, 2004

Hall B., "Business and Financial Method Patents, Innovation and Policy", Scottish Journal of Political Economy, 2009

Hall B., MacGarvie M., "The Private Value of Software Patents", Research Policy, 2010

Hirshleifer J., Riley, J., "The Analytics of Uncertainty and Information", Cambridge Surveys of Economic Literature, Cambridge University Press, 1992

Holman C., "Unpredictability in Patent Law and its Effect on Pharmaceutical Innovation", Mo. Law Review, 2011

Jovanovic B., Rousseau P., "General Purpose Technologies", Handbook of Economic Growth, 2005

Kraft H., Schwartz E., Weiss F., "Growth Options and Firm Valuation", European Financial Management, 2017

Lei Z., Wright B., "Why Weak Patents? Testing the Examiner Ignorance Hypothesis", Journal of Public Economics, 2017

Leahy J., Whited T., "The Effect of Uncertainty on Investment: Some Stylized Facts", Journal of Money, Credit and Banking, 1996

Lemley M., "Rational Ignorance at the Patent Office", Northwestern University Law Review, 2001

Lemley M., Sampat B., "Examiner Characteristics and Patent Office Outcomes", Review of Economics and Statistics, 2012

Lemely M., Shapiro C., "Probabilistic Patents", Journal of Economic Perspectives, 2005

Lu Q., Myers A., Beliveau S., "USPTO Patent Prosecution Research Data: Unlocking Office Action Traits", SSRN, 2017

Marco A., Sarnoff J., deGrazia C., "Patent Claims and Patent Scope", SSRN, 2016

Miller G., "WordNet: A Lexical Database for English", Communications of the ACM Vol. 38, No. 11: 39-41, 1995

Metsis V., Androutsopoulos I., Paliouras G., “Spam Filtering and Naïve Bayes – which Naïve Bayes”, Proceedings of CEAS, 2006

Meurer M., “Controlling Opportunistic and Anti-Competitive Intellectual Property Litigation”, Boston College Law Review, 2003

Paddock J., Siegel D., Smith J., “Option Valuation of Claims on Real Assets: The Case of Offshore Petroleum Leases”, Quarterly Journal of Economics, 1988

Puhani P., “The Treatment Effect, the Cross Difference, and the Interaction Term in Nonlinear Difference in Differences Models”, Economic Letters, 2012

Righi C., Simcoe T., “Patent Examiner Specialization”, NBER, 2017

Sampat B., Williams H., “How do Patents Affect Follow-on Innovation? Evidence from the Human Genome”, American Economic Review, 2019

Sarkar D., “Text Analytics with Python”, APress, 2016

Sherry E., Teece D., “Royalties, Evolving Patent Rights, and the Value of Innovation”, Research Policy, 2004

Taylor D., “Patent Eligibility and Investment”, SSRN, 2019 [article says Forthcoming]

Wang S., Manning C., “Baselines and Bigrams: Simple, Good Sentiment and Topic Classification”, Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics, 2012

Williams H., “Intellectual Property Rights and Innovation: Evidence from the Human Genome”, Journal of Political Economy, 2013

Wen W., Ceccagnoli M., Forman C., “Opening Up Intellectual Property Strategy: Implications for Open Source Software Entry by Start-up Firms”, Management Science, 2016

Appendix: Tables

Summary Statistics

	Alice-Insensitive USPC		*Alice-sensitive* USPC		All USPC	
	mean	std dev	mean	std dev	mean	std dev
Application Abstractness [before *Alice*]	0.017	0.01	0.0255	0.012	0.019	0.011
§ 101 Proportion [before *Alice*]	6.80%	25.30%	27.40%	44.60%	11.30%	31.70%
§ 101 Proportion [after *Alice*]	7.90%	27.00%	32.90%	47.00%	13.60%	34.30%
§ 101 Examiner Variance [before *Alice]*	0.031	0.057	0.109	0.064	0.036	0.06
§ 101 Examiner Variance [after *Alice]*	0.034	0.06	0.134	0.07	0.04	0.067

Table 1

Impact of *Alice* on the variance of USPC-level § 101 rejection intensity

	(1) § 101 Variance	(2) § 101 Variance	(3) § 101 Variance
Avg_Abs_i		0.248	0.286
		(0.307)	(0.308)
$PatGrp_{it} \times POST$	0.0223***	0.0224***	0.0558***
	(0.00505)	(0.00505)	(0.0212)
$PatGrp_{it} \times POST \times Avg\_Abs_i$			-1.382
			(0.850)
Ind. Claim Count	0.00269*	0.00256	0.00240
	(0.00158)	(0.00159)	(0.00159)
Length of Shortest Ind. Claim	-7.48e-06	-1.15e-05	-1.32e-05
	(2.73e-05)	(2.77e-05)	(2.77e-05)
# of Examiners in USPC	-8.67e-05***	-8.69e-05***	-8.34e-05***
	(2.74e-05)	(2.74e-05)	(2.75e-05)
# of Applications in USPC	4.36e-06	4.42e-06	4.09e-06
	(3.57e-06)	(3.57e-06)	(3.57e-06)
Avg. # of Apps per Examiner	-0.00243***	-0.00244***	-0.00244***
	(0.000348)	(0.000348)	(0.000348)
Avg. of § 101 Variance	0.119***	0.118***	0.115***
	(0.0330)	(0.0330)	(0.0330)
Var. of Examiner's Avg. Abs.	60.62***	56.68***	56.10***
	(15.42)	(16.17)	(16.18)
Constant	0.0335***	0.0303***	0.0304***
	(0.00776)	(0.00869)	(0.00869)
Observations	6,724	6,724	6,724
R-squared	0.470	0.470	0.470
Technology FE	USPC	USPC	USPC
Time FE	Half-year	Half-year	Half-year

*** p<0.01, ** p<0.05, * p<0.1

Table 2

Impact of *Alice* on the Probability of a rejection

	(1)	(2)	(3)
VARIABLES	Model 1	Model 2	Model 3
Abstractness		0.600***	0.316***
		(0.0360)	(0.0378)
$PatGrp_{it} \times POST$	0.0584***	0.0585***	0.0252***
	(0.00198)	(0.00198)	(0.00300)
$PatGrp_{it} \times POST \times Avg\_Abs_i$			1.335***
			(0.0976)
Ind. Claim Count	0.00841***	0.00828***	0.00831***
	(0.000294)	(0.000293)	(0.000293)
Length of Shortest Ind. Claim	-6.03e-05***	-6.57e-05***	-6.45e-05***
	(3.65e-06)	(3.79e-06)	(3.76e-06)
Constant	0.0977***	0.0886***	0.0926***
	(0.00906)	(0.00907)	(0.00907)
Observations	910,612	910,612	910,612
R-squared	0.314	0.315	0.315
Time FE	Month Year	Month Year	Month Year
Technology FE	USPC	USPC	USPC
Examiner FE	Yes	Yes	Yes

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3