Patent Term, Innovation, and the Role of Technology Disclosure Externalities*

Fabio Bertolotti†

November 2024

Abstract

I examine the impact of patent term on R&D and innovation in contexts where policy anticipation, common in real-world settings, is present. Using a difference-in-difference design, I exploit quasi-experimental variation in U.S. patent term across technological fields. Despite a general increase in *average* patent term, in most fields innovators faced a considerable probability of a patent term reduction for future innovations. Three key findings emerge: (1) R&D and innovation accelerate more in fields with a higher probability of patent term reduction, i.e., a shorter average patent term extension, before implementation. (2) This heightened activity persists for at least five years post-implementation, driven by indirect effects where the news-related acceleration fosters further innovation through technological externalities linked to cumulative knowledge creation. (3) Conversely, the *direct* effect of a shorter extension in patent term would stimulate relatively *less* innovation, absent the indirect effects of anticipation.

Keywords: Patent term, Innovation, Anticipation, Externality
JEL classification codes: O31, O33, O34, O38, O41

*I am deeply indebted to my supervisors Ricardo Reis and Xavier Jaravel and to my advisors Alessandro Gavazza and Ben Moll. I am extremely grateful to Andrea Lanteri for constant encouragement and insightful conversations. I also thank Andrea Alati, Philippe Aghion, Filippo Biondi, Thomas Brzustowski, Francesco Caselli, Eleonora Ciscato, Luca Citino, Jamie Coen, Maarten De Ridder, Gaia Dossi, Arnaud Dyevre, Martina Fazio, Nicola Fontana, Alessandro Gavazza, Jonathon Hazell, Ethan Ilzetzki, William Kerr, Joshua Krieger, Per Krusell, Andrea Lamorgese, Andrea Lanteri, Salvatore Lattanzio, Salvatore Lo Bello, Massimiliano Marcellino, Benjamin Moll, Kyle Myers, Francesco Nava,Tom Nicholas, Pietro Peretto, Micheal Peters, Akash Raja, Giorgio Saponaro, Mark Schankerman, John Van Reenen, Fabian Waldinger, Alwyn Young, Martina Zanella, and seminar participants at University of Rochester, CREi, Duke University, Max Planck Institute for Innovation, and Bank of Italy for their precious comments. I thank Antoine Dechezlepretre and Dennis Verhoeven for helping me access the data and Hwan Lin for sharing his numerical algorithms. Also, I thank ESRC for financial support.

†Bank of Italy and Laboratory for Innovation and Science at Harvard. Email: fabiobertolotti@hotmail.it

1 Introduction

Innovation is a key driver of long-term economic growth. Therefore, the economic literature has traditionally focused on the long-run effects of policies promoting Research and Development (R&D) and innovation. However, their implementation can also produce rich near-term dynamics, especially in the presence of policy anticipation, a pervasive phenomenon in real-world settings where policy-making involves negotiation. Due to the cumulative nature of knowledge creation, these transitional dynamics significantly impact long-term productivity and welfare, as highlighted by a growing body of literature (e.g., Comin and Gertler, 2006; Benigno and Fornaro, 2018; Anzoategui et al., 2019; Bianchi, Kung and Morales, 2019; Vinci and Licandro, 2020; Fornaro and Wolf, 2021).

This paper analyzes both the short- and long-term effects of innovation policy shocks in a unified setting. Specifically, I present new empirical evidence on the impact of an anticipated change in U.S. patent term, i.e., the duration of the legal monopoly granted by patents, on innovation and R&D. This setting is relevant for two reasons. First, despite being recognized as a key policy tool since Nordhaus (1967), empirical evidence on patent term is limited (Budish, Roin and Williams, 2016), and normative recommendations vary widely.[1] Second, policy anticipation may induce large short-term variation in innovative activity with an ex-ante ambiguous direction.

On one hand, news of a future patent term reduction may lead to a decline in innovation, as firms scale back R&D efforts in anticipation of lower future returns. On the other hand, the same news may accelerate R&D on existing projects, as innovators seek to benefit from the more favorable policy while it remains in effect. This temporary acceleration in activity could foster further innovations in the medium term due to the cumulative nature of knowledge creation (Romer, 1990).

The empirical analysis exploits anticipated quasi-experimental variation in effective patent term across technological fields resulting from the U.S. ratification of the Trade-Related Aspects of Intellectual Property Rights (TRIPs) agreement. TRIPs standardized intellectual property protection across future World Trade Organization (WTO) members, prompting the U.S. to change the expiry date of patents from 17 years after the grant date to 20 years after the application date, aligning with other advanced economies. Since legal monopoly is fully enforceable only after the patent is granted, the effective U.S. patent term changed from

[1] Normative models prescribe a patent term range that varies from zero (Boldrin and Levine, 2013) to infinite protection (Gilbert and Shapiro, 1990), and in most jurisdictions, the official patent term is determined by a rule-of-thumb approach. For example, US patent term was introduced in 1790 and set to 14 years after the grant date in line with English law, where it was based on the expected training period of two sets of apprentices, as reported by Nordhaus (1969).

17 years to 20 years *minus* the pending period, i.e., the time between application and grant, when the U.S. Patent and Trademark Office (USPTO) examines and processes applications.[2]

Identification exploits two sources of variation in a Difference-in-Difference (DiD) framework. First, cross-sectional variation in pending period across technological fields. Indeed, US patent applications classified in different fields are examined by distinct technical units within USPTO, which differ by technical examination complexity and congestion, resulting in heterogeneous examination times. Second, time variation stemming from two events related to the policy change: a news event at the end of 1992, when US innovators learned about the future policy intervention, and the implementation in June 1995, when the new rules entered into force.

Cross-sectional variation in pending period consists in heterogeneous distributions of patent-level pending period across fields, which I summarize by two key moments. First, the field-specific average pending period, which determines the change in average effective patent term, whose cross-field distribution is depicted in Figure 1a. Second, the field-specific fraction of patents with a pending period longer than three years, which represents the probability that each individual patent in a given field obtains a patent term reduction due to TRIPs. Figure 1b represents its cross-field distribution. Finally, Figure 1c illustrates the strong negative correlation between this probability and the change in average patent term.

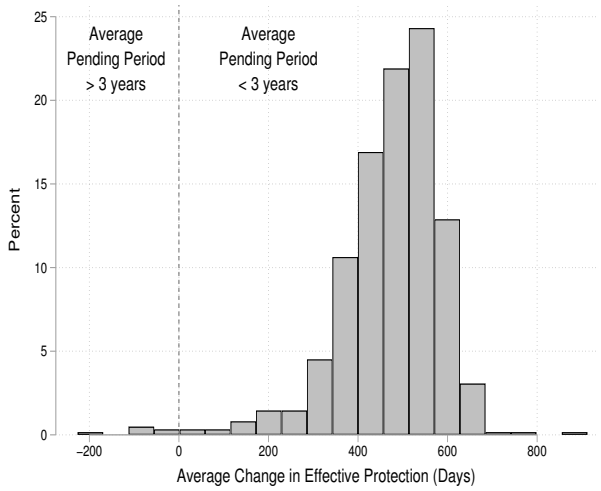
On average, most fields benefit from an extension in *average* patent term but variation across fields is significant. Notably, even in fields with an extension in average patent term longer than one year, individual applications could face up to a 25% probability of getting a patent term reduction. More broadly, a 10 percentage point higher probability of patent term reduction relates to around six-month shorter average patent term extension.

Various complementary analyses suggest that this cross-sectional variation is exogenous to potential confounders affecting innovation before, during, or after the policy.

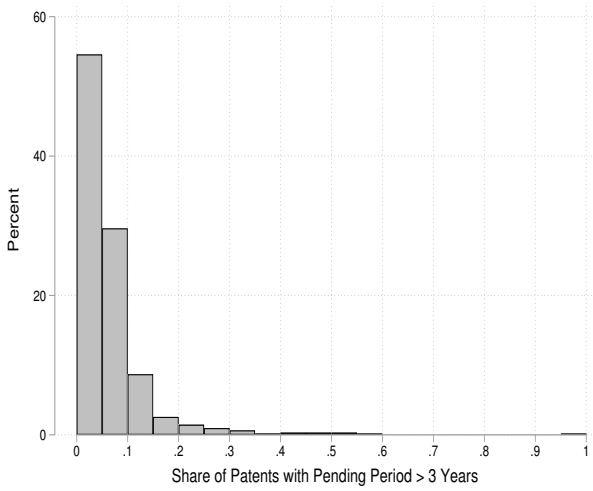
The empirical analysis reveals three key findings. First, following the policy news and before policy implementation, innovation accelerates more in fields with higher probability of patent term reduction and, thus, with a shorter average patent term extension (**Fact 1**). As in most fields the probability that individual patents obtained a patent term loss was significant, these estimates are consistent with the observed surge in aggregate patents applied for during the anticipation phase. I consider the number of granted patents counted by application date as the main innovation measure, but all the results hold for alternative quality-adjusted innovation measures, such as citations or patent value. Additionally, the analysis of firm-level R&D expenditures shows similar findings: firms more exposed to fields

[2]Firms can, in theory, commercialize innovations before grant, but they can sue for infringement only after the patent's publication, which occurred post-grant for patents filed before December 2000.

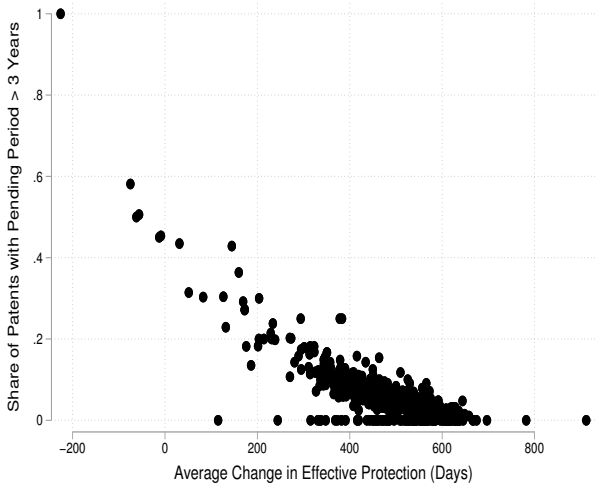
Figure 1: TRIPs-Related Change in Patent Term



(a) Average Patent Term Change Distribution



(b) Term Reduction Probability Distribution



(c) Correlation of Average Change and Reduction Probability

Notes: Panels (a) and (b) illustrate the distribution of the TRIPs-induced change in average effective patent term and the patent term reduction probability across technical fields, respectively. In panel (a), fields with an average pending period shorter than 3 years, denoted by positive values on the x-axis, faced an extension in patent term, on average. Panel (c) illustrates the correlation between the average change in effective patent term (x-axis) and the fraction of patents with a pending period longer than 3 years (y-axis) across technical fields. Each dot represents a 4-digit IPC technical field. Subsection 2.3 reports additional details on measurement.

with a higher probability of patent term reduction increased their R&D spending more during the anticipation phase. This suggests that the observed rise in patenting reflects actual innovation rather than mere adjustments in patenting strategies or mismeasurement.

As a second key finding, I show that the acceleration in R&D and innovation determined by the policy news persists for at least five years after policy implementation (**Fact 2**). This means that R&D and innovation remain relatively higher in fields with shorter average patent term extensions even after June 1995. The paper offers a simple framework to interpret these post-implementation DiD effects. If current innovation depends on both current R&D and past discoveries, the effects in Fact 2 reflect both the direct impact of the implementation of the new patent term and the indirect effects of changes in past knowledge driven by the initial news shock.

The analysis distinguishes between direct and indirect effects. Following the approach of Angrist and Pischke (2009), I adjust the baseline DiD specification to control for field-specific innovation histories, which capture the indirect effects of news-related changes on subsequent innovation outcomes. Consequently, the new DiD estimates for the policy variable reflect the *direct* impact of implementation, accounting for news effects.

As a direct effect, a *shorter* patent term extension leads to a relatively *smaller* increase in innovation (**Fact 3**). However, the indirect effects related to anticipation outweigh the direct effects and shape the "overall" positive post-implementation outcomes (Fact 2).

Estimates of the *direct* effects indicate a semi-elasticity of innovation to a one-month (one-year) increase in patent term of 1.8% (21.6%). Thus, longer patent protection effectively stimulates more innovation as a *direct* effect, consistent with previous literature. Specifically, Budish, Roin and Williams (2015) and Budish, Roin and Williams (2016) find that a one-year extension of patent monopoly increases R&D by 7% to 22% in the pharmaceutical industry. Additionally, Hémous et al. (2023)'s model implies that a one-month increase in U.S. patent term would increase U.S. innovation by 1.2%.

Moreover, DiD estimates for the anticipation phase indicate that a 1 percentage point higher probability of a patent term reduction for future patents corresponds to a 1.3% larger increase in patents before implementation (1.3% increase for a one-month shorter extension in average patent term). This semi-elasticity estimate is new to the innovation literature and suggests that anticipation effects are substantial.

Finally, the paper provides suggestive evidence of technological spillovers driving indirect post-implementation effects. The analysis uses heterogeneity in technological dependence across fields, proxied by patents backward citations intensity. I find that the initial acceleration in innovation following the news leads to a stronger persistence of higher innovation in fields where new inventions rely more on past innovations from the same field, consistence with Hedge, Herkenhoff and Zhu (2022)'s evidence on technology disclosure externalities. Additionally, an increase in time-varying technological dependence measures indicates that this is the primary driver of indirect post-implementation effects, while alternative mecha-

nisms—such as changes in competition or adjustments in firms' patenting strategies—lack comparable support in the data.

The empirical analysis underscores the importance of considering the intertemporal trade-offs in R&D and innovation decisions for both endogenous growth theories and real-world policymaking. Specifically, since incentives to innovate respond to temporary shocks affecting R&D returns (e.g., Barlevy, 2007; Argente, Lee and Moreira, 2018; Bertolotti, Gavazza and Lanteri, 2023), innovation-policy interventions influence both short- and long-run outcomes. With policy anticipation, these effects may push in opposite directions. Moreover, near-term variation in innovative activity is not only relevant in its own right but may drive medium- to long-term effects—consistent with a growing literature on medium-term business cycles (e.g., Comin and Gertler, 2006; Benigno and Fornaro, 2018; Anzoategui et al., 2019)—due to technological and knowledge externalities (Romer, 1990).

The remainder of the paper is organized as follows. Section 1.1 discusses the contributions to the literature. Section 2 presents the institutional setting of the policy. Section 3 discusses the data and preliminary descriptive evidence. Section 4 outlines a simple conceptual framework for the empirical analysis. Section 5 presents the main empirical evidence. Section 6 connects the empirical results to endogenous growth theory and investigates the technology externality mechanism. Section 7 concludes.

1.1 Related Literature

This paper contributes to several strands of the literature. First, a growing body of work studies persistent transitional dynamics of temporary shocks in macroeconomic models with endogenous innovation (e.g., Comin and Gertler, 2006; Benigno and Fornaro, 2018; Anzoategui et al., 2019; Bianchi, Kung and Morales, 2019; Vinci and Licandro, 2020; Fornaro and Wolf, 2021). Other papers provide empirical evidence on the pro-cyclical behavior of R&D and innovation (Barlevy, 2007; Argente, Lee and Moreira, 2018; Bertolotti, Gavazza and Lanteri, 2023) or link it to monetary policy shocks (Ma and Zimmermann, 2023). I contribute to this literature by providing novel evidence on the rich dynamics induced by policy shocks aimed at stimulating innovation in the long run. With anticipation, short-run effects may have opposite sign than long-term ones. Moreover, they are large because, differently from other types of shocks, these policies are specifically targeted to impact R&D and innovation.

Second, I contribute to a vast literature on the effects of patent policy on R&D and innovation (e.g., Sakakibara and Branstetter, 2001; Moser, 2005; Lerner, 2009; Kyle and McGahan, 2012; Moser and Voena, 2012; Galasso and Schankerman, 2015; Schankerman

and Schuett, 2017; Moscona, 2021; Acikalin et al., 2022) by focusing on a specific aspect, the patent term, which, since Nordhaus (1969), is commonly considered key for innovation but has received little empirical attention (Budish, Roin and Williams, 2016). Budish, Roin and Williams (2015) document that in the US pharmaceutical sector R&D is disproportionately directed towards treatments with shorter clinical trials, which implicitly offer longer effective protection. While their estimates capture both the impact of patent term as well as firms' preference for projects with faster return from investment, TRIPs variation allows me to isolate the effect of the policy. Abrams (2009) uses the same quasi-experimental strategy as this paper but assumes that the policy intervention was *unanticipated*, which leads to different econometric specifications and divergent reduced-form results. In contrast, I gather documental evidence from several sources that US firms anticipated the TRIPs and show the short- and medium-term importance of news effects.[3]

Third, several papers examine the effects of TRIPs on various outcomes related to U.S. and international patenting. Kyle and McGahan (2012) and Delgado, Kyle and McGahan (2013) analyze TRIPs' impact on patenting in developing countries, particularly in the pharmaceutical sector, while Bloomfield et al. (2022) investigate the diffusion of scientific knowledge. Lemus and Marshall (2018) explore how applicants are incentivized to respond more quickly to patent office inquiries. Hémous et al. (2023) assess the aggregate welfare effects of TRIPs in a model of international trade with endogenous innovation, technology diffusion, and patent protection. Caicedo and Pearce (2024) study how incentives to accelerate patent applications affect patent quality. My contribution focuses on a key aspect of patent policy—the patent term—and highlights the implications of policy anticipation.

Finally, the paper connects to the large empirical and theoretical literature on innovation-related spillovers, including: knowledge accumulation spillovers, at the core of Romer (1990), and recently re-examined by Aghion and Jaravel (2015) and Bloom et al. (2020); spillovers from basic to applied research (Akcigit, Hanley and Serrano-Velarde, 2020); geographic spillovers (Lychagin et al., 2016; Moretti, 2020; Lanahan and Myers, 2022); externalities at the inventor level (Bell et al., 2019; Akcigit et al., 2020); and spillovers in the technological space (Bloom, Schankerman and Van Reenen, 2013; Moretti, Steinwender and Van Reenen, 2019). I provide evidence of a technology disclosure externality acting through the diffusion of *novel* knowledge, which can be seen as a "standing on the shoulders of *young* giants"

[3]I discuss at length in Appendix B the consequences of ignoring anticipation and compare my work to Abrams (2009). He estimates a two-period DiD specification comparing patenting in a narrow window of data (6, 12, or 24 months) before and after the implementation shock of June 1995. Disregarding potential news effects may lead to an incorrect interpretation of implementation effects because the pre-implementation baseline considered for the DiD comparison is itself affected by the policy change due to news, thus violating the parallel trends assumption.

effect. This closely relates to Hedge, Herkenhoff and Zhu (2022), who document that a more timely publication of patent applications increases the rate of follow-up innovation. Moreover, I estimate the elasticity of future to current innovation and collect suggestive evidence on the half-life of this externality.

2 Nature and Timing of the TRIPs Policy Change

2.1 Content of the Policy Change

The empirical analysis exploits quasi-experimental variation in the US effective patent term following the adoption of the *Agreement on Trade-Related Aspects of Intellectual Property Rights* (TRIPs) in the US. TRIPs standardized intellectual property protection across trading partners as part of the Uruguay Round agreements that established the World Trade Organization (WTO). The main impact on the US patent system was a shift in patent expiry from *17 years after the grant date* to *20 years after the application date*.[4]

During the "pending period"—the time between application and grant, when the patent office examines applications—monopoly power is fully not enforceable.[5] As a result, the policy effectively changed the patent term from 17 years to 20 years *minus* the pending period. The identification strategy uses the interaction between policy timing and differences in pending periods across technological fields. Before discussing this heterogeneity in Section 2.3, I will argue that US innovators anticipated the TRIPs patent term change, leading to two policy phases: news (anticipation) and post-implementation.

2.2 Timing: News and Post-Implementation Phases

The Uruguay Round Agreements Act (URAA), enacted on December 8, 1994, officially ratified the TRIPs provisions in the US, with full implementation on June 8, 1995. Evidence from official documents and articles suggests that US firms were aware of the impending policy change well before its formal adoption.[6]

[4]The Uruguay Round Agreements Act, which ratified TRIPs in the U.S., introduced four major changes to U.S. patent law. First, the change in patent term, analyzed in this paper. Second, a non-discrimination rule for foreign inventors. Third, provisional applications were introduced, requiring conversion into formal filings within one year to avoid abandonment. Fourth, TRIPs expanded the scope of patentable subject matter in developing countries. Section 5.5 addresses possible confounding effects from these changes and supports the validity of the results.

[5]Firms can sell or use their innovations before grant, but they can sue for infringement only after the patent's publication. Before December 2000 publication coincided with grant.

[6]The final URAA included a retroactivity clause for patents filed before the June 1995 implementation, granting them longer protection under the new regime. However, this detail was unknown prior to the December 1994 signing and thus did not influence earlier incentives.

First, US businesses were directly involved in the TRIPs negotiations from the start of the Uruguay Round in 1986. According to Morgese (2009) and Matthews (2002), the US Advisory Committee on Trade Policy and Negotiations, which included executives from, e.g., IBM and Pfizer, significantly shaped the US delegation's stance. Second, the adjustment of the US patent term was mentioned in a draft circulated by the GATT Director-General in late 1991.[7] Third, as Montalvo (1996) notes, the Advisory Committee on Patent Law Reform took the first step toward this change in August 1992 by recommending a twenty-year patent term from the filing date.[8] This report, co-signed by representatives from IBM, 3M, P&G, Motorola, and others, explicitly referred to the 1991 TRIPs draft. Fourth, early legal articles, such as those by Reichman (1993), Martin and Amster (1994), and Doane (1994), discussed the TRIPs draft. Finally, a *New York Times* article from September 1992 also noted the proposed changes.[9] Therefore, US innovators were aware of the negotiations and could anticipate the policy change.

Historical records show that the signing of the Blair House Accord in November 1992 significantly reduced uncertainty about the agreements' adoption.[10] Therefore, the paper considers two distinct policy phases marked by a "news" event in November 1992 and an implementation event in June 1995.

Although some uncertainty remained after the Blair House Accord (Abrams, 2009), US innovators were aware of the potential policy change and could respond to the information. The flexible DiD specification in Section 4 captures potential anticipation effects without imposing ex-ante assumptions. Section 5.1 discusses the timing of news effects, their robustness to earlier dates, and their interpretation in light of remaining uncertainty. Appendix B further explores the consequences of neglecting news effects for the parallel trends assumption in DiD analysis and relates this paper to Abrams (2009), who assumes no policy anticipation.

[7] *GATT doc. MTN.TNC/W/FA, Draft Final Act Embodying the Results of the Uruguay Round of Multilateral Trade Negotiations, 20/12/91*

[8] *The Implementation of the Uruguay Round Agreement on Trade-Related Aspects of Intellectual Property–the TRIPs Agreement: Hearings on S.2368 and H.R. 4894 before the Subcomm. on Patents, Copyrights and Trademarks of the Senate Judiciary Comm. and the Subcomm. on Intellectual Property and Judicial Administration of the House Judiciary Comm., 103rd Cong., 2d Sess.*

[9] *Panel Proposes Patent Changes*, New York Times, Late Edition (East Coast); New York, 15 Sep 1992.

[10] This is reported by Morgese (2009) and at `https://en.wikipedia.org/wiki/Uruguay_Round`, where it reads: *"The round was supposed to end in December 1990, but the US and EU disagreed on how to reform agricultural trade and decided to extend the talks. Finally, In November 1992, the US and EU settled most of their differences in a deal known informally as the "Blair House accord" [...]"*

2.3 Variation in Patent Term across Technical Fields

As TRIPs changed the effective patent term from $T^{pre} = 17$ years to $T^{post} = 20$ years minus the pending period, the length of the pending period became crucial in determining the magnitude and direction of the policy's impact. A pending period shorter than three years resulted in a term extension, while a longer period led to a reduction.

To identify the effects of the TRIPs policy change, I exploit the interaction between the policy shock and pre-existing cross-sectional differences in pending periods across technical fields. This variation arises because US patent applications in different fields are reviewed by distinct USPTO technical units, which vary in complexity and backlog due to staffing or foreign filing intensity.[11]

Since historical data on pending periods by technical unit are unavailable, I use technical fields as proxies, defining fields as one of 621 4-digit patent classes under the International Patent Classification (IPC) scheme. I also experiment with narrower and broader field definitions in Section 5.5.

Because examination is patent-specific, cross-field heterogeneity consists in variation in the field-specific distributions of patent-level pending periods. I summarize this variation with two statistics from these distributions: the average pending period—which determines the change in average patent term—and the share of patents with a pending period exceeding three years—indicating the probability that a patent obtains a reduction in patent term. While the analysis primarily focuses on the average change, Section 3.2 discusses how the probability of patent term reduction helps interpret the main DiD estimates.[12]

Therefore, the change in average effective patent term and the patent term "loss" (reduction) probability in field j are

$$\Delta T_j = 20 \text{ years} \times 365 \; - \; \overline{PP}_j \; - \; 17 \text{ years} \times 365 \tag{1}$$

$$PL_j = N_j^{-1} \textstyle\sum_i \mathbb{I}(PP_{i,j}/365 > 3 \text{ years}) \tag{2}$$

where $\overline{PP}_j$ is the average pending period, in number of days, for patents classified in field j and granted before the TRIPs news, and PL_j is the fraction of such patents with a pending period longer than three years. N_j denotes the number of relevant patents and $\mathbb{I}(\cdot)$ is an indicator function equal to one if the pending period $PP_{i,j}/365$ of patent i exceeds 3 years.[13]

[11] Applicants propose relevant patent classes, but the USPTO may adjust them and assign the application to the appropriate unit. Since pending periods also depend on applicants' responsiveness (Lemus and Marshall, 2018), Section 5.5 shows no correlation between applicant responsiveness and pre-existing heterogeneity in average pending periods across fields.

[12] Higher moments of the pending period distribution at the field level have a negligible effect on patent outcomes once the average pending period is accounted for.

[13] I use PATSTAT (EPO, 2017) to compute $\overline{PP}_j$ and PL_j based on all *granted* US patent applications

The interactions of ΔT_j or PL_j with quarterly fixed effects constitute the treatment variables in the field-level DiD empirical analysis of Section 5.

Figure 1a illustrates the distribution of the average change, ΔT_j, across technical fields. Most fields were expected to obtain an anticipated increase in *average* effective patent term, whereas few fields faced a projected reduction. The mean of the distribution is +473 days, or about 15 months, with a standard deviation of 177 days.

However, many fields also featured a significant probability that individual patents obtained a pending period longer than three years, which would result in a patent term reduction (Figure 1b). Approximately 45% (15%) of fields faced a reduction probability larger than 5% (10%). Even fields with a positive average change in patent term had loss probabilities as high as 40%. Figure 1c highlights the strong negative correlation between the expected average patent term change (x-axis) and the loss probability PL_j (y-axis), as a higher fraction of patents with pending periods over three years leads to a smaller ΔT_j.

Section 5.5 explains that this cross-sectional variation is not correlated with observable or unobservable factors that could affect innovation differently across fields before and after the policy shock.

3 Data and Descriptive Evidence

3.1 Data

The empirical analysis is conducted at three levels: (i) technological field, (ii) firm, and (iii) NAICS 6-digit industry. Patent data is primarily sourced from PATSTAT (EPO, 2017), complemented with patent generality and originality measures from Hall, Jaffe and Trajtenberg (2001), patent quality and scope metrics from Marco, Sarnoff and deGrazia (2016), text-based novelty measures from Arts, Hou and Gomez (2021), and private economic value estimates from Kogan et al. (2017).[14] Data on disambiguated inventors and applicants are obtained from PatentsView (USPTO, 2023).

The quarterly panel spans the universe of 621 4-digit International Patent Classes (IPC), defining technical fields in this paper, from 1985Q1 to 2000Q4, covering the period around the TRIPs shocks. The sample ends in 2000 due to the introduction of the American Inventors Protection Act (AIPA), which further modified US patent law.[15] The average number of

that (i) belong to technical field j; (ii) whose earliest application is filed at the USPTO; and (iii) whose grant date is between January 1^{st}, 1990 and May 31^{st}, 1992.

[14] Higham, de Rassenfosse and Jaffe (2021) surveys alternative patent quality measures, and Kelly et al. (2021) compares economic value and text-based quality measures over the long run.

[15] AIPA mandated the publication of patent applications 18 months after filing, effective for patents filed after 11/29/2000. Analysis by Hedge, Herkenhoff and Zhu (2022) suggests the policy had no impact on

quarterly patents and 5-year forward citations-weighted patents are 36 and 195, respectively, with standard deviations of 136 and 1,070. Additional summary statistics are provided in Table A1 in Appendix A.

The firm-level dataset is a yearly panel of 2,421 listed US firms from the NBER-Compustat matched dataset by Hall, Jaffe and Trajtenberg (2001) covering the period 1985-2000, with balance sheet data sourced from Compustat (Standard&Poor's, 2022). On average, firms file 14 (granted) patents per year, and the average annual R&D expenditure is \$61 million, with standard deviations of 95 and \$359 million, respectively. Additional firm-level summary statistics are presented in Table A2.

Sectoral analyses use data on Total Factor Productivity (TFP), producer prices, and other aggregates from the NBER CES manufacturing database (Becker, Gray and Marvakov, 2021) for 428 6-digit NAICS industries from 1985-2000. Patent data are aggregated from the technical-field level using the "Algorithmic Links with Probabilities" crosswalks by Goldschlag, Lybbert and Zolas (2019), which map technological fields to industrial sectors. Table A3 provides summary statistics for the sectoral variables.

3.2 Descriptive Evidence

In this subsection, I present preliminary descriptive evidence on patenting dynamics around the TRIPs patent term change to guide the interpretation of DiD estimates of Section 5.

Figure 2a shows the evolution of granted patents *by application quarter* in three fields with similar pre-news trends but differing in average patent term change and loss probability. The solid line represents field C12P (in chemicals), which has a short average patent term extension ($\Delta T_j = 31$ days) and a high fraction of patents with a pending period over three years ($PL_j = 0.44$). The long-dashed line shows G01L (in physics), with a larger extension ($\Delta T_j = 130$ days) but a still sizable loss probability ($PL_j = 0.23$). The short-dashed line shows E05D (in construction), where the average extension is nearly two years ($\Delta T_j = 600$ days) and the loss probability is negligible ($PL_j = 0.01$).

In fields with a significant loss probability (C12P and G01L), later granted patent applications begin accelerating after the 1992Q4 news date (first vertical line) and peak in the two quarters before implementation in 1995Q2 (second vertical line). This acceleration is absent in E05D. These dynamics suggest that innovators expecting a reduction in patent protection—and therefore a shorter average term extension—accelerate patenting before the new policy takes effect.

The increase starts small but intensifies as implementation nears, with a peak in 1995Q1

patenting behavior prior to its implementation.

and 1995Q2 due to the URAA implementation details. Between its signing in December 1994 and full implementation in June 1995, applicants could choose the regime with the longest protection. This created a strong incentive to file during this transition period, especially in fields facing potential patent term reductions, explaining the observed bunching in 1995Q2.

After implementation, patenting remains higher than pre-news trends in the fields with the largest acceleration at news (C12P and G01L), suggesting persistent effects in fields with a shorter average patent term extension. In contrast, patenting remains stable and slightly above trend in E05D, the field with the largest average term extension.

These insights generalize beyond these three fields. Figure 2b shows the average deviation of patenting from its 1985Q1-1992Q3 trend during the news phase (excluding 1995Q1-Q2) by quintile of average patent term change ΔT_j (black dots, x-axis). The average loss probability (gray squares, right axis) is also reported across quintiles of ΔT_j.

Across all values of ΔT_j patenting increases at news, with the largest acceleration in fields with *shorter* patent term extensions and higher patent term reduction probabilities. Similar patterns are observed for the post-implementation period and citations-weighted patents.

In summary, the raw data suggest an overall increase in innovation, with stronger responses in fields facing a higher probability of patent term reduction and, thus, a shorter average extension. This descriptive evidence will guide the interpretation of the DiD estimates in Section 5

4 Conceptual Framework for the Empirical Analysis

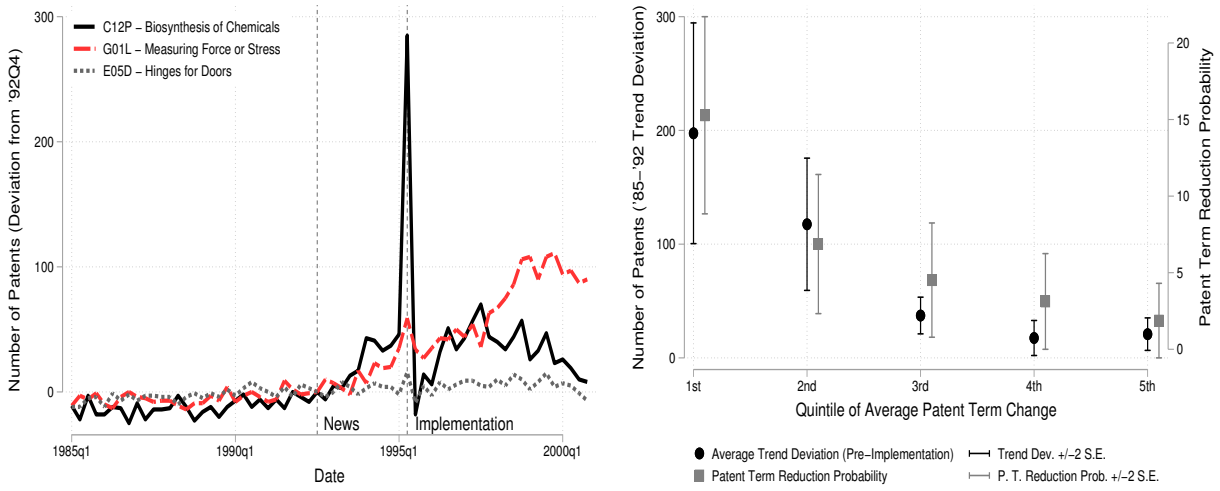
This section introduces a stylized framework to analyze news and policy implementation effects under anticipation. I first derive the causal marginal responses of interest, then link them to the paper's reduced-form DiD specifications. Identification concerns are discussed in Section 5.5.

4.1 Theoretical Effects of Interest with Anticipation

I assume that aggregate innovation at time t, denoted by I_t, is a function of (cumulative) past innovation, $CumI_{t-1}$, and current R&D effort $R\&D_t$, which is endogenously determined as a function of past innovation and a policy vector $\boldsymbol{\tau}_t$ collecting current and future expected policies. For the case of patent term T_t, $\boldsymbol{\tau}_t = [T_t, T^e_{t+1}, ...]$. Therefore, aggregate innovation itself is a function of past cumulative innovation and policy, i.e.,

$$I_t \equiv F(CumI_{t-1}, R\&D(CumI_{t-1}, \boldsymbol{\tau}_t)) = G(CumI_{t-1}, \boldsymbol{\tau}_t)$$

Figure 2: Descriptive Evidence on Patenting Across Fields



(a) Patenting across Fields (b) Pre-Implementation Acceleration

Notes: Panel (a) illustrates the evolution of the number of granted patent applications by application quarter in three technical fields: C12P with $\Delta T_j = 31$ days and $PL_j = 0.44$ (black solid line); G01L with $\Delta T_j = 130$ days and $PL_j = 0.23$ (red long-dashed line); E05D with $\Delta T_j = 600$ days and $PL_j = 0.01$ (gray short-dashed line). The first and second vertical lines mark the news quarter 1992Q4 and the implementation quarter 1995Q2, respectively. Panel (b) illustrates the relation between the average deviation of the number of granted patents applied for during the news phase 1992Q4-1994Q4 from its 1985Q1-1992Q3 trend (black dots, left y-axis) and the average change in effective patent term by quintile of the latter variable (x-axis). Moreover, gray dots (right y-axis) plot the average patent term loss probability—equal to 100 times the fraction of pre-news patents with a pending period longer than 3 years—by quintile of average change in effective patent term.

I consider three time periods: (i) $t-1 =$*pre-news*, when no news about future policy changes occurred; (ii) $t =$*news*, when the old policy is in place but news diffuses that a new policy will apply from $t+1$ onward; (iii) and $t+1 =$*implementation*, when the new policy enters into force.

The empirical analysis aims at estimating three effects of interest. First, the "news effect", i.e., the incremental change in innovation (or R&D) at $t =$*news* caused by a marginal change in the policy at $t+1 =$*implementation*. In the simple framework, this is:

$$\frac{\mathrm{d}I_{news}}{\mathrm{d}T_{implem}} = \frac{\partial G(\cdot,\cdot)}{\partial CumI_{pre-news}} \underbrace{\frac{\partial CumI_{pre-news}}{\partial T_{implem}}}_{=0} + \frac{\partial G(\cdot,\cdot)}{\partial T_{implem}} = \frac{\partial G(\cdot,\cdot)}{\partial T_{implem}} \tag{3}$$

The first term denotes the impact of the future policy change on news-period innovation through changes in pre-news innovation. This is null by the assumption that innovation is unaffected before policy news occurs (null pre-trends). The second term represents the

direct impact of the future policy change on current innovation upon news, due to changes in R&D.

The second effect of interest is the "overall" implementation effect, i.e., the incremental change in innovation at $t+1$ =*implementation* caused by the implementation of an *anticipated* marginal change in T_{implem}. As the following expression highlights, the "overall" effect results from the combination of an *indirect* effect, through changes in past innovation due to news effect (3), and a *direct* effect, due to changes in R&D caused by the actual implementation of the policy:

$$\frac{\mathrm{d}I_{implem}}{\mathrm{d}T_{implem}} = \underbrace{\frac{\partial G(\cdot,\cdot)}{\partial CumI_{news}} \frac{\partial CumI_{news}}{\partial T_{implem}}}_{Indirect} + \underbrace{\frac{\partial G(\cdot,\cdot)}{\partial T_{implem}}}_{Direct} \tag{4}$$

The third effect of interest is the "direct" impact of policy implementation, i.e., the second term of expression (4). It represents the incremental change in innovation caused by a concomitant marginal change in patent term absent anticipation effects.

4.2 Estimation of News and Overall Implementation Effects

To estimate news and "overall" implementation effects (3) and (4) on aggregate innovation, I leverage cross-sectional variation in average effective patent term change ΔT_j or patent term loss probability PL_j across technological fields in the DiD specification:

$$Y_{j,t} = \alpha_j + \sum_{\substack{k=1985Q1 \\ k \neq 1992Q3}}^{2000Q4} \gamma_k \mathbf{1}_{(t=k)} + \sum_{\substack{k=1985Q1 \\ k \neq 1992Q3}}^{2000Q4} \beta_k \mathbf{1}_{(t=k)} X_j + \varepsilon_{j,t} \tag{5}$$

where $Y_{j,t}$ is technical field-j and application quarter-t dependent variable (in levels), α_j are technical field fixed effects, $\mathbf{1}_{(t=k)}$ are quarter-specific dummy variables, with γ_k coefficients capturing the effect of any time-varying unobserved factor whose impact is common across fields, X_j is either ΔT_j or PL_j, and $\varepsilon_{j,t}$ is an idiosyncratic error term.[16]

The β_k coefficients denote the quarterly DiD effects of interest. They represent the marginal effect of an anticipated unit-change in X_j on the outcome variable, in *level* deviation from its baseline in 1992Q3 (excluded news quarter).

[16] In Appendix B, I show that quarterly specification (5) can flexibly capture potential news effects. In contrast, a standard two-period DiD analysis comparing outcomes before and after policy implementation would lead to biased treatment effect estimates.

I also analyze the following Poisson specification for positive count variables

$$Y_{j,t} = \exp\left\{\alpha_j + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \gamma_k \mathbf{1}_{(t=k)} + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \beta_k \mathbf{1}_{(t=k)} X_j + \varepsilon_{j,t}\right\}, \tag{6}$$

where DiD coefficients represent the marginal effect of X_j on the outcome in *percent* deviation from the pre-news baseline.[17]

The economic interpretation of the β_k coefficients varies across the pre-news, news, and implementation periods, as discussed in Section 4.1. For $k \leq 1992Q3$, β_k reflects the marginal effect of ΔT_j or PL_j before policy news, which should be zero, indicating the absence of pre-news differential trends in innovation (i.e., no pre-trends). During the news period ($1992Q4$–$1995Q2$), β_k represents the effect of news on future (granted) patent applications, capturing the marginal variation in X_j for quarter k. This aligns with the news effect in equation (3). In the post-implementation period ($k \geq 1995Q3$), β_k captures the quarter-specific marginal effect of the *anticipated* patent term change, accounting for both direct and indirect policy effects, as described in equation (4).

I consider several outcome variables to proxy aggregate innovation and quality at the technical field level. The "quantity" of innovation is measured by (i) the raw count of granted U.S. patents *by application quarter*, (ii) the number of U.S. patents weighted by 5-year forward citations (a common metric for scientific quality-adjusted innovation), or (iii) the private economic value of patents, as estimated by Kogan et al. (2017).

To proxy average quality, I use text-based measures of patent novelty (Arts, Hou and Gomez, 2021), measures of patent scope based on the type and length of claims (Marco, Sarnoff and deGrazia, 2016), average citations or economic value, as well as generality and originality measures (Hall, Jaffe and Trajtenberg, 2001).

For the analysis of the effect of patent term on R&D expenditures, I focus on firm-level data in Section 5.2. While patents capture innovation output, patent data alone do not provide a reliable measure of R&D input by field. In contrast, firm-level balance sheet data offer a more accurate measure of R&D expenditures.[18]

[17] Mullahy and Norton (2022) examine different models for non-negative and skewed outcome variables with large probability mass at zero. Differently from linear models with log- or inverse hyperbolic sine-transformation of the dependent variable, linear regressions or Poisson models with dependent variable in levels estimate correct marginal effects.

[18] I confirm that all patent-related results extend to a proxy for R&D effort by technical field, based on the number of unique inventors by technical field and quarter. These results are available upon request.

4.3 Recovering the Direct Policy Effect

The empirical analysis proceeds to distinguish between the direct and indirect effects within the overall implementation effect in equation (4). The direct effect measures how changes in patent term affect innovation, independent of anticipation and of the specific conditions of TRIPs implementation. In contrast, the indirect effects arise from how changes in past innovation, driven by anticipation, influence post-implementation innovation, reflecting the externalities tied to the cumulative nature of knowledge creation (Romer, 1990).

To isolate the direct policy effect, I modify DiD specification (5) to control for field-specific innovation histories, following Angrist and Pischke (2009)'s discussion of DiD with path-dependence. Therefore, direct effect estimates leverage variation in innovation that is orthogonal to past outcomes across technical fields with heterogeneous patent term changes. However, controlling for past outcomes requires omitting field fixed effects to avoid inconsistent DiD estimates (Nickell, 1981). The modified DiD specification is:

$$\begin{aligned} Y_{j,t} = &\sum_{k\neq'92Q3} \mathbf{Z}_{pre,j}\mathbf{1}_{(t=k)}\eta_k + \sum_{k\neq'92Q3} \gamma_k \mathbf{1}_{(t=k)} + \sum_{k\neq'92Q3} \phi_k \mathbf{1}_{(t=k)} X_j + \\ &+ \sum_{k\neq'92Q3} \psi_k \mathbf{1}_{(t=k)} \underbrace{\overline{Y}_{j,k-\mathcal{A}-1:k-1}}_{\equiv \frac{1}{\mathcal{A}}\sum_{q=k-\mathcal{A}-1}^{k-1} Y_{j,q}} + \psi_0 \underbrace{\overline{Y}_{j,t-\mathcal{A}-1:t-1}}_{\equiv \frac{1}{\mathcal{A}}\sum_{q=t-\mathcal{A}-1}^{t-1} Y_{j,q}} + v_{j,t}. \end{aligned} \tag{7}$$

To control for field-specific time-invariant characteristics, I introduce a vector $\mathbf{Z}_{pre,j}$, which includes pre-determined attributes (field size, average forward citations per patent, and average inventors per patent from 1980-1985) interacted with quarterly fixed effects. These substitute for the field fixed effects in specification (5).[19] The second line of (7) parsimoniously accounts for the quarter-specific impact of a field's innovation history. Specifically, $\overline{Y}_{j,k-\mathcal{A}-1:k-1}$ represents the average outcome over the $\mathcal{A}$ quarters before k, and ψ_k captures deviations from its baseline impact ψ_0. I set $\mathcal{A}$ to 10, corresponding to the TRIPs anticipation period, but confirm robustness with $\mathcal{A} = 8$ and $\mathcal{A} = 16$.

The new DiD coefficients ϕ_k estimate the effect of a marginal change in the average patent term change (or in the probability of a patent term reduction) on quarter-k innovation outcome, controlling for the effects of field-specific innovation histories. After policy implementation, these coefficients reflect the *direct* effect of equation (4). Before implementation, they coincide with the news effects, as per equation (3). Consistent with this, the empirical findings in Section 5.3 show that the direct effect closely aligns with the baseline

[19] Separate analyses confirm that replacing fixed effects with $\mathbf{Z}_{pre,j}$ controls does not alter the findings compared to the original TWFE DiD specification (5). Moreover, by estimating regression (7) including field fixed effects, I verify that Nickell (1981)'s bias is small, consistent with the fact that it scales with the inverse of the number of time periods in the sample, which is 64 in my sample.

DiD estimates during the news period. In contrast, the direct effect has opposite sign than the "overall" effect after implementation.

5 Effects of an anticipated change in patent term

This section presents the empirical findings of the paper. I begin with estimates of the news and overall implementation effects on patenting and patent quality by technical field. I then examine the response of R&D, a key input for innovation, using firm-level balance sheet data. Next, I present the *direct* implementation effects and compare their magnitude to previous literature. Finally, I address measurement and identification concerns in detail. Appendices C, D, and E provide additional results by field, firm, and on sectoral productivity and prices.

5.1 News and overall implementation effects by field

Figure 3 presents the results of DiD specification (5) for the number of granted patents by field and *application quarter* as outcome variable. Panel 3a shows the effect of a one-month *shorter* extension in average patent term ($X_j = \Delta T_j$), while panel 3b shows the effect of a one-percent *higher* patent term reduction probability ($X_j = PL_j$). Panels 5a and 5b display the same effects using the Poisson DiD specification (6). Dots represent point estimates, and the bands show 95% confidence intervals, with standard errors clustered by technical field and treatment phase. The figures highlight three key findings.

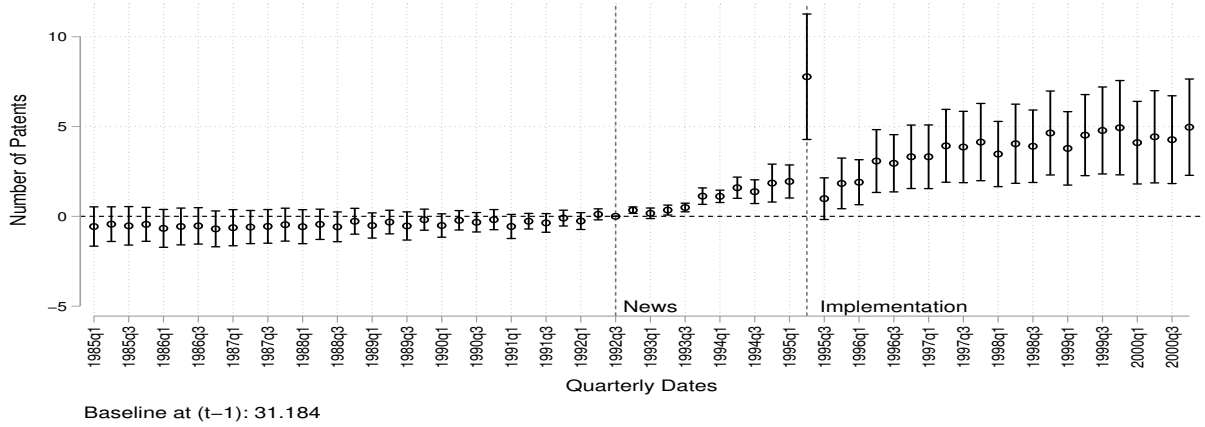
Pre-trends. First, the estimated effects of marginal changes in the regressors are close to zero before the news shock. Formal tests based on Roth (2022) strongly reject the presence of economically significant pre-trends.[20] Thus, the change in patent term is not correlated with unobserved, field-specific innovation patterns that existed prior to the policy news. Section 5.5 will discuss the absence of potential confounders occurring around or after the policy change

News effect. Second, during the news phase ($k \in [1992Q4; 1995Q2]$), the DiD estimates show a positive and statistically significant impact of both a one-month shorter extension in average patent term (Figure 3a) and a one-percent higher probability of patent term reduction (Figure 3b) on patenting. This aligns with the descriptive evidence in Section 3.2. A higher probability of patent term reduction—or a shorter extension in average patent

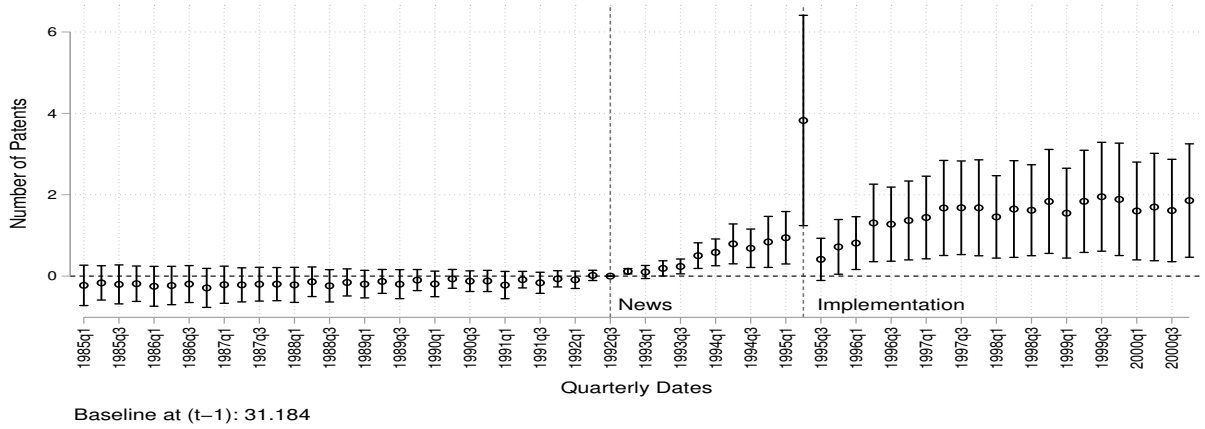
[20] I examine several alternatives for $X_j = \Delta T_j$. A linear trend with a quarterly slope of -0.005 (equal to the 2000Q4 point estimate divided by the number of post-news quarters) would be detected with a power of 1, while a trend with a slope of -0.001 (equal to the upper bound of the 2000Q4 confidence interval) would be detected with a power of 0.92.

Figure 3: DiD estimates by technical field – Patents by Application Quarter

(a) Effects of Average Patent Term Change



(b) Effects of Patent Term Reduction Probability



Notes: The figure represents the effects of the patent term change on field-level (granted) patent applications by application quarter from DiD specification (5). In panel (a), the regressor of interest is $X_j = \Delta T_j$ and the estimates represent the relative effect of a one-month *shorter* extension in the average patent term on the level of the outcome variable. This is computed as the DiD coefficients estimates $\widehat{\beta}_k$ times -30 (days). In panel (b), the regressor of interest is $X_j = PL_j$ and the estimates represent the relative effect of a one-percent *higher* probability of a patent term reduction on the level of the outcome variable. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

term—for post-implementation filings leads to an increase in (eventually granted) patent applications during the anticipation phase.

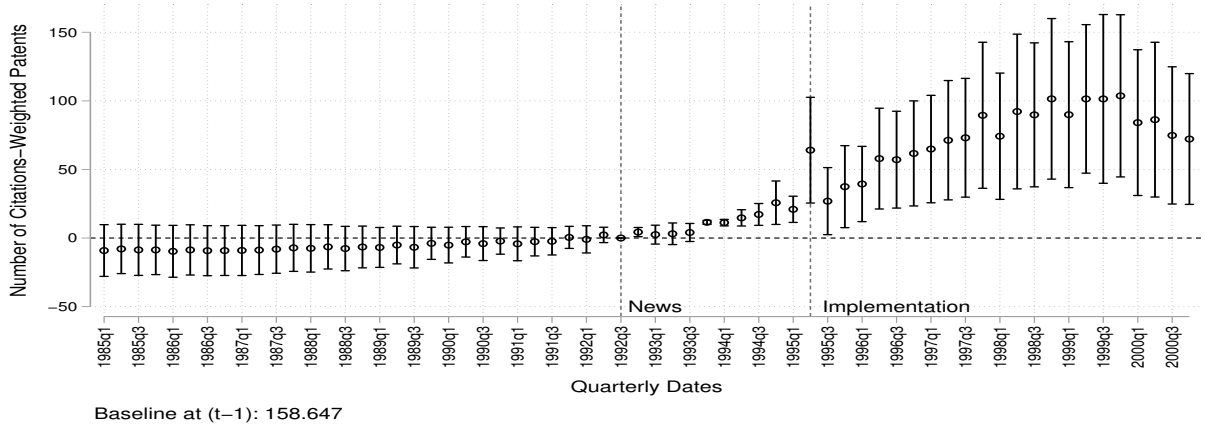
The effect starts small but grows as the policy implementation date approaches. For example, one year after the news and two years before implementation, a one-percent higher probability of patent term reduction leads to 0.2 additional (granted) applications per quarter (+0.5 percent in Poisson DiD estimates, Figure 5b). Similarly, in fields with a one-month

shorter extension, patenting increases by 0.5 additional units (+1.1 percent in Poisson estimates, Figure 5a). These effects nearly triple two years after the news and one year before implementation.

This timing reflects both the time required for R&D investments to generate innovation outputs as well as the gradual response of innovators to reduced uncertainty about the policy as its implementation neared. As to the first factor, R&D expenditures—on which I present empirical evidence in Section 5.2—have a gradual effect on project completion. The average R&D gestation lag, estimated by Pakes and Schankerman (1984) at about two years, aligns with the length of the anticipation phase. Moreover, the early increase in innovation may result from accelerating R&D on projects that were close to completion at news. As to the second factor, while the policy was anticipated, Abrams (2009) notes some uncertainty remained until 1994. Thus the effect may build as this uncertainty gradually resolves after the Blair House Accord in November 1992.

The DiD estimates show a sharp rise in 1995Q2, the last quarter before implementation. As discussed in Section 3.2, this is due to application bunching caused by URAA provisions, which allowed applicants to choose the most favorable policy regime for applications filed between December 1994 and June 1995. Consequently, 1995Q2 was the final quarter to avoid any potential patent term loss. The incentive to bunch was greater in fields with a higher probability of term reduction. However, the patenting increases observed in earlier quarters reflect genuine innovation changes, as supported by the consistent responses in R&D expenditures, quality-adjusted innovation measures, and productivity (Appendix E).

Figure 4: DiD estimates by technical field – 5-Year Forward Citations-Adjusted Patents



Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the level of 5-year citations-weighted (granted) patents by application quarter, based on estimates of DiD specification (5) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

I preview the effects of patent term on patent quality (Section 5.5) by estimating specification (5) using citations-weighted patents—i.e., granted patents applied for in quarter t and classified in field j weighted by their 5-year forward citations—as the dependent variable. This is a common proxy for scientific quality. Figure 4 shows the effect of a one-month shorter extension in average patent term on quality-adjusted patents, while Figure 5c provides similar estimates using the Poisson specification (6). Like in Figure 3a, there are no significant pre-trends. Positive estimates after the news confirm that fields with *shorter* extension in average patent term see a *larger* increase in quality-adjusted patenting during the anticipation phase. The effects are quantitatively similar to those for raw patent count, suggesting the patent term change did not significantly impact average patent quality.[21] Section 5.5 further examines alternative patent quality measures and links the lack of impact on average quality to selection into patenting and secrecy.

I summarize previous evidence on news effects as

Fact 1 At news, innovation and R&D increase relatively more in fields with a higher probability of a patent term reduction and, thus, with an shorter average patent term extension.

Section 5.2 presents the underlying evidence on the response of firm-level R&D expenditures to the patent term change.

Overall implementation effect. Figures 3 and 5 show the estimated post-implementation effects ($k \geq 1995Q3$) of the anticipated policy, combining both direct policy impacts and indirect effects from anticipation. A higher probability of a patent term reduction and a shorter average patent term extension are associated with higher innovation after implementation. Two years after the policy, a one-percent increase in the probability of a patent term reduction results in 1.7 additional quarterly patents (+3 percent in Poisson estimates), and a one-month shorter patent term extension corresponds to 3.9 more patents (+6.5 percent in Poisson results). Similar results are observed for citations-weighted patents in Figures 4 and 5c.

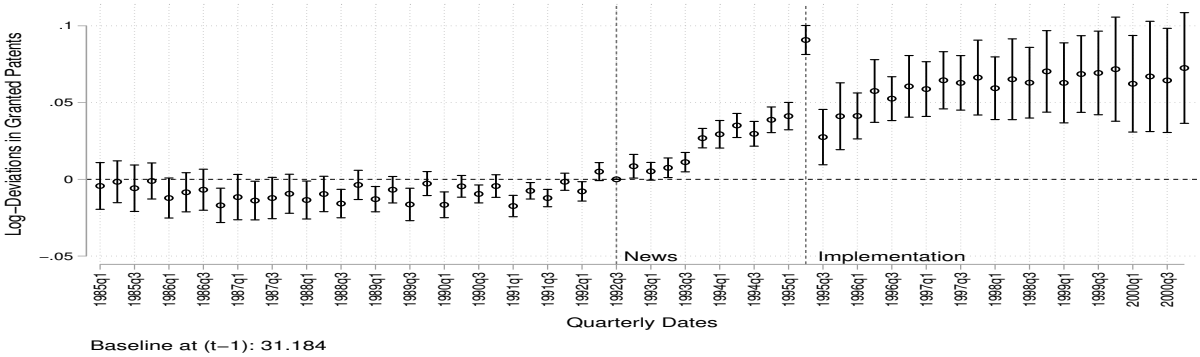
To determine if these post-implementation effects are temporary and persistent or permanent, I extend the sample to 2010. The main analysis ends in 2000 due to other changes in U.S. patent policy (Hedge, Herkenhoff and Zhu, 2022). Appendix C.1 shows that the effects are temporary, as the DiD estimates return to zero after 2000.

Therefore, I summarize evidence on overall post-implementation effect as

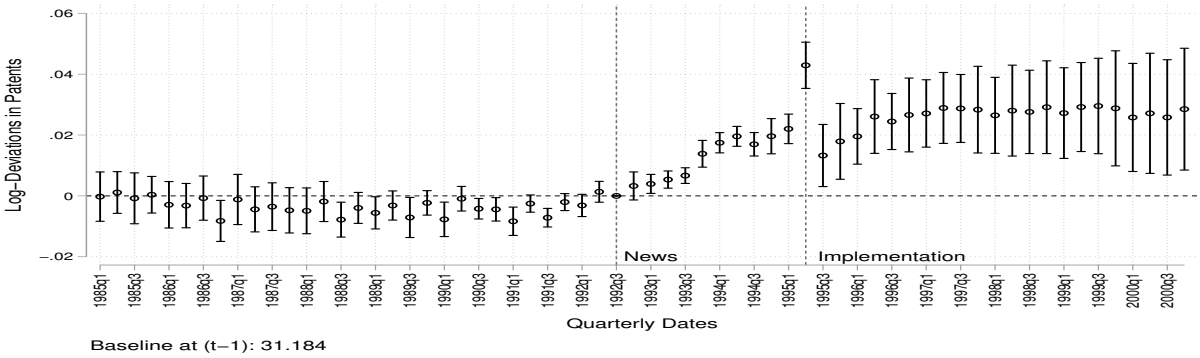
[21] For example, one year after the news and two years before implementation, a one-month shorter extension in patent term leads to a relative increase of around 4 citations-weighted granted patents per technical field-quarter, a +1.2% increase according to Poisson regression results (Figure 5c).

Figure 5: Poisson DiD estimates by field

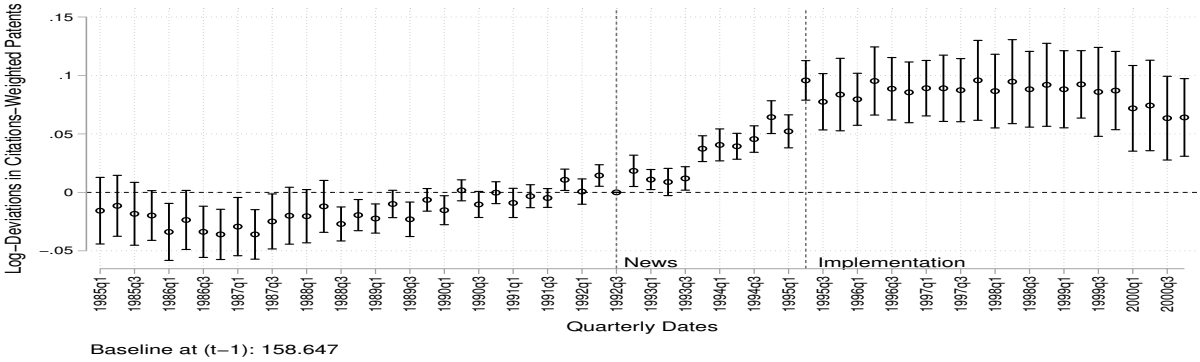
(a) Granted Patents and Average Patent Term Change



(b) Granted Patents and Patent Term Reduction Probability



(c) 5-Year Forward Citations-Weighted Patents and Average Patent Term Change



Notes: Panels (a) and (c) of the figure represent the *relative* effects of a one-month *shorter* extension in average patent term on log-deviations in (granted) patent applications (panel a) and citations-weighted patents (panel c) by application quarter, based on pseudo-Maximum Likelihood estimates of Poisson DiD specification (6). Panel (b) represents the *relative* effects of a one-percent *higher* patent term reduction probability on log-deviations in (granted) patent applications by application quarter, based on estimates of Poisson DiD specification (6). Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

Fact 2 After implementation, the sum of direct and indirect policy effects (overall effect) is such that innovation and R&D remain relatively higher in fields with a higher probability of patent term reduction and, thus, with an shorter average patent term extension.

Section 5.3 will estimate the *direct* post-implementation effect of the policy.

5.2 From Patents to Firm-level R&D

Next, I examine the impact of the patent term change on R&D expenditures, a key input to innovation. This analysis addresses concerns about using patent outcomes as proxies for innovation. Since some inventions are protected by secrecy rather than patents, changes in the patent term could influence both actual innovation and patenting decisions. I show that firm-level R&D expenditures respond to the patent term change in line with its effects on patent-based innovation outcomes, suggesting that the latter reflect, at least to some extent, genuine innovation changes.[22]

I use a direct measure of R&D expenditures from a yearly panel (1985-2000) of 2,421 listed U.S. firms, based on the NBER-Compustat dataset compiled by Hall, Jaffe and Trajtenberg (2001). For each firm i, the TRIPs-induced change in patent term, ΔT_i, is calculated as a weighted average of field-specific patent term changes, ΔT_j, with weights representing firm i's technological exposure to field j before TRIPs. This exposure is measured by the share of the firm's patents filed in field j from 1971 to 1991. I collect firm-level R&D expenditures (`xrd`), sales, and other variables from COMPUSTAT (Standard&Poor's, 2022). The firm-level analysis uses the following Poisson DiD specification.

$$Y_{i,t} = \exp\left\{\alpha_i + \sum_{\substack{k=1987 \\ k\neq 1991}}^{2000} \gamma_k \mathbf{1}_{(t=k)} + \sum_{\substack{k=1987 \\ k\neq 1991}}^{2000} \beta_k \mathbf{1}_{(t=k)} \Delta T_i + \theta' \mathbf{X}_{i,t} + \varepsilon_{i,t}\right\} \tag{8}$$

The regression compares R&D investment and patent outcomes across firms with varying exposure to the TRIPs-induced patent term change, both before and after the policy shocks. Identification assumes that, conditional on controls, the "shift" in patent term across technologies is quasi-random, despite the potential endogeneity of firms' technological exposure. This assumption is supported by several analyses in Section 5.5. Equation (8) includes firm fixed effects (α_i), year fixed effects, a vector of controls ($\mathbf{X}_{i,t}$) that incorporates firm age fixed effects, 3-digit SIC industry-year fixed effects, and a 3-digit SIC-specific quadratic trend in firm age. The error term is $\varepsilon_{i,t}$. Each DiD coefficient β_k captures the effect of an expected unit change in ΔT_i on the log deviation (or approximate percentage change) of R&D

[22] Section 5.5.1 further addresses issues related to innovation measures and secrecy. Appendix E explores how patent term changes affect sectoral Total Factor Productivity (TFP) and prices.

expenditures or innovation. I estimate the regression using pseudo-Maximum Likelihood.

Figures 6a and 6b illustrate the impact of a one-month shorter extension in average patent term on R&D expenditures and citations-weighted patents, respectively. The bands represent 95% confidence intervals, with standard errors clustered at the firm level.

The results align with the field-level findings from Section 5.1. During the news phase, a one-month shorter extension in patent term is associated with a 3.6% increase in yearly firm-level R&D (1993 estimate), equivalent to approximately $2.8 million. Similarly, news of a one-month shorter patent term extension corresponds to an 8% larger increase in citations-weighted patents at the firm level. The timing of these effects is also consistent with Section 5.1, as the change in citations-weighted patents is more gradual than the R&D response. This suggests that the quicker adjustment in R&D inputs leads to a slower change in innovation output.

Furthermore, for both R&D and innovation, the post-implementation effects maintain the same direction as the news phase, consistent with Fact 2. This indicates that firms with shorter patent term extensions experience relatively higher R&D and innovation not only during the news phase but also after implementation.

Appendix D.1 reports consistent results for firm-level patenting and patent value, and placebo analyses of firm-level variable costs, capital expenditures, and sales. Moreover, for the universe of US applicants, Appendix D.2 shows evidence of within-firm reallocation of patenting and inventors across technical fields that is consistent with field-level findings.

5.3 Direct Implementation Effect

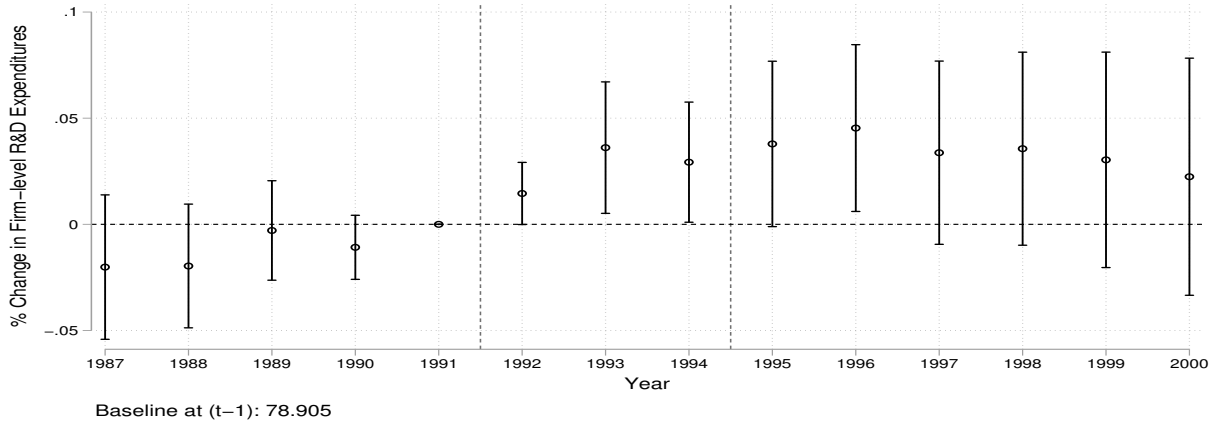
I now present evidence on the *direct* effect of patent term changes on innovation. As shown in expression (4), the post-implementation DiD estimates from Sections 5.1 and 5.2 reflect both indirect effects from policy anticipation and direct effects from the new patent term's actual implementation. The augmented DiD specification (7) isolates the direct effects by controlling for the influence of the news shock through field-specific innovation histories.

Figure 7 shows the direct effect of a one-month shorter extension in average patent term on the number of granted patents by field and application quarter, based on DiD estimates $(\widehat{\phi_k})$ from specification (7).

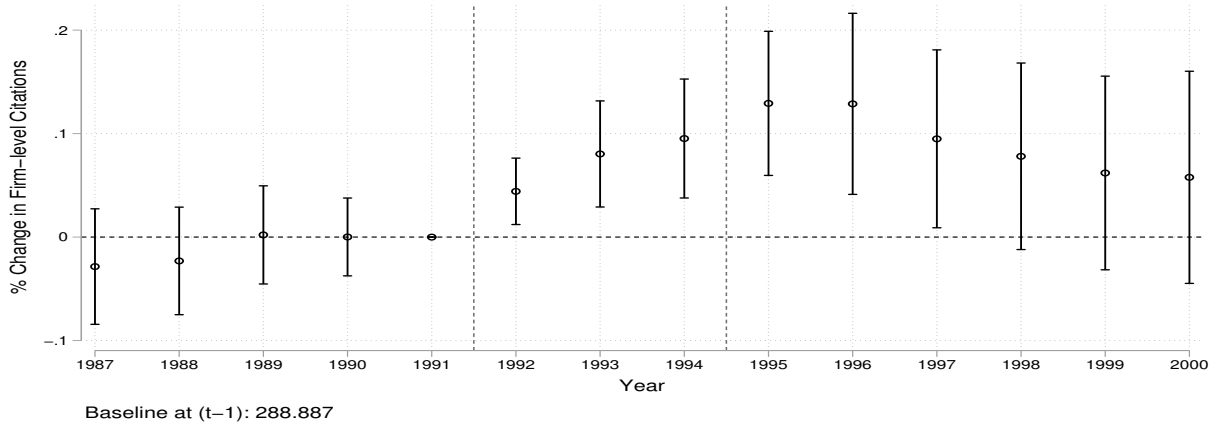
The figure highlights two main findings. First, the direct effect estimates remain positive *before* implementation, consistent with expression (3) stating that the news effect is mostly driven by the policy's direct impact. Thus, in fields with a relatively shorter patent term extension—and a higher probability of a patent term reduction—innovation accelerates more during the news phase (Fact 1).

Figure 6: Poisson DiD estimates for firm-level R&D and innovation

(a) R&D Expenditures



(b) Citations-Weighted Patents



Notes: The figure depicts the effect of a one-month *shorter* extension in average patent term at the firm level—based on the Poisson DiD specification (8)—on yearly R&D expenditures (panel a) and citations-weighted (granted) patents by application year (panel b). Point estimates represent the effect of the policy change on *percent deviations* of the outcome variable from its baseline value in 1991, i.e. \$78.9 million for R&D and 288.9 for citations. Standard errors are clustered at the firm level and bands represent 95% confidence intervals. The first and second vertical lines mark news and implementation events, respectively.

Second, the direct effect estimates turn *negative* and stable *after* policy implementation. A *shorter* extension in average patent term leads to a *smaller* increase in innovation, once controlling for innovation dynamics during the news phase. On average, a one-month shorter extension results in 0.64 fewer patents per quarter, or about a -1.8% reduction in the average quarterly number of patents in the post-implementation period. Using the field-specific probability of patent term reduction as the main regressor yields similar results.

Appendix C.10.1 provides consistent evidence for quality-adjusted innovation (citations-weighted patents) by field. Furthermore, Appendix C.10.2 shows that controlling for inno-

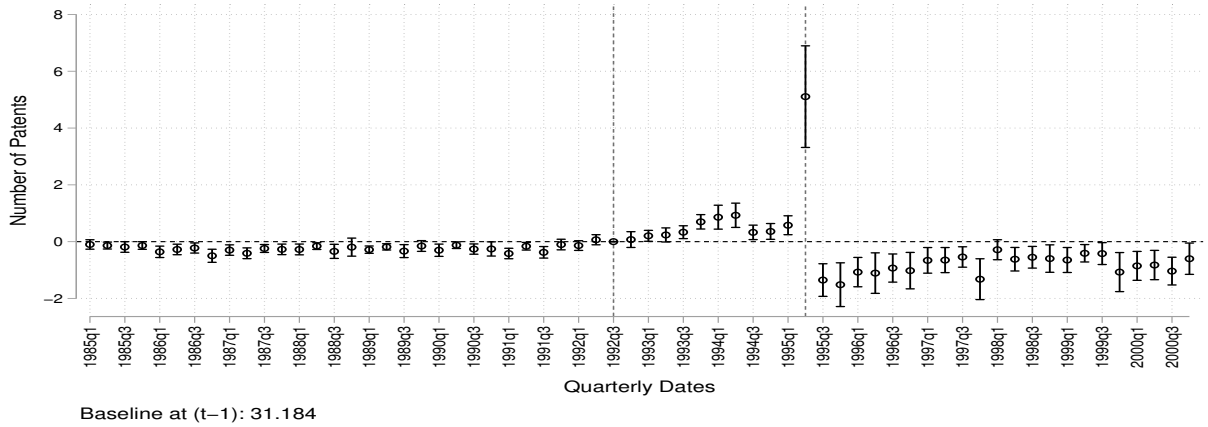
vation histories in related fields has a negligible impact on the results.

Therefore, absent anticipation effects, a shorter (longer) patent term hinders (stimulates) innovation as a direct effect, which I summarize as:

Fact 3 A *shorter* extension in average patent term determines a relatively *smaller* increase in innovation after policy implementation as a *direct* effect.

This result aligns with previous empirical studies on the effects of patent protection on innovation, which I will relate to the magnitude of my estimates next.

Figure 7: DiD estimates of direct effect controlling for anticipation



Notes: The figure depicts the relative *direct* effect of a one-month *shorter* extension in the average patent term on the level of (granted) patents by application quarter, based on estimates of DiD specification (7) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

5.4 Elasticity and Comparison with Previous Literature

Estimates of the direct post-implementation effect indicate that a one-month longer patent term increases innovation by 1.8%, implying a semi-elasticity of 21.6% for a one-year extension. This aligns with findings from Budish, Roin and Williams (2015) and Budish, Roin and Williams (2016), who estimate that a one-year extension in patent monopoly boosts R&D by 7% to 22% in the pharmaceutical industry. My estimates are at the upper end of this range. Similarly, Hémous et al. (2023)'s model predicts that a one-month increase in US patent terms would raise US innovation by 1.2%, close to my estimate of 1.8%.

These results are also consistent with the elasticity of TFP to market size estimated by Beerli et al. (2020), who found that a 1% increase in market size raises firm-level productivity by 0.46% using Chinese manufacturing data. Since a longer patent term expands the market size for innovations, I can map the results from Section 5.3 to their findings. In Appendix E,

I relate patenting variation due to the patent term change to sectoral TFP dynamics. From the direct post-implementation effects on patenting, I estimate an elasticity of aggregate TFP to patent term of around 0.33%, close to Beerli et al. (2020)'s estimate.

To summarize, I finally compute the elasticity of innovation to patent term based on *direct* post-implementation effects and news effects. A 1% *longer* extension in average patent term leads to a 3.6% *increase* in innovation post-implementation. In contrast, *news* of a 1% *shorter* anticipated extension for future inventions correlates with a 2.7% *increase* in patenting during the anticipation phase.[23] These findings on anticipation effects are novel to the literature and highlight the relevance of intertemporal trade-offs in the response of innovation to patent term changes.

5.5 Identification and Measurement Concerns

I conclude this section by addressing key identification and measurement concerns, focusing on the main DiD specification (5). This specification provides unbiased estimates of the news and implementation effects (3) and (4) under two conditions: accurate measurement of innovation and R&D, and the assumption that the policy-induced change in patent term across technological fields is as good as random. Specifically, the pre-news pending period across technical units must be unrelated, conditional on controls, to factors—observable or unobservable—that influence innovation across fields. To address potential violations of these conditions, I present complementary analyses.

5.5.1 Measurement of Innovation, Patent Quality, and Secrecy

The use of patent-based variables to measure aggregate innovation has known limitations. Not all inventions are patented, and changes in patent term can influence both the incentive to innovate and the choice between patenting or using secrecy to protect an innovation. Section 5.2 addressed this concern by showing that changes in patent term impacted firm-level R&D expenditures—a key input to innovation—in line with observed effects on patent outcomes. Appendix E further links changes in patent-based innovation measures to sectoral productivity and price variations.

In this subsection, I examine the effect of the patent term change on various measures of patent quality, which could be affected by shifts in selection into patenting. The direction of this effect is uncertain. In fields with a higher probability of patent term reduction,

[23] These estimates are derived from the average DiD estimates in the anticipation period (excluding 1995Q2, the pre-implementation quarter) divided by the average quarterly number of patents or inventors. This gives a semi-elasticity for a one-day change in patent term, which is then scaled by the number of days in the pre-TRIPs 17-year patent term.

innovators might opt to reduce quality to accelerate filings. Conversely, in fields with longer patent term extensions, the increased attractiveness of patent protection might lead more patenting of innovations that are harder to imitate, which commonly correlates with higher quality.[24] However, if few inventions hover around the quality selection threshold, selection may not be a quantitatively important driver of the observed results and average quality would not change significantly on aggregate.

I empirically investigate the policy's impact on several quality and value measures. Overall, I do not find consistent changes in average quality, suggesting that selection and patenting strategies had a limited effect on aggregate outcomes.

In Appendix C.2, I report no significant news or implementation effects on average patent originality and generality (Hall, Jaffe and Trajtenberg, 2001). Text-based quality metrics (Arts, Hou and Gomez, 2021), such as the introduction of new words or word combinations in patent texts, also show no news effects. Similarly, estimates of economic value (Kogan et al., 2017) and the average patent renewal rate are barely impacted by the policy. Average citations per patent, a proxy for scientific quality, and text similarity with preceding and subsequent patents (Arts, Hou and Gomez, 2021; Kelly et al., 2021) increase slightly after the implementation of a shorter patent term. These findings suggest that quality changes align with the overall increase in patenting and R&D activity.

Moreover, Appendix C.3 shows no significant effect on patent scope indicators (Marco, Sarnoff and deGrazia, 2016), such as the number or brevity of independent claims. This indicates that patenting strategies or claim formulations were not influenced by the patent term change. Innovators in fields anticipating a shorter patent term did not attempt to create patent thickets by filing narrower, multiple applications (Shapiro, 2001).

5.5.2 Technological Trends and Definition of Fields

Field- and quarterly fixed effects cannot fully account for field-specific technological trends that are independent of the policy but correlated with the pre-news average pending period. These trends could be influenced by macroeconomic factors, such as reduced defense spending after the Cold War, or technological shifts, like the rise of Information Technologies in the 1990s (Hall, 2004), or the emergence of complex technologies associated with longer average pending periods. I address this concern in several ways.

First, I test for the absence of significant pre-trends in the DiD estimates using the approach of Roth (2022). Second, I enrich specification (5) by adding quarter-by-3-digit

[24] Inventions that are more difficult to imitate are typically of higher quality and have less incentive to seek patent protection. Thus, a longer patent term may raise the quality threshold below which innovators choose to patent instead of relying on secrecy.

IPC fixed effects, controlling for technological trends across over 100 3-digit technical areas. This ensures that the DiD estimates reflect only the variation in patent term and innovation outcomes across 4-digit fields within the same 3-digit category.[25] Appendix C.4 shows that the DiD estimates are unaffected by this adjustment. Third, I verify in Appendix C.4 that the results remain robust when using 8-digit IPC subclasses as cross-sectional units of analysis while controlling for trends by 4-digit field. Lastly, I control for quarter-specific effects of technological complexity across fields, based on the classification by Galasso and Schankerman (2015). The results remain consistent with the baseline specification (5).

5.5.3 Heterogeneous Effects and Selection on Doses

In DiD settings with continuous treatments (dose), such as ΔT_j or PL_j, Callaway, Goodman-Bacon and Sant'Anna (2021) show that heterogeneous treatment effects across units, combined with selection on doses and the weighting scheme of the Two-Way Fixed Effects (TWFE) estimator, can distort the magnitude of TWFE DiD estimates for average causal responses (3) and (4). Appendix C.5 elaborates on this issue, using Callaway, Goodman-Bacon and Sant'Anna (2021)'s decomposition of TWFE to show that the DiD estimates in Figure 3 are constant across doses and generally stable. This indicates that fields more sensitive to patent term changes do not systematically experience longer or shorter patent term shifts, making the estimated effects likely representative of the average field.

However, the estimated effects of patent term changes do vary across fields along observable dimensions. Appendix C.6.1 shows that the policy's impact is larger in fields where a higher proportion of patents are renewed to the maximum term, suggesting greater responsiveness to patent term changes. Additionally, Appendix C.6.2 examines the role of uncertainty and finds that DiD estimates are larger in fields where the within-field dispersion in pending periods was smaller prior to the news shock, meaning inference on the TRIPs-induced change in patent term was less uncertain. Finally, Section 6 provides evidence on the heterogeneous effects of the policy based on field-specific competition intensity and technological dependence, which I identify as the main channel through which the news shock indirectly influences post-implementation innovation outcomes.

[25] For example, the 3-digit IPC C21, "Metallurgy of Iron," includes several 4-digit IPCs, such as C21B "Manufacture of Iron or Steel," C21C "Processing of Pig-Iron (...)," and C21D "Modifying the Physical Structure of Ferrous Metals (...)." Other examples include C25 "Electrolytic or Electrophoretic Processes," A43 "Footwear," and D03 "Weaving."

5.5.4 Endogenous Changes in Examination

While the TRIPs provisions did not formally modify the examination process at the USPTO, they may have influenced the rigor of examiners. If these changes correlate with the pre-TRIPs pending periods across technical units—potentially due to differential congestion or examination styles—the DiD coefficients in (5) could be biased for the marginal responses (3) and (4).

Since ungranted US patent applications were not published before 2001, assessing the impact of the patent term change on grant rates across fields is not feasible. However, in addition to finding inconclusive evidence on how the policy affected patent quality conditional on grants, Appendix C.7 confirms that the change in patent term did not correspond to differing trends in average pending period across fields, partly due to a reallocation of examiners to technical units where the number of applications increased the most due to TRIPs.[26] Finally, Appendix C.7 shows that DiD estimates are unaffected by using the change in patent term based on actual pending period by quarter as treatment in (5) instead of the average pre-TRIPs pending period.

5.5.5 Changes in Trade Regulation Concomitant to TRIPs

The Uruguay Round of Agreements included not only the adoption of TRIPs but also changes to maximum tariffs for countries joining the WTO. Coelli, Moxnes and Ulltveit-Moe (2022) demonstrate that tariff reductions positively influenced innovation in several countries. Therefore, specification (5) accurately estimates the news and implementation effects only if the field-specific impacts of tariff and trade regulation changes do not correlate with variations in patent term across technologies.

In Appendix C.8, I show that DiD estimates remain consistent when controlling for technology-specific changes in import tariff intensity in the US, European countries, and China during the period from 1996 to 2001. This consistency is due to the negligible correlation between tariff changes and the TRIPs-induced changes in patent term across fields.

5.5.6 Delayed and Uncertain Adoption of TRIPs in Developing Countries

The TRIPS Agreement required many Low- and Middle-Income Countries (LMICs) joining the WTO to adopt stronger patent protection, though delays were granted based on each country's economic development. These changes likely had opposing effects on U.S. patenting during the study period. While stronger protection may have benefited U.S. innovators by

[26]Moreover, the policy increased applicants' incentives to respond more swiftly to USPTO inquiries, as documented by Lemus and Marshall (2018) for the pharmaceutical sector.

improving market access in LMICs, uncertainty surrounding adoption may have deterred firms from investing in innovations targeting those markets.[27] If the U.S. patent term change correlates with these factors across fields, specification (5) may be biased. To address this, Appendix C.9 shows that the U.S. patent term change did not influence the intensity with which U.S. innovators sought patent protection in these LMICs across fields.

6 Interpretation and Transmission of News Shock

In this section I propose an interpretation for the documented news and post-implementation effects in light of existing theories. Then, I present empirical evidence on the transmission mechanism of the news shock to post-implementation innovation.

6.1 Interpretation

The estimated *direct* policy effect after implementation suggests that, without anticipation, a longer patent term would stimulate more innovation, aligning with previous empirical findings on patent protection and innovation (Section 5.4).

We can thus interpret news effects (Fact 1) consistent with the sign of the direct effect. Since innovators generally prefer a longer to a shorter patent term and because most fields featured a substantial probability of reduced effective patent duration under the new policy, innovation broadly increased. However, the motivation to speed up innovation and submit patent applications before the new policy took effect was stronger in fields with a higher probability of term reduction. Therefore, between news and policy implementation, R&D and innovation accelerated more in fields with a higher probability of a future patent term reduction (and a shorter average patent term extension). This increase in patenting reflects not only a re-timing of applications but also an intertemporal shift in actual R&D efforts, as seen in the firm-level analysis of R&D investment (Section 5.2). The gradual acceleration in patenting, leading up to policy implementation, is consistent with increased R&D activity over time.

These findings contribute to a growing literature on short- and medium-term R&D and innovation responses to temporary changes in innovation value, triggered by factors like aggregate shocks or demand shifts. Theoretically, multiple studies develop general equilibrium models where R&D is procyclical due to fluctuations in the market value of innovations (e.g., Comin and Gertler, 2006; Benigno and Fornaro, 2018; Anzoategui et al., 2019; Bianchi,

[27] Kyle and McGahan (2012) find that U.S. pharmaceutical firms increased innovation investments following TRIPS, benefiting from enforceable patents in new markets. Additional analyses, available upon request, yield similar results when excluding pharmaceutical and biotechnology-related fields.

Kung and Morales, 2019; Vinci and Licandro, 2020; Fornaro and Wolf, 2021; Bertolotti and Lanteri, 2024). Empirical studies also confirm procyclical R&D and innovation patterns (Barlevy, 2007; Argente, Lee and Moreira, 2018; Bertolotti, Gavazza and Lanteri, 2023). In this context, TRIPs policy news introduced a temporary shock, shifting the relative value of innovations obtained before and after implementation and driving the observed short-term R&D and innovation responses.

The empirical analysis also shows that news effects persist post-implementation (Fact 2). Specifically, in fields with shorter patent term extensions—where innovation accelerated the most following the news—patenting activity remains sustained for at least five years after implementation. This persistence is due to indirect effects that outweigh the direct policy effect, which would otherwise lead to relatively lower innovation in fields with shorter term extensions (Fact 3). This sustained shift in innovation patterns aligns with endogenous growth theory, as knowledge creation is cumulative; current innovations build on past advancements through a "standing on the shoulders of giants" effect (Romer, 1990). Temporary shocks to innovation may thus yield long-lasting impacts, especially since recent discoveries—such as those spurred by policy news—contribute more significantly than older knowledge to new projects and follow-up innovations, as evidenced by Hedge, Herkenhoff and Zhu (2022). Persistent productivity and output effects following temporary R&D and innovation changes also feature prominently in medium-term business cycle models like Comin and Gertler (2006).

6.2 Transmission Channels

In this subsection, I empirically examine the drivers of indirect effects—specifically, the transmission mechanism of the news shock to post-implementation innovation. Building on the theoretical insights of Section 6.1, I first focus on knowledge externalities. I then explore competition as an alternative channel in Appendix F.2.

The analysis of knowledge externalities leverages cross-field differences in the degree to which new inventions build on prior ones, using patent citations to measure technological dependence. Citations are commonly viewed as indicators of technological links between inventions.[28] Since most knowledge externalities occur within the same technology field (Liu and Ma, 2021), I focus on citations among patents classified in the same field. This approach is also supported by the analysis of direct effects conditional on cross-field innovation histories (Appendix C.10.2). Thus, I proxy technological dependence by the average number of

[28]Not all citations indicate genuine knowledge flows (see, e.g., Kuhn, Younge and Marco, 2020). However, citation-based tracking of technological dependence should still be reliable for this study's sample period despite limitations highlighted for more recent years.

backward citations made by patents in field j to earlier patents in the same field, published within three years of the citing patent's application. This time restriction reflects the length of TRIPs anticipation period. I denote this measure by $\overline{BB}_{jj}$.

I first analyze how the strength of post-implementation effects varies across fields based on the technological dependence of new innovations on prior ones. In fields with stronger technological dependence, I expect the acceleration in innovation during the news phase to result in stronger persistence of higher innovation levels after implementation, holding constant the average patent term change. To test this hypothesis, I estimate the triple-difference specification (9), which interacts within-field technological dependence $\overline{BB}_{jj}$—measured before the news shock—with the average patent term change ΔT_j.[29]

$$\begin{aligned} Y_{j,t} = \alpha_j + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \gamma_k \mathbf{1}_{(t=k)} + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \eta_k \mathbf{1}_{(t=k)} \overline{B}_{jj} \\ + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \beta_k \mathbf{1}_{(t=k)} \Delta T_j + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \theta_k \mathbf{1}_{(t=k)} \Delta T_j \times \overline{B}_{jj} + \varepsilon_{j,t} \end{aligned} \tag{9}$$

Like in previous specifications, α_j represents field fixed effects, $\mathbf{1}_{(t=k)}$ are quarter-specific dummy variables, and $\varepsilon_{j,t}$ is the idiosyncratic error term. In this triple-difference specification, the DiD coefficients β_k on the interactions of ΔT_j with quarterly dummies capture the relative effect of an anticipated one-day change in average patent term on the innovation outcome $Y_{j,k}$ *conditional* on zero technological dependence (i.e., $\overline{B}_{jj} = 0$). The new triple-difference coefficients θ_k indicate how the relative effect of the anticipated patent term change ΔT_j varies with the level of technological dependence, $\overline{B}_{jj}$.

Figure 8a shows the triple-difference effects of a one-month *shorter* average patent term extension, conditional on a unit-larger degree of technological dependence (i.e., an average increase of one within-field applicant citation, or $\overline{BB}_{jj} = 1$). Positive post-implementation effects suggest that stronger technological dependence correlates with a greater persistence of sustained levels of innovation in fields where innovation increased in the anticipation phase, in response to news of a shorter patent term extension (Fact 1). This finding aligns with theoretical expectations from endogenous growth models.

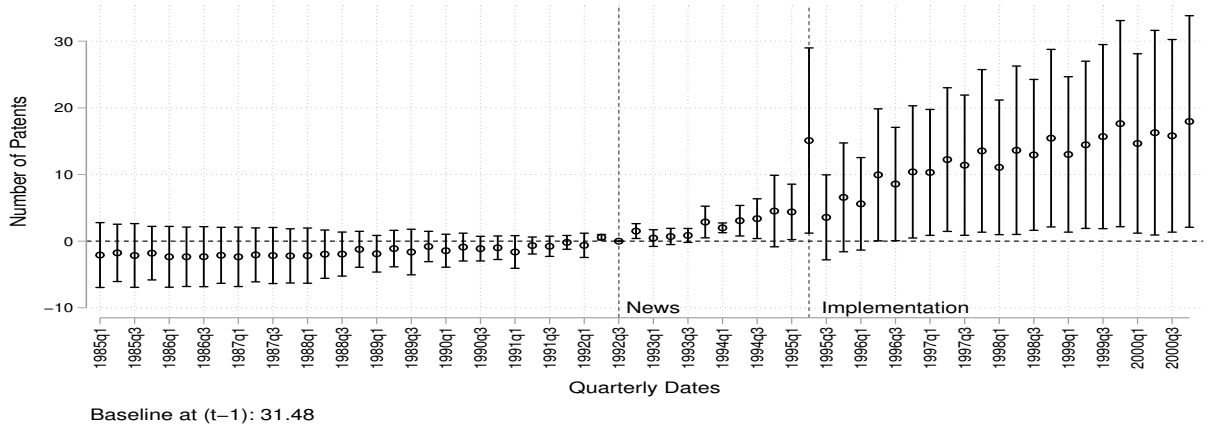
Furthermore, Figure 8b shows that, in the absence of technological dependence (i.e., conditional on $\overline{BB}_{jj} = 0$), a one-month *shorter* patent term extension results in *lower* post-implementation innovation compared to fields with a longer patent term extension. This result is consistent with the sign of the direct effect of patent term after policy implemen-

[29] The variable $\overline{BB}_{jj}$ is calculated using patents filed between 1980Q1 and 1989Q4.

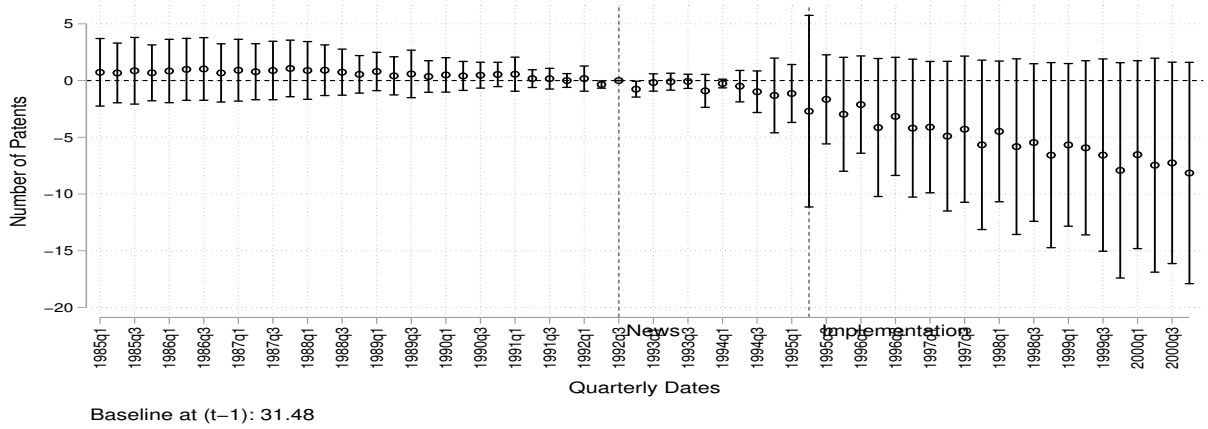
tation (Fact 3). Indeed, when past innovation has negligible impact on subsequent research productivity, post-implementation outcomes should be less affected by anticipation dynamics and thus reflect direct effects, consistent with the decomposition of equation (3).

Figure 8: Heterogeneity analysis based on within-field technological dependence

(a) Triple Difference Effects



(b) Difference-in-Difference Effects conditional on $\overline{BB}_{jj} = 0$



Notes: The figure depicts triple-difference (panel a) and difference-in-difference (panel b) effects of a one-month *shorter* extension in average patent term on the number (granted) patent applications by application quarter based on estimates of triple-difference specification (9). In panel (a), coefficients represents the relative change in the level of the outcome variable due to a one-month *shorter* extension in average patent term conditional on a unit-larger level of technological dependence, i.e., one more average within-field applicant-made citation ($\overline{BB}_{jj} = 1$). In panel (b), coefficients represent the relative change in the level of the outcome variable due to a one-month *shorter* extension in average patent term conditional on the absence of technological dependence—i.e., for $\overline{BB}_{jj} = 0$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

In a second analysis, I examine how within-field technological dependence endogenously

responds to the patent term change. If the post-implementation persistence of high innovation is driven by stronger knowledge diffusion, patents filed after implementation should increasingly cite recent inventions. Thus, measures of technological dependence should rise more in fields with shorter patent term extensions, where patenting accelerated the most during the news phase.

To test this hypothesis, I estimate DiD specification (5) using two alternative measures of technological dependence as dependent variables. First, I use the quarterly average number of backward citations made by applicants to prior patents within the same field, denoted as $\overline{B}_{jj,t}$. Second, I examine the share of patents filed in quarter t and classified in field j that include at least one applicant-made backward citation to previous patents in the same field, denoted as $S_{jj,t}$.

Panels (a) and (b) of Figure 9 show the DiD estimates for the effect of a one-month *shorter* average patent term extension on $\overline{B}_{jj,t}$ and $S_{jj,t}$, respectively. A shorter extension leads to an increase in time-varying measures of technological dependence post-implementation, consistent with theoretical expectations. The sustained high innovation rates following the implementation shock are largely driven by innovations building on recent discoveries, increasing technological dependence measures in fields with the highest acceleration during the news phase—i.e., those with shorter patent term extensions.

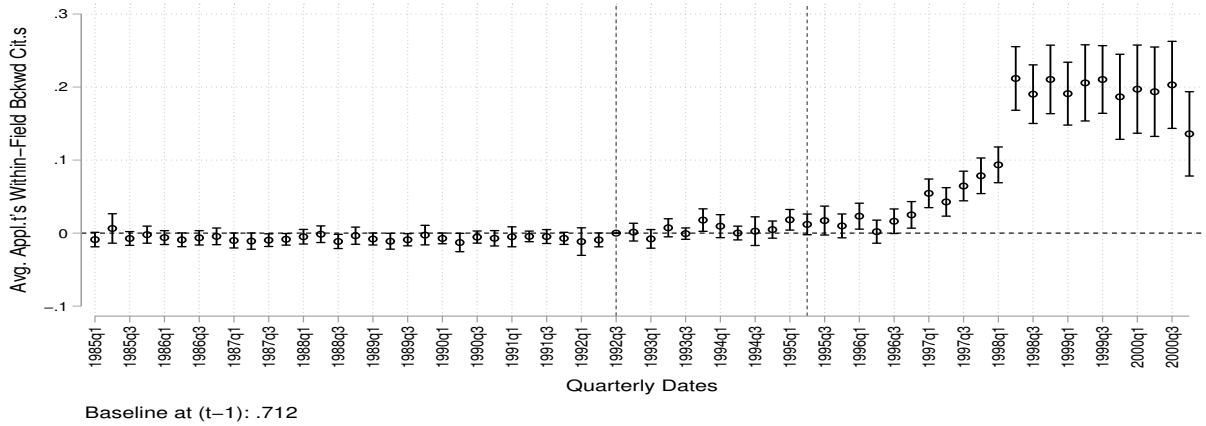
The timing of this effect, which peaks around four years after implementation, sheds light on the half-life of the externality. This four-year lag aligns with knowledge diffusion lags of approximately two years—proxied by the average time between application (invention) and patent publication around the TRIPs policy change—and an additional R&D gestation lag of roughly two years, close to Pakes and Schankerman (1984)'s estimates of a 1.5-year gestation period.

In additional analyses, I document direct citation links from patents applied for in the post-implementation phase to those filed during the news period. Furthermore, I decompose post-implementation indirect effects, finding that they mostly occur *between* different firms as opposed to *within* firms. This supports the interpretation of these effects as the result of an externality. These additional results are available upon request.

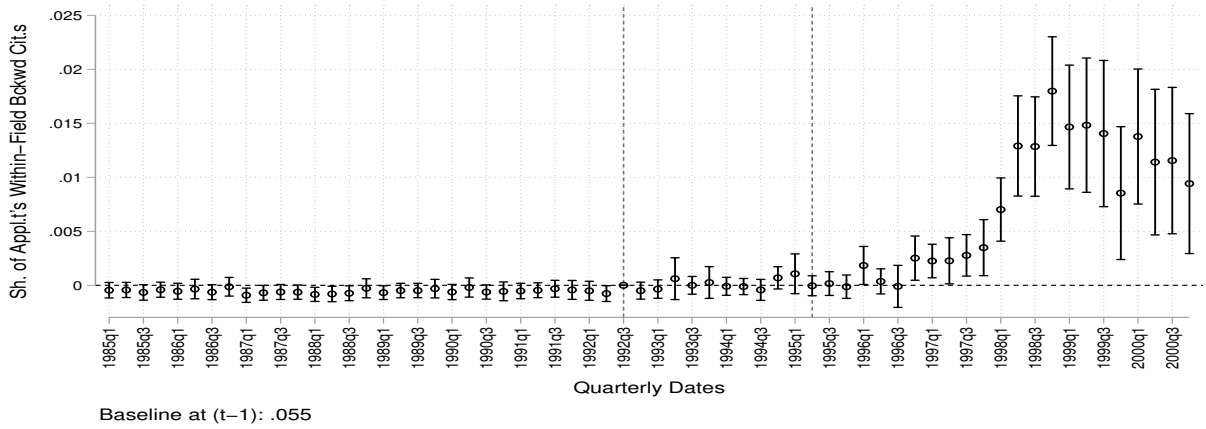
Finally, Appendix F.1 provides a synthetic measure of the strength of the observed knowledge externality, derived from an aggregated version of equation (7) across policy phases. The resulting elasticity of future innovation to a 1% positive shock to current innovation is 2.1, meaning that a 1% increase in innovation during the news phase is associated with a 2.1% rise in post-implementation innovation. This magnitude aligns with the indirect effects outweighing direct effects in the post-implementation period.

Figure 9: DiD estimates for within-field technological dependence by field

(a) Average Applicant-made Within-field Backward Citations



(b) Share of Patents with Applicant-made Within-field Backward Citation



Notes: The figure illustrates the relative effect of a one-month shorter extension in average patent term on the level of the average number of applicant-made backward citations per patent classified in field j and applied for in quarter t (panel a) and on the level of the share of patents classified in field j and filed in quarter t that have at least one applicant-made backward citation to patents also classified in field j (panel b). Effects are computed based on the estimates of DiD specification (5).Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

6.3 Alternative Channels

Appendix F.2 further examines alternative transmission mechanisms, specifically competition and strategic patenting by incumbent firms. The findings indicate that these alternative channels play a minimal role.

Regarding competition, the patent term change has a greater impact in more competitive fields, measured by the concentration of patents across applicants or the entry rate of new

applicants. When competitive pressure is higher, patent protection becomes more valuable, making innovators more responsive to changes in patent duration. However, the patent term change had no direct effect on competition metrics, suggesting that shifts in competition are not the primary driver of the indirect effects associated with Fact 2. Additionally, the average quality of patents from incumbents did not decrease relative to those from new entrants, indicating that established firms did not alter their patenting strategies for preemptive reasons in response to a shorter patent term extension.

7 Conclusions

When transitory shocks impact the economy, even temporary changes in R&D can lead to persistent shifts in innovation and productivity, significantly affecting output and welfare. Such effects are especially pronounced for innovation-policy shocks, which are specifically designed to influence R&D and innovation.

This paper provides quasi-experimental evidence on these effects for an anticipated change in patent term, an essential innovation policy tool on which we had limited empirical evidence. The role of policy anticipation makes this analysis particularly relevant, as real-world policy changes often involve prolonged negotiations and thus inevitable news effects.

The findings show that a *shorter* average patent term extension leads to a *smaller* increase in post-implementation innovation as a direct effect, with magnitudes consistent with prior literature on patent protection's impact on innovation. A key contribution of this paper, however, is to document a significant acceleration in R&D and innovation during the anticipation phase when there is news of a future shorter extension and a higher probability of patent term reduction. Due to the cumulative nature of innovation, this initial acceleration stimulates further discoveries through technological externalities, temporarily offsetting the direct negative effect of the policy after implementation.

The external validity of these results depends on whether indirect effects, rooted in the cumulative nature of innovation, are stronger than the policy-specific direct implementation effect. Given the sizable direct effect of patent term on innovation relative to other policy tools, these findings likely extend to alternative policy levers.

Thus, anticipated innovation policy shocks can create short-term dynamics that go in the opposite direction of their long-run impacts, with larger magnitudes than other shocks due to their specific targeting of R&D and innovation. To accurately assess the effects of innovation policies, endogenous growth theories should account for the short-run intertemporal incentives of R&D decisions, which can also influence medium- and long-term outcomes.

References

Abrams, David S. 2009. "Did TRIPS Spur Innovation? An Empirical Analysis of Patent Duration and Incentives to Innovate." *University of Pennsylvania Law Review*, 157(1): 1613.

Acikalin, Utku, Tolga Caskurlu, Gerard Hoberg, and Gordon M. Phillips. 2022. "Intellectual Property Protection Lost and Competition: An Examination Using Machine Learning." *NBER Working Paper No. w30671.*

Aghion, Philippe, and Xavier Jaravel. 2015. "Knowledge Spillovers, Innovation and Growth." *The Economic Journal*, 125(March): 533–573.

Akcigit, Ufuk, Doug Hanley, and Nicolas Serrano-Velarde. 2020. "Back to Basics: Basic Research Spillovers, Innovation Policy and Growth." *Review of Economic Studies*, Forthcoming.

Akcigit, Ufuk, Santiago Caicedo, Ernest Miguelez, Stefanie Stantcheva, and Valerio Sterzi. 2020. "Dancing with the Stars: Innovation through Interactions." *NBER Working Paper No. 24466.*

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion. Economics Books*, Princeton University Press.

Anzoategui, David, Diego Comin, Mark Gertler, and Joseba Martinez. 2019. "Endogenous Technology Adoption and R&D as Sources of Business Cycle Persistence." *American Economic Journal: Macroeconomics*, 11(3): 67–110.

Argente, David, Munseob Lee, and Sara Moreira. 2018. "Innovation and product reallocation in the great recession." *Journal of Monetary Economics*, 93: 1–20.

Arts, Sam, Jianan Hou, and Juan Carlos Gomez. 2021. "Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures." *Research Policy*, 50(2): 104144.

Barlevy, Gadi. 2007. "On the Cyclicality of Research and Development." *American Economic Review*, 97(4): 1131–1164.

Becker, Randy A., Wayne B. Gray, and Jordan Marvakov. 2021. "NBER-CES Manufacturing Industry Database (1958-2018, version 2021a)."

Beerli, Andreas, Franziska J. Weiss, Fabrizio Zilibotti, and Josef Zweimüller. 2020. "Demand forces of technical change evidence from the Chinese manufacturing industry." *China Economic Review*, 60(C).

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019. "Who Becomes an Inventor in America? The Importance of Exposure to Innovation." *Quarterly Journal of Economics*, 134(2): 647–713.

Benigno, Gianluca, and Luca Fornaro. 2018. "Stagnation Traps." *The Review of Economic Studies*, 85: 1425–1470.

Bertolotti, Fabio, Alessandro Gavazza, and Andrea Lanteri. 2023. "Dynamics of Expenditures on Durable Goods: the Role of New-Product Quality." *The Economic Journal*, 133(652): 1641–1656.

Bertolotti, Fabio, and Andrea Lanteri. 2024. "Capital Replacement and Innovation Dynamics." CEPR Press, Paris & London. 18869.

Bianchi, Francesco, Howard Kung, and Gonzalo Morales. 2019. "Growth, Slowdowns, and Recoveries." *Journal of Monetary Economics*, 101: 47–63.

Bloomfield, Micheal, , Anita M. McGahan, and Keyvan Vakili. 2022. "The Exportation of Science from Low- and Middle-Income Countries." *mimeo*.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review*, 110(4): 1104–44.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen. 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica*, 81(4): 1347–1393.

Boldrin, Michele, and David K. Levine. 2013. "The Case against Patents." *Journal of Economic Perspectives*, 27(1): 3–22.

Budish, Eric, Benjamin N. Roin, and Heidi Williams. 2015. "Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials." *American Economic Review*, 105(7): 2044–85.

Budish, Eric, Benjamin N. Roin, and Heidi Williams. 2016. "Patents and Research Investments: Assessing the Empirical Evidence." *American Economic Review*, 106(5): 183–187.

Caicedo, Santiago, and Jeremy Pearce. 2024. "Need for Speed: Quality of Innovations and the Allocation of Inventors." FRB of New York 1127.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna. 2021. "Difference-in-Differences with a Continuous Treatment."

Coelli, Federica, Andreas Moxnes, and Karen Helene Ulltveit-Moe. 2022. "Better, Faster, Stronger: Global Innovation and Trade Liberalization." *The Review of Economics and Statistics*, 104(2): 205–216.

Comin, Diego, and Mark Gertler. 2006. "Medium-Term Business Cycles." *American Economic Review*, 96(3): 523–551.

Delgado, Mercedes, Margaret Kyle, and Anita M. McGahan. 2013. "Intellectual Property Protection and the Geography of Trade." *Journal of Industrial Economics*, 61(3): 733–762.

Doane, Michael. 1994. "TRIPS and International Intellectual Property Protection in an Age of Advancing Technology." *American University International Law Review*, 9(2): 465–497.

EPO, European Patent Office. 2017. "2017 Autumn Edition."

Fornaro, Luca, and Martin Wolf. 2021. "Monetary Policy in the Age of Automation." Working Paper.

Galasso, Alberto, and Mark Schankerman. 2015. "Patents and Cumulative Innovation: Causal Evidence from the Courts." *The Quarterly Journal of Economics*, 130(1): 317–369.

Gilbert, Richard, and Carl Shapiro. 1990. "Optimal Patent Length and Breadth." *RAND Journal of Economics*, 21(1): 106–12.

Goldschlag, Nathan, Travis Lybbert, and Nikolas Zolas. 2019. "An 'Algorithmic Links with Probabilities' Crosswalk for USPC and CPC Patent Classifications with an Application Towards Industrial Technology Composition." *Economics of Innovation and New Technology*, 1–21.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg. 2001. "The NBER Patent Citations Data File: Lessons, Insights and Methodological Tools." *NBER Working Paper No. 8498.*

Hall, Brownyn H. 2004. "Exploring the Patent Explosion." *The Journal of Technology Transfer*, 35–48.

Hedge, Deepak, Kyle F. Herkenhoff, and Chenqi Zhu. 2022. "Patent Publication and Innovation." *NBER Working Papers 29770.*

Hémous, David, Simon Lepot, Thomas Sampson, and Julian Schärer. 2023. "Trade, Innovation and Optimal Patent Protection." *CESifo Working Paper No. 10777.*

Higham, Kyle, Gaétan de Rassenfosse, and Adam B. Jaffe. 2021. "Patent Quality: Towards a Systematic Framework for Analysis and Measurement." *Research Policy*, 50(4): 104215.

Jaffe, Adam B. 1986. "Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value." *American Economic Review*, 76(5): 984–1001.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run." *American Economic Review: Insights*, 3(3): 303–20.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *The Quarterly Journal of Economics*, 132(2): 665–712.

Kuhn, Jeffrey, Kenneth Younge, and Alan Marco. 2020. "Patent citations reexamined." *RAND Journal of Economics*, 51(1): 109–132.

Kyle, Margaret K., and Anita M. McGahan. 2012. "Investments in pharmaceuticals before and after TRIPS." *Review of Economics and Statistics*, 94(4): 1157–1172.

Lanahan, Lauren, and Kyle R. Myers. 2022. "Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy." *American Economic Review*, 112(7): 2393–2423.

Lemus, Jorge, and Guillermo Marshall. 2018. "When the Clock Starts Ticking: Measuring Strategic Responses to TRIPS's Patent Term Change." *Research Policy*, 47(4): 796–804.

Lerner, Josh. 2009. "The Empirical Impact of Intellectual Property Rights on Innovation: Puzzles and Clues." *American Economic Review*, 99(2): 343–48.

Liu, Ernest, and Song Ma. 2021. "Innovation Networks and R&D Allocation." National Bureau of Economic Research Working Paper 29607.

Lychagin, Sergey, Joris Pinkse, Margaret E. Slade, and John Van Reenen. 2016. "Spillovers in Space: does geography matter?" *Journal of Industrial Economics*, 64(2): 295–335.

Marco, Alan C., Joshua D. Sarnoff, and Charles deGrazia. 2016. "Patent Claims and Patent Scope." *USPTO Economic Working Paper 2016-04*, 104144.

Martin, Lisa B., and Susan L. Amster. 1994. "International Intellectual Property Protections in the New GATT Accord." *Journal of Property Rights*, 9.

Matthews, Duncan. 2002. *Globalising Intellectual Property Rights.* Routledge.

Ma, Yueran, and Kaspar Zimmermann. 2023. "Monetary Policy and Innovation." National Bureau of Economic Research Working Paper 31698.

Montalvo, Patricia. 1996. "How Will the New Twenty-Year Patent Term Affect You? A Look at the Trips Agreement and the Adoption of a Twenty- Year Patent Term." *Santa Clara High Tech. Law Journal*, 139.

Moretti, Enrico. 2020. "The Effect of High-Tech Clusters on the Productivity of Top Inventors." *Working Paper.*

Moretti, Enrico, Claudia Steinwender, and John Van Reenen. 2019. "The Intellectual Spoils of War? Defense R&D, Productivity and Spillovers." *Working Paper.*

Morgese, Giuseppe. 2009. *L'accordo sull'aspetto dei diritti di proprietà intellettuale attinenti al commercio (TRIPs).* Cacucci Editore.

Moscona, Jacob. 2021. "Flowers of Invention: Patent Protection and Productivity Growth in US Agriculture." *Working Paper.*

Moser, Petra. 2005. "How Do Patent Laws Influence Innovation? Evidence from Nineteenth-Century World's Fairs." *American Economic Review*, 95(4): 1214–1236.

Moser, Petra, and Alessandra Voena. 2012. "Compulsory Licensing: Evidence from the Trading with the Enemy Act." *American Economic Review*, 102(1): 396–427.

Mullahy, John, and Edward C Norton. 2022. "Why Transform Y? A Critical Assessment of Dependent-Variable Transformations in Regression Models for Skewed and Sometimes-Zero Outcomes." National Bureau of Economic Research Working Paper 30735.

Nickell, Stephen J. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica*, 49(6): 1417–1426.

Nordhaus, William D. 1967. "The Optimal Life of a Patent." Cowles Foundation for Research in Economics, Yale University Cowles Foundation Discussion Papers 241.

Nordhaus, William D. 1969. *Invention, Growth, and Welfare A Theoretical Treatment of Technological Change.* MIT Press.

Pakes, Ariel, and Mark Schankerman. 1984. "The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources." *R&D, Patents and Productivity*, Chicago: University of Chicago Press: 98–112.

Reichman, Jerome H. 1993. "The TRIPs Component of the GATT's Uruguay Round: Competitive Prospects for Intellectual Property Owners in an Integrated World Market." *Fordham Intellectual Property, Media & Entertainment Law Journal*, 9: 171–266.

Romer, Paul. 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5): S71–102.

Roth, Jonathan. 2022. "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." *American Economic Review: Insights*, 4(3): 305–22.

Sakakibara, Mariko, and Lee Branstetter. 2001. "Do Stronger Patents Induce More Innovation? Evidence from the 1988 Japanese Patent Law Reforms." *RAND Journal of Economics*, 32(1): 77–100.

Schankerman, Mark, and Florian Schuett. 2017. "Screening for Patent Quality." *CEPR Discussion Paper 11688.*

Shapiro, Carl. 2001. "Navigating the Patent Thicket: Cross Licenses, Patent Pools, and Standard Setting." In *Innovation Policy and the Economy, Volume 1. NBER Chapters*, 119–150. National Bureau of Economic Research, Inc.

Standard&Poor's. 2022. "COMPUSTAT."

USPTO. 2023. ""Data Download Tables." PatentsView. Accessed September 2023."

Vinci, Francesca, and Omar Licandro. 2020. "Switching-track after the Great Recession." University of Nottingham, Centre for Finance, Credit and Macroeconomics (CFCM) Discussion Papers.

APPENDICES

Appendix A Data description

The section describes the data sources for the empirical analyses. The field-level analysis mainly uses PATSTAT (EPO, 2017) for innovation variables. I complement PATSTAT with several other sources. For generality and originality measures, I rely on Hall, Jaffe and Trajtenberg (2001). The economic value of patents is taken from Kogan et al. (2017), which estimates the private economic value of patents by exploiting the stock market reaction to patent grants. Text-based measures of patent quality are taken from Arts, Hou and Gomez (2021), characteristics of patent claims are obtained from Marco, Sarnoff and deGrazia (2016), additional information on disambiguated applicants, inventors, and examiners are obtained from PatentsView (USPTO, 2023).

For the firm-level analysis, I leverage the NBER Patent Database for patent information and the matching of patents to firm identifiers. From COMPUSTAT (Standard&Poor's, 2022) I obtain balance-sheet information on R&D expenditures and other financials for a sample of listed firms responsible for a relevant share of aggregate US R&D, and innovation.

For sectoral analyses, the NBER CES manufacturing database (Becker, Gray and Marvakov, 2021) provides annual industry-level data over 1958-2011 by 6-digit NAICS industries for TFP, prices, and other aggregates. Finally, the 'Algorithmic Links with Probabilities' crosswalk by Goldschlag, Lybbert and Zolas (2019) is used to map technical field (4-digit IPC classes) into industries (6-digit NAICS) and vice versa, based on text-analysis of patents abstracts and descriptions of different sectors.

Tables A1, A2, and A3 report summary statistics by technical field, firm, and industry, respectively. Table A4 reports variation in the effects of TRIPs on average patent term change and patent term reduction probability across broad technological areas.

Appendix B Implications of Policy Anticipation

In this section, I illustrate the consequences of neglecting policy anticipation and highlight the differences of my analysis compared to the work of Abrams (2009).

In Section 2.2 I discuss several sources suggesting the presence of policy anticipation, i.e., that firms were aware of the patent term change approximately 2.5 years before its implementation. In contrast, Abrams (2009) assumes no anticipation based on the uncertainty surrounding TRIPs adoption until mid-1994. Moreover, Abrams (2009)'s empirical strategy focuses on narrow windows (6, 12, or 24 months) around policy implementation in June 1995

Table A1: **Summary statistics by technical field and quarter**

Variable	Mean	S.D.	10th p.	90th p.
Granted Patents	36.09	136.35	0	78
5 years Citations	195.15	1070.1	0	360
Patent value (Million Dollars)	351.01	3341.87	0	406.01
Pending Days (days)	1022.1	496.2	558.93	1660.94
Change in avg. patent term (days)	472.66	117.42	343.55	590.79
Standard dev. of Δ patent term (days)	37.13	38.93	11.53	72.55
Patents share w. p.p.¿3y.	.06	.08	0	.12
Share of patents renewed to max. term	.29	.26	0	.63
Share of entrant applicants	.54	.21	.28	.82
Share of patents granted to entrants	.49	.23	.2	.8
Patents-based HHI	1191.51	1982.33	116.35	2800
N. Patent with w.-field bckwd. cit.s	4.89	27.07	0	10
Sh. Patent with w.-field bckwd. cit.s	.19	.28	0	.59

The Table reports the summary statistics of different variables used in the paper by technical field (4-digit IPC class) and quarterly date. Each time a (granted) patent application is classified in more than one 4-digit IPC, I count it in each technical field.

Table A2: **Summary statistics by firm and year**

Variable	Mean	S.D.	10th p.	90th p.
Granted Patents	13.57	94.75	0	12
Citations-weighted Patents	199.02	1589.75	0	173.72
Patents value (Million Dollar)	287.16	3434.04	0	74.8
Expected change in protection time	445.02	118.57	273.29	571.4
Sales (Million Dollar)	2337.66	9997.57	2.59	4456.27
Age	14.73	13.88	1	37
R&D Expenditure (Million Dollar)	60.87	359.06	0	58.78

The Table reports the summary statistics of different variables used in the paper by firm (COMPUSTAT firms) and year

Table A3: **Summary statistics by industry and year**

Variable	Mean	S.D.	10th p.	90th p.
Granted Patents	199.23	591.3	.44	491.7
Citations-weighted Patents	1433.73	6930.68	1.4	2656.01
Patents value (Million Dollar)	3015.86	16684.91	.38	4595.21
Avg. TFP Growth (p.p.)	.39	6.43	-6.09	6.86
Avg. Inflation (p.p.)	1.89	4.87	-1.53	6.26

The Table reports the summary statistics of different variables used in the paper by industry (6-digit NAICS) and year

Table A4: Patent term change and variation across technical areas

Technical Section	Nb. Fields	$\overline{\Delta T_j}$	S.D.(ΔT_j)	$\overline{\mathbf{Pr}(\Delta T_j < 0)}$	S.D.($\mathbf{Pr}(\Delta T_j < 0)$)
A Human Necessities	83	484.99	104.01	.06	.06
B Operations, Transp.	163	489.22	91.87	.05	.05
C Chemistry, Metallurgy	84	376.41	158.87	.13	.12
D Textiles, Paper	37	465.31	98.84	.05	.06
E Fixed Constructions	29	554.15	44.22	.03	.02
F Mechanical Engineering	98	531.53	94.48	.04	.03
G Physics	78	434.05	123.6	.08	.12
H Electricity	49	462.77	75.43	.06	.04

Notes: The table reports the number fields, i.e., 4-digit IPC classes (column 2), the average change in patent term (column 2), the standard deviation in average pending period (column 4), the average probability of losing protection time from TRIPs (column 5), and the standard deviation of this probability (column 6) across 1-digit IPC technical areas (column 1).

in a two-period (pre-post) DiD specification including field-specific linear trends.

In the following analyses, I argue that this specification and the hypothesis of no anticipation determines violations of the parallel trends assumption underlying the DiD strategy. These violations are responsible for the opposite sign of Abrams (2009)'s DiD estimates compared to mine.

I start by replicating Abrams (2009)'s analysis in my sample with the following two-period DiD

$$Y_{j,t} = \alpha_j + Post_t + \delta \Delta T_j + \beta Post_t \Delta T_j + \chi_j t + \mathbf{X}_{j,t}\gamma + \varepsilon_{j,t} \tag{B1}$$

which is equivalent to Specification (2) in Abrams (2009). Subscript j indexes technical fields (4-digit IPC classes in my setting) and t identifies months. The outcome variable $Y_{j,t}$ represents either the number of patents or the number of citations-weighted patents by application month. Field fixed effects α_j are included, and the dummy variable $Post_t$ takes a value of 1 if month-t comes after the policy implementation of June 1995 and 0 otherwise. ΔT_j is the policy-induced change in average patent term, and $\chi_j t$ represents a field-specific monthly linear trend. Additionally, $\mathbf{X}_{j,t}$ includes field-specific controls, such as the average number of inventors per patent and the average number of claims per patent. The DiD coefficient of interest is β.

Table B1 reports DiD estimates. Consistent with Abrams (2009)'s original analysis, a longer extension in average patent term relates to a larger increase in patents or citations-weighted patents, which is opposite to the direction of my DiD estimates, where a *shorter* extension in patent term generates a larger acceleration in innovation at news that persists after implementation.[30]

[30]Differences in the magnitude of the coefficients between Table B1 and Abrams (2009)'s results are due to different sample restrictions (Abrams (2009) restricts the sample to technical fields with at least thirty

Table B1: **Replication of Abrams (2009)'s results**

	(1) Patents	(2) Patents	(3) Citations	(4) Citations
$Post_t$	-12.238***	-17.261***	-55.151**	-76.602**
	(3.106)	(4.051)	(22.098)	(31.091)
$Post_t \times \Delta Tj$	0.020***	0.029***	0.072*	0.109*
	(0.006)	(0.008)	(0.041)	(0.059)
Avg. Num. of Inventors		0.216***		4.121***
		(0.048)		(0.539)
Avg. Num. of Claims		-0.005		0.537***
		(0.004)		(0.080)
Constant	14.640***	16.980***	108.129***	111.448***
	(0.200)	(0.235)	(1.763)	(2.372)
Field F.E.	Y	Y	Y	Y
Field-specific Trend	Y	Y	Y	Y
Observations	14904	12603	14904	12603

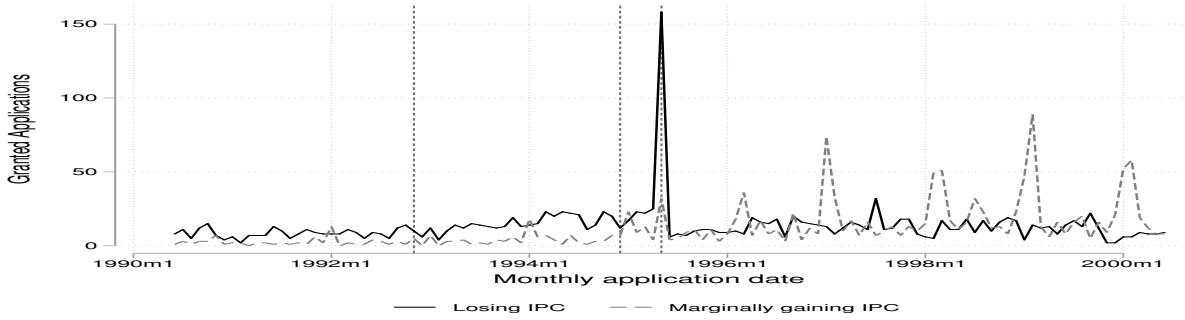
Notes: Columns (1) and (2) report the OLS estimates of specification (B1) using granted patents filed in month-t and classified in field-j as dependent variable. Columns (3) and (4) report the OLS estimates of specification (B1) using citations-weighted granted patents filed in month-t and classified in field-j as dependent variable. Standard errors are clustered by technical field. Statistical significance levels: * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

Next, I argue that the assumption of no anticipation in the context of the two-period DiD specification (B1) leads to violations of the parallel trends assumption. To this end, I provide a graphical intuition based on two technical fields, C12R and A01H, which experienced a slightly negative (-75 days) and positive (+50 days) change in average patent term due to TRIPs, respectively.

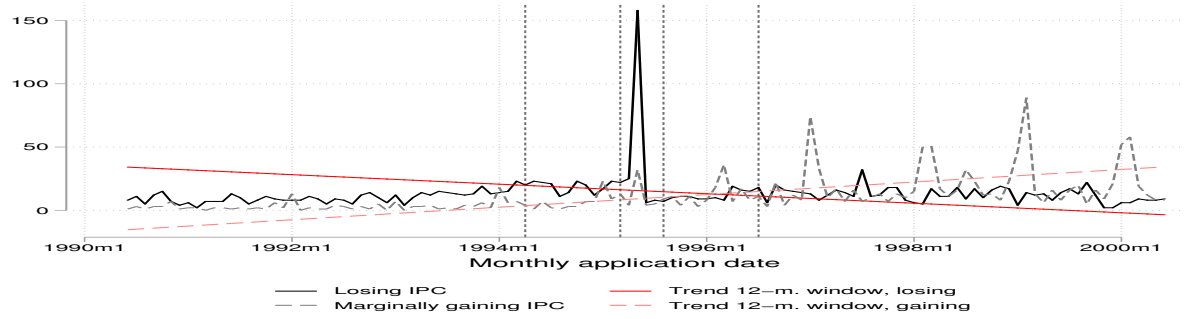
Figure B1a depicts the number of granted patents by application month over 1990-2000 in the two fields. The three vertical lines mark the news date (November 1992); the signing of the URAA (December 1994); and the implementation date (June 1995). The figure shows that the path of patenting was remarkably parallel in the two fields before news but started diverging after it, well before TRIPs ratification in December 1994. Therefore, field-specific time-trends estimated in a narrow sample around policy implementation, as implicit by specification (B1), provide an inaccurate representation of the field-specific path of innovation before the policy news. Figure B1b illustrates this point by plotting in red the estimated trends in the two fields. Vertical lines depict the limits of the estimation sample, i.e., the outer 12-month window around implementation date with an inner gap of 2 months before and after June 1995. The estimated trend inflates actual patent growth in the field exposed to a patent term reduction because innovation accelerates at news. The opposite

patents in every year) and slightly different definitions of fields (USPC classes instead of 4-digit IPC classes).

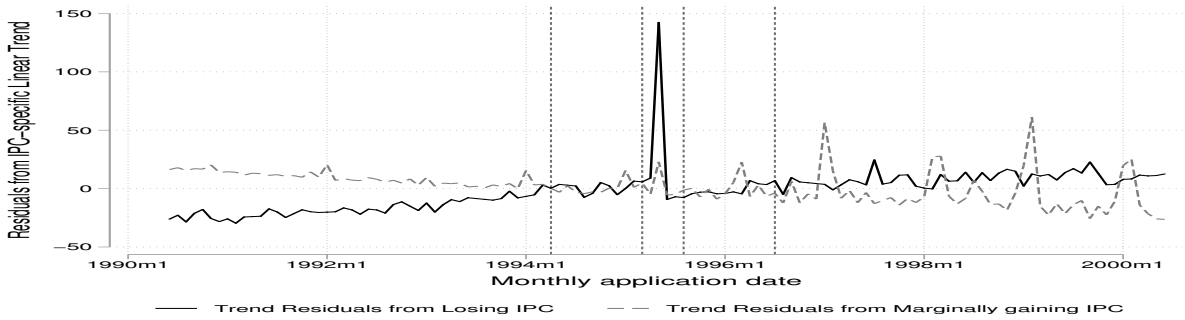
Figure B1: Descriptive comparison of two fields



(a) Patents in Losing vs. Gaining Field



(b) Fitted Trends – Narrow Window before Implementation



(c) Trend Deviations

Notes: Panel (a) shows the time series of granted patents applied for in month-t and classified in field C12R, the "losing field", and field A01H, the "gaining field". Panel (b) also plots estimated field-specific linear trends based on 12-month sample windows around policy implementation (excluding the two central months around June 1995), highlighted by vertical lines. Panel (c) plots the deviation of actual patents from estimated field-specific trends.

occurs for the field expected to increase its patent term.

This imposes a biased counterfactual 'trend" evolution of patenting in the two fields because estimated slopes are themselves endogenous to the policy, which significantly affects DiD estimates of (B1). By the Frisch-Waugh-Lowell theorem, β can be obtained from the residuals regression of deviation of patents from the field-specific time trends plotted in Figure B1b on the patent term change. Figure B1c represents these trend-deviations for the two fields C12R and A01H, which have clearly diverging patterns before the policy. This provides a graphical representation of why specification (B1) violates the parallel trends assumption when assuming no anticipation. After implementation, patenting in the field losing protection is below its inflated trend, while the opposite holds for the other field with a patent term extension. This accounts for the positive sign of Abrams (2009)'s DiD estimates.

Finally, in additional analyses available upon requests, I show that DiD estimates of specification (B1) are consistent with the empirical evidence of Section 5 if I extend the sample windows to June 1990-November 1994 and December 1995-June 2000 in order to more accurately represent pre-news patenting trends across fields.

Appendix C Additional empirical evidence by field

C.1 Extension of the analysis to 2010Q4

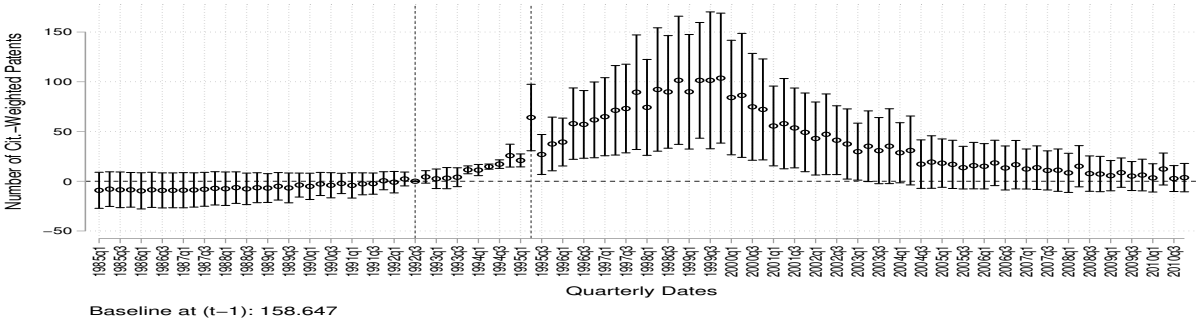
In this subsection, I provide suggestive evidence that the observed persistence of indirect post-implementation effects is temporary. To this end, I estimate DiD specification (5) on a sample extended to 2010Q4. Figure C1 illustrates that the overall post-implementation effect tends to mean-revert to zero at the end of the sample.

C.2 Evidence on patent quality

This subsection complements the discussion of Section 5.5.1 and presents empirical evidence on the effects of TRIPs patent term change on various measures of average patent quality.

Economic value. Figure C2 depicts the effect of a one-month *shorter* extension in average patent term on field-level patent value based on DiD estimates of specification (5). The dependent variable is computed as the sum of patents private economic value—from Kogan et al. (2017)—by patent application quarter and field. Results are consistent with those found for the number of patents and citations-weighted patents. The similar magnitudes to raw patent count indicate that the average economic value of patents responds just mildly

Figure C1: DiD estimates by field on extended sample to 2010

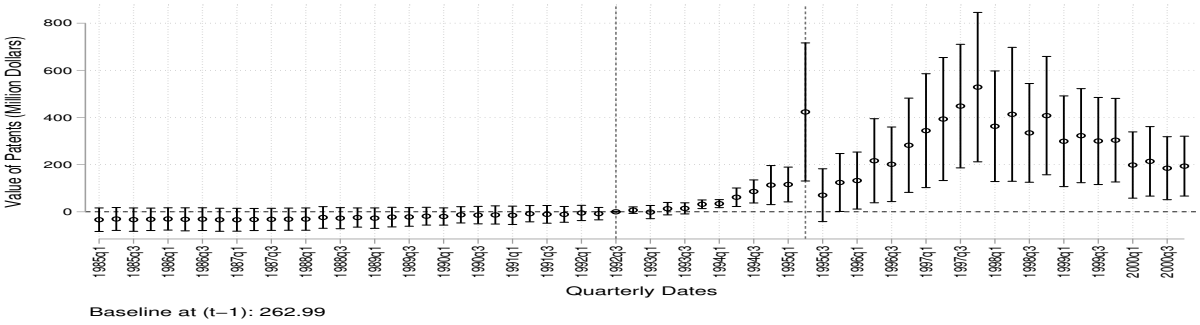


Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the level of citations-weighted (granted) patents by application quarter, based on estimates of DiD specification (5) with $X_j = \Delta T_j$. The sample is extended to 2010Q4. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

to the patent term change.

Moreover, in a separate analysis I show that the average renewal rate of patents up to the maximum patent term, an alternative proxy of patents economic value, does not vary in response to the patent term change. These results are available upon request.

Figure C2: DiD estimates for patents aggregate economic value by field



Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the level of patent value, based on estimates of DiD specification (5) with $X_j = \Delta T_j$. The dependent variable is computed as the field-specific sum of patents private economic value from Kogan et al. (2017) by application quarter. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

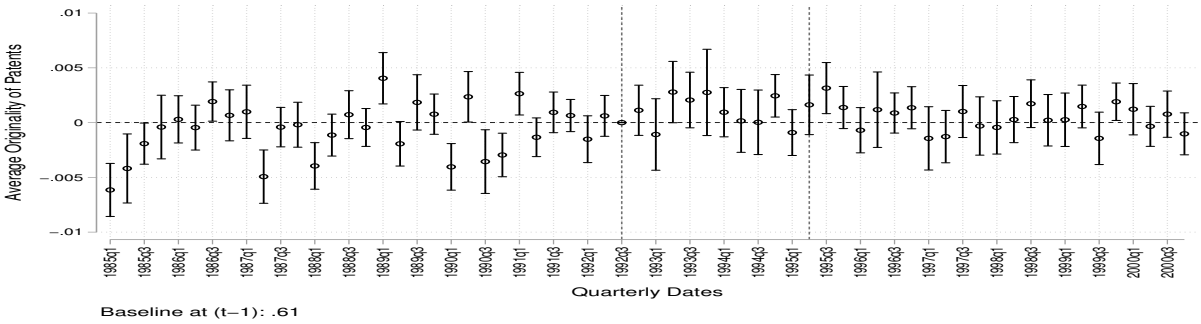
Originality and generality. Alternative metrics of patent quality are based on generality and originality (Hall, Jaffe and Trajtenberg, 2001), i.e., the heterogeneity of backward and forward citations in the technological space, respectively.[31]

[31]Following Hall, Jaffe and Trajtenberg (2001)'s definition, generality G_i and originality O_i for patent i are $G_i = 1 - \sum_{h=1}^{n} sf_{i,j}^2$ and $O_i = 1 - \sum_{h=1}^{n} sb_{i,j}^2$, respectively. $sf_{i,j}$ denotes the percentage of forward citations received by patent i from subsequent patents classified in field (patent class) h out of n. $sb_{i,j}$ denotes the percentage of backward citations made by patent i to previous patents classified in field (patent class) h out of n.

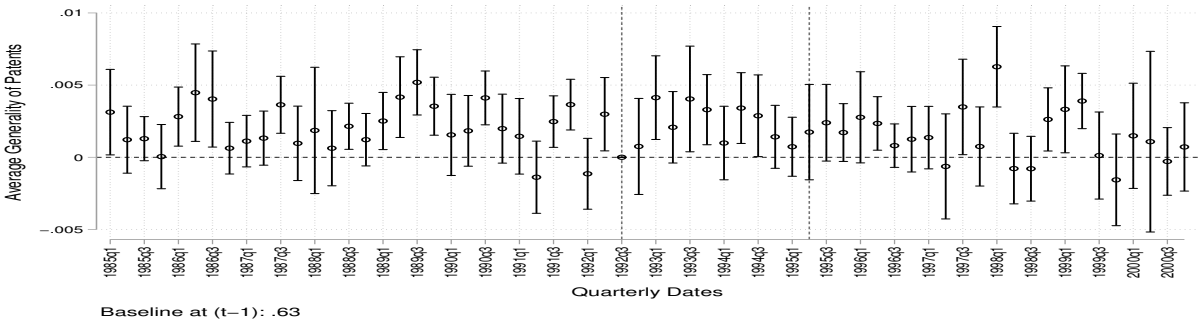
Figures C3b and C3a depict the effect of a one-month *shorter* extension in average patent term on field-level average originality and generality by patents application quarter, based on DiD estimates of specification (5). Results show that the patent term change has a negligible impact on these alternative measures of patent quality too.

Figure C3: DiD estimates for average generality and originality by field

(a) Originality



(b) Generality

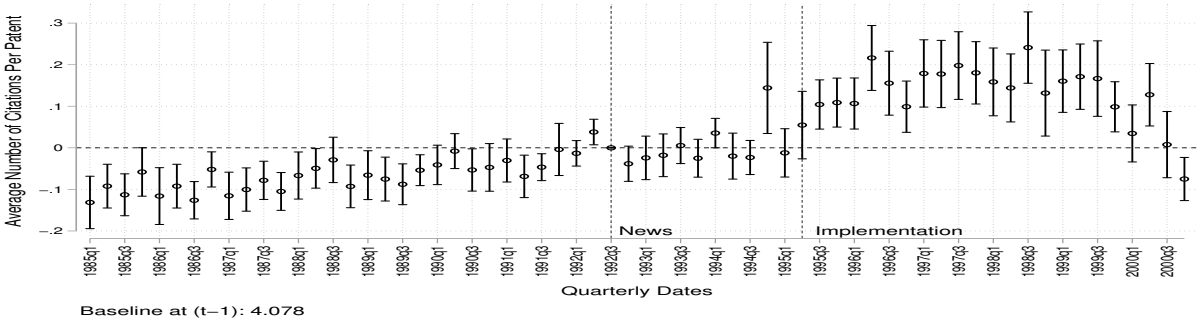


Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the level of average patent originality (panel a) and generality (panel b) by patents application quarter, based on estimates of DiD specification (5) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

Average scientific value. Figure C4 illustrates the effect of a one-month *shorter* extension in *average* patent term on field-level average 5-year forward citations per patent, by patents application quarter, based on DiD estimates of specification (5). The positive post-implementation effects highlight that average scientific quality increases in fields with a shorter extension in average patent term. Therefore, scientific quality co-moves with the "quantity" of innovation after implementation.

Text-based measures of novelty. Higham, de Rassenfosse and Jaffe (2021) explore various measures of patent quality, among which text-based proxies of novelty (see Kelly et al., 2021, for an example). Using text-based patent quality metrics by Arts, Hou and Gomez (2021), I analyze the effect of the patent term change on novelty based on (i) new words used

Figure C4: DiD estimates for average 5-year forward citations by field



Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the level of *average* 5-year citations-weighted (granted) patents by application quarter, based on estimates of DiD specification (5) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

in patents, (ii) new bigrams, (iii) new trigrams, (iv) new combinations of words, as well as cosine similarity with (iv) previous and (v) subsequent patents over 5-year horizons. For expositional convenience, I estimate a version of DiD specification (5) on an aggregate sample over four time-periods: pre-trends (1985Q1-1990Q4); pre-news (excluded, 1991Q1-1992Q3); news (1992Q4-1995Q2); and post-implementation (1995Q3-2000Q4). Across all regressions, the dependent variable is the field-specific average of each text-based novelty measure by patents application period.

Table C1 reports the effect of a one-month *shorter* extension in average patent term on text-based quality measures. The results do not highlight consistent patterns across measures. The average change in patent term does not correlate with variation in the use of new words, bigrams, trigrams, or combinations of words in the news period. In contrast, there is evidence of a mild negative effect of a shorter patent term extension on new words and bigrams after implementation. Moreover, average similarity with previous and subsequent patents increases with a shorter extension in ΔT_j both after news and implementation. This suggests that, on the one hand, novelty relative to previous knowledge declines, similarly to the effect based on the use of new words or bigrams, but, on the other hand, impact (based on forward similarity) increases. Overall, this suggests an ambiguous effect of patent term on novelty, as different measures move differently.

C.3 Evidence on patent scope

This subsection examines the effects of patent term on several proxies of patent scope by Marco, Sarnoff and deGrazia (2016). The analysis addresses the concern that observed changes in the number of patents or quality-adjusted patents may be driven by changes in

Table C1: Effect of Patent Term Change on Text-Based Quality

	Words	Bigrams	Trigrams	Word Comb.	Bcd. Cosine	Fwd. Cosine
$d_{pre} \times \Delta T_j \times (-30)$	0.0041	0.0201*	0.0180	2.3448	-0.0022	-0.0171***
	(0.0039)	(0.0117)	(0.0222)	(2.2022)	(0.0062)	(0.0064)
$d_{news} \times \Delta T_j \times (-30)$	-0.0043	-0.0114	0.0213	-0.6023	0.0142**	0.0130**
	(0.0035)	(0.0118)	(0.0217)	(1.4503)	(0.0057)	(0.0054)
$d_{implem} \times \Delta T_j \times (-30)$	-0.0116***	-0.0641***	-0.0191	-1.4396	0.0269***	0.0120***
	(0.0040)	(0.0130)	(0.0205)	(1.1955)	(0.0049)	(0.0046)
Constant	0.1609***	1.8258***	3.0808***	106.3141***	3.7125***	3.3756***
	(0.0369)	(0.1163)	(0.2361)	(14.8661)	(0.0552)	(0.0548)
Observations	2469	2469	2469	2469	2469	2469

Notes: The table reports the period-specific effects of a one-month shorter extension in average patent term on various measures of text-based patent quality based on DiD estimates of specification (5) estimated on an aggregate sample over 4 time-periods: pre-trends (1985Q1-1990Q4); pre-news (1991Q1-1992Q3, excluded); news (1992Q4-1995Q2); and post-implementation (1995Q3-2000Q4). Outcome variables are taken from Arts, Hou and Gomez (2021). Standard errors are clustered by 4-digit technical field. Statistical significance levels: * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

patenting strategies across fields. For example, the creation of patent thickets, i.e., nets of "overlapping" patents, may be used to extend the effective length of patent protection in fields with a relative shorter extension due to TRIPs.[32] However, this would result in narrower patent claims, which Marco, Sarnoff and deGrazia (2016) links to more detailed (longer) patent claims and fewer *independent* ones.

Therefore, I examine the effects of the anticipated patent term change on the average number of independent claims; the average share of independent claims on total claims; and claims average length in number of words or number of characters, by technical field and period of application filing. For expositional convenience, I consider a version of regression (5) estimated on a sample aggregated over 4 time-periods: pre-trends (1985Q1-1990Q4); pre-news (1991Q1-1992Q3, excluded); news (1992Q4-1995Q2); and post-implementation (1995Q3-2000Q4). Table C2 reports the period-specific effects of a one-month shorter extension in average patent term. Results indicate a negligible effect of patent term on all considered measures of patent scope both in the news and in the implementation phases.

C.4 Technological trends and alternative definition of fields

In this subsection, I present additional empirical analyses to address identification concerns related to differential trends in innovation that may confound DiD results of Section 5.1.

In the first analysis, I enrich DiD specification (5) with interactions of quarterly dummy

[32] This "evergreening" strategy is common in the pharmaceutical sector.

Table C2: Effect of Patent Term Change on Claims' Average Characteristics

	Indep. Claims	Sh. Indep. Claims	Avg. Nb. Words	Avg. Nb. Char.s
$d_{pre} \times \Delta T_j \times (-30)$	0.0203***	-0.1457	4.9613	28.6199
	(0.0070)	(0.1213)	(3.1920)	(19.3203)
$d_{news} \times \Delta T_j \times (-30)$	0.0082	0.0146	-0.6516	-4.7580
	(0.0090)	(0.0701)	(3.7195)	(23.3418)
$d_{implem} \times \Delta T_j \times (-30)$	-0.0002	0.0681	-2.5731	-15.8083
	(0.0069)	(0.1160)	(2.6505)	(16.6912)
Constant	2.7799***	25.0105***	906.4057***	5671.9151***
	(0.0748)	(0.6423)	(33.6988)	(207.7885)
Observations	2474	2474	2474	2474

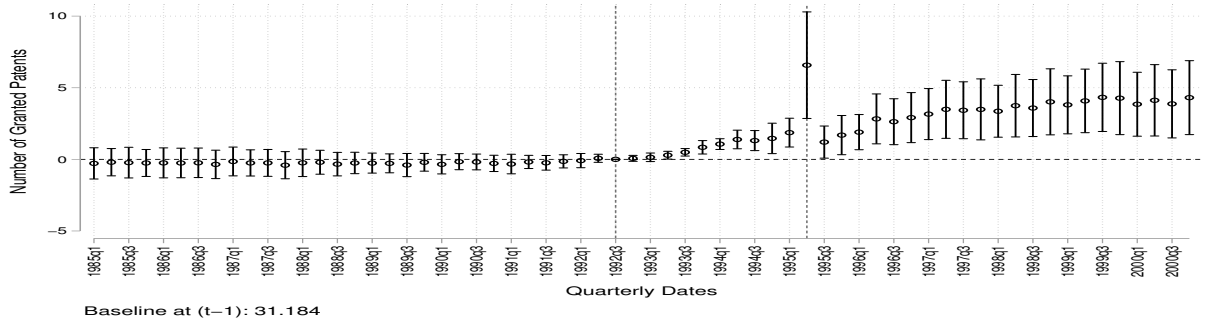
Notes: The table reports the period-specific effects of a one-month shorter extension in average patent term on several proxies of patent scope based on claims' characteristics, based on DiD estimates of specification (5) estimates on an aggregate sample over 4 time-periods: pre-trends (1985Q1-1990Q4); pre-news (1991Q1-1992Q3, excluded); news (1992Q4-1995Q2); and post-implementation (1995Q3-2000Q4). Outcome variables are taken from Marco, Sarnoff and deGrazia (2016). Standard errors are clustered by 4-digit technical field. Statistical significance levels: * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

variables with fixed effects by 3-digit IPC classes. These controls absorb rather granular variation in quarter-specific unobserved trends in innovation across more than 100 3-digit fields. Therefore, DiD effects are estimated by only exploiting variation in patent term change and patent outcomes across 4-digit fields within the same 3-digit technology.

Figure C5 depicts the relative effect of a one-month *shorter* extension in the average patent term on the number of (granted) patents by application quarter, based on DiD estimates of this richer specification. Results are fully consistent with Section 5.1. This suggests that quasi-experimental variation in average patent term is exogenous with respect to technological trends or macroeconomic shocks that may differentially affect innovation across more than 100 3-digit fields. Results available upon request document similar findings for an equivalent extension of Poisson DiD specification (6) and for citations-weighted patents as outcome.

In the second analysis, I use a more granular definition of field as cross-sectional unit of analysis, i.e., I define a field as an 8-digit rather than a 4-digit IPC class. This allows me to control for more than 600 4-digit IPC-specific technological trends through interactions of time fixed effects and 4-digit IPC fixed effects. Therefore, DiD estimates leverage variation in patenting and average patent term change across 8-digit classes within the same 4-digit. For expositional convenience, I aggregate the sample over 4 time-periods: pre-trends(1985Q1-1992Q2), pre-news (1992Q3, excluded), news (1992Q4-1995Q2), and post-implementation (1995Q3-2000Q4). Table C3 provides estimates of the period-specific effect of a one-month *shorter* extension in average patent term on several innovation outcomes and average innovation quality measures. Results are generally consistent with the findings of Section 5.1

Figure C5: DiD estimates by field controlling for 3-digit IPC-specific trends



Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the number of (granted) patents by application quarter, based on estimates of DiD specification (5) enriched by quarter-specific × 3-digit IPC specific fixed effects. These controls absorb granular variation in quarter-specific unobserved trends in innovation across more than 100 3-digit fields. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

and Appendix C.2, although post-implementation estimates are noisier.

In the third analysis, I control for differential innovation trends in technologically complex fields based on Galasso and Schankerman (2015)'s classification. To this end, I identify complex fields as A61B (Medical instruments); H03, H04, and H05 (Electronics); G06 and G16 (Computers and ICT); and A01H, A61K, C02F, C12M, C12N, C12P, C12Q, C12S, and G01N (Biotechnologies). I enrich DiD specification (5) with interactions of quarterly fixed effects with a dummy variable equal to one for complex fields. Figure C6 depicts the effect of a one-month shorter extension in average patent term on the number of granted patents by application quarter, based on estimates of this richer DiD specification. The results are consistent with the findings of Section 5.1.

C.5 Heterogeneous effects and selection on doses

In this subsection I address econometric concerns related to heterogeneous effects and the use of Two-Way Fixed Effects (TWFE) with a continuous treatment variable.

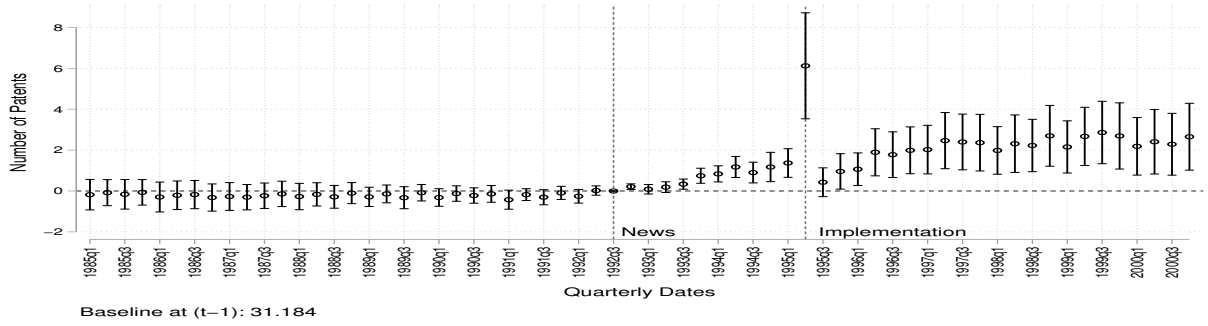
I build on Callaway, Goodman-Bacon and Sant'Anna (2021), who explore the relationship between average treatment effects, average causal responses, and the TWFE estimator in DiD settings where cross-sectional units receive a continuous treatment, henceforth "dose", denoted by d. Differently from my setting, however, they assume $d \geq 0$ for every unit and that a non-zero proportion of units are pure controls ($d = 0$). They derive an expression for the TWFE DiD estimator as a weighted average of average causal responses for the treated

Table C3: Linear Model – Effects of ΔT at the 8-digit IPC Class

	Patents	Citations	Value	Avg. Cit.s	Avg. Value
$d_{pre} \times \Delta T_j \times (-30)$	0.0066	0.0295	-0.1554	0.0105***	-0.0282
	(0.0065)	(0.1145)	(0.4854)	(0.0004)	(0.0289)
$d_{news} \times \Delta T_j \times (-30)$	0.0203**	0.2222***	1.2886**	0.0258***	0.0210**
	(0.0058)	(0.0330)	(0.3923)	(0.0044)	(0.0058)
$d_{implem} \times \Delta T_j \times (-30)$	0.0379*	0.6995	1.8838	0.0761***	0.0530
	(0.0142)	(0.3276)	(1.2356)	(0.0113)	(0.0924)
Time F.E:	Y	Y	Y	Y	Y
4-digit IPC F.E.	Y	Y	Y	Y	Y
Time × 4-digit IPC F.E.	Y	Y	Y	Y	Y
Observations	382016	382016	382016	244985	177949

Notes: The table reports the results of an alternative specification of DiD regression (5) taking 8-digit IPC fields as cross-sectional units of analysis, controlling for 4-digit field×time fixed effects, and using a sample aggregated over 4 time-periods pre-trends(1985Q1-1992Q2), pre-news (1992Q3, excluded), news (1992Q4-1995Q2), and post-implementation (1995Q3-2000Q4). Coefficients represent the period-specific effects of a one-month shorter extension in average patent term on several patent-based measures of innovation and innovation quality. Statistical significance levels: * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

Figure C6: DiD estimates by field controlling technological complexity



Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the number of (granted) patents by application quarter, based on estimates of DiD specification (5) enriched with interactions of quarter-specific effects with a dummy variable for complex fields according to Galasso and Schankerman (2015)'s classification. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

and average treatment for the treated, defined below, as:

$$\hat{\beta}^{twfe} = \int_{d_L}^{d_U} w_1(l)\Big[ACRT(l|l) + \frac{\partial ATT(l|h)}{\partial h}|_{h=l}\Big]dl + w_0 \frac{ATT(d_L|d_L)}{\partial d_L}$$

where $w_1(l)$ represents dose-specific weights on average causal response for the treated with dose l, defined as $ACRT(d|d) \equiv \frac{\partial \mathbb{E}[Y_t(l)|D=d]}{\partial l}|_{l=d}$ and $ATT(d|d) \equiv \mathbb{E}[Y_t(d) - Y_t(0)|D = d]$.

The TWFE estimator may be biased for the average causal response of interest—i.e., the effect on the outcome of a marginal change in the dose—for two reasons. First, if

there is selection in dose assignment, i.e., if units with different average treatment effects systematically receive different doses ($\frac{\partial ATT(l|h)}{\partial h}|_{h=l} \neq 0$). Second, because weights $w_1(l)$ mechanically assign more importance to observations with doses close to the mean of the treatment $\mathbb{E}[D]$ relative to the actual doses distribution.

I address these problems by extending Proposition 8 from Appendix B of Callaway, Goodman-Bacon and Sant'Anna (2021) to my setting with negative doses and no pure controls. The proposition provides an alternative decomposition of TWFE as a weighted average of all "Wald DiD" estimates, i.e., all the 2×2 difference-in-difference comparisons between units. The decomposition is

$$\hat{\beta}^{twfe} = \int_{\mathcal{D}} \int_{\mathcal{D}:h>l} w_{2\times 2}(l,h) \frac{m_\Delta(h) - m_\Delta(l)}{h-l} dh dl$$

where weights are $w_{2\times 2}(l,h) = (h-l)^2 f_D(h) f_D(l)/\mathbb{V}(D)$ and $m_\Delta(h) = Y_t(h) - Y_0(h)$ is the difference in outcome between t and pre-treatment $t' = 0$ for observations with dose h. This establishes a key role for 2×2 comparisons, which I leverage to investigate the role of selection of doses and weights.

I first analyze selection on doses. To this end, I examine the relationship between heterogeneity in treatment effects implied by individual Wald DiD comparisons and the magnitude of the average change in patent term. I illustrate the point by considering the data generating process $y_{i,t} = \alpha_i + \theta_t + \beta_{i,t}\Delta T_i + \varepsilon_{i,t}$, implicit to the TWFE estimator, where i identifies distinct cross-sectional units and t time. The outcome $y_{i,t}$ emerges from a time-invariant idiosyncratic component α_i, a common time-varying component θ_t, the individual effect $\beta_{i,t}$ of the dose ΔT_i assigned to i, and an error term $\varepsilon_{i,t}$. Hence, the within-field time difference $y_{i,t} - y_{i,0}$ is $\theta_t - \theta_0 + (\beta_{i,t} - \beta_{i,0})\Delta T_i$ and, assuming that the pre-treatment effect $\beta_{i,0} = 0$ for all units, the 2×2 comparison of units i and j is:

$$m_\Delta(\Delta T_i) - m_\Delta(\Delta T_j) = (y_{i,t} - y_{i,0}) - (y_{j,t} - y_{j,t}) = \beta_{i,t}\Delta T_i - \beta_{j,t}\Delta T_j + \Delta\varepsilon_{i,t} - \Delta\varepsilon_{j,t}$$

The selection problem arises if marginal effects correlate with the size of the dose, i.e., the change in patent term in this case. I propose focusing on pairs of observations with a change in patent term (dose) sufficiently close, where, even in the presence of selection, the β's can be assumed stable (on average) across observations. As a result, the value of the 2×2 building block of TWFE for doses close to some value ΔT_x becomes

$$\frac{m_\Delta(\Delta T_i) - m_\Delta(\Delta T_j)}{\Delta T_i - \Delta T_j} = \bar{\beta}_t(\Delta T_x)\frac{\Delta T_i - \Delta T_j}{\Delta T_i - \Delta T_j} + \frac{\Delta\varepsilon_{i,t} - \Delta\varepsilon_{j,t}}{\Delta T_i - \Delta T_j} \text{ for } (i,j) : (\Delta T_i, \Delta T_j) \in \mathcal{B}(\Delta T_x) \tag{C1}$$

Table C4: Magnitude of the Change in Patent Term and Average Local Causal Response

	News	Post-Implem.
ΔT	0.9454	2.5580
	(3.9469)	(13.5897)
ΔT^2	-0.0089	-0.0114
	(0.0188)	(0.0648)
ΔT^3	0.0000	0.0000
	(0.0000)	(0.0001)
ΔT^4	-0.0000	-0.0000
	(0.0000)	(0.0000)
Observations	19251	19251

Notes: The table reports results of a regression analysis that examines how the 2×2 DiD comparisons depend on the size of the average change in patent term, i.e., the treatment "dose". The first and second columns use as dependent variable the comparisons of patenting in the news and post-implementation periods, respectively, relative to pre-news period across pairs of fields with a change in patent term lower than 16.9 days. The regressors are the elements of a fourth-order polynomial in ΔT_i. Statistical significance levels: * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

where $\mathcal{B}(\Delta T_x)$ denotes a neighborhood of ΔT_x. Taking expectations across $(i,j) : (\Delta T_i, \Delta T_j) \in \mathcal{B}(\Delta T_x)$, the error differences disappear and (C1) is equal to $\bar{\beta}_t(\Delta T_x)$. Therefore, these restricted comparisons allow me to study the heterogeneity of local marginal effects $\bar{\beta}_t(\Delta T_x)$ depending on the size of the dose.

To empirically implement the proposed strategy, I compute all the 2×2 comparisons across fields aggregating patent counts for three specific periods: pre-news (1990Q1-1992Q3), news (1992Q4-1995Q5), and post-implementation (1995Q3-2000Q4). I focus on pairs of fields with an average change in patent term that differs by less than 16.9 days, which corresponds to the 10^{th} percentile of the difference distribution.[33] For each dose value ΔT_x, I calculate the average effect $\bar{\beta}_t(\Delta T_x)$ as per equation (C1) and then perform parametric and non-parametric tests to analyze how $\bar{\beta}_t(\Delta T_x)$ varies with ΔT_x. Table C4 shows the results of the parametric analysis based on a fourth-order polynomial in ΔT_i. The results suggest the absence of selection on doses, i.e., the effect of changes in average patent term on patenting is uniform across doses. An alternative analysis—whose results are available upon request—tests the same hypothesis non-parametrically using a set of dummy variables for each decile of ΔT_i and finds similar results.

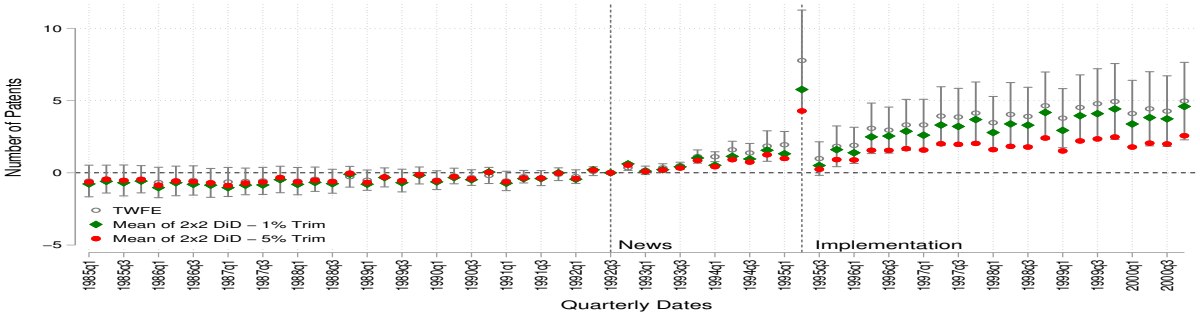
Next, I examine concerns related to the implicit weighting scheme of TWFE, which attributes more weight to observations with a dose close to its average compared to the empirical distribution of doses. Therefore, I re-compute estimates of aggregate DiD effects by averaging Wald DiD comparisons with equal weights—thus leveraging the empirical distribution of doses—and compare the new DiD effects to the ones based on TWFE. The average

[33] This defines $(i,j) : (\Delta T_i, \Delta T_j) \in \mathcal{B}(\Delta T_x)$ in previous paragraph.

of Wald DiD estimates exceeds in magnitude TWFE because the distribution of Wald DiD comparisons is skewed to the left, displaying some negative outliers. With equal weights, these outliers are given more relevance than with the TWFE weighting scheme. Specifically, the effect of a one-month shorter extension in average patent term is equal to +0.93 patents in the news period according to the unweighted average of 2×2 DiD comparisons but +0.64 patents according to TWFE estimates. Similarly, the average effect is +3.64 patents after implementation based on the unweighted average of 2×2 DiD comparisons and +3.24 according to TWFE estimates.

Besides selection, outliers in 2×2 comparisons may impact TWFE estimates potentially distorting its representativeness for the average field. Therefore, I re-compute causal effects based on equal-weights average of 2×2 DiD comparisons after trimming the upper and lower tails of their distribution. Trimming at the 1% level yields a news effect estimate of +0.47 patents and post-implementation effect of +3.11. Therefore, the effects are very close to TWFE estimates. Figure C7 provides a visual comparison of estimated effects of a one-month shorter extension in average patent term based on TWFE estimates and based on averaging individual 2×2 differences for fields with $\Delta T < 16.9$ days and trimming the distribution at 1% or 5% levels.

Figure C7: DiD TWFE estimates and average causal responses



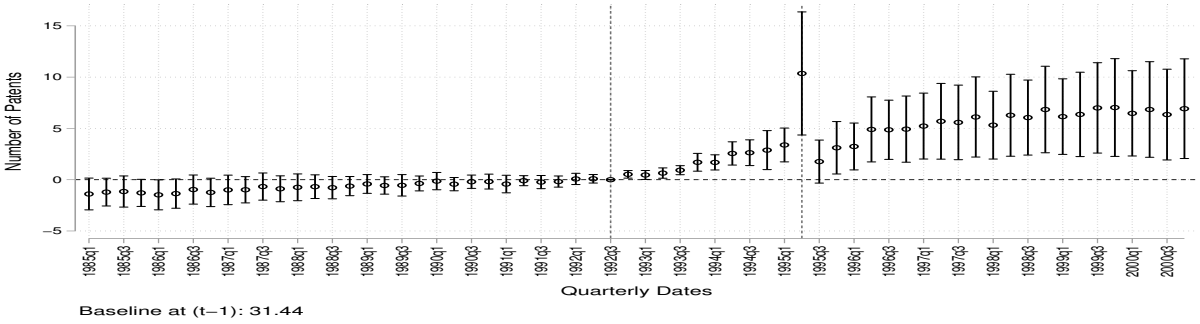
Notes: The figure displays the effect of a one-month *shorter* extension in average patent term on the number of (granted) applications by application quarters. Gray dots represent estimates based on Two-Way Fixed Effects (TWFE) estimation pf DiD specification (5). Red and green dots represent average quarterly causal effects obtained from averaging individual 2×2 differences for fields with a difference in $\Delta T < 16.9$ days and trimming the double differences distribution at the 1% or 5% levels, respectively. Bands depict 95% confidence bands of TWFE estimates, with standard errors clustered by technical field and treatment phase.

C.6 Heterogeneous effects based on observable dimensions

C.6.1 Heterogeneous effects by patent renewal rates

I analyze the heterogeneity of DiD estimates by expected responsiveness of the field to the maximum patent term. I proxy it by the pre-news share of patents for which maintenance

Figure C8: Triple difference estimates by renewal rate up to maximum term



Notes: The figure displays the triple-difference effect of a one-month *shorter* extension in average patent term on the number of patents (by application quarter) for fields more sensitive to maximum patent term ($d_{Ren_j>p50} = 1$) relative to less sensitive ones ($d_{Ren_j>p50} = 0$). Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

fees are paid in the 4^{th}, 7^{th}, and 11^{th} years after grant, thus extending patent life up to the maximum term. I estimate the heterogeneous effects of the change in average patent term based on triple-difference specification (9). I replace the interaction variable with a dummy $d_{Ren_j>p50}$ which takes value one for field-j if the share of patents renewed up to the maximum term exceeds the median (0.27) of the cross-field distribution.

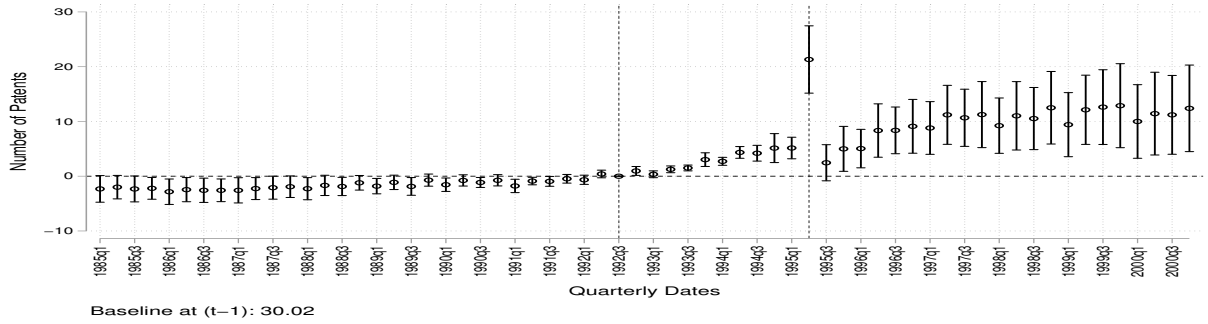
Figure C8 depicts the effect of a one-month *shorter* extension in average patent term on the number of patents (by application quarter) for fields more sensitive to maximum patent term ($d_{Ren_j>p50} = 1$) relative to less sensitive ones ($d_{Ren_j>p50} = 0$). As expected, the magnitude of the effect of the patent term change is larger in the former fields than in the latter.

C.6.2 Heterogeneous effects by dispersion around average patent term

Next, I analyze the heterogeneity of DiD estimates by the degree of dispersion of the patent-level pending period around its average by field. To this end, I compute the standard deviation of the patent-level pending period distribution by field before policy news. I then estimate the analogue of triple-difference specification (9) where I replace the interaction variable with a dummy $d_{\sigma_j \leq \sigma^m}$, which takes value 1 if the field-specific dispersion σ_j is smaller than the median across fields σ^m.

Figure C9 depicts the effect of a one-month *shorter* extension in average patent term on the number of patents (by application quarter) for fields with smaller dispersion of patent-specific pending period around the average ($d_{\sigma_j \leq \sigma^m} = 1$), i.e., with more "precise" average change in patent term. The magnitude of the effect is larger in fields with smaller dispersion.

Figure C9: Triple difference estimates by within-field dispersion of pending period



Notes: The figure displays the triple-difference effect of a one-month *shorter* extension in average patent term on the number of patents (by application quarter) for fields with a smaller dispersion in patent-specific pending period around its average by field ($d_{\sigma_j \leq \sigma^m} = 1$) relative to fields with a larger one ($d_{\sigma_j \leq \sigma^m} = 0$). Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

C.7 Endogenous changes in examination

In this subsection I present empirical analyses to complement the discussion of Section 5.5 on the role of endogenous changes in examination. I first document that the average patent term change—based on the average pending period before policy news—is not related to systematic variation in quarter- and field-specific average pending period (examination time) after news and implementation. Therefore, the policy did not impact average examination time, an important aspect of patent examination intensity. To this end, I estimate a version of DiD specification (5) on a sample aggregated into four time-periods—pre-trends (1985Q1-1989Q4), pre-news (1990Q1-1992Q3, omitted), news (1992Q4-1995Q2), and post-implementation (1995Q3-2000Q4)—with dependent variable being the period- and field-specific average pending period. The first column of Table C5 reports the period-specific effect of a one-month *shorter* extension in average patent term. The effect is null after both news and implementation.

Next, I examine why the quarter-specific average pending period does not systematically vary with ΔT_j across fields despite the surge in patent applications in fields with shorter patent term extensions. Using disambiguated examiner identifiers from PatentsView (USPTO, 2023), I analyze examiner reallocation across technical fields. I compute the field-specific examiner count over the four periods and estimate an aggregated DiD specification (5) with examiner headcount as the dependent variable. Column 2 of Table C5 shows that a one-month shorter patent term corresponds to a relative *increase* in examiners, both in the news phase and post-implementation. This reallocation offsets the impact of increased congestion on pending periods.

Finally, I verify that DiD estimates of specification (5) are fully robust to using as the

Table C5: Patent Term Change, Pending Period, and Examiners Reallocation

	Avg. Pending Period (Days)	Number of Examiners
$d_{pre} \times \Delta T_j \times (-30)$	-0.3397	-0.3459
	(1.9661)	(0.3726)
$d_{news} \times \Delta T_j \times (-30)$	1.2797	2.0492***
	(1.8581)	(0.3518)
$d_{implem} \times \Delta T_j \times (-30)$	0.2078	7.0015***
	(2.5557)	(1.7339)
Observations	2472	2472

Notes: The table reports the period-specific effects of a one-month shorter extension in average patent term on average pending period (column 1) and field-specific examiners headcount (column 2) based on DiD estimates of specification (5) estimated on an aggregate sample over 4 time-periods: pre-trends (1985Q1-1989Q4); pre-news (1990Q1-1992Q3, excluded); news (1992Q4-1995Q2); and post-implementation (1995Q3-2000Q4). Standard errors are clustered by 4-digit technical field. Statistical significance levels: * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

main regressor of interest the average change in patent term based on the actual quarter- and field-specific average pending period ($\Delta T_{j,k}$) rather than its counterpart based on the field-specific average pending period *before* policy news (ΔT_j). Results are available upon request.

C.8 Changes in trade regulation concomitant to TRIPs

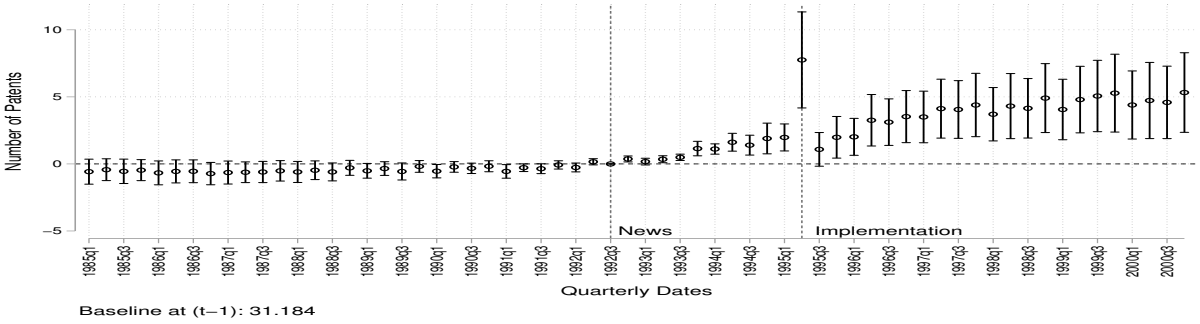
This subsection addresses concerns related to confounding effects of changes in tariffs and trade regulations in the context of the Uruguay Round of Agreements concomitant to TRIPs.

First, I show that controlling for the quarter-specific effect of changes in tariffs in DiD specification (5) does not alter the estimated effects of patent term on innovation (Figure 3). To this end, I build field-specific measures of tariff intensity based on WTO data on tariffs profiles. Specifically, I compute tariff intensity across 4-digit IPC technical fields as the weighted average of the share of HS-2002 product codes (linked to a specific technical field) with a reported import-tariff above 5% on WTO members.[34] I compute import-tariff intensity for the US, EU, and China and then jointly control for either their pre-TRIPs levels or their post-TRIPs 1996-2001 changes.

Figure C10 displays the effect of a one-month shorter extension in average patent term on the number of granted patents with this richer specification. The effects are fully consistent with the ones of Section 5. In a separate analysis I verify that the change in average patent term determined by TRIPs is uncorrelated with concomitant tariff changes across technical fields. Therefore, the latter do not act as confounders in the baseline DiD specification (5).

[34]Data were downloaded at the link http://tao.wto.org/ExportReport.aspx on 21/01/2022

Figure C10: DiD estimates by field controlling for Uruguay Round tariff changes



Notes: The figure depicts the effect of a one-month *shorter* extension in average patent term on the number of granted patents (by application quarter) once controlling in specification (5) for quarterly effects of the 2001-1996 change in import-tariff intensity for the US, Europe, and China. Bands represent 95% confidence intervals based on clustered standard errors by field and treatment phase. The two vertical lines mark news and implementation dates, respectively.

C.9 The role of developing countries in light of TRIPs

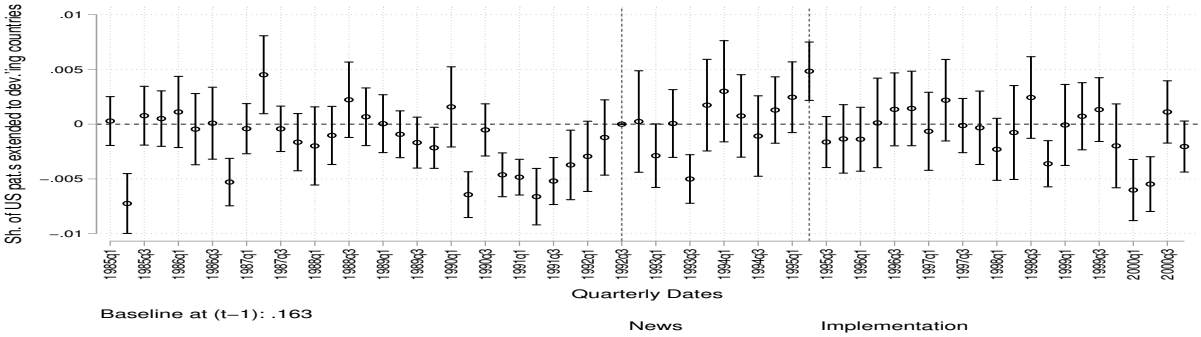
This subsection complements the discussion of Section 5.5 about the effect of TRIPs implementation in Low and Middle Income Countries (LMICs) on US innovation across fields. TRIPs strengthened intellectual property protection in LMICs, improving the access of US innovators to these markets. However, significant delays in TRIPs implementation in LMICs enhanced uncertainty. If these factors correlate with the average change in US patent term across fields, DiD estimates based on specification (5) would be biased.

To address this concern, I analyze the effect of US patent term on the propensity of US firms to also seek protection in LMICs. If observed variation in US innovation stems from changes in the access of US innovators to LMICs—rather than from changes in US patent term—the intensity with which they seek to extend their intellectual property rights in foreign markets after filing a US patent application should correlate with the patent term change. Leveraging Kyle and McGahan (2012)'s list of countries that significantly strengthened patent protection to join the WTO (e.g., China, India, Brazil, and South Korea), I compute the share of US patents—by field and application quarter—with at least one additional patent application to one of these destinations in the same patent family. I denote this variable as $S_{j,t}^{LMICs}$.

Figure C11 depicts the effect of a one-month shorter extension in average patent term on $S_{j,t}^{LMICs}$ based on estimates of baseline DiD specification (5). Results suggest a null effect of the patent term change on the propensity to seek patent protection in LMICs affected by TRIPs, both at news and after implementation. Similar findings—available upon request—hold when (i) focusing on individual countries, e.g., China, India, South Korea, and Brazil; (ii) when using a Poisson rather than a linear specification; (iii) using the share of citations-

weighted patents rather than the share of patents.

Figure C11: DiD estimates for the share of US patents extending protection to LMICs



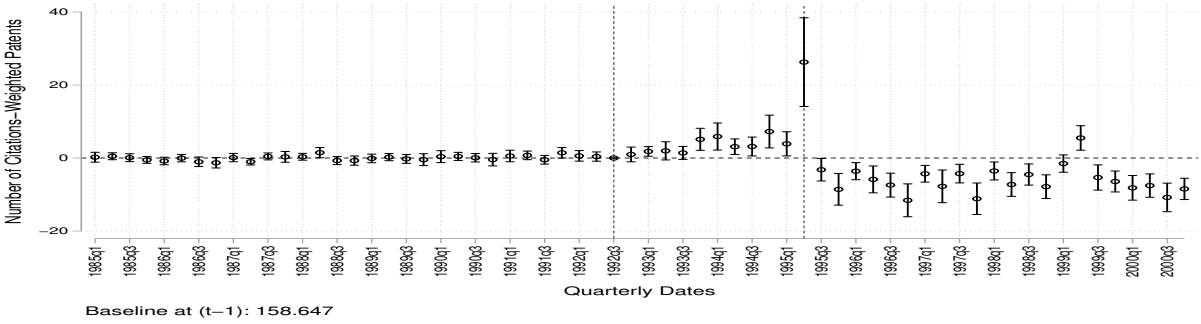
Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average US patent term on the share of US patents also seeking protection in LMICs increasing patent protection due to TRIPs. Effects are based on estimates of DiD specification (5) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

C.10 Additional evidence on direct effects by field

C.10.1 Direct effects for citations-weighted patents

Figure C12 illustrates direct effects of a one-month *shorter* extension in average patent term on citations-weighted patents, based on estimates of specification (7). Results are consistent with those for patents.

Figure C12: Direct effects of patent term by field – Citations



Notes: The figure depicts the relative *direct* effect of a one-month *shorter* extension in the average patent term on the level of citations-weighted (granted) patents by application quarter, based on estimates of DiD specification (7) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

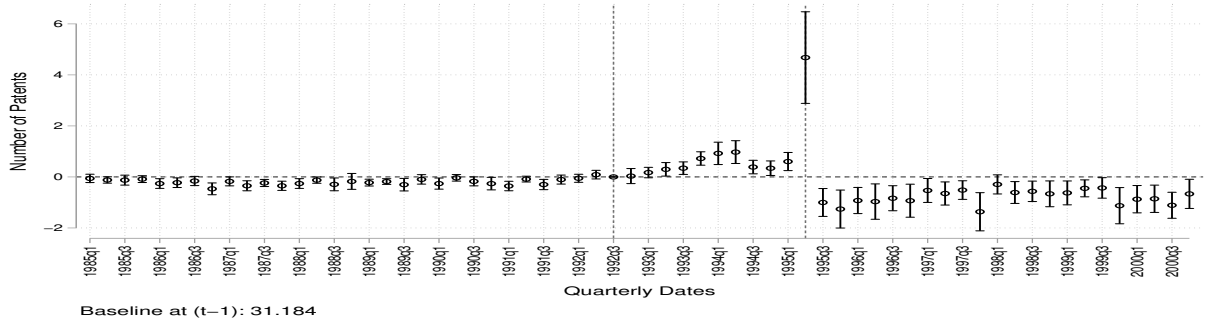
C.10.2 Influence of other innovation in other fields

Next, I extend the analysis of direct policy effects (Section 5.3) to allow for indirect news effects to influence post-implementation innovation *across* fields.

I enrich specification (7) to also control for the quarter-specific effects of past innovation in fields other than j—which I denote by $\overline{Y}_{-j,k-\mathcal{A}-1:k-1}$ and define as the weighted average of lagged innovation in other fields h with weights being the technological proximity $\rho_{h,j}$ between field j and h, i.e., $\overline{Y}_{-j,k-\mathcal{A}-1:k-1} = \sum_{h \neq j} \rho_{h,j} \overline{Y}_{h,k-\mathcal{A}-1:k-1}$. Following Jaffe (1986), technological proximity is $\rho_{h,j} = \mathbf{f}_h \mathbf{f}_j' \Big/ \sqrt{(\mathbf{f}_h \mathbf{f}_h')(\mathbf{f}_j \mathbf{f}_j')}$, where $\mathbf{f}_h$ ($\mathbf{f}_j$) is the vector that collects the field-specific number of backward citations from field h (j) patents. The greater the overlap between the distribution of citations from fields h and j, the closer $\rho_{i,j}$ is to one.

Figure C13 illustrates the *direct* effects of a one-month *shorter* extension in average patent term on the number of patents (by application quarter) once controlling for both within- and cross-field indirect news effects. The sign and magnitude of the estimates are almost identical to the ones of specification (7) (Figure 7), which controls for just within-field innovation histories. Therefore, indirect effects from innovation histories in other fields are economically negligible.[35]

Figure C13: Direct effects controlling for cross-field technological dependence



Notes: The figure depicts the relative *direct* effect of a one-month *shorter* extension in the average patent term on the number of (granted) patents by application quarter, based on estimates of DiD specification (7) enriched to control for the indirect effect of cross-field innovation histories. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

[35] For $k = 2000Q3$, a one-standard deviation increase in $\overline{Y}_{j,k-\mathcal{A}-1:k-1}$ is associated to +154.9 patents in quarter-k and field-j. A one-standard deviation increase in $\overline{Y}_{-j,k-\mathcal{A}-1:k-1}$ is associated to just +1.9 patents.

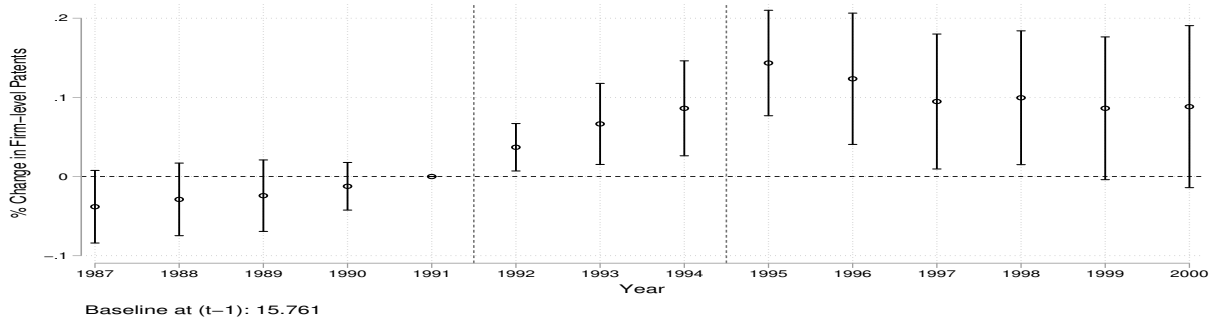
Appendix D Additional analyses at the firm-level

D.1 Additional firm-level outcomes for COMPUSTAT sample

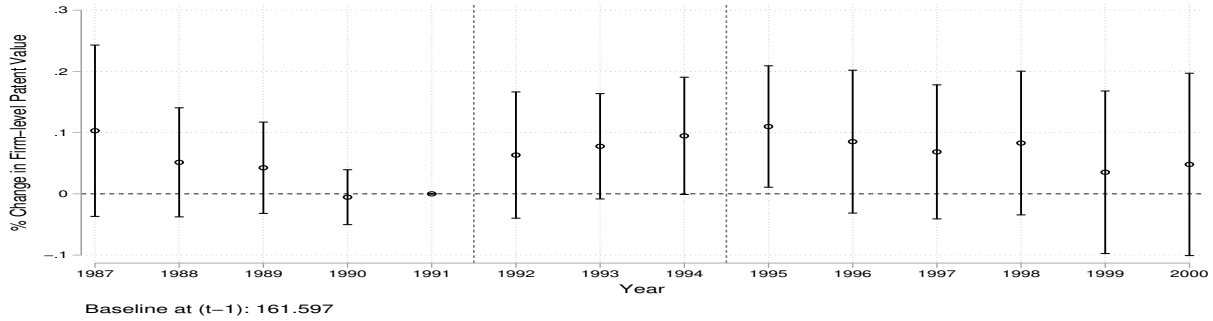
Alternative innovation outcomes. Figures D1a and D1b represent the effect of a one-month shorter extension in average patent term on log-deviations in firm-level granted patents and patent value from Kogan et al. (2017) (by year of application), respectively. I compute these effects using pseudo-MLE estimates of Poisson specification (8). Results are fully consistent with the ones reported Section 5.2 for R&D and citations-weighted patents.

Figure D1: Poisson DiD estimates for alternative firm-level innovation measures

(a) Number of Granted Patents by Firm and Application Year



(b) Patent Private Economic Value by Firm and Application Year

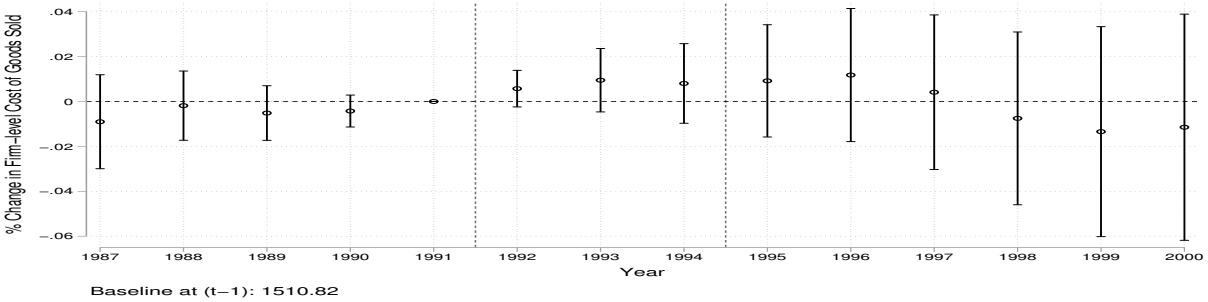


Notes: The figure depicts the effect of a one-month *shorter* extension in average patent term at the firm level—based on the Poisson DiD specification (8)—on number of patents (panel a) and patent value (panel b) by year of application. Point estimates represent the effect of the policy change on *percent deviations* of the outcome variable from its baseline value in 1991. Standard errors are clustered at the firm level and bands represent 95% confidence intervals. The first and second vertical lines mark news and implementation events, respectively.

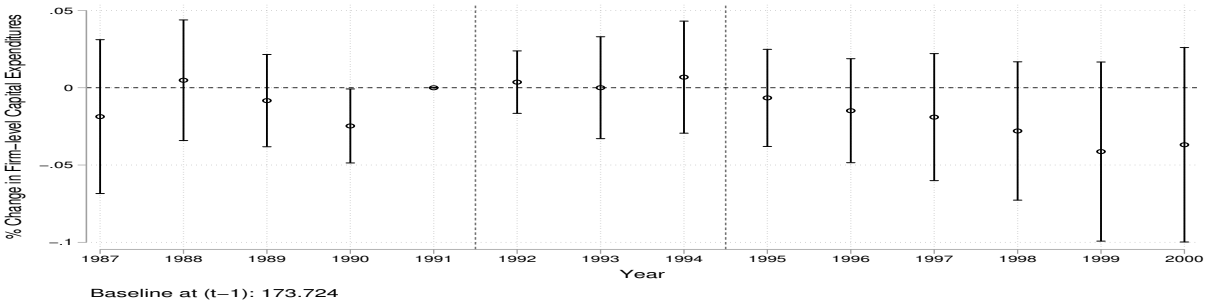
Placebo variables: Costs, Capital Expenditures, and Sales. Figures D2a, D2b, and D2c illustrate the results of three placebo analyses of the effect the change in patent term change on firm-level variables costs (`cogs` in COMPUSTAT), capital expenditures (`capx`), and sales (`sale`), respectively. In principle, the patent term change should have no (immediate) effect on previous variables, as the figures confirm.

Figure D2: Poisson DiD estimates for placebo firm-level outcomes

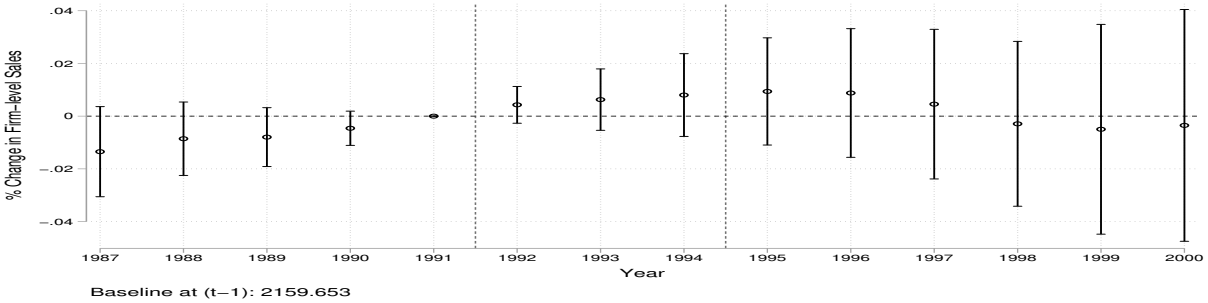
(a) Variable Costs



(b) Capital Expenditures



(c) Sales



Notes: The figure depicts the effect of a one-month *shorter* extension in average patent term at the firm level—based on the Poisson DiD specification (8)—on yearly cost of goods sold (panel a), capital expenditures (panel b), and sales (panel c). Point estimates represent the effect of the policy change on *percent deviations* of the outcome variable from its baseline value in 1991. Standard errors are clustered at the firm level and bands represent 95% confidence intervals. The first and second vertical lines mark news and implementation events, respectively.

D.2 Within-firm reallocation of R&D and innovation

This subsection provides additional evidence on the reallocation of R&D effort and innovation *within* firms *across* technological fields. The analysis employs PatentsView data (USPTO, 2023) on the universe of disambiguated assignees and inventors. I examine the effect of the change in the average field-specific patent term on innovation and the labor input to R&D—i.e., inventors—across fields within firms. Table D1 reports the results of the linear

Table D1: Linear Model - Patent Term Change and Firm × Field Outcomes

	Patents	Citations	Avg. Citations	Inventors	Sh. Inventors
$d_{pre} \times \Delta T_j \times (-30)$	-0.0016	-0.1070	-0.0106	0.0180	-0.0009
	(0.0091)	(0.0754)	(0.0203)	(0.0438)	(0.0008)
$d_{news} \times \Delta T_j \times (-30)$	0.0590***	0.4961***	0.0431**	0.0928***	0.0020***
	(0.0118)	(0.1063)	(0.0201)	(0.0177)	(0.0006)
$d_{implem} \times \Delta T_j \times (-30)$	0.2622***	3.8086***	0.0354	0.5312***	0.0019**
	(0.0699)	(1.2600)	(0.0410)	(0.1237)	(0.0009)
Assignee F.E.	Y	Y	Y	Y	Y
4-Digit IPC F.E.	Y	Y	Y	Y	Y
Assignee×4-Digit IPC F.E.	Y	Y	Y	Y	Y
Period F.E.	Y	Y	Y	Y	Y
Period×3-Digit IPC F.E.	Y	Y	Y	Y	Y
Observations	532368	532368	131892	532368	532368

Notes: The table reports the effect of a one-month shorter extension in average patent term on several measures of innovation and R&D effort *across* fields *within* firm, based on specification (D1). Parentheses report standard errors clustered by technical field (4-digit IPC). Statistical significance levels: * ($p < 0.05$), ** ($p < 0.01$), *** ($p < 0.001$)

specification

$$Y_{f,j,p} = \alpha_f + \gamma_j + \psi_{f,j} + \theta_p + \omega_{h \in j,p} \sum_p d_p \Delta T_j + \varepsilon_{f,j,p} \tag{D1}$$

where $Y_{f,j,p}$ is the outcome of interest for firm f in 4-digit field j and period p, where p is either 1985Q1-1989Q4; 1990Q1-1992Q3; 1992Q4-1995Q2; 1995Q3-2000Q4. α_f, γ_j, $\psi_{f,j}$, θ_p, and $\omega_{h \in j,p}$ are firm, 4-digit field, firm×4-digit field, period, and 3-digit field × period fixed effects, respectively.

Within firms, innovation—measured by patents and citations-weighted patents in columns 1 and 2 of Table D1—gets reallocated to fields with a *shorter* extension in average patent term, consistent with field-level evidence. Moreover, average patent quality across fields is mildly affected. (column 3). Finally, a shorter extension in average patent term is associated with a positive reallocation of R&D effort, proxied by the number and the share of inventors (columns 4 and 5, respectively) patenting in different fields for the same firm.

Furthermore, in results available upon request, I provide consistent evidence on the reallocation of innovative activity across fields at the inventor level within a given firm (assignee).

Appendix E Additional analyses at the industry level

In this section, I link changes in patenting induced by TRIPs patent term change to variation in welfare and Total Factor Productivity (TFP) at the sectoral level. To this end, I leverage the NBER CES manufacturing database (Becker, Gray and Marvakov, 2021) at the 6-digit

NAICS level. Data include estimates of TFP, inverse welfare measures based on producers' price indexes (value of shipments deflator), and other industry aggregates.

The analysis proceeds in two steps. First, I isolate variation in sectoral patent-based measures of innovation originating from the TRIPs-induced change in patent term. I estimate a sectoral version of specification (5) with 6-digit NAICS sectors as cross-sectional units and either the number of patents, citations-weighted patents, or patent value as outcome variables.[36] I control for the natural log of energy and material prices and 4-digit NAICS×year fixed effects.

In the second step, I use these fitted values from these regressions as excluded instruments for sector-s, year-t innovation $I_{s,t}$ in the second-stage panel regression:

$$y_{s,t} = \alpha_s + \sum_{k=1985}^{2000} \gamma_k \mathbf{1}_{(t=k)} + \beta I_{s,t} + \Xi X_{s,t} + \varepsilon_{s,t} \tag{E1}$$

which links sectoral outcomes $y_{s,t}$—either the natural logarithm of TFP or welfare—to patent-based measures of innovation. The specification controls for sectoral fixed effects α_s, yearly fixed effects $\mathbf{1}_{(t=k)}$, and the vector $X_{s,t}$ including the natural log of energy and material prices and 4-digit NAICS×year fixed effects.[37]

Table E1 reports 2SLS estimates of β in regression (E1) for different innovation measures (on the rows) and for TFP and price of shipments deflator (on the columns). An increase of 100 industry×year patents relates to a 3.3% rise in yearly TFP and 2.7% lower prices. This is approximately equivalent to a 0.1 elasticity of TFP to the raw patent count.[38] Similarly, an increase of 1,000 industry×year citations-weighted patents implies a 1.8% rise in productivity and 1.5% decline in prices. The first stage regressions have F-statistics that are always above 10.

The ratio of the effect of patent-based innovation measures on TFP and producers prices informs the average pass-through of productivity gains to higher consumer welfare. Depending on the innovation measure considered, the pass-through ranges between 76% and 84%.

[36] I compute sectoral innovation measures starting from field-level innovation measures and aggregating them using the probabilistic crosswalk by Goldschlag, Lybbert and Zolas (2019). Similarly, I compute the sector-specific change in average patent term ΔT_s based on field-specific changes ΔT_j as $\Delta T_s = \sum_j \Delta T_j \pi_{j|s}$, $\pi_{j|s}$ is the probability that a patent is classified in technical field j conditional on being linked to industry s.

[37] I estimate (E1) with by weighted 2SLS, with weights representing the pre-news size of the industry in terms of aggregate innovation. Specifically, I weigh sectors by the number of granted patents applied for in 1985.

[38] The sectoral average of industry×year patents in the pre-treatment year 1991 is 280. The elasticity is obtained with a back-of-the-envelope calculation as 0.033/(100/280)

Table E1: Sectoral Evidence on the Link between Patenting, TFP, and Welfare

	(1) TFP	(2) Price	(3) TFP	(4) Price	(5) TFP	(6) Price
Patents/100	0.033***	-0.027***				
	(0.008)	(0.007)				
Citations/1000			0.018***	-0.015***		
			(0.003)	(0.003)		
Patent value ($M)/1000					0.009**	-0.007**
					(0.004)	(0.003)
6-d NAICS f.e.	Y	Y	Y	Y	Y	Y
Year f.e.	Y	Y	Y	Y	Y	Y
3-d NAICS × Year f.e.	Y	Y	Y	Y	Y	Y
Observations	6684	6684	6684	6684	6684	6684

Notes: Columns (1), (3), and (5) report the OLS estimates of the β coefficient of specification (E1) having as dependent variable the natural logarithm of the 5-factors TFP, normalized to 100 in 1997. Columns (2), (4), and (6) report the 2-stage estimates of the IV regression. Standard errors are clustered by 3-digit NAICS × year. Statistical significance levels: $^{*}(p < 0.05)$,$^{**}(p < 0.01)$,$^{***}(p < 0.001)$

Appendix F Additional analyses on the mechanism

F.1 Elasticity of current innovation to past innovation

In this subsection, I provide a synthetic measure of the magnitude of the technology disclosure externality discussed in Section 6. To do so, I estimate a version of specification (7) on a field-level panel sample aggregated over four periods: (0) 1985Q1-1988Q4 (control period); (1) 1989Q1-1992Q3 (pre-news period); (2) 1992Q4-1995Q2 (post-news, pre-implementation period); and (3) 1995Q3-2000Q4 (post-implementation period). This aggregation is done to precisely distinguish innovation during the news and post-implementation periods. Table F1 reports the effect of a one-month shorter extension in average patent term and of past innovation based on the baseline DiD specification (5) (column 1) and on the augmented DiD specification (7) estimated by OLS (column 2) or 2SLS using fitted values from (5) as instruments for lagged innovation (column 3).

Estimates of the coefficients on the terms $P_{j,t-1}$ and $d_{implem} \times P_{j,t-1}$ are informative about the elasticity of current innovation to past innovation. Their point estimates imply that an increase of past (news) innovation by 1 patent leads to a current (post-implementation) innovation increase of approximately 5.1 patents. As the average number of patents in the news period is 408, +1 patent is equivalent to a +0.245 percent increase. As the average number of patents in the post-implementation period (1000), +5.1 patents is equivalent to a +0.51 percent increase. Therefore, the implied elasticity is approximately 2.1.

Table F1: Decomposition of the post-implementation effect

	(1) DiD	(2) Augmented DiD	(3) Augmented DiD – IV
$d_{pre} \times \Delta T_j \times$ (-30)	-3.09826		
	(2.22797)		
$d_{news} \times \Delta T_j \times$ (-30)	17.47662***	17.86102***	17.90820***
	(4.18025)	(4.11535)	(4.05163)
$d_{implem} \times \Delta T_j \times$ (-30)	93.94929***	-16.22455	-26.30657*
	(33.94374)	(12.01329)	(15.13318)
$P_{j,t-1}$		1.42867***	1.77714***
		(0.10966)	(0.18164)
$d_{news} \times P_{j,t-1}$		-0.04022	-0.00115
		(0.18083)	(0.26824)
$d_{implem} \times P_{j,t-1}$		3.05545***	3.35558***
		(0.89086)	(0.78411)
Period F.E.	Y	Y	Y
Field F.E.	Y		
Observations	2484	1856	1856

Notes: The table reports the effect of a one-month shorter extension in average patent term and of past innovation based on the baseline DiD specification (5) (column 1) and on the augmented DiD specification (7) estimated by OLS (column 2) or 2SLS using fitted values from (5) as instruments for lagged innovation (column 3).Standard errors are clustered by technical field. Statistical significance levels: $^{*}(p < 0.10)$,** $(p < 0.05)$,*** $(p < 0.01)$

F.2 Alternative mechanisms: Evidence on competition

In this subsection, I analyze the role of competition and incumbent firms' patenting for the transmission of indirect effects to post-implementation innovation.

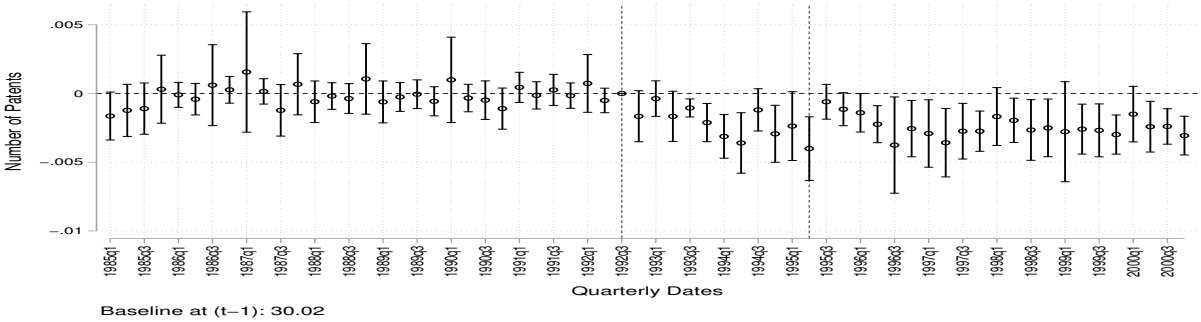
F.2.1 Field-level evidence on the role of competition

As a proxy of competition in the technological space, I consider the Herfindahl-Hirschman Index (HHI) of concentration based on the share of patents granted to different applicants within a given field. Results available upon request verify the robustness of the findings to an alternative proxy of entry intensity based on the share of patents granted to applicants that file a patent application in a specific field for the first time. I build the HHI using PATSTAT (EPO, 2017) and STAN harmonized applicant's identifiers from the EPO Worldwide Bibliographic Database (available in PATSTAT).

I first test whether the change in average patent term has stronger effects on innovation in more competitive fields. To this end, I estimate triple-difference specification (9) replacing as interaction variable with ΔT_j the field-specific HHI computed using field-j patents applied

for before policy news. Figure F1 illustrates the effect of a one-month *shorter* extension in the average patent term on the number of (granted) patents by application quarter, conditional on a unit-higher HHI, i.e., conditional on *higher* concentration. The negative sign of the effects confirms that the responsiveness of innovation to the average change in patent term is relatively higher when competitive pressure is stronger.

Figure F1: Heterogeneity analysis based on competition in patenting by field



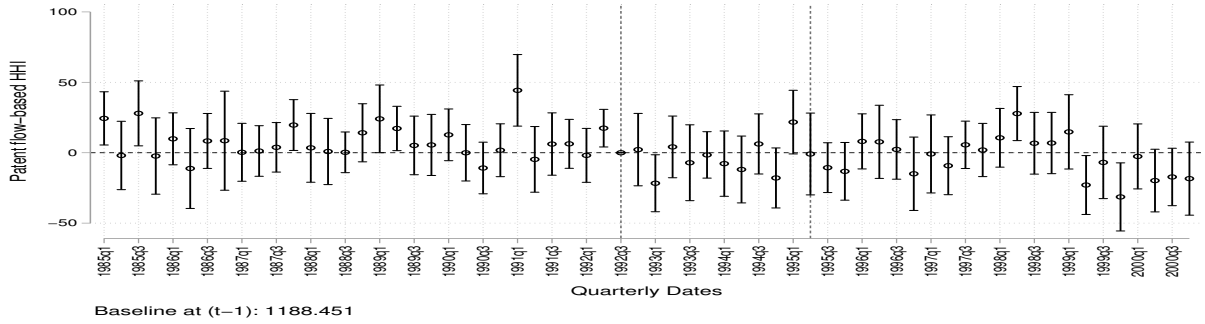
Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the number of (granted) patents by application quarter, conditional on a unit-higher Herfindahl-Hirschman Index (HHI) of concentration. The effects are based triple-difference estimates of specification (9) with HHI or entry intensity as interaction variables. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

Second, I analyze the effects of the patent term change on competition. To this end, I estimate DiD specification (5) with $HHI_{j,t}$–i.e., the Herfindahl–Hirschman Index based on the flow of patents applied for in quarter t and classified in field j—as dependent variable. Figure F2 illustrates the (null) effect of a one-month shorter extension in average patent term. If a relatively shorter extension in average patent term enhances post-implementation innovation due to more intense competition, the estimated effect should be negative. Therefore, this finding suggests that changes in competition are not the main driver of overall post-implementation effects. All the results, available upon request, are analogous for entry intensity.

F.2.2 Evidence on the average quality of incumbents' patents

Finally, I examine the effects of the patent term change on patenting behavior of incumbent applicant firms. A longer extension of patent term may indirectly affect innovation by inducing incumbents to use their patent rights to foreclose competition by other innovators. However, these patents would typically display lower quality. Therefore, I estimate specification (5) using the average quality of patents granted to incumbent applicants relative to new entrants as dependent variable. Figure F3 displays the effect of a one-month *shorter* extension in average patent term on the ratio of average citations for incumbents' patents and

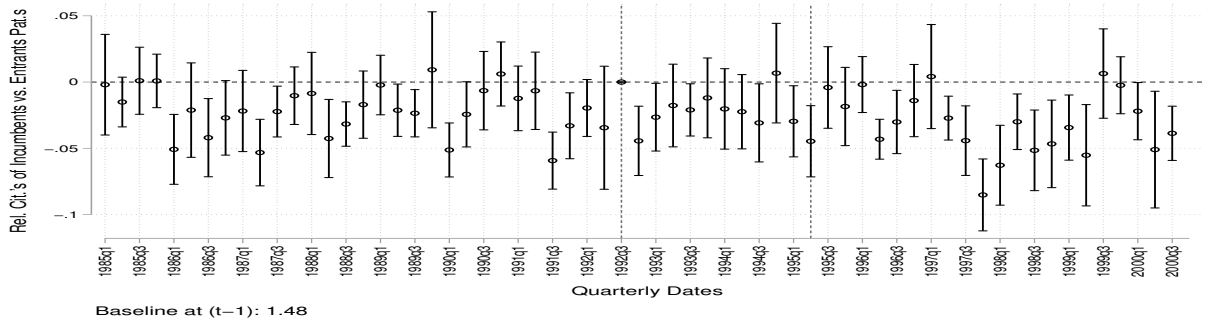
Figure F2: DiD estimates for proxies of competition in patenting by field



Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the level of quarter-t and field-j Herfindahl–Hirschman Index of concentration of granted patent applications across different applicants, based on estimates of DiD specification (5) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

entrants. The findings suggest the use of patents for anti-competitive purposes by incumbent innovators is not a key channel in driving overall post-implementation effects.

Figure F3: DiD estimates for relative quality of incumbents vs. entrants patents



Notes: The figure depicts the relative effect of a one-month *shorter* extension in the average patent term on the the relative quality of incumbents and entrants patents—measured by the average 5-year forward citations per patent—based on estimates of DiD specification (5) with $X_j = \Delta T_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.