
ARTICLE

DID TRIPS SPUR INNOVATION? AN ANALYSIS OF PATENT DURATION AND INCENTIVES TO INNOVATE

DAVID S. ABRAMS[†]

How to structure IP laws in order to maximize social welfare by striking the right balance between incentives to innovate and access to innovation is an empirical question. It is a challenging one to answer, both because innovation is difficult to value and because changes in IP protection are rare. The 1995 TRIPS agreement provides a unique opportunity to explore this question for two reasons. First, the implementation of the agreement was uncertain until shortly before adoption, making it a plausibly exogenous change in patent duration. Second, the nature of the law change meant that the patent-duration change was heterogeneous across patent classes. Using both patent counts and citation-weighted counts, I am able to exploit the TRIPS-induced law change to empirically evaluate the impact of patent duration on innovation. I find evidence for an increase in innovation due to patent-term extension following TRIPS. Both patent counts and citation-weighted counts increased more after TRIPS in those classes that received greater expected term extensions relative to classes receiving shorter extensions. While the precise calibration of innovation valuation is difficult, this Article provides the first attempt to empirically evaluate its response to a major change in patent duration from the TRIPS agreement.

[†] Assistant Professor of Law, University of Pennsylvania Law School. E-mail: dabrams@law.upenn.edu. Many thanks to Polk Wagner for the provision of the patent-count data and Marla Weinstein for her excellent research assistance. Thanks also to John Duffy, Josh Lerner, Gideon Parchomovsky, Devin Pope, Jeremy Tobacman, Polk Wagner, Joel Waldfogel, Maisy Wong, Christopher Yoo, and participants of the Symposium on Foundations of Intellectual Property Reform, held at the University of Pennsylvania Law School, for their many helpful comments.

INTRODUCTION	1614
I. BACKGROUND	1618
A. <i>The TRIPS Agreement</i>	1618
B. <i>Related Literature</i>	1622
C. <i>Theoretical Background</i>	1626
II. DATA.....	1627
III. ECONOMETRIC METHODOLOGY.....	1632
IV. RESULTS.....	1635
V. DISCUSSION.....	1639
CONCLUSION.....	1641
APPENDIX	1643

INTRODUCTION

In 1994, as part of the General Agreement on Tariffs and Trade (GATT) that created the World Trade Organization (WTO), the United States made the largest change in patent terms in over forty years.¹ In order to conform to the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS),² the United States changed the duration of patent protection from seventeen years from grant date to twenty years from application date.³ This change affords the opportunity to learn about one of the most basic issues in IP: the relationship between the quantity of innovation and the duration of IP protection.

In this Article, I explore this relationship between duration and quantity using U.S. patent application and citation data around the dates of the implementation of the TRIPS agreement. I use a difference-in-difference framework—exploiting the heterogeneous impact of the change across patent classes—based on Patent and Trademark Office (PTO) processing time. I find a statistically significant relationship between the magnitude of the term extension resulting from the TRIPS law change and patent count. Further, I find that this relationship persists when using citation-weighted patent counts as the dependent variable, which is arguably a better proxy for value of innova-

¹ The most significant change prior to 1994 was Congress's enactment of 35 U.S.C. § 154 in 1952. Patent Act of 1952, Pub. L. No. 82-593, 66 Stat. 792 (codified as amended in scattered sections of 35 U.S.C.).

² Agreement on Trade-Related Aspects of Intellectual Property Rights, Apr. 15, 1994, Marrakesh Agreement Establishing the World Trade Organization, Annex 1C, 108 Stat. 4809, 1869 U.N.T.S. 299, *available at* http://www.wto.org/english/docs_e/legal_e/27-trips.pdf [hereinafter TRIPS].

³ Uruguay Round Agreements Act, Pub. L. No. 103-465, § 532(a), 108 Stat. 4809, 4984-85 (1994) (codified as amended in scattered sections of 15, 17, 19, and 35 U.S.C.).

tion. I do not find a significant increase in the average number of citations to patents receiving longer extensions relative to those receiving shorter extensions following the change.

Standard theory argues that an increase in the duration of patent protection has two primary effects: an increased incentive to innovate (created by monopoly profits) and increased deadweight loss (due to exclusive rights).⁴ The optimal patent term is that point at which the marginal benefit from increased innovation is exactly offset by the marginal cost of the deadweight loss created by the patent right.

Determining the optimal patent term is extremely important from a policy perspective. If patent protection lasts too long, the monopoly-like deadweight loss, caused by the patent's conferral of exclusive rights, outweighs the additional innovation such rights will spur. On the other hand, a patent term that is too short will yield underproduction of innovation, leading to a decrease in productivity and growth. In order to find the optimum empirically, it is necessary to estimate the elasticity of production of innovation with respect to the duration of protection. Previously, the ability to empirically evaluate the optimal patent length was limited by a dearth of data, a lack of computing power, and an absence of a change to patent length with which to evaluate the elasticity.⁵

This Article works towards evaluating the elasticity of production of innovation by examining data from 1990 through 2000 in light of the 1995 change in patent protection. The change in duration provides the denominator for the elasticity calculation (which I calculate separately by patent class). Though the change was not completely unanticipated, it was part of a global-trade accord that encompassed myriad issues beyond intellectual property. Thus, the law change may be viewed as plausibly exogenous (this assumption will be examined more closely later in the Article) and allows for the identification of the impact of a modification in patent duration.

⁴ See WILLIAM D. NORDHAUS, *INVENTION, GROWTH, AND WELFARE: A THEORETICAL TREATMENT OF TECHNOLOGICAL CHANGE* 76 (1969) (noting that a longer patent life "increases invention" but also increases "losses from inefficiencies associated with monopoly").

⁵ Elasticity is a mathematical term often used in economics that is similar to the derivative in calculus. But whereas the derivative measures the absolute change in the dependent variable with respect to the independent variable, the elasticity is defined as the percentage change in the dependent variable with respect to the percentage change in the independent variable.

In this Article, I am interested in the elasticity of production of innovation with respect to patent duration. The elasticity I wish to estimate is percent change in innovation divided by percent change in patent duration.

The numerator in the elasticity calculation should ideally be the percentage change in the total social welfare derived from patented ideas—an exceedingly difficult quantity to measure. In this Article, I take two approaches to approximating this quantity.⁶ The first approximation makes use of simple patent counts as a proxy for the value of innovation, examining the number of patents obtained prior to and after the 1995 law change by patent class. While this method should yield the correct sign for the elasticity, the magnitude could be substantially off unless the value of the marginal patentable innovation is constant, which seems highly unlikely. Nevertheless, using patent counts as a measure of innovation has a long history and is informative about the response of innovators to incentives.⁷

Previous research has shown that patents vary widely in value.⁸ A closer approximation of total patent value may be obtained by weighting patents according to how many citations they receive from subsequent patents.⁹ This method accounts for some of the heterogeneity in

⁶ There is also a third approach—based on patent-renewal data—which I describe *infra* note 10. However, a full analysis under this approach is beyond the scope of this Article.

⁷ See, e.g., Rebecca Henderson, Adam B. Jaffe & Manuel Trajtenberg, *Universities as a Source of Commercial Technology: A Detailed Analysis of University Patenting, 1965–1988*, 80 REV. ECON. & STAT. 119, 119-22 (1998) (exploring the use of university patents as an indicator of innovation). See generally Zvi Griliches, *Introduction to R & D, PATENTS, AND PRODUCTIVITY* 1, 14 (Zvi Griliches ed., 1984) (reflecting on the general theme of the conference, that “patent counts are, after all, one of the few direct quantitative glimpses into the innovation process available to us”).

⁸ See, e.g., Jean O. Lanjouw, Ariel Pakes & Jonathan Putnam, *How To Count Patents and Value Intellectual Property: The Uses of Patent Renewal and Application Data*, 46 J. INDUS. ECON. 405, 411-18 (1998) (recognizing variance in patent value across time, industries, nationalities, and technology categories); Ariel Pakes, *Patents as Options: Some Estimates of the Value of Holding European Patent Stocks*, 54 ECONOMETRICA 755, 767-80 (1986) (analyzing the returns on European patent stocks and finding that patent value varies by country); Ariel Pakes & Margaret Simpson, *Patent Renewal Data*, 1989 BROOKINGS PAPERS ON ECON. ACTIVITY: MICROECONOMICS 331, 355-64 (noting that “patents vary greatly in both their private and social values” while considering how patent-renewal data can be used to analyze patent value).

⁹ See Bronwyn H. Hall, Adam Jaffe & Manuel Trajtenberg, *Market Value and Patent Citations*, 36 RAND J. ECON. 16, 19 (2005) (describing how the valuable economic and technological information provided by patent citations may help gauge the “value” of patents”); Dietmar Harhoff et al., *Citation Frequency and the Value of Patented Inventions*, 81 REV. ECON. & STAT. 511, 515 (1999) (finding that “patents reported to be relatively valuable by the companies holding them are more heavily cited in subsequent patents”); Adam B. Jaffe, Manuel Trajtenberg & Rebecca Henderson, *Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations*, 108 Q.J. ECON. 577, 583 (1993) (discussing how patent citations carry more weight than academic citations); Adam B. Jaffe, Manuel Trajtenberg & Michael S. Fogarty, *Knowledge Spillovers and Patent Citations: Evidence from a Survey of Inventors*, 90 AM. ECON. REV. PAPERS & PROC. 215, 217 (2000)

patent values and allows for weaker assumptions about the value of the marginal patentable innovation.¹⁰ I make use of the National Bureau of Economic Research (NBER) patent database, which contains patent citations, in order to obtain a better estimate of the numerator.¹¹

The rest of the Article proceeds as follows: Part I discusses the background of the TRIPS agreement and the subsequent law change, related literature, and theoretical framework used to analyze its impact. The data are found in Part II and the econometric methodology

(noting how authors increasingly utilize “the total number of citations received by a patent as an indicator of the relative significance of patents”); Jean O. Lanjouw & Mark Schankerman, *Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators*, 114 ECON. J. 441, 446-47 (2004) (developing a composite index of patent quality using multiple factors, including backward and forward citations); Manuel Trajtenberg, *A Penny for Your Quotes: Patent Citations and the Value of Innovations*, 21 RAND J. ECON. 172, 173-75 (1990) (suggesting that patent citations should be used as an index of the value of patents); Jean O. Lanjouw & Mark Schankerman, *The Quality of Ideas: Measuring Innovation with Multiple Indicators* 1-5 (Nat’l Bureau of Econ. Research, Working Paper No. 7345, 1999), available at <http://www.nber.org/papers/w7345> (using backward and forward citations as indicators of patent quality).

¹⁰ In other work, I take a third approach to estimating the change in production of innovation in response to the implementation of TRIPS: using actual renewal decisions by patent holders to estimate the distribution of the private value of patents. In most countries, patent holders must pay a maintenance fee to keep their patents in force. This fee implicitly forces patent holders to make a calculation regarding the future expected returns from the patent.

Various scholars have used this patent-renewal approach in analyzing data from European countries. *E.g.*, Jean Olson Lanjouw, *Patent Protection in the Shadow of Infringement: Simulation Estimations of Patent Value*, 65 REV. ECON. STUD. 671 (1998) (West Germany); Pakes & Simpson, *supra* note 8 (Norway and Finland); Mark Schankerman & Ariel Pakes, *Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period*, 96 ECON. J. 1052 (1986) (United Kingdom, France, and Germany); Mark Schankerman, *How Valuable Is Patent Protection? Estimates by Technology Field*, 29 RAND J. ECON. 77 (1998) (France).

The *Patents as Options* methodology has rarely been used on U.S. data in the academic literature for three primary reasons. First, the United States only began charging maintenance fees for options in 1982. Because this research originated shortly thereafter, there was not enough data available at the time to perform the estimation. Second, U.S. maintenance fees tend to be somewhat lower than those in Europe, which makes estimating the upper tail of the distribution more dependent on functional-form assumptions. Last, maintenance fees are only required to be paid three times in the United States (as opposed to annually in Europe), making the estimates on U.S. data less precise. There has, however, been one previous publication making use of patent-renewal data in the United States. *See* Jonathan A. Barney, *A Study of Patent Mortality Rates: Using Statistical Survival Analysis To Rate and Value Patent Assets*, 30 AIPLA Q.J. 317 (2002) (examining a particular sample of patents granted at a particular point in time).

My analysis (forthcoming) will be one of the first applications of the *Patents as Options* approach to U.S. data, and it will allow for the most sophisticated evaluation of the law change resulting from the TRIPS agreement.

¹¹ For a description of the data used in this Article, see *infra* Part II.

is explained in Part III. Part IV contains the main results from both the patent-count and citation approaches. Part V discusses the results and their limitations.

I. BACKGROUND

A. *The TRIPS Agreement*

The TRIPS agreement grew out of the Uruguay Round of GATT negotiations, which lasted from 1986 through 1994 and led to the creation of the WTO.¹² TRIPS covered many aspects of intellectual property and required harmonization of IP laws among the developed signatory countries. Within the realm of patents, the key requirements of TRIPS were that patents be made available without discrimination to citizens of TRIPS signatory nations on both products and processes¹³ and that the protection extend for a minimum of twenty years.¹⁴

At the time of TRIPS's passage, U.S. patent law provided for seventeen years of patent protection, as measured from the patent grant date. Thus, TRIPS necessitated a significant modification of U.S. law. The seventeen years of patent protection in the United States was not derived from an economic calculation, as advocated in this Article. Initially, U.S. law was modeled on English law, which set an initial fourteen-year patent term based on the expected training period for two sets of apprentices.¹⁵ Nordhaus points out that "[a]fter some further compromise it was decided for the United States that 2.43 apprentices, or 17 years, would be the proper length."¹⁶

For the purposes of this analysis, there are three relevant dates to consider that could have potentially affected innovative activity, as displayed below in Figure 1. The first is the date that the final package negotiated through the Uruguay Round was signed, April 15, 1994.¹⁷ The second is the date the U.S. Congress ratified the package, De-

¹² TRIPS, *supra* note 2.

¹³ *Id.* art. 27.

¹⁴ *Id.* art. 33.

¹⁵ See NORDHAUS, *supra* note 4, at 82 n.18 (quoting SUBCOMM. ON PATENTS, TRADEMARKS & COPYRIGHTS OF THE SENATE COMM. ON THE JUDICIARY, 85TH CONG., AN ECONOMIC REVIEW OF THE PATENT SYSTEM 9 (Comm. Print 1958) (prepared by Fritz Machlup)), which explains the basis for U.S. patent terms.

¹⁶ *Id.*

¹⁷ Final Act Embodying the Results of the Uruguay Round of the Multilateral Trade Negotiations, Apr. 15, 1994, 33 I.L.M. 1125 (1994), *available at* http://www.wto.org/english/docs_e/legal_e/03-fa.pdf.

ember 8, 1994.¹⁸ The third is the date that the change became effective in the United States, June 8, 1995.¹⁹ Based on news reports at the time, it is clear that while the April 15, 1994, signing of the agreement was symbolically important, it was far from certain that the agreement would be effectuated.²⁰ It was not until ratification by Congress that TRIPS was expected to be implemented.

Figure 1: Key Dates for the Uruguay Round of GATT



Difference in difference is an economic technique that may be used to isolate the effect of a law change such as that brought on by the TRIPS agreement. One major application of the technique is to avoid ascribing a causal effect to a law change that is really due to a contemporaneous change or a time trend. This is a significant problem with the single difference, or before-after analysis. With the difference-in-difference approach, one takes the before-after difference for a group that should be affected by the law change and takes a second difference with the before-after difference of a group that should

¹⁸ Uruguay Round Agreements Act, Pub. L. No. 103-465, 108 Stat. 4809 (1994) (codified as amended in scattered sections of 15, 17, 19, and 35 U.S.C.).

¹⁹ See *id.* § 534 (mandating that “the amendments made by section 532 take effect on the date that is 6 months after the date of enactment of this Act”).

²⁰ See, e.g., Keith Bradsher, *Foes Line Up To Do Battle over GATT*, N.Y. TIMES, Oct. 3, 1994, at D1 (elaborating on the battle over GATT); Helene Cooper & John Harwood, *World Trade Pact Limp Toward a Showdown, Bruised by Talk Radio, Sovereignty Issue, Politics*, WALL ST. J., Nov. 23, 1994, at A16 (listing a variety of roadblocks in GATT’s path); Bob Davis & Lawrence Ingrassia, *Trade Acceptance: After Years of Talks, GATT Is at Last Ready To Sign Off on a Pact*, WALL ST. J., Dec. 15, 1993, at A1 (recognizing that congressional ratification of GATT still faced opposition from American industry); Bob Davis, *Unexpected Obstacles Are Threatening To Delay or Derail Congressional Approval of GATT Pact*, WALL ST. J., Apr. 8, 1994, at A16 (detailing the problems in getting congressional approval for GATT); Thomas L. Friedman, *Dole Explains Trade Treaty Stand*, N.Y. TIMES, Sept. 14, 1994, at D2 (explaining the Senate Minority Leader’s opposition to parts of GATT); David Sanger, *After Years of Talk, Trade Pact Now Awaits Congressional Fate*, N.Y. TIMES, Nov. 27, 1994, at A1 (detailing the remaining questions regarding GATT).

be unaffected (or less affected). This is akin to creating an experimental and a control group.

For example, consider a law that becomes effective in January 2010 that reduces the cost of bankruptcy but that is phased out for individuals earning above \$100,000. A comparison of bankruptcy rates between 2009 and 2010 shows an overall increase of ten percent, but the concern is that this may be due to factors other than the law. Now consider the difference-in-difference approach, comparing the before-after change in bankruptcy rates between those making over \$100,000 with those making less. One finds that bankruptcy rates increased by eleven percent for those earning over \$100,000 and increased nine percent for those earning under \$100,000. The difference in difference is $9\% - 11\% = -2\%$. Thus it seems likely that the law did not increase bankruptcy rates at all, as in fact bankruptcy rates slightly decreased for individuals impacted by the law, relative to those unaffected.

For the difference-in-difference approach used in this Article, it is important to understand the exact nature and timing of the changes in expectations of prospective innovators and how one might expect those to be reflected in patenting decisions. Prospective filers of patents could respond to the knowledge of the law change (December 8, 1994), to the change itself, or to both. There are two subsections of 35 U.S.C. § 154 (2006) that are relevant to this inquiry. The first is the subsection specifying the patent term in the transition period:

The term of a patent that is in force on or that results from an application filed before the date that is 6 months after the date of the enactment of the Uruguay Round Agreements Act shall be the greater of the 20-year term as provided in subsection (a), or 17 years from grant, subject to any terminal disclaimers.²¹

The second is the subsection specifying the patent term after the change:

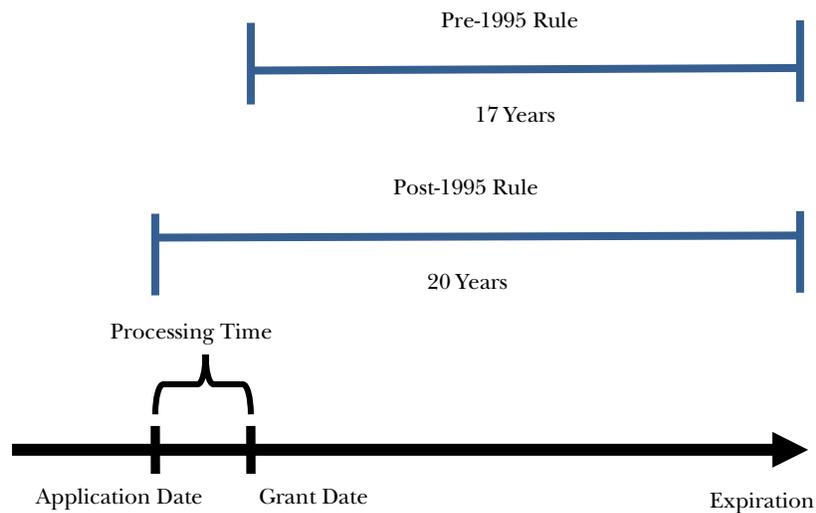
Subject to the payment of fees under this title, such grant shall be for a term beginning on the date on which the patent issues and ending 20 years from the date on which the application for the patent was filed in the United States or, if the application contains a specific reference to an earlier filed application or applications under section 120, 121, or 365(c) of this title, from the date on which the earliest such application was filed.²²

²¹ *Id.* § 154(c)(1).

²² *Id.* § 154(a)(2).

Thus, the change to the patent term had two dimensions. The first was a change in length from seventeen to twenty years. The second was a change in how the length of protection was measured—from grant date to application date, as illustrated by Figure 2. This had the important effect of varying the change in expected patent duration *depending on the processing time between application and grant date*. This is the variation that makes the difference-in-difference approach used in this Article possible. While patent protection increased on average, it only did so for those patent applications that required less than three years processing time. Patent applications that took longer than three years between application and grant date ended up with a shorter term of protection under the new law than previously.

Figure 2: Impact of TRIPS on Patent Duration



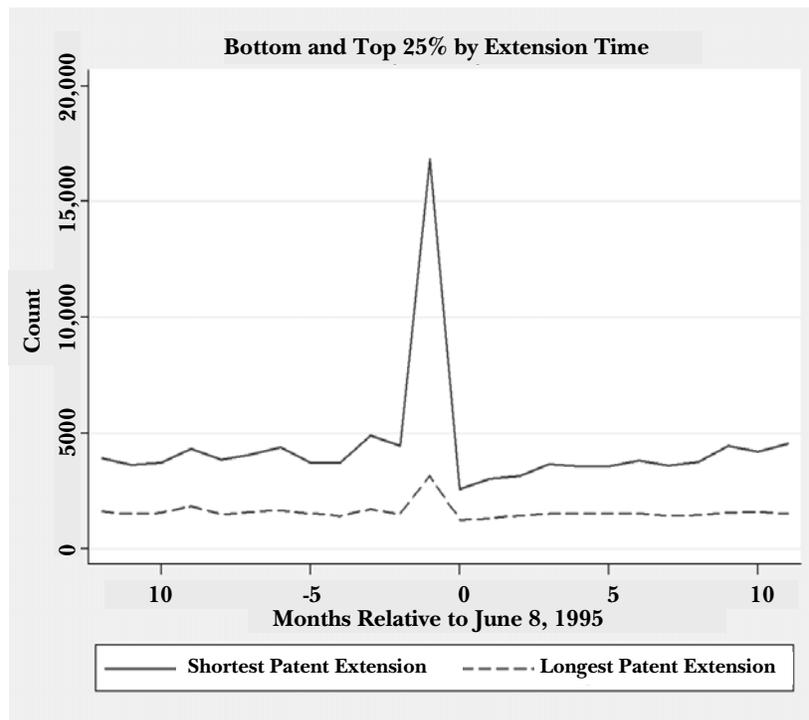
The Uruguay Round Agreements Act was enacted on December 8, 1994, making June 8, 1995, the date that the patent term changed.²³ Note that the wording of the law implicitly grants an option to those who filed their application before June 8, 1995, to receive the longer of seventeen years from grant date and twenty years from application date.²⁴ Thus, individuals who expected to have a decrease in patent

²³ See *supra* notes 18-19.

²⁴ See *supra* note 21 and accompanying text.

duration should have applied ahead of the June 8, 1995, change. Those who expected to have an increase in duration did not need to wait until after the change because they could opt into the new rule even before June 8, 1995. This should lead to a spike in applications prior to the rule change that is most pronounced among patent classes with the lowest expected extension. This prediction is borne out by the data, as can be seen in Figure 3.

Figure 3: Monthly U.S. Patent Counts, 1994–1997



B. *Related Literature*

The framework for much economic analysis of patents is laid out by William Nordhaus in his 1969 work, *Innovation, Growth, and Welfare*.²⁵ In this framework (described further in Part II), policymakers set patent terms to equalize the welfare gains from innovation due to longer patents with the welfare losses due to the grant of exclusivity. The estimation performed in this Article may be used to inform the

²⁵ NORDHAUS, *supra* note 4.

policymakers from Nordhaus's setup. Half of the calculation is provided here; the deadweight-loss calculation is left for future research.

While the Nordhaus framework is the dominant approach, there are a number of papers that suggest alternative perspectives on the function of patents. In Nordhaus's model, time of patent expiration is a monotonically increasing function of patent duration.²⁶ This is a consequence of the fact that, in the model, innovation cannot be sped up in response to changes to patent duration. John Duffy proposes a model that allows for varying innovation times that can be impacted by patent duration.²⁷ He shows that the relationship between time of patent expiration and patent duration will then have a U-shape, and thus, that the monopoly/innovation tradeoff need not exist if the patent duration is chosen optimally.²⁸ While this Article will not distinguish between Duffy's approach and that of Nordhaus, it is worth noting that Duffy's perspective leads to a different interpretation of the estimates found here.

In *Rewards Versus Intellectual Property Rights*, Steven Shavell and Tanguy Van Ypersele argue that an ideal rewards system would be superior to any kind of patent system because it would not suffer from the deadweight loss of monopoly.²⁹ They acknowledge that the practical difficulty with such a system is that the government must determine the correct award sizes.³⁰ Instead, they propose that an optional rewards system layered onto a patent system would be superior to a simple patent system in encouraging innovation.³¹ The difficulty of setting proper incentives for innovation, either through rewards or subsidies to researchers, is discussed by John Duffy and Louis Kaplow in separate papers.³² Kaplow states as the main argument for the patent system over a rewards system that it is thought to be too difficult to determine the appropriate level of reward fairly and accurately on a

²⁶ *Id.* at 71.

²⁷ John F. Duffy, A Minimum Optimal Patent Term 2-3 (Jan. 9, 2003) (unpublished manuscript), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=354282 (presenting graphs showing that "an increase in patent terms does not necessarily involve this tradeoff").

²⁸ *Id.*

²⁹ Steven Shavell & Tanguy Van Ypersele, *Rewards Versus Intellectual Property Rights*, 44 J.L. & ECON. 525, 530 (2001).

³⁰ *Id.* at 541-42.

³¹ *Id.* at 537-41.

³² John F. Duffy, *The Marginal Cost Controversy in Intellectual Property*, 71 U. CHI. L. REV. 37 (2004); Louis Kaplow, *The Patent-Antitrust Intersection: A Reappraisal*, 97 HARV. L. REV. 1813 (1984).

“case-by-case basis.”³³ Michael Abramowicz suggests a way to make a system of prizes more practical by obfuscating the determination of reward and thereby reducing the inefficient rent-seeking behavior that prizes can cause.³⁴

Several recent papers suggest possible shortcomings of the basic economic model of patents. Alexander Tabarrok points out that although modern patent systems are premised on the notion of paying back innovators’ high fixed costs through grants of monopoly, in reality the law ignores fixed costs.³⁵ He proposes modifications to patent law to address this disconnect.³⁶ Nancy Gallini reviews the literature on the issue of the relationship between strength of patents and quantity of innovation.³⁷ She discusses recent theories that differ from the Nordhaus model by accounting for follow-on inventions, and she does not predict an unambiguously positive relationship between patent duration and innovation. She notes that “[e]xtending patent life may increase an entrant’s incentives to introduce an imitation during the patent period,” with the result being that “incentives to innovate may decline with increases in patent life.”³⁸ Gideon Parchomovsky and R. Polk Wagner suggest another shortcoming of the standard model in the modern environment: its failure to account for portfolios of patents.³⁹ They argue that the relationship between value and number of patents is nonlinear among holders of patent portfolios, and thus even negative-expected-value patents could be pursued by such entities.⁴⁰ In this environment, it is unclear what impact increased duration of protection would have.

Empirical research on the incentive effects of patent duration is quite difficult due to the scarcity of policy variation. Josh Lerner has assembled a large cross-country dataset covering 150 years of IP

³³ Kaplow, *supra* note 32, at 1844.

³⁴ Michael Abramowicz, *Perfecting Patent Prizes*, 56 VAND. L. REV. 115, 218-24 (2003).

³⁵ Alexander Tabarrok, *Patent Theory Versus Patent Law*, CONTRIBUTIONS TO ECON. ANALYSIS & POL’Y, 2002, at 1-2, <http://www.bepress.com/cgi/viewcontent.cgi?article=1039&context=bejeap>.

³⁶ *Id.* at 8-11.

³⁷ Nancy T. Gallini, *The Economics of Patents: Lessons from Recent U.S. Patent Reform*, 16 J. ECON. PERSP. 131 (2002).

³⁸ *Id.* at 136.

³⁹ Gideon Parchomovsky & R. Polk Wagner, *Patent Portfolios*, 154 U. PA. L. REV. 1 (2005).

⁴⁰ *Id.* at 28-29.

changes.⁴¹ Although he finds little impact of strength of patent protection, the study suffers from the standard difficulties of cross-country comparisons and endogeneity. He attempts to confront the endogeneity concern by instrumenting for patent law changes with international agreements (including TRIPS) and still finds no significant impact on the number of patents filed.⁴²

Petra Moser also uses historical data in one of the most creative recent papers on innovation, in which she studies nineteenth century World's Fairs.⁴³ By collecting data on the exhibits at the World's Fairs and nineteenth Century IP laws, she attempts to determine the impact that the laws have on field and magnitude of innovation. She finds significant cross-country differences in the distribution of innovation by technological field that grow stronger with increased patent duration.⁴⁴ She also finds some evidence that there is a diminishing marginal incentive effect of patent duration.⁴⁵

Empirical work on patent duration is relatively scant,⁴⁶ and work focusing on the impact of TRIPS in the United States is almost nonexistent save for an early piece by Mark Lemley.⁴⁷ In that piece, he collects data from a single issue of the PTO's *Official Gazette* shortly after passage of the TRIPS legislation (but before its implementation) in order to evaluate the likely impact of the law change on patent duration. He finds that the mean patent duration should increase, although not uniformly across classes (a finding that is subsequently borne out in the data).⁴⁸ He further predicts that processing time should decrease due to increased incentives for patent attorneys to re-

⁴¹ Josh Lerner, *Patent Protection and Innovation over 150 Years* (Nat'l Bureau of Econ. Research, Working Paper No. 8977, 2002), available at <http://www.nber.org/papers/w8977>.

⁴² *Id.* at 26-28.

⁴³ Petra Moser, *How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs*, 95 AM. ECON. REV. 1214 (2005).

⁴⁴ *Id.* at 1216.

⁴⁵ *Id.* at 1224.

⁴⁶ Adam Jaffe also examines recent developments in the U.S. patent system and surveys both the theoretical and empirical literature. While he touches on patent scope, he does not discuss any recent work on patent duration or the impact of the TRIPS agreement, perhaps because there has been relatively little. Adam B. Jaffe, *The U.S. Patent System in Transition: Policy, Innovation, and the Innovation Process*, 29 RES. POL'Y 531 (2000).

⁴⁷ Mark A. Lemley, *An Empirical Study of the Twenty-Year Patent Term*, 22 AIPLA Q.J. 369 (1994).

⁴⁸ *See id.* at 392-93.

spond to office actions more quickly.⁴⁹ This prediction does not appear to be borne out by the data, perhaps due to an offsetting increase in PTO processing times. Nevertheless, the paper is an important early contribution to the understanding of the impact of TRIPS in the United States.

C. Theoretical Background

So far the discussion has focused on the number of innovations, but we now turn to the magnitude. Following Nordhaus,⁵⁰ one may write down a simple model of the profit-maximizing inventor. Under some rough assumptions, we find that magnitude of innovation should increase with increasing duration of patent protection. The exact timing of the response depends on the lag in time between research inputs and patentable outputs. If the lag is less than six months, then we should expect to see an increase in value of patent applications beginning on June 8, 1995. If the lag is longer, we should expect to see the increased value of patents occurring later. This leads to a testable prediction (although one not examined in this Article): those industries with longer lags should see an increase in citation-weighted patents later than those in industries with shorter lags. But the key testable implication is straightforward: there should be an increase in the level of innovation—by whatever measure—following the implementation of TRIPS due to the overall term extension.

Another important approach to evaluating the shift in innovation around the TRIPS law change is by making use of citation data. Previously, the analysis has made use of patent counts alone, which may be a good proxy for aggregate innovation under certain assumptions. For example, if we assume that there is no correlation between innovation quantity and marginal patent value, then patent counts are an adequate proxy for aggregate value. However, it is easy to imagine instances in which this would not be the case, such as when a firm can either produce two patentable goods worth one util each or one patentable good worth two utils. In this case, marginal patent value is inversely correlated with quantity of innovation, and aggregate value is uncorrelated with patent count.

⁴⁹ See *id.* at 415-16.

⁵⁰ See NORDHAUS, *supra* note 4, at 70-75 (establishing a framework for the economics of patents).

One indicator that should correlate with a patent's value is how frequently it is cited by subsequent patents. Presumably, patents with greater value, whether private or social, will have a greater impact on future inventors, and this should be reflected in their citations. There is literature on the relationship between patent citations and value that goes back at least to Manuel Trajtenberg's 1990 RAND paper.⁵¹ In the paper, he relies on his previous research, in which he uses a structural approach to estimate the demand system for attributes of Computed Tomography (CT) scanners.⁵² The estimated parameters of the demand system are used to calculate a social value for innovation in CT-scanner technology. He tests the correlation between these values and simple patent counts as well as citation-weighted patent counts, and he finds a substantially stronger correlation with the latter.⁵³ He further finds that the correlation increases when the citation index used is slightly convex, indicating an increasing return to citations.⁵⁴ The empirical analysis in this Article focuses on both patent counts and citation-weighted counts.

II. DATA

All data originally come from the U.S. Patent and Trademark Office, but was obtained for this research from two sources. Polk Wagner kindly made available data on all U.S. patents granted between 1976 and fall 2008, with information including patent class, processing time, and application date. Through the NBER-patent citation database, I obtained data on all U.S. patents granted between 1963 and 2002, along with citations made to those patents between 1975 and 2002.⁵⁵ These data include such fields as patent category, number of

⁵¹ Trajtenberg, *supra* note 9. For further literature, see other sources cited *supra* note 9.

⁵² See *id.* at 177-78 (discussing demand for CT scanners as analyzed in Manuel Trajtenberg, *The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners*, 97 J. POL. ECON. 444 (1989)).

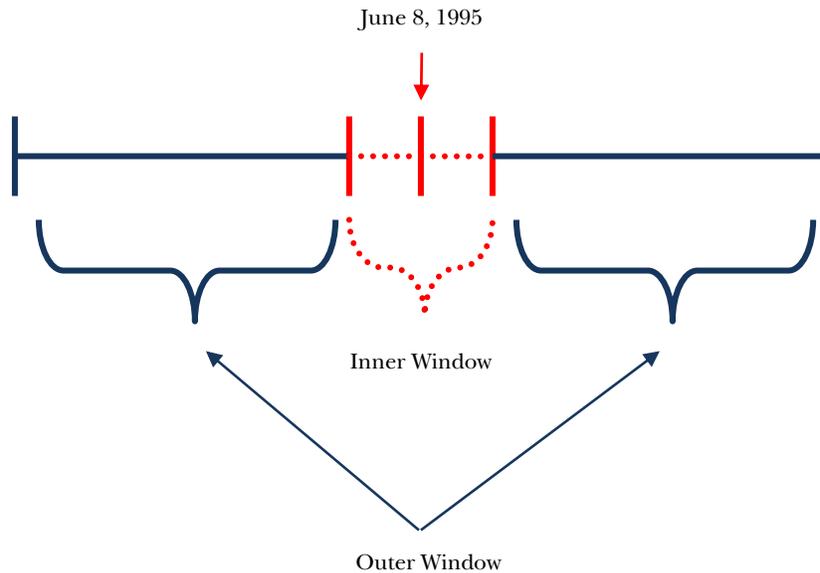
⁵³ See *id.* at 180-81.

⁵⁴ That is, when using a weighting factor of the form c^n , Trajtenberg finds a higher correlation between the citation index and the measure of patent value when $n > 1$. *Id.* at 182.

⁵⁵ To view this data file in detail, see The NBER U.S. Patent Citations Data File, <http://www.nber.org/patents> (last visited Apr. 15, 2009). For the source of this database, as well as a description of its enormity and an analysis of its implications, see Bronwyn H. Hall, Adam B. Jaffe & Manuel Trajtenberg, *The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools* (Nat'l Bureau of Econ. Research, Working Paper No. 8498, 2001), available at <http://www.nber.org/papers/w8498>.

claims, number of citations, and information on the assignee, inventor, and year (but not date) of application. Because I am interested in a change occurring in 1995, I chose various windows around that year for the analysis, the largest of which ranges ten years around June 8, 1995, as depicted in Figure 4.⁵⁶ Summary statistics of the data can be found in Tables 1 and 2 in the Appendix.

Figure 4: Windows for Patent Count/Citation Analysis

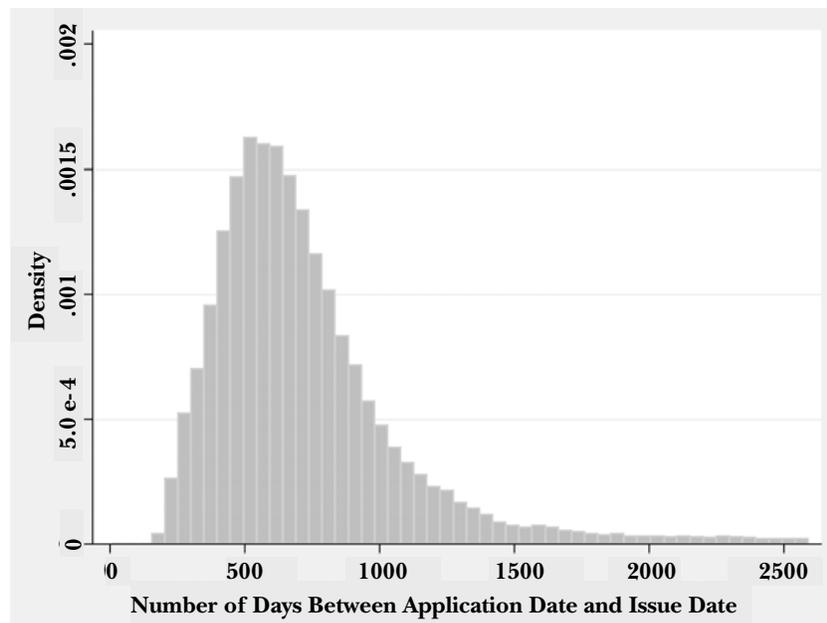


The data summarized and used in the regressions include over one million patent applications that were submitted between 1990 and 2000 and subsequently granted. I excluded patent classes that did not receive at least thirty subsequently granted applications in each year between 1990 and 2000 to ensure that the calculations are not overly influenced by outliers. Imposing this restriction excludes approximately 200,000 observations. The vast majority of these observations are actually in patent classes that are either established or abolished in the decade of interest, which results in zero observations for some years. These naturally tend to be classes relating to new or obsolete

⁵⁶ This Article only reports data for outer windows up to two years, but the results from larger windows are available from the author.

technology,⁵⁷ and patenting behavior in these classes may be very different from behavior with regard to established technology. In particular, it is likely that many other factors beyond duration of patent protection have substantial influence in these fields; thus, excluding them allows for a more focused examination of the phenomenon in question. There may also be a truncation-bias concern: namely, that patents with long processing times would be represented at a lower rate because the data were obtained in fall 2008. However, as Figure 5 indicates, a very small fraction of patents have processing times as long as eight years, so this should not have a significant impact on the results.

Figure 5: Processing Time Two Years Prior to Law Change



One can see in Table 1 that the number of claims per patent application increases over the 1990s, with the median patent application having one extra claim in the latter half of the decade.⁵⁸ “Forward citations” refers to the number of times that a patent was cited by subsequent applications. I have data on citations contained in applications submitted

⁵⁷ See *infra* app. tbl.3.

⁵⁸ Note that there are approximately one-third the number of observations for claims for the five years following the law change as there were before. This is an artifact of the NBER data set, which is missing claims data for most of the later observations.

through December 31, 2002. This measure will monotonically increase with time, and thus it is not surprising to note that the mean number of forward cites decreases from 6.4 to 1.8 from the period before the law change to the period afterwards. The earlier group had an average of ten years to collect citations, versus five years for the later group.

In order to compare the value of patents, something must be done to normalize the citations data. The approach I take in this Article is a fixed-effects approach, predicated on the assumption that the expected present value of all patents in a month is time invariant. Thus, the mean patent issued in July 1999 should have the same value as that issued in February 1991. This should be reflected by their having the same number of citations, if the number of citations were computed n years after the respective grant dates for each mean patent. Of course, at any point in time, the mean 1991 patent will have more citations to it than the mean 1999 patent simply because it has had more time to accumulate them. In the data, the citation counts are based on the total number of citations each patent had received as of December 31, 2002. Using the fixed-effects approach, I normalize the citation counts by dividing the actual number of patents by the mean number of citations received by all patents in that month. The advantage of this approach is that it allows for a nonparametric functional form for the citation-generating process. One disadvantage is that the renormalized distribution of patents will not necessarily have the same shape for all subsets, largely due to the fact that patents with zero citations will still have zero after renormalization.

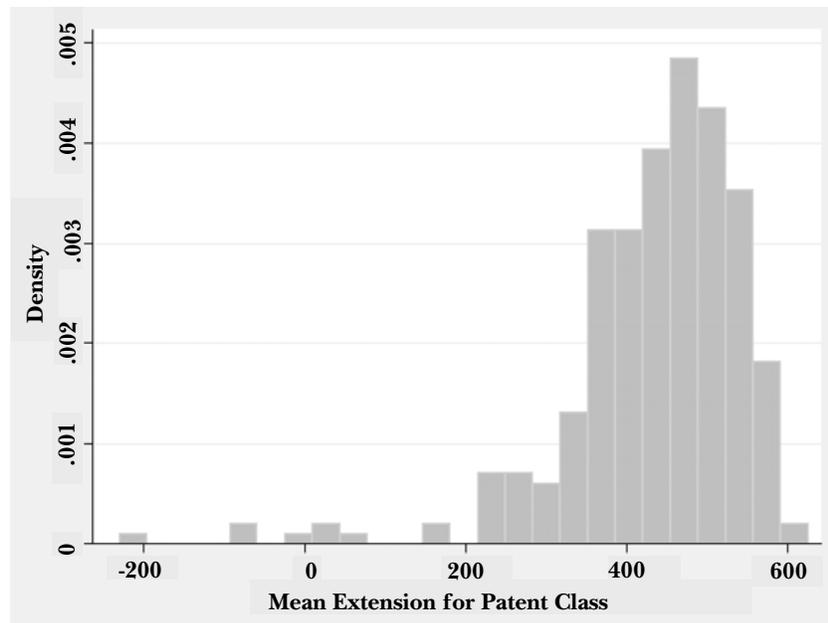
A comparison of the summary statistics for the new measure of forward citations in panels A and B of Table 1 bears this out. By construction, the mean number of normalized forward citations is the same for both periods, but the median is substantially lower after the law change. This is a direct consequence of the aforementioned problem with zeroes.

The key independent variable in this analysis is the expected length of term extension for a prospective innovator.⁵⁹ It is important that there is enough dispersion in this variable that there is meaningful variation with which one may hope to identify an impact. Panel A of Table 1 indicates that the standard deviation of the term extension is approximately one year, so this study will mostly be limited to identifying the impact of relatively small percentage changes in patent duration (around five to ten percent). This is a potential concern of the study design and is discussed further in Part V. Figure 6 shows the full

⁵⁹ Its construction is described in detail *infra* Part III.

left-skewed distribution of the term extension by patent class. Almost all patent classes receive an increase in expected duration due to TRIPS. Even though there are few classes that receive actual reductions in expected duration, the variation in the regression analysis comes from *differences* across classes in patent extension, so the actual magnitude of the extension is unimportant.

Figure 6: Patent Extension



A comparison between the processing time for the full data set and the subset from after the law change indicates that PTO processing time is increasing. This should not present a major problem for this analysis unless the increase in processing time is anticipated and correlated with expected extension. That is a real possibility, however, and its impact on this analysis is discussed in Part V.

The researchers who assembled the NBER dataset have created a cruder categorization of patents than of classes, which are determined by the PTO and number close to one thousand.⁶⁰ These categories are used as controls for some of the analyses, as well as to look at broad subsets of the data. In Table 2, summary statistics are presented by

⁶⁰ See Hall, Jaffe & Trajtenberg, *supra* note 55, at 41-42 app. 1 (listing the categorizations of patent classes).

category on term extension, citations, and simple counts. One point that is apparent from the table is that there is substantial variation in the impact of TRIPS even across very broad technological categories (as measured by term extension). There is also substantial heterogeneity in citations by category, with computer and communication patents receiving the most on average. Finally, the categories—and the patent classes—differ in the total number of patents and their growth rate over the period examined. To account for this, some of the specifications described in the next section include class-specific time trends. This should isolate the changes in innovation that are due to the heterogeneous influence of TRIPS on patent-term extensions.

III. ECONOMETRIC METHODOLOGY

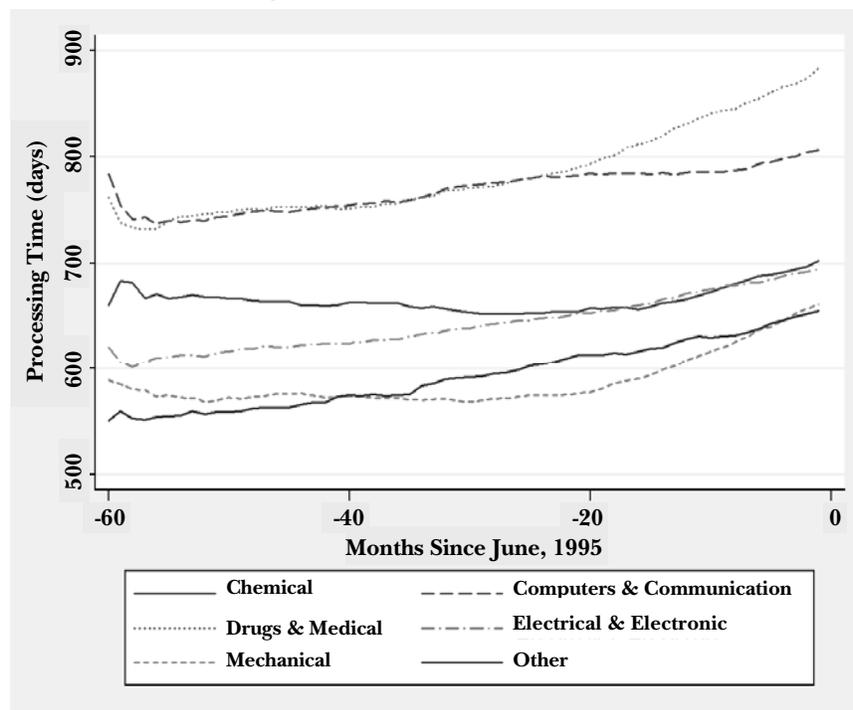
The specifications used to empirically evaluate the impact of the TRIPS law change follow from the theoretical framework discussed *supra* Section I.C. The first approach uses patent count as a proxy for innovation. An increase in duration of patent protection should lead to an increase in the number of patents following implementation of the TRIPS agreement. Further, the increase in patents should be greatest in those classes that receive the longest extensions. To test this, I make use of a difference-in-difference framework. The first difference is time: that before and that after the June 8, 1995, change of patent term (as indicated below by the *After_m* dummy variable). The second difference is cross-sectional and makes use of the variation across patent classes with regard to the expected effect of the law change on patent duration. Innovators with patents in classes with long processing times can expect a short extension or even a decreased length of protection in extreme cases. Those with patents in classes with rapid PTO processing times should expect a substantial extension because of the law change.

A question then arises: how do innovators form expectations about patent duration after the law change? The answer must be a function of PTO processing time, patent class, and perhaps other idiosyncratic effects. In this Article, I make the simplifying assumption that the expected patent duration is solely a function of recent PTO processing times within a given patent class. Thus, I am able to calculate the expected patent extension (or reduction) due to the law change.

Even making this assumption, there are numerous ways to calculate the expected extension. To gain some insight into the problem, it is instructive to understand the stability of PTO processing times. Figure 7 presents processing time by technological category averaged

over the previous two years.⁶¹ The figure indicates that, for most categories, there is little change in the mean PTO processing time, with the possible exception of the Drugs and Medical category. This category experienced the most rapid increase in processing time of any category over the full period examined, and it occurred in the two years prior to the implementation of TRIPS. Even this rapid increase amounted to about one hundred days over two years. Nevertheless, the categories keep their relative ordering over this period. To calculate the expected extension, I first calculate the mean processing time (averaged over a lagged two-year period) for each patent class. The value of this variable on December 8, 1994, is used for all classes.⁶² This is then subtracted from the three-year change in patent protection to obtain the variable *Extension_c*. The distribution of this variable can be found in Figure 6.

Figure 7: Processing Time by Technological Category Averaged over the Previous Two Years



⁶¹ That is, each point is the average of the PTO processing time for all subsequently granted patents applied for within the previous two years.

⁶² This date is used because it is the earliest time in which inventors could have reasonably responded to the forthcoming change in the law and thus the earliest that they could have formed their expectations.

For the count regressions, I estimate the following base specification on monthly patent counts by patent class, which I will refer to as Specification (1):

$$P_{mc} = \alpha + B_1 \text{After}_m + B_2 \text{Extension}_c + B_3 \text{After}_m * \text{Extension}_c + B_4 t + X_{mc} + E_{mc} .$$

Here m denotes month and c denotes patent class. P_{mc} is the patent count and After_m is a dummy variable that is zero before June 8, 1995, and one otherwise. Extension_c is the expected extension by class, t controls for a linear time trend, and X_{mc} is a set of time-varying controls (not included in the base specification). These controls include the type of entity to which the patent was assigned, generality of patent,⁶³ number of claims made, patent subcategory, and a quadratic function of time that can vary by pre-post period.

There may be a concern that there are differences across patent classes in their rates of growth of patenting. To address this concern, I run a specification of the following form, deemed Specification (2), which allows for class-specific time trends:

$$P_{mc} = \alpha + B_1 \text{After}_m + B_2 \text{Extension}_c + B_3 \text{After}_m * \text{Extension}_c + B_c \text{Class}_c t + X_{mc} + E_{mc} .$$

The regressions are run on variable windows of data, with both an inner and outer window. The inner window is used to exclude data right around the change, which is likely to be impacted by the short-term effects discussed above. The outer window is varied in order to allow for tradeoffs between greater data and greater potential contamination from long-run secular trends (which may vary by patent class).

The second major approach to analyzing the impact of TRIPS on innovation uses patent citations, rather than simple counts, as an indicator of value. Citation data poses a challenge not present with simple-count data, because, while each granted patent has a weight of one at any time, citations are monotonically increasing over time. In Part II, I discussed the approach that I have taken to computing a normalized citation count.

Once the renormalized citations are computed, I take two different approaches to analyzing the impact of TRIPS using citations as an indicator of patent value. The first approach is analogous to Specifi-

⁶³ See Hall, Jaffe & Trajtenberg, *supra* note 55, at 21-23 (suggesting the generality of a patent as an indicator of its widespread impact in the field).

citations (1) and (2), run at the month-class level, where citations replace patent count as the dependent variable. These specifications lose the patent-level variation, so I also analyze the impact of the law change on citations using patent-level data in Specification (3) below:

$$Cite_{ic} = \alpha + B_1 After_i + B_2 Extension_c + B_3 After_i * Extension_c + B_4 t + X_{ic} + E_{mc}.$$

Although the relationship between citations and extension is likely to be somewhat nonlinear, I use a linear model here because it is likely to be a good approximation over the time period examined. The final specification combines information from citations with the count data. I run regressions using counts as the dependent variable (as in (1) and (2)) but weighted by citations. I discuss the results of the estimations of these regressions below.

IV. RESULTS

The main findings from the regression analysis show that there was a statistically significant change in the number of patents applied for following the TRIPS agreement, but no significant difference in mean citations per patent. Additionally, when examining citation-weighted patent counts, the impact of TRIPS is found to be significant.⁶⁴ These results are reported in Table 4 (located in the Appendix) and Figures 8 through 10 below.

Table 4 contains results from three types of regressions: those with patent counts as the dependent variable, patent counts weighted by citations, and citations. The first two columns in Table 4 present results from the most basic specifications, leaving out almost all controls. In the first column (as with the next five), the dependent variable is monthly patent counts, by patent class. This regression is run on a dataset with a six month outer window and two month inner window. The coefficient of interest (on the *After * Extension* interaction term) is significantly different from zero ($p < 0.01$). This indicates that patent classes with longer extensions due to TRIPS tended to have a greater increase in patents following TRIPS than those classes with shorter extensions. The second column reports results from the same regression run with a twelve-month outer window and

⁶⁴ Here and elsewhere I may use causal language about “the impact of the TRIPS agreement.” This language is used for simplicity, since the changes are not necessarily due to TRIPS. This is discussed further *infra* Part V.

two-month inner window. The findings are substantially the same, although the coefficient of interest is somewhat smaller in this specification (but still significantly different from zero).

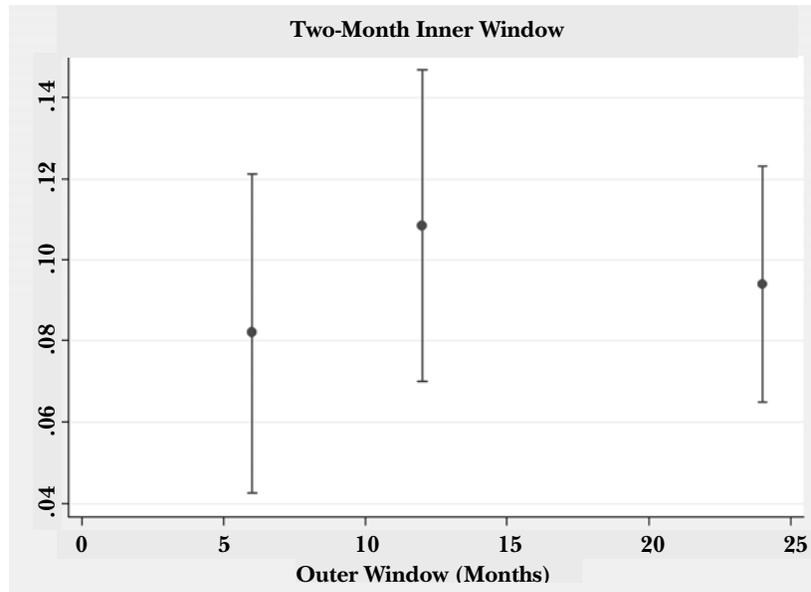
The coefficient on *After* in all of the regressions with patent count as the dependent variable is negative and significantly different from zero at $p < 0.01$. At first blush, this negative coefficient may be counterintuitive. But as might be expected from the option-like nature of the law change, there was a spike in applications just prior to the implementation of the law, followed by a subsequent dip, as shown in Figure 3. This is what leads to the negative coefficient on *After*; it is precisely this window upon which the analysis is focused. The magnitude of the extension also has a negative relationship with patent counts. In some specifications it is statistically significant, and in others it is not. The interpretation here must be that patent classes that receive longer extensions tend to have lower patent counts.

Column three of Table 4 adds a larger set of controls, as compared to columns one and two. The controls include measures of the citing patents, including the mean year of the citing patents and the mean number of cites to the citing patents (a measure of their importance). There are two controls that make use of a measure of patent generality: one of the citing patents and one of the patents in question.⁶⁵ Two additional controls are included: the number of claims made on the patent and the number of “parents.” The control variables increase the R^2 of the regression from 0.16 to 0.46 for the twelve-month outer window. The finding from this specification is the same as the first two: there is a statistically significant coefficient on the difference-in-difference interaction term.

Columns four and five present regressions of the form described above by Specification (2). The main difference with the previously described regressions is the addition of class-specific time trends. These trends account for the possibility that classes have varying trends in patenting that are unconnected to the TRIPS agreement. The difference-in-difference coefficients estimated in these specifications are somewhat larger than those estimated in the base specification. These coefficient estimates (along with one for a twenty-four month outer window) are presented visually in Figure 8. The point estimates and standard errors are plotted and are relatively consistent.

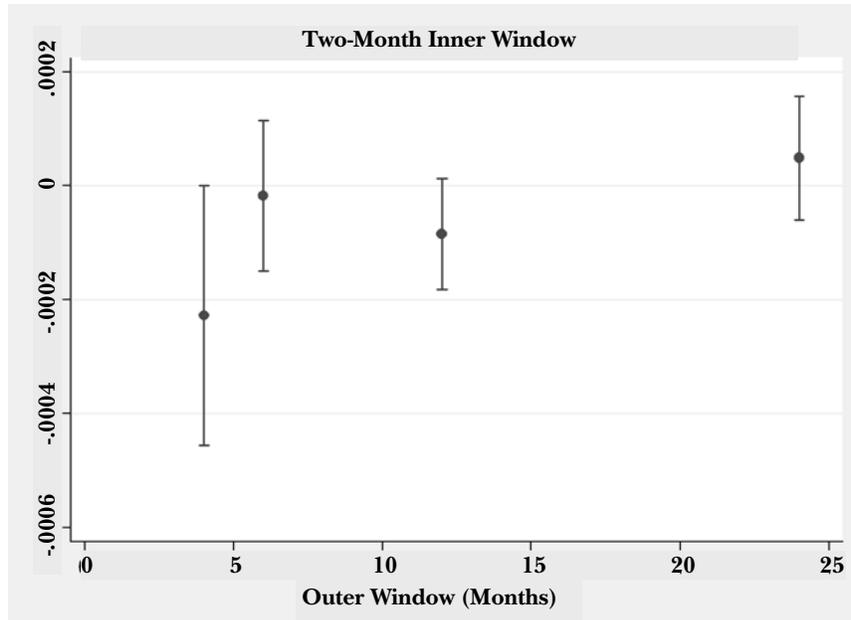
⁶⁵ See Hall, Jaffe & Trajtenberg, *supra* note 55, at 21 (suggesting that a high generality score indicates widespread impact of the patent).

Figure 8: Coefficient Estimates



Specification (3) above describes the regressions examining the impact of the TRIPS agreement, not on patent counts, but on citations per patent. As discussed above, this is likely a better measure of the main quantity of interest—innovative output.⁶⁶ Columns seven and eight in Table 4 and Figure 9 present results from the estimation of Specification (3). Unlike the regressions discussed previously, column seven presents results from analysis of patent-level data. For the six-month outer window (and all other windows not reported), the coefficient on the *After * Extension* interaction term is statistically indistinguishable from zero. This is also the case when the data are analyzed at the class-month level, as in column eight. Figure 6 displays the coefficient and the standard error from this regression, along with three others with varying outer windows. All point estimates are insignificantly different from zero, with standard errors for most around 0.0001. In Part V, I return to these findings and discuss the implication of the zero coefficients.

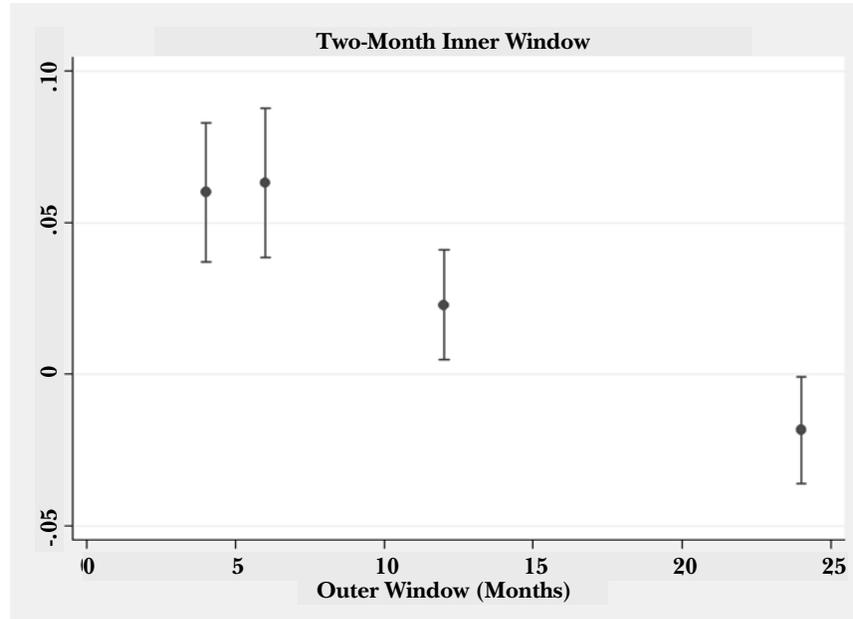
⁶⁶ See *supra* note 7 and accompanying text.

Figure 9: Results from the Estimation of Specification (3)

The final approach that I take to valuing innovation using patents employs citation-weighted patent counts. This approach is similar to those described in columns one through five of Table 4. The dependent variable is still patent count, but now each patent is weighted by the number of citations it had received as of December 31, 2002 (normalized according to the procedure described *supra* Part II). This is the preferred specification because it should get closest to the object of interest here: a measure of the value of innovation.

The results from the citation-weighted regressions are presented in column six of Table 4 and in Figure 10. When we examine relatively smaller windows, as in Table 4 or the first two observations in the figures, we find a statistically significant coefficient on the interaction term of around 0.06. The figure indicates that the size of the window is important, because the coefficient loses statistical significance as the outer window increases beyond half a year. This mixed finding is discussed further in Part V.

Figure 10: Results from Citation-Weighted Regressions



V. DISCUSSION

In order to gain a better understanding of the regression results, it is useful to perform some back-of-the-envelope calculations in order to understand the magnitude of the law changes' impact. The preferred specification is the citation-weighted regression, shown in column six of Table 4. The coefficient on the interaction term is 0.063, meaning that a one-day-greater extension for a patent category is associated with a 0.063 more patents (citation weighted) after the implementation of TRIPS. Although Trajtenberg found somewhat increasing returns to citations, to be conservative I will assume a linear relationship between value of innovation and citation-weighted counts.⁶⁷

Let us consider the increase in value of innovation due to a one-standard-deviation increase in patent-term extension. The standard deviation of the term extension (by class) is 114 days (see Figure 6 for the full distribution). Multiplying this by the coefficient above, we

⁶⁷ Trajtenberg, *supra* note 9, at 182-83.

find that a one-standard-deviation increase in patent term extension is associated with an increase of about seven monthly patents. From a mean of approximately thirty-four, in percentage terms, this comes to a twenty-one percent increase in value of innovation—a very substantial increase. It seems unlikely that the deadweight loss due to exclusive rights would be enough to offset this considerable gain, suggesting that an increase in patent terms could lead to greater welfare.

There are several impediments to drawing clear policy implications from this exercise. These include concerns about outliers, unobserved variation, misspecification, and external validity. The magnitude of the estimated effect seems inordinately high, given that the extensions are relative to base protections of seventeen years, and, thus, the total extension is only on the order of seven percent (which would imply an elasticity of around three).

One potential explanation for the large point estimates is that the results are driven by a few highly affected classes. Preliminary analysis indicates that it is likely that biological patents are responsible for the bulk of the observed impact of TRIPS. This finding is strengthened by the fact that the pharmaceutical industry was one of the industries most ardently opposed to TRIPS due to concern over shortened patent duration. A more detailed analysis of the impact of TRIPS on pharmaceutical patents is beyond the scope of this Article, but is the subject of future research.

Unobserved variables are another potential roadblock to drawing a causal inference of the impact of patent-term extension. I attempt to control for as many patent characteristics as possible, but there is still a relatively large amount of unexplained variation. Of particular concern is if there is a variable that is correlated with both term extension and patent counts because the law change came with six month's advance notice. One mechanism by which this might occur is if innovators in patent classes with short expected extensions rushed to patent under the old law, thus creating a spike in those applications prior to the change. Then those classes would show an artificially low number of applications immediately after the law change, as many that would have been filed at that time had already been filed. This is exactly the pattern we see in Figure 3. Of course, this is the rationale behind excluding an inner window around the law change. But if the inner window is not long enough, the estimate will be biased upward.

Another potential problem with the interpretation of the findings is the possibility that innovators will form expectations about the patent-term extension differently from that assumed. In order to test the

explanatory power of the lagged processing-time variable, I regressed actual term extension on the computed extension. I found that the computed extension increases explanatory power by eight percent relative to a regression on the other control variables (R^2 increased from 0.114 to 0.195). If the deviations between innovators' beliefs about the extensions and the calculated extensions used in the analysis are correlated, then the estimates provided will overstate the magnitude of the impact in equilibrium, where we assume that people have rational expectations. Alternatively, people may have rational expectations even in the transition, and PTO processing times may be changing in a way that is correlated with prior processing times (i.e., classes with longer initial processing times also have a greater *increase* in processing time than those with shorter initial processing times).

There are two final concerns about how well this exercise meets the aim of shedding light on the relationship between patent duration and the value of innovation. The first concern is whether the data used are on *granted* patent applications. Ideally, one would prefer data on all applications. If the grant rate changed differentially by magnitude of extension around the time of the TRIPS agreement, this could lead to a spurious result. A second concern is how well patent counts—or even citation-weighted counts—actually correspond to value of innovation. Much of the literature making use of citation-weighted counts is based on a small number of studies that try to obtain external measures of patent value, generally on a small dataset. To the extent that this relationship is not general, one cannot draw any strong conclusions about value of innovation here.

CONCLUSION

Understanding the incentive effects of patent protection is a core issue in intellectual property scholarship, about which almost nothing is currently known. This Article seeks both to advance our knowledge of the relationship between value of innovation and duration of patent protection and to point the way toward further research.

The TRIPS agreement was the biggest change in U.S. patent protection in over forty years.⁶⁸ It altered not only the mean expected length of patent protection but also the method by which it is calculated—from grant date to application date.⁶⁹ This aspect of the law

⁶⁸ See *supra* note 1.

⁶⁹ See *supra* text accompanying notes 21-22.

had a heterogeneous impact across patent classes, depending on the expected PTO processing time. I use this heterogeneity and compare patent counts, citations, and citation-weighted counts before and after the implementation of TRIPS.

My findings indicate that patent classes with longer extensions are associated with a statistically significant increase—relative to patent classes with shorter extensions—in patent counts and citation-weighted counts after the law change. There does not, however, appear to be any statistically significant association with mean citations per patent. Although this finding is striking, it must be understood in light of several potential confounds that may otherwise cause it to be overstated. While this study cannot conclusively determine the relationship between the duration of patent protection and the value of innovation, the insights gained here point toward even more powerful analyses for future research.

APPENDIX

Table 1: Summary Statistics

Panel A: Before Law Change				
Variable	Mean	Median	Std. Dev.	Observations
Number of Claims	13.51	11.00	11.05	496,920
Forward Cites (through 2002)	6.40	4.00	9.85	509,672
Forward Cites (FE adjustment)	1.00	0.56	1.51	509,672
PTO Processing Time (years)	1.89	1.66	1.08	510,890
Term Extension (years)	1.11	1.34	1.08	510,890

Panel B: After Law Change				
Variable	Mean	Median	Std. Dev.	Observations
Number of Claims	14.18	12.00	10.87	167,577
Forward Cites (through 2002)	1.81	1.00	3.65	605,481
Forward Cites (FE adjustment)	1.00	0.29	1.93	605,481
PTO Processing Time (years)	2.32	2.06	1.15	660,330

Panel C: All Data				
Variable	Mean	Median	Std. Dev.	Observations
Number of Claims	13.68	11.00	11.01	664,497
Forward Cites (through 2002)	3.91	2.00	7.54	1,115,153
Forward Cites (FE adjustment)	1.00	0.47	1.75	1,115,153
PTO Processing Time (years)	2.14	1.88	1.14	1,171,220

Note: Summary statistics for all patents granted from 1990 through 2000, in classes that received at least thirty applications per year that were eventually granted. Some claims data are missing in original NBER patent files, with no claims data reported after 1998. 290 total patent classes were used.

Table 2: Summary Statistics by Technological Category

	Term Extension (Years)	Forward Citations	Observations	
			Before	After
Drugs and Medical	0.65	4.50	64,476	79,137
Computers and Communication	0.83	5.53	49,715	87,177
Chemical	1.13	3.29	99,909	95,747
Electrical and Electronic	1.17	4.29	87,751	111,541
Other	1.32	3.29	103,951	117,390
Mechanical	1.33	3.36	103,740	114,347

Note: Summary statistics by broad technological category (using the NBER patent-classification system) for all patents granted from 1990 through 2000 in classes that received at least thirty applications per year that were eventually granted. Term extension is calculated based on mean PTO processing time in that category prior to the law change. “Before” and “After” are relative to date of patent-term change, June 8, 1995.

**Table 3: Patents with Fewer Than Thirty Applications in
at Least One Year Between 1990 and 2000**

Current Title	PTO Class	Patents from 1990 to 2000	Year Established	Year Abolished
Information-Processing-System Organization	395	26,831	1991	2000
Semiconductor-Device Manufacturing: Process	438	18,654	1997	-
Electrical Computers and Data Processes	364	12,392	1977	1999
Incremental Printing of Symbolic Information	347	8236	1994	-
Electrical Computers and Digital-Processing Systems: Multicomputer Data Transferring	709	7735	1999	-
Semiconductor Device Manufacturing: Process	437	7322	1987	1997
Data Processing: Database and File Management or Data Structures	707	7214	1997	-
Miscellaneous Active Electrical Nonlinear Devices, Circuits, and Systems	327	7121	1994	-
Electrical Computers and Digital-Processing Systems: Memory	711	5083	1997	-
Electrophotography	399	4921	1996	-
Error Detection/Correction and Fault Detection/Recovery	714	4490	1999	-
Electrical Computers and Digital-Processing Systems: Support	713	4149	1999	-
Data Processing: Financial, Business Practice, Management, or Cost/Price Determination	705	4090	1997	-
Electrical Computers and Digital Data-Processing Systems: Input/Output	710	4088	1999	-

Current Title	PTO Class	Patents from 1990 to 2000	Year Established	Year Abolished
Photography	396	4061	1996	-
Abrading	451	3965	1994	-
Games Using Tangible Projectile	473	3916	1993	-
Liquid-Crystal Cells, Elements and Systems	349	3619	1996	-
Data Processing: Vehicles, Navigation, and Relative Location	701	3559	1997	-
Data Processing: Speech-Signal Processing, Linguistics, Language Translation, and Audio Compression/Decompression	704	3332	1997	-
Cleaning Compositions for Solid Surfaces, Auxiliary Compositions, or Processes of Preparing the Compositions	510	3079	1996	-
Electronic Digital-Logic Circuitry	326	3000	1994	-
Data Processing: Measuring, Calibrating, or Testing	702	2869	1998	-
Error Detection/Correction and Fault Detection/Recovery	371	2802	1980	1999
Photography	354	2536	1973	1996
Television-Signal Processing for Dynamic Recording or Reproducing	386	2509	1996	-
Data Processing: Generic Control Systems or Specific Applications	700	2358	1999	-
Electrical Computers and Digital-Processing Systems: Processing Architectures and Instruction Processing	712	1987	1999	-
Etching a Substrate: Processes	216	1630	1995	-

Note: This table presents a partial list of patent classes dropped due to having insufficient observations in a given year. All of the most populous dropped classes were either created or discontinued in the 1990s.

Table 4: Impact of TRIPS on Patent Counts

Dependent Variable	Patent Count						Citations	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
After	-41.8 (13.3)**	-32.2 (9.9)**	-37.1 (10.2)**	-41.0 (18.7)*	-58.2 (18.0)**	-31.8 (11.5)**	0.027 (0.105)	0.063 (0.081)
Extension (days)	-0.18 (0.048)**	-0.15 (0.041)**	-0.037 (0.021)	-0.033 (0.024)	-0.048 (0.021)*	-0.069 (0.027)*	0.00032 (.00025)	0.00014 (0.00016)
After*Extension	0.085 (0.029)**	0.049 (0.020)*	0.039 (0.019)*	0.082 (.039)*	0.108 (0.038)**	0.063 (0.025)**	-1.76E-05 (.00013)	-9.36E-05 (0.00012)
Time (months)	-1.21 (0.37)**	0.89 (0.14)	5.77 (1.39)**	-1.11 (0.75)	1.33 (0.28)**	-1.04 (0.76)	0.012 (.021)	0.0056 (0.016)
Generality of Citing Patents	- -	- -	-17.2 (8.71)*	-4.79 (3.61)	-1.54 (1.56)	-24.43 (15.2)	-0.092 (.082)	-0.65 (0.16)**
Mean Year of Citing Patents	- -	- -	4.93 (1.87)**	0.40 (0.71)	-0.43 (0.42)	7.08 (2.71)**	0.34 (.027)**	0.12 (0.026)**
Mean Cites to Citing Patents	- -	- -	10.01 (3.40)**	0.78 (0.91)	0.03 (0.38)	14.52 (4.46)**	0.44 (.032)**	0.45 (0.044)**
Claims	- -	- -	0.22 (0.20)	0.02 (.10)	-0.003 (0.044)	0.01 (0.34)	0.016 (.0021)**	0.02 (0.0032)**
Generality Index	- -	- -	-1.37 (5.74)	-1.87 (1.85)	-1.39 (0.88)	-10.69 (8.17)	0.052 (.025)*	0.58 (0.065)**
Number of Patents	- -	- -	9.01 (3.76)**	-2.11 (1.76)	0.46 (1.10)	10.56 (5.14)*	0.041 (.017)*	0.08 (0.036)*
Observation Level	Class-Month	Class-Month	Class-Month	Class-Month	Class-Month	Class-Month	Patent	Class-Month
Class-Specific	no	no	no	yes	yes	no	no	no
Time Trend								
Citation-Weighted	no	no	no	no	no	yes	no	no
Outside Window	6	12	12	6	12	6	6	6
Inside Window	2	2	2	2	2	2	2	2
Observations	1155	2882	2856	1148	2856	1148	21050	1148
R²	0.17	0.16	0.46	0.98	0.95	0.50	0.19	0.59

Note: All data from PTO spanning different windows around June 8, 1995, as indicated (windows in months). Citations are total received per patent as of December 31, 2002. Extension is the expected increase in patents due to TRIPS, calculated by patent class. All standard errors are clustered by patent class. Dummies for technology subcategory are included in all regressions. “*” indicates significance at $p < 0.05$ and “**” indicates significance at $p < 0.01$.