Printed in U.S.A.
Vol. 118, No. 5

DO PATENTS DRIVE INVESTMENT IN SOFTWARE?

James Hicks

ABSTRACT—In the wake of a quartet of Supreme Court decisions which disrupted decades of settled law, the doctrine of patentable subject matter is in turmoil. Scholars, commentators, and jurists continue to disagree sharply over which kinds of invention should be patentable. In this debate, no technology has been more controversial than software. Advocates of software patents contend that denying protection would stymie innovation in a vital industry; skeptics argue that patents are a poor fit for software, and that the social costs of patents outweigh any plausible benefits. At the core of this disagreement is a basic problem: the debate is predicated on various claims about how and whether patents incentivize innovation, but like much of patent law, these claims rest on meager empirical foundations. This Article bolsters these foundations by testing one important claim: that patents serve to attract investment in new inventions.

Using a novel quasi-experimental approach and an original dataset, I investigate whether the grant of a patent makes a business-methods software startup more likely to attract early-stage venture capital investment. In contrast to prior scholarship, I find no evidence that patents play a role in channeling investment to these startups, nor that they lead to more successful downstream outcomes such as acquisitions and initial public offerings.

These findings have important implications for both patent policy and scholarship. First, this Article provides new evidence on the perennial controversy over whether business-methods software should be patent-eligible. As Congress continues to contemplate new legislation to clarify the law of patentable subject matter, the results call into question a leading justification for granting patents in this area. Second, in light of previous scholarship which finds a relationship between patents and investment in other areas of technology, I demonstrate the importance of developing industry-specific evidence on the role that patents play in stimulating innovation.

AUTHOR—Academic Fellow, Columbia Law School. For helpful conversations, I thank Ash Ahmed, Abhay Aneja, Shyam Balganesh, Mala Chatterjee, José Argueta Funes, Adam Badawi, Ben Chen, Hal Edgar, Joan Farre-Mensa, Kristelia García, Talia Gillis, Jill Grennan, Sonia Katyal,

Clarisa Long, Justin McCrary, Tejas Narechania, Anthony Niblett, Dave Schwartz, Neel Sukhatme, Eric Talley, Molly Van Houweling, Matt Wansley, Brian Wright, Sam Zyontz, and, in particular, Bobby Bartlett and Julian Nyarko. I also thank participants at workshops at Georgetown, George Mason, and Berkeley; PatCon 11 at Northwestern; the 2022 CLEA meeting at the University of Toronto; the 2022 ALEA meeting at Columbia; and WIPIP 2022 at Saint Louis University. Finally, I'm grateful to Don Bowen and Gerard Hoberg for sharing insights and code from their own patent–VC matching project.

INTRODUCTION 1278
I. THE PATENTABLE-SUBJECT-MATTER CONTROVERSY 1284
II. PATENTS AND INVESTMENT: WHAT DO WE KNOW? 1289
A. Surveys 1289
B. Correlational Studies 1290
C. The Holy Grail: Causality 1292
III. HOW CAN WE ISOLATE THE EFFECT OF PATENTS? 1294
A. The Empirical Design: Alice v. CLS Bank 1294
B. Constructing the Dataset 1297
IV. RESULTS: DO PATENTS AFFECT INVESTMENT OUTCOMES? 1302
A. First Stage: Did Alice *Matter?* 1302
B. Funding 1303
C. Acquisitions and IPOs 1309
V. PROBING ALTERNATIVE EXPLANATIONS 1313
A. What About Bilski*?* 1313
B. Did Investment Decline After Alice*?* 1316
C. Were the Applicants Who Received Their Decision After Alice *Different?* 1319
VI. IMPLICATIONS AND SOME CAVEATS 1320
A. Limitations 1320
B. What Next for Patentable Subject Matter? 1321
CONCLUSION 1325
APPENDIX 1326

INTRODUCTION

Patent law is in a strange state. More than two centuries after the first patent statute was passed, patentable subject matter—the threshold question of what kinds of inventions should be protected at all—is a central controversy of the field. It hasn't always been this way. For years, patent

eligibility was a quiet cul-de-sac of patent doctrine, with broad statutory language and a handful of mostly well-settled judicial exceptions.[1] But in the wake of four Supreme Court opinions in as many years, eligibility law is in tumult.[2]

At the core of the discontent is the Court's new test for patentability. This test prescribes that when a patent claim is directed to an "abstract idea," it must be married with an additional "inventive concept" to render it patent-eligible.[3] In the years since the test's creation, a clear definition of abstract ideas has proved elusive, and no one seems to know what constitutes an inventive concept.[4] The result has been, to put it mildly, disarray. The Patent and Trademark Office (PTO) is launching studies,[5] Congress is holding hearings,[6] and courts are striking down patents.[7] The Federal Circuit is nonplussed: a recent high-profile petition for rehearing en banc resulted in a 6–6 split and five separate opinions.[8] But despite all this, the Supreme Court has rejected multiple entreaties to clarify its jurisprudence.[9] And although Congress has proposed new statutory language to resolve the issue, stakeholders' views differ on how the law should be changed.[10] Put simply:

---

[1] *See infra* Part I.

[2] Bilski v. Kappos, 561 U.S. 593 (2010); Mayo Collaborative Servs. v. Prometheus Lab'ys, Inc., 566 U.S. 66 (2012); Ass'n for Molecular Pathology v. Myriad Genetics, Inc., 569 U.S. 576 (2013); Alice Corp. v. CLS Bank Int'l, 573 U.S. 208 (2014).

[3] *Alice*, 573 U.S. at 217–18; *see infra* Part I.

[4] *See, e.g.*, Talha Syed, *Reconstructing Patent Eligibility*, 70 AM. U. L. REV. 1937, 1940 (2021) ("Everyone now knows there is an *Alice* two-step test, but no one knows quite what it means."); Jeffrey A. Lefstin, *The Three Faces of Prometheus: A Post-*Alice *Jurisprudence of Abstractions*, 16 N.C. J.L. & TECH. 647, 649 (2015) ("[T]here is now less clarity on the basic question of patent eligibility than at almost any other time in American patent law.").

[5] USPTO Patent Eligibility Jurisprudence Study, 86 Fed. Reg. 36,257 (July 9, 2021).

[6] *The State of Patent Eligibility in America: Hearing Before the Subcomm. on Intell. Prop.*, 116th Cong. (2019); Patent Eligibility Restoration Act of 2022, S. 4734, 117th Cong. (2022).

[7] *See* Mark A. Lemley & Samantha Zyontz, *Does* Alice *Target Patent Trolls?*, 18 J. EMPIRICAL LEGAL STUD. 47, 49–50 (2021) (finding that patentable-subject-matter litigation increased dramatically after 2014).

[8] Am. Axle & Mfg., Inc. v. Neapco Holdings LLC, 966 F.3d 1347 (Fed. Cir. 2020) (per curiam).

[9] Despite the pleas of both the appellate court and the Solicitor General, the Supreme Court declined to take up the *American Axle* case and provide a guide out of the chaos. Petitioner's Reply Brief at 8, Am. Axle & Mfg., Inc. v. Neapco Holdings LLC, 142 S. Ct. 2902 (2022) (No. 20-891) (noting the Solicitor General's call for clarity in § 101 doctrine); *Am. Axle*, 142 S. Ct. at 2902 (denying certiorari); *see also* Athena Diagnostics, Inc. v. Mayo Collaborative Servs., LLC, 927 F.3d 1333 (Fed. Cir. 2019) (per curiam) (per curiam denial of rehearing en banc, accompanied by eight separate opinions, many calling for clarity in regards to patentable subject matter).

[10] Victoria T. Carrington & Jorge L. Contreras, *Assessing Responses to the PTO's 2021 Patent Eligibility Study*, PATENTLY-O (Feb. 1, 2022), https://patentlyo.com/patent/2022/02/assessing-responses-eligibility.html [https://perma.cc/D8ZR-9LSW] (finding significant disagreement amongst stakeholders in a 2021 Patent and Trademark Office (PTO) survey regarding recent eligibility jurisprudence); Kevin

in important areas of the economy, we're not sure what is eligible for a patent, and we can't agree what should be.

A central feature of this current malaise is one particular area of invention: business methods. Two of the Supreme Court's recent quartet of eligibility cases involved such inventions.[11] Business-methods patents typically claim a software-based method of applying a well-known business or financial practice.[12] (Amazon's patent on "1-Click" shopping is a canonical example.)[13] As finance and e-commerce have become increasingly sophisticated—and central to the modern economy—interest in these patents has grown.[14] But they are also enormously controversial.

Indeed, the patentability of business methods—and software more generally—has long provoked fierce debate.[15] Influential commentators argue that making software unpatentable would chill innovation in an economically important industry.[16] Others contend that patents, with their long term of protection and protracted examination process, are simply

D. Rodkey & David C. Reese, *Senate Legislation Proposes New Changes to Patent Eligibility and PTAB Proceedings*, FINNEGAN (June 27, 2023), https://www.finnegan.com/en/insights/ip-updates/senate-legislation-proposes-changes-to-patent-eligibility-and-ptab-proceedings.html [https://perma.cc/K8NG-S56Q].

[11] Bilski v. Kappos, 561 U.S. 593, 599 (2010); Alice Corp. v. CLS Bank Int'l, 573 U.S. 208, 212 (2014).

[12] In principle, a business-methods patent need not involve software—and indeed the application at issue in *Bilski* did not. *Bilski*, 561 U.S. at 599. However, the vast majority of contemporary business-methods patents are directed towards business practices applied using a computer, and they have become a significant part of modern software patenting. *See, e.g.*, Bronwyn H. Hall, *Business and Financial Method Patents, Innovation, and Policy*, *in* PERSPECTIVES ON PATENTABLE SUBJECT MATTER 248, 250 (Michael Abramowicz, James E. Daily & F. Scott Kieff eds., 2015) (noting that while there is no precise definition of business methods, many scholars treat them as interchangeable with "internet" and "software" patents). For the purposes of this study, I treat business methods as a subset of software patents.

[13] Method and Sys. for Placing a Purchase Ord. via a Commc'ns Network, U.S. Pat. No. 5,960,411 (filed Sept. 12, 1997) (issued Sept. 28, 1999); *see also* Robert P. Merges, *As Many as Six Impossible Patents Before Breakfast: Property Rights for Business Concepts and Patent System Reform*, 14 BERKELEY TECH. L.J. 577 (1999) (discussing more examples).

[14] For background, see generally John F. Duffy, *Why Business Method Patents?*, 63 STAN. L. REV. 1247 (2011).

[15] *See infra* Part I.

[16] *See, e.g.*, Maureen K. Ohlhausen, *Patent Rights in a Climate of Intellectual Property Rights Skepticism*, 30 HARV. J.L. & TECH. 103, 106 (2016) ("Nevertheless, there is ample evidence that patents serve a materially valuable role in promoting innovation in at least some settings."); Gene Quinn, *The Road Forward for Software Patents Post*-Alice, IPWATCHDOG (Feb. 25, 2015), http://www.ipwatchdog.com/2015/02/25/the-road-forward-for-software-patents-post-alice [https://perma.cc/825X-PST4] (discussing a practitioner's view that "if your primary interest is in a strong patent system that protects and rewards innovation, [*Alice* is] arguably not a good decision"); Daniel F. Spulber, *Should Business Method Inventions Be Patentable?*, 3 J. LEGAL ANALYSIS 265, 268 (2011) ("First, patents for business method inventions are important for entrepreneurship and for the commercialization of many scientific and technological inventions.").

unnecessary to promote innovation in an industry characterized by low capital costs; rapid, iterative development; and short product lives.[17] Justice John Paul Stevens channeled the latter sentiment in his concurrence in *Bilski*, in which he argued that business methods in particular should be categorically ineligible for a patent: "I find it hard to believe that many of our entrepreneurs forwent business innovation because they could not claim a patent on their new methods."[18] Justice Stevens's claim is a fundamental one which goes to the very core of innovation policy. But whether he was correct remains an open question.[19]

The problem at the heart of this indeterminacy is a lack of evidence. We simply do not know enough about how the patent system actually works.[20] The dominant justification for intellectual property is its "incentive" function, by which a temporary grant of legal exclusivity is intended to encourage the creation of valuable new expressive and technological innovation.[21] But pinning down the incentive effect of any particular intellectual property regime has long been a difficult—verging on intractable—empirical task.[22]

---

[17] *See, e.g.*, Brief Amici Curiae of Professors Peter S. Menell & Michael J. Meurer in Support of Respondent at 31, Bilski v. Kappos, 561 U.S. 593 (2010) (No. 08-964) ("The social value of patents tends to be lower in fields, including business method innovations, characterized by cumulative innovation and fuzzy patent boundaries."); Joshua D. Sarnoff, *Patent-Eligible Inventions After* Bilski*: History and Theory*, 63 HASTINGS L.J. 53, 98–99 (2011) (discussing ways in which non-patent incentives can promote innovation). More generally, software patents are also accused of being poorly drafted, low-quality, and primarily a tool for "patent trolls"—though there is dispute on these assertions, too. *See, e.g.*, John R. Allison, Mark A. Lemley & Joshua Walker, *Extreme Value or Trolls on Top? The Characteristics of the Most-Litigated Patents*, 158 U. PA. L. REV. 1, 12–20 (2009) ("The most-litigated patents are overwhelmingly likely to be software patents."); James Bessen & Michael J. Meurer, *The Direct Cost from NPE Disputes*, 99 CORNELL L. REV. 387, 406 (2014) (noting how it often costs more to litigate a patent dispute than the patent is worth). *But see* David L. Schwartz & Jay P. Kesan, *Analyzing the Role of Non-Practicing Entities in the Patent System*, 99 CORNELL L. REV. 425, 434 (2014) (disputing Bessen and Meurer's calculations).

[18] *Bilski*, 561 U.S. at 651 (Stevens, J., concurring).

[19] *See, e.g.*, Maya M. Durvasula, Lisa Larrimore Ouellette & Heidi L. Williams, Comments on Patent Eligibility Jurisprudence Study, 86 Fed. Reg. 36,257 (Sept. 7, 2021) ("Many have argued—largely based on anecdotes or descriptive data—that recent changes in patent eligibility caselaw have either increased or decreased innovation . . . we argue that neither view is supported by the available empirical evidence.").

[20] *See, e.g.*, John M. Golden, Robert P. Merges & Pamela Samuelson, *The Path of IP Studies: Growth, Diversification, and Hope*, 92 TEX. L. REV. 1757, 1768 (2014) (heralding a rise in empirical studies of intellectual property and calling for more).

[21] U.S. CONST. art. 1, § 8, cl. 8 ("[Congress shall have the power . . .] to promote the progress of science and useful arts, by securing for limited times to authors and inventors the exclusive right to their respective writings and discoveries.").

[22] *See, e.g.*, ROBERT P. MERGES, JUSTIFYING INTELLECTUAL PROPERTY 2–4 (2011) (arguing that current data about the optimal intellectual property system are "maddeningly inconclusive"); Hall, *supra* note 12, at 271 ("[B]road evidence that the patent system encourages innovation always and everywhere is hard to come by.").

In this Article, I introduce new evidence on whether patents affect innovation in this controversial area of software. I focus on one particularly important mechanism: whether patents play a role in directing investment to new companies.[23] For innovative young startups, attracting early-stage venture capital funding is a crucial step in their development,[24] and it's widely argued that patents are an important tool for startup companies seeking such investment.[25] The empirical reality, however, is less clear.

There are certainly plausible theoretical reasons to think that patents might play a role in directing investment, including as a signal of quality and as a promise of market exclusivity.[26] But the evidence is murky. Despite a growing number of empirical studies in this area, it has proved very difficult to isolate the distinctive impact of intellectual property on financing. Surveys of both venture capitalists and entrepreneurs suggest that patents do help companies attract investment, with variation across industries.[27] This finding is reinforced by quantitative studies, which generally find a correlation between patenting activity and the receipt of venture capital investment.[28] But these latter studies are bedeviled by a common problem: there are typically many factors which drive both the likelihood of patenting *and* the chance of investment (for example, the underlying quality of the invention or team). Without a way to hold these other hard-to-observe factors constant,

---

[23] *See* Mark A. Lemley, *Reconceiving Patents in the Age of Venture Capital*, 4 J. SMALL & EMERGING BUS. L. 137, 144 (2000) (arguing that scholars should turn their attention from litigation to the broader uses of patents in innovative companies and markets).

[24] *See* Josh Lerner & Ramana Nanda, *Venture Capital's Role in Financing Innovation: What We Know and How Much We Still Need to Learn*, 34 J. ECON. PERSPS. 237, 248 (2020); Peter Lee, *Enhancing Venture Capital*, 24 YALE J.L. & TECH. 611, 636 (2022) (noting the centrality of venture capital funding to the contemporary American innovation landscape).

[25] *See, e.g.*, JONATHAN BARNETT, INNOVATORS, FIRMS, AND MARKETS: THE ORGANIZATIONAL LOGIC OF INTELLECTUAL PROPERTY 38 (2020) ("[P]atents can sometimes provide a signal (albeit with considerable noise) of innovation quality, which in turn can result in improved valuations of a firm's equity or increase the likelihood that a start-up secures outside financing."); Ohlhausen, *supra* note 16, at 140–41 (noting a pervasive view amongst founders that venture capitalists consider patents important to investment decisions); John R. Harris, *The Patent System Is Under Assault—Startups, Should You Care? Ten Things About Patents That Startups Need to Consider*, 44 AIPLA Q.J. 27, 33 (2016) (arguing that having an enforceable patent helps a startup to protect themselves against demands from patent assertion entities); Jay P. Kesan & Runhua Wang, *Eligible Subject Matter at the Patent Office: An Empirical Study of the Influence of* Alice *on Patent Examiners and Patent Applicants*, 105 MINN. L. REV. 527, 545 (2020) (surveying views amongst scholars as to whether the *Alice* decision deters innovation); Stuart J.H. Graham, Robert P. Merges, Pamela Samuelson & Ted M. Sichelman, *High Technology Entrepreneurs and the Patent System: Results of the 2008 Berkeley Patent Survey*, 24 BERKELEY TECH. L.J. 1255, 1255 (2009) (exploring the competitive advantages early-stage companies might seek when they patent, including protection against copying, securing financing, and enhancing their reputation).

[26] *See infra* Part II.

[27] *See infra* Section II.A.

[28] *See infra* Section II.B.

it is impossible to assess whether patents play a distinctive role in mediating investors' choices or whether money simply flows to the best ideas regardless of intellectual property. The question remains: do patents matter for investment?

In this Article, I introduce a new quasi-experimental approach to this problem. In *Alice Corp. v. CLS Bank International*—the final case in the eligibility quartet, and a landmark patent decision—the Supreme Court struck down a patent on a method to implement the concept of escrow using computer software.[29] Despite its ostensible vagueness, the test laid out in *Alice* was widely perceived to limit patent eligibility for a large class of software inventions.[30] And indeed, soon after *Alice* was handed down, the PTO dramatically tightened its examination standards for affected applications, particularly in the technology area at issue in the case: business-methods software.[31] The result was patent application rejections—lots of them. As I show in this Article, grant rates for business-methods patent applications fell by more than 75% in a matter of weeks.[32] This created two groups of software startups, which were divided by an arbitrary line with serious implications. One set of applicants—those whose applications were examined after the decision came down—was subject to far stricter patentability standards.

Using this sudden change in patentability standards to divide applicants, I track the patent allowance rates and financial outcomes for both groups of companies. I confirm that allowance rates did indeed drop precipitously for business-methods software applications in the wake of *Alice*, beginning almost immediately after the decision. Despite a dramatic difference in patenting outcomes between the two groups, however, I find no difference in either investment or subsequent rates of acquisition. In particular, I find no evidence that a startup company's first patent grant has an effect on subsequent venture capital investment and no evidence that the grant increases the likelihood that a company will later be acquired or launch an initial public offering.[33] I also check for—and rule out—various alternative explanations for these findings, including the possible impact of an earlier

---

[29] 573 U.S. 208, 212 (2014).

[30] Rob Merges, *Symposium: Go Ask* Alice — *What Can You Patent After* Alice v. CLS Bank*?*, SCOTUSBLOG (June 20, 2014), https://www.scotusblog.com/2014/06/symposium-go-ask-alice-what-can-you-patent-after-alice-v-cls-bank/ [https://perma.cc/TSB8-KZUU].

[31] *See* Memorandum from Andrew H. Hirshfeld, U.S. Patent & Trademark Off., Preliminary Examination Instructions in View of the Supreme Court Decision in *Alice Corporation Pty. Ltd. v. CLS Bank International, et al.* 1–2 (June 25, 2014), https://www.uspto.gov/sites/default/files/patents/announce/alice_pec_25jun2014.pdf [https://perma.cc/7EPR-VWRP].

[32] *See infra* Section IV.A.

[33] *See infra* Part IV.

Supreme Court case and potential changes in the overall investment climate.[34]

This Article makes two key contributions. First, I demonstrate empirically that patent grants have no apparent effect on the receipt of investment or on subsequent acquisitions and initial public offerings (IPOs) for business-methods software startups. This result is quite surprising in light of previous scholarship, much of which argues that an important feature of patents is the signal that they provide to investors.[35] Second, I provide rare empirical evidence for the ongoing policy debate over patent-eligibility reform. As Congress continues to contemplate changes to the law, it is vital to develop more nuanced, industry-specific evidence about the varied roles that patents play in incentivizing new innovation. The findings in this Article lend support to those who have called for Congress to carve business-methods software out of patentable subject matter.

The remainder of this Article proceeds as follows. I begin in Part I by describing the history of the debate over patentable subject matter, the role played by software, and the seismic effect of the Court's recent eligibility quartet. In Part II, I summarize the previous research on patenting and investment. I draw on literatures in law, management, and economics to explain what we know—and what we don't—about the role that intellectual property (IP) plays in connecting inventions with finance. In Part III, I explain how the Supreme Court's opinion in *Alice* provides a unique opportunity to test the effect of patents on venture capital investment choices, and I introduce a novel dataset that links startup patent applicants with financing and other business outcomes. In Parts IV and V, I present the main results and consider several possible alternative explanations. Finally, in Part VI, I turn to the implications of these findings for ongoing debates in patent policy and scholarship.

I. THE PATENTABLE-SUBJECT-MATTER CONTROVERSY

The threshold requirement for patentability is that the claimed invention constitutes patentable subject matter.[36] This has not generally presented a significant hurdle for would-be applicants. The patent statute defines patent-eligible subject matter broadly, providing that a patent may issue to "[w]hoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement

[34] *See infra* Part V.

[35] *See infra* Part II.

[36] 35 U.S.C. § 101.

thereof."[37] And courts have historically adopted a broad construction, abiding by the maxim that Congress intended patents to be available for "anything under the sun that is made by man."[38]

Notwithstanding this broad construction, courts have long held three types of discovery to be categorically ineligible for protection: abstract ideas, laws of nature, and naturally occurring phenomena.[39] In two important cases, the Supreme Court held that algorithms were not patent eligible, evincing in both cases a concern about tying up fundamental, abstract knowledge.[40] But that view softened several years later in *Diamond v. Diehr*, when the Court allowed a patent on a software-implemented process for curing rubber—thus giving license to a new generation of software patents.[41]

This was, essentially, the doctrinal landscape until 1998, when the Federal Circuit opened the door to business-methods patenting.[42] In *State Street Bank & Trust Co. v. Signature Financial Group, Inc.*, the Court upheld a patent on a piece of software that implemented a method for tax-efficient administration of a portfolio of mutual funds on the grounds that the invention produced "a useful, concrete and tangible result."[43] Business-methods patents were (and remain) very controversial, and many academics and commentators raised concerns about the *State Street* decision at the time.[44] But the Supreme Court declined to intervene. Over time, business methods, and software more generally, became an established—if still contentious—part of the patent firmament.[45]

---

[37] *Id.*

[38] Diamond v. Chakrabarty, 447 U.S. 303, 309 (1980) (quoting S. REP. NO. 82-1979, at 2399 (1952)).

[39] *See, e.g.*, Mayo Collaborative Servs. v. Prometheus Lab'ys, Inc., 566 U.S. 66, 77 (2012) ("The Court has long held that this provision contains an important implicit exception. '[L]aws of nature, natural phenomena, and abstract ideas' are not patentable." (quoting Diamond v. Diehr, 450 U.S. 175, 185 (1981))).

[40] Gottschalk v. Benson, 409 U.S. 63 (1972); Parker v. Flook, 437 U.S. 584 (1978) (finding that a mathematical algorithm is only patentable to the extent that it has some inventive concept in its application).

[41] 450 U.S. at 175.

[42] State St. Bank & Tr. Co. v. Signature Fin. Grp., Inc., 149 F.3d 1368 (Fed. Cir. 1998), *cert. denied*, 525 U.S. 1093 (1999). The "useful, concrete, and tangible" standard was later superseded by the "machine-or-transformation" test in the Federal Circuit's *Bilski* opinion, and then again by the Supreme Court's (rather ambiguous) "inventive concept" test in *Alice*. *See In re* Bilski, 545 F.3d 943, 959–60 (Fed. Cir. 2008); Alice Corp. v. CLS Bank Int'l, 573 U.S. 208, 221 (2014).

[43] 149 F.3d at 1373.

[44] *See, e.g.*, Rochelle Cooper Dreyfuss, *Are Business Method Patents Bad for Business?*, 16 SANTA CLARA COMPUT. & HIGH TECH. L.J. 263 (2000); Michael J. Meurer, *Business Method Patents and Patent Floods*, 8 WASH. U. J.L. POL'Y 309 (2002); Duffy, *supra* note 14.

[45] *See, e.g.*, Julie E. Cohen & Mark A. Lemley, *Patent Scope and Innovation in the Software Industry*, 89 CALIF. L. REV. 1, 4 (2001) (arguing that scholars should spend less time debating whether software should be patentable and more on cabining the scope of the right granted by the patent).

In the last decade, however, this sleepy doctrinal backwater has taken center stage.[46] In a series of important cases beginning in 2010, the Supreme Court began to cabin the scope of patentable subject matter.[47] The first case in the Court's recent foray into patent eligibility was also the first time that it addressed modern business-methods patenting. In *Bilski v. Kappos*, the petitioner claimed an application of the financial practice of hedging.[48] The Court held that the claims were directed to an "unpatentable abstract idea," but its opinion provided little in the way of guidance for future litigants or the PTO, and the majority rejected any categorical exception for business-methods patents.[49] Then, in *Mayo Collaborative Services v. Prometheus Laboratories, Inc.*, the Court found that a method of administering a drug (and calibrating the right dosage) to treat autoimmune diseases was merely a routine application of a natural law.[50] In *Association for Molecular Pathology v. Myriad Genetics, Inc.*, it ruled that a particular sequence of naturally occurring human DNA—in this case, a sequence of genes that happened to strongly predict breast cancer—was a natural phenomenon, and its discovery was therefore outside the scope of patentable subject matter.[51]

In the final case in this quartet, *Alice Corp. v. CLS Bank International*, the Court turned back to business methods.[52] It found that a patent which claimed a software implementation of intermediated financial settlement—in other words, escrow—was again merely an implementation of an abstract idea, and not patentable subject matter.[53] But more important than the particular patent at issue was the Court's formalization of a two-step test for assessing subject matter eligibility. Under *Alice* (and *Mayo*), courts and the PTO are instructed to ask two questions. First, is the patent claim directed to an abstract idea, a law of nature, or a naturally occurring phenomenon? If so,

---

[46] *See, e.g.*, Mark A. Lemley, Michael Risch, Ted Sichelman & R. Polk Wagner, *Life After* Bilski, 63 STAN. L. REV. 1315, 1317 (2011) ("[T]he Supreme Court got out of the business of patentable subject matter for nearly thirty years.").

[47] Jeffrey A. Lefstin, Peter S. Menell & David O. Taylor, *Final Report of the Berkeley Center for Law & Technology Section 101 Workshop: Addressing Patent Eligibility Challenges*, 33 BERKELEY TECH. L.J. 551, 559 (2018).

[48] 561 U.S. 593, 611–12 (2010).

[49] *Id.* Four Justices would have held business methods to be categorically ineligible for a patent. *Id.* at 614 (Stevens, J., concurring) ("The wiser course would have been to hold that petitioners' method is not a 'process' because it describes only a general method of engaging in business transactions—and business methods are not patentable.").

[50] 566 U.S. 66, 77 (2012).

[51] 569 U.S. 576, 579–80 (2013).

[52] 573 U.S. 208, 212 (2014).

[53] *Id.* ("[M]erely requiring generic computer implementation fails to transform that abstract idea into a patent-eligible invention.").

second, does it nevertheless apply the idea, law, or phenomenon in a sufficiently inventive way as to render the claim patent-eligible?[54]

To say these decisions have been controversial would be to understate the firestorm. *Alice*, in particular, has come under fire from all sides. Many commentators rue the vague and incoherent nature of the new rule.[55] The bulk of the criticism is directed at the Court's new two-part test for the patentability of so-called applied abstract ideas, which Professor Rob Merges calls "brief, yet somehow baroquely obscure."[56] (Amongst the questions on which we have little guidance: What *is* an abstract idea? What transformation rises to the level of an "inventive concept"?) Others argue that the decisions will have a chilling effect on innovation in important areas of economic activity, including medical diagnostics, scientific discovery, and software.[57]

In the wake of *Alice*'s controversial impact, Congress has shown a renewed interest in clarifying the law of patentable subject matter.[58] Proposals range from a categorical bar on business-methods patenting to a full legislative overruling of *Alice*, and various points in between. This is a politically fraught issue. But it's also not obvious what the right approach should be.[59] The relationship between intellectual property and innovation is

[54] *Id.* at 217–21.

[55] *See, e.g.*, Michael Risch, *Nothing Is Patentable*, 67 FLA. L. REV. F. 45, 53 (2015) ("[W]e apply a standard so vague that it would invalidate patents throughout history and, by extension, many otherwise meritorious patents today."); David O. Taylor, *Amending Patent Eligibility*, 50 U.C. DAVIS. L. REV. 2149, 2151 (2017) ("[T]he Court's two-part test has created a significant risk of reduced incentive to invent."); Lefstin, *supra* note 4 ("[G]iven the Court's reluctance to provide specific guidance, there is little agreement on how the analysis of patent eligibility should be structured.").

[56] Merges, *supra* note 30; *see also* Risch, *supra* note 55, at 45 (characterizing the test as "a foggy standard cloaked as a rule").

[57] *See* David O. Taylor, *Confusing Patent Eligibility*, 84 TENN. L. REV. 157, 240 (2016); Kevin Madigan & Adam Mossoff, *Turning Gold into Lead: How Patent Eligibility Doctrine Is Undermining U.S. Leadership in Innovation*, 24 GEO. MASON L. REV. 939, 941 (2017). On the other hand, forthcoming work indicates that this chilling effect has not come to pass in the area of medical diagnostics. Colleen V. Chien, Jenna Clark & Arti K. Rai, *Molecular Diagnostic Patenting After* Mayo v. Prometheus*: An Empirical Analysis* (Duke L. Sch. Pub. L. & Legal Theory Series, Working Paper No. 2023-69, 2023), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4648623 [https://perma.cc/AP74-8RL6].

[58] *See supra* note 6; Patent Eligibility Restoration Act of 2023, S. 2140, 118th Cong. (2023).

[59] Senator Thom Tillis, who cosponsored recent legislative efforts, sums up the challenge in an interview: "If someone doesn't like the compromise that I've come up with so far, tell me. But don't just tell me that you hate it—give me ideas on how to write a better bill that will, in your opinion, result in a better outcome that promotes innovation and restores certainty." Gene Quinn & Eileen McDermott, *Tillis Addresses Criticism of His Eligibility Reform Bill*, IPWATCHDOG (Aug. 31, 2022), https://ipwatchdog.com/2022/08/31/tillis-addresses-criticism-eligibility-reform-bill-warns-wd-tx-not-backtrack-standing-order [https://perma.cc/DE6D-Q544].

complex, in both theory and practice.[60] The optimal intellectual property regime balances incentives to innovate against tying up knowledge (both for consumers and future innovators), and facilitating market entry against fortifying the position of powerful incumbent companies.[61] But in practice, achieving the right balance in any particular industry is challenging.

This is an old problem. Software—of which business methods is one flavor—has a particularly long and fraught history in intellectual property law.[62] As a theoretical matter, patents are an odd fit with the economics of software. Patents are expensive and time-consuming to obtain. This may be a sensible trade-off for inventions with high fixed costs and lasting value—pharmaceuticals are famously the leading example—but software is characterized by rapid evolution and iterative improvement. Scope also presents a challenge. Business-methods software has notoriously fuzzy boundaries, resulting in notice failures and opportunistic litigation. There is also some evidence that "patent thickets" play a role in delaying entry and financing in business-methods and related software.[63] Put simply, the theoretical case for patents in this area isn't very strong.[64]

Still, many commentators argue that an important role remains for software patents: attracting financing to new inventions.[65] These commentators contend that limiting patentability will inhibit research-and-development investment and, since that investment is the lifeblood for most early-stage companies, will depress innovation. Do patents serve this role for

---

[60] *See* Hall, *supra* note 12, at 248–50; Peter S. Menell & Suzanne Scotchmer, *Economic Models of Innovation: Stand-Alone and Cumulative Creativity*, *in* 1 RESEARCH HANDBOOK ON THE ECONOMICS OF INTELLECTUAL PROPERTY LAW 119 (Ben Depoorter & Peter S. Menell eds., 2019).

[61] Hall, *supra* note 12, at 249.

[62] The literature on this topic is truly voluminous, but for some particularly important early works, see Peter S. Menell, *Tailoring Legal Protection for Computer Software*, 39 STAN. L. REV. 1329, 1353–72 (1987), which argues that the economic characteristics of software make it ill-suited for traditional modes of intellectual property protection; Pamela Samuelson, Randall Davis, Mitchell D. Kapor & J.H. Reichman, *A Manifesto Concerning the Legal Protection of Computer Programs*, 94 COLUM. L. REV. 2308, 2310–12 (1994), which argues for a sui generis form of protection for computer software; and Cohen & Lemley, *supra* note 45, at 56, which argues that "applying existing patent doctrine to software patents threatens to create exclusionary rights that are extraordinarily broad even by patent standards."

[63] Iain M. Cockburn & Megan J. MacGarvie, *Patents, Thickets and the Financing of Early-Stage Firms: Evidence from the Software Industry*, 18 J. ECON. & MGMT. STRATEGY 729, 736 (2009). A patent thicket is "a dense web of overlapping intellectual property rights that a company must hack its way through in order to actually commercialize new technology." Carl Shapiro, *Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting*, *in* 1 INNOVATION POLICY AND THE ECONOMY 119, 120 (Adam Jaffe, Josh Lerner & Scott Stern eds., 2000).

[64] As Professor Bronwyn Hall argues: "We know that patents are not considered essential for capturing the returns to innovation in many industries, and there seems no reason to think that financial patents are any different." Hall, *supra* note 12, at 264.

[65] *See supra* note 25.

business-methods software? If so, this would be a strong argument for revisiting the *Alice* status quo. This is the question that this Article takes on.

II. PATENTS AND INVESTMENT: WHAT DO WE KNOW?

At first blush, whether patents facilitate access to financing seems like a straightforward question to answer. There are many compelling theoretical reasons to think that patents have a role to play.[66] First, a patent could send a signal about the product's stage of development or the competence of the team.[67] In this way, as Professor Clarisa Long argues, intellectual property might help to mitigate the information asymmetry between startup teams and potential investors.[68] Second, a patent might indicate that the company has already carved out some area of exclusivity within its market niche, therefore increasing the expected profitability of the invention. And third, an issued patent may promise some salvage value. If the company fails, as startups commonly do, the investors would still own a patent (or a portfolio of patents) which could be sold or licensed. But while each of these theories is plausible, the empirical reality is much less clear. Scholars have taken various approaches to studying the role of patents in early-stage investment, including surveys, correlational studies, and some limited causal analysis. In this Part, I survey this empirical literature—which encompasses research from law, economics, and management—and provide context for the novel approach that I develop in Part III.

A. Surveys

The first approach to these questions, which is predominantly the province of legal scholarship, involves extensive surveys of entrepreneurs and venture capitalists. For example, a team from Berkeley conducted a major national survey of over 1,300 technology startups in order to understand the myriad reasons why young companies do (and don't) seek patents.[69] With respect to investment, the authors see significant variation between industries but find that, on the whole, most startups believed patents to be at least somewhat important to their professional investors. For example, 97% of biotechnology and 59% of software firms reported that they thought patents were important for attracting venture capital.[70] In a

[66] *See generally* Bronwyn H. Hall & Dietmar Harhoff, *Recent Research on the Economics of Patents*, 4 ANN. REV. ECON. 541 (2012) (finding suggestive evidence that patents can help to obtain financing).

[67] *See* Mark A. Lemley, *Rational Ignorance at the Patent Office*, 95 NW. U. L. REV. 1495, 1505–06 (2001).

[68] *See* Clarisa Long, *Patent Signals*, 69 U. CHI. L. REV. 625, 672 (2002).

[69] Graham et al., *supra* note 25, at 1260, 1272.

[70] *Id.* at 1306–08.

companion paper, two of the authors report that software startups found patents to be "moderately important" for improving their chances of attracting investment.[71] Similarly, in an earlier survey of firms in the semiconductor-design industry, respondents reported using patents as a means to secure venture capital investment in the startup phase.[72]

Professor Ronald Mann focuses specifically on the software industry, interviewing investors, developers, and lawyers.[73] Mann finds his respondents to be somewhat apathetic about the importance of patents, though he reports that the value of IP increases for companies that survive to later-stage financing rounds.[74] A more recent study by Professor David Taylor asks investors explicitly about their awareness of changes to patent eligibility doctrine, on the theory that knowledgeable venture capitalists (VCs) would be expected to change their investment priorities if patenting became more difficult for particular technologies.[75] One of the interesting takeaways from Taylor's work is that he finds a surprisingly low rate of awareness of actual changes to patent law.[76] But for the purposes of the project at hand, several specific findings stand out from the article: 72% of investors in software report that patent eligibility is an important consideration when making investment choices, and 35% say that the elimination of patents would cause them to reduce their investment in the affected industry.[77]

B. Correlational Studies

The second set of scholarship is comprised of studies that use various statistical methods to tease out the relationship between IP and funding based on observed, quantitative data. These papers focus on a range of different

---

[71] Ted Sichelman & Stuart J.H. Graham, *Patenting by Entrepreneurs: An Empirical Study*, 17 MICH. TELECOMM. TECH. L. REV. 111, 158 (2010).

[72] Bronwyn H. Hall & Rosemarie Ham Ziedonis, *The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979–1995*, 32 RAND J. ECON. 101, 104 (2001).

[73] *See* Ronald J. Mann, *Do Patents Facilitate Financing in the Software Industry?*, 83 TEX. L. REV. 961 (2005).

[74] *Id.*

[75] David O. Taylor, *Patent Eligibility and Investment*, 41 CARDOZO L. REV. 2019, 2060, 2095–96 (2020).

[76] *See id.* at 2051 (reporting that 62% of responding investors were not familiar with any recent patent eligibility cases decided by the Supreme Court).

[77] *Id.* at 2059, 2066. Another line of survey work considers the role of patents in creating a portfolio to defend against licensing demands. *See, e.g.*, Robin Feldman, *Patent Demands & Startup Companies: The View from the Venture Capital Community*, 16 YALE J.L. & TECH. 236, 281 (2014) (finding, based on a survey of more than 200 venture capitalists (VCs), that they believe licensing demands have a negative impact on entrepreneurs and the venture-backed community); Colleen Chien, *Startups and Patent Trolls*, 17 STAN. TECH. L. REV. 461, 483 (2014) (discussing how growing companies "can buy patents on the open market and quickly bulk up their portfolios 'on demand'").

industries and use somewhat varied empirical approaches but share a common jumping-off point: they generally start with a sample of venture-backed companies and try to deduce the characteristics which relate to more rounds of funding or larger investments. Much of this work appears in the economics and management literatures. For example, Professors David Hsu and Rosemarie Ziedonis find that patent applications in the semiconductor industry are positively related to the company's valuation at successive rounds of financing, particularly in early-stage financing.[78] Professors Iain Cockburn and Megan MacGarvie explore software and business methods directly, leveraging the expansion of patentability in those areas in the 1990s.[79] They find that the existence of patent applications leads to quicker financing rounds.[80] In the legal literature, Professor Michael Risch reports that the rate of venture capital financing is *ten times* higher for patent-holders than non-patentees in the Kauffman Firm Survey (a survey that follows the progression of early-stage firms over time).[81] Finally, Professors Ronald Mann and Thomas Sager test some of the survey findings from Mann's earlier work on software and find a strong association between patents and both the number of rounds of financing and the total amount invested.[82] However, as the authors acknowledge, it is difficult to infer a causal relationship from their results.[83]

In some respects, revealed preferences can be more compelling than stated ones—in other words, watch what VCs do, not just what they say. But this second set of papers suffers from a difficulty in credibly isolating the distinctive effects of intellectual property. This is true for two reasons. First, there are many unobservable factors related to quality of inventions and teams that may drive both patenting *and* investment. Indeed, if both IP and

---

[78] David H. Hsu & Rosemarie H. Ziedonis, *Resources as Dual Sources of Advantage: Implications for Valuing Entrepreneurial-Firm Patents*, 34 STRATEGIC MGMT. J. 761, 772–73 (2013).

[79] Cockburn & MacGarvie, *supra* note 63, at 729.

[80] *Id.* at 761. As Professor Bronwyn Hall observes in a recent review essay, several papers in the literature suggest that patent *applications* may be more important than patent grants for driving investments. *See* Bronwyn H. Hall, *Is There a Role for Patents in the Financing of New Innovative Firms?*, 28 INDUS. & CORP. CHANGE 657, 663–64 (2019). She speculates that the relatively high grant rates at the PTO imply that any application is a reasonable predictor of a future patent grant, and therefore the grant may provide little additional information to investors. *Id.* at 663; *see also* Carolin Haeussler, Dietmar Harhoff & Elisabeth Mueller, *How Patenting Informs VC Investors – The Case of Biotechnology*, 43 RSCH. POL'Y 1286, 1286 (2014) (arguing that applications can provide a signal of quality to potential investors). However, regardless of the general effect of applications, it seems unlikely to hold in the field of business-methods software, since even before *Alice* this was not an area with a high patent-allowance rate.

[81] *See* Michael Risch, *Patent Troll Myths*, 42 SETON HALL L. REV. 457, 492 (2012).

[82] Ronald J. Mann & Thomas W. Sager, *Patents, Venture Capital, and Software Start-Ups*, 36 RSCH. POL'Y 193, 194 (2007).

[83] *Id.* at 199–200.

venture capital work as advertised, then we should expect to see that the best inventions attract both financing and patent protection. And second, they struggle to answer the question about when patents enter the picture. Assuming that IP and finance are correlated, does money follow the patenting, or does patenting follow the money? The former explanation would suggest that there is an ex ante role for patents in signaling quality or providing collateral for investments. By contrast, the latter would suggest that funders see IP primarily as a channel to protect existing investments (or even just a standard part of their due diligence process). Of course, either would be interesting from the perspective of patent theory and policy, but the two explanations reflect different institutional roles for patents.[84]

C. The Holy Grail: Causality

Of the quantitative studies exploring the relationship between patents and finance, only one can claim to truly isolate the *causal* effect of a patent grant on investment. Professors Joan Farre-Mensa, Deepak Hegde, and Alexander Ljungqvist leverage the fact that PTO examiners, who are (or used to be) randomly assigned to review applications, vary dramatically in their propensity to grant patents.[85] This introduces some external variation—in other words, variation that is unrelated to the applicant or the invention—into the application outcomes. (Intuitively, if you draw a lenient examiner, you're more likely to receive a patent than an applicant with an otherwise identical invention who is assigned to a much tougher one.)[86] The authors

[84] *See also* Hall, *supra* note 80, at 661–64 (reviewing this literature in detail).

[85] Joan Farre-Mensa, Deepak Hegde & Alexander Ljungqvist, *What Is a Patent Worth? Evidence from the U.S. Patent "Lottery,"* 75 J. FIN. 639, 650–51 (2020).

[86] The so-called examiner (or judge) leniency instrument has become increasingly popular amongst empirical scholars of the patent system because it introduces a credible source of randomness to a system with very little of it. *See, e.g.*, Alberto Galasso & Mark Schankerman, *Patents and Cumulative Innovation: Causal Evidence from the Courts*, 130 Q.J. ECON. 317 (2015) (finding a substantial increase in citations to patents that are invalidated at the Federal Circuit); Bhaven Sampat & Heidi L. Williams, *How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome*, 109 AM. ECON. REV. 203, 232 (2019) (finding no substantively important effect of gene patents on follow-on innovation); MICHAEL D. FRAKES & MELISSA F. WASSERMAN, SELECTION AND DECISION IN JUDICIAL PROCESS AROUND THE WORLD 20 (Yun-chien Chang ed., 2019) (finding "a large sensitivity in the likelihood of patent litigation to the inherent leniency of the associated examiner"); Abhay Aneja, Oren Reshef & Gauri Subramani, Try, Try, Try Again? Differential Responses to Rejection & the Gender Innovation Gap 2 (July 2021) (unpublished manuscript), https://sites.bu.edu/tpri/files/2021/07/GSubramani_Try_try_try721.pdf [https://perma.cc/HT6P-NRDA] (estimating the causal effect of rejection on patent continuations). In the venture capital context, Professor Patrick Gaulé also uses the examiner leniency instrument to assess the effect of a patent grant on VCs' ability to successfully exit their investments through acquisitions and initial public offerings (IPOs). Patrick Gaulé, *Patents and the Success of Venture-Capital Backed Startups: Using Examiner Assignment to Estimate Causal Effects*, 66 J. INDUS. ECON. 350, 351–52 (2018) (using a similar approach to find that patent approval had a substantively large

track subsequent funding outcomes and find that a patent grant causes a 50% increase in the likelihood of attracting early-stage investment over the five years following the first PTO action, as well as an increase in the likelihood of a subsequent IPO.[87] More striking still, the authors explore the variation across industries and find that the impact of a patent grant is stronger in information technology (IT) than other industries.[88] (Though for present purposes, it's important to note that the authors do not count business-methods software as part of their "IT" category, instead combining it with "other industries.") For IT patents, the first patent grant *doubles* the likelihood of receiving investment whereas for biochemistry, the effect is not significantly different from zero.[89] This last result is quite surprising in light of both the findings in the survey literature and the conventional wisdom amongst many commentators that biotechnology is a particularly important site of patent protection.[90]

In summary, the overall takeaway is mixed. Surveys suggest a role for patents in attracting venture capital investment. The strength of this role varies across industries, but even investors and entrepreneurs in software report that patents are important. Most quantitative evidence finds a clear correlation between patenting and funding—and sometimes quite a large one—but suffers from important methodological shortcomings. The very best evidence we have indicates that the approval of a patent does drive subsequent investment, though the authors find notable differences between industries (and don't consider software explicitly). In short, we need more evidence, particularly with respect to individual technologies.

---

impact on successful outcomes—IPOs and acquisitions—for companies in the life sciences sectors, but not in information technology).

[87] One aspect of the Farre-Mensa, Hegde, and Ljungqvist paper's empirical design is somewhat unusual. Although the authors compare startup companies based on whether their first patent was ultimately granted, they measure investment from the date of the *first office action* received by the applicant. Farre-Mensa et al., *supra* note 85, at 644–45. In most cases, the first action is a "nonfinal rejection," and it is usually received several years before the ultimate disposition of the application. (Following the first action, applicants generally exchange multiple rounds of "rejection," correspondence, and amendment with the PTO before a final outcome.) The paper's design therefore requires a very strong assumption: that the content of the first action letter allows investors to predict (with certainty, or something close to it) whether an application will be granted. The authors provide some suggestive evidence that this is true, but it is clearly not in the case of pending applications whose patentability standard was changed by *Alice*.

[88] *Id.* at 676.

[89] *Id.*

[90] *See, e.g.*, Graham et al., *supra* note 25, at 1279–80 ("[E]arly-stage biotechnology companies are much more likely to use, and to see utility in using, the patent system."). *But see* Rachel E. Sachs, *The Uneasy Case for Patent Law*, 117 MICH. L. REV. 499, 501–02 (2018) (arguing that a rich innovation ecosystem exists in areas of biotechnology that are not readily patentable).

III. HOW CAN WE ISOLATE THE EFFECT OF PATENTS?

In a researcher's ideal world, patents would be randomly assigned to applicants, allowing us to fully isolate the effect of the IP right.[91] But while this might appeal to scholars, it obviously would not find favor with patent applicants. In lieu of such an experiment, we need a next-best alternative: a quasi-experiment in which the grant (or not) of a patent is as close to random as possible. In this Part, I describe a unique legal setting that provides such an opportunity.

A. The Empirical Design: Alice v. CLS Bank

The fourth and final case in the eligibility quartet—*Alice Corp. v. CLS Bank International*—has been the most influential. The controversy that accompanied the decision was matched by the immediacy of its impact. Following the case, a wave of patentable-subject-matter litigation swept federal courts. In the succeeding five years, hundreds of patent invalidations were handed down, spawning complaints from litigants of inconsistency and unpredictability in patentable-subject-matter jurisprudence.[92] The response from the PTO—the agency responsible for examining patent applications—was even swifter. On June 25, 2014, six days after *Alice* was handed down, the PTO issued its initial informal guidance to examiners.[93] Following the guidance, there was a precipitous increase in application rejections in the affected PTO examination units. As I show in Part IV, the average allowance rate in the three business-methods technology centers fell from around 45% to less than 10% in the months following the decision.[94]

This legal setting provides a unique opportunity to observe the effects of a stark and unexpected change in the PTO's examination practices. Applicants who received a final decision before June 25, 2014 were subject to a far more lenient patentability standard than those who received a decision afterwards. But patent applications typically take several years to be resolved, and applicants cannot control their examination dates with great precision.[95] In other words, in a small window of time around the PTO's

---

[91] For a nuanced discussion of the benefits of taking a more experimental approach to innovation policy—though of course she does not advocate a purely random system—see Lisa Larrimore Ouellette, *Patent Experimentalism*, 101 VA. L. REV. 65, 68 (2015).

[92] Lemley & Zyontz, *supra* note 7, at 49–50.

[93] Hirshfeld, *supra* note 31, at 1, 3 (noting that "mere instructions to implement an abstract idea on a computer" would not constitute "significantly more" than an unpatentable abstract idea).

[94] This is in line with prior work. *See* Colleen V. Chien & Jiun Ying Wu, *Decoding Patentable Subject Matter*, PATENTLY-O PAT. L.J. 10, 15 (2018); Kesan & Wang, *supra* note 25, at 535; Colleen V. Chien, Nicholas Halkowski, Maria He & Rodney Swartz, *Parsing the Impact of* Alice *and the PEG*, PATENTLY-O PAT. L.J. 20 (2020).

[95] The average application pendency in my sample is 2.6 years. *See* Table 1.

change in standards, the side of the line that an applicant fell on was effectively random. By tracking patenting and funding outcomes for these two groups of applicants—who differ, on average, *only* in the likelihood that they were granted a patent—we can isolate the role that the patent itself plays in securing subsequent funding.

To put this quasi-experiment into practice, I use a statistical tool known as regression discontinuity (RD).[96] The idea behind an RD is simple and intuitive. It is often difficult to estimate the effect of a specific law or policy simply by comparing those who were subject to it and those who were not, because the two groups likely differ in many *other* respects. Some of these other factors may be related to the likelihood of being targeted by the law in the first place, while others may not be easy to observe. However, in a situation where the policy of interest is implemented at some well-defined threshold that the affected groups can't easily control—for example, the value of a credit score or income, or a particular date—the people or companies close to the threshold on either side are likely to be similar on average in all ways *except* for the fact of treatment. By taking the difference in outcomes at the threshold for the groups on the left and right of that cutoff, we can obtain a credible estimate of the effect of the intervention alone. In this way, discontinuity designs can provide compelling quasi-experimental evidence of the effect of a law or policy change.[97]

In the study at hand, the treatment of interest—the "policy," in other words—is the *outcome of a startup's first-decided patent application*. I define this to be the first application for which a company receives a final outcome: either the application is granted by the PTO, or it is abandoned following a series of rejections. I use the first-decided application for two reasons. First, prior work suggests that the first grant—that is, the fact of owning *any* patent rather than not owning a patent—is the most likely to be consequential for early-stage investors.[98] Second, by using the first-decided application, I minimize any concern that a patentee whose application was abandoned following *Alice* might nevertheless hold a broader patent

[96] For explanations of this tool, see generally David S. Lee & Thomas Lemieux, *Regression Discontinuity Designs in Economics*, 48 J. ECON. LIT. 281 (2010), and MATIAS D. CATTANEO, NICOLÁS IDROBO & ROCÍO TITIUNIK, A PRACTICAL INTRODUCTION TO REGRESSION DISCONTINUITY DESIGNS: FOUNDATIONS (2019). The design used in the present study (like all RDs) is subject to some important but technical assumptions, which I discuss in Appendix B.

[97] Note that while I use quasi-randomization to motivate the design, I analyze the data using both the continuity and local randomization approaches to RD. *See infra* note 134 and accompanying text; Matias D. Cattaneo, Rocío Titiunik & Gonzalo Vazquez-Bare, *Comparing Inference Approaches for RD Designs: A Reexamination of the Effect of Head Start on Child Mortality*, 36 J. POL'Y ANALYSIS & MGMT. 643, 644 (2017).

[98] Farre-Mensa et al., *supra* note 85, at 30.

portfolio (not accounted for in this study) which could affect their chance of receiving funding. To measure companies' success, I use two types of outcome: first, the incidence and amount of investment and second, a startup's chance of being acquired or launching an IPO following the disposition of the application.

To divide the two groups, I use the date that the PTO released its informal *Alice* guidance to examiners: June 25, 2014.[99] This was sent one week after the decision came down. An application that was decided before this date is "pre-*Alice*," while an application decided after the threshold is "post-*Alice*," and was subject to more stringent examination.[100]

To establish the date of final disposition for applications that were granted, I use the date of issue of the patent.[101] For applications that were abandoned, the choice of a final decision date is more ambiguous. In practice, patent applications aren't denied; rather, they are considered abandoned if an applicant fails to respond to a notice of rejection from the PTO within six months.[102] This means that there is a six-month period during which the applicant could have made the actual decision to abandon an application, and so any date in that window is a potential option. However, because I assume that applicants learn of the newly heightened approval

[99] Memorandum from Andrew H. Hirshfeld, *supra* note 31.

[100] Note that my design is more complicated than the stylized one described above, because the discontinuity is not perfect. In settings where the threshold marks a crisp, deterministic assignment into treatment, the design is known as a "sharp" discontinuity. Consider, for example, a setting where every student with a GPA of 3.8 or greater receives academic honors, and every other student does not. That is not the case here. In the pre-*Alice* period, not all applications were granted, and following *Alice*, not all were rejected. In other words, while applicants were much less likely to be granted a patent after *Alice*, it was not impossible. In such cases, where the probability of treatment changes, we call the empirical setup a "fuzzy" discontinuity and adopt a slightly modified empirical approach. *See infra* Section IV.A; Matias D. Cattaneo & Rocío Titiunik, *Regression Discontinuity Designs*, 14 ANN. REV. ECON. 821, 827–29 (2022).

[101] The PTO also records the date on which it sends the applicant a "Notice of Allowance," which informs them that their patent has been granted. However, until the patentee has paid the issue fee, the PTO retains the power to revoke its allowance—a power which it can and does exercise, including after the *Alice* decision. *See, e.g.*, 37 C.F.R. § 1.313 (2022) (permitting applications to "be withdrawn from issue for further action at the initiative of the Office"); Gene Quinn, Alice v. CLS *Reality: PTO Pulling Back Notices of Allowance*, IPWATCHDOG (July 25, 2014, 4:44 PM), https://ipwatchdog.com/2014/07/25/alice-v-cls-reality-pto-pulling-back-notices-of-allowance [https://perma.cc/X573-W8PK] (observing a trend of "granted claims . . . being pulled back into prosecution only to be rejected" after *Alice*). For this reason, it would be inappropriate to use the grant date in this study.

[102] 35 C.F.R. §§ 1.134–.135 (2022). An applicant can also affirmatively abandon an application (though this is far less common). 37 C.F.R. § 1.138 (2022). In these cases, I treat the "decision date" as the day the PTO receives the request from the applicant.

standards by way of a rejection notice from the PTO, I assume that applications are abandoned as of the day that notice is mailed.[103]

B. Constructing the Dataset

This study requires a combination of data from multiple disparate sources. For data on patent applications, I start with the PTO's Patent Examination Dataset (PatEx), which contains information on every published application at the PTO.[104] I am interested in utility patent applications that received a final disposition in a window around the PTO's implementation of the *Alice* decision—June 25, 2014—but of course applications take some time to move through the examination process. So, in order to cast a wide net, I begin my data collection with decided applications that were filed as early as January 2007.[105] I end the sample with applications filed the day before *Alice*—June 18, 2014—so as to avoid any confounding from new applicants who might have behaved differently following the decision.

Next, I build an "assignment history" for every application in the broad sample.[106] Unfortunately, PatEx does not include information about the owner of the patent. So, to gather data on patent assignees, I combine various sources. To identify the initial assignee, I begin with PatentsView's "disambiguated assignees" dataset, which does a good job reconciling the same applicant across multiple applications.[107] In the event that PatentsView has no assignment recorded, I check whether the PTO has a record of an assignment to an employer in its database.[108] Together, these two sources account for 95% of the applicants in the sample. For the remaining 5%, I use

---

[103] The results using the *end* of the six-month window for abandoned applications are substantively similar.

[104] *See* Richard D. Miller, *Technical Documentation for the 2019 Patent Examination Research Dataset (Patex) Release* (U.S. Patent & Trademark Off., Econ. Working Paper No. 2020-4, 2020), https://www.uspto.gov/sites/default/files/documents/PatEx-2019-Technical-Doc.pdf [https://perma.cc/SBN7-8CPY].

[105] I also exclude any application with a priority date before January 1, 2007. *See* U.S. PATENT & TRADEMARK OFF., MANUAL OF PATENT EXAMINING PROCEDURE § 210 (9th ed. 2019, rev. Feb. 2023).

[106] In general, my procedure draws on those described in Donald E. Bowen III, Laurent Frésard & Gerard Hoberg, *Rapidly Evolving Technologies and Startup Exits*, 69 MGMT. SCI. 940, 942–44 (2023), and Farre-Mensa et al., *supra* note 85, at 645–47.

[107] *PatentsView Project Disambiguation Algorithms*, PATENTSVIEW, https://patentsview.org/disambiguation [https://perma.cc/4BHE-NHSS]. PatentsView is a PTO-affiliated website that provides a platform for visualization of various types of patent data. The process of "disambiguating" patent applicants is very important (and challenging) because the application data are generated by the applicants themselves. As a result, typographical errors are common, as are different spellings of the same name (for example, "1-CLICK L.L.C." and "1 CLICK LLC").

[108] *Patent Assignment Dataset*, U.S. PAT. & TRADEMARK OFF., https://www.uspto.gov/ip-policy/economic-research/research-datasets/patent-assignment-dataset [https://perma.cc/3KKJ-MWWL].

a mix of data from Google Patents[109] and the PTO's assignment database. Finally, because applicants often change names in the course of prosecution (and between different applications), I check the PTO's name-change data, and reconcile every name used by unique assignees.

During the course of examination, applications may be reassigned to other parties. To confirm the assignee of a granted patent, I again turn to PatentsView. For the small number of patents for which PatentsView is missing data, I supplement with data from Google Patents as well as a recent academic study, both of which report the assignee name as recorded on the issued patent.[110] For abandoned applications, I use the PTO and Google databases described above in order to identify the most recent assignee before disposition of the application. In most cases this is simply the initial assignee. Where a pending application has been reassigned to a party other than the initial applicant, I remove it from the data. I also remove any applications that are assigned to a non-U.S. entity.[111]

My focus is on young companies likely to be seeking venture capital investment—in other words, startups. Unfortunately, companies are not identified in this way in the PTO's data. Instead, I infer their status using a multistage process. I start by limiting the applications to those which were initially filed by PTO-designated "small entities" or "micro entities."[112] Next, I remove noncorporate entities. In many cases the only data available about the assignee are their name and country. So, in order to identify the type of identity, I use an algorithm to parse the names into one of several categories. I discard individuals, nonprofits, governments, universities, and other institutions (such as medical centers). To confirm that the remaining applicants are *standalone* companies, I check each name against a list of public companies and their subsidiaries, and exclude any company that either went public or became a subsidiary of a public company before the

[109] GOOGLE PATENTS, https://patents.google.com/ [https://perma.cc/X6FC-R5A4].

[110] *See* Noah Stoffman, Michael Woeppel & M. Deniz Yavuz, *Small Innovators: No Risk, No Return*, 74 J. ACCT. & ECON. 1, 5 (2022). The data are maintained and updated by Professor Mike Woeppel and are available at *United States Patent and Trademark Office Data*, MICHAEL WOEPPEL (Mar. 2023), https://www.mikewoeppel.com/data [https://perma.cc/66LN-EBUX].

[111] Of course, many non-U.S. applicants seek U.S. patents, but non-U.S. startups solicit investment from different sources and operate in legal contexts that may differ in many ways from their U.S. counterparts (not least the availability of domestic patent protection for software). To avoid any confounding, I focus on the United States in this study.

[112] For the purposes of the PTO, a "small entity" is a nonprofit, an individual, or a corporate entity with fewer than 500 employees (including controlling affiliates). 37 C.F.R. § 1.27(a) (2022). A "micro entity" is a small entity whose income was not higher than three times the median household income in the preceding year. 35 U.S.C. § 123.

disposition of its first application.[113] In order to zero in on the remaining applicants who are most likely to be new companies, I exclude any applicant who filed a patent application in the fifteen years preceding January 1, 2007.[114]

Finally, because my focus is on *Alice*-impacted applications, I further limit this sample to companies whose first-disposed application was for a business-methods patent. To identify the business-methods applications, I use the PTO's internal art unit assignments.[115] Once submitted to the PTO, an application is routed to one of several "technology centers," each of which contains a set of "art units"—groups of examiners who specialize in the relevant technology.[116] I filter to applications in the core business-methods technology centers: TC3620, 3680, and 3690.[117]

After all these steps, my sample comprises the set of published,[118] first-decided applications in a business-methods art unit for all standalone U.S.

---

[113] For companies that were acquired before the disposition of their application, I cannot separate the effect of a patent from that of being acquired. Consider, for example, that such a company would be unlikely to seek VC funding after disposition, but for reasons unrelated to whether their patent was approved or not.

[114] Note that this implicitly excludes any pre-2007 continuation applications.

[115] These are not the only technologies that were affected by *Alice*. In their study of the effect of *Alice* on application rejections, Kesan and Wang use a broader conception of *Alice*-impacted technologies, which also encompasses bioinformatics and some other computing-related areas. Kesan & Wang, *supra* note 25, at 533. Games and educational methods (AU 3714–3718) were also subject to significantly more strenuous examination after *Alice*. However, these art units have relatively few applications—particularly after subsetting to a narrow window around the *Alice* decision—which limits our ability to make statistical inference, and so I focus on business methods in this study.

[116] USPTO, MANUAL OF PATENT EXAMINING PROCEDURE § 903.08(a) (2018).

[117] *Business Methods*, U.S. PAT. & TRADEMARK OFF. (July 26, 2019), https://www.uspto.gov/patents/basics/types-patent-applications/utility-patent/patent-business [https://perma.cc/KC52-32G6]; *see also* Chien et al., *supra* note 94, at 22–23; Kesan & Wang, *supra* note 25, at 552–53 (both using similar approaches to identify business methods patents). However, this is not an incontestable choice; identifying all the patents for a particular area of invention is tricky. *See, e.g.*, John R. Allison & Ronald J. Mann, *The Disputed Quality of Software Patents*, 85 WASH. U. L. REV. 297, 304–13 (2007) (discussing the challenges of identifying a broad sample of software patents). One possible complication with using art units is that sophisticated parties—perhaps those with better inventions—might be more likely to write their patent claims in such a way as to try to avoid the more "difficult" business-methods art units. Indeed, there is some evidence for this in the literature (at least post-*Alice*). *See* Lefstin et al., *supra* note 47, at 591 ("[A]pplicants focus a lot of their strategic effort in drafting their patents so that they will be assigned to a technology center with a higher eligibility proclivity."). As a robustness check, I repeated my analysis using the broader Cooperative Patent Classification code (CPC G06Q) to identify business-methods applications. The results are substantively the same.

[118] I do not have access to *unpublished* patent applications. To explain how this could potentially bias the results requires a brief introduction to patent examination procedure. Since the passage of the American Inventors Protection Act (AIPA), patent applications are presumptively public documents. American Inventors Protection Act of 1999, 35 U.S.C. § 122(b) (2022). Under the AIPA, applications automatically publish eighteen months after they are filed, unless either the applicant affirmatively

startup companies who filed a patent application between January 2007 and June 2014. To zero in on the effect of *Alice*, I focus my analysis below on applications with a final resolution date no more than two years before or after the decision.

To obtain data about investments and company outcomes, I turn to Pitchbook, a proprietary online database of venture capital and private equity funding rounds.[119] For each round of investment, the Pitchbook data include funding dates, the stage of funding, and other information. The data also include the current status of the company, including whether it was acquired, was liquidated, or launched an IPO. Unfortunately, investment databases have some limitations. Because VC transactions are private investments in private companies, there is no standardized reporting structure. Most data sources therefore rely on the deal being reported by the target company, the investors, or a fund.[120] Nevertheless, previous research indicates that the leading data providers—including Pitchbook—do a reasonable job at tracking funding rounds.[121] Like most providers of investment data, Pitchbook tracks a variety of different funding sources, but not all are relevant to this study. In order to focus on equity investment, I exclude government grants and collateralized debt.

Combining patent application and funding data is a nontrivial enterprise.[122] The assignment datasets and Pitchbook do not share a common

requests nonpublication at the time of filing (in doing so, the applicant must forgo any foreign patent application) or abandons the application before the eighteen-month mark. Prior work suggests that across entity size and technology type, few applicants elect for nonpublication at the time of filing, and that around 15% of applications are ultimately abandoned without publication. Deepak Hegde & Stuart Graham, *Disclosing Patents' Secrets*, 347 SCI. 236, 236–37 (2015) (noting that "pregrant disclosure is preferred overwhelmingly in every technology sector"); Farre-Mensa et al., *supra* note 85, at 645. Of course, if abandoned applications that are unpublished are *systematically* different to those that are published, my sample could have some selection bias. It is difficult to know exactly what the magnitude of that bias would be—if any—but I note that by limiting my sample to applications before *Alice*, I avoid capturing any strategic increase in the use of nonpublication requests following the decision.

[119] PITCHBOOK, https://www.pitchbook.com [https://perma.cc/AHV5-Z7F8].

[120] The database providers also file FOIA requests with public pension funds and draw from publicly available sources, such as the SEC's Form D filings, which provide notice of the sale of securities, or S-1 statements in the event that a startup subsequently goes public.

[121] Steven N. Kaplan & Josh Lerner, *Venture Capital Data: Opportunities and Challenges* (Nat'l Bureau of Econ. Rsch., Working Paper No. 22500, 2016). The challenge (and cost) of tracking this opaque information is an important reason for the relative quality of proprietary over open-source databases. Note also that to the extent that Pitchbook data exhibit some kind of selection bias, it should be the same for both the pre- and post-*Alice* groups in my study. In other words, there's no reason to think that any selection bias would differ between the groups. Andre Retterath & Reiner Braun, Benchmarking Venture Capital Databases 27 (Sept. 17, 2020) (unpublished manuscript), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3706108 [https://perma.cc/6XGJ-BK7G].

[122] Josh Lerner & Amit Seru, *The Use and Misuse of Patent Data: Issues for Corporate Finance and Beyond*, 35 REV. FIN. STUD. 2667, 2672 (2022).

identifier for companies, so pairing applicant data with funding outcomes involves matching the text of companies' names and locations. Although the Pitchbook data are clean and consistent, the assignment data are mostly applicant-generated and typographical errors are common.[123] I take two approaches to these problems. First, I clean the patent assignee and Pitchbook "target company" names to standardize corporate entity descriptions, removing punctuation and suffixes.[124] Using these cleaned names, I produce lists of unique applicants and funding target companies and perform an initial "exact" match. For the remaining applicants and companies, I perform a fuzzy match, computing the similarity between the names of the assignees and target companies using a statistical algorithm.[125] In summary, I find at least one venture capital transaction for 624 out of 2,234 total patent applicants. Table 1 summarizes the data.

TABLE 1: SUMMARY STATISTICS FOR FIRMS' FIRST-DECIDED APPLICATIONS

	Rejected	**Granted**
Number of Unique Applicants	1,747	487
Earliest Disposition	June 25, 2012	June 26, 2012
Latest Disposition	June 24, 2016	May 24, 2016
Application Pendency (Average)	2.6 years	2.8 years
Application Pendency (Standard Deviation)	1.3 years	1.4 years
Number of Companies Receiving *Any* Funding	474	150
Percent Receiving Pre-Disposition Funding	24.2%	25.9%
Percent Receiving Post-Disposition Funding	13.4%	19.5%
Percent Subsequently Acquired	11.6%	14.6%
Average Post-Disposition Investment	$4.4m	$10.6m

[123] *See supra* note 107 (discussing the process for cleaning applicant data).

[124] I discuss the cleaning and matching processes in detail in the Appendix. *See infra* Appendix A.2.

[125] *See infra* Appendix A.3.

IV. RESULTS: DO PATENTS AFFECT INVESTMENT OUTCOMES?

In this Part, I turn to my central question: do patents drive investments in startups? I begin by examining whether, and to what degree, the *Alice* decision impacted patent approval rates. Next, I use these findings to analyze the impact of patents on the startups' success on two fronts: first, the incidence and amount of investment, and second, the likelihood that a startup was acquired or launched an initial public offering. Despite the suggestive evidence in the prior literature, I find no evidence that patent grants drive either venture capital investments or later success.

A. First Stage: Did Alice *Matter?*

The first step is to estimate the impact of *Alice* on the treatment. In other words, if *Alice* had an effect on patent examination for business-methods inventions, we would expect to see a decline in allowance rates following the decision. Figure 1 shows the average allowance rate organized by date of final disposition over a window of two years before and after *Alice*.[126] Each point reflects the average rate over approximately two weeks of dispositions. The difference is immediately apparent: average allowance rates for first-time business-methods applicants fell from 45% to around 10% in the months following *Alice*. The decision had an unambiguous—and very fast—impact on business-methods applicants' chances of receiving a patent.

[126] In econometric parlance, this is called the "first stage." *See* Lee & Lemieux, *supra* note 96, at 300–01.

FIGURE 1: THE EFFECT OF *ALICE* ON PATENT GRANT RATES FOR BUSINESS-METHODS APPLICANTS' FIRST-DECIDED APPLICATION

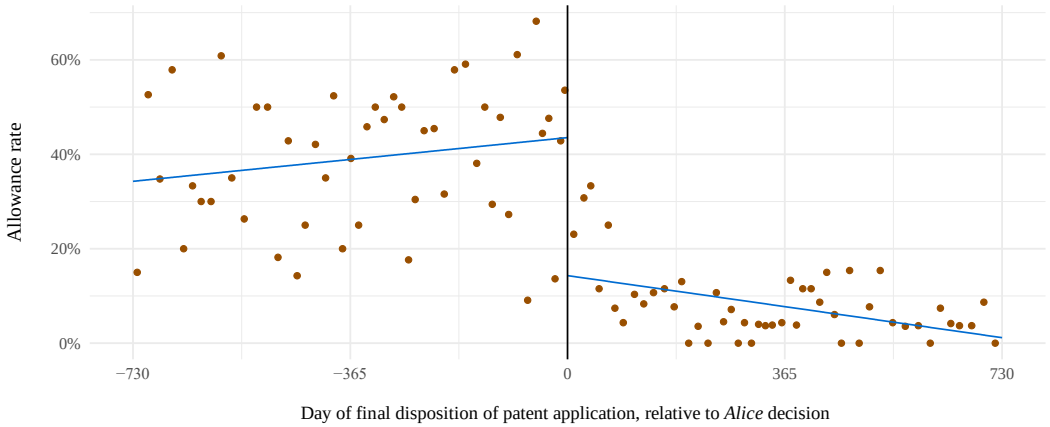

Note. Each point reflects the average of approximately two weeks of observations.

B. Funding

Having established that *Alice* caused a sharp decline in patent grants, I turn to outcomes. If it is true that patent ownership affects investment decisions, then we would expect to see a corresponding drop in the share of companies that received venture capital funding after *Alice*. To test this, Figure 2 shows the rate of post-disposition funding for applicants, again grouped into two-week bins. Although the average likelihood of funding is somewhat noisy, ranging from 0%–30% of companies in each month, there is no evidence of any *change* in investment, discontinuous or otherwise, in the wake of the PTO's new *Alice* standards. Indeed, the rate of funding appears to be steady throughout the four-year window, regardless of when an application was finally decided.

FIGURE 2: SHARE OF BUSINESS-METHODS APPLICANTS RECEIVING INVESTMENT AFTER THE RESOLUTION OF THEIR FIRST-DECIDED APPLICATION, GROUPED BY THE DATE OF DECISION

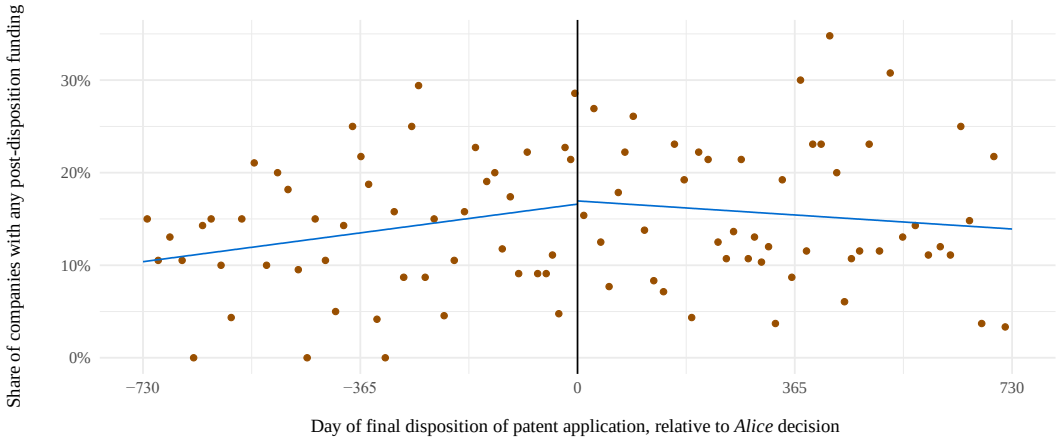

Note. Each point reflects the average of approximately two weeks of observations.

Perhaps, however, the difference is not in the rate of funding, but in the amount invested? Although I don't observe any change in the likelihood of investment, it's possible that startups whose patent applications are granted raise more money in absolute terms. To test this alternative, I narrow my focus to the subset of startups who received at least *some* funding after the disposition of their first application. In Figure 3, I plot the average total investment for these companies, once again grouped by the two-week period in which their application was decided. Figure 3 shows that, amongst companies who received post-decision investment, the average amount invested ranged between $1 million and $100 million. But the graph also tells the same story as with the incidence of funding: there's no evident change in the average amount of funding after *Alice*.[127] I find no evidence that a patent grant leads to a change in the average amount invested in startups.

[127] Amongst companies that received at least some funding after disposition of their application, the average total post disposition funding was $46.9 million for pre-*Alice* dispositions and $42.6 million for post-*Alice* dispositions, with standard deviations of $94.8 million and $87.2 million, respectively.

FIGURE 3: INVESTMENT FOR BUSINESS-METHODS APPLICANTS AFTER THE RESOLUTION OF THEIR FIRST-DECIDED APPLICATION, GROUPED BY THE DATE OF DECISION

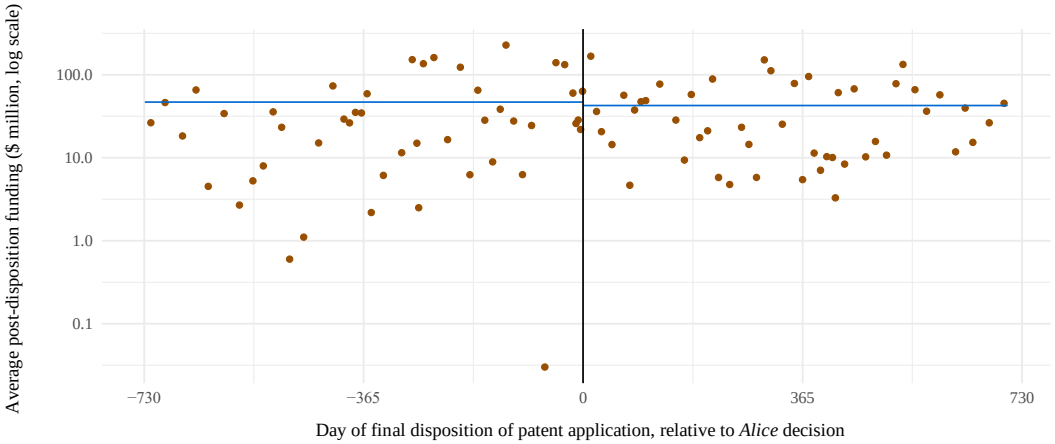

Note. Companies that received no funding are dropped. Each point reflects the average of approximately two weeks of observations.

In the remainder of this Section, I repeat these analyses with more formal statistical tests. But the graphical results convey the core story: despite a precipitous fall in allowance rates for business-methods patents, the likelihood of funding and the amount invested remained unchanged compared to the companies whose applications were resolved before *Alice*. Again, I find no evidence that patents affect investment in business-methods startups.

Table 2 provides formal tests of Figures 1 and 2. The four columns show results across a range of different specifications and estimation windows. In practice, the implementation is as follows: I choose various numbers of months before and after the *Alice* threshold, and then fit a regression line on each side. The difference between these curves at the point of the *Alice* decision forms the estimate. Column (1) shows the results for a linear regression fit across the full pre- and post-*Alice* windows. Columns (2)–(4) show the results of difference-in-means estimates for windows of six, twelve, and twenty-four months. The first row, which is analogous to Figure 1, shows the results of a regression of patent approval on post-*Alice* disposition. If *Alice* diminished patent allowances, we should see a strong and negative relationship—and indeed we do: on the order of negative 30 percentage points. The second row, which is analogous to Figure 2, shows results of a regression of investment likelihood on post-*Alice* disposition.[128] Again, if patent ownership drives investment, we should

[128] In technical terms, this is the "reduced form" or "intention-to-treat" parameter.

expect to see a drop in the likelihood of funding amongst the post-*Alice* group, who are much less likely to hold a patent. But here, on the other hand, the effect is approximately zero.

Note, though, that this row is an estimate of *Alice*'s effect on the chance of receiving investment—that is, it compares companies who received their disposition before the decision with those who received it after. Because not every company received a patent before *Alice* (and vice versa), we need to make a further adjustment in order to estimate the effect of a patent grant.

TABLE 2: THE EFFECT OF *ALICE* ON PATENT APPROVAL AND LIKELIHOOD OF FUNDING

	(1)	(2)	(3)	(4)
	Linear	Difference-in-Means		
Window (months)	24	6	12	24
First Stage: Patent Approval	–0.292 (0.035)	–0.263 (0.036)	–0.333 (0.024)	–0.312 (0.017)
Reduced Form: Likelihood of Funding	0.003 (0.030)	–0.002 (0.031)	–0.001 (0.020)	0.018 (0.015)
First-Stage F-Statistic	69	53.9	193.2	322
Observations	2234	593	1222	2234

Note. This Table shows formal tests of Figures 1 and 2 over a range of different windows. The first row regresses "first patent approval" on "post-*Alice* disposition" and shows a significant decline in patent approval after *Alice*. The second row regresses "receipt of funding after disposition of the application" on "post-*Alice* disposition" and shows no evidence that applications decided after *Alice* were less likely to receive funding. Robust standard errors are in parentheses.

In Table 3, I show the results of the fuzzy RD, which uses the estimates above in order to zero in on the role played by patents.[129] Again, the columns contain the results of several different regression models. In every case, the dependent variable is the same: whether the startup company received any venture capital equity investment following the first decision on the company's patent applications. The first row—"patent approved"—reports the key result: it shows the effect of obtaining a patent on receiving post-disposition investment.

The first column, labeled "OLS," provides some context: it reports the simple relationship between the presence of post-disposition funding and approval of a company's first patent application.[130] This is a replication of the correlational findings in the prior literature, which are described in Section II.B. As expected, I find a strongly positive and statistically significant relationship. The coefficient implies that having the first patent application approved is associated with a 6.1 percentage-point (roughly 40%) increase in the likelihood of subsequently receiving funding. But we know from the earlier discussion that this estimate is biased in the absence of other information about the quality of the invention, the team, and so on.[131]

Columns (2)–(5) show the results of the (fuzzy) regression-discontinuity approach. Each column uses a slightly different specification, but the goal is always the same: to estimate the effect of the first patent approval on the likelihood of subsequent investment. As above, columns (2)–(5) show results for a range of different specifications and windows.

The results in columns (3) and (4) indicate an effect of approximately zero, while columns (2) and (5) are actually slightly negative.[132] But more importantly, none of these RD estimates are significantly different from zero.[133] This accords with expectations from Figure 2 and Table 2 above: as

---

[129] In practice, the fuzzy RD estimate is simply the ratio of the reduced form and the first stage estimates. Intuitively, the idea is to take the post-*Alice* change in investment and adjust it to account for the fact that not all the companies' application outcomes were changed by the decision (some would have been rejected regardless of the legal regime; a small handful would have been approved regardless). Of course, to the extent that the effect of *Alice* (the numerator) is zero, the ratio is also zero.

[130] OLS refers to ordinary least squares. The equation is $y_i = \alpha + \beta x_i + \epsilon_i$, where y is a 0/1 indicator for whether funding was received after the disposition of the first application, and x is a 0/1 indicator for the approval of the first patent application.

[131] *See supra* Section II.B.

[132] The variability between the point estimates in Table 3 is a consequence of the noise in the outcome: because there is a good deal of variation in funding outcomes over time, each different model is sensitive to the particular observations that are included as the window size changes.

[133] I report heteroskedasticity-robust standard errors at the 95% level. Recent results imply that these standard errors are too small when $F < 104.7$ but applying an adjustment factor to columns (2) and (3)

those figures suggest, there is no evidence of a change in funding outcomes at the point of the *Alice* decision. This flows through to the fuzzy RD estimate: the results show no evidence that patent approval has an effect on obtaining funding for business-methods startups.

TABLE 3: FUZZY RD ESTIMATES: THE EFFECT OF PATENT APPROVAL ON THE LIKELIHOOD OF POST-DISPOSITION FUNDING

	(1)	(2)	(3)	(4)	(5)
	OLS	Linear	Difference-in-Means		
Window (months)	24	24	6	12	24
Patent Approved	0.061 (0.020)	–0.012 (0.102)	0.006 (0.117)	0.002 (0.062)	–0.057 (0.049)
Observations	2,234	2,234	593	1,222	2,234
Mean of Outcome	14.7%	14.7%	16.7%	14.9%	14.7%

Note. In every case, the dependent variable is an indicator (0/1) for whether a startup received any funding after the disposition of their first patent application. Column (1) shows the simple (i.e., noncausal) relationship between a patent grant and investment. Columns (2)–(5) show various specifications for the discontinuity (causal) approach, estimated using two-stage least squares. Robust standard errors are in parentheses.

One drawback of the conventional RD approach is that its estimates can be noisy, particularly with small sample sizes, which leads to the need to extrapolate by using data far from the threshold. A natural robustness check uses the local randomization interpretation of RD.[134] Local randomization

would not change the substantive conclusion. *See* David S. Lee, Justin McCrary, Marcelo J. Moreira & Jack Porter, *Valid* t*-Ratio Inference for IV*, 112 AM. ECON. REV. 3260, 3270–72 (2022).

[134] The "local randomization" interpretation of discontinuity designs is conceptually distinct from the more commonly invoked "continuity" interpretation. Jasjeet S. Sekhon & Rocío Titiunik, *On Interpreting the Regression Discontinuity Design as a Local Experiment*, 38 ADVANCES IN ECONOMETRICS 1 (2017). Continuity, which I have thus far assumed, requires only that the potential outcomes are smooth at the threshold. In other words, in the counterfactual world *without* the *Alice* decision, we assume that none of the outcomes would have changed suddenly in mid-2014. Local randomization, on the other hand, requires something more: that the value of the running variable (in our case, the date of application disposition) is as-if randomly assigned to each applicant within a small window around the threshold. In most settings, this is a very strong assumption. (To see why, consider the relationship between students' GPA and test scores. We would expect average test scores to increase along with average GPA, and so at any arbitrary threshold the students on one side would not be

requires a stronger assumption about the research design, but in return offers hypothesis testing that is reliable in small samples.[135] In Tables 4 and 5, I show the results of tests of the sharp null hypothesis of no difference between the rate and amount of funding before and after *Alice*.[136] In line with the results above, I cannot reject the sharp null hypothesis of no effect, regardless of the window of observation.

TABLE 4: LOCAL RANDOMIZATION TEST OF NO DIFFERENCE IN INCIDENCE OF FUNDING BETWEEN PRE- AND POST-*ALICE* DISPOSITIONS

Window (days)	60	90	120	150	180
p-value	0.858	0.756	0.587	0.907	0.911
Observations (before, after)	(125, 101)	(158, 163)	(210, 207)	(263, 256)	(286, 305)

Note. The numbers in parentheses indicate the number of included observations before and after *Alice*, respectively. *p*-values are calculated using Fisherian exact inference.

TABLE 5: LOCAL RANDOMIZATION TEST OF NO DIFFERENCE IN AMOUNT OF FUNDING BETWEEN PRE- AND POST-*ALICE* DISPOSITIONS

Window (days)	60	90	120	150	180
p-value	0.822	0.712	0.865	0.718	0.607
Observations (before, after)	(21, 15)	(24, 26)	(29, 34)	(35, 38)	(41, 46)

Note. The companies are subset to those that received at least one funding round after disposition. *p*-values are calculated using Fisherian exact inference.

C. Acquisitions and IPOs

Of course, the amount of investment is not the only way in which a patent might affect the success of a startup company. If a patent grant

"exchangeable" with the students on the other.) In the setting at hand, however, randomization is very plausible. As discussed above, I assume that the precise day that an application was decided was indeed effectively random—dependent on examiner workload, application filing date, and myriad other administrative factors.

[135] Cattaneo et al., *supra* note 97, at 677.

[136] *See* PAUL R. ROSENBAUM, DESIGN OF OBSERVATIONAL STUDIES 30 (2010).

increases the latent value of a company—whether by signaling team competence, the stage of product development, or the existence of a legal monopoly—we might also expect the company to be a more appealing target for acquisition or an initial public offering.[137] Because most startups fail, these events are in a sense the true outcome of interest. Capital raised is a proxy for success, but what investors seek is a profitable "exit" from their investments. The simple comparison of successful and unsuccessful patent applicants in Table 1 suggests that there is a modest difference in post-disposition acquisition activity: 14.6% of startups whose first patent was granted were subsequently acquired versus only 11.6% for unsuccessful companies. But once again, there's good reason to think that this simple relationship doesn't tell the full story in the absence of more information about the quality of the inventions, the product teams, and so on. We can use the same approach as in Section IV.A to test whether patent approval has a *causal* impact on exit outcomes.

Figure 4 shows the percentage of startup applicants who were acquired or entered into a merger after the disposition of their first application.[138] Recall that Figure 1 revealed a sharp drop in patent grants after *Alice* was decided. Just as before, if it's true that patents drive acquisition activity, then we would expect to see an accompanying decline in that activity immediately following the decision. But once again, there is no evident effect at the threshold. Although acquisition activity ranges between 0% and 20%, the overall trend is steady regardless of when the application was decided. Tables 6a and 6b show the results of formal statistical tests, which support the null result.

[137] Despite the popular attention lavished on IPOs, acquisitions by another company are by far the most common path to exit for startup investors. *See, e.g.*, Brian Broughman & Jesse M. Fried, *Carrots and Sticks: How VCs Induce Entrepreneurial Teams to Sell Startups*, 98 CORNELL L. REV. 1319, 1322 (2013) ("[T]rade sales [by VCs] are actually much more common than IPOs and, in the aggregate, are likely to be almost as financially important to VCs." (citing Darian M. Ibrahim, *The New Exit in Venture Capital*, 65 VAND. L. REV. 1, 12 (2012))).

[138] I exclude management or leveraged buyouts orchestrated by private equity firms, which are closer to further investment rounds than "true" corporate acquisitions or mergers. *See generally* Steven N. Kaplan & Per Strömberg, *Leveraged Buyouts and Private Equity*, 23 J. ECON. PERSPS. 121 (2009) (distinguishing leveraged buyouts from venture capital firm investment because, in part, the former involves acquiring majority control while the latter does not).

FIGURE 4: SHARE OF COMPANIES ACQUIRED AFTER DISPOSITION OF THEIR FIRST-DECIDED APPLICATION

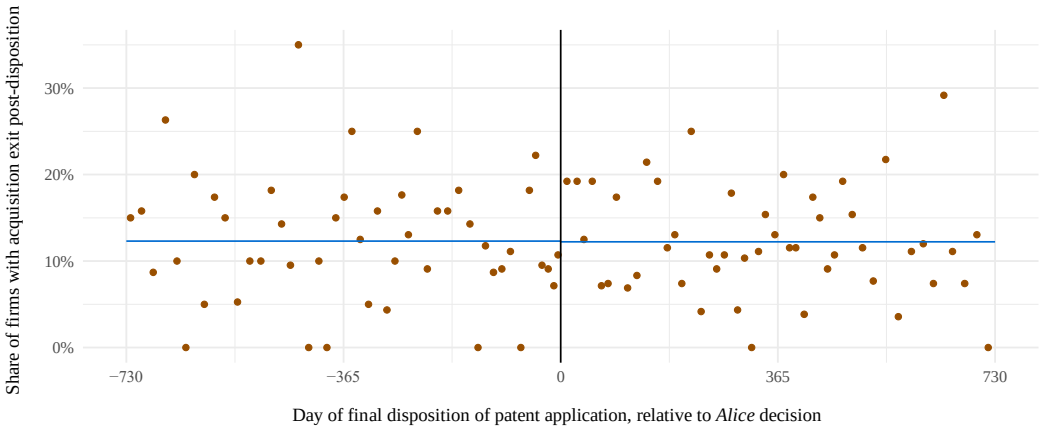

Note. Each point reflects the average of approximately two weeks of observations.

TABLE 6A: THE EFFECT OF *ALICE* ON SUBSEQUENT ACQUISITION

Window (months)	6	12	24
Reduced Form: Likelihood of Acquisition	0.032	0.003	–0.001
	(0.027)	(0.019)	(0.014)
Observations	593	1222	2234

Note. This Table shows the effect of *Alice* on subsequent acquisition (a formal test of Figure 4). Robust standard errors are in parentheses.

TABLE 6B: THE EFFECT OF PATENT APPROVAL ON SUBSEQUENT ACQUISITION

Window (months)	6	12	24
Fuzzy Discontinuity: Patent Approval	–0.120	–0.008	0.003
	(0.105)	(0.057)	(0.045)
Observations	593	1222	2234

Note. This Table shows the effect of patent approval on subsequent acquisition using two-stage least squares (the fuzzy discontinuity approach). Robust standard errors are in parentheses.

At the same time, the mere fact of acquisition may hide important variation. Not all startup acquisitions are alike: some are essentially fire sales—liquidations by another name—while others are true value-enhancing purchases.[139] If it's the case that acquisitions continued at the same rate after *Alice*, but that they changed in *kind*, then Figure 4 might miss an important effect. To check for this, I subset the companies to those that experienced a "successful" exit event after the decision on their application. I define success as either launching an IPO or being acquired for more than the sum of previously raised capital.[140] Because relatively few startups actually experience such an exit, the number of observations is much lower, and the graphical presentation would be very noisy. Instead, Table 7 shows the results of the same permutation test as before and also includes an estimate of the average difference between the rates of successful exit before and after *Alice*. Once again, as the large *p*-values indicate, I cannot reject the null hypothesis of no difference between "successful" downstream outcomes for the two groups of applicants.

In summary, there is no evidence that *Alice* has an effect on exit outcomes for business-methods software firms. Neither the likelihood of being acquired nor the chance of a successful exit appears to be driven by the approval of a firm's first patent.

[139] There are, of course, many reasons that a startup might be an appealing acquisition target, including the novelty of the product, the quality of the team (i.e., "acquihires"), market conditions, their IP portfolio, and more. But again, the empirical design here allows me to abstract away from these details on the assumption that within a window around *Alice*, the applicant-companies are *on average* similar to each other in all ways except for their patent outcomes.

[140] This is a conservative measure of success, but it may still be underinclusive if Pitchbook's data on funding amounts are missing or inaccurate in some nonrandom way. This is possible; indeed, in my sample 15% of the transactions do not report the amount invested. On the other hand, to the extent that there is a selection bias in what VC databases report, it is generally thought to be in favor of larger deals. *See* Retterath & Braun, *supra* note 121, at 17. This should mitigate the effect of any bias when trying to identify a successful outcome.

TABLE 7: LOCAL RANDOMIZATION TEST OF NO DIFFERENCE IN LIKELIHOOD OF "SUCCESSFUL" EXIT BETWEEN PRE- AND POST-*ALICE* DISPOSITIONS

Window (months)	60	90	120	150	180
Average Difference After *Alice*	0.005	0.000	0.010	0.000	–0.024
p-value	0.917	1.000	0.854	0.860	0.606
Observations (before, after)	(46, 37)	(54, 54)	(74, 67)	(98, 86)	(109, 104)

Note. A "successful" exit is defined as either launching an IPO or being acquired for more than the sum of invested capital. See main text for further discussion. *p*-values are calculated using Fisherian exact inference.

V. PROBING ALTERNATIVE EXPLANATIONS

The results in Part IV provide strong evidence that patents do not affect the course of investment in early-stage software companies, but they rely on several substantive assumptions which have thus far been implicit. In this Part, I surface these assumptions, exploring and testing three plausible alternative explanations for the null findings. First, I consider whether a previous Supreme Court case, *Bilski v. Kappos*, may have been the true "shock" to the software-patent system.[141] Second, I check for potential changes in the broader investment landscape that might lead to an underestimate of the effect of the patent. And third, I consider possible selection effects in applicants' behavior.

A. What About Bilski*?*

The most obvious challenge to the empirical design in this Article is that *Alice* was not the Supreme Court's first foray into business-methods software. The Court's first decision in the recent patent-eligibility quartet, *Bilski v. Kappos*, addressed very similar subject matter. The majority in *Bilski* found the claimed invention, which was based on an application of the financial practice of hedging, to be an unpatentable abstract idea.[142] Following the decision, some commentators claimed that applicants began

[141] 561 U.S. 593, 611–12 (2010).

[142] *Id.* at 611–12.

to engage in strategic games, drafting patents so as to try to avoid business-methods examination art units at the PTO (or simply switching to nonpatent sources of protection altogether). If this is true, then it is conceivable that *Bilski* rendered patents functionally irrelevant for software inventors—and, by implication, meant that investment decisions were already agnostic to patents in this industry. But while *Bilski* was a momentous decision for patent lawyers and commentators, there's less evidence that it was such a seismic event for investors—or indeed represented a material change for patent applicants.

It's true that the seeds of the confusion generated by *Alice*'s abstract-idea test were already planted in *Bilski*,[143] but the consequences of the two decisions differed significantly. First, the *Bilski* Court's explicit rejection of any sort of categorical exclusion for business methods[144] seems to have left the door open for significant activity in business-methods patent applications. As Figure 5 indicates, the number of applications to business-methods art units climbed steadily in the years following *Bilski*, reaching its peak at the *Alice* decision (before declining in the wake of the PTO's new *Alice* standards).[145]

[143] *See* Kevin Emerson Collins, Bilski *and the Ambiguity of "An Unpatentable Abstract Idea,"* 15 LEWIS & CLARK L. REV. 37, 38–39 (2011).

[144] *See* Lemley et al., *supra* note 46, at 1316.

[145] For similar data going slightly farther back in time, see Dennis Crouch, *The Business Method Patent Art Units*, PATENTLY-O (July 18, 2010), https://patentlyo.com/patent/2010/07/business-method-patents.html [https://perma.cc/Y53J-RW97], which provides data on patents issued by business-methods art units between 2005 and 2010.

FIGURE 5: NUMBER OF NEW APPLICATIONS BY SMALL ENTITIES IN BUSINESS-METHODS ART UNITS

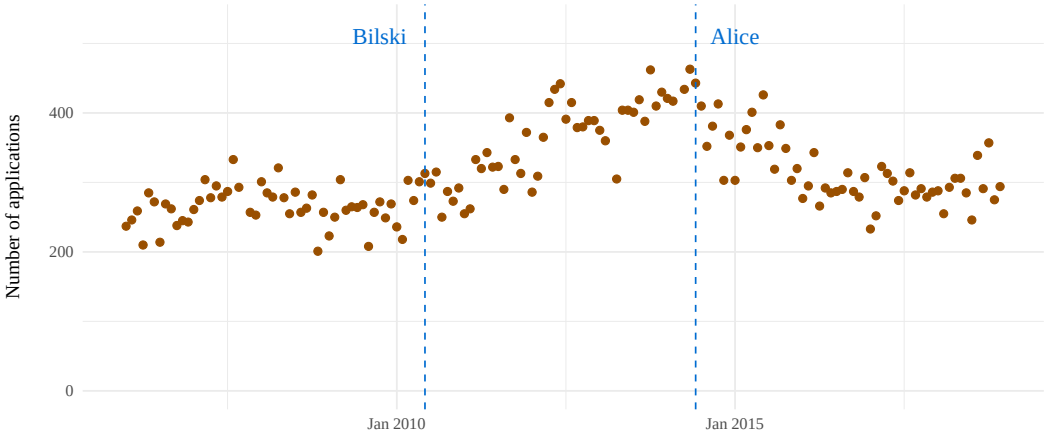

Notes. Each point reflects one month of applications The blue dotted lines indicate the dates of *Bilski v. Kappos* and *Alice v. CLS Bank*.

Most importantly, unlike *Alice*, *Bilski* did not lead to widespread rejections of patent applications. To put some data behind this, Figure 6 shows the share of business-methods office actions which contained a rejection on the basis of § 101 patentable subject matter.[146] If, in the wake of *Bilski*, it was the case that better inventions were able to select out of business-methods art units through strategic claim drafting, then we would expect to see rejection rates increase in those art units. However, the evidence suggests quite the opposite: there was a steady decline in rejections based on patentable subject matter in the post-*Bilski* period, which reached a low point just before *Alice*.[147]

[146] The noisy data in 2008 reflects the limitations of the office-action dataset. The PTO has only released granular office-action data for applications filed from January 2008 onwards. *See* Qiang Lu, Amanda F. Myers & Scott Beliveau, *USPTO Patent Prosecution Research Data: Unlocking Office Action Traits* 2 (U.S. Patent & Trademark Off., Working Paper No. 2017-10, 2017). Since the average application takes some time to receive an initial response from the PTO, the data at the beginning of that period are very sparse.

[147] Note that Figures 5 and 6 do not imply a causal relationship between the *Bilski* decision and what followed. The decision was not the only word on software patents in this period. Both the Federal Circuit and the PTO (through updated guidance to examiners and internal examination practices) potentially shaped the law and practice of software patenting. *See, e.g.*, Charles Duan, *Examining Patent Eligibility*, 96 ST. JOHNS L. REV. (forthcoming 2023) (manuscript at 24, 71–72), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4388598 [https://perma.cc/J3AE-MBYT] (arguing that updates to USPTO guidance on patent claim eligibility from 2009 to 2012 and the Federal Circuit's "increasing friendliness to software-implemented business methods" led to an increase in filing of these applications and higher success rates at overcoming § 101 rejections). It could also be the case that applicants simply learned how to draft more successful claims as they internalized the lessons of *Bilski*. However, what *is* clear from this evidence is that business-methods inventions continued to be a vibrant area of patent activity after *Bilski*.

FIGURE 6: SHARE OF ALL OFFICE ACTIONS FOR APPLICATIONS IN BUSINESS-METHODS ART UNITS THAT CONTAINED A REJECTION BASED ON § 101 PATENTABLE SUBJECT MATTER

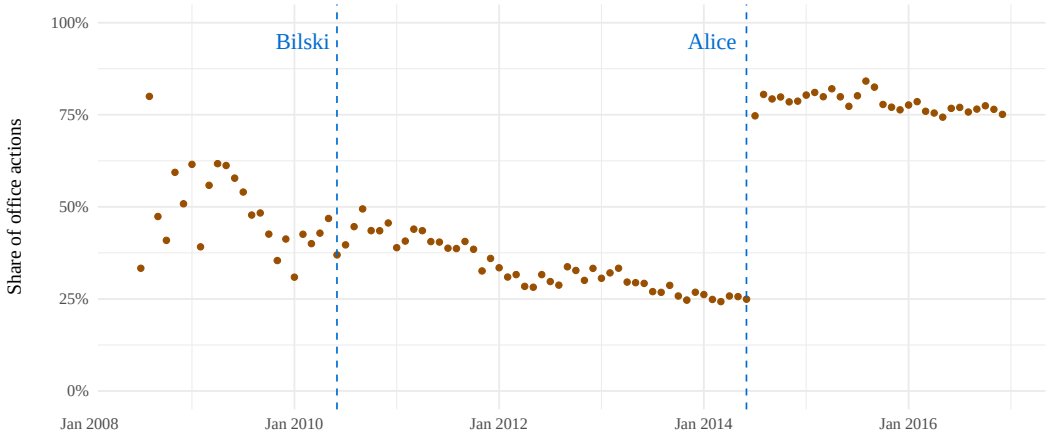

Note. The blue dotted lines indicate the dates of *Bilski v. Kappos* and *Alice v. CLS Bank*.

In summary, there's little evidence that *Bilski* was the true "shock," and grant rates in the affected art units did not decline after decision. On the contrary, the transformation in the PTO's approach to business-methods patents is visible only after *Alice*.

B. Did Investment Decline After Alice*?*

A related, but slightly different, explanation for my findings is that if investment *did* decline after *Alice*, it might have done so in a way that would not be detected by this empirical design. The potential problem is as follows: investment (and acquisitions) take place over a relatively long timeframe, but I focus on patent application dispositions that happened either just before or soon after *Alice* was decided. As a result, many of the actual VC investment choices (and nearly all of the exit outcomes) for both groups of companies happened *after* the decision. If the Supreme Court's opinion had the effect of chilling overall investment in software startups, then it would stand to reason that we would fail to detect any effect of patent approval because downstream investment would be suppressed for *all* applicants, regardless of when they received a decision.[148]

The most direct way to assess this concern is to look at the broader investment environment. Separate from the specific applicants in question here, was there a general move away from investment in business-methods

[148] Indeed, some of the results in Professor Taylor's survey of investors might prompt this concern. *See* Taylor, *supra* note 75, at 2055 (reporting that 59% of responding investors said that they "strongly agreed" or "somewhat agreed" that they would be less likely to invest in an industry if patents became harder to obtain).

software? To answer this, Figures 7 and 8 show the annual number of transactions (orange dots) and total invested capital (blue bars) for investments in financial software. (These are only two of several possible subindustries, of course, but it is broadly representative of investments in business-methods software.) In short, investments did not decline after 2014—in fact they increased quite dramatically. Aggregate investment grew every year, nearly doubling between 2014 and 2015 in the case of fintech. For both subcategories, the number of completed deals rose year-on-year throughout the decade.[149] Although we cannot observe the counterfactual investments that would have taken place in a world without the Court's decision, we can rule out any general decline that might confound the results of the study.

A more nuanced approach to changes in investment behavior is to look at the stage of funding reached by companies on either side of the *Alice* line. If, for example, overall investment remained the same (or increased) despite the disappearance of the patent signal, it may have been because venture capitalists looked to alternative signals of quality. One obvious possibility is that they focused on later-stage financing (where companies have a more developed track record). If post-*Alice* investment targets were more likely to be later-stage companies, then investment may have remained high, thus disguising the effect of patent approval.

FIGURE 7: AGGREGATE INVESTMENT IN "FINANCIAL SOFTWARE"

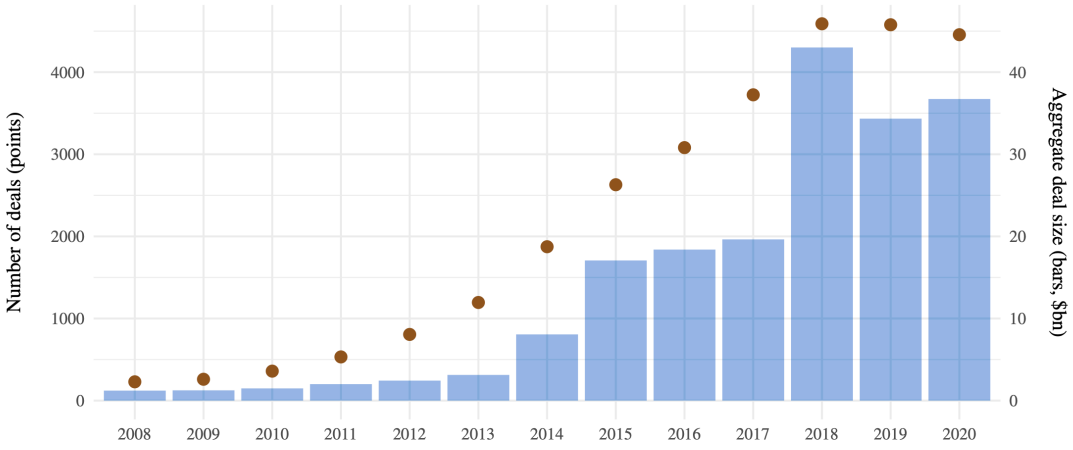

Note. Orange dots reflect the number of deals, and blue bars reflect the total invested capital for that year.

[149] Of course, there is no reason to think this *increase* in investment activity is causally related to *Alice* (or patent law more broadly).

FIGURE 8: AGGREGATE INVESTMENT IN "BUSINESS AND PRODUCTIVITY SOFTWARE"

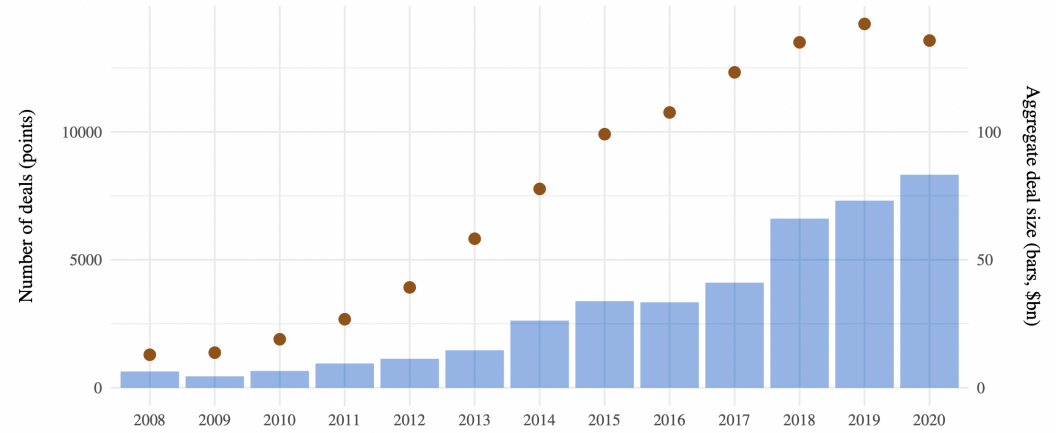

Note. Orange dots reflect the number of deals, and blue bars reflect the total invested capital for that year.

To check this, Table 8 shows the distribution of the highest stage of financing that was reached before each company's first patent application was decided. The Table reveals that amongst startups who received investment before their patent application was decided, around 35%–40% received only pre-seed or seed funding (this includes angel investors, incubators, crowdfunding, and similar sources), and a further 30% had received a Series A funding round. Crucially, there is no apparent difference between those who had their applications decided before *Alice* and those who had their applications decided after.

TABLE 8: HIGHEST ROUND OF FUNDING RAISED BEFORE THE FIRST APPLICATION WAS DECIDED (LIMITED TO APPLICANTS WHO RECEIVED PRE-DISPOSITION FUNDING)

	Pre-*Alice*	Post-*Alice*
Pre-Series	34.2%	38.3%
Series A	31.5%	27.9%
Series B	11.7%	12.9%
Series C	3.2%	5.2%
Later Stage	12.6%	9.8%
Private Equity	3.6%	4.0%
Missing Data	3.2%	1.8%

C. Were the Applicants Who Received Their Decision After Alice *Different?*

Finally, we might be concerned that the results are driven by some kind of selection effect amongst applicants.[150] Although patent applicants cannot control the PTO's examination schedule with any precision, they are in control of when and whether to abandon their application.[151] For example, if more sophisticated applicants (and investors) were aware of the Court's jurisprudence, they might have behaved differently, perhaps by abandoning their applications more quickly.

One useful way to assess any potential selection problem is to consider a covariate whose value was determined *before Alice*. The decision could not affect such a variable, so there should be no discontinuous change in its average value across the threshold. In this context, a particularly interesting covariate to check is whether applicants received any pre-*Alice* funding because applicants who already have VC backing (and therefore potential access to professional and legal advice) might be expected to have responded differently to the PTO's heightened rejection standards. Figure 9 shows the rate of pre-*Alice* funding, plotted against the date of disposition of the applicant's first application. If this value changed sharply after *Alice*, it would indicate that applicants with prior investment were behaving differently in response to the new PTO review standards. But in fact, there is no evidence of a change at the *Alice* threshold, lending support to the idea that applicants are similar in this characteristic regardless of when their applications were actually decided. In a similar vein, the average amount of pre-disposition funding was $10.6 million for applications decided before *Alice* versus $10.4 million for those decided after, while the average number of pre-disposition funding rounds were 1.98 and 2.29, respectively.

[150] In addition to the robustness checks in this Section, I include the results of a formal density test for sorting in Appendix B.

[151] For further elaboration of an applicant's role in the patent examination process, see *supra* note 102 and accompanying text.

FIGURE 9: SHARE OF COMPANIES WHO RECEIVED AT LEAST ONE FUNDING ROUND BEFORE *ALICE*, GROUPED BY TWO-WEEK WINDOWS

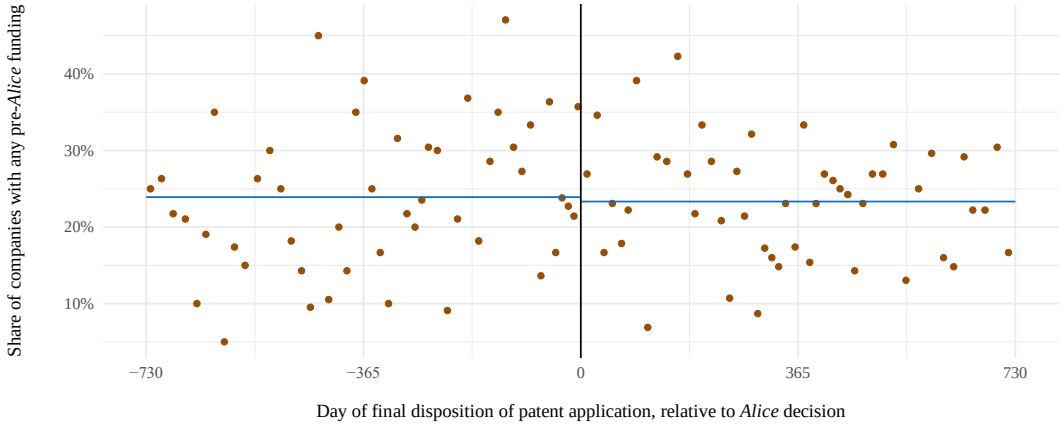

VI. IMPLICATIONS AND SOME CAVEATS

In summary, I find no evidence that patent approval affects either venture capital investment or subsequent acquisition activity for business-methods patent applicants. Between the applicants who received decisions before and after *Alice*, there is no difference in the likelihood or amount of investment, nor in the likelihood of a subsequent acquisition or a successful exit event. In this Part, I provide some context for these findings. I begin by noting several limitations to the empirical design and then go on to explore the implications for the ongoing policy debate over what kinds of invention should be eligible to receive a patent.

A. Limitations

There are a handful of important limitations to the results. My empirical design allows me to disentangle the distinct effect of a patent in a credible way—put differently, the design has strong "internal validity." But what I gain in specificity, I lose in generality.

First, my results apply specifically to the group of companies whose application outcomes were different as a result of *Alice*. This may appear to be a trivial observation, but given that the pre-*Alice* allowance rate did not exceed 50% in the affected art units, there clearly exist a large number of business-methods applicants who would not have been granted a patent under any PTO regime. Fortunately, the applicants whose counterfactual outcome was "switched" by *Alice* are conceptually the most interesting,

because theirs are the inventions at the borderlands of patentability, where we might expect the patent signal to provide the most information.[152]

Second, I'm addressing a specific question: what is the role of the *first* patent? This allows me to estimate the effect on funding of holding *some* intellectual property versus holding *none*, but it does not necessarily tell us what role IP plays for companies that build larger patent portfolios. However, my focus here is on startups. Companies that have built large patent portfolios are far less likely to be early-stage startups and are also less likely to be seeking venture capital investment, particularly in the software industry. Software startups, for the most part, simply don't have many patents.[153] In other words, portfolios present an interesting, but quite different, question about the role of patents in the software industry, and the empirical design used here is not well suited to answer it.

B. What Next for Patentable Subject Matter?

How should patentable subject matter adjust in light of these findings? This is a particularly timely question, as Congress is again considering a legislative fix to § 101 of the patent statute.[154] Software is once more a central issue in this debate, given its role in the Court's recent abstract-ideas jurisprudence. And business methods, in particular, seem to have become the proverbial "tail wagging the dog" of eligibility doctrine. Further, a rich literature argues that the Court (or Congress) should restore the implicit business-methods exception to patent eligibility that prevailed before the Federal Circuit's decision in *State Street*.[155] The findings in this Article cast doubt on one of the few plausible justifications for eligibility, and in that respect lend empirical support to those calling for a renewed statutory bar.

---

[152] Indeed, Farre-Mensa and his collaborators find that the switchers in their study come from the "middle quintile of the distributions of sales, employment, sales growth, and employment growth." Farre-Mensa et al., *supra* note 85, at 665. It would be interesting to investigate in more granular detail those applicants who still received a patent after *Alice*. Given how few applications were successful in that period, one might think the receipt of a patent would send a stronger signal about the underlying invention than it would have pre-*Alice*. For now, I reserve this question for later work.

[153] In the sample of startup companies studied here, the mean number of applications per firm is 1.79, and the median is 1.00. Relatedly, one might be concerned that applicants with post-*Alice* rejections would file continuations—effectively sending a signal to investors that the prosecution process was ongoing. However, in the Appendix, I show that rates of continuation filings are very low amongst applicants in this sample. *See infra* Appendix C.

[154] Quinn & McDermott, *supra* note 59.

[155] *See, e.g.*, Michael J. Meurer, Bilski *and the Information Age a Decade Later*, *in* THE CAMBRIDGE HANDBOOK OF INVESTMENT-DRIVEN INTELLECTUAL PROPERTY 114, 131–32 (Enrico Bonadio & Patrick Goold eds., 2023) ("The case for categorical exclusion of business method patents is stronger today than it was in 2010."); Meurer, *supra* note 44, at 310; David S. Olson, *Taking the Utilitarian Basis for Patent Law Seriously: The Case for Restricting Patentable Subject Matter*, 82 TEMP. L. REV. 181, 240 (2009).

If it were the case that *Alice had* depressed investment in business-methods software, there would be good reason to be circumspect about doing away with patents in the field. But the absence of any such evidence weakens the case for patents in this area, which is already theoretically tenuous—it has long been unclear whether patents are a good fit for software in general.[156] The relatively low cost of research and rapid pace of development in the industry are at odds with the expensive, protracted patent application process.[157] Meanwhile, the incentive to innovate is unlikely to be diminished by competitive imitations when complementary economic tools such as network effects and first-mover advantage can confer significant market power.[158] Indeed, in some respects, other forms of intellectual property protection seem to be a closer fit to the economics of the industry. For example, as modern software continues to coalesce around artificial intelligence and software-as-a-service cloud computing, trade secrecy may offer an increasingly plausible alternative avenue for formal protection.

Still, if business-methods patents have little apparent value, it's reasonable to wonder whether patent eligibility really matters in this area. If software investors don't care about patents because they are not central to most companies' business models, then one might suppose those patents will be economically insignificant. The problem is that there are social costs to issued patents, regardless of their quality or private value. For example, patents can be used by incumbents to build large defensive portfolios or by nonpracticing entities to extract rents—and there is evidence of both in the software industry.[159] They might also stifle competitive entry.[160] Indeed, venture capital companies report that they are unlikely to invest in a startup that has an existing patent demand against it.[161] And the fuzzy boundaries of many business-methods patents mean that would-be infringers are often not on meaningful notice of the exclusive right.[162] In the best of circumstances, patents exact a tax on innovation: the question in any given industry is

---

[156] *See supra* note 12 and accompanying text.

[157] Pamela Samuelson & Jason Schultz, *"Clues" for Determining Whether Business and Service Innovations Are Unpatentable Abstract Ideas*, *in* PERSPECTIVES ON PATENTABLE SUBJECT MATTER, *supra* note 12, at 8, 21–23.

[158] *See, e.g.*, Dreyfuss, *supra* note 44, at 275 (arguing that business-method patents may not be needed to incentivize innovation).

[159] *See* Stuart J.H. Graham & Ted Sichelman, *Why Do Start-Ups Patent?*, 23 BERKELEY TECH. L.J. 1063, 1080–81 (2008); Cockburn & MacGarvie, *supra* note 63, at 769 ("[P]atent thickets appear to offer significant protection to incumbent firms . . . ."); Bessen & Meurer, *supra* note 17, at 390.

[160] Hall, *supra* note 12, at 269.

[161] Feldman, *supra* note 77, at 280.

[162] *See* Peter S. Menell & Michael J. Meurer, *Notice Failures and Notice Externalities*, 5 J. LEGAL ANALYSIS 1, 33 (2013).

whether the benefits surpass that cost—and there are reasons to think those costs are particularly salient in the business-methods context.

Patent-eligibility law is, of course, not the only tool available to policymakers. Although it is the threshold requirement for patentability, there are several other doctrinal levers that could be used to tune patent policy. For example, rather than categorical exceptions, we might instead apply more rigorous assessments of nonobviousness or novelty to potentially "abstract" inventions, or narrow the scope of patents that are granted.[163] These are sensible suggestions—certainly more robust cabining of the scope of business-methods software claims would be a salutary development.[164] The latter approach in particular could also allow more judicial flexibility in responding to new challenges and new technologies as they arise.[165] On the other hand, to the extent that the PTO has a role in increased policing of obviousness and novelty, there's reason to be skeptical about whether the structure of the examination process or the PTO's resources are designed to produce better outcomes.[166] And as a practical matter, the potential for these patents to be asserted in socially costly, rent-seeking ways seems likely to remain high.[167] In any case, for the moment we are at a legislative crossroads, and the question presented is whether or not to codify a business-methods

---

[163] *See, e.g.*, Kristen Osenga, *Still Aiming at the Wrong Target: A Case for Business Method and Software Patents from a Business Perspective*, *in* PERSPECTIVES ON PATENTABLE SUBJECT MATTER, *supra* note 12, at 29, 42 ("The concern about opportunistic litigation is made much more palpable when the patent in question is weak . . . . [This is] better addressed under the other patentability requirements, namely §§ 102, 103, and 112."); Lemley et al., *supra* note 46, at 1329 ("[T]he abstract ideas doctrine is not about finding a conceptual category of inventions that is entitled to no protection at all . . . . Instead, it is about encouraging cumulative innovation . . . by preventing patentees from claiming broad ownership over fields of exploration . . . .").

An alternative, stronger version of this argument would do away with most eligibility bars altogether, leaving other doctrines to do the work of filtering deserving inventions. *See* Sepehr Shahshahani, Patent Law's Abstract-Ideas Problem (2023) (unpublished manuscript) (on file with author) (arguing that there's no principled case for the categorical exclusion of abstract ideas that are otherwise novel, nonobvious, and useful); Athena Diagnostics, Inc. v. Mayo Collaborative Servs., LLC, 927 F.3d 1333, 1335 (Fed. Cir. 2019) (en banc) (Lourie, J., concurring) ("If I could write on a clean slate, I would write as an exception to patent eligibility, as respects natural laws, only claims directed to the natural law itself . . . .").

[164] In the past, the PTO has attempted several quality control procedures for business-methods patents, including the "second pair of eyes" review (an internal PTO procedure) and Covered Business Method review (an administrative venue for an alleged infringer to challenge a granted patent), though of course the latter focused on invalidating the patent *in toto* rather than clarifying or limiting scope. *See, e.g.*, John R. Allison & Starling D. Hunter, *On the Feasibility of Improving Patent Quality One Technology at a Time: The Case of Business Methods*, 21 BERKELEY TECH. L.J. 729, 734 (2006); Leahy-Smith America Invents Act, Pub. L. No. 112-29, § 18, 125 Stat. 284, 329 (2011).

[165] Dan L. Burk & Mark A. Lemley, *Policy Levers in Patent Law*, 89 VA. L. REV. 1575, 1639 (2003).

[166] *See* Lemley, *supra* note 67, at 1495, 1531.

[167] *See supra* note 159; Meurer, *supra* note 157, at 128–29 (discussing how business-methods patents can be asserted in socially harmful ways).

exception.[168] It's important to be careful when interpreting the empirical results in this study, but in an industry where the benefits of patents are demonstrably negligible and the costs nontrivial, there is little support for the case that business methods should be patent eligible.[169]

Finally, this study's most important caveat is both a limitation and a strength of the empirical design. These findings are relevant to a specific group of patents—business-methods software—and it would be a mistake to extrapolate too far outside this context. There is every reason to think that patents play different roles in different industries, as a variety of the studies discussed in Part II indicate. Indeed, when these results are put next to those from prior studies, important distinctions begin to emerge.[170] Farre-Mensa, Hegde, and Ljungqvist find that the average effect of patent approval on the likelihood of funding is 50%, but when contrasted, the results of the two studies suggest that there is important heterogeneity between industries.[171] Startups in certain sectors of information technology are substantially more likely to receive investment after their first patent is approved (in their study), while in biochemistry (also in their study) and business-methods software (in this one), patent approval appears to have no effect.[172] Given these differences, any "one-size-fits-all" reform proposal is inapposite.[173]

---

[168] Quinn & McDermott, *supra* note 59.

[169] Olson, *supra* note 155, at 234. The advocates of a scope-limiting approach also argue, more provocatively, that categorical exclusions are a "pointless" exercise (or at least unwise) because of the impossibility of drawing clear conceptual lines around categories of technology. Lemley et al., *supra* note 46, at 1317, 1326–27. But it's not clear that this is such a difficult exercise—Congress has legislated with respect to business-methods patents, and the PTO has extensive experience with them. *See supra* text accompanying note 163. Like most technologies, the outer borders of business methods are not crisp, and any line-drawing exercise will inevitably be somewhat under and overinclusive. The question is whether the benefits of exclusion exceed the costs of the "false positives" (i.e., those innovations which would have been incentivized by patent protection, had it been available). *See* Olson, *supra* note 155, at 234–36 (making similar arguments).

[170] *See supra* Section II.C.

[171] Farre-Mensa et al., *supra* note 85.

[172] The gulf between these results and previous ones raises a broader and rather thorny question: how uniform should the patent system be? The findings indicate that, at least with respect to driving investment, patents play quite different roles in different industries. Uniformity is an enormously complex theoretical topic, and a full treatment is outside the scope of this Article. *See, e.g.*, Dan L. Burk & Mark A. Lemley, *Is Patent Law Technology-Specific?*, 17 BERKELEY TECH. L.J. 1155, 1156 (2002) (arguing that, in important ways, patent doctrine is not (or at least was not) technology-neutral in its application, especially as to biotechnology and computer software); Michael Carroll, *Tailoring Intellectual Property Rights to Reduce Uniformity Cost*, *in* RESEARCH HANDBOOK ON THE ECONOMICS OF INTELLECTUAL PROPERTY LAW 377 (Peter Menell, Ben Depoorter & David Schwartz eds., 2019) (discussing the undesirable incentives created by uniformity). Still, the evidence introduced here helps to flesh out this conversation, unearthing evidence of the on-the-ground diversity in patent usage across industries.

[173] Of course, this is not a new observation. *See, e.g.*, John R. Allison & Mark A. Lemley, *Who's Patenting What? An Empirical Exploration of Patent Prosecution*, 53 VAND. L. REV. 2099, 2146–47

Of course, all this is only one piece of the puzzle: we must interpret these findings in concert with other evidence on the role that patents play in the software industry more broadly. Early-stage financing is an important avenue by which patents might facilitate innovative activity, but it is not the only one. For example, while patents appear to play no role in channeling investment for business-methods software, we also need evidence about their role in stimulating new research, facilitating a robust licensing market, and so on.[174] This study provides concrete evidence about a particularly important mechanism, but more research is warranted.[175]

CONCLUSION

To return to my core question: does the grant of a patent make a startup more likely to receive venture capital investment? Using a novel quasi-experimental approach, I provide new evidence on this question. The findings are striking: contrary to previous empirical studies, my results suggest that patents have little impact on either early-stage financing for business-methods software or on subsequent outcomes for the startups involved.

The results here shed light on the role played by patents in a crucial part of the innovation ecosystem and help to inform the ongoing debate over patentable-subject-matter policy. Business-methods software is an area of significant economic activity and investment, but it is also one where the role of patents has been, and continues to be, very controversial. As policymakers seek to repair the law of patent eligibility, these findings provide important evidence on one part of the puzzle.

(2000) ("Objections and proposals for reform that are tailored to the needs of one industry may not fit another well at all.").

[174] Though to the extent that there *is* empirical work on some of these questions, it does not provide much support to the case for software patenting. *See, e.g.*, James Bessen & Robert M. Hunt, *An Empirical Look at Software Patents*, 16 J. ECON. & MGMT. STRATEGY 157 (2007) (finding that software patents are associated with manufacturing firms in industries known for strategic—that is, defensive—patenting); Bronwyn H. Hall & Megan MacGarvie, *The Private Value of Software Patents*, 39 RES. POL'Y 994, 1006 (2010) (finding no change to the valuation of pure software firms after software patents became available); Brian Love, *An Empirical Study of Patent Litigation Timing: Could a Patent Term Reduction Decimate Trolls Without Harming Innovators?*, 161 U. PA. L. REV. 1309, 1347 (2013) ("[M]ore than 80% of NPE-filed suits assert high-tech patents generally, and more than 65% have software-related claims."). *But see* Colleen V. Chien, *Software Patents as a Currency, Not Tax, on Innovation*, 31 BERKELEY TECH. L.J. 1669, 1699 (2016) (arguing that many software patent sales from public companies reflect meaningful technology transfer and exchange).

[175] *See* Golden et al., *supra* note 20, at 1767 ("A new, policy-oriented synthesis can be robust precisely because it is built on many solid, discrete studies.").

APPENDIX

A. Data Cleaning

1. Entity Detection

The focus of this study is on startup companies. Unfortunately, a patent application contains very little metadata about its applicant, which makes it difficult to identify the applicants of interest. So, in order to infer the entity type of each applicant, I use a multistage algorithm based on the applicant's name.

For example, if a name contains words such as "museum," "hospital," "university," "foundation," or "trustee," I classify it as an institution and remove it from the sample.[176] If an applicant's name contains the pattern word-word-letter (in other words: last name-first name-middle initial) or honorifics such as "dr," "esq," or "phd," I classify them as an individual and remove them.

If an applicant's name contains a common corporate suffix, including "corp," "co," "intl," "ind," and "llc," I classify it as a company and retain it. For the remaining applicants, I use a list of common company names to identify any remaining businesses, including terms such as "finance," "agri," "pharma," "travel," "digital" (and many more).

2. Standardization and Disambiguation

Patent application data are primarily applicant-generated. This presents at least two problems. First, the applicants' names take many different forms (for example, "IBM" and "International Business Machines") or contain simple typographical errors, both which make it challenging to consolidate multiple applications by the same applicant. To address this problem, I rely on the disambiguation data provided by PTO PatentsView.[177] These assignments are generated by a series of machine-learning algorithms which cluster applicants based on their names and locations.[178] In a handful of cases where I identify an error in these data, I manually correct it. I use the same disambiguation data when identifying corporate name changes in the assignment data, ensuring that if a company uses multiple names over its lifespan, those names are resolved to a single company identifier.

Second, the company names are not standardized in any way, making it difficult to match applicants to investments. (For example, a patent

[176] A full list of all the matching patterns and standardization routines discussed in this Part is available in the supplemental code.

[177] PATENTSVIEW, *supra* note 107.

[178] NICHOLAS MONATH, CHRISTINA JONES & SARVO MADHAVAN, PATENTSVIEW: DISAMBIGUATING INVENTORS, ASSIGNEES, AND LOCATIONS (2021) (on file with author).

application might be filed for "Ariad Pharmaceuticals," while Pitchbook might report data under "ARIAD PHARM.") In order to make the names comparable, I apply a standardization routine to both patent applicants and Pitchbook target companies. My approach draws on the Stata program "stnd_compname" as well as the routines developed for the Patent Data Project.[179] This routine corrects myriad spelling errors (e.g., "compnay," "scientfc"), removes special characters and corporate identifiers (e.g., "corp," "ltd," "llc"), and standardizes common suffixes (e.g., "dotcom" and "com" become ".com").

3. Matching Applicant Companies with Investments

After standardizing and cleaning the names of both patent assignees and Pitchbook target companies, I use a matching algorithm to pair them together. The algorithm has the following steps:

(1) Where the company names and states are identical in both databases, treat that as a match (this accounts for 75% of matches).
(2) For the remaining non-matched assignees, if the company names are identical but the states are different (or one state is missing), output for manual inspection (23% of matches).
(3) For the remaining non-matched assignees, perform a fuzzy match between the company names using Jaro-Winkler (J-W) distance scores.[180] I review the results of this match by hand for every pair with a J-W score of less than 0.2 (2% of matches).
(4) To catch any instance where Pitchbook has multiple companies with the same name in the same state, manually review any assignee that matches more than one Pitchbook ID and keep only the correct match.

Following this process, I end up with at least one Pitchbook record for each of 1,305 unique patent assignees (though only 624 of these firms had a VC funding round or exit event).

B. Regression Discontinuity Assumptions

As with any regression discontinuity, the empirical design in this project rests on a set of important assumptions. Although the specifics of these assumptions are technical, assessing their credibility relies mainly on substantive knowledge of the area.

---

[179] *See* Nada Wasi & Aaron Flaaen, *Record Linkage Using Stata: Preprocessing, Linking, and Reviewing Utilities*, 15 STATA J. 672, 681 (2015); PATENT DATA PROJECT, https://sites.google.com/site/patentdataproject/Home [https://perma.cc/4VCK-RRNH].

[180] *See* William W. Cohen, Pradeep Ravikumar & Stephen E. Fienberg, *A Comparison of String Metrics for Matching Names and Records*, 3 KDD WORKSHOP ON DATA CLEANING & OBJECT CONSOLIDATION 73 (2003).

First, I assume that in the absence of the *Alice* decision, the outcomes would have crossed the threshold smoothly, with no evidence of a jump (the continuity of potential outcomes assumption). In other words, this assumption means that if *Alice* had not occurred, then there would have been no sharp change in June 2014.

Second, I assume that applicants can't "sort" around the threshold—in other words, that applicants were not able to maneuver to have applications decided in advance of the Court's decision. In practice, the pendency of applications is quite variable, and averages just under three years, leaving little opportunity for applicants to game the process by strategically timing their applications.[181] (Although applicants could potentially have abandoned their applications in anticipation of the result in *Alice*, it isn't clear what incentive they would have to do so.)[182] The graph below tests this more formally, and shows no evidence of sorting.

FIGURE A1

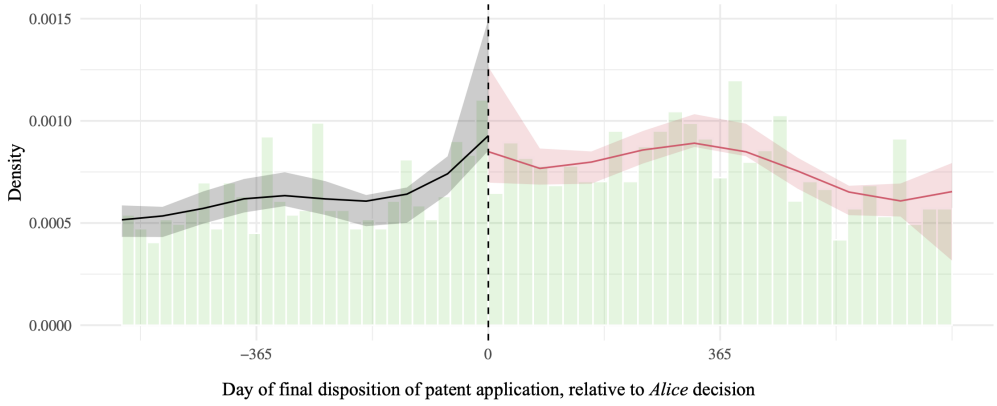

[181] Note also that I limit my sample to applications filed before the *Alice* decision.

[182] It's also not at all clear that entrepreneurs or their investors were paying attention. *See, e.g.*, Taylor, *supra* note 75, at 2077–78 (finding that investors have a generally low level of knowledge about changes in patent doctrine). Furthermore, in a prior case, the Court affirmatively declined to set forth a sweeping rule that might destabilize business-methods patenting. *See* Bilski v. Kappos, 561 U.S. 593, 606–09 (2010). In any case, under the assumption that applicants were aware that the Court was considering an important patentable-subject-matter case and sensed that it was likely to use the case in order to create a particularly unfavorable standard, the date that any Supreme Court decision will be handed down is not known ex ante.

With respect to the "fuzzy RD" estimates in Tables 3 and 6b, I also assume that the treatment affects all applicants in the same direction (the monotonicity assumption). In other words, I assume that no startup who received a patent would be *less* likely to receive funding as a result. This seems, in this context, to be a reasonable assumption.

Finally, I assume that receiving a disposition after *Alice* affects the likelihood of receiving funding only by decreasing the likelihood that the patent will be granted. On this point, we might reasonably suppose that if overall investment activity in business methods was chilled by the decision, then the effect of having an application decided after *Alice* could be to depress both the likelihood that the application was allowed and the likelihood of receiving funding. If this were the case, we would be unable to separately estimate the effect of the patent approval. However, there is reason for optimism here. As show in Parts IV and V of the main Article, I find no evidence that overall investment did, in fact, decline in response to *Alice*.[183]

C. Continuations

One esoteric feature of patent prosecution, known as continuations, might lead us to question how many of the applications in this sample were *truly* abandoned.[184] The problem is as follows. In the face of a rejection—or indeed for many other reasons—applicants have the opportunity to file various kinds of "continuing" applications. A continuing application is a new application which is in some way a descendant of an existing application.[185] They take different forms, including continuations (where the specification is the same but new claims are added), continuations-in-part (where the specification is similar, but with some added subject matter), and divisionals (where a narrower invention is carved out of a broader parent).[186] In different ways, each provides an opportunity for applicants to try to "write around" examiner rejections. Given the widespread rejections on the basis of patentable subject matter after *Alice*, it's reasonable to ask whether many of the apparently abandoned applications were in fact supplanted by continuations with slightly different claim language, but similar inventive

[183] I explore this assumption in more detail in Section V.B.

[184] I owe particular thanks to Hal Edgar for his thoughtful conversations about this issue.

[185] 37 C.F.R. § 1.53.

[186] *Id.* Note that continuations are distinct from a "request for continued examination" (RCE), which is simply an administrative procedure to reopen the *same* (rejected) application. U.S. PATENT & TRADEMARK OFF., *supra* note 105, § 706.07(h). For the purposes of this study, I treat RCEs as part of the same process as the original application.

subject matter. If this were the case, it might be inappropriate to treat the unsuccessful group as "failed" applicants.

The PTO provides a continuity dataset which links continuation applications to their "parent" application.[187] Of the 1,745 abandoned applications in my sample, 316 (18%) have at least one continuation. If we look only at applications with a post-*Alice* disposition date, the figure is 213 out of 1,149 (18.5%), which suggests that there was little change in the rate of continuation filings in response to *Alice*. Relatedly, across the sample, only half of the continuations were filed *after* disposition of the original application. In other words, only 9% of the applications had continuations that were even potentially in response to a final rejection. In short: whether because of the perceived low value of these patents, or the challenge of evading examiner rejections in business-methods art units, there's little in the way of continuation practice in this area.

[187] *See* Miller, *supra* note 104, at 36–38.