Does the U.S. Patent & Trademark Office Grant Too Many Bad Patents?: Evidence from a Quasi-Experiment

67 STAN. L. REV. ___ (forthcoming 2015)

Michael D. Frakes[*]
Melissa F. Wasserman[**]

Many believe the root cause of the patent system's dysfunction is that the U.S. Patent & Trademark Office (PTO or Agency) is issuing too many invalid patents that unnecessarily drain consumer welfare. Concerns regarding the Agency's over-granting tendencies have recently spurred the Supreme Court to take a renewed interest in substantive patent law and have driven Congress to enact the first major patent reform act in over sixty years. Policymakers, however, have been modifying the system in an effort to increase patent quality in the dark. As there exists little to no compelling empirical evidence the PTO is actually over-granting patents, lawmakers are left trying to fix the patent system without even understanding the root causes of the system's shortcomings.

This Article begins to rectify this deficiency, advancing the conversation along two dimensions. First, it provides a novel theoretical source for a granting bias on the part of the Agency, positing that the inability of the PTO to finally reject a patent application may create an incentive for a resource-constrained Agency to allow additional patents. Second, this Article attempts to explore, through a sophisticated natural-experiment framework, whether the Agency is in fact acting on this incentive and over-granting patents. Our findings suggest that the PTO is biased towards allowing patents. Moreover, our results suggest the PTO is targeting its over-granting tendencies towards those patents from which it stands to benefit the most through

[*] Associate Professor of Law, Northwestern University School of Law; Faculty Research Fellow, National Bureau of Economic Research; Faculty Fellow, Institute for Policy Research, Northwestern University. J.D., Ph.D.

[**] Associate Professor of Law; Richard and Anne Stockton Faculty Scholar; & Richard W. and Marie L. Corman Scholar, University of Illinois College of Law. J.D., Ph.D.

The work was funded in part by University of Illinois at Urbana Champaign Research Board, Award 12088 and the Cornell Institute for the Social Sciences Small Grant Award. We thank Michael Abramowicz, Rochelle Dreyfuss, Stuart Graham (former Chief Economist at the PTO), Paul Heald, Mark Lemley, Oskar Liivak, Alan Marco (Acting Chief Economist at the PTO), Peter Menell, W. Nicholson Price II, Michael Risch, David Schwartz, Saurabh Vishnubhakat and seminar participants at Chicago-Kent College of Law Workshop on Empirical Methods in Intellectual Property, 2013 American Law and Economics Association Annual Meeting, faculty workshop at the Notre Dame School of Law, Michigan Intellectual Property Workshop, and DePaul Law School Manzo Patent Law Seminar for useful comments.

allowing. Our findings provide policymakers with much needed evidence that the PTO is indeed over granting patents. Our results also suggest that the literature has overlooked a substantial source of Agency bias and hence recent fixes to improve patent quality will not achieve their desired outcome of extinguishing the PTO's over-granting proclivities.

INTRODUCTION

In recent years the patent system has come under trenchant criticism. There is a general consensus that the system is both "broken" and a "failure."[1] Many believe the root cause of the patent system's dysfunction is that the United States Patent & Trademark Office (PTO or Agency) is issuing too many invalid patents that unnecessarily drain consumer welfare, stunt productive research, and unreasonably extract rents from innovators.[2] Such sentiments have even been the subject of multiple reports by the National Academies and the Federal Trade Commission.[3] Concerns regarding the Agency's over-granting tendencies have been so pressing that they have spurred the Supreme Court to take a renewed interest in substantive

---

[1] James Bessen & Michael J. Meurer, PATENT FAILURE: HOW JUDGES, BUREAUCRATS, AND LAWYERS PUT INNOVATORS AT RISK (2008); Dan L. Burk & Mark A. Lemley, THE PATENT CRISIS AND HOW THE COURTS CAN SOLVE IT (2009); B. JAFFE & JOSH LERNER, INNOVATION AND ITS DISCONTENTS (2004).

[2] *See infra* notes 13, 18-22 and accompanying text.

[3] *See, e.g.*, FEDERAL TRADE COMMISSION, TO PROMOTE INNOVATION: THE PROPER BALANCE OF COMPETITION AND PATENT LAW AND POLICY (2003), *available at* http:///www.ftc.gov/os/2003/10/innovationrpt.pdf; NAT'L RESEARCH COUNCIL, A PATENT SYSTEM FOR THE 21ST CENTURY (Stephen A. Merrill et al. eds., 2004); NAT'L RESEARCH COUNCIL, REAPING THE BENEFITS OF GENOMIC AND PROTEOMIC RESEARCH: INTELLECTUAL PROPERTY RIGHTS, INNOVATION AND PUBLIC HEALTH (Stephen A. Merril & Anne-Marie Mazza eds., 2006).

patent law,[4] while likewise driving Congress to enact the first major patent reform act in over sixty years.[5]

Despite the centrality of patent quality to the health of the patent system and the plethora of reasons put forward as to why the PTO may be biased towards allowing patents, there exists little to no compelling empirical evidence that the Agency is *actually* over-granting patents.[6] To date, the debate regarding patent policy has been driven by anecdotes of a few infamously issued patents or data that alone is insufficient to deduce that the PTO is overly permissive in its granting decisions.[7] As a result, policymakers have been modifying the system in an effort to increase patent quality in the dark. Without guidance as to which features of the system may actually be causing the PTO to over-grant patents or even compelling evidence that the Agency is biased towards allowing patents, policymakers are left trying to fix the patent system without even understanding the root causes of the system's shortcomings.

This Article begins to rectify this deficiency and advances the conversation along two dimensions. First, it provides a novel theoretical source for why the Agency may allow too many invalid patents. This newly addressed source of bias stems from an oddity of the U.S. patent system—the inability of the PTO to finally reject a patent application. The capacity of aggrieved patent applicants to continuously restart the examination process upon rejection by filing repeat applications can potentially overwhelm the existing examination infrastructure. The fact that the Agency currently faces crushing backlog of over 600,000 patent applications of which close to 40% constitute repeat filings suggests this has already occurred. The PTO can effectively turn off the spigot of repeat filings by allowing—i.e., granting—patents earlier on in the examination process.[8] That is, the PTO could attempt to decrease the incentives of applicants to file repeat applications (and hence concomitantly decrease its backlog of applications) by biasing its grant rate upwards.

To be clear, repeat filings has been the subject of a growing academic discourse.[9] These debates, however, have mostly focused on applicant behavior rather than agency incentives. This Article aims to build on and fill various gaps in the literature by refocusing the debate on the PTO's inducements.

---

[4] Mark A. Lemley & Bhaven Sampat, *Is the Patent Office a Rubber Stamp?*, 58 EMORY L.J. 181, 185 (2008).

[5] *Id.* For a discussion of the recent major changes to the patent system enacted by the judiciary and by the legislature—in each case in an effort to combat the Agency's perceived penchant for allowing patents—see infra notes 26-27 and accompanying text.

[6] *See infra* Part II.B.

[7] *Id.*

[8] As a bonus, by inflating its granting tendencies in this matter, the PTO will also generate additional revenue from post-allowance fees—fees that the PTO collects *only when it grants a patent*.

[9] *See, e.g.*, Mark A. Lemley & Kimberly A. Moore, *Ending Abuse of Patent Continuations*, 84 B.U. L. REV. 63, 63 (2004) (detailing the applicant misuse of continuation applications).

Second, this Article attempts to explore, through a sophisticated methodological design, whether the Agency is in fact operating on this incentive and actually over-granting patents. We begin this exercise by setting forth a theoretical model that predicts that the PTO will only inflate its granting tendencies in order to diminish repeat filings during times in which the PTO faces binding resource constraints—i.e., during times in which it cannot meet its expected examination demand through available resources. To empirically test whether a resource-constrained PTO will grant additional patents, we amassed a rich database of previously unavailable patent data through the filing of Freedom of Information Act Requests to the PTO. This novel database includes patent processing statistics on all 4,733,263 patent applications filed before the PTO over an approximately twenty-year period.[10] We then employ a natural-experiment framework, wherein we compare the Agency's grant rate before and after a period of time in which the PTO experienced deterioration in its net resources. Of course, grant rates may change over time for a number of reasons unrelated to the story at hand—e.g., the quality of the underlying applications may change from year to year. As such, we do not rely solely on this before-after comparison as a mechanism to isolate the effect of repeat filings on its grant rates. Rather, we explore how a negative resource "shock" to the Agency differentially affects a treatment group and a control group, much like a randomized medical experiment would. For the purposes of this design, our treatment group is represented by those technologies that have historically exhibited a stronger disposition to partake in the costly activity of interest in this Article—that is, a stronger disposition towards filing repeat applications. It is with respect to these applications that the PTO would be especially inclined to act on the above-specified incentive. Likewise, our control group is captured by those technologies that have not historically exhibited such a tendency and with respect to which the PTO may be less inclined to inflate its grant rate upon experiencing a downturn in its resources that leaves the Agency unable to process all of the patent applications that are awaiting review.

By netting out the granting experiences of the control group over this relevant period of time, this methodology allows us to hone in on the effects on the Agency's grant rate stemming from the theorized incentive itself. In other words, this "quasi"-experimental framework (supplemented by a number of robustness exercises) enables us to better make *causal* inferences—that is, to more convincingly trace the Agency's observed increase in grant rates to the potentially biasing feature itself and thus to rule out potentially confounding explanations for the documented grant-rate patterns.

We find evidence that the PTO is indeed over-granting patents during times in which the Agency lacks sufficient resources to meet its

[10] Details on the database are provided in Part III.

expected demand of examination. Moreover, our findings suggest that the PTO is preferentially granting patents with respect to which it stands to benefit the most through granting—those in high-continuation-filing-rate technologies.

What are we to make of the evidence that the PTO is granting patents in an effort to diminish repeat filings? On the one hand it suggests that policymakers are generally correct in suggesting that the Agency is allowing too many bad patents. Importantly, however, it also suggests that policymakers have been barking up the wrong tree. By ignoring the import that repeat filings play in influencing PTO decision making, legal scholarship has overlooked a substantial source of Agency bias. As a result, the recent "fixes" to the patent system to improve patent quality fail to address the substantial over-granting bias identified in this Article.[11] Additionally, and equally as worrisome, policymakers have also failed to appreciate that the PTO may not just be biased towards allowing patents in general but instead biased towards allowing particular types of patents. Accordingly, the harms recognized with the granting of bad patents have been under inclusive. To the extent that the PTO is extending preferential treatment to technologies based on backlog incentives and not based on legitimate social interests in intervening in certain industries, the Agency may also be undesirably distorting the allocation of resources across different sectors of the economy.

The rest of the Article is organized as follows. Part I briefly describes the primary complaint that scholars and stakeholders have registered against the patent system: the PTO is issuing too many bad patents. This Part also describes how the literature to date has failed to produce any persuasive empirical evidence demonstrating that the Agency is in fact biased towards allowing patents. Part II theorizes the PTO's possible incentive to grant patents, introducing a model of agency behavior under resource constraints, which predicts that the PTO will act on the incentive to grant additional patents when it faces binding resource constraints—i.e., the Agency cannot meet its expected examination demand through available resources. Part III refines this model of agency behavior by examining how the PTO's incentive to grant additional patents to curtail repeat filings likely varies across patent types. Part III culminates with the delineation of several testable hypotheses that will guide our empirical analysis. Part IV describes the data set and the methodology employed to test our hypotheses. The results of our empirical analysis are presented in Part V. In Part VI, we explore the mechanisms of action that may underlie our primary results. Finally, Part VII evaluates the implications of our results and assesses potential methods to reduce the PTO's tendency towards issuing patents.

---

11 *See infra* Part I.A.

I. HARMS WITH PTO OVER-GRANTING AND INSUFFICIENCY OF EVIDENCE

A. Over-Granting Explanations and Harms

There is near universal agreement that the United States patent system suffers from rampant shortcomings and failures.[12] Many believe the root cause of the system's malfunction is that the United States Patent & Trademark Office (PTO) is allowing too many invalid patents to issue.[13]

The PTO's perceived penchant for issuing invalid patents has been attributed to a multitude of causes. Some contend that the Agency suffers from poor management, while others argue it is hamstrung by the strength of its patent examiner union. Both have been blamed for the structure of the PTO's examiner compensation system which favors patent grants over denials.[14] Others have argued the PTO's dearth of resources precludes it from spending sufficient time inspecting patent applications.[15] Still others contend the Agency's lack of expertise, especially in emerging fields of science, is the root cause of its over-granting proclivities.[16] Finally, some have

---

[12] JAMES BESSEN & MICHAEL J. MEURER, PATENT FAILURE: HOW JUDGES, BUREAUCRATS, AND LAWYERS PUT INNOVATORS AT RISK (2008); DAN L. BURK & MARK A. LEMLEY, THE PATENT CRISIS AND HOW THE COURTS CAN SOLVE IT (2009); B. JAFFE & JOSH LERNER, INNOVATION AND ITS DISCONTENTS (2004).

[13] *See, e.g.*, Mark A. Lemley, *Rational Ignorance at the Patent Office*, 95 NW. U. L. REV. 1495, 1495 n.1 (2001); Dough Lichtman & Mark A. Lemley, *Rethinking Patent Law's Presumption of Validity*, 60 STAN. L. REV. 45, 47 n.5 (2007); FEDERAL TRADE COMMISSION, TO PROMOTE INNOVATION: THE PROPER BALANCE OF COMPETITION AND PATENT LAW AND POLICY (2003), *available at* http:///www.ftc.gov/os/2003/10/innovationrpt.pdf; John R. Thomas, *Collusion and Collective Action in the Patent System: A Proposal for Patent Bounties* 2001 U. ILL. L. REV. 305, 316-22; 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, 589-91 (1999); Mark A. Lemley & Bhaven Sampat, *Is the Patent Office a Rubber Stamp?*, 58 EMORY L.J. 181, 181 (2008); R. Polk Wagner, *Understanding Patent-Quality Mechanisms*, 157 U. PA. L. REV. 2135, 2139-45 (2009); Roger Allan Ford, *Patent Invalidity versus Noninfringement*, 99 CORNELL L. REV. 71, 87 (2013); B. JAFFE & JOSH LERNER, INNOVATION AND ITS DISCONTENTS (2004); JAMES BESSEN & MICHAEL J. MEURER, PATENT FAILURE: HOW JUDGES, BUREAUCRATS, AND LAWYERS PUT INNOVATORS AT RISK (2008).

[14] THOMAS H. STANTON ET AL., NAT'L ACAD. OF PUBLIC ADMIN., U.S. PATENT & TRADEMARK OFFICE: TRANSFORMING TO MEET THE CHALLENGES OF THE 21ST CENTURY 102 (2005) (noting that the productivity schedule is "highly biased toward early allowances"); Clarisa Long, *The PTO and the Market for Influence in Patent Law*, 157 U. PA. L. REV. 1965, 1991 (2009) ("Internal PTO practices create a bias in favor of granting patents."); 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, 607 (1999) ("Consequently, the only way to earn bonus points with confidence is to allow a patent application."); John R. Thomas, *Collusion and Collective Action in the Patent System: A Proposal for Patent Bounties* 2001 U. ILL. L. REV. 305, 324-25.

[15] B. JAFFE & JOSH LERNER, INNOVATION AND ITS DISCONTENTS 130-33 (2004) (describing the PTO's budgetary woes); Mark A. Lemley, *Rational Ignorance at the Patent Office*, 95 NW. U. L. REV. 1495, 1500 (2001) (noting that examiners spend on average only 18 hours reviewing a patent application).

[16] Jeffrey R. Kuester, *Risks Associated with Restricting Business Method and E-Commerce Patents*, 17 GA. ST. U. L. REV. 657, 581 (2001) (noting "[e]very new technology

assigned responsibility to the Agency's culture noting that the PTO once famously asserted that its "primary mission" is "to help customers get patents."[17]

Irrespective of the exact cause, invalidly issued patents impose significant costs on society. Such patents can be utilized by nonpracticing entities or patent trolls to opportunistically extract licensing fees from innovators;[18] they can also impede competitors from entering markets[19] and stunt further innovation.[20] Erroneously allowed patents can inhibit the ability of start-ups to obtain venture capture, especially if the patent in question is held by a dominant player in the market.[21] Finally, such patents can also compromise business relations of market entrants, as customers may be deterred from transacting with a company out of fear of a contributory patent infringement suit.[22]

More generally, invalid patents can result in supra-competitive pricing and diminished quantity. While society may be willing to accept these consequences when the patent is properly issued; an invalid patent imposes these costs upon society without providing the

---

presents the PTO with challenges of creating a sufficient prior art database and channeling the expertise necessary to evaluate the prior art. Internet business methods patents are similar, in this respect, to biotechnology and software"); Kevin M. Baird, *Business Method Patents: Chaos at the USPTO or Business as Usual*, 2001 U. ILL J.L. TECH & POL'Y 347, 364 ("The lack of prior art references and examiner training has led to the issuance of many invalid business method patents resulting in more patent litigation and greater uncertainty in the patent system."); Lucas Osborn, *Tax Strategy Patents: Why the Tax Community Should not Exclude the Patent System*, 18 ALB. L.J. SCI. & TECH 325, 357 ("Mirroring the criticism of business method patents generally, critics assert that patent examiners lack training and expertise to analyze tax law-intensive patent applications.").

17 PATENT & TRADEMARK OFFICE, U.S. DEP'T OF COMMERCE, FY2001 CORPORATE PLAN 23 (2001), *available at* http://www.uspto.gov/web/offices/com/corpplan/pt04.pdf.

18 James Bessen, *The patent troll crisis is really a software patent crisis*, WASH. POST. Sept. 3 2013.

19 *See* FTC, TO PROMOTE INNOVATION 14 (2003), Executive Summary, at 3, *available at* http://www.ftc.gov/sites/default/files/documents/reports/promote-innovation-proper-balance-competition-and-patent-law-and-policy/innovationrpt.pdf (noting that allowing patents on obvious inventions can thrawt competition); Christopher R. Leslie, *The Anticompetitive Effects of Unenforced Invalid Patents*, 91 MINN. L. REV. 101, 119-24 (2008).

20 Michael A. Heller & Rebecca S. Eisenberg, *Can Patents Deter Innovation? The Anticommons in Biomedical Research*, 280 SCI. 698 (1998); Suzanne Schotchmer, *Standing on the Shoulders of Giants: Cumulative Research and the Patent Law*, 5 J. ECON. PERSPECTIVE 29, 32 (1991) (noting that overly broad patent protection "can lead to deficient incentives to develop second generation products").

21 *See* FTC, TO PROMOTE INNOVATION 14 (2003), ch. 2, at 8, *available at* http://www.ftc.gov/sites/default/files/documents/reports/promote-innovation-proper-balance-competition-and-patent-law-and-policy/innovationrpt.pdf ("The threat of being sued for infringement by an incumbent [patent holder]—even on a meritless claim—may scare . . . away venture capital financing." (internal quotation marks omitted)).

22 Christopher R. Leslie, *The Anticompetitive Effects of Unenforced Invalid Patents*, 91 MINN. L. REV. 101, 125-26 (2008). Dow Chem. Co. v. Exxon Corp., 139 F.3d 1470, 1472, 1477 (Fed. Cir. 1998), discussed in In re Ciprofloxacin Hydrochloride Antitrust Litig., 363 F. Supp. 2d 514, 544 (E.D.N.Y. 2005) ("Dow alleged that Exxon had threatened to sue actual and prospective Dow customers for patent infringement, even though Exxon allegedly had no good-faith belief that Dow infringed the patent when Exxon made the threats and had allegedly obtained the patent by inequitable conduct.").

commensurate benefits.[23] That is, a PTO that is applying the patentability standards in a patent-protective manner is likely to be routinely granting patents on inventions that were either already known or represent only a trivial advancement over the existing scientific knowledge.[24] As a result, a grant-biased PTO is likely to systematically issue patents that fail to provide any innovation benefit.[25]

The quality of issued patents has become such an important and visible issue that both Congress and the judiciary have attempted to rectify the Agency's perceived over-granting tendencies. The Supreme Court has recently taken a renewed interest in substantive patent law, wherein, among other things, it has strengthened the doctrine of nonobviousness in an effort to make it easier for the Agency to reject invalid patents.[26] Just last year Congress enacted the first major patent reform bill in over six decades. The Agency was granted new adjudicatory authorities and the ability to set its own fees in an effort to improve patent quality.[27]

B. In Search of Empirical Evidence

Despite the fact that major changes to the patent system are being driven by concerns that the Agency is allowing too many invalid patents to issue, there exists little to no compelling empirical evidence that any particular feature of the system is, in part, driving the PTO to over grant. In fact, the literature generally fails to convincingly establish that the PTO is even biased towards allowing patents more generally.

Consider for instance the frequently cited statistic that nearly half of all litigated patents that make it to a final judgment are

---

[23] WILLIAM D. NORDHAUS, INVENTION, GROWTH, AND WELFARE: A THEORETICAL TREATMENT OF TECHNOLOGICAL CHANGE 76 (1969); Ian Ayres & Gideon Parchomovsky, *Tradable Patent Rights*, 60 STAN. L. REV. 863, 867 (2000); Keith Leffler and Christofer Leffler, *Efficiency Trade-offs in Patent Litigation Settlements: Analysis Gone Astray?*, 39 U.S.F. L. REV. 33, 33 (2004).

[24] Not surprisingly, the patentability standards reflect a careful balance between encouraging innovation and drains on consumer welfare. In order for an invention to be patent eligible it must be both new and represent a non trivial advancement over current scientific understanding. 35 U.S.C. § 103 (2010). If an invention was obvious to the person of ordinary skill in the art or was already in the public domain, the invention would have likely arisen without the patent incentive. In contrast, an invention that represents a significant advancement in the art may not have arisen but for the patent inducement.

[25] Mark A. Lemley & Carl Shapiro, *Probabilistic Patents*, 19 J. ECON. PERSPECTIVES 2 (2005).

[26] KSR Int'l Co. v. Teleflex Inc., 550 U.S. 398 (2007).

[27] America Invents Act, § 10, at 316-20 (to be codified at 35 U.S.C. § 41) (fee-setting authority); § 6, at 299-313 (to be codified at 35 U.S.C. §§ 311-319, 321-329) (post-grant review proceedings), § 12, at 325-327 (to be codified at 35 U.S.C. § 257) (supplemental examination), § 18, at 329 (to be codified at 35 U.S.C. § 321) (Transitional program for covered business-method patents); H.R. REP. No. 112-98, pt. 1, at 39-40 (noting that the primary purpose of the America Invents Act is to "prove patent quality").

invalidated by courts.[28] For this statistic to demonstrate that the PTO is allowing too many invalid patents to issue, litigated patents must not differ substantially from non-litigated patents. Yet, we know that litigated patents are a highly select sample of patents whose characteristics vary substantially from allowed patents in general.[29]

Long term trends in the number of patents issued, another oft-cited statistic, also fail to compellingly establish that the PTO is over-granting patents.[30] It is possible that the Agency is issuing more patents because it has developed a bias towards over granting but it is also conceivable that the number of issued patents has increased because the PTO's examination capacity has grown (which it has dramatically) or that the quality of patent applications has increased.[31] The same holds for the PTO's grant rate.[32] The Agency's allowance rate by itself simply tells us nothing about how well the PTO is scrutinizing patent applications.[33] A 90% grant rate is not too high if almost all patent applications filed are meritorious whereas a 10% grant rate is not too low if none merit allowance.

Perhaps the study that has come closest to providing compelling evidence that the PTO is over granting patents (beyond our prior work discussed below) is the groundbreaking research by Mark Lemley and Bhaven Sampat. Lemley and Sampat explore how examiner experience affects the outcomes at the PTO for nearly 10,000 patent applications filed in January 2001.[34] They find that more experienced examiners have a higher grant rate than their junior colleagues and that the examiner grant rate monotonically increases

---

[28] *See, e.g.*, Roger Allan Ford, *Patent Invalidity Versus Noninfringement*, 99 CORNELL L. REV. 71, 88-89 (2013) (suggesting because half of all litigated patents that make it to final judgment are invalidated the PTO is issuing too many invalid patents); Christopher R. Leslie, *The Anticompetitive Effects of Unenforced Invalid Patents*, 91 MINN. L. REV. 101, 105-06 (same); Michael A. Carrier, *Post-Grant Opposition: A Proposal and a Comparison to the America Invents Act*, 45 U.C. DAVIS L. REV. 103, 108-09 (2011) (same); Christopher R. Leslie, *Patents of Damocles*, 83 IND. L.J. 133, 136-37 (same); Tun-Jen Chiang, *The Upside-Down Inequitable Conduct Defense*, 107 NW. U.L. REV. 1243, 1248-49 (2013) (same).

[29] For instance, litigated patents have more claims, spend more time in review at the PTO, cite more prior art and come disproportionately from certain industries than non-litigated patents. John R. Allison, et al., *Valuable Patents*, 92 GEO. L.J. 435, 438-39 (2004).

[30] *See, e.g.*, John M. Golden, *Proliferating Patents and Patent Law's "Cost Disease*," 51 HOUS. L. REV. 455 (2013) (linking the patent crisis to the proliferation of patents and arguing that recent increases in the rate of patenting is not new but reasonably modest in comparison to the growth rate of the nineteenth century).

[31] *Compare* U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT FOR FISCAL YEAR 2013 at 188 tbl.1 (noting that the PTO disposed of over 600,000 patent application in the fiscal year of 2013) [hereinafter 2013 PERFORMANCE AND ACCOUNTABILITY REPORT], *with* U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT FOR FISCAL YEAR 1994 at 62 tbl.1 (noting that the PTO disposed of approximately 180,000 patent applications in the fiscal year of 1993).

[32] *See, e.g.*, B. JAFFE & JOSH LERNER, INNOVATION AND ITS DISCONTENTS 142 (2004) (noting that the PTO's allegedly high grant rate is evidence of declining U.S. patent quality).

[33] We are not the first to make this point. *See, e.g.*, Mark A. Lemley & Bhaven Sampat, *Is the Patent Office a Rubber Stamp?*, 58 Emory L.J. 181, 186 (2008).

[34] *See* Mark A. Lemley & Bhaven Sampat, *Examiner Characteristics and the Patent Office Outcomes*, 94 REV. OF ECONOMICS & STATISTICS 817, 817 (2012).

as a function of examiner experience.[35] Importantly, Lemley and Sampat provide some evidence that it is senior examiners that are allowing too many patents rather than junior examiners that are granting too few.[36] By showing that there is at least some group within the Agency that is perhaps allowing too many patents, their analysis necessarily suggests that the Agency as a whole may be allowing too many patents. Nonetheless, Lemley and Sampat's analysis does have some shortcomings that limit the degree to which their findings are helpful for policymakers. Mainly, their analysis is unable to causally link their results to a particular feature of the Agency that may even be modified in the first place (such as the analysis we set forth below) or even causally link their results to a particular feature of a given examiner. Given that their analysis is only looking at a single snapshot in time, their methodological design itself (as admitted by the authors) cannot completely determine whether their results are being driven by a true experience effect or from a selective retention story—that is, whether their findings arise from the act of an individual examiner gaining more experience over time or from a story in which examiners with inherently permissive granting tendencies happen to be those that stay with the Agency the longest[37] In a current working paper, we demonstrate the validity of these concerns, wherein we adopt a methodological approach designed to separate these various stories. In part, signaling the presence of selection effects, this working paper demonstrates that the true effect of experience itself does not follow the monotonically increasing relationship that Lemley and Sampat report but instead may follow either an inverse-U pattern or even a strictly *negative* relationship.[38]

One of the few studies that has provided compelling empirical evidence that the PTO is over granting patents, is our prior work which utilized a natural experiment methodology to explore the influence of the Agency's fee-structure on its decision to grant patents.[39] Because the PTO garners over 50% of its patent operating budget through post-allowance fees—i.e., fees the Agency collects only if it grants patents—we posited that a financially constrained PTO

---

35 *Id.*

36 For instance, they find that senior examiners are citing less prior art than junior examiners, are more likely to allow a patent on the first office action, and are more likely to grant applications that the European Patent Office rejected than junior examiners. *Id.* at 820–21, 824.

37 Although Lemley and Sampat estimate empirical specifications suggesting that their pattern of results are not impacted by whether or not the given examiner departs the Agency within five years, *id.* at 824–25, this alternative specification cannot capture all possible ways in which a selective-retention story may transpire. In a current working paper, we demonstrate that this variable does not fully account for possible selection patterns. *See infra* note 38.

38 Michael D. Frakes and Melissa F. Wasserman, *Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents?: Evidence from Micro-Level Application Data* (working paper) (on file with author).

39 Michael D. Frakes and Melissa F. Wasserman, *Does Agency Funding Affect Decisionmaking?: An Empirical Assessment of the PTO's Granting Patterns*, 66 VAND. L. REV. 67 (2013).

may grant additional patents in an effort to generate additional funds.[40] Utilizing novel patent data we examined whether the Agency did in fact allow additional patents during times of financial distress.[41] Our results suggest that the back-end fee structure of the Agency biased a financially constrained PTO towards allowing patents.[42] Moreover, our findings suggest that the Agency was targeting its distortionary granting practices at the types of patents that it profited the most from allowing—patent applications filed by large entities and patent application that historically had high renewal rates.[43]

To be clear, the behaviors under investigation in the present Article differ from that studied in our prior research and stem from distinct underlying motivations on the part of the PTO—i.e., a desire to cut-off a never-ending, burdensome cycle of continuations rather than a desire to generate revenues. In addition to being conceptually differentiated, our analysis also demonstrates that the continuation-avoidance findings emphasized in this Article cannot simply be explained by a correlated desire on the part of the Agency to over-grant in an effort to raise additional revenues. In other words, as discussed further in the Online Appendix and in Part V.B.5, we further demonstrate that the stories studied in these respective investigations are empirically independent of each other.

Given the dearth of convincing empirical evidence, policymakers have been attempting to "fix" the broken patent system with little guidance as to what structures may be the actual cause of the system's failures. It may be that strengthening the patentability standards and granting the PTO new adjudicatory authority improves patent quality. It might also be that these adjustments fail to target the principal biasing features of the system or, even more insidiously, that these recent changes are making matters worse not better.[44] Because any policy solution to curbing the PTO's over granting proclivities should start with targeting the sources to the problem policymakers are in dire need of studies that utilize convincing empirical strategies. Moreover, given the magnitude of harms associated with a potential over-granting tendency of the PTO, it is of the utmost importance that we amass as much evidence as possible as to whether the PTO is actually biased towards granting patents, along with as much convincing evidence as possible on the sources behind such biases. To this end, this Article utilizes a strategy specifically

---

[40] *Id.* at 79–80.

[41] *Id.* at 84–85.

[42] *Id.* at 119.

[43] *Id.* This Article, however, differs from our previous work on multiple dimensions, including exploring a different possible source of PTO bias towards granting patents and empirically examining the mechanism by which the Agency would utilize to preference the granting or certain patent types over others.

[44] John F. Duffy, The Big Government Patent Bill, *Patentlyo*, http://patentlyo.com/media/docs/2011/06/the-big-government-patent-bill.pdf (June 23, 2011) (arguing that the American Invents Act will increase the cost and complexity of the American Patent system).

designed to allow causal inferences to be drawn—that is, that will allow us to discern whether particular features of the Agency actually *cause* the PTO to over-grant patents. Of course, the first step in this analysis entails identifying those features of the Agency that theoretically cause it to act in the hypothesized manner. As such, we now turn our focus to a discussion on how the PTO's inability to finally dispose of a patent application may bias a resource constrained Agency towards allowing patents.

II. A NOVEL THEORY FOR WHY THE PTO IS BIASED TOWARDS ALLOWING PATENTS

The PTO's primary task is to promptly determine which inventions merit the award of a patent.[45] The Agency's ability to perform its mission, however, may be jeopardized by a peculiar feature of the U.S. patent system—the fact that the PTO can never finally reject a patent application. The ability of aggrieved patent applicants to continuously restart the examination process upon rejection by filing repeat applications could severely undermine the examination process. This section demonstrates that a PTO may be induced to allow patents in an effort to cut off the never ending stream of repeat filings. It then examines when and if the PTO will act on this incentive to allow patents by introducing a model of Agency behavior.

A. Never Ending Stream of Repeat Filings

The United States' patent filing and examination system is peculiar. Unlike its foreign counterparts, the PTO can never truly reject of an application. That is, an aggrieved patent applicant can always choose to start the examination process over by filing a "repeat" application. Repeat applications fall generally in one of two categories: continuation applications and requests for continued examination (RCE). While a continuation application is technically a new application whereas a RCE is effectively a continuation of the same application, they are largely used for the same purpose: providing the applicant who has been denied the coverage she seeks with an additional chance for her patent application to be allowed.[46]

---

[45] 2013 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 31, at 3 (noting the PTO's "mission is to foster innovation . . . by delivering high-quality and timely examination of patent . . . applications"). We use the term patents in this Article to refer to "utility" patents. Utility patent protects the way an article is used and works. 35 U.S.C. § 101.

[46] Technically a patent application can file either a continuation application, 35 U.S.C. 120 (2006) (allowing continuations to use the filing date of the original patent application), or a request for continuing examination (RCE), 35 U.S.C. § 132(b) (2006); 37 C.F.R. § 1.114(a) (2006). Although both a RCE and a filing of a continuation application result in the restarting of the examination process of a rejected patent application, the former results in the continued prosecution of the existing application whereas the latter results in the filing of an entirely new application that is identical to the rejected application. *See* U.S. PAT. & TRADEMARK OFFICE, U.S.

To illustrate the incentives posed by repeat applications, we proceed initially by treating the PTO's decision regarding patentability as a binary one. In other words, the full scope of the invention an applicant seeks—that is, all of the claims of the patent application—are either allowed or rejected.[47] Because the vast majority of repeat applications are filed when all of the claims are rejected this assumption closely mirrors actual practice.[48] Nevertheless, for completeness we relax this assumption and explore the incentives created when only a subset of claims is allowed in Section III below.

Repeat filings have the potential to seriously undermine the examination process. Currently there is no limit on the number of repeat applications an applicant can file, thus a patent application can churn through the PTO indefinitely.[49] Because the PTO must process all patent applications that are filed, repeat filings can overwhelm the existing infrastructure. There is growing evidence that this is already occurring. Repeat filings have increased from 11% of filed applications in the fiscal year 1980 to 40% in the fiscal year of 2012. Thus, a substantial portion of the 600,000 applications currently

---

DEP'T OF COMMERCE, MANUAL OF PATENT EXAMINING PROCEDURE, § 706.07(h)("An RCE is not the filing of a new application. Thus, the Office will not convert an RCE to a new application such as an application filed under 37 C.F.R. [§] 1.53(b) . . . ."); Mark A. Lemley & Kimberly A. Moore, *Ending Abuse of Patent Continuations*, 84 B.U. L. REV. 63, 68 n.14 (2004). As these differences are immaterial for the purposes of this Article we refer continuation applications and request for continued examination collectively as RCEs. The fact that RCE is simply a request to continue the examination of the original application, a RCE enters the examination queue in the place where the parent application sat whereas a continuing application is a new application and goes to the back of the examination queue. *See* U.S. PAT. & TRADEMARK OFFICE, U.S. DEP'T OF COMMERCE, MANUAL OF PATENT EXAMINING PROCEDURE, § 708.

47 Claims represent the scope of the invention an applicant seeks to patent. *See* 35 U.S.C. § 112 (2006) ("The specification shall conclude with one or more claims particularly point out and distinctly claiming the subject matter which the applicant regards as his invention.").

48 During our period of study 63% of repeat filings were RCEs. From 2000 onwards close to 75% of repeat filings were RCEs. RCEs are utilized almost exclusively when an applicant has had her entire patent scope rejected. *See* Mark A. Lemley & Bhaven Sampat, *Examining Patent Examination*, 2010 STAN. TECH. L. REV. 2, ("RCEs keep the whole case pending in the office, so they can't be used to take a patent on narrow claims and to continue to fight for broad claims; rather, RCEs are primarily useful to continue fighting with an examiner who is reluctant to grant claims").

49 In 2007 the PTO utilized rulemaking in an attempt to limit the number of continuation applications and RCEs an applicant could file. However, the Agency ultimately rescinded the regulations amidst court challenges. *See* Changes to Practice for Continued Examination Filings, Patent Applications Containing Patentably Indistinct Claims, and Examination of Claims in Patent Applications, 72 Fed. Reg. 46716, 46716-46843 (Aug. 21, 2007) (to be codified at 37 C.F.R. pt. 1); Tafas v. Dudas, 511 F. Supp. 2d 652 (E.D. Va. 2007) (granting a preliminary injunction preventing the PTO from implementing changes to the continuation practice on the eve of their implementation); Tafas v. Dudas, 541 F. Supp. 2d 805 (2008) (granting summary judgment against the PTO); Tafas v. Doll, 559 F.3d 1345 (Fed. Cir. 2009) vacated, 328 Fed. Appx. 658 (Fed. Cir. 2009) (holding that both the claim and continuation rules were procedural in nature and within the agency's rulemaking authority and that the continuations rule were inconsistent with patent law); Tafas v. Doll, 328 Fed. Appx. 658, 658-59 (Fed. Cir. 2009) (granting petition to rehear the case en banc, vacating the panel opinion); Press Release, PTO, USPTO Rescinds Controversial Patent Regulations Package Proposed by Previous Administration, PTO Press Release #09-21 (Oct. 8, 2009), available at http://www.uspto.gov/news/09_21.jsp.

awaiting substantive review comprise repeat filings.[50] The PTO has noted that repeat filings are "having a crippling effect on the Office's ability to examine "'new' (i.e., non-continuing) applications."[51] In fact the Agency's large patent pendency times, which in the fiscal year of 2012 topped thirty-three months,[52] have led at least one commentator to quip that the PTO has become the "burial ground" for patent applications.[53]

Of course, repeat filings do not necessarily need to wreak havoc on the examination system. The PTO has been effectively fully user-fee funded since 1991 and applicants pay an examination fee for every application filed, whether initial or repeat. If the PTO collected enough in examination fees to cover the costs associated to review an application, any uptick in application filing rate could theoretically be addressed by expanding the Agency's examination capacity.[54] This, however, is not the case. While 30% of the PTO's patent operating budget is garnered from examination fees, these fees cover less than one-third of the costs incurred by the Agency to evaluate applications.[55] As a result, the Agency may lack the funds necessary to address the onslaught of repeat filings through additional hiring efforts.

How can the Agency address its growing backlog of patent applications? Lacking the legal authority to abolish repeat applications all together, the PTO could attempt to combat the logjam of applications by decreasing the incentives of applicants to file repeat applications in the first instance.[56] One way to do so, conceivably, is through its granting practices. By allowing additional patents early

---

[50] U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT FOR FISCAL YEAR 2012, at 17 [hereinafter 2012 PERFORMANCE AND ACCOUNTABILITY REPORT] (noting the PTO has a backlog of 608,283 patent applications).

The incidence of continuation practice has grown significantly since the PTO became fully user-fee funded in 1991. Omnibus Budget Reconciliation Act of 1990, Pub. L. No. 101-508, § 10101, 104 Stat. 1388, 1388-391 (1990). In the early 1990s approximately fifteen percent of patent filings comprised repeat applications whereas in the late 2000s this number skyrocketed to forty percent. These figures were derived from the data that we collected from the PTO and that are discussed in further detail in Part IV. 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 50, at 17 (noting the PTO has a backlog of 608,283 patent applications).

[51] Changes to Practice, 72 FED. REG. at 46,718.

[52] 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 50, at 3.

[53] Terry Carter, *A Patent on Problems*, ABA JOURNAL, March 2010, *available at* http://www.abajournal.com/magazine/article/a_patent_on_problems/.

[54] There are, however, questions as to whether the PTO would be actually able to hire enough new patent examiners. *See, e.g.*, U.S. GOV'T ACCOUNTABILITY OFFICE, GAO-07-1102, U.S. PATENT AND TRADEMARK OFFICE: HIRING EFFORTS ARE NOT SUFFICIENT TO REDUCE THE PATENT APPLICATION BACKLOG (2007), *available at* http://www.gao.gov/new.items/d071102.pdf.

[55] U.S. PTO, Proposed Rule Setting and Adjusting Patent Fees, 77 Fed. Reg. 55028, 55034 (proposed Sept. 6, 2012) (to be codified at 37 C.F.R. 1, 41, and 42) (stating 70% of the PTO patent operational costs stem from examination expenses). The PTO estimated that the average cost of examining a patent application was approximately in the fiscal year of 2011 was $3,600, 2011 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 33, at 17 (stating that in 2011 that the average patent cost the $3,597 to examine), yet, the examination fee in 2011 was set at only $1,090, 37 C.F.R. § 1.16(1)(1) (basic filing fee is $330); 37 C.F.R. § 1.16(k) (utility search fee is $540); 37 C.F.R. § 1.16(o) (utility examination fee is $220).

[56] *See supra* note 49.

on in the examination process the Agency extinguishes the incentive of patent applicants to refile.[57] That is, by biasing its grant rate in an upwards direction the PTO can turn off the spigot of repeat filings and hence diminish (or at least slow the growth of) its backlog of patent applications.[58] The extent to which the PTO would act on this inducement depends, in part, on the objectives of the Agency and its needs.

B. The PTO Objectives

We contend that high level officials at the PTO are largely mission minded. That is, the Agency will first and foremost attempt to faithfully carry out its mission of promoting innovation by providing

---

[57] Importantly, this may not fully extinguish a patent applicant's incentive to file a repeat application. There is likely some subset of patent applicants that are utilizing repeat applications to "build out" the patent claims over time to either cover their competitor's product, to acquire a subsequent patent to hold-up the licensee, or to ensure that their marketed product is covered by an independent claim. Pharmaceutical companies, for instance, are likely to be practicing this latter tactic. In any of these scenarios, a patent applicant may file a request for continued examination even if the PTO allows all of the claims because the applicant wants to wait until some later time when it has acquired additional information to modify its patent claims and obtain a patent.

[58] An alternative way to conceptualize the incentives created by repeat filings focuses almost exclusively on their financial consequences. Repeat filings are especially costly to the PTO to review. Because examination fees are set substantially below the costs incurred to the agency to review applications, each time a patent application is rejected and then refiled the mismatch between fees and costs only compiles further. Although the cost of examining a repeat filing is, on average, less than the cost of examining an initial filing, the savings do not reach the amount required to align fees with examination expenses. The examination fees associated with RCEs cover, on average, only half of the costs incurred by the Agency to review these applications. 37 C.F.R. § 1.17(e) (in the fiscal year of 2010 the examination fees for a RCE were $ 810 for a large entity and $ 405 for a small entity). *See* PTO, Detailed Appendices: Patent Fee Proposal at 61, *available at* http://www.uspto.gov/aia_implementation/fee_setting_-_ppac_hearing_appendices_7feb12.pdf (PTO estimates the historical cost of examining RCE is approximately $1700). *See also* Melissa F. Wasserman, *The PTO's Asymmetric Incentives: Pressure to Expand Substantive Patent Law*, 72 OHIO ST. L.J. 379, 409-10 (2011). Thus, net costs to the PTO of examining an application grow in the aggregate with each additional cycle through the Agency.

The question becomes: can the Agency break this cycle? The PTO could increase its granting proclivities. By allowing patents earlier in the examination process the PTO can forestall the filing of costly repeat filings. As a bonus, by increasing its inclinations to allow patents in this manner, the PTO will also garner additional fee income. After all, historically over fifty percent of the Agency's budget has been garnered through post-allowance fees—i.e., fees that the PTO only collects if it grants a patent. 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, supra note 50, at 72 (stating that approximately 53.4% of total patent income comes from maintenance fees and issue fees); U.S. PATENT & TRADEMARK OFFICE, ANNUAL REPORT, FISCAL YEAR 1994, at 59 (1995) (stating that approximately 54% of total patent income comes from maintenance fees and issue fees). Thus, allowing additional patents not only forestalls the filing of costly repeat applications, which results in decreasing the Agency's average cost of examining an application, but it also results in increasing the fees the PTO collects per application processed.

both timely and high quality examination of patent applications.[59] Because of the inadequacies of patent examination fees, the PTO must subsidize patent examination through other activity that is associated with fee collections. This other activity has largely been the issuance of patents, which is associated with an issuance fee, and the renewal of patents, which is associated with a maintenance fee.[60]

A mission-minded PTO would not necessarily bias its grant rate in an upwards direction in order to diminish repeat filings. If the Agency had enough resources—i.e., the Agency's grant rate and renewal fee income are sufficient—to cover the costs associated with reviewing all awaiting patent applications, the PTO would not be bound by any resource constraints. In such an instance, the Agency could process what was expected of it and hence the PTO would have no incentive to rid itself of repeat applications.

At other times, however, the Agency may lack the resources—i.e., the PTO's grant rate combined with the amount of renewal fees collected may be insufficient—to cover the Agency's examination expenses associated with both initial and repeat filings.[61] Although

---

[59] 2013 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 31, at 3 (noting the PTO's "mission is to foster innovation . . . by delivering high-quality and timely examination of patent . . . applications").

[60] Patent do not automatically remain enforceable for their entire twenty-year term; instead, they must be renewed to remain enforceable. *See* 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 50, at 73 (noting that "renewals [fees] recoup costs incurred during the initial patent process"). These post-allowance fees, which comprise over 50% of the Agency's patent budget, 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 50, at 72 (stating that approximately 53.4% of total patent income comes from maintenance fees and issue fees); U.S. PATENT & TRADEMARK OFFICE, ANNUAL REPORT, FISCAL YEAR 1994, at 59 (1995) (stating that approximately 54% of total patent income comes from maintenance fees and issue fees), are typically larger than the examination fees. In the fiscal year of 2011, the issuance fee was set at $1,510, and the maintenance fees which are due at 3 ½, 7 ½, and 11 ½ years from the date the patent issues were $980, $2,480, and $4,110 respectively. 37 C.F.R. § 1.18(a) (2011) (utility fee); C.F.R. § 1.20(e-g) (2011) (maintenance fees). Again, small entities pay half these amounts. More importantly, the expenses associated with issuing and maintaining a patent are minimal. U.S. GOV'T ACCOUNTABILITY OFFICE, GAO-97-113, FEES ARE NOT ALWAYS COMMENSURATE WITH THE COSTS OF SERVICES 26 (1997) (noting that "only 8.6 percent of the costs associated with an individual patent were attributable to the actual issue of the patent and 0.1 percent were attributable to its maintenance"). Thus, the PTO has become heavily reliant on these post-examination fees to fund the examination of patent applications. *See* 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 50, at 73 (noting that "renewals [fees] recoup costs incurred during the initial patent process").

[61] There are various factors beyond the control of the PTO that may disrupt the equilibrium reached between the Agency's fee revenue and its operational costs. The PTO is more likely to be resource constrained and encounter an imbalance between its fee income and examination or other operational costs under two broad scenarios: (1) when the PTO's operational costs increase without a corresponding increase in fee income or (2) when the stream of fee income decreases without a corresponding reduction in operational costs.

This first scenario may arise if (1) the aggregate examination costs rise due to a shift in patent applications towards more complex technology (to which the PTO allocates more examination hours) or (2) patent examinations demanded of the PTO increase relative to the existing stock of patents from which the PTO may collect post-allowance fees. The second scenario may materialize for several reasons: (1) the quality of the stream of incoming patents may deteriorate, leaving the PTO otherwise inclined to grant patents less frequently; (2) patentees may elect to pay their maintenance fees at a lower rate; or (3) the aggregate incidence of small-entity applicants may rise.

the PTO could lobby Congress for supplemental funding or increased fee levels,[62] these approaches are unlikely to yield routine success. How will the mission-minded PTO respond to this resource constraint? It is possible that a resource-bound PTO would simply focus on the proportion of its applications that it can afford to process. Because the PTO would not be able to process all applications expected of it at that time, the Agency would be forced to grow a backlog of applications awaiting review. The PTO, however, may not be willing to accept such an outcome. The Agency is under extreme pressure to decrease its mounting backlog of patent applications.[63] Moreover, the PTO has identified that its single biggest challenge is to decrease its patent pendency—the time between filing a patent application and receiving substantive communication from the agency regarding its patentability.[64]

As a result, we posit that a resource-constrained PTO may bias its grant rate in an upward direction in an effort to diminish repeat filings and decrease its backlog. Obviously, if the Agency distorts its granting behavior in this manner, it is sacrificing some patent quality. A PTO, however, that is concerned with decreasing its backlog may be willing to make such a sacrifice in an effort to better maintain its application throughput. This may be especially true because the Agency's backlog and patent pendency is a highly visible, easily measured metric whereas the Agency's non-biased grant rate is difficult to measure.[65]

III. THE DIFFERENTIAL IMPACT OF THE PTO'S INCENTIVES TO ALLOW PATENTS

The previous section established that the PTO can diminish repeat filings and concomitantly decrease (or at least slow the growth

---

In each such case, the indicated development will decrease the ratio between fees collected by the PTO and the obligatory operational costs. Thus, all else equal, these developments increase the likelihood that the Agency's fee collections will fail to cover its examinational expenses.

[62] This course of action would allow the PTO to increase its budget without distorting its own granting behavior.

[63] 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 50, at 17 (noting the PTO has a backlog of 608,283 patent applications).

[64] U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT FISCAL YEAR 2008, at 33, *available at* http://www.uspto.gov/about/stratplan/ar/USPTOFY2008PAR.pdf [hereinafter 2008 PERFORMANCE AND ACCOUNTABILITY REPORT] (noting the PTO's "biggest challenge is to address the growth of pendency and the backlog of patent applications waiting to be examined . . . ."). Jon W. Dudas, Message from the Under Secretary of Commerce for Intellectual Property and Director of the U.S. PTO, Message, *available at* http://www.uspto.gov/about/stratplan/ar/2005/02_message_director.jsp (noting that "volume and complexity of patent applications continues to outpace current capacity to examine them" and the PTO has "backlog of historic proportions").

[65] *See, e.g.*, Christi J. Guerrini, *Defining Patent Quality*, 82 FORDHAM L. REV. (forthcoming 2014) (arguing that there is no clear definition of patent quality).

of) its backlog by allowing additional patents. In order to better understand when the Agency would act on this bias of granting patents, it also introduced a model of agency behavior: a mission-minded-but-resource-constrained PTO. This section further refines this model of Agency behavior by considering an additional nuance: although the PTO can reduce the incentives to file repeat applications by inflating its grant rate, not all patent grants are equally likely to forestall the filing of a continuation application. As a result, a resource-constrained PTO may find that it will best achieve its objectives by granting more patents with respect to some types of patents relative to others. This section examines how the propensity of applicants' continuation filings bears on the PTO's ability to diminish repeat filings (and concomitantly its backlog of patent applications) by issuing additional patents.

A. Varying Propensity in the Filing of Continuation Applications

While patent applicants can always restart the examination process by filing a repeat application, not every applicant will elect to do so at the same rate. That is, some patent applications, if rejected, are much more likely to be refiled as continuations than others. The PTO stands to gain more by granting patent applications to these types of applications relative to those are more likely to be abandoned upon rejection.[66] The likelihood that the examination process would be extended upon the receipt of a final rejection, however, is not readily apparent to the Agency upon the filing of a patent application. Nevertheless, the propensity of repeat filings varies substantially across technological fields. Thus, the PTO may assess the incidence of continuation practice by using relevant historical data on continuation rates associated with patents within the same technology category of the application.[67]

Assuming that the resource-constrained PTO is mission minded, how does the differential in the propensity of continuation filings affect its decision on granting additional patents? As discussed in the previous section, the PTO may inflate its granting rate in an effort to diminish repeat filings and hence concomitantly its backlog of patent applications. However, rather than increase patenting across

---

[66] Wasserman, *supra* note 58, at 412.

[67] The PTO may assign these categorical likelihoods using a relatively coarse classification of technology types (e.g., chemical applications, electrical device applications, etc.) or, perhaps, using the more fine-grained, internal classification system that the PTO uses to instruct its examination search process. Every patent application that is filed before the PTO is assigned a classification before it enters examination. The agency utilizes classifications to funnel patent applications to examiners with the prerequisite scientific knowledge to review the application. With respect to the technology, the PTO is well suited to differentiate across patent applicants using this fine-grained internal classification system (as opposed to a broader technological classification) given that the complexity measures used to allocate examination hours (and thus examiner pay) are determined in the first instance with reference to the applicant's patent class.

the board, a resource-constrained Agency may grant relatively more patents in technologies with high-continuation-filing rates than patents in technologies with low-continuation-filing rates, as the Agency stands to decrease repeat filings more by granting the former over the latter.[68] Under the assumption of mission mindedness, the PTO will likely wish to minimize the degree to which it distorts its behavior away from its social objectives of providing timely and high quality patent review. As a result, a resource-constrained PTO that is attempting to diminish repeat filings would prefer to satisfy this goal by granting a few extra patents in technology categories with respect to which it will gain the biggest reduction—i.e., those with high-continuation-filing rates—rather than a larger number of extra patents in technology categories where this reduction in repeat filings is smaller—i.e., those with low-continuation-filing rates.

Hypothesis 1: *A resource-constrained PTO will grant patents at an incrementally higher rate for patents within technology categories that generally have high-continuation-filing rates relative to patents within categories that generally have low-continuation-filing rates.*

B. Per-Claim Allowance versus Application Allowances

Up to this point, we have treated the granting decision as a binary one: the PTO either allows or rejects the entire scope of the invention. We now turn to relaxing this assumption and considering the situation in which the PTO finds that only a subset of the proposed claims or scope of invention meet the patentability rules. While the vast majority of patent continuations are filed after *all* the proposed claims are rejected, patent applicants, at times, elect to file a continuation application upon a receipt of notice of allowance. This may happen in situations where at least some of the claims were rejected. That is, a patent applicant elects for the subset of allowed claims to matriculate into an issued patent and chooses to refile, rather than abandon, the subset of the proposed claims that has been finally rejected. The rest of this subsection demonstrates a second testable hypothesis that emerges from this richer model of PTO and applicant behavior.

In approaching this more nuanced analysis, let us first dispense with the simple scenario in which the PTO confronts an application with respect to which it would have allowed *none* of the proposed claims under a non-biased application of the patentability rules. In this situation, a resource-constrained PTO may choose to allow some or all of these claims in an effort to discourage repeat

68 Note the PTO will generally generate the same amount of post-allowance fees by granting a patent in a high-continuation-filing-rate technology as it would in a low-continuation-filing-rate technology.

filing.[69] Ultimately, this scenario is but a simple extension of the above binary model. Because a resource-healthy PTO would have rejected the patent application in question, any response in this fashion will result in distorting the Agency's *grant rate* in an upwards direction.[70] Thus, Hypothesis 1 will capture the Agency's distortionary response in this scenario.

However, now consider the previously unaddressed situation in which the Agency's non-biased application of patentability standards would result in a portion, though less than all, of the proposed claims being allowed.[71] In this instance, we contend that the PTO may respond to a situation of resource constraint by allowing some of those claims within this application that it would have otherwise rejected. This practice may at least forestall the possibility of a continuation application being filed upon these otherwise rejected claims, thereby possibly diminishing the agency's backlog of patent applications. Note that this distortionary response will not be captured by Hypothesis 1. Although a resource-constrained PTO may still attempt to diminish repeat filings by allowing additional claims, this act will no longer result in inflating its grant rate—a resource-healthy PTO would have allowed some of these claims anyway, in which case this disposition would have been reflected in the grant-rate calculation already. Accordingly, this scenario leads to a second testable hypothesis.[72]

Hypothesis Two: *A resource-constrained PTO will allow an incrementally higher number of claims for patents within technology categories that generally have high-continuation-filing rates relative to patents within categories that generally have low-continuation-filing rates.*

---

69 Of course, to succeed in achieving these cost reductions, it would need to allow enough claims (or breadth of patent scope) to discourage the applicant from refiling the application to cover the rejected claims.

70 The Agency, however, will only be able to diminish its operational costs to the extent it can forestall the filing of a continuation application. As a result, a budget strapped PTO has the incentive not to allow just one narrow claim but as many claims (or breadth of patent scope) as necessary to discourage the refiling of the application.

71 We recognize that the number of claims is an only proxy for the scope of the invention. It is possible that a resource-constrained PTO that is seeking to diminish repeat filings would allow the same number of claims as a resource-healthy PTO but the scope of the allowed claim would be larger with the former than the latter. Unfortunately, our analysis is not able to test for this scenario.

72 We may capture this particular distortionary response by testing for an increase in the average number of claims among allowed patents. Note, however, that this measure will also be affected by the fact that the base itself may grow as a result of the relevant resource pressures. After all, as Hypothesis 1 is meant to capture, the number of granted patents may also increase at these times. This overall grant-rate response may attenuate any effect of the resource shock on the average claims number calculation to the extent that the number of claims allowed *within these newly granted patents* is less than the average number of claims allowed among all patents. Of course, to the extent that we are able to document an increase in this average claims measure even in the face of this attenuation bias, we may feel confident in reaching a conclusion that the PTO likewise responds to financial woes by allowing more claims within otherwise allowed patents.

IV. DATA AND METHODOLOGY

A. Data

When the PTO lacks the resources to examine the applications expected of it, we ask whether the Agency effectively takes the easy way out and grants at higher rates to those technologies that would otherwise be inclined to file repeat applications if the Agency rejected their initial applications. To explore this question, not only is it necessary to collect technology-specific data on the PTO's patent processing practices, but it is also necessary to observe such information over a sufficiently long time period in order to draw upon enough variation in the resource condition of the Agency. Prior investigations into the PTO's granting practices across different types of patents were limited to data available from only 2001 onwards.[73] Relying upon the post-2001 period would preclude our ability to draw upon the swing from favorable resource times to unfavorable times that transpired within the PTO over the course of the 1990s. As such, through the filing of a series of Freedom of Information Act Requests, we collected previously unavailable patent processing statistics on all 4,733,263 patent applications filed before the PTO from 1991 to 2010.[74] These data include information on such measures (among others) as the number of allowances, disposals,[75] initial filings, continuation filings, and RCE filings for each technology-by-year combination. These counts are captured out of the full universe of utility patent applications filed over the relevant time period.

To begin, we use these data to calculate patent grant rates specific to given technology-by-year combinations (e.g., for genetic patents in 1995). Such rates constitute our primary outcome measure in this analysis. Consistent with the PTO's own representation of its granting practices, we calculate grant rates as the number of patents granted by the PTO divided by the number of patent applications

---

[73] *See*, *e.g.*, Lemley & Sampat, *supra* note 11, at 187–89.

[74] We limit the analysis to the period from 1991 onwards. It was at this time that the Agency became essentially fully funded through user fees. Given the role of fees in the resource sustainability analysis, this time period will allow us to target swings in resource sustainability within an environment marked by a consistent financing structure.

[75]. In the data received from the PTO, disposals include patent applications that have been allowed and abandoned. Abandoned patent applications include those that have been rejected and those that have been abandoned for business reasons. The PTO data includes a finally rejected patent application if an applicant subsequently files a continuation application in its disposals. In contrast, the disposal data provided by the PTO does not include a finally rejected patent application if an applicant subsequently files an RCE in its disposals. Nonetheless, using the separate RCE filings data we received from the PTO, we construct alternative grant rates in which we built RCE filing counts into the denominator of the grant rate calculation. We discuss these and related findings in the Online Appendix, while also discussing why the preferred specification uses the PTO-provided disposal counts as the relevant denominator.

disposed of by the PTO.[76] Under this calculation, the PTO's grant rate was roughly 70 percent over this sample period. For the purposes of calculating such rates at a technology-specific level and consistent with our previous work, we categorize technology groups according to the technological subcategories (delineating thirty-seven different groups) specified by Bronwyn Hall, Adam Jaffe, and Manuel Trajtenberg and developed for the Patent Data Project of the National Bureau of Economic Research.[77]

Critical to our analysis is the assignment to each technology group of a propensity to keep pursuing patent applications following a rejection by the PTO. As set forth above and as fleshed out further below, we predict that the Agency will be more likely to increase its grant rate in order to forestall costly continuation practices with respect to those types of applicants that are more likely to file repeat applications in the first place. This repeat-filing-proclivity (or continuation-proclivity) measure represents the second key variable of significance in our analysis. Consistent with the relevant empirical precedent, our primary approach makes this assignment according to the continuation-filing rate exercised by the relevant technology at the beginning of the analysis period.[78] In the Online Appendix, we demonstrate the general robustness of the findings presented below to various alternative allocation schemes.

Note that for the purposes of calculating a continuation rate in this base period, we take the ratio of the number of continuation applications filed to the total number of applications filed at that time (including both continuation and initial filings). In the Online Appendix, we provide a breakdown of continuation rates across each of

---

76. *See, e.g.*, EUROPEAN PATENT OFFICE, JAPANESE PATENT OFFICE, KOREAN INTELLECTUAL PROPERTY OFFICE & U. S. PATENT OFFICE, FOUR OFFICE STATISTICS REPORT 2010 (2010), *available at* http://www.trilateral.net/statistics/tsr/fosr2010/annex2.pdf.

77. Bronwyn H. Hall et al., *The NBER Patent-Citations Data File: Lessons, Insights, and Methodological Tools*, *in* PATENTS, CITATIONS, & INNOVATIONS: A WINDOW ON THE KNOWLEDGE ECONOMY 403, 434–37 (Adam B. Jaffe & Manuel Trajtenberg eds., 2002). The PTO classifies patents into nearly five hundred different technology classes. This classification scheme, however, changes somewhat over time as new classes are added or as others are divided. These compositional changes (particularly divisions) potentially complicate an empirical analysis that tracks within-category changes in PTO behavior over time. For these reasons (and to facilitate a more manageable regression framework), in our preferred specifications we group patents into the relatively coarser technology classification system set forth by Hall, Jaffe, and Trajtenberg. As demonstrated by the Online Appendix, however, the results are nearly identical when using regression specifications based on the PTO classifications themselves. In any event, this approach may constitute a more appropriate specification to the extent that the PTO elects to differentiate its granting practices (as hypothesized) at a relatively coarser level. Moreover, if the PTO does indeed differentiate all the way to the PTO classification level, any such differential response should still be observable at the more aggregated level assuming some amount of correlation of continuation proclivities across PTO classes within National Bureau of Economic Research ("NBER") subcategories, as is borne out by the data.

78 This base-period assignment helps avoid a concern that the allocation into the various continuation-proclivity bins is confounded by the Agency-level granting response we are trying to target with this analysis. In other words, this base-period assignment better ensures that we maintain stable "treatment" and "control" groups throughout the course of this natural-experiment analysis.

the 37 technology categories. This breakdown evidences a meaningful level of variation in continuation proclivities across technologies, providing support for the methodological framework discussed below, which draws upon this variation in order to help tease out the independent influence of the PTO's health on its granting practices (as distinct from the influence of potentially confounding factors) and thus to help push us in the direction of a causal interpretation of the documented relationship. We now turn to an explanation of this methodological approach.

B. Methodology

1. Basic Framework

In order to prove that the PTO is issuing too many patents, one must confront two essential obstacles. First, one must identify a source behind the PTO's propensity to issue excessive patents—that is, to allow some applications that would otherwise fail to meet proper patentability standards. While we have endeavored with this paper to broadly challenge the literature to date in failing to adequately support the proposition of excessive PTO granting through convincing empirical strategies, it is beyond the scope of this paper to attempt to rectify these methodological deficiencies with respect to all possible sources of PTO granting bias. Rather, we will attempt to empirically investigate one potentially significant bias stemming from the fact that a rejection is never really final, as set forth above. This analysis will allow us to explore in detail certain important implications arising specifically from the PTO's practice of condoning repeat filings. At the same time, this more specific exercise facilitates a general demonstration of how causal inference techniques may bring us closer to inferring that the Agency's granting practices are indeed excessive in nature.

Ideally, to draw such inferences, we would compare the PTO's actual grant rates under its present structure with those grant rates that it would have applied in a hypothetical world in which all would be the same but for the particular bias of interest—that is, in a world in which we could remove the bias stemming from the PTO's desire to curtail the burdens associated with repeat filings. This ideal comparison group, of course, is not attainable. We cannot roll back the clock and observe how the PTO would have behaved under an alternative system. As such, the second obstacle in proving that the PTO is issuing too many patents entails the construction of a counterfactual environment meant to come as close as possible to replicating this ideal comparison group.

We began the discussion above by theorizing that constrained resource conditions of the PTO induce it to attempt to diminish repeat filings by granting patent applications more permissively. At first blush, one might think to create a set of comparison environments to explore the validity of this theoretical prediction by testing whether grant rates do in fact increase following negative resource shocks to the PTO (as captured by relevant markers). Such a simple test, however, would be highly problematic. Primarily, a mere before-and-after comparison of this nature would fail to disentangle the effects on its grant rates of this negative resource shock to the PTO from the effects of other time-varying factors that likewise bear on observed grant rates—e.g., variations in average application quality, general economic conditions, etc. Possible developments of this latter nature simply render it difficult to isolate our story of interest.[79]

As such, in constructing a convincing counterfactual, not only do we use a time period characterized by favorable resource conditions as a comparison group by which to assess whether the PTO's need for resources causes it to grant at higher rates in order to diminish its expected demand of examination but we also construct a counterfactual along a separate key dimension. That is, we attempt to capture a comparison group that may be subject to all of those forces that shape the Agency's granting practices *other than the bias of interest in our present analysis*. By observing the experiences of this latter comparison group over time, we may be able to determine how grant rates would have trended across favorable and unfavorable time periods as a result of these additional, spurious influences of concern—e.g., as a result of trends in underlying application quality. In other words, if we are concerned that simply looking at grant rates before and after a negative resource shock to the Agency will be confounded by "other developments," we can attempt to find a strategy to effectively estimate such developments. Having done so, we may then be in a better position to isolate the true effect on grant rates that arise from a negative shock to the PTO's resources.

Let us restate this general strategy in terms of the mathematical steps needed to achieve this isolation. With respect to those types of applicants that are subject to all of the Agency's grant-rate influences—including our bias of interest—we calculate the difference in grant rates before and after the period of time in which the Agency experiences a negative resource shock. We then do the same for those types of applicants that are *subject to all influences other than the bias of interest*. Finally, we take the difference in these two separate estimates. This "difference-in-difference" calculation[80] should allow us to net out the effect of unobservable drivers of PTO practices and target our inquiry on the story at hand.

The question, of course, becomes: how do we find a comparison

---

[79] To be clear, our theory does not predict that grant rates will increase in an absolute sense following a tightening of the PTO's finances. Rather, we predict that grant rates would rise above what they otherwise would have been absent the deterioration in the PTO's financial state, where other influences could have nonetheless shaped granting practices in the absence of any such deterioration.

[80] For an overview of difference-in-difference estimation, *see* JOSHUA D. ANGRIST & JÖRN-STEFFEN PISCHKE, MOSTLY HARMLESS ECONOMETRICS (2009).

group of this latter variety? For these purposes, we draw upon further aspects of the theory set forth above and test for a differential grant-rate response to a negative shock to the Agency's resources—i.e., to a shock that leaves the PTO's fee revenue *insufficient* to enable the Agency's to examine what is expected of it—across applicants with varying propensities to continue rejected applications. To simplify the demonstration of this approach, assume that there are two types of applicants: (1) those in technologies with high propensities to file repeat applications upon rejection of their applications and (2) those in technologies with low propensities to engage in such practices. We can effectively view the high-continuation-propensity group as our "treatment" group. It is with respect to these applicants that we expect the PTO will want to increase their grant rates during times of resource strain. Similarly, we can view those low-continuation-propensity applicants as our "control" group—i.e., those applicants with respect to which the PTO may not feel especially inclined to grant to at excessive rates during strained times. Since the negative resource shock to the Agency should not be expected to alter the granting practices applied to these control applicants, any change in the grant rates felt by these applicants should be a reflection of those other time-varying factors that likewise bear on observed grant rates—i.e., factors other than the present mechanism of interest. Accordingly, by applying the above-specified difference-in-difference approach to these treatment and control groups, we may effectively subtract out the effect of such spurious influences of concern and leave us with a targeted estimate of the relationship of interest—that is, of the relationship between the PTO's resource condition and the extent to which it grants applications.[81]

2. Expansions beyond Basic Framework

This process of forming a counterfactual through two dimensions of differentiation—in this case, by looking before and after negative resource shocks and by looking across technologies with

[81] As indicated previously, in our primary specifications, we use observed continuation rates at the beginning of the sample period in order to determine a technology's continuation propensity. One may be concerned, however, that simply looking at the observed continuation filing rate for a technology may not truly reflect its propensity to file a continuation application in the face of a *given* rejection of an application. After all, a higher continuation rate may in part be a reflection of a higher number of opportunities to file a continuation application in the first place due to a higher baseline level of rejections. We address this concern in great detail in the Online Appendix. We first note that the distribution of baseline continuation rates across technologies do not closely match up with the distribution of baseline (i.e., beginning of sample period) grant rates across technologies. In other words, among those technologies with low levels of observed base-period continuation filing rates include those with high and low levels of observed grant rates at that time period. Moreover, in the Online Appendix, we present results of certain empirical specifications that are designed to allocate continuation proclivities across technologies in a way that better separates true proclivities from greater *opportunities* for filing continuation applications. We defer the reader to the Online Appendix for more specifics regarding this separation exercise. As demonstrated by the Online Appendix, the results presented below are entirely robust to this alternative approach.

different continuation proclivities—is a frequently used causal-inference technique in the empirical economics and law and economics literatures. The empirical specifications that we estimate below, however, are slightly richer than your basic difference-in-difference framework. We defer the reader to the Online Appendix for a discussion regarding the intricacies of this richer design.

For now, we simply wish to stress that our methodology differs—beneficially so—from the simple description set forth above insofar as we do not merely rely upon one treatment group and one control group. That is, we do not simply group all technologies into one high-continuation group and one low-continuation group. Rather, we consider technologies along a continuum of repeat-filing proclivities. We begin by assigning each technology group a given continuation rate. Then, rather than comparing just one high-repeat-filing technology with one low-repeat-filing technology, we explore how the PTO's grant rate changes *on average* as we move along this spectrum of continuation rates. Ultimately, the inquiry becomes one in which we determine whether the PTO responds to a negative resource development on average by extending higher and higher grant rates as it confronts technology groups with stronger and stronger inclinations to file repeat applications.

This more nuanced framework softens concerns—which may otherwise be paramount with just one treatment group and one control group—that the observed findings may be attributable to a spurious and coincidental development that happens to strike a particular technology during the treatment period. After all, *by chance alone* (and not as a result of the story we are telling), we could happen to observe an increase in grant rates in a particular high-continuation-rate technology during the period in which the Agency experiences a negative resource shock. However, with respect to another particular high-continuation-rate technology, again by chance alone, we could observe a relative decease in such grant rates. As such, by pulling from enough groups and looking at an average relationship over a number of groups, it is likely that all such spurious developments may wash each other out and average to zero, easing concerns that the estimated differential outcome between the treatment and control groups is merely attributable to some unknown and unobservable story. Ultimately, the ability to wash away the influence of such confounding factors leaves us with greater confidence that our estimates are truly reflective of the impact of the PTO's resource condition on its granting practices.

The benefits deriving from a rich natural experiment framework cannot be emphasized enough. Even those studies that are able to find a viable control group to evaluate the impact of some reform or some shock will be limited in their abilities to bring us closer to a causal interpretation of the observed findings when they are simply looking at an event that is felt by just one group and not felt by another—i.e., the most basic difference-in-difference framework. In effect, such analyses draw upon only 4 data points: treatment group prior to the event, treatment group after the event, control group prior to the event and control group after the event. Drawing reliable inferences from such a small number of effective observations is highly problematic. Consider, for instance, a recent publication in the *Stanford Law Review* by David Abrams and Polk Wagner[82] in which they investigate the impact on small inventors' patenting behavior of the shift from a first-to-invent system to a first-to-file system, using Canada's 1989 reform of this nature as the driver of this experimental framework and using the experiences of the U.S. to form the control analysis. Their analysis demonstrates the substantial challenges associated with a sparse research design of this nature. Contemporaneous with Canada's move to a first-to-file system was, among other things, their adoption of a deferred examination system, where parties could file applications but request that the Canadian Patent Office not review them until a later period. Such contemporaneous developments within Canada, which likely also impact patenting behavior, challenge one's ability to tease out the independent influence of the reform of interest—i.e., the move to a first-to-file system.

3. Causal Assumptions

What should be evident from the above discussion is that the results generated by empirical methodologies of this nature may only be interpreted in causal terms under a certain set of assumptions, the most important of which can be thought of as a parallel-trends assumption.[83] That is, the primary concern with such an approach is one in which the control group and the treatment group may have happened to trend in different directions over time for reasons other than the story at hand.

As suggested already, the fact that we evaluate this story over 37 distinct technology groups—not just 2 simple groups—makes this assumption more palatable. By looking at an average relationship over a spectrum of repeat-filing proclivities, we dampen the influence of unknown factors that could otherwise challenge this parallel trends assumption. As such, what we are left with in order to infer causation is a potentially milder assumption—i.e., that there are no factors missing from our analysis that are more *systematically* (as distinct from coincidentally) related to swings in the PTO's ability to meet its expected examination demand that might otherwise explain any

---

[82] David S. Abrams & R. Polk Wagner, *Poisoning the Next Apple? The America Invents Act and Individual Inventors*, 65 STAN. L. REV. 517 (2013).

[83] Alberto Abadie, *Semiparametric Difference-in-Differences Estimators*, 72 REV. ECON. STUD. 1, 1-2 ("In particular, the conventional [Difference-in-Difference] estimator requires that in absence of the treatment, the average outcomes for treated and controls would have followed parallel paths over time.").

diverging granting patterns across technologies with different repeat-filing tendencies upon rejections. In presenting the results below, we walk the reader through various ways in which we further support the validity of this assumption and in which we demonstrate that the observed findings are consistent with an actual effect on granting policies stemming from the PTO's incentive to reduce repeat filings.

4. Resource Sustainability Metrics

An essential component of this experimental design is the identification of times in which the PTO lacks sufficient resources to examine the patent applications that have been filed with the Agency. That is to identify when the Agency's fee income is insufficient to cover the costs of examination demanded of the Agency. We perform this task by constructing an empirical marker indicative of the PTO's resource health. As set forth more completely in the Online Appendix, we predict that the PTO would be more likely to be resource constrained upon various *aggregate* developments, including (1) an increase in the Agency's backlog of examinations, (2) a decrease in its annual renewal fee collections, (3) an increase in the average complexity of its examinations and (4) a decrease in the incidence of large-entity applicants. Each such development would disrupt any balance between the proportion of incoming fees and the examination costs demanded of the Agency, leaving the PTO more likely to be resource constrained.

Using data on annual fluctuations over time in these various factors, along with information on parameters of the PTO's fee schedule, we construct a composite sustainability measure for each year in our sample. This "sustainability score" is constructed so as to simulate the impact of these various factors on the PTO's resource health in a manner consistent with the empirically relevant influence of each such factor—that is, if fluctuations in the aggregate backlog truly have the largest impact on the PTO's ability to process the applications demanded of the Agency, the backlog factor will be given an appropriately higher weight in the determination of this score. More specifically, using data on these factors, we simulate an annual measure equal to the ratio between (1) the issuance and post-issuance fees generated by the existing stock of patents at a given point in time and (2) the net costs associated with the examinations demanded of the PTO at that time. A higher simulated sustainability score is suggestive of fewer resource pressures, while a decline in this score suggests a negative shock to the PTO's resources. The score is meant to provide a sense of the ease with which the PTO may use its stream of incoming funds to satisfy the substantial costs associated with all of those examinations presently awaiting the PTO, which we estimate using the Agency's aggregate backlog of pending applications.[84]

Key to the methodological design just laid out is to observe how technology-specific examination decisions depend upon fluctuations in the Agency's aggregate resource condition. An inherent assumption in this quasi-experimental framework is that one can view such resource fluctuations as plausibly "exogenous"—i.e., that the sustainability score is effectively external to the technology-specific grant-rate decisions other than through the hypothesized story. Another way to state this assumption is that there are not omitted factors that simultaneously impact both the sustainability score and the technology-specific grant rate. The fact that the score is a highly aggregated measure that already abstracts away from technology-specific factors is encouraging in this respect. Also helpful is the fact that much of the drivers of this score are factors outside of the control of the Agency.[85] Nonetheless, we set forth below a range of robustness exercises that are meant to instill even greater confidence that the results cannot simply be explained by omitted factors and that they are likely attributable to the story under investigation.

V. RESULTS

To provide a quick recap of the hypotheses under investigation, we first predict that the PTO will respond to negative shocks to the Agency's resources by (1) granting at incrementally higher rates to applicants in those technologies with stronger proclivities to continue

[84] It is worth emphasizing that the sustainability score is not meant to reflect the actual profits accruing to the PTO in a given year. Rather, this simulation is meant to capture the extent of the external cost pressures facing the Agency and the extent to which its existing patent stock is capable of generating fees to withstand that pressure.

[85] One might contend that there is a slight mechanical relationship between the outcome variable of this analysis—grant rates—and the independent variable of interest—the sustainability score—insofar as an increase in granting will bring in greater fee revenues for the Agency and thus improve its resource position. This observation is unlikely to explain our results for several reasons. To begin, we are not simply correlating grant rates with the sustainability score; rather, we are correlating swings in the sustainability score with the *differential* in grant rates across various technologies. Across-the-board increases in grant rates would improve the PTO's sustainability score, but would not contribute to this differential. Second, any such mechanical relationship would actually work in the opposite direction. That is, such a relationship would predict that an increase in grant rates would lead to an increase in the sustainability score. The fact that we predict and observe the opposite relationship in the face of this possible mechanical result suggests that our findings cannot be explained by such forces. Finally, since the dependent variable is specific to a technology, any impact of this granting decision on the aggregate sustainability score is likely to be small. As such, the observation of a strong relationship between the sustainability score and the differential grant rate is more likely to emanate from the causal chain hypothesized—that is, from aggregate fluctuations in the PTO's resource health causing the Agency to alter its examination practices—as opposed to the other way around. To add even greater confidence to this discussion, we note that the tabular regression results presented in the Online Appendix remain virtually unchanged when we attempt to sever this slight mechanical relationship by calculating a new sustainability score for each observation in the analysis that removes the contribution of the technology group associated with that observation—e.g., a sustainability score applied to the Resins category that removes all Resins-related factors contributing to the score (results available upon request).

the application process upon rejection relative to applicants in those technologies with weaker such proclivities and (2) allowing a greater number of claims within each allowed application in those high-repeat-filing-prone technologies relative to low-repeat-filing-prone technologies. By responding along both the extensive and intensive margins of allowance in this manner—i.e., both allowing more frequently on an application-by-application basis and allowing more claims per application—the Agency may find itself better able to diminish repeat filings and concomitantly decrease (or at least slow the growth of) its backlog of patent applications.

A. Primary Results: Graphical Analysis

1. Description of Graphical Approach

In the Online Appendix, we present tabular regression results for the basic difference-in-difference approach mapped out above. In the main text, however, we aim to confront this analysis in a more visible manner. As such, we focus on a pictorial depiction of our key findings. If anything, such graphical analyses are more comprehensive insofar as they allow us to observe how the story materializes dynamically on a year-by-year basis.

Figure 1: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies Overlaid with Sustainability Score

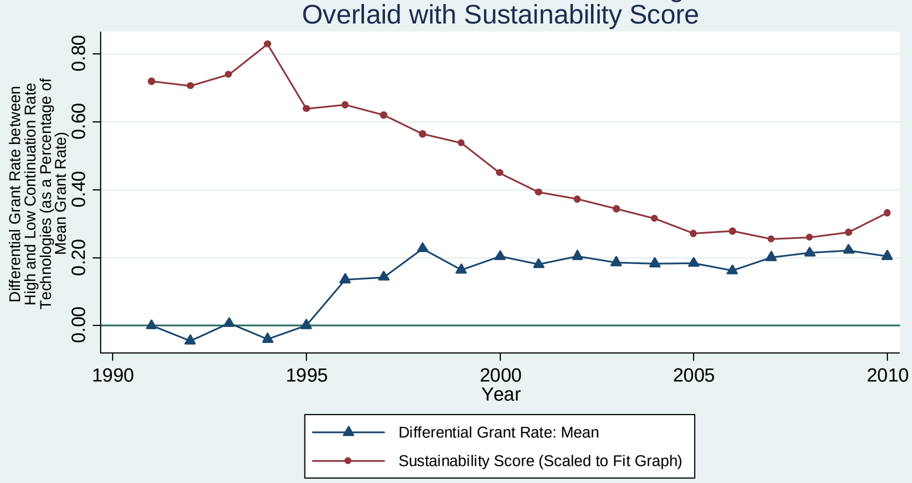

Notes: the blue-line with triangular markers represent means of the estimated coefficients of a dynamic difference-in-difference regression specification that interacts the technology-category continuation rate with indicator variables representing each year in the 1991–2010 period. This collection of estimated dynamic coefficients can be interpreted as a time trend in the differential grant rate between high and low continuation rate technologies. This differential is normalized to zero in 1991, representing the reference year. Each regression includes technology category and year fixed effects. Patent-processing data are from the PTO.

More specifically, in Figure 1, we plot a year-to-year time trend in the differential grant-rate between technologies that are highly prone to repeat a filing upon rejection and technologies that are not especially prone to file such repeat applications. We plot this trend over the 1991 to 2010 period. To state things more clearly, each value in this time trend captures the degree to which the grant rates extended to the high-repeat-filing groups exceed those extended to the low-repeat-filing groups at that time. However, rather than presenting a trend in the absolute differences in the grant rates between such groups, we normalize such differences such that they equal zero in the base period of 1991 and thereafter plot how such normalized differences evolve over each subsequent year. For instance, assume that the grant rate in the high group equals 70 percent in 1991 and the grant rate of the low group at that time equals 65 percent. We nonetheless plot a difference equal to 0 in 1991. Our interests are in learning how this baseline differential itself trends over the subsequent time period. As such, hypothetically, if the low-continuation-prone group's grant rate stays at 65 in 1992, but the high-continuation group's rate increases to 72, then we would hypothetically plot a value of 2 for 1992 since the difference in those rates would have grown by 2 percentage points over that 1-year period.[86] Similarly, if the high continuation group's rate decreases to 67, then we would hypothetically register a value of -3 for 1992.

By normalizing the differential grant rate to zero in the base period, we effectively account for and disregard any fixed, time-invariant differences in granting tendencies that apply to applicants across technologies.[87] We stress that our methodological framework is not simply capturing any inherent disparities in grant rates between high- and low-continuation-prone technologies. Such disparities are, if anything, irrelevant sources of information for our analysis. What is relevant for our purposes is how those differences (whatever level they may be at) *change over time in connection with an alteration in the resources of the PTO*. That is, when the Agency experiences a negative shock to its resources, do we see the grant-rate difference between high- and low-continuation-prone groups *grow*? This graphical and methodological framework allows us to hone in on that association of interest.[88] To facilitate a visual representation as to how trends in the differential grant rate between high- and low-continuation-prone applicants correlate with corresponding trends in the resource health of the Agency, we simply overlay the differential grant-rate time trend—marked by the blue line with triangles in Figure 1—with a

---

[86] Technically, each figure plots the differential trend in grant rates in percentage terms, as distinct from percentage-point terms. As such, a 2 percentage-point increase in the differential grant rate of interest would represent a 2.9 percent increase, considering a mean grant rate over the sample of roughly 70 percent (2 / 70 = 2.9).

[87] We acknowledge that different technologies may see their applications allowed at different rates as a fundamental matter. For instance, applicants in a particular technology field—perhaps even those with high continuation-filing proclivities—may tend generally to file higher quality applications and thus garner higher grant rates.

[88] Moreover, by doing so graphically, we arguably explore that association in a more meaningful fashion than simply estimating a single regression coefficient capturing the general association between swings in the Agency's resource sustainability and swings in these differential grant rates.

time trend in the annual "sustainability scores"—marked by the red line with circles.

Note that this approach of plotting the time trend in the difference in grant rates across our "treatment" groups and our "control" groups also allows us to neutralize any across-the-board time trends in grant rates. As will be discussed momentarily, the differential grant rate between the treatment and control groups plotted in Figure 1 begins to rise when the Agency's resource state deteriorates. This rise is not merely a reflection of a general upward trend in granting rates across the whole Agency over those years—e.g., a trend in which the grant rates rise by 5 percentage points for both the high- continuation groups and the low-continuation groups. Since, in each year, we are taking the difference between the grant rates of these respective groups, we are effectively netting out any national trends in PTO practices.

Ultimately, one can see that this graphical approach, by accounting both for fixed differences in granting tendencies across technologies and for general / across-the-board differences in granting tendencies across years, effectuates the dual layers of differencing characteristic of the "difference-in-difference" experimental design discussed in Part IV. Differencing out across-the-board trends in granting practices (by making comparisons across technologies) and differencing out inherent differences in granting practices across technologies (by making comparisons with respect to the baseline reference period), we are able to isolate the theoretical prediction of interest. In other words, this experimental design allows us to rule out a number of potentially confounding explanations for the observed patterns, leaving us with greater confidence that we are identifying the causal effect of the PTO's ability to meet its expected examination demand on its tendencies to grant excessively.[89]

Before interpreting the graphical findings depicted in Figure 1, we note that Figure 2 replicates the grant-rate analysis from Figure 1 for the case of our alternative dependent variable—that is, the average number of allowed claims in each technology-by-year group. That is, Figure 2 plots a year-by-year trend in the difference between the average number of claims in allowed applications of the high-continuation-prone technologies and the average number of claims in allowed applications of the low-continuation-prone technologies, where this difference is normalized to zero in the 1991 base period (as above). We likewise overlay this differential claims-number trend with the year-by-year trend in the Agency's sustainability score, allowing us to visualize whether the PTO will begin to allow more claims precisely when the Agency's resource status begins to deteriorate.

---

[89] This demonstration also highlights the critical importance of collecting relevant data along both a time-series and a cross-sectional—i.e., across technology units—dimension. Studies relying on just one dimension or the other, such as the Lemley and Sampat analysis discussed above, s*ee supra* note 34, will be limited in their ability to isolate causal stories.

2. Description of Findings

With this experimental framework set out, what do Figures 1 and 2 tell us? Overall, these visuals are remarkably consistent with the theoretical predictions of our analysis. To begin, consider the early years in the sample—that is, the early 1990s. At this time, the sustainability score is rising, suggesting that the Agency's resources (and hence its ability to process the applications demanded of it) may have been improving at such time. In part, this may have been due to the fact that the PTO began to collect the substantial 12-year maintenance fees for the first time during these years.[90] With such favorable conditions, one might predict that the Agency would find itself in little need of taking any distortionary actions—i.e., granting excessively—in order to diminish repeat filings. Consistent with these predictions, over these early years, the Agency does not appear to be altering the manner in which it treats the high-continuation-prone groups relative to the low-continuation-prone groups either in terms of whether it allows applications at all or in terms of how many claims to allow.

Figure 2: Time Trend in Differential Issued Claims Quantity (Logged) between Treatment and Control Technologies Overlaid with Sustainability Score

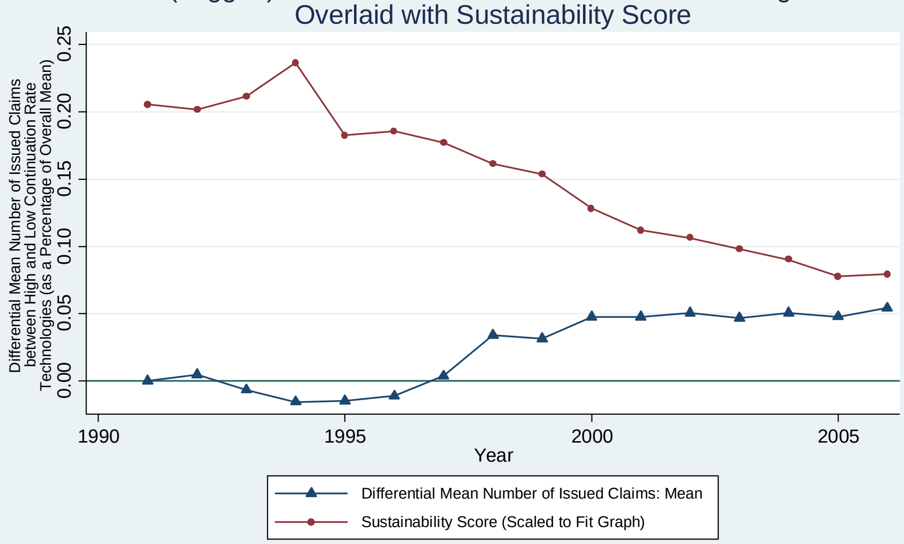

Notes: this graph replicates that of Figure 1, except using the mean number of issued claims for the technology-by-year group as the relevant outcome variable.

90 U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT, FISCAL YEAR 1994, at 29 ("In the fiscal year 1994, the PTO was just beginning to receive the full effects of the third stage renewal."). The PTO saw a substantial jump in renewal fee income in the fiscal year 1994. *Compare id.* at 59 (noting that 32% of patent fee collections resulted from maintenance fees) *with* U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT, FISCAL YEAR 1993, at 29 & fig.9 (noting that 26% of patent fee collections resulted from maintenance fees).

The fact that the differential granting outcomes stay roughly flat around the zero line over such years suggests that these various outcomes were trending in the same manner in the treatment technologies as they were in the control technologies at that time (again acknowledging that they may have been doing so at different baseline *levels*). This is encouraging for a number of reasons. Not only is it consistent with our predictions that the differential grant rate (or average number of allowed claims) is near zero on average during times when the Agency is flush with resources, but it is especially encouraging in that it remains at this near-zero level throughout the full duration of such good times. A key assumption underlying this quasi-experimental framework is that the treatment groups and the control groups would have trended in the same direction *but-for the experimental shock of interest*—in this case, but-for the negative resource shocks to the Agency that are about to follow. While it is technically impossible to test that assumption (since it is not possible to roll back the clock and adjust the PTO's resource landscape), it is at least reassuring that this "parallel trends" assumption appears to hold in the period of time leading up to the negative resource shocks that will come to the Agency's health beginning in the mid-1990s.

Now, consider the period of time following the mid-1990s. At this time, the Agency's sustainability score, calculated as above, begins to fall, suggesting that the PTO began to encounter greater and greater difficulties in marshalling enough fee income to cover the costs associated with all of those obligatory examinations awaiting review. Between 1994 and 1995, the Agency's sustainability score fell by roughly 23 percent. It thereafter fell by an additional 30 percent by the year 2000. This decline in the Agency's financial position actually continued throughout the subsequent decade, with its 2009 nearly 39 percent lower than its 2000 level.[91]

[91] Contemporaneous with this documented decline in the Agency's financial balance is naturally a substantial increase in the Agency's backlog of examinations awaiting review (and awaiting completion of examination). While this backlog grew only 14 percent throughout the first five years of the 1990s, it grew a staggering 114 percent over the subsequent five years. It then grew a further 191 percent over the course of 2000s.

Figure 3: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score with Confidence Intervals

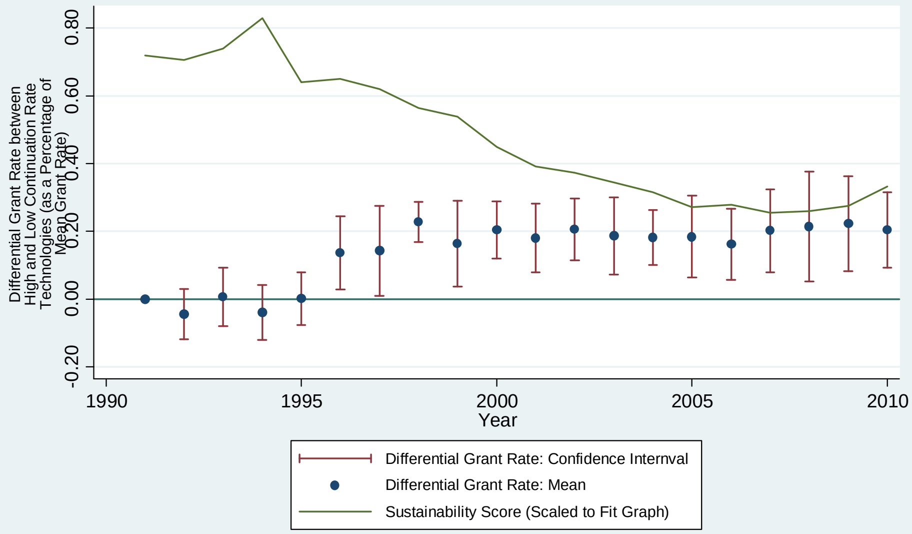

Notes: the vertical bars represent 95% confidence intervals for the estimated coefficients presented in Figure 1

The theory set forth above would suggest that upon these negative resource developments, the PTO would begin to grant at incrementally higher rates (or begin to allow an incrementally higher number of claims) to those technologies that are inherently prone to file more continuation applications relative to those that are not so prone to do as such. The results presented in Figures 1 and 2 are indeed consistent with those expectations, demonstrating the theorized up-tick in the differential granting patterns corresponding with a decline in the PTO's ability to generate enough fee income to cover the costs associated with the examination demanded of the Agency. Note that Figure 3 replicates Figure 1 but adds 95 percent confidence intervals for the associated estimates. As demonstrated by Figures 1 and 3, both the mean point estimate for the differential grant-rate of interest and the full 95 percent confidence interval for that estimate remain consistently above zero over the entire extent of this downward trajectory (i.e., there is little noise in this association), bolstering our confidence in the contention that the slide in the PTO's health lead it to grant at higher rates to the inherently high repeat-filings technologies.

Finally, in the last year of the sample (2010), it appears as if the Agency's resources begin to make a turn for the better. Ideally, it would be beneficial to observe a longer period of improving health to fully test whether we document a retreat from such distortionary granting practices as the Agency's resources begins to improve again. Nonetheless, working with this limited improvement we are able to observe in the data, we do document a slight reduction in this distortionary granting behavior in the final sample year, as arguably expected.

3. Interpreting the Magnitude of the Findings

Not only do these results suggest "some" level of distortionary granting practices in response to negative developments in the PTO's ability to process the volume of applications demanded of it, but they also suggest a substantial distortion in such practices. More specifically, we find that the difference in the grant rate between the high-continuation-prone technologies and the low-continuation-prone technologies grew by an amount equal to roughly 15 percent of the mean overall grant rate—or by roughly 11 percentage points—between 1994 and 1995. By the time the PTO's health hit bottom in 2008, the difference in the grant rate between the treatment group and the control group—relative to the baseline difference between such groups in 1991—hit a level equal to roughly 23 percent of the mean grant rate (representing roughly a 16 percentage point spread in the grant rate across these technology groups).[92]

Under an assumption that the differential grant rate between the treatment and control groups would have continued to hover around the zero line throughout the whole sample period as it did when the PTO's health was strong, these findings suggest that declines in the PTO's health combined with the Agency's inability to finally reject a patent applications may have indeed *caused* the Agency to grant a substantial number of otherwise unnecessary patents. Considering that the PTO's grant rate may have risen or fallen over the years for a number of additional reasons, these findings imply that the grant rates we have in fact observed may be substantially higher than what they otherwise would have been absent the resource considerations emphasized in this Article. *As such, unlike much of the literature to date, these findings provide direct*

[92] Both the differential grant rate across high and low-continuation-prone technologies and the differential number of allowed claims across such technologies increased slightly over the extent of this gradual deterioration in PTO health (as just demonstrated in the case of grant rates), arguably consistent with any expectations that the severity of the granting distortion would rise as the depth of the funding woes of the Agency intensified. However, it appears that in response to the Agency's negative resource slide, the PTO may have relatively quickly pushed itself to distort its granting practices across the indicated technologies close to the limits of what it was willing to endure. That is, the rate of growth in this differential rate slowed considerably after this differential's quick emergence. All else equal, after all, we do contend that the PTO may place some negative value on extending differential treatments to different types of applicants. As matters got worse and worse, it may have left itself little room to continue exhibiting any greater preference for the high continuation-prone groups. Some of the other specifications that we estimate, however, including those that focus on technologies with greater than median examiner experience (see Figure 6) and those that include control variables for the characteristics of the underlying patents (Figure 5) exhibit a more noticeable upwards trend in this differential grant rate over the course of the financial decline.

evidence in support of the contention that the PTO is indeed biased towards granting additional, unnecessary patents.

The results also suggest that the PTO will likewise begin to expand the scope of its allowed claims at the same moments. Data limitations only allow us to track average allowed claim measures until 2006. However, by that time, the difference in the average number of allowed claims between the treatment group and the control group grew by an amount (relative to the baseline difference between such groups in 1991) equal to roughly 6 percent of the mean number of allowed claims. This growth represents roughly 1.3 additional claims in the treatment group relative to the control group stemming from the deterioration in the PTO's ability to process the patent applications awaiting review.

4. Related Support

As hypothesized, the above results do indeed document an increase in the granting tendencies of the Agency—targeted on those technologies that are more likely to pose continuation threats—during times of resource constraints. If the underlying theory is correct, one might further predict in connection with such findings that the Agency would contemporaneously succeed in curtailing some amount of repeat filings in this process. This is, after all, the impetus behind the inflationary granting strategy. We test this prediction in Figure 4.

The structure of Figure 4 matches that of Figures 1–3, except that we now plot a year-by-year time trend in the differential continuation rate (which includes RCE filings) between those technologies that are generally highly prone to file continuations and those technologies that are generally not highly prone to file continuations. Recall that these "general" tendencies are identified in a fixed, time-invariant manner—that is, they do not change year-by-year.[93] We note at the outset that the year-to-year fluctuations that we do observe in the technology-specific continuation rates are not striking enough to alter our general allocation of technologies into inherently high-tendency groups and inherently low-tendency groups.[94] It is important to emphasize that this allocation process itself—i.e., designating stable treatment and control groups—is an essential part of the experimental design. With this inherent allocation of continuation proclivities laid out, we then explore how continuation rates do, in fact, change over time within these separate groups. After all, general tendencies aside, we predict that the Agency will attempt to influence these continuation patterns through its granting behavior.

---

[93] *See supra* Part IV.

[94] That is, as we proceed throughout the sample period, this allocation does not change substantially (discussed further in the Online Appendix), which is encouraging from the perspective of maintaining a stable experimental framework.

Figure 4: Time Trend in Differential Continuation Rate (Logged) between Treatment and Control Technologies Overlaid with Sustainability Score

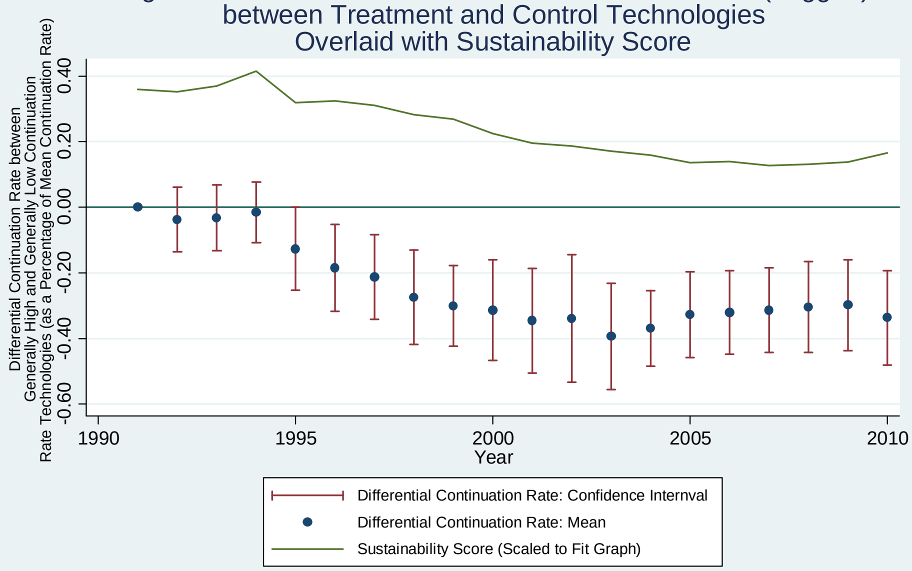

Notes: this figure replicates that of Figure 3, except that its outcome variable is now the year-by-year continuation rate for the underlying technology.

With these structural considerations in mind, what does Figure 4 tell us? As predicted by the theory, when the Agency begins to experience a decline in its ability to generate fee income to cover the volume of examination demanded of it, the time-varying continuation rate in the generally high-continuation-prone groups begins to fall relative to the generally low-continuation-prone groups. This pattern, combined with the patterns demonstrated in Figures 1–2, lend support to the idea that the Agency may grant more permissively during times of resource strain *in order to diminish repeat filings*, with such practices being concentrated among those types of applicants for which this strategy will deliver the largest bang-for-the-buck. In other words, the patterns in Figures 1–2 and 4 suggest that the PTO can decrease the rate of repeat filings by inflating its granting proclivities.

On a final note, Figure 4 is further encouraging insofar as it rules out another potentially confounding explanation for the findings in Figures 1–3. One may have otherwise been concerned that the pattern of increasing grant rates in the high-continuation-prone technologies following the mid-1990s could have been a reflection of a spurious uptick in the degree to which such technologies acted on their general tendencies and thus an uptick in the degree to which they filed repeat applications (especially RCEs) upon rejections. This repeat-filings strategy may improve an applicant's chances of receiving an ultimate allowance on its application, as the likelihood an application may be allowed is likely to increase with each additional cycle through the examination process. Thus, the high observed grant rate may be attributable not to the PTO's distortion of its granting proclivities but instead to the spurious uptick in repeat filings in our treatment groups. Could this explain the above findings? In demonstrating that the differential utilization of continuation/RCE filings, in fact, *fell* for the treatment groups relative to the control groups upon negative resource shocks to the Agency, not only does Figure 4 we lend support to the idea that the granting patterns from Figures 1–3 are reflective of a strategy in which the PTO sought to stem off costly repeat filings by granting excessively—as discussed in the preceding paragraph—but it also rules out any alternative explanation in which the observed granting patterns stem from a temporal increase in the rate of repeat filings in the treatment technologies.

5. Inclusion of Control Variables

In describing Figure 1, we simply noted that it depicted the trend in the differential grant rates between the treatment group and the control group. Technically, the points along these trends are derived through regression analysis, rather than simply hand-pulling these differences out of the data. To avoid unnecessary confusion at this stage of the discussion, we defer the reader to the Online Appendix for a more technical explanation of this approach. We bring this up now only to emphasize that by estimating this differential grant-rate trend through a regression framework, we afford ourselves the ability to control for other time-varying characteristics of the various technologies. That is, if certain characteristics of the patent landscape are likewise changing within technologies over time, we may be concerned that changes in such characteristics are responsible for the observed granting patterns, rather than the story we are claiming to isolate. To the extent that we have data on these characteristics, we may test for this directly by including such measures as additional right-hand-side variables in the underlying regression specification. Doing so allows us to partial out and neutralize the influence of these characteristics when estimating the time trend in the differential grant rate between high- and low-continuation-prone technologies. To the extent that the pattern of results presented in the above figures persists after this exercise, it is clear that variations over time in these observable characteristics of the various technologies cannot be responsible for the patterns emphasized above.

Figure 5: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score with Technology-by-Year Control Variables

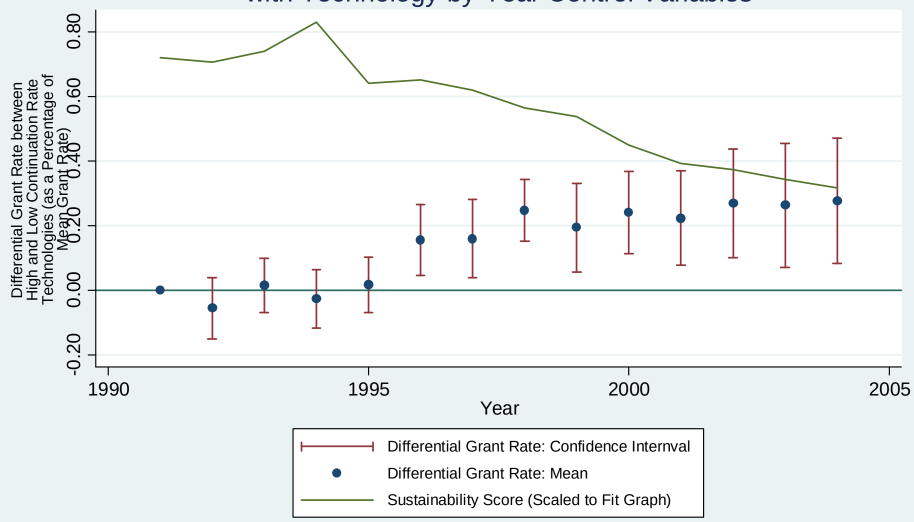

Notes: this figure replicates that of Figure 3, except that it controls for various technology-by-year patent characteristics, along with initial application filing rates.

For the purposes of this exercise, we turn to data provided by the National Bureau of Economic Research (NBER) bearing on various features of the patents issued in each year and across each technology.[95] Such features include the number of citations made to these patents—a possible indicator of patent quality[96]—along with the percentage of patents in the technology-by-year group that fall into each of eight different classifications of patentee types (i.e., individual, corporate or government; foreign or domestic).[97] Unfortunately, these NBER measures are collectively available until 2004 only. In Figure 5, we replicate Figure 3 but include these identified measures as control variables in the underlying regression. As can be observed, the pattern of results from Figure 5 closely matches that presented in Figure 3.[98]

6. Number of Technology Groups

To clarify one final matter regarding the above graphical methodology, we note that the regression underlying these graphs

---

[95] The NBER patent data can be accessed from https://sites.google.com/site/patentdataproject/Home.

[96] Manuel Trajtenberg, *A Penny for Your Quotes: Patent Citations and the Value of Innovations*, 21 RAND J. ECON. 172 (1990); Deepak Hedge and Bhaven N. Sampat, *Applicant Citations, Examiner Citations, and the Private Value of Patents*, 5 ECON. LETTERS 287 (2009).

[97] In those specifications building on Figures 1, 3 and 4, we also include the average number of allowed claims as an additional control variable.

[98] In the Online Appendix, we demonstrate the same insensitivity of the findings presented in Figure 2 to the inclusion of control variables.

does not group technologies into simple high-continuation and low-continuation bins; rather, it allows each technology to carry a unique continuation rate and asks how the impact of this negative resource shock to the Agency causes a divergence in grant rates as we move along this distribution of continuation proclivities. We quantify and calibrate this distributional movement by reference to a 2-standard-deviation shift in continuation rates across technologies—i.e., more technically, we plot the time trend in differential grant rate between technologies with continuation rate X and those whose continuation rate equals two standard deviations below X (effectively averaging over different levels of X). A shift of this magnitude is meant to capture a meaningful level of separation across technologies, facilitating our interpretation of the results as capturing the differential grant rates of "high-continuation-prone" versus "low-continuation-prone" technologies.

B. Additional Robustness Checks

1. Overview

Throughout this empirical analysis, we have attempted to illuminate certain *causal* explanations for the PTO's excessive granting patterns. A crucial feature of causal analyses involves ruling out other explanations for the observed relationship, thereby honing in on the precise story of interest. In this case, that means establishing that the PTO's response to its resource woes was indeed responsible for the observed granting patterns, rather than some additional story. The above analysis has already attempted to rule out many such stories. The essence of the basic difference-in-difference design, after all, was meant to rule out the influence of general, across-the-board time trends in granting patterns, along with inherent differences in granting tendencies across technologies. The previous subsection thereafter sought to rule out the influence of certain factors with respect to which data is available at the requisite levels—e.g., average forward-looking citations to patents issued in a particular technology and in a given year, an often emphasized metric of patent quality. In this subsection, we indicate a number of other exercises we have undertaken to rule out a range of additional explanations for the observed findings.

2. Substantive Patent Law Changes

One may be concerned that certain substantive developments in underlying patent doctrines may be responsible for the observed differential trends in the PTO's granting behavior. For instance, this may occur if, during the sample period, the law expands what constitutes patentable subject matter within technologies that happen to have high-continuation proclivities. Patent scholars have noted that patent-eligible technology has expanded largely to include inventions in the field of biotechnology, software, and business methods.[99] However, most of the legal developments of this potentially expansionary nature with respect to biotechnology occurred in the early- to mid-1980's prior to the estimation sample frame.[100] While the mid- to late-1990's likewise experienced expansions in patentable subject matter that likely targeted software and business method patents,[101] the above results are not a reflection of these developments, as demonstrated by Figure B10 in the Online Appendix. The estimates remain virtually unchanged when we remove those technology categories implicated by the relevant legal developments. In a related demonstration of the robustness of the above findings, we note that the above pattern of results cannot be explained by the contributions of any one technology alone. Dropping each technology field one-by-one, we confirm that the estimated relationship, both in terms of magnitude and in terms of statistical precision, remain virtually identical to that estimated in the above figures.

3. Selection-Effects Concerns

In this subsection, we evaluate a potential concern stemming from changes over time in initial application filings behavior across technologies. Primarily, one may be concerned that the relative increase in grant rates following the mid-1990s in the high-repeat-filing-prone technologies could be a reflection of a story in which those particular technologies started to file fewer original applications during the post-mid-1990s period. This response might leave behind a selected sample of higher-quality applications, resulting in the PTO granting more patents with respect to these technologies. We alleviate this concern in various ways. First, we note that Figure 5 above likewise includes controls for the share of the original applications filed with the PTO in a given year that are attributale to each technology. Moreover, in Figure B14 of the Online Appendix, we test directly for, and refute the existence of, any such differential decline in original application rates in high-continuation-prone technologies relative to low-continuation-prone technologies.

99. *See, e.g.*, Wasserman, *supra* note 58, at 381 (describing the "dramatic expansion of the scope of patentable subject matter").

100. *See* Diamond v. Chakrabarty, 447 U.S. 303, 305 (1980) (holding that "human-made, genetically engineered bacterium" is patentable subject matter); *see also* Policy Statement on the Patentability of Animals, 1077 Off. Gaz. Pat. & Trademark Office 24 (Apr. 21, 1987), *reprinted in* DONALD S. CHISUM, 9 CHISUM ON PATENTS app. 24-1 (2005) ("[T]he Patent and Trademark Office now considers nonnaturally occurring, nonhuman multicellular living organisms, including animals, to be patentable subject matter.").

101. *See, e.g.*, State St. Bank & Trust v. Signature Fin. Grp., 149 F.3d 1368, 1373 (Fed. Cir. 1998) (enlarging patent subject matter to include anything that provides a "useful, concrete, and tangible result"), *abrogated by In re* Bilski, 545 F.3d 943 (2008); Proposed Examination Guidelines for Computer-Implemented Inventions, 60 Fed. Reg. 28,778, 28,778 (proposed June 2, 1995).

4. Difference-in-Difference-in-Difference

In the Online Appendix, we likewise estimate what are called "triple-differences" specifications, which effectively draw upon one additional dimension of control within the above quasi-experimental framework. With this approach, we rely upon the observed fact (confirmed by our data) that "large-entity" applicants likewise are generally more prone to file continuation applications than "small-entity" applicants. Small entity applicants are individuals, nonprofit corporations, or small businesses.[102] As such, building on the theory set forth in Part III, if one really thought that the PTO would target its distortionary granting practices on those types of applicants that delivered the most "bang-for-the-buck," one might predict that they would focus mostly on large-entity applicants within high-continuation-prone technologies, or at least to a relatively stronger degree than they would increase their grant rates for small-entity applicants within such technologies. This motivates a strategy whereby we essentially evaluate whether the primary difference-in-difference findings are stronger in the case of large entities relative to small entities. Our results suggest that this is indeed the case, as set forth in the Online Appendix.

A fundamental benefit of this approach is that it effectively allows us to rule out a much larger number of additional, unobservable stories as a potential explanation for our key findings. Among other things, this approach still allows us to identify the causal link between the PTO's ability to process the patent applications awaiting review and its granting practices even if it is the case that some unrelated story that we have failed to capture in the data caused grant-rates to spike in our treatment technologies in the post-mid-1990s period. As long as that unobserved story impacted large and small entities alike within those technologies, this triple-differences strategy would still be able to neutralize the impact of any such potentially confounding explanation. The fact that our results persist in the face of this triple-differences approaches leaves us with substantially greater confidence that we are indeed identifying a grant-rate response to the PTO's dire resource conditions.

5. Other

In the Online Appendix, we further demonstrate that the pattern of results presented above is generally robust a number of additional exercises, including, among others, (1) the inclusion of control variables for the average application examination pendency associated with the relevant technology and year,[103] (2) the inclusion

---

[102] 37 C.F.R. § 1.27(a)(1)-(3).

[103] We thank Alan Marco, Acting Chief Economist at the PTO for suggesting this robustness check. This control rules out a story in which the increase in grant rates for the

of a set of controls to rule out a story in which the results can be explained by the passage of the Trade-Related Aspects of Intellectual Property Rights ("TRIPS") agreement in 1995 and by its resulting impacts on effective patent lengths,[104] and (3) the specification of technology groups according to the more fine-grained PTO Classification System, rather than the technology subcateogries specified by Hall, Jaffe, and Trajtenberg.[105]

Finally, it is critical to note that the findings presented in this Article do not simply derive from the phenomenon invesigated in our previous research whereby we found that the PTO begins to grant at higher rates to high fee-generating patent types (e.g., high maintenance-rate technologies) during times of financial strain.[106] In other words, the above findings cannot be explained by any overlap between the continuation proclivity of a technology group and the maintenance proclivity of that technology (or the incidence of large entity applicants in that technology), as demonstrate more fully in the Online Appendix, whereby we estimate specifications that effectively fit both stories simultaneously. The results of this exercise suggest that these are, in fact, distinct phenomena that emanate from distinct motivations of the Agency.

C. Results: Concluding Remarks

As far as we have endeavored to take the reader towards reaching a definitive *causal* interpretation of the findings, we acknowledge that we cannot bring the reader 100 percent to that goal. To be sure, virtually no empirical analysis can achieve such a feat. Even randomized trials—emphasized by some as the "gold standard" of causal analysis—have their limitations.[107] To emphasize certain additional virtues of the observational, quasi-experimental approach undertaken in this Article, let us say just a few final words comparing our approach to that which could arise from a controlled and randomized trial attempting to derive answers to the questions posed by our theory.

The first obvious pitfall with simply relying on the results of randomized investigations is the political impracticality of administering such trials in the first place, especially large-scale,

treatment technologies is due to a reduction in the initial examination delay by the Agency, acknowledging that any such reduction in examination pendency would make it less likely that an applicant would abandon their application and thus more likely that the application will be allowed.

104 *See* David S. Abrams, *Did TRIPS Spur Innovation? An Analysis of Patent Duration and Incentives to Innovate*, 157 U. PA. L. REV. 1613, 1615 (2009).

105 *See supra* note 77.

106 Frakes and Wasserman, *supra* note 39.

107 *See*, *e.g.*, Lisa Larrimore Ouellette, *Patent Experimentalism*, __ VA. L. REV. ___ (forthcoming), *available at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2294774 ("the gold standard [for creating comparison groups] is random assignment of the policy treatment of interest.").

expensive and long-term trials of the sort necessary to evaluate some of the most pressing policy disputes. Politics aside, randomized experiments are not without methodological flaws of their own. Consider, for instance, the often discussed "Hawthorne" effect,[108] whereby the experiment may be confounded by the knowledge of the subject that they are being evaluated in an experimental setting. In the present context, any randomization of PTO financing structures in an effort to explore how the Agency's financial situation impacts its examination decisionmaking would unlikely be immune from Hawthorne-effect concerns. How the PTO responds to a congressionally imposed experiment of which the Agency is fully aware may differ substantially from how it would respond absent such knowledge and absent the associated congressional monitoring.

Quasi-experimental studies of the sort utilized in this Article allow us to observe how the PTO's practices respond to financial woes in an actual, *real-world* setting—i.e., a setting in which the Agency is not tempering its granting behavior in order to alter the outcomes of an acknowledged experimental analysis. At the same time, of course, our analysis has provided us with an opportunity to form constructive treatment and control groups that are meant to at least approximate what such groups would look like if truly randomized. A chief challenge with a quasi-experimental approach, of course, is in showing that the treatment and control groups emphasized are otherwise comparable absent the intervention of interest—in other words, that they are *as good as randomized*. To do so will almost universally require the reader to make some assumption. Our goal with this analysis is to make the required leap of faith as small as possible and to make any such assumptions as palatable as possible. While policy analysis would ideally rest on firmer ground and require no leap at all, one must ask oneself what guidance policymakers would have without any such information at all.

VI. MECHANISMS

The previous Part set forth evidence that the PTO is indeed over granting patents during times of resource distress. Moreover, it demonstrated that the Agency is targeting its over-granting proclivities at those types of patents it stands to benefit the most from allowing—patents that are in historically high-continuation-prone technologies. One may, however, wonder just how the PTO preferences granting certain patent types over others. After all, the Agency is not a monolithic actor but instead currently employs over 8,000 patent examiners.[109] The current literature's emphasis on the

108 *See generally* John G. Adair, *The Hawthorne Effect: A Reconsideration of the Methodological Artifact*, 69 J. APPLIED PSYCH. 334 (1984) (addressing the idea that subjects may behave differently because they are aware that they are in an experiment).

109 2013 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 31, at 9 & fig.2.

autonomous nature of patent examiners and the difficulties involved in monitoring examiners casts some doubt on whether the PTO would be able to pull off such an endeavor.[110] As a result, this Part begins to explore the possible mechanisms by which the Agency would be able to target its inflationary granting tendencies.

To help illuminate the mechanism underlying the Agency's preferential treatment of certain patent types we conducted a series of interviews with former patent examiners that served as supervisors to an Art Unit, an administrative unit within the PTO comprising thirteen to twenty examiners who examine applications in the same technological field. The former examiners we had interviewed at the time this Article went to print were aware both that the continuation filings varied substantially across technologies and of the general resource health of the Agency.[111] And at least one former examiner stated that a resource-constrained PTO would attempt to increase its net resources by granting patents faster (i.e., allowing patents earlier on in the examination process).[112] Although these interviews do not definitely prove the mechanism by which the PTO is granting certain patent types over others they do shed light on the feasibility that the PTO is biasing its grant rate in an upwards direction in an effort to increase its net resources.

We believe there are at least two different channels by which the PTO can favor certain patent types over others. The first is a top-down channel, wherein high-level officials who are responsive to the cost sensitivity of repeat filings instruct examiners to preferentially grant patents filed in high-continuation-rate technologies. The PTO's ability to extend such categorical or technology-specific instructions to examiners is facilitated by the Agency's organizational structure, which is itself largely based on technological divisions. Before a patent application enters examination, it is sorted by technology type and routed to an Art Unit wherein the application is randomly

---

110 *See, e.g.*, Michael Abramowicz & John F. Duffy, *Ending the Patenting Monopoly*, 157 U. PA. L. REV. 1541, 1544, 1551, 1559–60, 1563–64 (2009) (noting the PTO's difficulties in controlling patent examiner's output despite strict rules and oversight); Meurer, *supra* note 17, at 700 (detailing the difficulties associated with implementing reforms affecting patent examiners).

Several scholars have found that patent examiner characteristics have an effect on patent outcomes. *See, e.g.*, Iain M. Cockburn et al., *Are All Patent Examiners Equal? Examiners, Patent Characteristics, and Litigation Outcomes*, *in* PATENTS IN THE KNOWLEDGE BASED ECONOMY 19 (Wesley M. Cohen & Stephen A. Merrill eds., Nat'l Acads. Press 2003) (finding that differences in examiners explain a significant percentage of the variation in the characteristics of issued patents, and that some examiners are more likely than others to have their patents upheld in court); Lemley & Sampat, *supra* note 34, at 817 (finding that more experienced examiners cite less prior art and are more likely to grant patents); Douglas Lichtman, *Rethinking Prosecution History Estoppel*, 71 U. CHI. L. REV. 151 (2004) (finding that certain examiners more systematically required applicants to narrow the scope of their patents).

111 Telephone Interview with Former Supervisory Patent Examiner # 1, at 5–6 (July 8-9, 2013) (on file with author).

112 Telephone Interview with Former Supervisory Patent Examiner # 1, at 7 (July 8-9, 2013) (on file with author).

assigned to a patent examiner.[113] Art Units are likewise aggregated into larger parcels and are eventually aggregated into one of nine technology centers.[114] This hierarchical structure creates a situation in which patents of a particular technology are consistently examined by a targeted population of examiners (i.e., applications of a particular technology are not randomly assigned among a large number of Art Units; rather, they are assigned to one or a few Art Units).[115] This consistency makes it easier for top-level officials within the Agency to coordinate with and direct examiners to grant more patents in one technology category over another. The former examiners we interviewed noted there were many top-down initiatives that occurred during their employment at the PTO.[116]

If the mechanism the PTO is utilizing to inflate its granting tendencies is a top-down channel, one would suspect that the elevated grant rate of a technology parcel (i.e., Art Unit or aggregation of several Art Units) with high-repeat filings would be distributed across all patent examiners within that parcel. That is, both senior examiners, those with more experience, and junior examiners, those with less experience, in a technology parcel with high-repeat filings would demonstrate an inflated grant rate during times the Agency's resources were insufficient to meet its expected examination demand. There is little reason to think that only experienced examiners rather than the entire group of examiners within a technology parcel would be instructed to inflate their grant rate.

The second conduit for favoring certain patent types over others is an examiner-focused channel, whereby patent examiners themselves, without necessarily receiving prompting from supervisors, respond to negative implications of repeat filings. In contrast to the top-down channel, the examiner channel, we propose, is more likely to manifest within senior patent examiners than within their less experienced counterparts. Senior examiners may be more likely to internalize the negative impact of resource shortfalls to the Agency than junior examiners for several reasons. First, senior examiners have by definition been employed by the PTO for a longer duration. As a result, they may be more likely to buy-in to the mission of the Agency and respond to its needs, including decreasing the Agency's backlog of unexamined patent applications. Even if senior examiners lack such altruistic motivations, they may nevertheless respond to the

113 *See* Lemley & Sampat, *supra* note 34, at 818, 822.

114 *Patent Technology Centers*, U.S. PAT. & TRADEMARK OFFICE, http://www.uspto.gov/about/contacts/ phone_directory/pat_tech/index.jsp (last modified Feb. 17, 2010) (listing the nine patent technology centers within the U.S. Patent and Trademark Office).

115 Art Units may be assigned patent applications from one class, a portion of a class, or from several classes involving closely related technology. *See Patent Classification: Classes Arranged by Art Unit*, U.S. PAT. & TRADEMARK OFFICE, http://www.uspto.gov/patents/resources/classification/art/index.jsp (last modified Oct. 3, 2012).

116 Telephone Interview with Former Supervisory Patent Examiner # 1, at 7 (July 8-9, 2013) (on file with author).

Agency's resource crunch by favoring certain types of patents over others because they are more likely to recognize the negative impact of resource shortfalls to their daily life. More experienced examiners likely have firsthand knowledge of how overtime may be eliminated and how special examiner programs may be cut when the Agency's fee collections are low.

Second, senior examiners have far more discretion in their granting decisions than their junior cohorts. While senior examiners' patentability decisions are rarely revisited, junior examiners work is always reviewed by a senior examiner within their Art Unit.[117] This greater leeway enables senior examiners to change their granting proclivities without fearing negative retribution. As a result, senior patent examiners may possess sufficient motivation and opportunity to preferentially grant patents in high-repeat-filing technologies relative to others when the PTO is resource constrained. It may be that such examiners are consciously inflating their grants rates but it is also possible that these examiners are subconsciously responding to the Agency's needs. We do not purport to shed light on whether senior examiners are knowingly and intentionally responding to the Agency's resource distress by inflating their grant rates but only on whether the primary channel is one that is top down or examiner driven.

Ideally, to test whether or not the story we have set forth in this Article plays out more strongly with respect to senior versus junior examiners—as hypothesized under an examiner-driven versus top-down-driven story—we would need to collect examiner-seniority-specific granting rate. That is, we would need grant rates disaggregated not only at the technology-by-year level over the full sample period, but also at the technology-by-year-by-examiner-seniority level. While we have amassed data of this nature from the early-mid 2000s onwards, this information is not readily or publicly available over the full sample period, as would be required in order to take advantage of the experimental framework set forth above. As a second-best approach, we filed additional Freedom of Information Act Requests with the PTO in order to collect data by which we could at least determine the average experience level for a given examiner within each technology over the full sample period. With this distribution of average examiner seniority across the various technologies, we can determine which technologies are generally populated by more experienced examiners and which by less experienced examiners. We perform this allocation by looking above and below the median of this seniority distribution across technologies. Finally, we may then explore whether the primary story set forth in Figure 3 is more pronounced among those generally senior technologies than it is among those generally junior technologies, which one would naturally predict if we thought that this story would be more pronounced among senior examiners themselves.

117 *See* Lemley & Sampat, *supra* note 34, at 818–19.

We conduct this exercise in Figures 6 and 7, with Figure 6 replicating Figure 3 in the case of the above-median-seniority technologies—that is, plotting the time trend in the differential grant rate between high- and low-continuation-prone technologies, confining the analysis to the set of technologies that fall into the top half of the examiner seniority distribution across technologies. Figure 7 does the same for the below-median-seniority technologies. As can be demonstrated, this analysis suggests that the PTO's inclination to grant preferentially to continuation-prone technologies is one that plays out to a noticeably stronger degree with respect to the underlying set of technologies that is typically populated by more senior examiners. Note that these contrasting patterns are even more striking when using the PTO classification system to identify technology groups, as distinct from the coarser NBER technology sub-categories, as we demonstrate in Figures 8 and 9. These findings suggest that the primary story set forth in Figure 1 (and 3) of this Article is likely one that is more relevant among the pool of senior examiners and hence that at least one channel by which the PTO is inflating its grant rate during times of resource stress emanates from the examiners themselves (as opposed to being solely driven by top-down forces).

This additional exercise not only sheds light on the mechanism by which the Agency achieves an elevated grant rate but also serves as yet another robustness check to our primary result—a resource-constrained PTO grants additional patents in an effort to decrease repeat filings. As discussed above, there is a potentially confounding explanation for the observed elevated grant rate: the underlying quality of applications in high repeat-filings technologies has been improving over time whereas the quality in low repeat-filings technologies remained unchanged. Although we have ruled out this potentially concerning explanation multiple times throughout our analysis, this additional study provides yet another way to discredit the idea that a differential change in the underlying quality of applications (or similar spurious story) is contributing to our results. Because patent applications are randomly assigned to examiners within an Art Unit, if our findings were really being determined by an improvement in the quality of applications filed in high-repeat-filing technologies then one would expect that all examiners in a high-repeat-filing Art Unit would demonstrate an elevated grant rate.[118] The fact we do not find such a universal increase in grant rates but instead find that the senior examiners are more likely than junior examiners to grant additional patents lends further support that the PTO is in fact biasing its grant rate in an upwards direction in an effort to decrease repeat filings and combat its backlog of patent applications.

[118] *See* Lemley & Sampat, *supra* note 34, at 822 (noting that patent applications are randomly assigned within at Art Unit).

Figure 6: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
Above-Median Experience Technologies

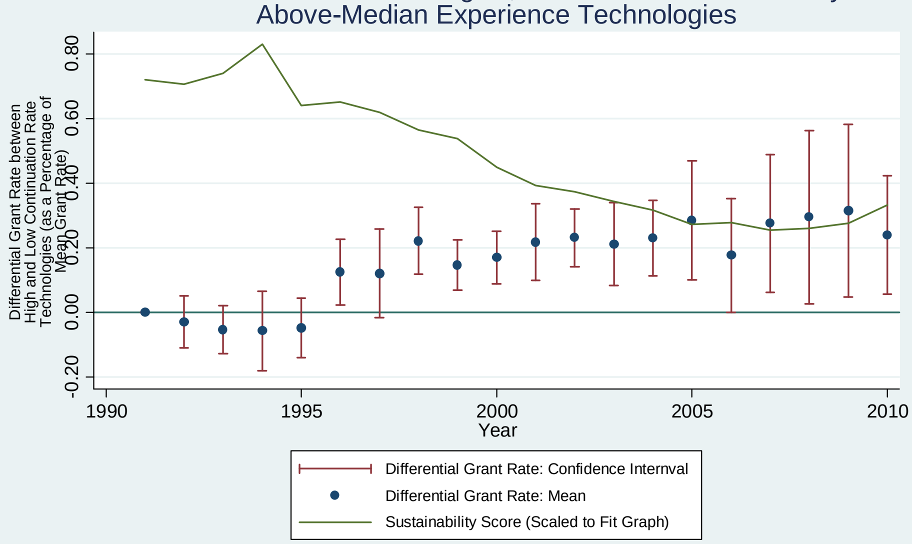

Notes: this figure replicates that of Figure 3, except that it focuses only on the set of technologies with generally above-median examiner experience levels.

Figure 7: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
Below-Median Experience Technologies

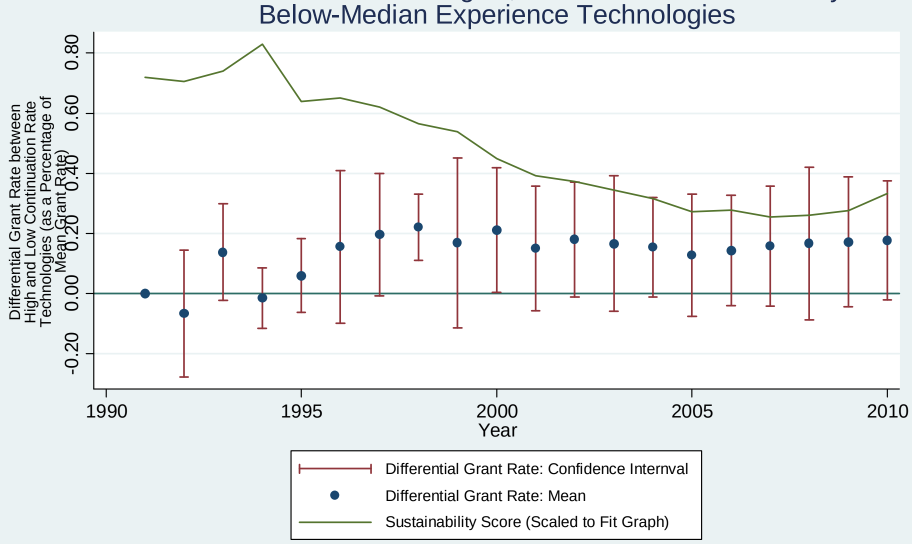

Notes: this figure replicates that of Figure 3, except that it focuses only on the set of technologies with generally below-median examiner experience levels.

Figure 8: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
Above-Median Experience Technologies
Using PTO Classes

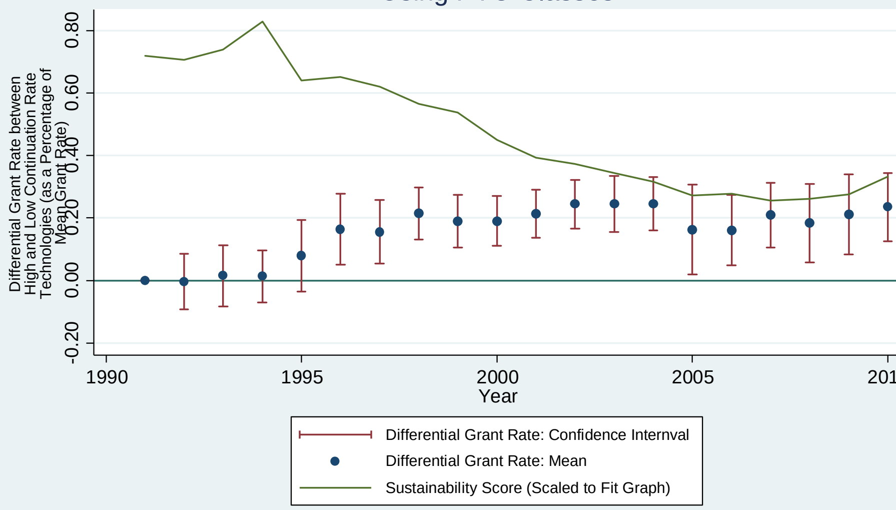

Notes: this figure replicates that of Figure 6, except that it uses PTO Classes to identify technologies

Figure 9: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
Below-Median Experience Technologies
Using PTO Classes

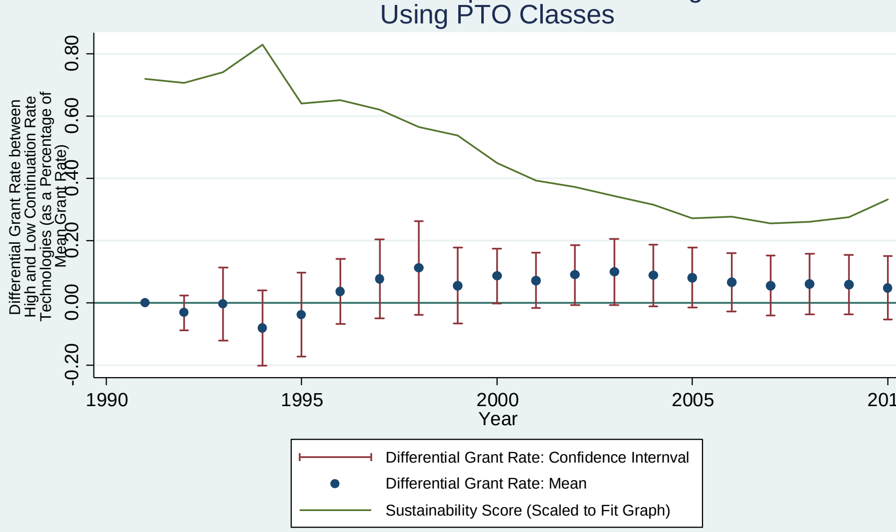

Notes: this figure replicates that of Figure 7, except that it uses PTO Classes to identify technologies.

VII. IMPLICATIONS AND REDUCING THE PTO'S INCENTIVE TO GRANT PATENTS

This section begins by exploring the implications of our findings and then sketches a possible mechanism for reducing the

Agency's incentive to grant additional applications in high-continuation-filing technologies.

A. Implications

The results presented in this Article have significant implications for the debate regarding how best to fix the broken patent system. Our findings, suggest that the existing policy debates have ignored an important and significant bias in the PTO's decision making that is essential for a complete positive and normative assessment of the patent system.

To begin, our results imply that policymakers have failed to appreciate that the PTO is not just biased towards allowing patents in general but instead biased towards allowing particular types of patents. To the extent that the Agency is extending preferential treatment to technologies in an effort to decrease repeat filings and not from legitimate social interests in intervening in certain industries, the Agency may also be undesirably distorting the allocation of resources across different sectors of the economy. As a result, the harms recognized with the granting of bad patents are likely larger than what commentators have recognized to date.

Additionally, by ignoring the import that repeat filings play in influencing PTO decision making, legal scholarship has overlooked a substantial source of Agency bias. The easiest way to eliminate the bias in PTO decision making is to extinguish its source. As a result, the recent "fixes" to the patent system, such as granting the PTO new adjudicatory authority and strengthening the nonobviousness standards in an effort to make it easier for the PTO to reject invalid patents will not curtail the substantial over granting bias identified in the Article. Our findings suggest that the Agency's incentives desperately need to be restructured.

B. Reducing the PTO's Incentives to Grant Patents in High-Repeat-Filing Technologies

Our analysis intimates that the inability of the PTO to finally reject a patent application biases the Agency towards allowing patents during times the PTO lacks sufficient resources to review the applications demanded of the Agency. Our results also suggest that a resource-constrained PTO will target its inflationary granting tendencies towards those patents with respect to which allowances which it stands to benefit from the most—patents in high-repeat-filing technologies. The PTO's incentive to grant patents could be significantly reduced by abolishing repeat filings or restructuring the Agency's fee structure either (1) to enable the PTO to expand its examination capacity to process repeat applications or (2) to minimize the risk that the PTO's fee collections will be insufficient to cover its operational costs in the first place.

To begin, the most obvious way to eliminate the Agency's incentive to allow additional patents is to abolish repeat filings altogether. If patent applicants could not continuously refile applications, the chance the PTO's existing examination infrastructure would be overwhelmed would be severely diminished. More importantly, however, if the PTO found itself resource constrained the Agency could no longer slow or diminish its backlog of applications by granting additional patents. Several scholars have suggested that continuation filings should be curtailed. Most prominently, Mark Lemley and now Judge Kimberly Moore have argued that continuation filings should be limited as many patent applicants abuse continuation practice by, among other things, modifying claim language to cover after-arising technology or obtaining multiple patents covering the same invention.[119] At the same time, commentators generally agree that the extent continuation applications are utilized in an abusive manner varies across technologies. For instance, repeat filings are generally believed to be utilized by pharmaceutical and biotechnology companies in a non-abusive manner.[120] Nevertheless, the PTO, largely out of concerns of its growing patent pendency and backlog, did attempt to limit repeat filings in 2007.[121] After a protracted court battle, in which the Agency's authority to promulgate such regulations was questioned, the PTO ultimately rescinded the regulations. Thus, it is unclear whether the PTO has the authority to limit repeat filings or whether Congress would have to amend the Patent Act to effect such a change.

Additionally, the PTO's incentive to allow patents in high-repeat-filing technologies could be nearly abolished by increasing the examination fees of repeat applications to cover the expenses associated with examining these filings. In this scenario, the examination fees of an initial application would still be set below the cost of reviewing that application, thus preserving the low examination fees of the historical fee schedule for the majority of applications. These low fees have traditionally been justified under access concerns—that is, low upfront fees increase access to the system, thereby spurring innovation.[122] Yet because examination fees of repeat filings would cover the costs associated with reviewing these applications, the PTO would have sufficient fee income to expand its examination capacity to process these applications. This assumes, of

---

[119] Lemley & Moore, *supra* note 9, at 76–79, 81.

[120] These companies often file patent application early in the development stage of their products and then utilize repeat filings to narrow their claims to ensure they have sufficient coverage of the lead compound or specific delivery mechanism. Stephen T. Schreiner & Patrick A. Doody, *Patent Continuation Applications: How the PTO's Proposed New Rules Undermine an Important Part of the U.S. Patent System with Hundreds of Years of History*, 88 J. PAT. & TRADEMARK OFF. SOC'Y 556, 557 (2006).

[121] *See supra* note 49.

[122] U.S. PTO, Final R ule Setting and Adjusting Patent Fees, 78 FED. REG. 4213 (Jan. 18, 2013) (noting that by "keeping front-end fees below the cost of application processing . . . the final fee schedule continues to . . . ease access to the patent system . . . .").

course, that the Agency would be able to hire and retain enough patent examiners to process the growing number of application filings. Although the Agency has significantly expanded its examination corps in the past five years, some have questioned whether a PTO flush with resources could hire enough examiners to dissipate its backlog of applications.[123]

Alternatively, because the PTO appears to only act upon the incentive to grant additional patents when its fee revenue fails to cover the examination expenses demanded of the Agency, the PTO's distortionary granting tendencies could be substantially diminished by increasing the PTO's overall health. One way to put the PTO is a better position it to alter its fee schedule so that the chances the Agency's fee collections are insufficient to cover the costs associated with all of those obligatory examinations awaiting review is substantially decreased.[124] The most straightforward approach to increasing the Agency's health is to decrease the PTO's reliance upon post-allowance fees to fund patent examination. By bringing examination fees to the level of examination costs, the Agency need not be concerned that application filings will grow out of step with the stock of existing patents from which it derives post-allowance fees. It will come closer to always ensuring that it has sufficient fee revenues to cover the costs of those applications demanded of the Agency.

The PTO recently obtained the ability to set its fees by rulemaking, wherein the Agency's fees cannot exceed its aggregate costs.[125] Just last year the PTO exercised this new authority and promulgated an updated fee schedule that modestly increases examination fees,[126] substantially decreases issuance fees,[127] and substantially increases renewal fees.[128] Unfortunately, the Agency's

---

[123] *See, e.g.*, UNITED STATES GOVERNMENT ACCOUNTABILITY OFFICE, U.S. PATENT AND TRADEMARK OFFICE HIRING EFFORTS ARE NOT SUFFICIENT TO REDUCE THE PATENT APPLICATION BACKLOG (2007) (concluding that "it is unlikely that the agency will be able to reduce the growing backlog simply through its hiring efforts").

[124] Alternatively, the PTO's financial risk might be diminished if the Agency was no longer dependent on fee income but instead funded through tax revenue. Congress, however, would have to provide the PTO with a sufficient budget to cover its growing examination demands, which is hardly a given. In the past, Congress has routinely utilized PTO fees to fund other governmental activities, even when the Agency's resource health was in question. *See supra* Part II.A. Further, mounting concerns over the federal government's fiscal cliff suggest that funding the Agency through taxes may not result in PTO receiving sufficient resources to process it growing backlog of patent applications. One of the primary drivers behind Congress's increased reliance on user fees to finance agencies has been belief that such funding mechanism increases the resource sustainability of agency, especially in the current environment of deficit containment.

[125] America Invents Act sec. 10.

[126] U.S. PTO, Final Rule Setting and Adjusting Patent Fees, 78 FED. REG. 4,212, 4,224 (Jan. 18, 2013) (to be codified at 37 C.F.R. 1, 41, and 42) (increasing examination fees by 27%, which equates to a $340 increase).

[127] *Id.* at 4,225 (decreasing issuance fees by 46%, which equates to an $810 decrease).

[128] *Id.* (increasing the first stage maintenance fee due at 3.5 years by 39%, which equates to a $450 increase; increasing the second stage maintenance fees due at 7.5 years by 24%, which equates to a $700 increase; increasing the third stage maintenance fees due at 11.5 years by 54%, which equates to a $2,590 increase).

new fee structure does not eliminate the risk that the PTO's fee collections will fail to cover its operational expenses. The ratio of post-allowance to examination fees remains virtually unchanged.[129] As a result, the PTO will remain heavily dependent on issuance and renewal fees to fund its examination process. As long as the PTO is subsidizing the examination process through fees generated by other activity, the health of the examination process will continue to be threatened, as the payment of these post-allowance fees may grow out of step with the examination demands on the PTO. Thus, under the new proposed fee structure, the Agency's risk that it will lack sufficient resources to processing patent applications remains.

The chances that the PTO's fee collections will fail to cover its costs associated with all of those obligatory examinations awaiting review could be significantly diminished by increasing the examination fees so that they cover the costs of patent processing—both initial and repeat filings. If the examination fees are sufficient to meet the costs of reviewing applications, then the PTO should be able to address repeat filings by expanding its examination capacity.[130] Of course, the PTO's fee structure should not only try to optimize the incentives of the Agency but also the social welfare of society more generally. Perhaps the biggest drawback from this proposal is that a substantial increase to examination fees could have a negative effect on the number of patent applications filed. Yet, studies to date have suggested that small increases to patent examination fees have a negligible effect on the volume of patent filings.[131] Moreover, because the actual fees paid to the PTO for the examination of a patent application are a fraction of the overall cost of securing a patent,[132] there is reason to believe that even a two- or three-fold increase in examination fees may not have substantial negative social welfare costs. However, more empirical investigations are needed before any definite conclusions can be drawn.

Finally, because the AIA requires that the PTO's fees must be set to a level that will recoup its operational costs—and not more than the Agency's expenses—issuance and renewal fees will likely need to

---

[129] Under the proposed fee schedule the ratio of post-allowance to examination fees is approximately 8.5 (960+1600+3600+7400)/1600 = 8.5) whereas under the current fee schedule the ratio of post-allowance to examination fees is approximately 8.6 ((2040+1130+2850+4730)/1250 = 8.6). *Id.* at 4,224–4,225.

[130] Again this assumes the PTO would be able to hire enough patent examiners to combat its growing backlog. At least one commentator has questioned whether such a use of labor force is normatively desirable. *See* Golden, *supra* note 30.

[131] *See* Gaetan de Rassenfosse & Bruno van Pottelsberghe de la Potterie, *On the Price Elasticity of Demand for Patents*, 74 OX. BULL. OF ECON. AND STAT. 58, 71-72 (2012) (finding that the demand for patents is responsive to price, but relatively inelastic); Timothy K. Wilson, *Patent Demand – A Simple Path to Patent Reform*, 2 INT'L IN-HOUSE COUNS. J. 806, 810-12 (2008) (arguing that filing fees need to be raised significantly in order to reach the elastic portion of the demand curve).

[132] A utility patent application of minimal complexity costs on average $ 10,000 to prepare and prosecute while a utility patent application of relatively high complexity costs on average $ 16,000 to prepare and prosecute. AIPLA, Report of Economic Survey 2009, at I-110-12 (2009).

be decreased if examination fees are substantially increased.[133] It is not, however, readily apparent that decreasing post-allowance fees, especially renewal fees, will increase overall social welfare. Maintenance fees are believed to serve an important social welfare function, as they effectively shorten the lifetime of a patent and hence may decrease the static costs associated with patents. As a result, any fee schedule must strike a careful balance between the incentives of the PTO and general social welfare concerns. If this becomes too difficult, it may be time to consider lifting the restriction that the aggregate PTO fees collections not exceed the Agency's operational costs. To be clear, we are not advocating that the PTO should be able to spend these excess fees; rather, we only suggest that finding the optimal fee structure may require more flexibility in setting fee levels than the AIA currently allows.

To the extent that access to the patent system remains a concern should examination fees be set to recoup examination costs, any excess post-allowance fees collected could be utilized to help subsidize the patent examination process, albeit in a less direct method than they are now. One possibility would be for the federal government to keep these post-allowance fees, add them to the general tax revenue and in turn create a refundable tax credit to patent applicants. Alternatively, the PTO may be allowed to retain these fees but instead of using them to fund its operational activity, the Agency might be required to distribute these fees as a rebate to patent applicants. Both approaches effectively avoid using any given fee to satisfy both funding and non-funding objectives, instead directing separate fees to their separate functions and diminishing the need for compromising social-welfare tradeoffs. Moreover, although these structures may achieve the same result of subsidizing the examination process through post-allowance fees, importantly they shift the financial risk from the Agency to patent applicants. When post-allowance fees fail to cover examination demands, patent applicants' refundable tax credit or rebate would be lowered, but the Agency's health would no longer suffer.

CONCLUSION

This Article presents compelling empirical evidence that the PTO is allowing too many invalid patents. Our findings suggest that the inability of the PTO to finally reject a patent application biases the Agency towards granting patents. Moreover, we find the Agency's distortions are more likely to manifest with respect to patents that it stands to benefit the most from allowing—i.e., those patents in high-repeat-filing technologies. In addition to their theoretical implications, our findings also speak to policy issues concerning patent law. Prior to our study, commentators failed to recognize the extent to

133 America Invents Act sec. 10.

which repeat filings distort the Agency's granting decisions. As a result, recent fixes to the patent system are unlikely to address the significant granting bias identified in this Article.

Does the U.S. Patent & Trademark Office Grant Too Many Patents?: Evidence from a Quasi-Experiment

Online Appendix A: Methodology

Basic Difference-in-Difference Specification

In exploring whether the Agency begins to grant at higher rates to high continuation-prone technologies upon a deterioration in the Agency's financial health, we estimate the following specification (with a unit of observation at the technology-year level):[1]

$$LOG\_GR_{ct} = \alpha + \gamma_c + \lambda_t + \beta_1 (SUSTAINABILITY_t * CONT\_RATE_c) + \beta_2 X_{ct} + \varepsilon_{ct} \quad (1)$$

where LOG_GR_{ct} is the natural log of the grant rate for technology category c (based on the National Bureau of Economic Research (NBER) patent sub-categories) in year t (grant rates are calculated as specified below). Category fixed effects and year fixed effects are specified by γ_c and λ_t, accounting for fixed differences in grant rates across technologies and years. $SUSTAINABILITY_t$ captures the annual sustainability score described below, a metric indicative of the PTO's health, with a high score representing a more favorable resource balance. Technology-specific measures of the general propensity to file continuation applications are (discussed further below) are captured by $CONT\_RATE_c$. X_{ct} includes certain time-varying covariates specific to technology categories, allowing us to account for fluctuations over time in certain characteristics of the underlying patent applications (at least out of those applications that ultimately result in a patent issuance). Such covariates include the average number of patent claims (in the grant-rate specifications), the average number of citations to the relevant patents, and the percentage of patentees representing various inventor types (e.g., individual, corporate, government, etc.).[2] Included in X_{ct} are also controls for the share of initial / original application filings for that

1. This specification is modeled after the approach taken by Daron Acemoglu and Amy Finkelstein in their investigation into the differential response across hospitals with varying levels of Medicare representation to the imposition of Medicare's Prospective Payment System. *Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector*, 116 J. POL. ECON. 837 (2008).

2. Data on claims, assignee types, and citations by technology category were obtained from the National Bureau of Economic Research Patent Data Project, *available at* https://sites.google.com/site/patentdataproject/Home.

year attributable to the given technology to address concerns that fluctuations in the grant rates between high and low continuation-prone technologies is driven by differential trends in underlying filing rates (and perhaps underlying qualities of applications). Standard errors are clustered at the technology level to account for auto-correlation in grant rates within technologies over time.

The coefficient of interest is represented by β_1, capturing the degree to which an improvement (deterioration) in the state of the PTO's resources, as captured by the sustainability score, is associated with a reduction (increase) in the differential grant-rate between high-continuation-prone and low-continuation-prone technologies. Perhaps more specifically, this coefficient captures the degree to which the PTO's grant-rate response to a negative shock to the resources of the Agency increases or decreases as we move along the distribution of continuation rates across technologies. An increase of this nature would correspond to an estimated negative coefficient of the interaction term between the sustainability score and technology-specific continuation-rate variable—i.e., the difference-in-difference variable of interest—and would be consistent with the theoretical predictions set forth in the text. This coefficient can be interpreted as an effect of the financial shock under an assumption of conditional mean independence ($E[\varepsilon | \text{SUSTAINABILITY*CONT\_RATE}, X, \gamma, \lambda, \alpha] = E[\varepsilon | X, \gamma, \lambda, \alpha]$) – that is, under an assumption that there are no unobservable shocks in granting patterns that are correlated with being in a high-continuation-prone category in the post-resource-deterioration period.

Dynamic difference-in-difference specifications. In more dynamic specifications, we endeavor to track how the differential grant-rate between high-continuation-prone and low continuation-prone technologies trends on a year-by-year basis (allowing us to visually track how such trends match up with corresponding year-by-year trends in the sustainability score). To estimate these differential grant-rate trends, we estimate the following dynamic specification:

$$LOG\_GR_{ct} = \alpha + \gamma_c + \lambda_t + \sum_{t=1991}^{2010} (\beta_t (d_t * COST_c)) + X_{ct} + \varepsilon_{ct} \quad (2)$$

where d_t is an indicator variable for year t. Since the technology-year covariates from the NBER are only available in the pre-2006 period, our primary specifications consider the full 1991-2010 period and simply drop these covariates.

Number of Claims Specifications

In those specifications, in which we estimate how a negative shock to the Agency's resources differentially impacts—across our treatment and control groups—the average number of allowed claims among issued patents (in order to explore the PTO also responds to hard times by granting a greater number of claims, as opposed to simply increasing grant rates more extensively), we estimate specifications identical to those above, except that we use the natural log of the average number of allowed claims among issued patents for each technology-year group as the dependent variable (and naturally, avoid including this measure among X).

Triple differences. Finally, we explore a richer specification premised on the assumption that the PTO will target their distortionary granting practices even more intensely on large entity applicants within high continuation-prone technologies. Behind this strategy is the idea that granting patent applications filed by large entities may be especially beneficial to a resource-constrained Agency because large entities are more prone relative to small entities to file continuation applications, as reflected in our data. As such, those applicant types that the PTO may wish to target *in order to generate the most bang-for-the-buck* with their over-granting tendencies are not just the high continuation-prone technologies, but the large entity applicants within those particular technologies.

Think of this another way. Begin by actually considering a low continuation-rate technology. In light of the general benefit of granting to large entity applicants (as above), even within this low continuation technology, the PTO may wish to extend a preferential grant rate to large entities following a negative shock to the Agency's resources. What the assumption underlying this triple-differences approach presumes is that this large-entity grant-rate bump is even higher with respect to large entities within high continuation-prone technologies.

Why would this be so? Why is it that this combination of application types will receive the real emphasis here? Perhaps, because the PTO is trying to limit its distortionary practices to only those few areas where it can really benefit the most from granting additional patents.

Accordingly, to test the hypothesis that the PTO extends even higher grant rates following a negative shock to the Agency's resources to large entities within high continuation-prone technologies, even after accounting for across-the-board preferential responses for large entities and for high-continuation-prone technologies independently, we estimate the following difference-in-difference-in-difference (AKA, triple differences) specification:[3]

$$\begin{aligned} GR_{e,c,t} = \alpha + \gamma_c + \lambda_t + \partial_{c,t} + \delta_{e,t} + \epsilon_{c,e} + LARGE \\ + \beta_1(SUSTAINABILITY_t * LARGE) \\ + \beta_2(SUSTAINABILITY_t * CONT\_RATE_c) \\ + \beta_3(LARGE * CONT\_RATE_c) \\ + \beta_4(SUSTAINABILITY_t * LARGE * CONT\_RATE_c) \\ + \varepsilon_{e,c,t} \end{aligned} \quad (3)$$

The coefficient of interest is captured by β_4, representing the degree to which the PTO extends preferential treatment following a sustainability score change to this interaction between being a large entity and being in a high continuation-prone technology. A negative coefficient confirms this more targeted granting story—that is, the negative coefficient from the difference-in-difference specification is predicted to play our more intensely with respect to large-entity, relative to small-entity, applicants (after all, a stronger negative effect entails a negative value for the triple interaction). Effectively, this specification builds on the basic difference-in-difference specification above by layering in another dimension of control—that is, taking the basic difference-in-difference story and predicting that it will play out more strongly in the case of large-entity applicants relative to small entity applicants.

A key benefit of this approach is that it also allows us to account for / rule out a much larger number of potentially confounding stories. For instance, one concern with the basic difference-in-difference specification set forth above is that there may be something we are not accounting for in the data—and that is unrelated to our theory—that can explain why grant rates rise following negative shocks to the Agency's resources in our particular "treatment" technologies. Something special may be happening in these technologies during these particular time periods. By incorporating another dimension of control in our analysis, we can hope to dismiss any such concerns (at least partially). If we assume that these potentially confounding /

---

[3] The general triple-differences methodology is motivated by Gruber. Jonathan H. Gruber, *The Incidence of Mandated Maternity Benefits*, 84 AM. ECON. REV. 622, 627 (1994).

unrelated stories play out over time with respect to large entities and small entities alike, then we may effectively use the experiences over time of the small entities—with respect to which we predict the PTO would be less inclined to over-grant—as a way to estimate any general grant-rate time trends within each of the various technologies, trends which may be impacted by a range of influences (not just the PTO's state of resources). Accordingly, by seeing how the main difference-in-difference story unfolds differentially across large and small entities, we may be able to net out these general technology-specific time trends and therefore come up with an even better way to isolate a grant-rate response to the negative shock to the Agency's resources. *That is, this triple differences approach will allow us to capture the effect of a resource shock on the PTO's granting behaviors even if turns out that some other factor—e.g., time trends in the quality of applications flowing to particular technologies—could otherwise explain why high-continuation-prone technologies happened to experience an increase in grant rates following the mid-1990s.*

Stated even more broadly, to the extent that our assumption is correct and that the PTO will indeed wish to target its granting distortions on applicants with both large-entity status and a high-continuation technology designation, the tripled-differences approach will allow us to flexibly control for unobservable factors that are specific to (1) given years / entity-size combinations, (2) given years / technology combinations and (3) given entity-size / technology combinations.

Variables.

Grant rates. Patent grant rates for each technology category-year cell are calculated as the number of patent allowances within the relevant cell divided by the number of patent disposals within that cell. Patent disposals, in turn, equal the number of patents allowed plus the number of patents abandoned. The data received from the PTO does not treat requests for continued examinations (RCEs) as abandonments. Nonetheless, we have received counts of the number of RCEs filed in each technology-by-year group. As such, in alternative specifications, we calculate grant rates where we also include RCEs in the denominator. This alternative rate is imperfect insofar as we do not know when the RCE filings are actually disposed of. We simply know when they are filed. A more accurate version of this alternative grant rate would include RCE disposals in the denominator.

Continuation rates:

Continuation rates are calculated for each technology-by-year group by dividing the number of continuation filings and RCE filings (and RCE predecessors) by the total number of filings for the relevant group.

In the preferred specification, following the relevant difference-in-difference precedent (based on a differential response to a national reform across institutional types),[4] we assign continuation rates to technology categories—to form our treatment and control groups—according to the continuation rates observed at the beginning of the sample period (in our case, we chose the year prior to the implementation of the PTO's user-fee funding structure). Keeping the continuation-proclivity allocations fixed over time in this manner (and using an early sample period allocation) will ensure stability in the treatment and control groups throughout the analysis. Clearly, there will be some amount of within-technology fluctuations in continuation rates over time—in fact, this is predicted by the theory itself. However, our data suggests stability in the allocation of those groups that fall into high- and low-continuation-rate bins throughout the full sample period. We demonstrate this in several ways. First, we find a correlation coefficient of 0.6 between the baseline 1990 continuation rate for a given technology and the average continuation rate (including RCEs in the rate) over the full sample period (not just the baseline year) for that technology. Second, if we take those above-median continuation rate technologies in 1990, we find that nearly ¾ of them remain above-median in the distribution of continuation rates across technologies when we calculate continuation rates over the full sample period, with those technologies that fell below the median generally just missing the cut. Finally, we find that only 13 percent of the variation in continuation rates over the sample period can be accounted for by variation within technologies over time in continuation rates (with the rest of the variation being accounted for by general across-the board time trends in such rates and by inherent different across technologies).

---

[4] See, e.g., Acemoglu and Finkelstein, *supra* note 1; see also David S. Abrams, *Did TRIPS Spur Innovation? An Analysis of Patent Duration and Incentives to Innovate*, 157 U. PA. L. REV. 1613, 1615 (2009).

Nonetheless, in Online Appendix B below, we demonstrate that the pattern of results presented in the key figures in the text are robust to alternative allocations of continuation rates across the technologies—mainly, using the average continuation rate (including RCEs) over the full sample period and not just the baseline 1990 period.

In Table A1, we set forth the various continuation rates discussed above, by the NBER technology sub-category.

Sustainability score.

The key driver of our empirical framework are fluctuations in the aggregate resource health of the Agency. To capture such fluctuations, we construct a metric called the sustainability score, which we now describe.

We predict that the PTO will be more likely to trigger its sustainability / break-even constraint as the ratio between its incoming post-allowance fee collections to outgoing examination expenditures falls. As predicted initially in our prior research,[5] we expect that this would be more likely to occur upon the following developments: an increase in the PTO's backlog, a decrease in its annual maintenance fee collections, an increase in its average examination complexity (i.e., the average number of hours allocated to each examination disposed of in a given year) and a decrease in the percentage of patentees that are large entities. We construct the follow sustainability measure, which captures the impact of each such development on the PTO's aggregate financial balance in a manner that facilitates across-factor comparisons in such impacts.

Broadly, the sustainability score in a given year equals the amount of incoming post-allowance fees for that year divided by the net examination costs associated with all of the patent applications awaiting examination at that time. More specifically:

$$SUST_t = \frac{MAINT\_FEE\_COLLECTIONS_t + ISSUANCE\_FEE\_COLLECTIONS_t}{DEMAND(LE)_t * NET\_COST(LE)_t + DEMAND(SE)_t * NET\_COST(SE}$$

where

$$MAINT\_FEE\_COLLECTIONS_t = (ISSUANCES(LE)_{t-12} * MAINT\_RATE(12yr)_t * \text{MAINT\_FEE}(12\text{yr})) +$$

[5] *See* Frakes & Wasserman, *supra* note **Error! Bookmark not defined.**.

$\left(ISSUANCES(SE)_{t-12} * \ MAINT\_RATE(12yr)_t * \frac{1}{2} * \right.$
$\left. \text{MAINT\_FEE(12yr)} \right) + \ (ISSUANCES(LE)_{t-8} * \ MAINT\_RATE(8yr)_t *$
$\text{MAINT\_FEE(8yr)}) + \ \left(ISSUANCES(SE)_{t-8} * \ MAINT\_RATE(8yr)_t * \right.$
$\left. \frac{1}{2} * \text{MAINT\_FEE(8yr)} \right) +$
$(ISSUANCES(LE)_{t-4} * \ MAINT\_RATE(4yr)_t * \text{MAINT\_FEE(4yr)}) +$
$\left(ISSUANCES(SE)_{t-4} * \ MAINT\_RATE(4yr)_t * \frac{1}{2} * \right.$
$\left. \text{MAINT\_FEE(4yr)} \right);$

where

$$ISSUANCE\_FEE\_COLLECTIONS_t = ISSUANCES(LE)_t * ISSUE\_FEE + ISSUANCES(SE)_t * ISSUE\_FEE * 1/2 \ ;$$

where

$$NET\_COST(LE)_t = AVG\_EXAM\_COST \left(\frac{COMPLEXITY_t}{COMPLEXITY(REFERENCE)} \right) - EXAM\_FEE;$$

where

$$NET\_COST(LE)_t = AVG\_EXAM\_COST \left(\frac{COMPLEXITY_t}{COMPLEXITY(REFERENCE)} \right) - \ 1/2 * EXAM\_FEE;$$

where

$$DEMAND(LE)_t = BACKLOG_t * LARGE\_ENTITY\_FILING\_RATE_t;$$

and where

$$DEMAND(SE)_t = BACKLOG_t * SMALL\_ENTITY\_FILING\_RATE_t.$$

As we have stated previously,[6] the above score is meant to simulate how variations in *the above-mentioned factors* (keeping all other factors fixed) affect the ability of the PTO to process all of its awaiting patent applications. That is, it provides a meaningful and empirically

---

[6] *See* Frakes & Wasserman, *supra* note **Error! Bookmark not defined.**.

relevant way of assessing the relative contributions to the PTO's resources of each of these factors. While an actual annual resource measure may consider the costs associated with those applications disposed of during a given year, this sustainability measure considers the costs associated with all of those applications awaiting examination at that time—i.e., the backlog. The resources associated with examining the backlog represent a better sense of the external pressures being placed upon the PTO, as opposed to the resources associated with those patents the PTO elected to dispose of during the year. Our goal is then to evaluate how these external pressures to the agency's resources induce it to take certain actions.

The sustainability score keeps fixed over time the per-application fee amounts themselves (based on the 2011 amounts), again focusing only on variations in the above-mentioned factors. Though, our preliminary extensions of this score based on our current understanding of fee amendments suggest that the regression results persist under this extension. Fee level fluctuations over the sample period were minor (and were out of the hands of the PTO itself). Likewise, in calculating the average cost per examination, the only factor changing over time is the average complexity of the patents disposed of during that year. To calculate net costs, we multiply the average examination cost in 2010 by the ratio of the average examination complexity for the given year (based on the distribution of patents disposed of during the year) to the average examination complexity of 2010, the reference year.

TABLE A1. CONTINUATION / RCE RATES BY TECHNOLOGY CATEGORY

	(1)	(2)
Technology Category	1990 Rate	Mean Rate over 1991-2010
Agriculture, Food, Textiles	16.9	22.6
Coating	12.5	23.4
Gas	11.5	15.7
Organic Compounds	18.6	22.3
Resins	17.2	26.9
Miscellaneous Chemical	15.6	22.7
Communications	10.5	24.9
Computer Hardware & Software	18.3	28.4
Computer Peripherals	2.5	28.9
Information Storage	14.5	30.0
Electronic business methods and software	11.1	32.4
Drugs	21.5	34.0
Surgical and Medical Instruments	14.6	29.6
Genetics	18.9	26.5
Miscellaneous Drugs and Medical	15.3	32.9
Electrical Devices	7.4	17.2
Electrical Lighting	7.4	16.7
Measuring & Testing	8.6	16.0
Nuclear & X-rays	9.4	19.2
Power Systems	8.1	17.6
Semiconductor Devices	20.8	26.4
Miscellaneous Electrical	10.9	20.9
Mat. Proc & Handling	12.6	19.3
Metal Working	11.9	19.7
Motors & Engines & Parts	6.8	13.0
Optics	8.8	21.1
Transportation	9.1	14.4
Miscellaneous Mechanical	10.5	17.5
Agriculture, Husbandry and Food	11.7	19.1
Amusement Devices	6.9	21.0
Apparel & Textile	8.9	13.1
Earth Working & Wells	7.9	17.6
Furniture, House Fixtures	7.1	13.8
Heating	9.3	14.8
Pipes & Joints	15.0	19.1
Receptacles	9.1	16.8
Miscellaneous Other	8.7	18.2

Online Appendix B: Additional Results and Robustness Checks

Difference-in-Difference Results

In Table A2, we present regression results for our estimation of speciation (1) set forth above. These simple tabular results complement the dynamic specifications (see equation (2) above), which form the heart of our analysis and which we present in graphical form. Both the sustainability score and the continuation rate variables are scaled such that an increase in 1 in the measure reflects a two-standard deviation shift in the respect measure.

TABLE A2: DIFFERENCE-IN-DIFFERENCE RESULTS

	(1)	(2)	(3)	(4)
Panel A: Grant-Rate Specification				
Difference-in-Difference Coefficient Estimate (Estimate of SUSTAINABILITY*CONT_RATE)	-0.13*** (0.04)	-0.08*** (0.03)	-0.23*** (0.08)	-0.23*** (0.08)
Number of observations				
Panel B: Number of Claims Specification				
Difference-in-Difference Coefficient Estimate (Estimate of SUSTAINABILITY*CONT_RATE)	-0.05*** (0.02)	-1.02*** (0.36)	-0.04* (0.02)	-0.04* (0.02)
Dependent Variable Logged?	Y	N	Y	Y
Including Technology-Year Covariates from the NBER?	N	N	Y	Y
Including controls for rate of initial application filings for that year attributable to given technology	N	N	N	Y

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within technology categories over time. All regressions include technology fixed effects and year fixed effects. The sustainability score and technology-specific continuation rates are each normalized such that an increase in "1" of each normalized variable represents an increase of 2 standard deviations of the respective measure. Regressions are weighted by the number of disposals with each technology-year cell.

Robustness Checks: Graphical Analyses

Below, we set forth various figures replicates the main figures presented in the text, but subject to the alterations identified in the notes.

Figure B1: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
First Alternative Continuation-Proclivity Allocation

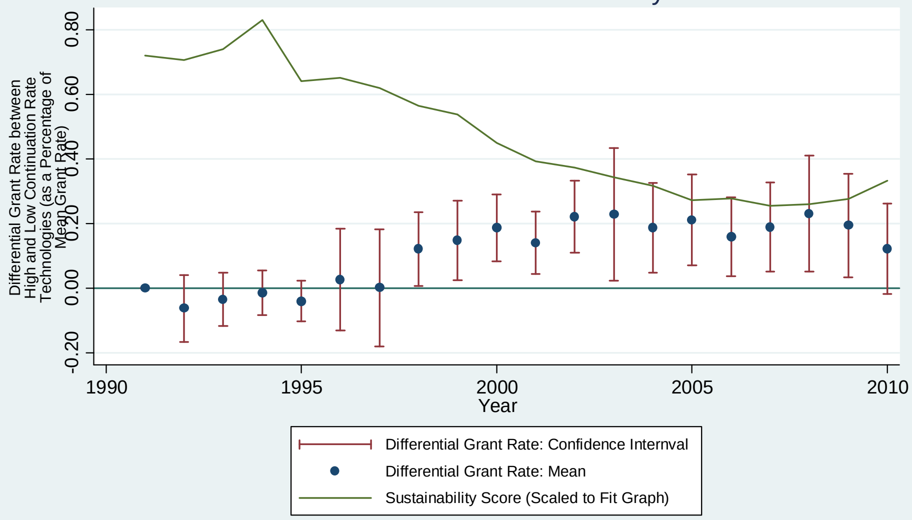

Notes: this figure replicates that of figure 3 in the text, except that it forms treatment and control groups based on continuation rates between 1986 and 1990, not just 1990

Figure B2: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
Second Alternative Continuation-Proclivity Allocation

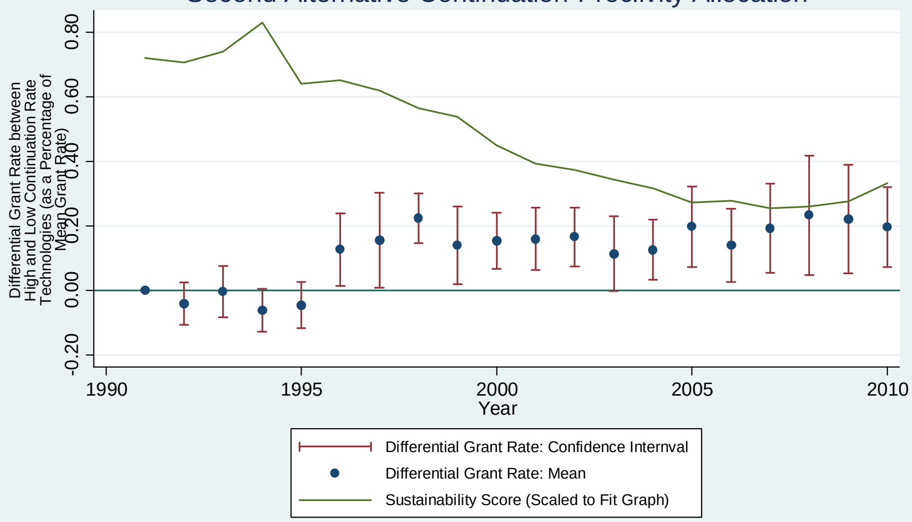

Notes: this figure replicates that of figure 3 in the text, except that it forms treatment and control groups based on continuation rates prior to 1995 (not just 1990)--that is, prior to the PTO's financial downturn.

Figure B3: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Third Alternative Continuation-Proclivity Allocation

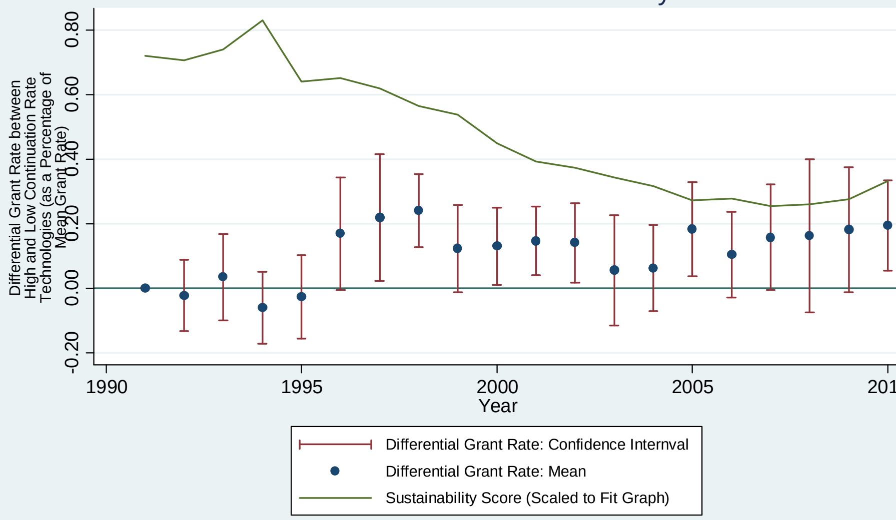

Notes: this figure replicates that of figure 3 in the text, except that it forms treatment and control groups based on mean continuation rates over the entire sample period.

Figure B4: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Using PTO Classes

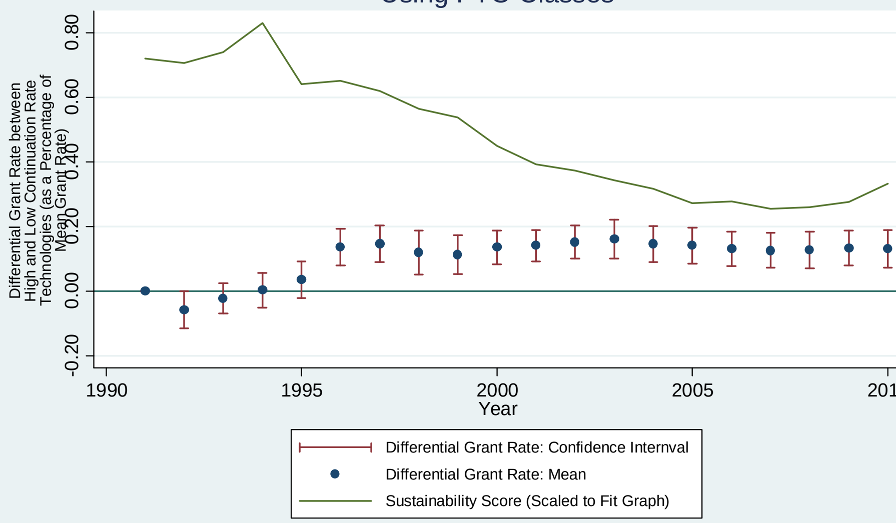

Notes: this figure replicates that of Figure 3 except that it determines technologies according to PTO classes as opposed to NBER sub-categories

In the tabular regression approach, the claims-number specification exhibits a statistically significant relationship in the predicted direction. The dynamic regressions are more taxing insofar as they require estimating the differential grant rate across technologies on a year-by-year basis, using up more degrees of freedom. As such, it is not surprising that this will induce slightly more noise into the estimates. This extra noise is enough to push the confidence bounds for the coefficients that map out the differential grant rate trends to cover the zero line (see Figure B5). Nonetheless, the pattern of the mean coefficients themselves follow that presented for the grant-rate specifications, lending support to the general idea that the PTO appears to be more permissive in its granting to the high continuation-prone technologies during times of resource strain. The confidence bounds, however, do exceed the zero line when specifying technologies according to the more fine-grained PTO classification system, an approach that naturally provides more variation in the data (see Figure B6).

Figure B5: Time Trend in Differential Issued Claims Quantity (Logged) between Treatment and Control Technologies Overlaid with Sustainability Score with confidence intervals

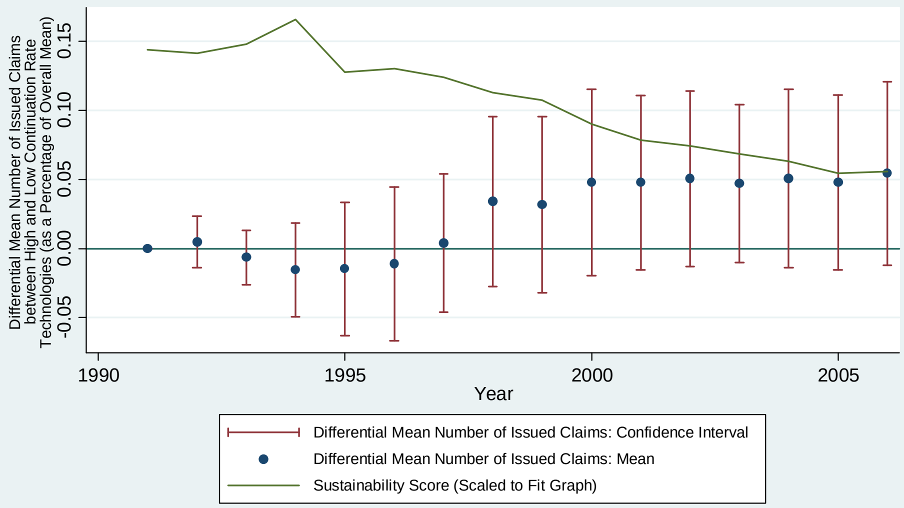

Notes: this graph replicates that of Figure 2, except that it includes the 95 percent confidence intervals.

Figure B6: Time Trend in Differential Issued Claims Quantity (Logged) between Treatment and Control Technologies Overlaid with Sustainability Score with confidence intervals, using PTO Classes

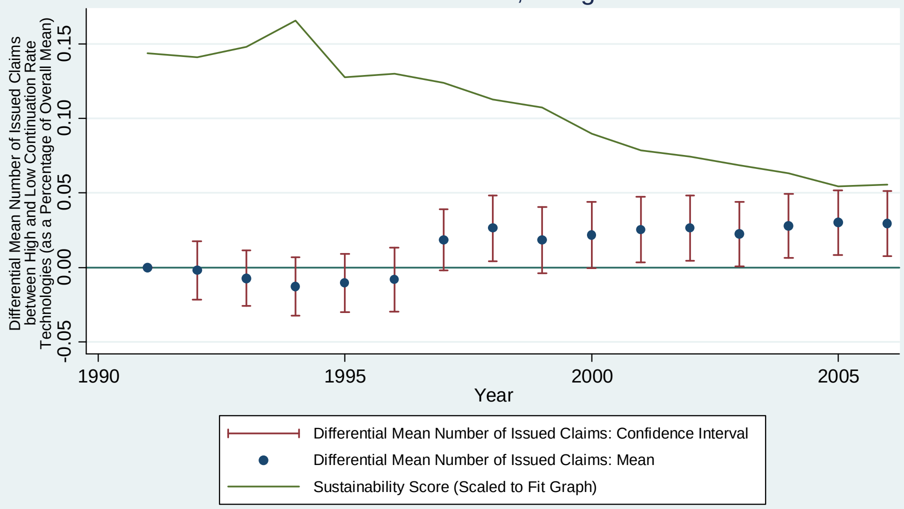

Notes: this graph replicates that of Figure B5, except that it uses PTO Classes to identify technologies.

Figure B7: Time Trend in Differential Issued Claims Quantity (Logged) between Treatment and Control Technologies Overlaid with Sustainability Score with technology-year control variables

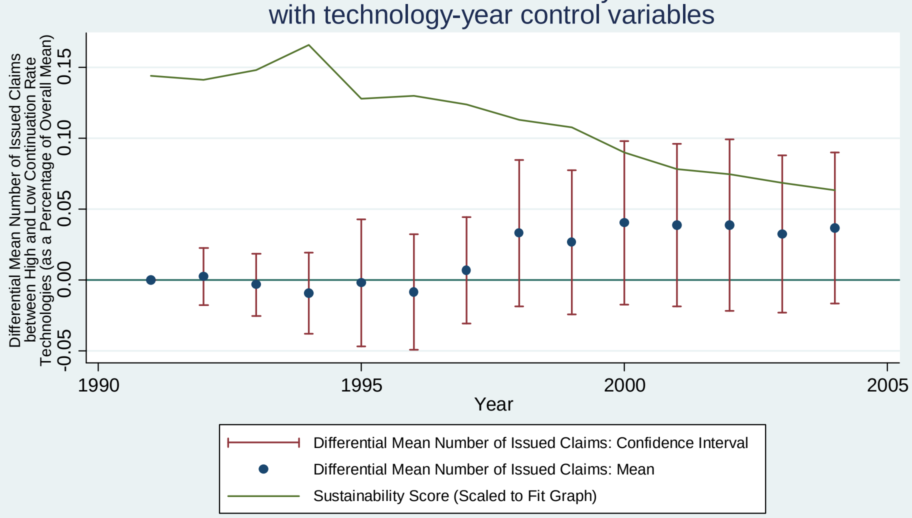

Notes: this graph replicates that of Figure B5, except that it includes technology-year covariates from the NBER, along with controls for initial application filing shares for the given technology.

Figure B8: Time Trend in Differential Issued Claims Quantity (Logged) between Treatment and Control Technologies Overlaid with Sustainability Score with technology-year control variables, using PTO Classes

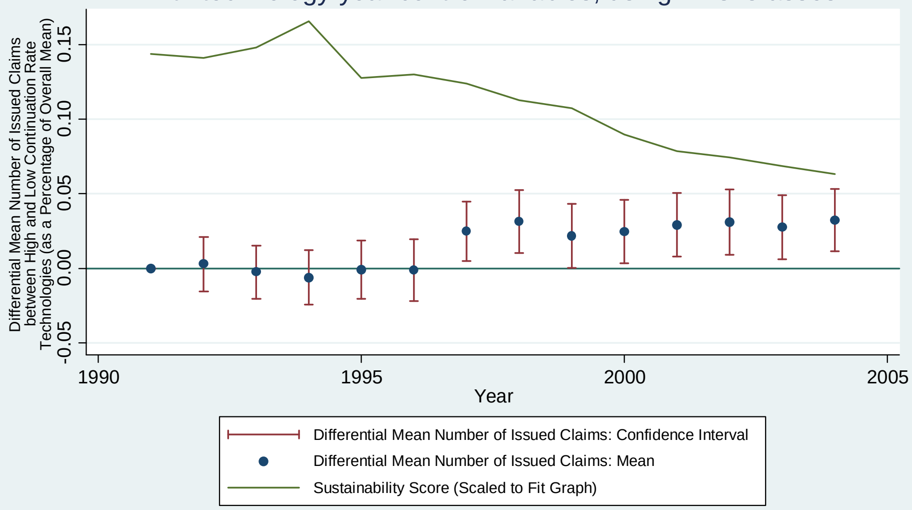

Notes: this graph replicates that of Figure B7, except that it uses PTO Classes to identify technologies.

Figure B9: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Alternative Sustainability Score

Notes: this figure replicates that of Figure 3 in the text, except that the sustainability score captures examination demand using the backlog of all applications awaiting final review, not simply the blacklog of applications awaiting a first office action.

Figure B10: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Dropping Software and Business Methods Patents

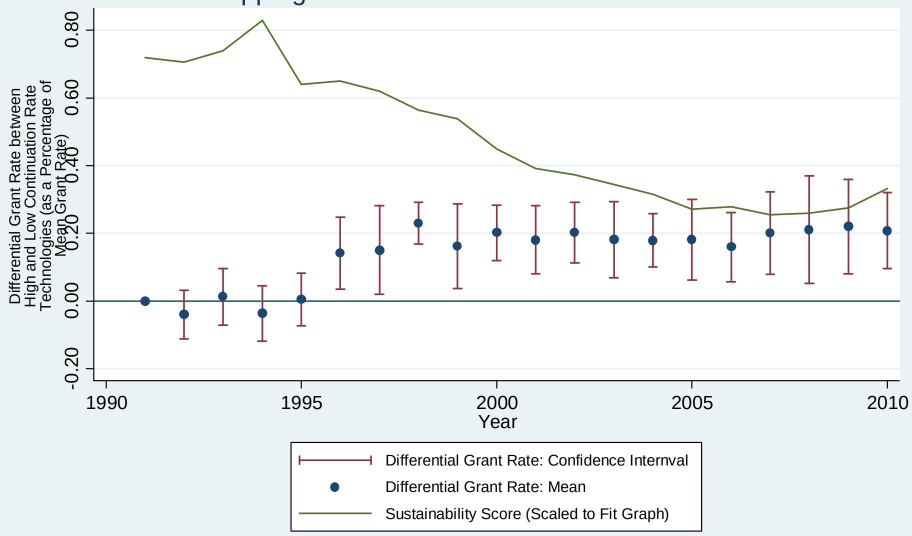

Notes: this figure replactes that of Figure 3, except that it excludes software and business method patents.

Figure B11: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Dropping Software and Business Methods Patents, using PTO Classes

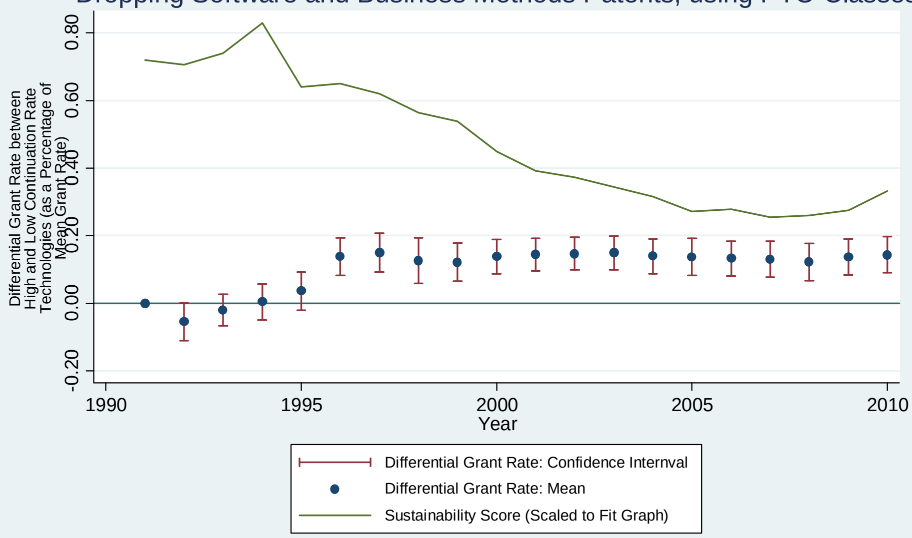

Notes: this figure replactes that of Figure B10, except that it uses PTO Classes to identify technologies.

Figure B12: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
Include RCE Filings in Denominator of Grant Rate

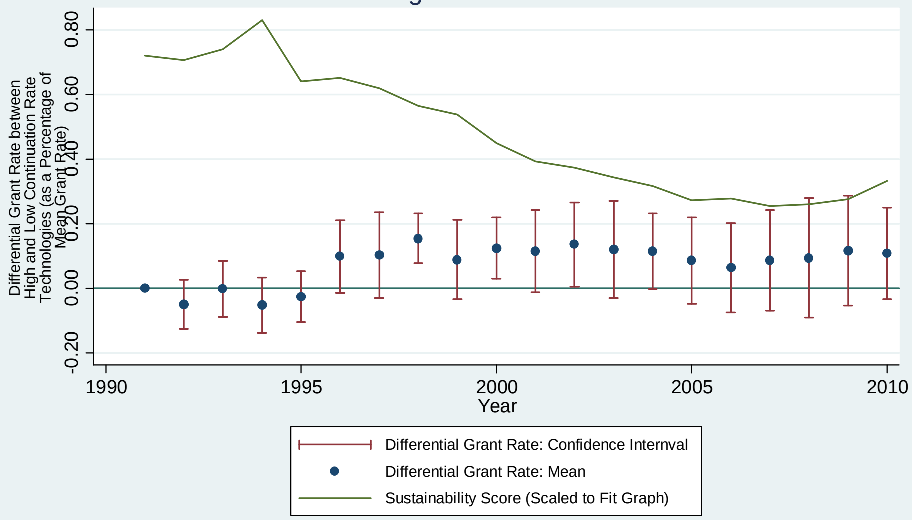

Notes: this figure replactes that of Figure 3, except that it includes request for continued examination filings in the denominator of the grant rate.

Figure B13: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score
Include RCE Filings in Denominator of Grant Rate
Using PTO Classes

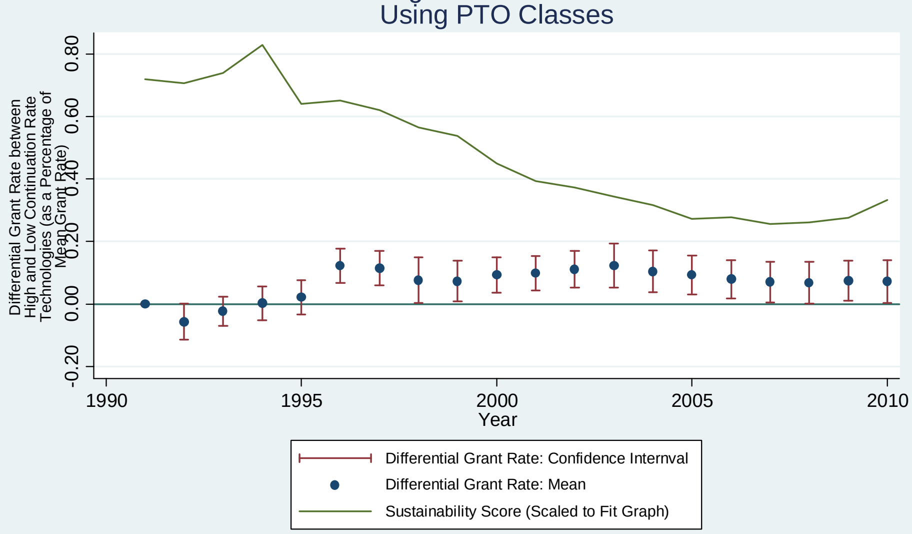

Notes: this figure replactes that of Figure B12, except that it uses PTO Class to identify technologies.

Figure B14: Time Trend in Differential Filings Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score

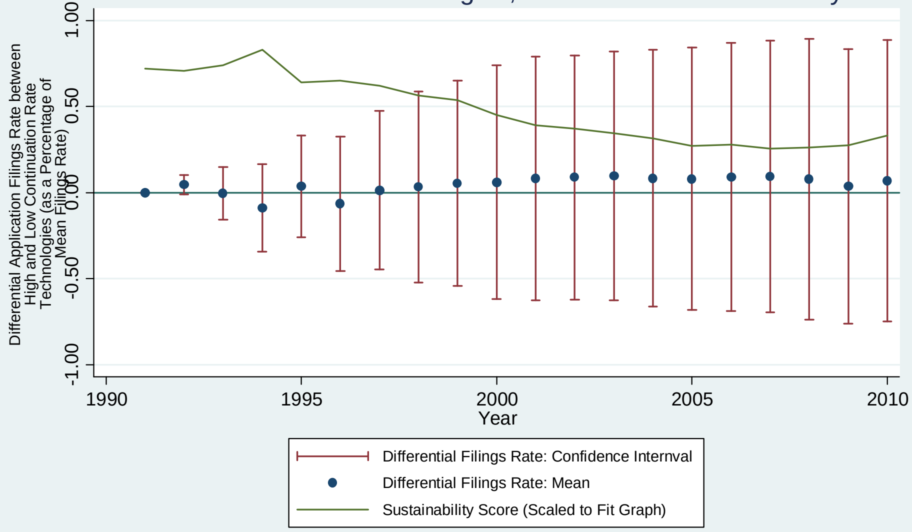

Notes: this figure replactes that of Figure 3, except that it replaces the outcome variable with the differential filings rate (share of original filings for that year attributable to given technology) between the treatment and control technologies.

Figure B15: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Include Controls for Examination Pendency

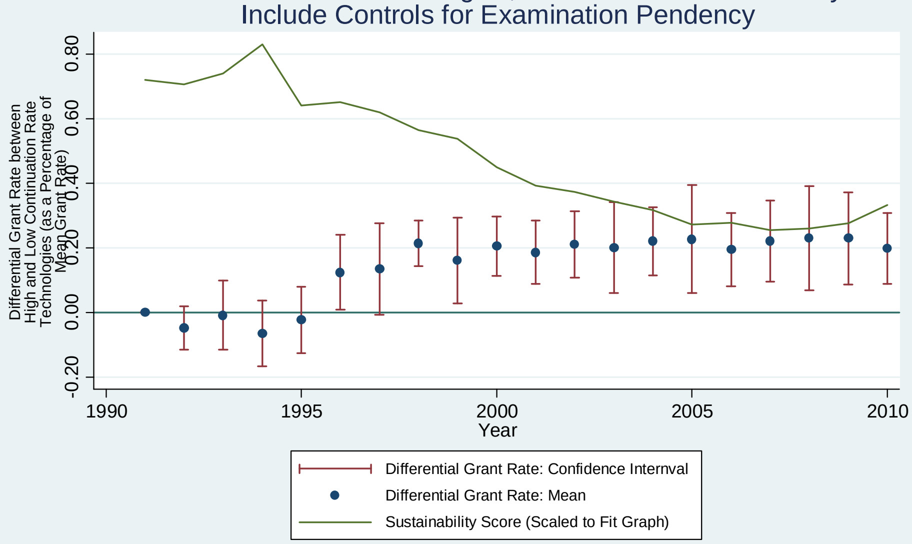

Notes: this figure replicates that of Figure 3, except that it includes a control for the technology-by-year specific examination pendency duration, as measured by the mean time to completion of a first office action.

Figure B16: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Include Controls for Effect of TRIPS

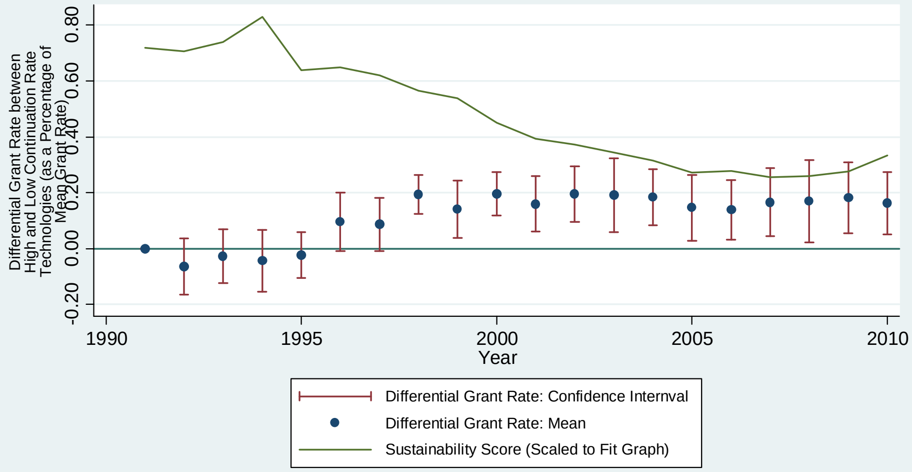

Notes: this figure replicates that of Figure 3, except that it includes controls for a time trend in the differential grant rate between those technologies that received the biggest patent term extension due to TRIPS and those that received the smallest extension due to TRIPS. Simulatenaously fitting this alternative grant-rate trend with the differential grant rate trend between the high- and low-continuation prone technologies, allows us to explore whether the primary results presented in the text can be explained by an alternative story based on the passage of TRIPS.

Figure B17: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Include Controls for Baseline Grant Rates

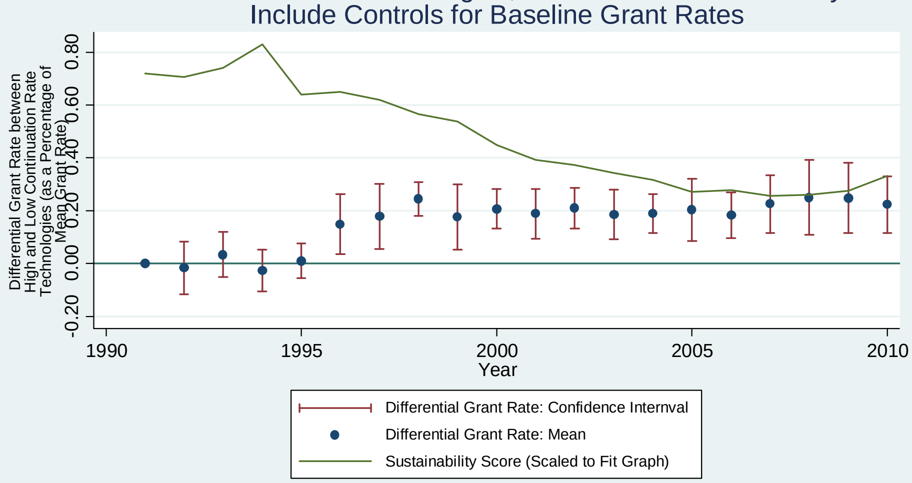

Notes: this figure replicates that of Figure 3, except that it includes controls for a time trend in the differential grant rate between those technologies that started the sample period with initially high grant rates and those technologies that began the sample period with initially low grant rates. That is, we simultaneously fit this alternative year-by-year grant-rate trend with the differential grant rate trend between the high- and low-continuation-prone technologies. Though they are simultaneously fit in the regression, the Figure only presents the estimated trend in the differential grant rate between high- and low-continuation-prone technologies.

Figure B18A: Time Trend in Differential Grant Rate (Logged) between Treatment and Control Technologies, Overlaid with Sustainability Score Including Controls for Maintenance Rate Effects

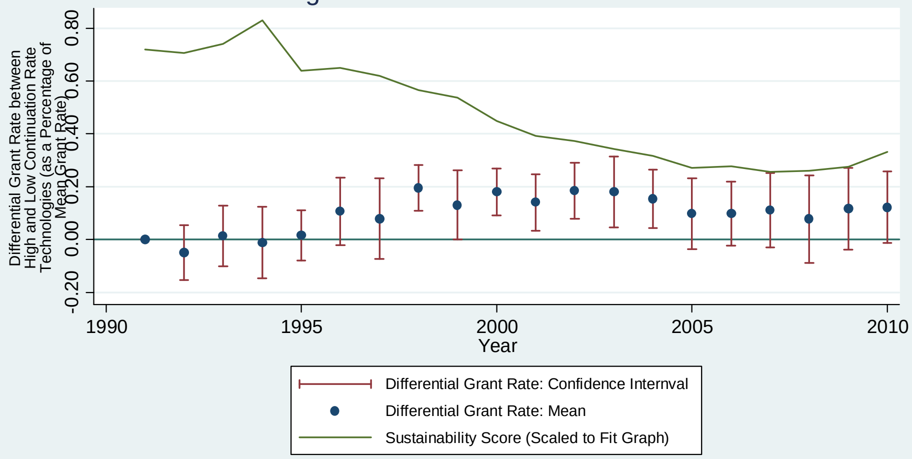

Notes: this figure replicates that of Figure 3, except that it includes controls for a time trend in the differential grant rate between those technologies that started the sample period with initially high levels of four year maintenance rates relative to those with initialy low levels of maintenance rates. That is, we simultaneously fit this alternative year-by-year grant-rate trend with the differential grant rate trend between the high- and low-continuation-prone technologies. Though they are simultaneously fit in the regression, the Figure only presents the estimated trend in the differential grant rate between high- and low-continuation-prone technologies.

Figure B18B: Time Trend in Differential Grant Rate (Logged) between High and Low Maintenance-Rate Technologies, Overlaid with Sustainability Score Including Controls for Continuation Rate Effects

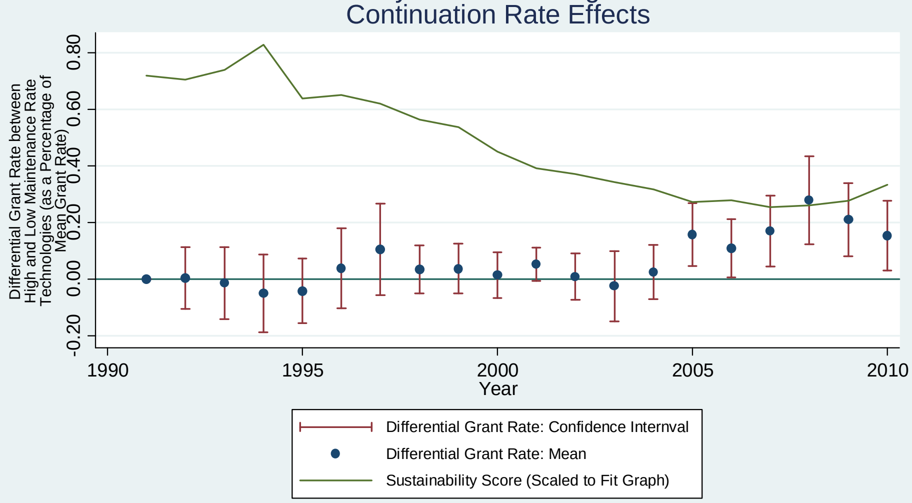

Notes: this figure builds on Figure B18A and presents results from the same regression, except that now, we present the time trend in the differential grant rate between high and low maintenance rate technologies, controlling for the corresponding time trend in the grant rate between high and low continuation-prone technologies.

As discussed in the text, one may be concerned with identifying a technology's disposition towards filing continuations or RCE's by reference to its continuation / RCE filing rate. The essence of the concern is that the continuation rate itself may also be a function of the inherent grant rate (i.e., the grant rate not subject to bias by the PTO) associated with a technology, which may itself be a reflection of the inherent level of application quality for that technology. After all, with more inherent rejections, there may be a greater opportunity to file continuations in the first place. A higher continuation rate based on this enhanced opportunity may not represent a truly higher disposition towards filing continuations. We appease this concern in a couple of ways. First, we actually observe a negative 0.32 correlation between the Agency's grant rate in the baseline year (1990—i.e., the year prior to the beginning of our observation period, which commences in 1991 when the Agency enacts its current funding structure) and the Agency's continuation rate at that time. That is, there is not anything close to a 1-to-1 mapping of continuation rates onto baseline grant rates. Among the technologies with low grant rates at the beginning of the sample, some have high continuation rates and have some have low. As such, with a just simply observation of the data, it does not appear that there is much to be concerned with in regards to the idea that continuation rates are merely picking up variations in inherent grant rates. Nonetheless, to address this concern more directly, we also estimate specifications in which we effectively attempt to separate that portion of continuation rates that are attributable to the baseline grant rate, honing in more specifically on the continuation proclivity itself. To achieve this, we estimate a specification, in which, effectively, we simultaneously fit both the time trend in the differential grant rate across technologies with high and low continuation rates and the time trend in the differential grant rate across technologies with high and low baseline grant rates (we achieve this by including in the specification a set of interactions between each year dummy and the technology-specific continuation rate in 1990, along with a set of interactions between each year dummy and the technology-specific grant rate in 1990). In Figure B17, we present the estimated differential grant rate between high and low continuation rate technologies—i.e., replicating Figure 3 of the text—except that the underlying regression simultaneously fits the second differential time trend mentioned above. As demonstrated, the Figure is virtually unchanged.

On a similar note, one may also be concerned that technology-specific continuation rates are also simply picking up another feature of the relevant technologies that may likewise generate a similar grant-rate response. Primarily, one may be concerned that continuation-rates are correlated with those patent types emphasized in our previous research[7] as being high-fee generators for the Agency. With respect to high-fee-generating patent types relative to low-fee-

[7] *See id.*

generating patent types, the Agency may likewise wish to increase grant rates in response to resource constraints, as we have demonstrated with our previous work. In Figure B18, we rule out a concern in which the present findings are merely picking up the findings of the previous research due to some correlation between the technology-specific continuation rate and the technology-specific maintenance rate. Similar to Figure B17, we estimate a specification in which we simultaneously fit a time trend in the differential grant rate across technologies with high and low continuation rates and the time trend in the differential grant rate across technologies with high and low maintenance rates. We achieve this by including in the specification a set of interactions between each year dummy and the technology-specific continuation rate in 1990, along with a set of interactions between each year dummy and the technology-specific maintenance rate in 1990. As before, while the underlying regression includes all such interactions, Figure B18A presents the results only for the trend in the differential grant rate between high and low continuation rates, while Figure B18B presents the results for the trend in the differential grant rate between high- and low-maintenance-rate technologies.

As is demonstrated by Figure B18A, the increase in the Agency's grant rate for high-continuation-rate technologies following the decline in the Agency's sustainability score in the mid-1990s cannot be explained simply by any correlation between such technologies and their maintenance rates. Collectively, Figures B18A and B18B suggest that the PTO responded to financial declines in the period between the mid-1990s by turning heavily towards granting towards high continuation-prone technologies, which would both cut costs for the Agency and increase fee revenue. Following 2004, the Agency arguably appears to have turned more heavily towards granting at higher rates to high maintenance-rate technologies. At this time, congressional fee diversion policies weakened leaving the Agency with greater ability to retain its fee collections. As hypothesized in our previous work,[8] this 2004 policy switch will leave the Agency especially inclined to turn heavily to fee-generating tactics in order to deal with financial woes, making the idea of granting to high fee-generating patent types such as high maintenance-rate technologies especially appealing. During this latter period, the Agency appears to have relied slightly less heavily on granting at higher rates to high continuation-rate types, though it still appears to be relying on this approach to some degree. The bottom line is that both stories—i.e, a story in which the Agency grants more heavily to high fee-generating patent types and to high repeat-filing-prone patent types—appear to be independently present. In other words, the results presented in this Article and in our previous research do not appear to collapse to just one story or the other.

---

[8] *See id.*

Table A3 presents estimates of specification (3) set forth above. We present the results of the triple-differences coefficient only, indicating the extent to which the primary difference-in-difference results are stronger for the case of large relative to small entity patent applicants. We predict a negative value for this triple-differences coefficient considering: (1) that the predicted coefficient of the primary difference-in-difference is negative in value and (2) the primary difference-in-difference finding is predicted to be stronger for large entities relative to small. After all, a relatively stronger negative relationship would show up as a negative coefficient of this interaction term (while a weaker negative relationship would be identified with a positive interaction term coefficient).

TABLE A3. DIFFERENCE-IN-DIFFERENCE-IN-DIFFERENCE RESULTS

	(1)	**(2)**
Triple Differences Coefficient (coefficient of β_4 from specification (3))	-0.03** (0.01)	-0.05* (0.03)
Grant-Rate Logged	NO	YES

* significant at 10%; ** significant at 5%; *** significant at 1%. Standard errors are reported in parentheses and are clustered to correct for autocorrelation within patent-category / entity-size combinations over time. All regressions include patent-category fixed effects, year fixed effects and entity-size fixed effects, along with technology-by-year fixed effects, technology-by-entity size fixed effects, and entity-size-by-year fixed effects. Regressions are weighted by the number of disposals used to form each observation's grant rate. Data was obtained from the PTO.