

DOES THE U.S. PATENT AND TRADEMARK OFFICE GRANT TOO MANY BAD PATENTS?: EVIDENCE FROM A QUASI- EXPERIMENT

Michael D. Frakes* & Melissa F. Wasserman**

Many believe the root cause of the patent system's dysfunction is that the U.S. Patent and Trademark Office (PTO or Agency) is issuing too many invalid patents that unnecessarily drain consumer welfare. Concerns regarding the Agency's overgranting 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 overgranting 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 the resource-constrained Agency to allow additional patents. Second, this Article attempts to explore, through a sophisticated natural experiment framework, whether the Agency is in

* 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.

** Associate Professor of Law, Richard and Anne Stockton Faculty Scholar, and Richard W. and Marie L. Corman Scholar, University of Illinois College of Law.

The work was funded in part by the University of Illinois at Urbana-Champaign Research Board, Award 12,088, 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, Katherine Strandburg, Saurabh Vishnubhakat, and seminar participants at the 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, DePaul Law School Manzo Patent Law Seminar, Intellectual Property Scholars Conference, and Junior Scholars in Intellectual Property Law Workshop for useful comments.

fact acting on this incentive and overgranting patents. Our findings suggest that the PTO is biased toward allowing patents. Moreover, our results suggest the PTO is targeting its overgranting tendencies toward those patents it stands to benefit from the most—that is, those patent applications directed toward technologies that have historically had high repeat-filing rates, such as information, computer, and health-related technologies. Our findings provide policymakers with much-needed evidence that the PTO is indeed overgranting patents. Our results also suggest that the literature has overlooked a substantial source of Agency bias; hence, recent fixes to improve patent quality will not achieve their desired outcome of extinguishing the PTO’s overgranting proclivities.

INTRODUCTION	615
I. HARMS FROM PTO OVERGRANTING AND INSUFFICIENCY OF EVIDENCE.....	618
A. <i>Overgranting Explanations and Harms</i>	618
B. <i>In Search of Empirical Evidence</i>	621
II. A NOVEL THEORY FOR WHY THE PTO IS BIASED TOWARD ALLOWING PATENTS	625
A. <i>Never-Ending Stream of Repeat Filings</i>	625
B. <i>PTO Objectives</i>	629
III. THE DIFFERENTIAL IMPACT OF THE PTO’S INCENTIVES TO ALLOW PATENTS	631
A. <i>Varying Propensity in the Filing of Continuation Applications</i>	632
B. <i>Per-Claim Allowances Versus Application Allowances</i>	633
IV. DATA AND METHODOLOGY	635
A. <i>Data</i>	635
B. <i>Methodology</i>	637
1. <i>Basic framework</i>	637
2. <i>Expansions beyond basic framework</i>	640
3. <i>Causal assumptions</i>	642
4. <i>Resource sustainability metrics</i>	642
V. RESULTS	644
A. <i>Primary Results: Graphical Analysis</i>	645
1. <i>Description of graphical approach</i>	645
2. <i>Description of findings</i>	649
3. <i>Interpreting the magnitude of the findings</i>	652
4. <i>Related support</i>	653
5. <i>Inclusion of control variables</i>	655
6. <i>Number of technology groups</i>	657
B. <i>Additional Robustness Checks</i>	657
1. <i>Overview</i>	657
2. <i>Substantive patent law changes</i>	658
3. <i>Selection-effects concerns</i>	659
4. <i>Difference-in-difference-in-difference</i>	659
5. <i>Alternative resource-health metric</i>	660
6. <i>Other</i>	661
C. <i>Results: Concluding Remarks</i>	662
VI. THE MECHANISM BY WHICH THE PTO PREFERENCES GRANTING CERTAIN TYPES OF PATENTS OVER OTHERS	663

VII. IMPLICATIONS AND WAYS TO REDUCE THE PTO’S INCENTIVE TO GRANT
 PATENTS 671
 A. *Implications for Patent Policy*..... 671
 B. *Reducing the PTO’s Incentives to Grant Patents in High-Repeat-Filing Technologies*..... 671
 CONCLUSION 675

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 U.S. Patent and 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 overgranting tendencies have been so pressing that they have spurred the Supreme Court to take a renewed interest in substantive patent law,⁴ while likewise driving Congress to enact the first major patent reform act in nearly 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 toward allowing patents, there exists little to no compelling empirical evidence that the Agency is *actually* overgranting patents.⁶ To date, the debate regarding patent policy has been driven by anecdotes of a few infamously issued patents

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); ADAM B. JAFFE & JOSH LERNER, *INNOVATION AND ITS DISCONTENTS: HOW OUR PATENT SYSTEM IS ENDANGERING INNOVATION AND PROGRESS, AND WHAT TO DO ABOUT IT* (2004).

2. See FED. TRADE COMM’N, *TO PROMOTE INNOVATION: THE PROPER BALANCE OF COMPETITION AND PATENT LAW AND POLICY* 5-7 (2003); see also BESSEN & MEURER, *supra* note 1; JAFFE & LERNER, *supra* note 1; *infra* notes 13, 18-22 and accompanying text.

3. See, e.g., FED. TRADE COMM’N, *supra* note 2; 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. Merrill & Anne-Marie Mazza eds., 2006).

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

5. Leahy-Smith America Invents Act, Pub. L. No. 112-29, 125 Stat. 284 (2011) (codified as amended in scattered sections of 35 U.S.C.). 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 notes 26-27 and accompanying text below.

6. See *infra* Part I.B.

or data that alone are 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 overgrant patents or even compelling evidence that the Agency is biased toward 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 describes a novel theoretical source of bias that may cause the Agency to grant 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 reject a patent application with finality. 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 a crushing backlog of over 600,000 patent applications, of which close to forty percent 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 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 upward.

To be clear, repeat filings have 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.

Second, this Article attempts to explore, through a sophisticated methodological design, whether the Agency is in fact operating on this incentive and actually overgranting 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—that is, during times in which it cannot meet its expected examination demand with 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 (FOIA) requests to the PTO. This novel database includes patent processing statistics on all 4,733,263 patent applications filed with the PTO

7. See *infra* Part I.B.

8. As a bonus, by inflating its granting tendencies in this manner, the PTO will also generate additional revenue from postallowance 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, 64 (2004) (detailing applicant misuse of continuation applications).

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; for example, 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 the Agency's 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 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 constitutes 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 subtracting the granting experiences of the control group from the treatment 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 overgranting patents during times in which the Agency lacks sufficient resources to meet its expected demand for examination. Moreover, our findings suggest that the PTO is preferentially granting those patents it stands to benefit the most from—those in high-continuation-filing-rate technologies, such as information and communication technologies, which include software, business methods, and information storage, and health-related technologies, which include surgical and medical instruments and genetics.

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 allows too many bad patents. Importantly, however, it also suggests that policymakers have been barking up the wrong tree. By ignoring the important role that repeat filings play in influencing PTO decisionmaking, legal scholarship has overlooked a substan-

10. Details on the database are provided below in Part IV.

tial source of Agency bias. As a result, the recent “fixes” to the patent system to improve patent quality fail to address the substantial overgranting 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 toward allowing patents in general but also be biased toward allowing particular types of patents. Accordingly, recognition of the harms from the granting of bad patents has been underinclusive. 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. That is, because the PTO is preferentially granting patents in high-continuation-rate technologies, such as information and communication and health-related technologies, society may be overinvesting resources in these technological fields.

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 issues 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 toward allowing patents. Part II theorizes the PTO’s possible incentive to grant patents, introducing a model of agency behavior under resource constraints that predicts that the PTO will act on the incentive to grant additional patents when it faces binding resource constraints—that is, when 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 dataset 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 toward issuing patents.

I. HARMS FROM PTO OVERGRANTING AND INSUFFICIENCY OF EVIDENCE

A. *Overgranting Explanations and Harms*

There is near-universal agreement that the U.S. patent system suffers from rampant shortcomings and failures.¹² Many believe the root cause of the system’s malfunction is that the PTO grants too many invalid patents.¹³

11. See *infra* Part I.A.

12. See sources cited *supra* note 1.

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 that 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 overgranting proclivities.¹⁶ Finally, some have 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

13. See, e.g., BESSEN & MEURER, *supra* note 1; FED. TRADE COMM'N, *supra* note 2; JAFFE & LERNER, *supra* note 1; Roger Allan Ford, *Patent Invalidity Versus Noninfringement*, 99 CORNELL L. REV. 71, 87-88 (2013); Lemley & Sampat, *supra* note 4, at 181; Mark A. Lemley, *Rational Ignorance at the Patent Office*, 95 NW. U. L. REV. 1495, 1495 n.1 (2001); Doug Lichtman & Mark A. Lemley, *Rethinking Patent Law's Presumption of Validity*, 60 STAN. L. REV. 45, 47 n.5 (2007); 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); John R. Thomas, *Collusion and Collective Action in the Patent System: A Proposal for Patent Bounties*, 2001 U. ILL. L. REV. 305, 316-22; R. Polk Wagner, *Understanding Patent-Quality Mechanisms*, 157 U. PA. L. REV. 2135, 2139-45 (2009).

14. NAT'L ACAD. OF PUB. ADMIN., U.S. PATENT AND 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."); Merges, *supra* note 13, at 607 ("Consequently, the only way to earn bonus points with confidence is to allow a patent application."); Thomas, *supra* note 13, at 324-25.

15. JAFFE & LERNER, *supra* note 1, at 130-33 (describing the PTO's budgetary woes); Lemley, *supra* note 13, at 1500 (noting that examiners spend on average only eighteen hours reviewing a patent application).

16. Jeffrey R. Kuester & Lawrence E. Thompson, *Risks Associated with Restricting Business Method and E-Commerce Patents*, 17 GA. ST. U. L. REV. 657, 681 (2001) ("Every new technology presents the PTO with the challenges of creating a sufficient prior art database and channeling the expertise necessary to evaluate the prior art. Internet business method patents are similar, in this respect, to biotechnology and software."); Lucas Osborn, *Tax Strategy Patents: Why the Tax Community Should Not Exclude the Patent System*, 18 ALB. L.J. SCI. & TECH. 325, 357 (2008) ("Mirroring the criticism of business method patents generally, critics assert that patent examiners lack training and expertise to analyze tax law-intensive patent applications."); Kevin M. Baird, Note, *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.")

17. U.S. PATENT & TRADEMARK OFFICE, CORPORATE PLAN—2001: PATENT BUSINESS 23 (2001), available at <http://www.uspto.gov/web/offices/com/corpplan/pt04.pdf>; see also Long, *supra* note 14, at 1967 (positing that the PTO has invited capture by patent applicants in an effort to increase its own stature).

Irrespective of the exact cause, invalidly issued patents—that is, patents that are granted even though the invention fails to meet the patentability requirements—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 seeking to enter markets¹⁹ and stunt further innovation.²⁰ Erroneously issued patents can inhibit the ability of start-ups to obtain venture capital, especially if a dominant player in the market holds the patent in question.²¹ Finally, such patents can also compromise business relations for 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 supracompetitive pricing and diminished quantity. While society may accept such consequences for a properly issued patent, an invalid patent imposes these costs upon society without providing the 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 either were already known or represent only a trivial advancement over existing scientific knowledge.²⁴ As a result, a grant-

18. James Bessen, *The Patent Troll Crisis Is Really a Software Patent Crisis*, WASH. POST SWITCH (Sept. 3, 2013), <http://wapo.st/1F3yMP1>.

19. See FED. TRADE COMM'N, *supra* note 2, at 3 (noting that allowing patents on obvious inventions can thwart competition); Christopher R. Leslie, *The Anticompetitive Effects of Unenforced Invalid Patents*, 91 MINN. L. REV. 101, 119-25 (2006).

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

21. See FED. TRADE COMM'N, *supra* note 2, at 8 (“The threat of being sued for infringement by an incumbent [patent holder]—even on a meritless claim—may ‘scare . . . away’ venture capital financing.” (alteration in original) (quoting *Competition and Intellectual Property Law and Policy in the Knowledge-Based Economy: Hearings Before the Fed. Trade Comm’n* 189 (Feb. 20, 2002) (statement of Josh Lerner, Professor, Harvard Business School), available at http://www.ftc.gov/sites/default/files/documents/public_events/competition-ip-law-policy-knowledge-based-economy-hearings/020220trans.pdf)).

22. See, e.g., *In re Ciprofloxacin Hydrochloride Antitrust Litig.*, 363 F. Supp. 2d 514, 544 (E.D.N.Y. 2005) (noting, in a discussion of *Dow Chemical Co. v. Exxon Corp.*, 139 F.3d 1470 (Fed. Cir. 1998), that “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”); Leslie, *supra* note 19, at 125-27.

23. Mark A. Lemley & Bhaven Sampat, *Examiner Characteristics and Patent Office Outcomes*, 94 REV. ECON. & STAT. 817, 817 (2012).

24. Not surprisingly, the patentability standards reflect a careful balance between encouraging innovation and avoiding drains on consumer welfare. In order for an invention to be patent eligible it must be both new and represent a nontrivial advancement over current scientific understanding. 35 U.S.C. § 103 (2013). If an invention was obvious to the person

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 the judiciary and Congress have attempted to rectify the Agency's perceived overgranting 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 In 2011, 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 Policymakers, however, have been making changes to the patent system absent empirical evidence to help illuminate the actual problems at hand.

B. *In Search of Empirical Evidence*

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

Consider, for instance, the frequently cited statistic that courts invalidate nearly half of all litigated patents that make it to final judgment.²⁸ For this statistic to demonstrate that the PTO allows too many invalid patents to issue, litigated patents must not differ substantially from nonlitigated patents. Yet we

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. Lemley & Sampat, *supra* note 23, at 817.

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

27. See Leahy-Smith America Invents Act, Pub. L. No. 112-29, § 6, 125 Stat. 284, 299-313 (2011) (codified as amended at 35 U.S.C. §§ 311-19, 321-29) (providing for post-grant review proceedings); *id.* § 10, 125 Stat. at 316-20 (codified as amended at 35 U.S.C. § 4) (providing for fee-setting authority); *id.* § 12, 125 Stat. at 325-27 (codified as amended at 35 U.S.C. § 257) (providing for supplemental examination); *id.* § 18, 125 Stat. at 329-31 (codified as amended at 35 U.S.C. § 321) (providing for a transitional program for covered business method patents); H.R. REP. NO. 112-98, pt. 1, at 39-40 (2011) (noting that the primary purpose of the America Invents Act is to "improve patent quality").

28. See, e.g., 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) (suggesting that because half of all litigated patents that make it to final judgment are invalidated, the PTO is issuing too many invalid patents); Ford, *supra* note 13, at 87-89 (same); Leslie, *supra* note 19, at 105-06 (same).

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 overgrants patents.³⁰ It is possible that the Agency issues more patents because it has developed a bias toward overgranting, 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 tells us nothing about how well the PTO scrutinizes 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 overgrants patents (beyond our prior work discussed below) is the groundbreaking research by Mark Lemley and Bhaven Sampat. Lemley and Sampat explored how examiner experience affected the outcomes at the PTO for nearly 10,000 patent applications filed in January 2001.³⁴ They found that more experienced examiners have a higher grant rate than their junior colleagues and that the examiner grant rate monotonically increases as a function of examiner experience.³⁵ Importantly, Lemley and Sampat provided some evidence that senior examiners may allow too many patents while junior examiners grant too few.³⁶ By showing that there is at least some group within the

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 as compared with nonlitigated patents. John R. Allison et al., *Valuable Patents*, 92 GEO. L.J. 435, 438 (2004).

30. See, e.g., John M. Golden, *Proliferating Patents and Patent Law's "Cost Disease"*, 51 HOUS. L. REV. 455, 457-58 (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 light of economic growth).

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

32. See, e.g., JAFFE & LERNER, *supra* note 1, at 142 (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., Lemley & Sampat, *supra* note 4, at 186.

34. See Lemley & Sampat, *supra* note 23, at 817.

35. *Id.* at 821-22.

36. For instance, they found that senior examiners cite less prior art, 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 820-22, 824.

Agency that perhaps allows too many patents, their analysis necessarily suggests that the Agency as a whole may allow 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 (as our analysis below does) or causally link their results to a particular feature of a given examiner. Given that their analysis only looks at a single snapshot in time, their methodological design (as they admit) cannot completely determine whether their results are 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.

One of the few studies that has provided compelling empirical evidence that the PTO overgrants patents is our prior work using a natural experiment methodology to explore the influence of the Agency's fee structure on its decision to grant patents.³⁸ Because the PTO garners over fifty percent of its patent operating budget through postallowance fees—for example, fees the Agency collects only if it grants patents—we posited that a financially constrained PTO might grant additional patents in an effort to generate additional funds.³⁹ Employing the same novel patent data utilized in this Article, we examined whether the Agency did in fact allow additional patents during times of financial distress.⁴⁰ Our results suggest that the back-end fee structure of the Agency biased a financially constrained PTO toward allowing patents.⁴¹ Moreover, our find-

37. Although Lemley and Sampat estimated empirical specifications suggesting that their results are not impacted by whether or not the given examiner departs the Agency within five years, *id.* at 824-25, this alternate 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. Michael D. Frakes & Melissa F. Wasserman, *Is the Time Allocated to Review Patent Applications Inducing Examiners to Grant Invalid Patents?: Evidence from Micro-Level Application Data* (Nat'l Bureau of Econ. Research, Working Paper No. 20,337, 2014), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2472794.

38. Michael D. Frakes & Melissa F. Wasserman, *Does Agency Funding Affect Decisionmaking?: An Empirical Assessment of the PTO's Granting Patterns*, 66 *VAND. L. REV.* 67, 96 (2013).

39. *Id.* at 79-80.

40. *Id.* at 84-85.

41. *Id.* at 119.

ings suggest that the Agency was targeting its distortionary granting practices on the types of patents that it profited the most from allowing—patent applications filed by large entities and patent applications of the type that historically had high renewal rates.⁴²

To be clear, the behaviors under investigation in the present Article differ from those studied in our prior research and stem from distinct underlying motivations on the part of the PTO—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 overgrant 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 overgranting biases identified and studied in these respective investigations are empirically independent of each other.⁴³

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 actually cause 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 for curbing the PTO’s overgranting proclivities should start with targeting the sources of the problem, policymakers are in dire need of studies that use convincing empirical strategies. Moreover, given the magnitude of harms associated with a potential overgranting 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 toward granting patents, along with as much convincing evidence as possible on the sources behind such biases. To that end, this Article utilizes a strategy specifically designed to allow causal inferences to be drawn and to discern whether particular features of the Agency actually

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

43. Michael D. Frakes & Melissa F. Wasserman, *Does the U.S. Patent and Trademark Office Grant Too Many Bad Patents?: Evidence from a Quasi-Experiment: Online Appendix*, STAN. L. REV. 22 fig.B18A, 23 fig.B18B (Mar. 2015), http://www.stanfordlawreview.org/sites/default/files/Frakes_Wasserman_67_Stan_L_Rev_Online_Appendix.pdf [hereinafter *Online Appendix*].

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

cause the PTO to overgrant 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 of how the PTO's inability to dispose of a patent application with finality may bias the resource-constrained Agency toward allowing patents.

II. A NOVEL THEORY FOR WHY THE PTO IS BIASED TOWARD 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 Part demonstrates that the 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 an application. That is, an aggrieved patent applicant can always choose to start the examination process over by filing a repeat application. Repeat applications generally fall in one of two categories: continuation applications and requests for continued examination (RCEs). While a continuation application is technically a new application and an 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 To illustrate the incentives posed by repeat applications,

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 patents protect the way an article is used and works. *See* 35 U.S.C. § 101 (2013).

46. Technically, a patent applicant can file either a continuation application, 35 U.S.C. § 120 (allowing continuations to use the filing date of the original patent application), or an RCE, 35 U.S.C. § 132(a)-(b). *See* 37 C.F.R. § 1.114(a) (2014). Although both an RCE and 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* MPEP § 706.07(h) (9th ed. Mar. 2014) ("An RCE is not the filing

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 Part 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,

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)"); Lemley & Moore, *supra* note 9, at 68 n.14. As these differences are immaterial for the purposes of this Article, we refer to continuation applications and requests for continued examination collectively as RCEs. The fact that an RCE is simply a request to continue the examination of the original application means that an 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* MPEP, *supra*, § 708.

47. Claims represent the scope of the invention an applicant seeks to patent. *See* 35 U.S.C. § 112 (“The specification shall conclude with one or more claims particularly pointing 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 onward, 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., no. 2, ¶ 24, <http://journals.law.stanford.edu/sites/default/files/stanford-technology-law-review-stlr/online/lemley-sampat-examining-patent.pdf> (“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. Changes to Practice for Continued Examination Filings, 72 Fed. Reg. 46,716 (Aug. 21, 2007) (codified at scattered sections of 37 C.F.R.). However, the Agency ultimately rescinded the regulations amidst court challenges. *See* Press Release, U.S. Patent & Trademark Office, USPTO Rescinds Controversial Patent Regulations Package Proposed by Previous Administration (Oct. 8, 2009), https://web.archive.org/web/20091012011042/http://www.uspto.gov/news/09_21.jsp (accessed via the Internet Archive index); *see also* *Tafas v. Dudas*, 511 F. Supp. 2d 652, 671 (E.D. Va. 2007) (granting a preliminary injunction preventing the PTO from implementing changes to the continuation practice on the eve of the rule’s implementation); *Tafas v. Dudas*, 541 F. Supp. 2d 805, 817 (E.D. Va. 2008) (granting summary judgment against the PTO); *Tafas v. Doll*, 559 F.3d 1345, 1356-57, 1359-61 (Fed. Cir. 2009) (holding on appeal that both the claim and the continuation rules were procedural in nature and within the Agency’s rulemaking authority and that the continuations rule was inconsistent with patent law); *Tafas v. Doll*, 328 F. App’x 658, 658 (Fed. Cir. 2009) (granting petition to rehear the case en banc and vacating the panel opinion); *Tafas v. Kappos*, 586 F.3d 1369 (Fed. Cir. 2009) (en banc) (dismissing the appeal as moot after the PTO rescinded the rules).

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 fiscal year 1980 to 40% in fiscal year 2012.⁵⁰ Thus, repeat filings constitute a substantial portion of the 600,000 applications currently awaiting substantive review.⁵¹ 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.”⁵² In fact, the Agency’s long patent pendency times, which in fiscal year 2012 topped thirty-two months,⁵³ have led at least one commentator to quip that the PTO has become the “burial ground” for patent applications.⁵⁴

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.⁵⁵ If the PTO collected enough in examination fees to cover the costs associated with reviewing an application, any uptick in the application filing rate could theoretically be addressed by expanding the Agency’s examination capacity.⁵⁶ This, however, is not the case. While thirty percent of the

50. These numbers were calculated from our patent data obtained from the PTO through FOIA requests, which include information on which application filings were original versus repeat.

51. See U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT FOR FISCAL YEAR 2012, at 17 (2012) [hereinafter 2012 PERFORMANCE AND ACCOUNTABILITY REPORT] (noting that 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. See Omnibus Budget Reconciliation Act of 1990, Pub. L. No. 101-508, § 10101, 104 Stat. 1388, 1388-91 (establishing a structure to fund the PTO through user fees). In the early 1990s, approximately 15% of patent filings comprised repeat applications; in the late 2000s, this number skyrocketed to 40%. These figures were derived from the data that we collected from the PTO and that are discussed in further detail in Part IV below.

52. Changes to Practice for Continued Examination Filings, 72 Fed. Reg. at 46,718.

53. 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 51, at 3.

54. Terry Carter, *A Patent on Problems*, A.B.A. J. (Mar. 1, 2010, 11:40 AM CST), http://www.abajournal.com/magazine/article/a_patent_on_problems.

55. See 35 U.S.C. § 41(a)(3) (2013). Technically, an applicant pays filing, search, and examination fees (collectively referred to as examination fees). During the time of this study, the fees for examining an RCE were set below the examining fees for a new application, and the fees for examining a continuation application were the same as for an original application. 37 C.F.R. § 1.16(a)(1) (2011) (setting the basic filing fee at \$330 for a large entity and \$165 for a small entity); *id.* § 1.16(k) (setting the utility search fee at \$540 for a large entity and \$270 for a small entity); *id.* § 1.16(o) (setting the utility examination fee at \$220 for a large entity and \$110 for a small entity); *id.* § 1.17(e) (setting the examination fees for an RCE in fiscal year 2010 at \$810 for a large entity and \$405 for a small entity). The PTO defined a “small” entity as any individual, nonprofit corporation, or corporation that qualifies as a small business under the Small Business Act. *Id.* § 1.27(a)(1)-(3).

56. There are, however, questions as to whether the PTO would actually be able to hire enough new patent examiners. See, e.g., U.S. GOV’T ACCOUNTABILITY OFFICE, GAO-07-

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.⁵⁷ 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 altogether,⁵⁸ the PTO could attempt to combat the logjam of applications by decreasing the incentives of applicants to file repeat applications in the first instance. One way to do so, conceivably, is through its granting practices. By allowing additional patents early in the examination process, the Agency extinguishes the incentive of patent applicants to refile.⁵⁹ That is, by biasing its grant rate in an upward 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.⁶⁰ The extent to

1102, U.S. PATENT AND TRADEMARK OFFICE: HIRING EFFORTS ARE NOT SUFFICIENT TO REDUCE THE PATENT APPLICATION BACKLOG (2007).

57. See *Setting and Adjusting Patent Fees*, 77 Fed. Reg. 55,028, 55,034 (proposed Sept. 6, 2012) (stating that seventy percent of the PTO's patent operational costs stem from examination expenses). The PTO estimated that the average cost of examining a patent application in fiscal year 2011 was approximately \$3600. U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT, FISCAL YEAR 2011, at 17 (2011) [hereinafter 2011 PERFORMANCE AND ACCOUNTABILITY REPORT] (stating that the average patent cost \$3594 to examine in fiscal year 2011). The examination fee in 2011 for applicants other than small entities, however, was set at only \$1090—a \$330 basic filing fee, 37 C.F.R. § 1.16(a)(1), a \$540 utility search fee, *id.* § 1.16(k), and a \$220 utility examination fee, *id.* § 1.16(o).

58. See *supra* note 49.

59. 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 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 practice this latter tactic. See 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). In any of these scenarios, a patent applicant may file an RCE even if the PTO allows all of the claims because the applicant may want to wait until some later time, when it has acquired additional information, to modify its patent claims and obtain a patent.

60. 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 by the Agency to review applications, the mismatch between fees and costs grows each time a patent application is rejected and then refiled. Melissa F. Wasserman, *The PTO's Asymmetric Incentives: Pressure to Expand Substantive Patent Law*, 72 OHIO ST. L.J. 379, 409-10 (2011). 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 those applications. Compare 37 C.F.R. § 1.17(e) (setting RCE examination fees for fiscal year 2010 at \$810 for a large entity and

which the PTO would act on this inducement depends, in part, on the objectives of the Agency and its needs.

B. *PTO Objectives*

We contend that high-level officials at the PTO are largely mission minded.⁶¹ The Agency will first and foremost attempt to faithfully carry out its mission of promoting innovation by providing both timely and high-quality examination of patent applications.⁶² Because of the inadequacies of patent examination fees, the PTO must subsidize patent examination through other activities that relate to 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.⁶³

\$405 for a small entity), with U.S. Patent & Trademark Office, Detailed Appendices: Patent Fee Proposal 61 (Feb. 7, 2012), http://www.uspto.gov/sites/default/files/aia_implementation/fee_setting_-_ppac_hearing_appendices_7feb12.pdf (estimating the historical cost of examining an RCE to be \$1696). 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 costly repeat filings. As a bonus, by increasing its inclination to allow patents in this manner, the PTO will also garner additional fee income. After all, over fifty percent of the Agency's budget historically has been garnered through postallowance fees—that is, fees that the PTO only collects if it grants a patent. 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 51, at 72 (stating that approximately 53.4% of total patent income comes from maintenance fees and issue fees); WORKING FOR OUR CUSTOMERS, *supra* note 31, at 59 (showing 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 decreases the Agency's average cost of examining an application, but also increases the fees the PTO collects per application processed.

61. This assumption is supported by our previous work, wherein we found that the PTO distorted its grant rate in an effort to raise additional revenue only when the Agency was financially constrained. Thus, our previous work suggests that the PTO's objectives are more likely to align with those of an agency that is mission minded rather than budget maximizing. Frakes & Wasserman, *supra* note 38, at 75.

62. 2013 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 31, at 3 (noting that the PTO's "mission is to foster innovation . . . by delivering high-quality and timely examination of patent . . . applications").

63. Patents do not automatically remain enforceable for their entire twenty-year term; instead, they must be renewed to remain enforceable. *Cf.* 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 51, at 73 (noting that renewal fees "recoup costs incurred during the initial patent process"). These postallowance fees, which comprise over fifty percent of the Agency's patent budget, *see id.* at 72 (stating that approximately 53.4% of total patent income comes from maintenance fees and issue fees); WORKING FOR OUR CUSTOMERS, *supra* note 31, at 59 (showing that approximately 54% of total patent income comes from maintenance fees and issue fees), are typically larger than the examination fees. In fiscal year 2011, the issuance fee was \$1510, and the maintenance fees, which are due 3.5, 7.5, and 11.5 years from the date the patent issues, were \$980, \$2480, and \$4110, respectively.

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

At other times, however, the Agency may lack the resources—that is, 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.⁶⁴ Although the PTO could lobby Congress for supplemental funding or increased fee levels,⁶⁵ 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, it 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

See 37 C.F.R. § 1.18(a) (2011) (specifying the issuance fee); *id.* § 1.20(e)-(g) (specifying the maintenance fees). Again, small entities pay half these amounts. More importantly, the expenses associated with issuing and maintaining a patent are minimal. U.S. GEN. ACCOUNTING OFFICE, GAO/RCED-97-113, INTELLECTUAL PROPERTY: 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 postexamination fees to fund the examination of patent applications.

64. 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 toward 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 postallowance 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.

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 examination expenses.

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

extreme pressure to decrease its mounting backlog of patent applications.⁶⁶ 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.⁶⁷

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 sacrifices some patent quality—that is, the Agency allows some patents that should be denied. A PTO that is concerned with decreasing its backlog, however, 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 are highly visible, easily measured metrics, whereas the Agency's nonbiased grant rate is difficult to measure.⁶⁸

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

The previous Part posited that the PTO can reduce repeat filings and concomitantly decrease (or at least slow the growth of) its backlog by allowing additional patents. In order to better understand when the Agency would act on this bias toward granting patents, it also introduced a model of agency behavior: a mission-minded but resource-constrained PTO. This Part 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

66. See 2012 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 51, at 17 (noting that the PTO has a backlog of 608,283 patent applications); see also *Commerce, Justice, Science, and Related Agencies Appropriations for 2012: Hearings Before the H. Comm. on Appropriations Subcomm. on Commerce, Justice, Sci., & Related Agencies*, 113th Cong. 3-4 (2013) (statement of Todd Zinser, Inspector Gen., U.S. Department of Commerce) (noting that addressing the patent backlog is one of the top challenges facing the Department of Commerce); *Inspector General's Top Management Challenges Facing the USPTO*, U.S. PAT. & TRADEMARK OFF., http://www.uspto.gov/about/stratplan/ar/2011/oai_01.html (last modified Jan. 3, 2012) (noting the management challenge of reducing the patent application backlog).

67. U.S. PATENT & TRADEMARK OFFICE, PERFORMANCE AND ACCOUNTABILITY REPORT, FISCAL YEAR 2008, at 33 (2008) [hereinafter 2008 PERFORMANCE AND ACCOUNTABILITY REPORT] (noting that 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 USPTO: Fighting Piracy and Counterfeiting by Protecting Intellectual Property Rights*, U.S. PAT. & TRADEMARK OFF. (Nov. 2, 2005), http://www.uspto.gov/about/stratplan/ar/2005/02_message_director.jsp (noting that "the volume and complexity of patent applications continues to outpace current capacity to examine them" and that the PTO has a "backlog of historic proportions").

68. See, e.g., Christi J. Guerrini, *Defining Patent Quality*, 82 *FORDHAM L. REV.* 3091 (2014) (arguing that there is no clear definition of patent quality).

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 Part examines how the propensity of applicants' continuation filings bears on the PTO's ability to reduce repeat filings (and, concomitantly, its backlog of patent applications) by issuing additional patents in technologies that have historically had high repeat-filing rates.

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 than others to be refiled as continuations. The PTO stands to gain more by granting patents to these types of applications relative to those that are more likely to be abandoned upon rejection. 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.⁶⁹

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 Part, the PTO may inflate its grant rate in an effort to diminish repeat filings and hence concomitantly its backlog of patent applications. However, rather than increase patenting across the board, a resource-constrained Agency may grant relatively more patents in technologies with high continuation-filing rates than patents in technol-

69. The PTO may assign these categorical likelihoods using a relatively coarse classification of technology types (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 with 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 application's patent class.

It is also well known that continuation filing rates vary across technology. *See* Lemley & Moore, *supra* note 9, at 86-88 (documenting the different use of repeat filings across technologies). Furthermore, the former supervisory patent examiners we interviewed were aware that repeat filings varied substantially across technologies. *See infra* note 129 and accompanying text.

ogies with low continuation-filing rates, as the Agency stands to decrease repeat filings more by granting the former over the latter.⁷⁰ 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 reduce repeat filings would prefer to satisfy this goal by granting additional 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 in which this reduction in repeat filings will be smaller (i.e., those with low continuation-filing rates). This analysis leads to the following testable hypothesis.

Hypothesis 1: *Upon deterioration in the PTO's resource condition, the Agency will begin to grant at a relatively higher rate patents within technology categories that generally have high continuation-filing rates in comparison to patents within categories that generally have low continuation-filing rates.*

B. *Per-Claim Allowances 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 relax this assumption and consider the situation in which the PTO finds that only a subset of the proposed claims or scope of the invention meets 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 receipt of a notice of allowance.⁷¹ This may happen in situations in which at least some of the claims were rejected. That is, a patent applicant may elect for the subset of allowed claims to matriculate into an issued patent and choose to refile, rather than abandon, the subset of the proposed claims that has been rejected. The rest of this Subpart 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 nonbiased 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 filing.⁷² Ultimately, this scenario is but a simple extension of the binary

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

71. See *supra* note 48.

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

model discussed above. Because a resource-healthy PTO would have rejected the patent application in question, any response in this fashion will distort the Agency's grant rate in an upward direction.⁷³ Thus, Hypothesis 1 will capture the Agency's distortionary response in this scenario.

Now, however, consider the previously unaddressed situation in which the Agency's nonbiased application of patentability standards would result in a portion, though less than all, of the proposed claims being allowed.⁷⁴ In this instance, we contend that the PTO may respond to constrained resources by allowing some of the application's claims 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 reducing the Agency's backlog of patent applications. Note that Hypothesis 1 will not capture this distortionary response. Although a resource-constrained PTO may still attempt to reduce repeat filings by allowing additional claims, this act will no longer inflate its grant rate; a resource-healthy PTO would have allowed some of these claims anyway, in which case this disposition would have already been reflected in the grant-rate calculation. Accordingly, this scenario leads to the second testable hypothesis.⁷⁵

Hypothesis 2: Upon deterioration in the PTO's resource condition, the Agency will begin to allow a relatively higher number of claims for patents within technology categories that generally have high continuation-filing rates in comparison to patents within categories that generally have low continuation-filing rates.

73. 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 to allow not just one narrow claim but as many claims (or breadth of patent scope) as necessary to discourage the refiling of the application.

74. We recognize that the number of claims is only a 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.

75. 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 number of claims 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, it is necessary not only to collect technology-specific data on the PTO's patent-processing practices, but also 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 onward.⁷⁶ 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. We limit the analysis to the period from 1991 onward, as 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. As such, through the filing of a series of FOIA requests, we collected previously unavailable patent processing statistics on all 4,733,263 patent applications filed with the PTO from 1991 to 2010.⁷⁷ These data include information on such metrics (among others) as the number of allowances, disposals,⁷⁸ initial filings, continuation filings, and RCE filings for each technolo-

76. See, e.g., Lemley & Sampat, *supra* note 4, at 187.

77. Our FOIA request was for the following information, on a yearly basis from 1979 to the present, broken down by patent class: (1) the number of patent applications filed; (2) the number of patent applications filed by small entities; (3) the number of patent applications allowed; (4) the number of patent applications filed by small entities that were allowed; (5) the number of patent application disposals; (6) the number of patent applications filed that were continuations or RCEs; (7) the number of patent applications filed by small entities that were continuations or RCEs; and (8) the number of patent applications filed by small entities that were disposed of.

78. 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 include a finally rejected patent application if an applicant subsequently filed 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 constructed 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. *Online Appendix*, *supra* note 43, at 6.

gy-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 disposed of—that is, granted or abandoned—by the PTO.⁷⁹ Under this calculation, the PTO's grant rate was roughly seventy 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 (NBER).⁸⁰

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, 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.⁸¹ This repeat-filing-proclivity (or continuation-proclivity)

79. See, e.g., EUR. PATENT OFFICE ET AL., FOUR OFFICE STATISTICS REPORT, 2010 EDITION 77 (2011), available at <http://www.trilateral.net/statistics/tsr/fosr2010/fullreport.pdf>.

80. See Bronwyn H. Hall et al., *The NBER Patent-Citations Data File: Lessons, Insights, and Methodological Tools*, in PATENTS, CITATIONS, AND INNOVATIONS: A WINDOW ON THE KNOWLEDGE ECONOMY 403, 452 tbl.9 (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. We note that the Hall et al. classification scheme is the most commonly employed in empirical patent scholarship. As demonstrated in the Online Appendix, however, the results are nearly identical when using regression specifications based on the PTO classifications themselves. *Online Appendix, supra* note 43, at 14 fig.B4, 16 fig.B6, 17 fig.B8, 18 fig.B11, 19 fig.B13. Thus, the results appear to be insensitive to the precise classification scheme employed, suggesting that such results may be robust to alternative schemes not considered by the authors. 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 NBER subcategories, as is borne out by the data.

81. See *supra* Part III.

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 sample period.⁸² In the Online Appendix, we demonstrate the general robustness of the findings presented below to various alternative allocation schemes.⁸³

Note that for the purposes of calculating a continuation rate at the beginning of the sample 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 the thirty-seven technology categories.⁸⁴ 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 issues 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 Article to broadly challenge the literature to date for failing to adequately support the proposition of excessive PTO granting through convincing empirical strategies, it is beyond the scope of this Article to attempt to rectify these methodological deficiencies with respect to all possible sources of PTO granting bias. Rather, we will 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

82. This beginning-sample-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.

83. *Online Appendix*, *supra* note 43, at 7, 13 figs.B1 & B2, 14 fig.B3.

84. *Id.* at 10 tbl.A1.

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 alternate system. As such, the second obstacle in proving that the PTO issues too many patents requires the construction of a counterfactual environment meant to come as close as possible to replicating the 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—for example, variations in average application quality or general economic conditions. Possible developments of this latter nature simply render it difficult to isolate our story of interest.⁸⁵

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 for 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—for example, as a result of trends in underlying application quality. In other words, if we are concerned that simply

85. 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.

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 arises 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” calculation⁸⁶ 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 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—a shock that leaves the PTO’s fee revenue *insufficient* to enable the Agency 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 its grant rates during times of resource strain. Similarly, we can view those low-continuation-propensity applicants as our “control” group—that is, 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 those applicants should be a reflection of other time-varying factors that likewise bear on observed grant rates—factors other than the Agency’s tendency to inflate its grant rate during times of resource strain. 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 secure a targeted estimate of the relationship of interest—that is, of the

86. For an overview of difference-in-difference estimation, see JOSHUA D. ANGRIST & JÖRN-STEFFEN PISCHKE, *MOSTLY HARMLESS ECONOMETRICS* 227-42 (2009).

relationship between the PTO's resource condition and the extent to which it grants applications.⁸⁷

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 different continuation proclivities—is a frequently used causal-inference technique in the empirical economics and law and economics literatures.⁸⁸ The empirical specifications that we estimate below, however, are slightly richer than a basic difference-in-difference framework. We refer the reader to the Online Appendix for a discussion regarding the intricacies of this richer design.⁸⁹

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

87. 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 (the continuation rate at the beginning of the sample period) across technologies does not closely match up with the distribution of baseline (i.e., at the beginning of the sample period) grant rates across technologies. See *Online Appendix, supra* note 43, at 23. 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 refer 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. See *id.* at 21 fig.B17, 23-24.

88. ANGRIST & PISHCKE, *supra* note 86, at 227-42.

89. *Online Appendix, supra* note 43, at 1-3.

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 might 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 experienced a negative resource shock. However, with respect to another particular high-continuation-rate technology, again by chance alone, we could observe a relative decrease 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.⁹⁰

90. 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 by another—that is, the most basic difference-in-difference framework. In effect, such analyses draw upon only four 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 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 United States to form the control analysis. David S. Abrams & R. Polk Wagner, *Poisoning the Next Apple? The America Invents Act and Individual Inventors*, 65 *STAN. L. REV.* 517 (2013). Their analysis demonstrates the substantial challenges associated with a research design of this basic nature. Contemporaneous with Canada's move to a first-to-file system was, among other things, its adoption of a deferred examination system, wherein parties could file applications but request that the Canadian Patent Office not review them until a later period, and its implementation of maintenance fees assessed upon patentees over time. 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—that is, the move to a first-to-file system. *Id.* at 553-59. Abrams and Wagner, of course, acknowledge the limitations of their difference-in-difference approach and carefully present separate empirical markers that they contend are more consistent with the first-to-file story than the deferred-examination or maintenance-rate story. *See id.* at 557-59.

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.⁹¹ The primary concern with such an approach is that 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 already suggested, the fact that we evaluate this story over thirty-seven distinct technology groups—not just two 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, in order to infer causation, we are left with a potentially milder assumption—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, which might otherwise explain any divergent 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 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 when 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 more likely be resource constrained upon various *aggregate* developments, including (1) an increase in the Agency’s backlog of examinations awaiting review relative to the number of existing patents from which the Agency collects renewal fees, (2) a decrease in the rate at which existing patent holders renew their patents, (3) an increase in the average complexity of its examinations, and (4) a decrease in the incidence of large-

91. Alberto Abadie, *Semiparametric Difference-in-Differences Estimators*, 72 REV. ECON. STUD. 1, 1-2 (2005) (“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.”).

entity applicants.⁹² Each such development would disrupt any balance between the incoming fees and the examination costs demanded of the Agency, leaving the PTO more likely to be resource constrained. Consider for instance the first factor, which our data suggest is likely the most critical on the list. That is, consider what happens when the number of applications awaiting review by the Agency grows at a rate substantially higher than the stock of existing patents from which the Agency collects maintenance fees. Since each incoming application comes with considerable examination costs, and since examination fees are insufficient to cover these examination costs, the PTO is heavily reliant upon maintenance fees to stay afloat.⁹³ As such, to the extent that the rate of growth of those patents from which the Agency collects maintenance fees lags behind the rate of growth in incoming applications (or, in a more cumulative sense, the rate of growth of its backlog), the user-fee-funded PTO will face concerns about its ability to generate enough resources to cover this growing workload.

Using data on annual fluctuations over time in those various factors identified in the preceding paragraph, along with information on the 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 (as our data suggest), the backlog factor will be given an appropriately higher weight in the determination of this score. More specifically, using data on these factors (annual maintenance rates, backlog levels, incidence of large-entity applicants, etc.), we simulate an annual measure equal to the ratio between (1) the issuance and postissuance 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.⁹⁴

92. *Online Appendix, supra* note 43, at 7-9.

93. See Michael D. Frakes & Melissa F. Wasserman, *The Failed Promise of User Fees: Empirical Evidence from the U.S. Patent and Trademark Office*, 11 J. EMPIRICAL LEGAL STUD. 602, 606 (2014).

94. 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.

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"—that is, 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 no 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 many of the drivers of this score exist outside of the control of the Agency.⁹⁵ 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 (Hypothesis 1) granting patents at relatively higher rates to applicants in those technologies with stronger proclivities to continue the application process

95. 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—for instance, a sustainability score applied to the resins category that removes all resins-related factors contributing to the score (results available upon request).

upon rejection in comparison with applicants in those technologies with weaker such proclivities, and (Hypothesis 2) allowing a greater number of claims within each allowed application in those technologies prone to high repeat filing relative to technologies not prone to high repeat filing. By responding along both the extensive and intensive margins of allowance in this manner—by both issuing patents 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.

Our findings are consistent with both of the above hypotheses. That is, our results indicate that during times of resource constraint the PTO grants patents in high-repeat-filing technologies, which include, among others, information and communication and health-related technologies, at a relatively higher rate than patents in low-repeat-filing technologies, which include, among others, furniture, house fixtures, and apparel textiles. Additionally, our results indicate that upon deterioration in resources, the PTO increases the number of allowed claims in patents in technologies with stronger proclivities to continue the application process upon rejection compared to those technologies with weaker such proclivities. Thus, our findings provide compelling empirical evidence that the PTO is indeed, in certain circumstances, biased toward granting patents. Moreover, our findings suggest that the Agency is targeting its granting proclivities in certain fields over others.

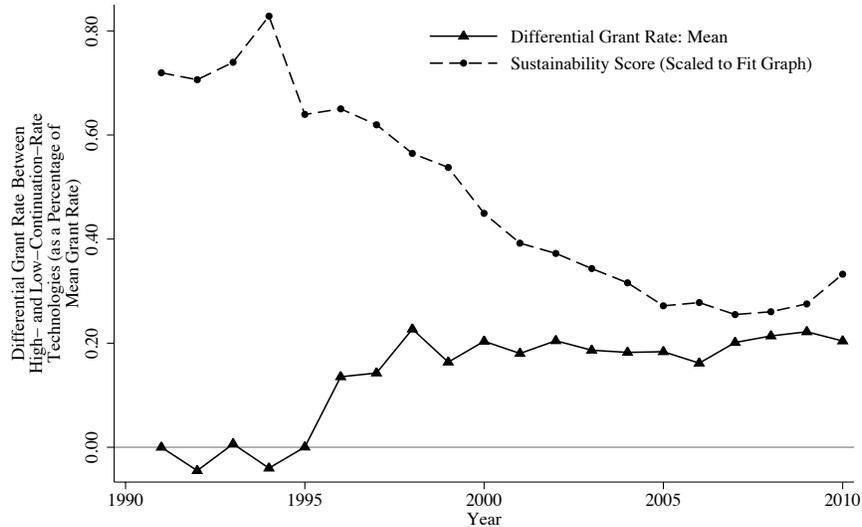
A. *Primary Results: Graphical Analysis*

1. *Description of graphical approach*

In the Online Appendix, we present tabular regression results for the basic difference-in-difference approach mapped out above.⁹⁶ 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.

96. *Online Appendix, supra* note 43, at 12 tbl.A2.

FIGURE 1
Time Trend in Differential Grant Rate (Logged) Between Treatment and Control Technologies, Overlaid with Sustainability Score



Notes: The solid line with triangular markers represents 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. The 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—which include, among others, software, business methods, medical and surgical devices, and genetics—and technologies that are not especially prone to file such repeat applications, which include, among others, furniture, house fixtures, and apparel textiles.⁹⁷ We plot this trend over the 1991 to 2010 period. Each value in this time trend captures the degree to which grant rates in the high-continuation-prone groups exceed those in the low-continuation-prone 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 the normalized differences evolve over each subsequent

97. See *id.* at 10 tbl.A1 (providing a complete listing of the repeat filing rates of technologies).

year. For instance, assume that the grant rate of the high-continuation-prone group equals 70% in 1991 and the grant rate of the low-continuation-prone group at that time equals 65%. We nonetheless plot a difference equal to 0 in 1991. Our interest is 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-prone group's rate increases to 72%, then we would plot a value of 2 for 1992 because the difference in those rates would have grown by two percentage points over that one-year period.⁹⁸ Similarly, if the high-continuation-prone 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.⁹⁹ 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.¹⁰⁰ To facilitate a visual representation demonstrating 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 solid line with triangles in Figure 1) with a time trend in the annual “sustainability scores” (marked by the dashed 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

98. 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% increase, considering a mean grant rate over the sample of roughly 70% ($2 \div 70\% = 2.9$).

99. 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.

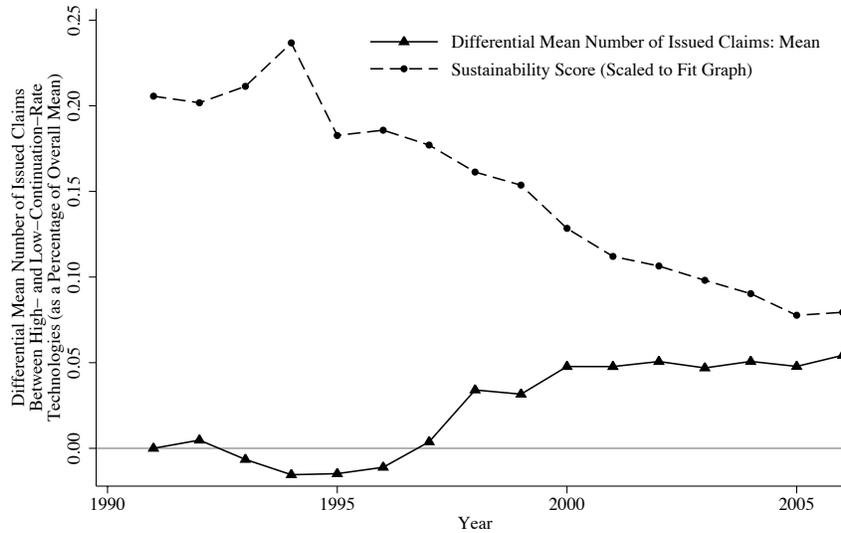
100. Moreover, by doing so graphically, we arguably explore that association in a more meaningful fashion than by simply estimating a single regression coefficient capturing the general association between swings in the Agency's resource sustainability and swings in these differential grant rates.

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 grant rates across the whole Agency over those years—for example, a trend in which the grant rates rise by five percentage points for both the high-continuation-prone group and the low-continuation-prone group. Because 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 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.¹⁰¹

101. This demonstration also highlights the critical importance of collecting relevant data along both a time-series and a cross-sectional (i.e., across technologies) dimension. Studies relying on just one dimension or the other, such as Lemley & Sampat, *supra* note 34, will be limited in their ability to isolate causal stories.

FIGURE 2
Time Trend in Differential Issued Claims Quantity (Logged) Between Treatment and Control Technologies, Overlaid with Sustainability Score



Note: This Figure replicates Figure 1, except that it uses the mean number of issued claims for the technology-by-year group as the relevant outcome variable.

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—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. 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.

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. That is, our findings are consistent with the PTO responding to negative shocks to the Agency’s resources by (Hypothesis 1) granting patents at relatively higher rates to applicants in those technologies with stronger proclivities to continue the application process upon rejection in comparison to ap-

plicants in those technologies with weaker such proclivities, and (Hypothesis 2) allowing a greater number of claims within each allowed application in those technologies prone to high repeat filing relative to technologies not prone to high repeat filing. To begin, consider the early years in the sample—the early 1990s. At this time, the sustainability score rises, suggesting that the Agency’s resources (and hence its ability to process the applications demanded of it) may have been improving. In part, this may have been due to the fact that the PTO began to collect the substantial twelve-year maintenance fees for the first time during these years.¹⁰² With such favorable conditions, one might predict that the Agency would find itself in little need of taking any distortionary actions—namely, granting excessively—in order to diminish repeat filings. Consistent with these predictions, over these early years, the Agency does not appear to have altered the manner in which it treated the high-continuation-prone groups relative to the low-continuation-prone groups either in terms of whether to allow applications at all or in terms of how many claims to allow.

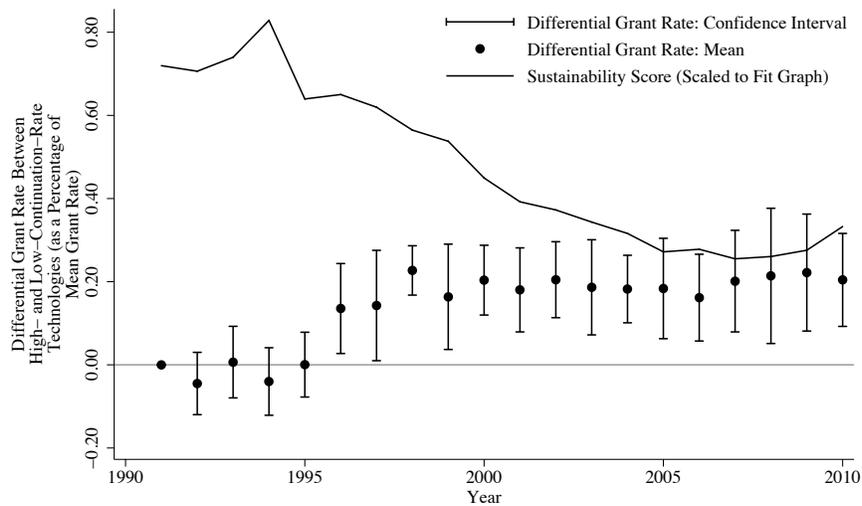
The fact that the differential granting outcomes stay roughly flat around the zero line over the early 1990s suggests that these various outcomes were trending in the same manner in the treatment technologies as they were in the control technologies at that time (again acknowledging that they may have been doing so at different baseline *levels*). This is encouraging for a number of reasons. Not only is it consistent with our prediction 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 also 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 group and the control group 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 were 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 came 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 begins to fall, suggesting that the PTO began to encounter greater and greater difficulties in marshaling enough fee income to cover the costs associated with all of those obligatory examinations awaiting

102. WORKING FOR OUR CUSTOMERS, *supra* note 31, at 29 (“In 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 fiscal year 1994. *Compare id.* at 59 (noting that 32% of patent fee collections came from maintenance fees in fiscal year 1994), with U.S. PATENT & TRADEMARK OFFICE, ANNUAL REPORT, FISCAL YEAR ’93, at 34 fig.9 (1994) (noting that 26% of patent fee collections came from maintenance fees in fiscal year 1993).

review. Between 1994 and 1995, the Agency’s sustainability score fell by roughly 23%. It thereafter fell by an additional 30% between 1995 and 2000. This decline in the Agency’s financial position actually continued throughout the subsequent decade, with its 2009 level nearly 39% lower than the 2000 level. Contemporaneous with this documented decline in the Agency’s resource 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% throughout the first five years of the 1990s, it grew a staggering 114% over the subsequent five years. It then grew a further 190% over the course of the 2000s.

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



Note: 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 prone to file more continuation applications relative to those that are not as prone to do so. The results presented in Figures 1 and 2 are indeed consistent with those expectations, demonstrating that the theorized uptick in the differential granting patterns corresponds with a decline in the PTO’s ability to generate enough fee income to cover the costs associated with the examinations demanded of the Agency. Note that Figure 3 replicates Figure 1 but adds 95% 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% 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 led it to grant at higher rates to the high-continuation-prone technologies.

Finally, in the last year of the sample (2010), it appears as if the Agency's resources began to take 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 begin to improve again. Nonetheless, working with the limited improvement we are able to observe in the data, we do document a slight reduction in the distortionary granting behavior from 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% of the mean overall grant rate—or by roughly eleven percentage points—between 1994 and 1995. By the time the PTO's resource 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 the groups in 1991—hit a level equal to roughly 23% of the mean grant rate (representing roughly a sixteen percentage point spread in the grant rate across these technology groups).¹⁰³

103. 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 Agency's funding woes 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 the 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 above-median examiner experience, *see infra* Figure 6, and those that include control variables for the characteristics of the underlying patents, *see infra* Figure 5, exhibit a more noticeable upward trend in this differential grant rate over the course of the financial decline.

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 application 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 evidence in support of the contention that the PTO is indeed biased toward granting additional patents (of presumably marginal quality) in high-repeat-filing technologies.

The results also suggest that the PTO will likewise begin to expand the scope of its allowed claims at the same moments of resource strain. Data limitations only permit us to track average allowed claim measures until 2006.¹⁰⁴ 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 the groups in 1991) equal to roughly six percent of the mean number of allowed claims (as shown in Figure 2). 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—targeting those technologies that are more likely to pose continuation threats—during times of constrained resources. 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.

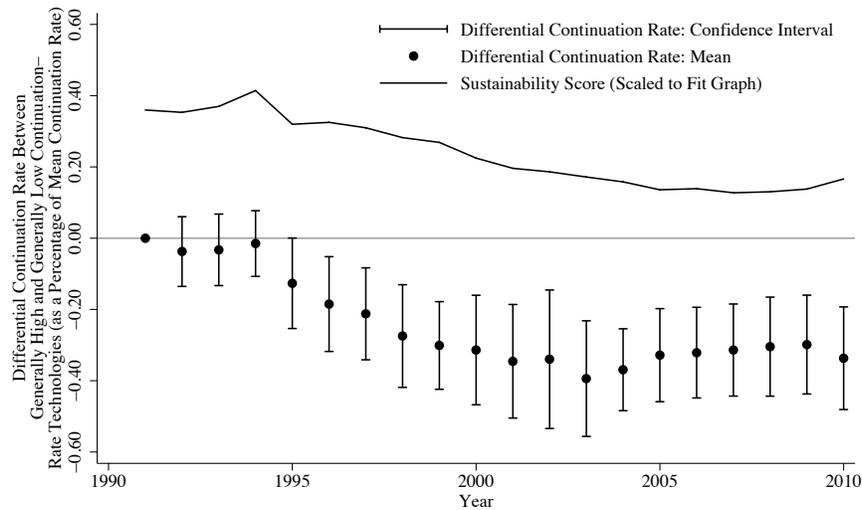
The structure of Figure 4 matches that of Figures 1-3, except that we now plot a year-by-year time trend in the differential continuation rate (which includes RCE filings) between those technologies that are generally highly prone to file continuations and those technologies that are generally not highly 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.¹⁰⁵ We

104. Data on the number of claims allowed per patent were obtained from the NBER patent database. This information was available only for the pre-2006 period.

105. *See supra* Part IV.

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.¹⁰⁶ It is important to emphasize that this allocation process itself—that is, 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.

FIGURE 4
Time Trend in Differential Continuation Rate (Logged) Between Treatment and Control Technologies, Overlaid with Sustainability Score



Note: This Figure replicates 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

106. That is, as we proceed throughout the sample period, this allocation does not change substantially, which is encouraging from the perspective of maintaining a stable experimental framework. See *Online Appendix, supra* note 43, at 6-7.

groups. This pattern, combined with the patterns demonstrated in Figures 1 and 2, lends 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 that an application will be allowed is likely to increase with each additional cycle through the examination process. Thus, the high observed grant rate might 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? By demonstrating that the differential utilization of continuation and RCE filings in fact *fell* for the treatment groups relative to the control groups upon negative resource shocks to the Agency, Figure 4 not only lends support to the idea that the granting patterns in Figures 1-3 reflect a strategy in which the PTO sought to stem off costly repeat filings by granting excessively—as discussed in the preceding paragraph—but 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.

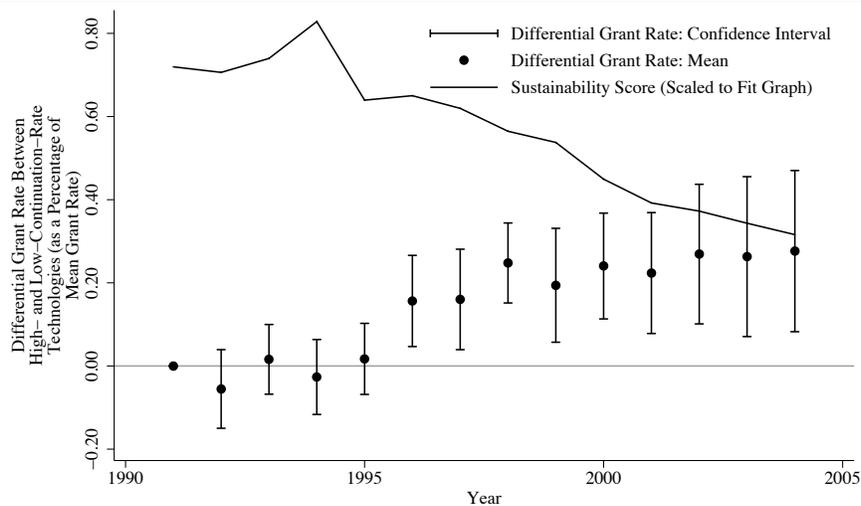
5. *Inclusion of control variables*

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 refer the reader to the Online Appendix for a more technical explanation of this approach.¹⁰⁷ 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 characteris-

107. *Id.* at 1-3.

tics, rather than the story we are claiming to isolate, are responsible for the observed granting patterns. 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 parse 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



Note: This Figure replicates 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 NBER bearing on various features of the patents issued in each year and across each technology.¹⁰⁸ Such features include the number of citations made to these patents—a possible indicator of patent quality¹⁰⁹—and the percentage of patents

108. The NBER patent data can be accessed at NBER PAT. DATA PROJECT, <https://sites.google.com/site/patentdatapoint> (last visited Feb. 23, 2015).

109. See, e.g., Deepak Hegde & Bhaven Sampat, *Examiner Citations, Applicant Citations, and the Private Value of Patents*, 105 *ECON. LETTERS* 287 (2009); Manuel Trajtenberg,

in the technology-by-year group that fall into each of eight different patentee-type classifications (e.g., individual, corporate, or government; foreign or domestic).¹¹⁰ 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 3.¹¹¹

6. *Number of technology groups*

To clarify one final matter regarding the above graphical methodology, we note that the regression underlying these graphs 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 two-standard-deviation shift in continuation rates across technologies; 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 pattern of excessive granting. 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 other story. The above analysis has already attempted to rule out many such stories. The essence

A Penny for Your Quotes: Patent Citations and the Value of Innovations, 21 RAND J. ECON. 172 (1990).

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

111. In the Online Appendix, we demonstrate the same insensitivity of the findings presented in Figure 2 to the inclusion of control variables. *Online Appendix, supra* note 43, at 16 fig.B7.

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 Subpart sought to rule out the influence of certain factors with respect to which data are available at the requisite levels—for example, 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 Subpart, 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 significantly expanded to include inventions in the fields of biotechnology, software, and business methods.¹¹² However, most of the legal developments of this potentially expansionary nature with respect to biotechnology occurred in the early to mid-1980s, prior to the estimation sample frame.¹¹³ While the mid- to late 1990s likewise experienced expansions in patentable subject matter that likely targeted software and business method patents,¹¹⁴ 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.¹¹⁵

112. See, e.g., Wasserman, *supra* note 60, at 381 (describing the “dramatic expansion of the scope of patentable subject matter”).

113. See *Diamond v. Chakrabarty*, 447 U.S. 303, 305, 310 (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) (“The Patent and Trademark Office now considers nonnaturally occurring non-human multicellular living organisms, including animals, to be patentable subject matter . . .”).

114. See, e.g., *State St. Bank & Trust Co. v. Signature Fin. Grp.*, 149 F.3d 1368, 1373 (Fed. Cir. 1998) (enlarging patentable subject matter to include anything that provides “a useful, concrete, and tangible result” (quoting *In re Alappat*, 33 F.3d 1526, 1544 (Fed. Cir. 1994)) (internal quotation marks omitted)), *abrogated by In re Bilski*, 545 F.3d 943 (Fed. Cir. 2008) (en banc); Request for Comments on Proposed Examination Guidelines for Computer-Implemented Inventions, 60 Fed. Reg. 28,778, 28,778-80 (June 2, 1995) (clarifying that computer-implemented inventions constitute patentable subject matter).

115. These results are available from the authors upon request.

Dropping each technology field one by one, we confirm that the estimated relationship, both in terms of magnitude and in terms of statistical precision, remains virtually identical to that estimated in the above Figures.

3. *Selection-effects concerns*

In this Subpart, we evaluate a potential concern stemming from changes over time in initial-application filing behaviors across technologies. Primarily, one may be concerned that the relative increase in grant rates in the period following the mid-1990s in the technologies prone to high repeat filing could be the result of those particular technologies filing fewer original applications during that period. This response might leave behind a select 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 attributable to each technology. Moreover, in Figure B14 of the Online Appendix, we directly test 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.

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.¹¹⁶ With this approach, we rely upon the observed fact (confirmed by our data) that large-entity applicants are generally more prone to file continuation applications than small-entity applicants. Small-entity applicants are individuals, nonprofit corporations, or small businesses.¹¹⁷ 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 the Agency would increase its grant rates for large-entity applicants within high-continuation-prone technologies, at least to a relatively greater degree than 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.¹¹⁸

116. *Online Appendix, supra* note 43, at 3-5.

117. 37 C.F.R. § 1.27(a)(1)-(3) (2011).

118. *Online Appendix, supra* note 43, at 25, 26 tbl.A3.

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 some unrelated story that we have failed to capture in the data caused grant rates to spike in our treatment technologies during the period following the mid-1990s. 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 approach leaves us with substantially greater confidence that we are indeed identifying a grant-rate response to the PTO's dire resource conditions.

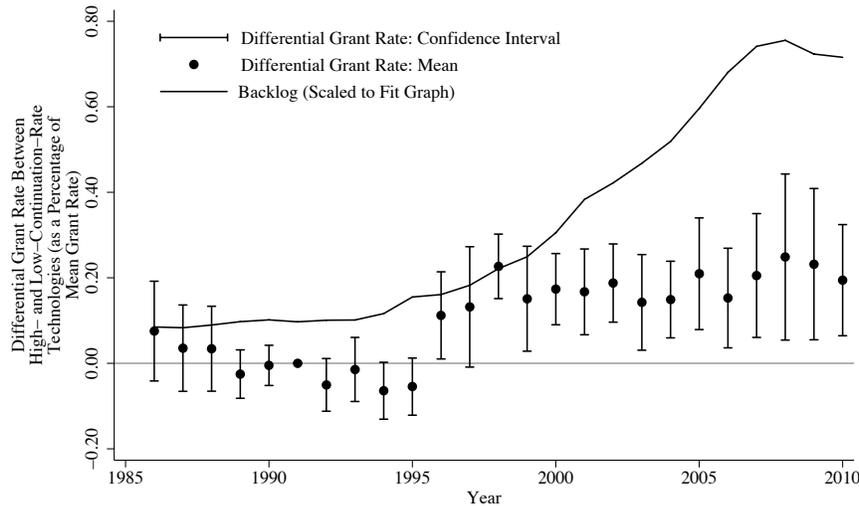
5. *Alternative resource-health metric*

When the Agency finds itself unable to collect enough fees to satisfy the expenses associated with its examination demand—that is, when the Agency finds itself resource constrained—one natural consequence is that the Agency will grow a backlog of applications awaiting examination. Indeed, one may even view fluctuations in the Agency's application backlog as a metric in and of itself by which to evaluate the resource health of the Agency. With this simple notion in mind, in Figure 6, we replicate the above exercise but overlay the differential grant rate between high- and low-continuation-rate technologies with the annual time trend in the Agency's backlog (scaled to fit the graph). As demonstrated by this Figure, the Agency's backlog remained low and stable until the mid-1990s, at which point it began to grow steadily, consistent with the decline in the Agency's sustainability score documented in the prior Figures. Again, it is at this time that the Agency began to grant at relatively higher rates to the high-continuation-prone technologies.

A key advantage of this alternative approach is that it allows us to push the analysis period into the mid-1980s (when we first are able to collect technology-specific grant rates). Since the Agency did not become fully user-fee funded until 1991, the concept of a sustainability score—which reflects a balance of fees and costs—had little meaning until that point. The backlog, however, is arguably a valid metric of the resource condition of the Agency regardless of the method of financing. This longer time period is beneficial in that it allows us to view PTO behavior over a longer period of time in which it had positive resource health—that is, the mid-1980s until the mid-1990s. Encouragingly, over this nearly ten-year period of low and stable backlogs, the differential in the grant rates between the treatment group and the control group hovered consistently around zero. It was not until the Agency's backlog began to increase substantially that the Agency began to alter its granting practices in the hypothesized manner. This lends further credibility to the quasi-experimental research design and to its inherent assumption that the treatment and the control groups

would have otherwise trended in the same manner but for the resource shock of interest.

FIGURE 6
Time Trend in Differential Grant Rate (Logged) Between Treatment and Control Technologies, Overlaid with Backlog Trend with Confidence Intervals



Note: This Figure replicates Figure 3, except that it uses the Agency backlog as the underlying metric of sustainability.

6. Other

In the Online Appendix, we further demonstrate that the pattern of results presented above is generally robust through 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,¹¹⁹ (2) the inclusion of a set of controls to rule out a story in which the results can be explained by the passage of the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS) in 1995¹²⁰ and its resulting effect

119. 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 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 his application and thus more likely that the application will be allowed. *Id.* at 20 fig.B15.

120. Agreement on Trade-Related Aspects of Intellectual Property Rights, Apr. 15, 1994, Marrakesh Agreement Establishing the World Trade Organization, Annex 1C, 1869 U.N.T.S. 299.

on effective patent lengths,¹²¹ and (3) the specification of technology groups according to the more fine-grained PTO classification system, rather than the NBER technology subcategories specified by Hall, Jaffe, and Trajtenberg.¹²²

Finally, it is critical to note that the findings presented in this Article do not simply derive from the phenomenon investigated in our previous research, in which 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.¹²³ 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); we demonstrate this more fully in the Online Appendix, where 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.¹²⁴

C. Results: Concluding Remarks

As far as we have endeavored to bring the reader toward a definitive *causal* interpretation of the findings, we acknowledge that we cannot bring the reader one hundred 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” for causal analysis¹²⁵—have their limitations. 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 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, 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. Consider, for

121. *Online Appendix*, *supra* note 43, at 21 fig.B16; 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).

122. *Online Appendix*, *supra* note 43, at 14 fig.B4, 16 fig.B6, 17 fig.B8, 18 fig.B11, 19 fig.B13; see *supra* note 80.

123. See Frakes & Wasserman, *supra* note 38, at 124.

124. *Online Appendix*, *supra* note 43, at 22 figs.B18A, 23 fig.B18B, 24-25.

125. See, e.g., Lisa Larrimore Ouellette, *Patent Experimentalism*, 101 VA. L. REV. 65, 87 (2015) (“The gold standard for [creating] comparison groups is random assignment of the policy treatment of interest.”).

instance, the often-discussed “Hawthorne effect,”¹²⁶ whereby the experiment may be confounded by the subject’s knowledge that she is 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 not likely be immune from the Hawthorne effect. 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—that is, 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 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 assumptions. 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.

VI. THE MECHANISM BY WHICH THE PTO PREFERENCES GRANTING CERTAIN TYPES OF PATENTS OVER OTHERS

The previous Part set forth evidence that the PTO is indeed overgranting patents during times of resource distress. Moreover, it demonstrated that the Agency is targeting its overgranting proclivities at those types of patents it stands to benefit the most from allowing—patents that are in historically high-continuation-prone technologies, which include, among others, information and software and health-related 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 8000 patent examiners.¹²⁷ The current literature’s emphasis on the 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.¹²⁸

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

127. 2013 PERFORMANCE AND ACCOUNTABILITY REPORT, *supra* note 31, at 9 & fig.2.

128. 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 difficul-

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 supervisory patent examiners (SPEs), individuals who 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 majority of former SPEs we had interviewed at the time this Article went to print were aware both that continuation filings varied substantially across technologies and of the Agency's general resource health.¹²⁹ And at least one former SPE stated that a resource-constrained PTO would attempt to increase its net resources by granting patents faster (i.e., allowing patents earlier in the examination process).¹³⁰ This former SPE further noted that the Agency would likely utilize some top-down initiatives to effectuate an increase in grant rate.¹³¹ That is, the pressure to increase the grant rate would originate from high-level bureaucrats and trickle down to patent examiners.

We found some evidence that is at least suggestive of such a top-down initiative through a FOIA request to the PTO associated with another project. More specifically, an e-mail sent from the Office of Patent Training to individuals who appear to supervise SPEs requested that SPEs be given "compact prosecution training" and then train examiners in their Art Unit accordingly.¹³²

ties in controlling patent examiners' output despite strict rules and oversight); Michael J. Meurer, *Patent Examination Priorities*, 51 WM. & MARY L. REV. 675, 700 (2009) (noting 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, 19-20, 52-53 (Wesley M. Cohen & Stephen A. Merrill eds., 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, 151-55 (2004) (finding that certain examiners more systematically required applicants to narrow the scope of their patents).

129. Telephone Interview with Former Supervisory Patent Exam'r No. 1, U.S. Patent & Trademark Office (July 7-8, 2014); Telephone Interview with Former Supervisory Patent Exam'r No. 2, U.S. Patent & Trademark Office (July 11, 2014); Telephone Interview with Former Supervisory Patent Exam'r No. 4, U.S. Patent & Trademark Office (July 22, 2014); Telephone Interview with Former Supervisory Patent Exam'r No. 5, U.S. Patent & Trademark Office (July 22, 2014).

130. Telephone Interview with Former Supervisory Patent Exam'r No. 1, *supra* note 129.

131. *Id.*

132. E-mail from Patrick J. Nolan, Manager, Training & Dev. Program, Office of Patent Training, to Joseph Del Sole et al. (Apr. 15, 2013) (on file with authors).

The e-mail states that the “genesis of this training was a request from the ADCs [Associate Deputy Commissioners] to update our current compact prosecution training for the Patent Corps in an effort to help reduce RCE filings.”¹³³ The attached training slides indicate that compact prosecution training includes, among other things, “[i]dentifying allowable subject matter in an effort to expedite prosecution.”¹³⁴ The e-mail and attached training materials could be read simply as a call to improve the quality of review during the early stages of the prosecution of an application in order to better identify those applications that in fact meet the patentability requirements. Of course, the mandate to “[i]dentify[] allowable subject matter” more quickly may also capture pressure by the Agency to grant patents of marginal validity in order to discourage the filing of RCEs in those contexts.¹³⁵ We note of course that it is unlikely that the PTO would ever be so explicit in clarifying its intentions if, in fact, this latter interpretation were behind its training initiatives. While our interviews with former SPEs, the above-quoted e-mail, and training materials 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 upward direction in an effort to decrease repeat filings.

More specifically, we believe there are at least two different channels by which the PTO can favor certain patent types over others. We note that these channels, however, are not necessarily mutually exclusive. The first is a top-down channel, whereby 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 assigned to a patent examiner.¹³⁶ Art Units are likewise aggregated into larger parcels and are eventually aggregated into one of nine technology centers.¹³⁷ This hierarchical structure creates a situation in which a targeted population of

133. *Id.* There are three deputy commissioners, who are high-ranking officials, at the PTO. *USPTO Organization*, U.S. PAT. & TRADEMARK OFF., <http://www.uspto.gov/about/offices/index.jsp> (last modified Oct. 29, 2009).

134. USPTO Patent Training Acad., *Compact Prosecution: Clarity, Correctness and Completeness—Your Role in Compact Prosecution* at 3 (Apr. 8, 2013), Attachment to E-mail from Patrick J. Nolan, *supra* note 132.

135. It is also possible, even if not the intention of the PTO, that the allowance of invalid patents is a side effect of compact prosecution. That is, examiners who feel pressure to make decisions quicker may also make more errors than examiners who do not feel such pressure.

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

137. *Patent Technology Centers*, U.S. PAT. & TRADEMARK OFF., 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 PTO).

examiners consistently examines patents of a particular technology. (For example, applications of a particular technology are not randomly assigned among a large number of Art Units; rather, they are assigned to a single Art Unit or to only a few.)¹³⁸ 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.

If the mechanism the PTO is utilizing to inflate its granting tendencies is a top-down channel, one might suspect that the elevated grant rate of a technology parcel (i.e., Art Unit or aggregation of several Art Units) with high repeat-filing rates would be distributed across all patent examiners within that parcel. That is, both senior examiners and junior examiners in a technology parcel with high repeat-filing rates would demonstrate an inflated grant rate during times when the Agency's resources were insufficient to meet its expected examination demand. Nevertheless, we acknowledge that it is possible that high-level bureaucrats may only instruct a trusted few—most likely senior examiners—to increase their grant rate rather than directing all the examiners within that parcel to increase their allowance rates. Thus, this first channel for favoring certain patent types over others may result in all examiners or only senior examiners within a technology parcel with high repeat-filing rates biasing their grant rates in an upward direction.

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 into 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 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 on 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.

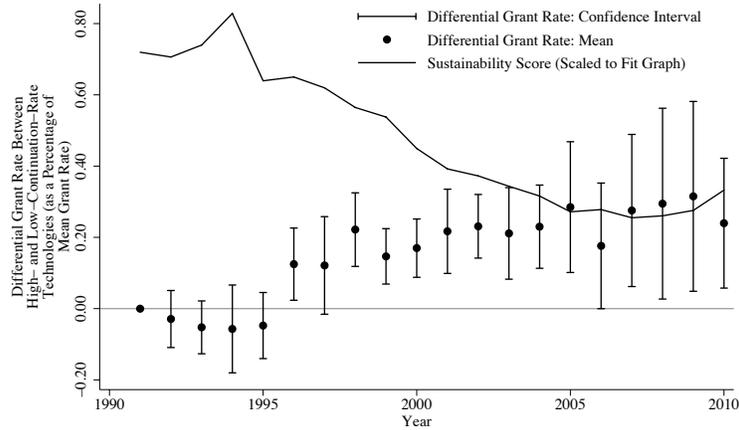
138. Art Units may be assigned patent applications from one class, from a portion of a class, or from several classes involving closely related technology. See *Patent Classification: Classes Arranged by Art Unit*, U.S. PAT. & TRADEMARK OFF., <http://www.uspto.gov/patents-application-process/patent-search/understanding-patent-classifications/patent-classification.html> (last modified Feb. 5, 2015).

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.¹³⁹ This greater leeway enables senior examiners to change their granting proclivities without fearing 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 grant 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, we would need to collect examiner-seniority-specific grant rates. 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 onward, 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 FOIA requests with the PTO in order to collect data through 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 are populated by less experienced examiners. We perform this allocation by looking above and below the median of this seniority distribution across technologies. 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.

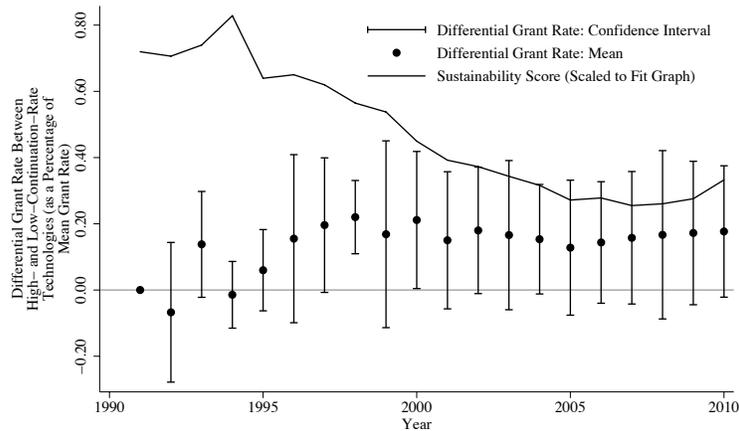
139. See Lemley & Sampat, *supra* note 34, at 818-19.

FIGURE 7
 Time Trend in Differential Grant Rate (Logged) Between Treatment and Control Technologies, Overlaid with Sustainability Score, Above-Median-Experience Technologies



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

FIGURE 8
 Time Trend in Differential Grant Rate (Logged) Between Treatment and Control Technologies, Overlaid with Sustainability Score, Below-Median-Experience Technologies



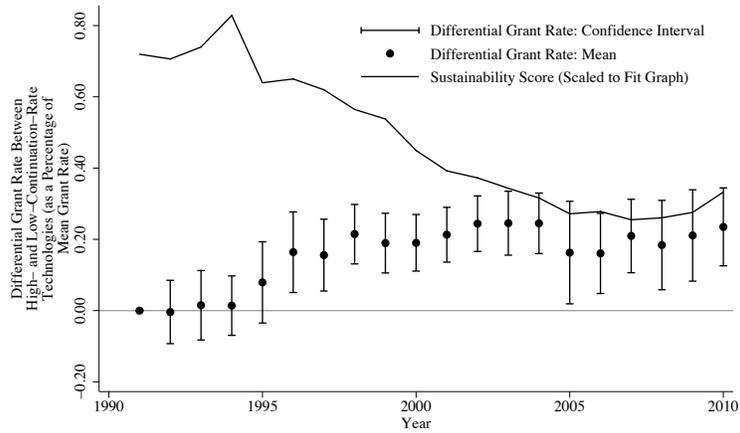
Note: This Figure replicates Figure 3, except that it focuses only on the set of technologies with generally below-median examiner experience levels.

We conduct this exercise in Figures 7 and 8, with Figure 7 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 8 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. These contrasting patterns are even more striking when using the PTO classification system to identify technology groups, as distinct from the coarser NBER technology subcategories, as we demonstrate in Figures 9 and 10. These findings suggest that the primary story set forth in Figure 1 (and Figure 3) of this Article is likely one that is more relevant among the pool of senior examiners and hence rules out the explanation that the channel by which the PTO is inflating its grant rate during times of resource stress is solely driven by top-down forces that target all examiners in the technology parcel. We note, however, that our results do not definitely point to one channel over the other, as it is possible that both top-down and examiner-driven initiatives may result in the elevation of the grant rate of senior examiners only. Moreover, it is possible that both channels might be working in tandem.

Finally, this additional exercise also serves as yet another robustness check on our primary result—that 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-filing technologies has been improving over time, whereas the quality in low-repeat-filing technologies has remained unchanged. Although we have ruled out this potentially concerning explanation multiple times throughout our analysis, this additional exercise provides yet another way to discredit the idea that a differential change in the underlying quality of applications (or a similar story) is contributing to our results. Because patent applications are randomly assigned to examiners within an Art Unit, if our findings were really being driven 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.¹⁴⁰ The fact that 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 to our conclusion: the PTO is in fact biasing its grant rate in an upward direction in an effort to decrease repeat filings and combat its backlog of patent applications.

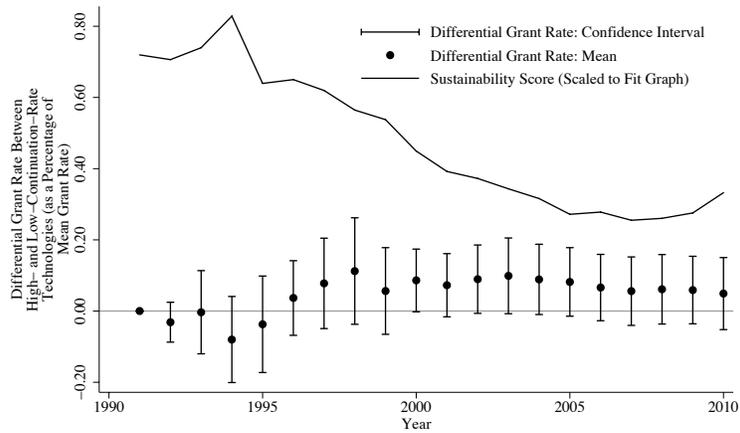
140. *See id.* at 822 (noting that patent applications are for the most part randomly assigned within an Art Unit).

FIGURE 9
 Time Trend in Differential Grant Rate (Logged) Between Treatment and Control Technologies, Overlaid with Sustainability Score, Above-Median-Experience Technologies Using PTO Classes



Note: This Figure replicates Figure 7, except that it uses PTO classes to identify technologies rather than NBER subcategories.

FIGURE 10
 Time Trend in Differential Grant Rate (Logged) Between Treatment and Control Technologies, Overlaid with Sustainability Score, Below-Median Experience Technologies Using PTO Classes



Note: This Figure replicates Figure 8, except that it uses PTO classes to identify technologies rather than NBER subcategories.

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

This Part 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 for Patent Policy*

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 decisionmaking 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 biased just toward allowing patents in general, but also toward 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 due to 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. That is, because the PTO is preferentially granting patents in high-continuation-rate technologies, such as information and communication and health-related technologies, society may be overinvesting resources in these technological fields. As a result, the harms recognized from granting bad patents are likely larger than what commentators have acknowledged to date.

Additionally, by ignoring the import that repeat filings play in influencing PTO decisionmaking, legal scholarship has overlooked a substantial source of Agency bias. The easiest way to eliminate the bias in PTO decisionmaking 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 overgranting 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 PTO's inability to finally reject a patent application biases the Agency toward 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 toward those patents 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 to either (1) enable the PTO to expand its examination capacity to process repeat applications or (2) 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.¹⁴¹ At the same time, commentators generally agree that the extent to which continuation applications are used in an abusive manner varies across technologies. For instance, pharmaceutical and biotechnology companies are generally believed to use repeat filings in a nonabusive manner.¹⁴² Nevertheless, the PTO, largely out of concern for its growing patent pendency and backlog, did attempt to limit repeat filings in 2007.¹⁴³ After a protracted court battle, in which the Agency's authority to promulgate such regulations was questioned, the PTO ultimately rescinded the regulations.¹⁴⁴ Thus, it is unclear whether the PTO has the authority to limit repeat filings or whether Congress would have to amend the patent laws 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 by access concerns: low up-front fees increase access to the system, thereby spurring innovation.¹⁴⁵ Yet be-

141. Lemley & Moore, *supra* note 9, at 76-79, 81.

142. These companies often file patent applications 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. Schreiner & Doody, *supra* note 59, at 557.

143. *See supra* note 49.

144. *See supra* note 49.

145. Setting and Adjusting Patent Fees, 78 Fed. Reg. 4212, 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").

cause 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 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.¹⁴⁶

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 in a better position is to alter its fee schedule to substantially decrease the chances that the Agency's fee collections are insufficient to cover the costs associated with all of those obligatory examinations awaiting review.¹⁴⁷ The most straightforward approach to increasing the Agency's health is to decrease the PTO's reliance upon postallowance fees to fund patent examinations. By bringing examination fees to the level of examination costs, the Agency would no longer need to be concerned that application filings will grow out of step with the stock of existing patents from which it derives postallowance fees. The Agency would thereby come closer to always ensuring that it has sufficient fee revenues to cover the costs of those applications demanded of it.

The PTO recently obtained the ability to set its fees by rulemaking, although the Agency's fees cannot exceed its aggregate costs.¹⁴⁸ In 2013, the PTO exercised this new authority and promulgated an updated fee schedule that modestly increases examination fees,¹⁴⁹ substantially decreases issuance

146. See, e.g., U.S. GOV'T ACCOUNTABILITY OFFICE, *supra* note 56, at 5 (concluding that "it is unlikely that the agency will be able to reduce the backlog simply through its hiring efforts").

147. 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. Frakes & Wasserman, *supra* note 38, at 77-78. Furthermore, mounting concerns over the federal government's fiscal cliff suggest that funding the Agency through taxes may not result in the PTO receiving sufficient resources to process its growing backlog of patent applications. One of the primary drivers behind Congress's increased reliance on user fees to finance agencies has been the belief that such a funding mechanism increases the resource sustainability of the Agency, especially in the current environment of deficit containment.

148. See Leahy-Smith America Invents Act, Pub. L. No. 112-29, § 10(a), 125 Stat. 284, 316 (2011), *reprinted in* 35 U.S.C. § 41 note (2013).

149. See Setting and Adjusting Patent Fees, 78 Fed. Reg. at 4224 tbl.4 (increasing examination fees for large entities by \$340, or 27%).

fees,¹⁵⁰ and substantially increases renewal fees.¹⁵¹ Unfortunately, the Agency's new fee structure does not eliminate the risk that the PTO's fee collections will fail to cover its operational expenses. The ratio of postallowance fees to examination fees remains virtually unchanged.¹⁵² 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 postallowance 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 process 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 also 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.¹⁵³ Of course, the PTO's fee structure should try to optimize not only the Agency's incentives but also the social welfare of society more generally. Perhaps the biggest drawback from this proposal is that a substantial increase in 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.¹⁵⁴ 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,¹⁵⁵ there is

150. *See id.* (decreasing issuance fees for large entities by \$810, or 46%).

151. *See id.* (increasing the first-stage maintenance fee due at 3.5 years for large entities by \$450, or 39%; increasing the second-stage maintenance fees due at 7.5 years for large entities by \$700, or 24%; and increasing the third-stage maintenance fees due at 11.5 years for large entities by \$2590, or 54%).

152. Under the proposed fee schedule, the ratio of postallowance to examination fees is approximately 8.5 $((960 + 1600 + 3600 + 7400) / 1600 = 8.5)$; under the current fee schedule, the ratio of postallowance to examination fees is approximately 8.7 $((2070 + 1150 + 2900 + 4810) / 1260 = 8.7)$. *See id.*

153. 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 the labor force is normatively desirable. *See* Golden, *supra* note 30, at 486-89.

154. *See* Gaetan de Rassenfosse & Bruno van Pottelsberghe de la Potterie, *On the Price Elasticity of Demand for Patents*, 74 OXFORD BULL. ECON. & 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).

155. A utility patent application of minimal complexity costs on average \$6500 to prepare and prosecute, while a utility patent application of relatively high complexity costs on

reason to believe that even a two- or threefold increase in examination fees may not have substantial negative social welfare costs. However, more empirical investigation is needed before any definite conclusions can be drawn.

Finally, because the America Invents Act allows the PTO to set fees only to recoup its operational costs,¹⁵⁶ issuance and renewal fees will likely need to be decreased if examination fees are substantially increased. It is not, however, readily apparent that decreasing postallowance 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 PTO's incentives and general social welfare concerns. If this becomes too difficult, it may be time to consider lifting the restriction that the aggregate PTO fees collected 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 America Invents Act 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 postallowance 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 postallowance fees, add them to the general tax revenue, and in turn create a refundable tax credit to patent applicants. Alternatively, the PTO could 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 nonfunding 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 postallowance fees, they importantly shift the financial risk from the Agency to patent applicants. When postallowance 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

Patent quality has been at the heart of the patent reform debate. Despite the general consensus that the PTO allows too many invalid patents to issue that

average \$10,000 to prepare and prosecute. AIPLA, REPORT OF THE ECONOMIC SURVEY 2013, at 27 (2013).

156. See Leahy-Smith America Invents Act, Pub. L. No. 112-29, § 10(a)(2), 125 Stat. 284, 316 (2011), reprinted in 35 U.S.C. § 41 note (2013).

are unnecessarily draining consumer welfare and stunting innovation, there exists little compelling evidence that the PTO is actually biased toward allowing patents. As a result, policymakers have been modifying the patent system in an effort to increase patent quality in the dark. As there exists little compelling empirical evidence the PTO is actually overgranting patents, lawmakers are left trying to fix the patent system without even understanding the root causes of the system's dysfunction.

This Article begins to rectify this deficiency in the literature by providing lawmakers with much-needed guidance as to what features of the Agency are actually inducing the PTO to overgrant patents. In doing so, we provide not only a novel source of the Agency's overgranting proclivities—the PTO's inability to finally reject a patent application—but also some of the first compelling empirical evidence that the Agency is in fact biased toward granting patents. Our results suggest that the inability of the PTO to finally rid itself of an application biases it toward 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—namely, those patents in high-repeat-filing technologies, such as information and communication technologies (which include software, business methods, and information storage, among others) and health-related technologies (which include surgical and medical instruments and genetics, among others). Thus, our findings should begin to help policymakers make informed decisions on how best to modify the patent system so as to increase patent quality.