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

Fabio Bertolotti†

July 2025

Abstract

I examine the impact of patent term on R&D and innovation in the presence of policy anticipation, common in real-world settings. Using a difference-in-difference design, I exploit quasi-experimental variation in U.S. patent term across technological fields due to the ratification of TRIPs agreements in 1995. Despite a general increase in *average* patent term, in most fields innovators faced a considerable probability of 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 thank the Editor and three anonymous referees for their comments, which helped me substantially improved the paper. 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, help, 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, Jonathon Hazell, Ethan Ilzetzki, William Kerr, Joshua Krieger, Per Krusell, Andrea Lamorgese, Salvatore Lattanzio, Salvatore Lo Bello, Massimiliano Marcellino, Kyle Myers, Francesco Nava,Tom Nicholas, Pietro Peretto, Micheal Peters, Akash Raja, Giorgio Saponaro, Mark Schankerman, Lorenzo Stanca, 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. For this reason, the economic literature has traditionally focused on the long-run effects of policies promoting Research and Development (R&D) and innovation. However, the implementation of such policies 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 may impact long-term productivity and welfare, as highlighted by a growing body of literature (e.g., Comin and Gertler, 2006; Benigno and Fornaro, 2018; Bianchi, Kung and Morales, 2019; Vinci and Licandro, 2020; Fornaro and Wolf, 2021)

This paper jointly analyzes the short- and long-term effects of innovation policy shocks. Specifically, I present new empirical evidence on the impact of an *anticipated* change in U.S. patent term—the duration of the legal monopoly granted by patents—on innovation and R&D.

Although patent term has been recognized as a key policy tool since Nordhaus (1967), empirical evidence on its effects remains limited (Budish, Roin and Williams, 2016), and normative recommendations vary widely.[1]

Moreover, the short-term effects of policy anticipation on innovation are ex-ante ambiguous. News of a future patent term reduction may lead to an immediate decline in innovation, as firms scale back R&D efforts in anticipation of lower future returns. This behavior would align with the long-run effects of patent protection documented by previous literature. However, perhaps surprisingly, the same news may also temporarily *accelerate* R&D on existing projects, as innovators seek to benefit from the more favorable policy while it remains in effect.

The empirical analysis of the paper 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 a patent is granted, the effective U.S. patent term changed from 17 years to 20 years *minus* the pending period, i.e., the time between application and grant, during which 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.

[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, U.S. 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.

[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.

First, cross-sectional variation in pending period across technological fields. Indeed, U.S. 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 U.S. 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 using two key moments. First, the field-specific average pending period, which determines the change in the average effective patent term. Second, the field-specific share 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.

On average, most fields benefit from an extension in *average* patent term, but variation across fields is significant, approximately ranging between one-year reductions and 2.5-year extensions. Notably, even in fields with an extension in average patent term longer than one year, individual applications could face up to 25% probability of getting a patent term reduction. More broadly, there is a strong negative correlation between the two moments: a 10-percentage-point increase in the probability of patent term reduction corresponds to an approximately four-month shorter average patent term extension.

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

The empirical analysis reveals three key findings. First, following the policy news and prior to its implementation, innovation accelerates more in fields with a higher probability of patent term reduction and, thus, with a shorter average patent term extension (**Fact 1**). As in most fields the probability of individual patents experiencing a patent term loss was substantial, these estimates are consistent with the observed surge in aggregate patent filings during the anticipation phase, suggesting that innovators primarily responded to the event of a patent term reduction.

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 with a higher probability of patent term reduction increased their R&D spending more during the anticipation phase. This evidence 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**). Thus, R&D and innovation remain relatively higher in fields with a higher fraction of patents with shorter patent term even after June 1995.

The positive response of innovation to the news of a potential patent term reduction—and its persistence following implementation—may appear surprising. Prior literature documents that stronger patent protection tends to stimulate innovation as a direct effect, which might suggest that news of a *reduction* in patent term should *deter*, rather than encourage, innovative activity. To address this apparent inconsistency, I first conduct a series of robustness checks to rule out several confounders. These include potential mismeasurement of innovation, unrelated technological trends, concurrent shifts in international trade and patent policy, and changes in the rigor of patent examination. Moreover, I validate that the predicted variation in patent term prior to the policy news closely matches the realized variation across technological fields after implementation.

I then argue that the observed news effects are in fact consistent with innovators favoring a longer over a shorter patent term. By accelerating R&D and innovation in the anticipation phase, inventors more exposed to a patent term reduction can file applications before the implementation of the new policy, obtaining a longer patent term under the old regime while it remains in effect. Supporting this interpretation, complementary analyses exploiting within-field dispersion in pending periods suggest that firms were particularly responsive to scenarios involving adverse policy changes, consistent with loss aversion.

Moreover, the post-implementation persistence of the effect can be understood through the lenses of endogenous growth models with knowledge spillovers, where current innovation depends not only on contemporaneous R&D but also on past discoveries. In this framework, the initial acceleration in innovation determined by the policy news may propagate over time, fostering additional inventions. As a result, the post-implementation estimates associated with Fact 2 capture both the direct effect of the policy change and the indirect effects of knowledge diffusion induced by the earlier news shock. If the latter channel temporarily outweighs the former, the reduced-form estimates may display a persistent positive response.

The analysis then 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. Thus, 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**), consistent with prior literature. 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 1.7% (20.9%). Thus, longer patent protection effectively stimulates more innovation as a *direct* effect and the magnitude aligns with previous literature. Specifically, Budish, Roin and Williams (2015, 2016) find that a one-year extension of patent monopoly increases R&D by 7% to 22% in the pharmaceutical industry. Additionally, the model in Hémous et al. (2023) 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 one-percentage-point higher probability of patent term reduction for future patents corresponds to a 1.4% larger increase in patents before implementation (2.9% 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 exploits heterogeneity in technological dependence across fields, proxied by patents backward citations intensity. I find that the initial acceleration in innovation following the policy news leads to a stronger persistence of higher innovation in fields where new inventions rely more on past innovations from the same field, consistent with evidence by Hegde, Herkenhoff and Zhu (2023) on technology disclosure externalities. Additionally, an increase in time-varying technological dependence measures indicates that this channel is the primary driver of indirect post-implementation effects, while alternative mechanisms—such as changes in technological competition or adjustments in patenting strategies—lack comparable support in the data.

The empirical analysis underscores the importance of considering the intertemporal tradeoffs 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, innovation-policy interventions influence both short- and long-run outcomes. With policy anticipation, these effects may push in opposite directions and, crucially, near-term variation in innovative activity may itself drive medium- to long-term effects due to technological and knowledge externalities (Romer, 1990; Hegde, Herkenhoff and Zhu, 2023). Thus, the rich dynamics of my empirical estimates provide valuable inputs for normative analysis, as they can inform the calibration of key parameters of innovation-based growth models used to study the optimal patent term.

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. Sections 3 and 4 discuss the data and preliminary descriptive evidence, respectively. Section 5 outlines a simple conceptual framework for the empirical analysis. Section 6 presents the main empirical evidence. Section 7 investigates the mechanisms and connects the empirical results to endogenous growth theory. Section 8 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; Bertolotti and Lanteri, 2024). 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, 2021; Moscona, 2021; Acikalin et al., 2022) by focusing on a specific aspect, the patent term, which, since Nordhaus (1967), 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 U.S. 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 U.S. firms anticipated the TRIPs and show the short- and medium-term importance of news effects. In Section 4.2 and Appendix B I detail the differences between the two analyses in light of policy anticipation.[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

[3] In summary, Abrams (2009) estimates a two-period DiD specification comparing patenting before and after the implementation shock of June 1995. Disregarding potential news effects may lead to an imprecise 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 DiD assumptions.

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., 2018); 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" effect. This finding closely relates to Hegde, Herkenhoff and Zhu (2023), who document that a stable increase in the speed of knowledge diffusion permanently increases the rate of follow-up innovation. My setting features transitory effects from the same mechanism as knowledge diffusion accelerates only temporarily during the news phase. Additionally, 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 U.S. 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 U.S. 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

[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 6.5 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.

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 U.S. 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 U.S. firms were aware of the impending policy change well before its formal adoption.[6]

First, U.S. 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 U.S. Advisory Committee on Trade Policy and Negotiations, which included executives from, e.g., IBM and Pfizer, significantly shaped the stance of the U.S. delegation. Second, the adjustment of the U.S. 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, U.S. 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

[6]The URAA included a clause extending patent term for active patents filed before June 8, 1995, if the new rules implied a later expiry date. However, this detail was likely unanticipated. Historical records of U.S. policy debate—e.g., hearings of the Senate Advisory Committee on Patent Law Reform—never mention it, and it was absent from TRIPs implementation in other countries, like Canada. Moreover, even if anticipated, it would not affect the main findings. In fact, the clause applied only to patents benefiting from a longer term under the new policy, while, as discussed in Section 4.1, U.S. innovators primarily reacted to the probability of a *reduction* in patent term. I further discuss the role of possible anticipation of this clause in Section 7.1.

[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

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), U.S. innovators were aware of the potential policy change and could respond to the information. The flexible DiD specification in Section 5 captures potential anticipation effects without imposing ex-ante assumptions. Section 6.1 discusses the timing of news effects, their robustness to earlier dates, and their interpretation in light of remaining uncertainty. Section 4.2 and Appendix B further explore the consequences of neglecting news effects for the parallel trends assumption in DiD analysis and relate this paper to Abrams (2009).

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 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 U.S. 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 experiment with narrower field definitions in Section 6.5.

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 the 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. The analysis primarily focuses on patent term reduction probability, as descriptive evidence and complementary analyses in Sections 4.1 and 7.1, respectively, suggest that it is the main driver of DiD estimates.

reads: *"The round was supposed to end in December 1990, but the U.S. and EU disagreed on how to reform agricultural trade and decided to extend the talks. Finally, In November 1992, the U.S. and EU settled most of their differences in a deal known informally as the "Blair House accord" [...]"*

[11] Applicants propose relevant patent classes following detailed pre-filing guidelines and the USPTO verifies them to assign the application to the appropriate examination unit. This verification suggests that there is limited scope for strategic choice of patent classes by applicants. Moreover, since pending periods also depend on applicants' responsiveness (Lemus and Marshall, 2018), Section 6.5 shows no correlation between applicant responsiveness and pre-existing heterogeneity in average pending periods across fields.

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

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

$$\Delta T_j = 20 \text{ years} \times 365 \; - \; \overline{PP}_j \; - \; 17 \text{ years} \times 365 \tag{2}$$

where PL_j is the fraction of patents classified in field j and granted before the TRIPs news with a pending period longer than three years, and $\overline{PP}_j$ is the average pending period, in number of days, for such patents. 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.[12] The interactions of PL_j or ΔT_j with quarterly fixed effects constitute the treatment variables in the field-level DiD empirical analysis of Section 6.

Figure 1a illustrates the distribution of patent term reduction probability, PL_j, across technical fields. Many fields featured a significant probability that individual patents obtained a pending period longer than three years: 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%.

In fact, most fields were expected to obtain an anticipated increase in *average* effective patent term, whereas few fields faced a projected reduction. Figure 1b depicts the distribution of the average change, ΔT_j, across technical fields. The mean of this distribution is +473 days, or about 15.5 months, with a standard deviation of 118 days.

Finally, 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 6.5 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.

2.4 Relation between Expected and Realized Patent Term Changes

The two proposed measures of TRIPs effect on patent term are based on pre-news variation in pending periods across fields. In this subsection, I show that this ex ante variation consistently predicts realized heterogeneity in patent term across fields *after* the policy shocks.

To this end, I estimate the relationship between the ex ante treatment variable X_j and its realized quarterly counterpart X_{jt}—where j indexes fields and t quarters—through the following

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

Figure 1: TRIPs-Related Change in Patent Term

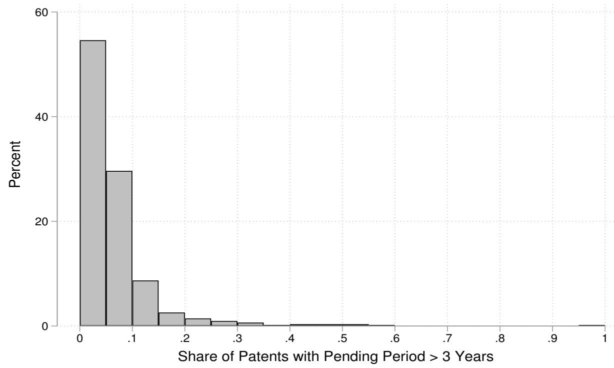

(a) Patent Term Reduction Probability Distr.

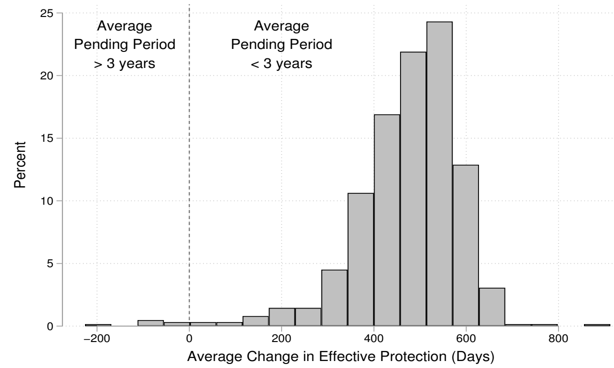

(b) Average Patent Term Change Distr.

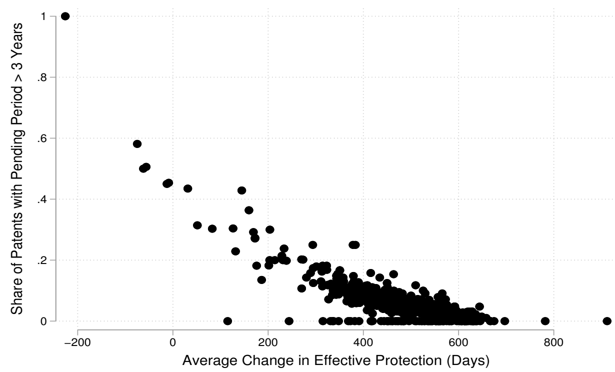

(c) Correlation of Average Patent Term Change and Reduction Probability

Notes: Panels (a) and (b) illustrate the distribution of the TRIPs-induced patent term reduction probability and the change in average effective patent term across technical fields, respectively. In panel (b), 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.

regression:

$$X_{jt} = d_t + \rho X_j + u_{jt} \tag{3}$$

where X_{jt} is either: (i) the share of patents filed in quarter t and field j with a realized pending period longer than three years (PL_{jt}) or (ii) the difference (in days) between the realized field-specific average pending period and three years (ΔT_{jt}). The term d_t denotes quarter fixed effects capturing common shifts in the relationship between X_{jt} and X_j, and u_{jt} is an error term. I estimate this specification over two periods: the post-news sample (1992Q4–2000Q4) and the post-implementation sample (1995Q3–2000Q4).

The coefficient of interest, ρ, captures how variation in the ex ante proxy X_j maps into its realized counterpart X_{jt} over time. A coefficient of one implies that field-level differences in

Table 1: Relation Between Pre- and Post-News Patent Term Variables

	PL_{jt} Post 92Q3	PL_{jt} Post 95Q2	ΔT_{jt} Post 92Q3	ΔT_{jt} Post 95Q2
Pre-News PL_j	0.957	1.002		
	(0.070)	(0.080)		
	[0.819,1.094]	[0.846,1.159]		
Pre-News ΔT_j			1.053	1.104
			(0.059)	(0.069)
			[0.938,1.169]	[0.969,1.239]
Average of d_t	0.140	0.164	-269.681	-337.969
	(0.005)	(0.006)	(30.170)	(35.441)
	[0.131,0.150]	[0.152,0.175]	[-328.931,-210.431]	[-407.570,-268.368]
Observations	18084	12029	18084	12029

Notes: The table reports point-estimates, standard errors (in parentheses), and 95% confidence intervals (in square brackets) of ρ and averages of quarterly dummies d_t in specification (3) for patent term reduction probability $X_j = PL_j$ (columns 1 and 2) and the average patent term change $X_j = \Delta T_j$ (columns 3 and 4). Regression (3) is estimated on a sample of field-quarter observations with more than one granted patent application, to reduce noise in estimates of realized X_{jt}, and on post-news quarters $1992Q4 - 2000Q4$ (columns 1 and 3) or post-implementation quarters $1995Q3 - 2000Q4$ (columns 2 and 4). Results are robust to including all fields. Standard errors are clustered by 4-digit technical field.

the ex ante variable translate one-for-one into realized variation after the shocks, apart from quarter-specific uniform level shifts captured by time dummies.

Table 1 reports point estimates, standard errors (in parentheses), and 95% confidence intervals (in square brackets) for ρ and the average of quarterly dummies (d_t), using $X_j = PL_j$ (columns 1 and 2) and $X_j = \Delta T_j$ (columns 3 and 4). Across both patent term measures and sample periods, the estimates imply that predicted variation in patent term changes translates almost one-for-one into realized variation across fields.

The average of quarterly dummies indicates a general increase in examination time, with the probability of a pending period exceeding three years rising by about 14 percentage points after the policy news. This change reflects a broad upward trend in examination duration, but importantly, one that appears *uncorrelated* with pre-news pending periods across fields.

Appendix C.1 provides a graphical representation of the uniform increase in realized pending periods across levels of PL_j or ΔT_j. Furthermore, Section 6.5 discusses that examination time did not endogenously respond to the policy and that DiD estimates are unaffected when using realized rather than expected measures of patent term change.

3 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

Table 2: Summary Statistics by Technical Field × Quarter

Variable	Mean	S.D.	10th Perc.	90th Perc.
Number of Patents	37.73	139.20	0.00	81.00
5-Year Citation-Weighted Patents	205.58	1097.17	0.00	378.00
Patent Value (Million Dollars)	366.96	3416.24	0.00	435.56
Patent Term Reduction Prob. PL_j	0.06	0.08	0.00	0.12
Change in Avg. Patent Term ΔT_j (Days)	472.66	117.42	343.55	590.79
Within-Field S.D. of ΔT_j (days)	37.13	38.93	11.53	72.55
Share of Patents Renewed to Max. Term	0.29	0.26	0.00	0.63
Sh. Patents w. Within-Field Backwd. Cit.s	0.19	0.28	0.00	0.58

Notes: The table reports sample summary statistics at the 4-digit IPC field × quarter level.

(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).[13] 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 U.S. patent law.[14] The average number of quarterly patents and 5-year forward citations-weighted patents are 38 and 206, respectively, with standard deviations of 139 and 1,097. Table 2 provides field-level summary statistics.

The firm-level dataset is a yearly panel of 2,410 listed U.S. 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 13 (granted) patents per year, and the average annual R&D expenditure is \$85 million, with standard deviations of 88 and \$424 million, respectively. Approximately 70% of firm-year observations display non-missing R&D, with less than 3% of them being zero, and 44% have zero granted applications. Table 3 presents summary statistics for the firm-level sample.

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 A1 provides summary statistics for the sectoral variables.

[13] 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.

[14] AIPA mandated the publication of patent applications 18 months after filing, effective for patents filed after 11/29/2000. Analysis by Hegde, Herkenhoff and Zhu (2023) suggests the policy had no impact on patenting prior to its implementation, which I confirm in complementary analyses available upon request.

Table 3: Summary Statistics by Compustat Firm × Year

Variable	Mean	S.D.	10th Perc.	90th Perc.
Number of Patents	12.61	88.08	0.00	11.00
5-Year Citation-Weighted Patents	182.35	1468.60	0.00	162.56
Patent Value (Million Dollars)	247.53	3136.16	0.00	62.33
R&D (Million Dollars)	85.13	424.24	0.23	103.00
Sales (Million Dollars)	2312.02	10094.70	2.87	4304.07
Age	14.76	13.63	1.00	36.00
Patent Term Reduction Prob.	0.08	0.06	0.03	0.14
Average Patent Term Change (Days)	449.56	93.73	343.95	545.64

Notes: The table reports sample summary statistics at the Compustat firm × year level.

4 Descriptive Evidence and the Role of Anticipation

4.1 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 6.

Figure 2a shows the evolution of granted patents *by application quarter* in three fields with similar pre-news trends but differing in patent term reduction probability and average patent term change. 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 ($PL_j = 0.43$). The long-dashed line shows G10L (in physics), with a larger extension ($\Delta T_j = 132$ 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 ($\Delta T_j = 599$ days) and the loss probability is negligible ($PL_j = 0.01$).

In fields with a significant loss probability (C12P and G10L), 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 thus a shorter average term extension—accelerate patenting before the new policy takes effect. Specifically, historical records of the U.S. policy debate on TRIPs adoption highlight that the risk of patent term reductions was salient.[15]

The increase starts small but intensifies as implementation nears, with a peak in 1995Q1 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 option created a strong incentive to file during the transition period, especially in fields facing potential patent term reductions, explaining the observed bunching in 1995Q2.

[15] The "Public Comments" subsection of Section II "Patent Term" (p. 75) of the 1994 Congressional hearings on S.2368 and H.R.4894 notes: "The following disadvantages [...] were noted: [...] (4) in some cases, the present term of protection (17 years from grant) can be shorter than a term of 20 years from the filing date."

After implementation, patenting remains higher than pre-news trends in the fields with the largest acceleration at news (C12P and G10L), suggesting persistent effects in fields with a larger fraction of patents obtaining a patent term reduction. In contrast, patenting remains stable and slightly above trend in E05D, the field with the lowest reduction probability.

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 patent term reduction probability PL_j. Across all quintiles of PL_j, patenting increases at news, with the largest acceleration in fields with *higher* patent term reduction probability. 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—which appears the main driver of news effects—and a shorter average extension. This descriptive evidence will guide the interpretation of the DiD estimates in Section 6.

Figure 2: Descriptive Evidence on Patenting Across Fields

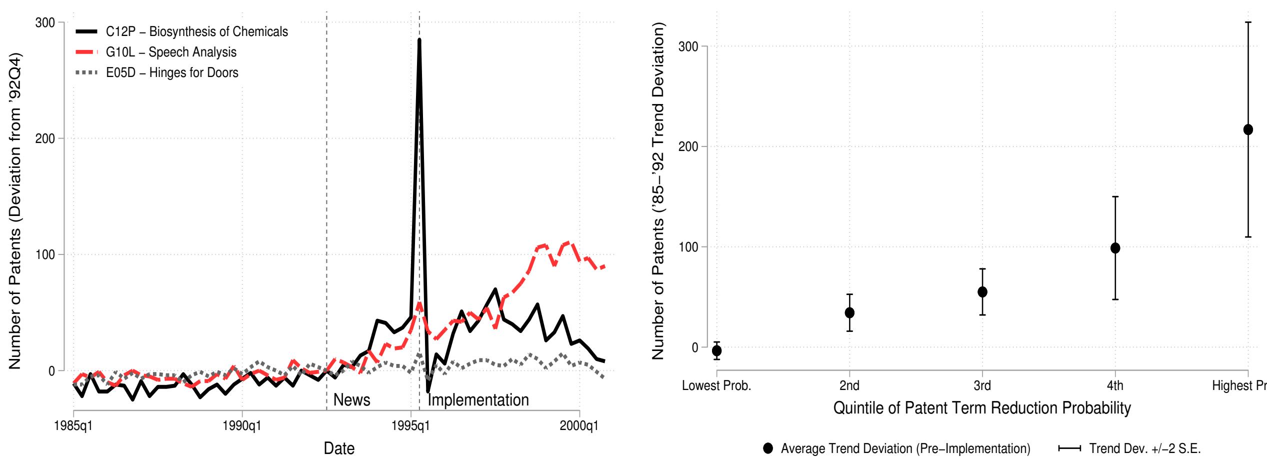

(a) Patenting across Fields (b) Pre-Implementation Acceleration

Notes: Panel (a) illustrates the evolution of the number of granted patent applications by application quarter in three technical fields: C12P with $PL_j = 0.435$ (black solid line); G10L with $PL_j = 0.229$ (red long-dashed line); E05D with $PL_j = 0.011$ (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 and patent term reduction probability—equal to the share of pre-news patents with pending period longer than 3 years—by quintile of the latter variable (x-axis).

4.2 The Role of Policy Anticipation

Previous descriptive evidence shows that innovation dynamics markedly changed across fields in response to policy news. This subsection argues that ignoring these anticipation effects can lead to an imprecise interpretation of DiD estimates based solely on comparing outcomes immediately before and after policy implementation. This distinction is key to understanding how my approach differs from that of Abrams (2009), who analyzes the same policy episode under the assumption of no anticipation of TRIPs and, as a result, reaches different conclusions.

First, anticipation violates a core assumption of the DiD identification strategy: that pre-implementation outcomes serve as an unaffected baseline for the pre-post implementation DiD comparison. Because at news innovation increases the most in fields facing higher reduction probability, pre-implementation levels are artificially inflated in these fields and thus do not provide a reliable reference for the DiD comparison around implementation. For instance, between April 1994 and March 1995—the pre-implementation window used by Abrams (2009)—the average monthly number of patents in field C12P was 15.1 units higher than in the quarter before policy news, compared to just 2.4 units in field E05D. As a result, post-implementation innovation appears relatively *weaker* in fields like C12P compared to E05D when considering the inflated pre-implementation baseline, despite still being *stronger* when considering the unaffected pre-news innovation level (Figure 2a). Thus, the sign of post-implementation DiD estimates reverses depending on the considered reference point.[16]

Second, Abrams (2009) estimates field-specific linear time trends over narrow windows (6, 12, or 24 months) before and after implementation, excluding a four-month gap around June 1995. These trends are intended to control for underlying innovation dynamics unrelated to the policy. However, because innovation responds to policy news, the estimated trends are based on data influenced by the shock. As a result, the constructed counterfactual trajectories diverge from actual pre-news dynamics, violating the parallel trends assumption. In contrast, with the DiD specification presented in Section 5 I find no evidence of pre-trends.

Appendix B provides further details and illustrates these differences with an example.

5 Conceptual Framework for the Empirical Analysis

This section introduces a stylized framework to analyze news and implementation effects under anticipation. I first derive the causal marginal responses of interest, then link them to the reduced-form DiD specifications of the paper. Section 6.5 discusses identification concerns.

5.1 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

[16] From August 1995 to July 1996, average monthly patents in C12P (E05D) were +6.7 (+1.8) units relative to the news quarter, but -8.3 (-0.6) units compared to the pre-implementation window used by Abrams (2009). Since C12P features a higher patent term reduction probability than E05D, the DiD estimate for the post-implementation period would be *positive* when considering the unaffected pre-news reference point, but *negative* when assuming no anticipation and using the endogenous pre-implementation baseline.

a function of past innovation and a policy vector $\boldsymbol{\tau}_t$ collecting current and future expected policies. This reduced-form representation is consistent with most endogenous growth theories. For the case of patent term T, $\boldsymbol{\tau}_t = [T_t, T^e_{t+1}, ...]$, where T_t denotes the current policy and T^e_{t+1} the expected one at $t+1$. Therefore, aggregate innovation itself is a function of past cumulative innovation and policy, i.e.,

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

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:

$$\frac{\mathrm{d}I_{news}}{\mathrm{d}T_{implem}} = \frac{\partial G(\cdot,\cdot)}{\partial CumI_{pre-news}} \underbrace{\frac{\partial CumI_{pre-news}}{\partial T_{implem}}}_{=0} + \frac{\partial G(\cdot,\cdot)}{\partial T_{implem}} = \frac{\partial G(\cdot,\cdot)}{\partial T_{implem}} \tag{4}$$

The first term denotes the impact of the future policy change on news-period innovation through changes in pre-news innovation. This term 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 (4), and a *direct* effect, due to changes in R&D caused by the actual implementation of the policy:

$$\frac{\mathrm{d}I_{implem}}{\mathrm{d}T_{implem}} = \underbrace{\frac{\partial G(\cdot,\cdot)}{\partial CumI_{news}} \frac{\partial CumI_{news}}{\partial T_{implem}}}_{Indirect} + \underbrace{\frac{\partial G(\cdot,\cdot)}{\partial T_{implem}}}_{Direct} \tag{5}$$

The third effect of interest is the "direct" impact of policy implementation, i.e., the second term of expression (5). It represents the incremental change in innovation caused by a concomitant marginal change in patent term absent anticipation effects.

5.2 Estimation of News and Overall Implementation Effects

To estimate news and "overall" implementation effects (4) and (5) on aggregate innovation, I leverage cross-sectional variation in patent term loss probability PL_j or average effective patent term change ΔT_j across technological fields in the DiD specification:

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

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 PL_j or ΔT_j, and $\varepsilon_{j,t}$ is an idiosyncratic error term.

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).

I also analyze the following Poisson specification for positive count variables

$$Y_{j,t} = \exp\left\{ \alpha_j + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \gamma_k \mathbf{1}_{(t=k)} + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \beta_k \mathbf{1}_{(t=k)} X_j + \varepsilon_{j,t} \right\}, \tag{7}$$

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 5.1. For $k \leq 1992Q3$, β_k reflects the marginal effect of PL_j or ΔT_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 about the future policy on (granted) patent applications, capturing the marginal variation in X_j for quarter k. This effect aligns with the news effect in equation (4). 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 (5).

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).

[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.

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 de-Grazia, 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 6.2. 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]

5.3 Recovering the Direct Policy Effect

The empirical analysis proceeds to distinguish between the direct and indirect effects within the overall implementation effect in equation (5). 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 (6) to control for field-specific innovation histories, following the discussion on DiD with path-dependence in Angrist and Pischke (2009). 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). Hence, the modified DiD specification is:

$$\begin{aligned} Y_{j,t} = & \sum_{k \neq' 92Q3} \mathbf{Z}_{pre,j} \mathbf{1}_{(t=k)} \eta_k + \sum_{k \neq' 92Q3} \gamma_k \mathbf{1}_{(t=k)} + \sum_{k \neq' 92Q3} \phi_k \mathbf{1}_{(t=k)} X_j + \\ & + \sum_{k \neq' 92Q3} \psi_k \mathbf{1}_{(t=k)} \underbrace{\overline{Y}_{j,k-\mathcal{A}-1:k-1}}_{\equiv \frac{1}{\mathcal{A}} \sum_{q=k-\mathcal{A}-1}^{k-1} Y_{j,q}} + \psi_0 \underbrace{\overline{Y}_{j,t-\mathcal{A}-1:t-1}}_{\equiv \frac{1}{\mathcal{A}} \sum_{q=t-\mathcal{A}-1}^{t-1} Y_{j,q}} + \upsilon_{j,t}. \end{aligned} \tag{8}$$

To control for field-specific time-invariant characteristics, I replace field fixed effects by 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.[19] The second line of (8) parsimoniously accounts for the quarter-specific impact of a field 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 $\mathcal{A}$ equal to 10, corresponding to

[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.

[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 (6).

the TRIPs anticipation period, but confirm robustness with $\mathcal{A}$ between 7 and 16 quarters.

The new DiD coefficients ϕ_k estimate the effect of a marginal change in the probability of a patent term reduction (or in the average patent term change) 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 (5). Before implementation, they coincide with the news effects, as per equation (4). Consistently, the empirical findings in Section 6.3 show that the direct effect closely aligns with the baseline DiD estimates during the news period. In contrast, the direct effect has opposite sign than the "overall" effect after implementation.

6 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.

6.1 News and overall implementation effects by field

Figure 3 presents the results of DiD specification (6) for the number of granted patents by field and *application quarter* as outcome variable. Panel 3a shows the effect of a one-percentage-point *higher* patent term reduction probability ($X_j = PL_j$), while panel 3b shows the effect of a one-month *shorter* extension in average patent term ($X_j = \Delta T_j$). Panels 5a and 5b display the same effects using the Poisson DiD specification (7). 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 6.5 will discuss the absence of confounders occurring around or after the policy change

[20]For $X_j = PL_j$, a linear trend with a quarterly slope of 0.056 (equal to the 2000Q4 point estimate divided by the number of post-news quarters) would be detected with a power of 0.99, while a trend with a slope of 0.028 (equal to half of the 2000Q4 point estimate divided by the number of post-news quarters) would be detected with a power of 0.51. For $X_j = \Delta T_j$, a linear trend with a quarterly slope of 0.15 (equal to the 2000Q4 point estimate for a one-month shorter extension divided by the number of post-news quarters) would be detected with a power of 1, while a trend with a slope of 0.075 (equal to half of the 2000Q4 point estimate for a one-month shorter extension divided by the number of post-news quarters) would be detected with a power of 0.79.

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

(a) Effects of Patent Term Reduction Probability

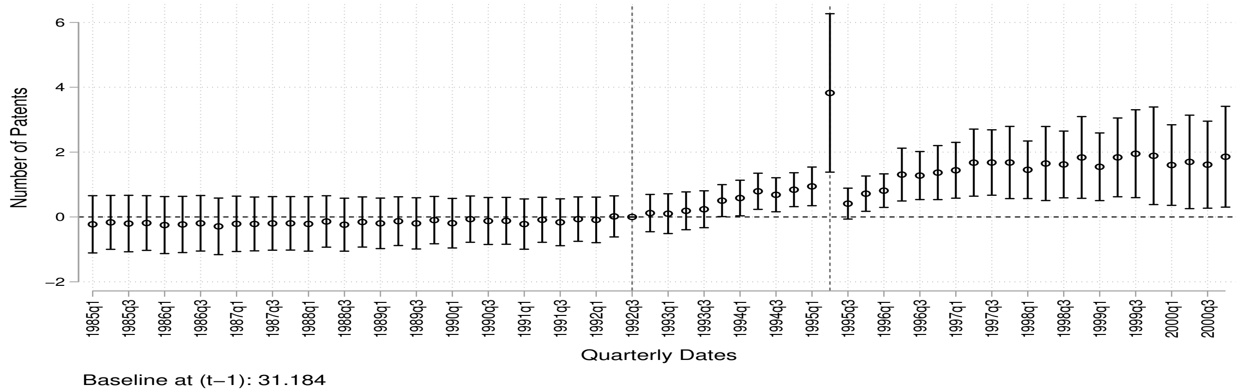

(b) Effects of Average Patent Term Change

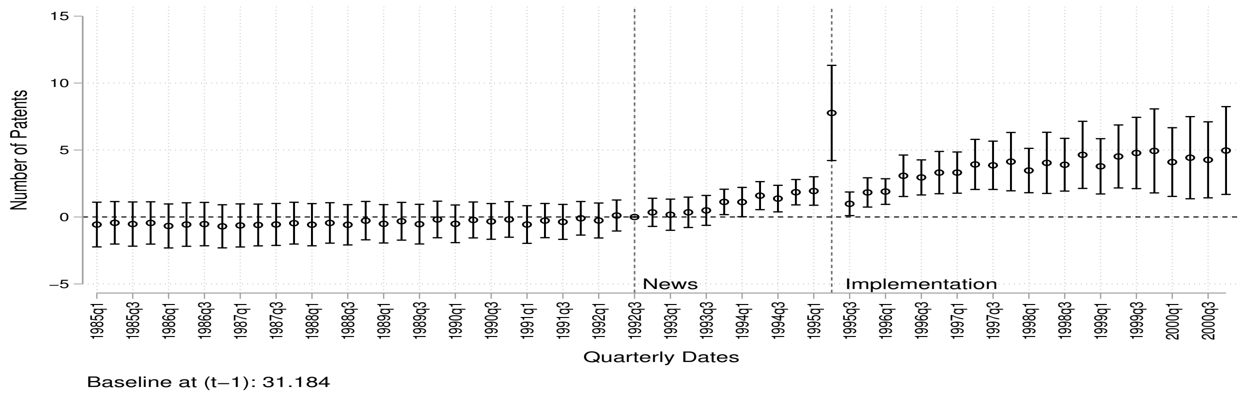

Notes: The figure represents the effects of the patent term change on field-level (granted) patent applications by application quarter from DiD specification (6). In panel (a), the regressor of interest is $X_j = PL_j$ and the estimates represent the relative effect of a one-percent *higher* probability of a patent term reduction on the level of the outcome variable. In panel (b), 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 effect is computed as the DiD coefficients estimates $\widehat{\beta}_k$ times -30 (days). 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.

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-percentage-point higher probability of patent term reduction (Figure 3a) and a one-month shorter extension in average patent term (Figure 3b) on patenting. This result aligns with the descriptive evidence in Section 4.1. A higher probability of patent term reduction—or a shorter extension in average patent 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 policy implementation approaches. For instance, one year after the news and two years before implementation, a one-percentage-point higher probability of patent term reduction leads to 0.19 additional (granted) applications per quarter (+0.5 percent in Poisson DiD estimates, Figure 5a). Similarly, in fields with a one-month shorter extension, patenting increases by 0.35 additional units (+0.8 percent in Poisson estimates, Figure 5b). These effects nearly triple two years after the news and one year before implementation.

The positive relationship between patenting and news of a higher probability of patent term reduction may appear surprising given prior evidence that stronger patent protection stimulates innovation as a direct effect. However, as I discuss in the Introduction, my findings are consistent with prior literature: innovators who are more exposed to a future patent term reduction and want to obtain the longest term of protection have stronger incentives to accelerate innovation and file applications before policy implementation. Descriptive evidence and complementary analyses of Sections 4.1 and 7.1 support this interpretation.

The dynamics of news effects reflect 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 6.2—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 4.1, this spike 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).

I preview the effects of patent term on patent quality (Section 6.5) by estimating specification (6) 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-percentage-point higher probability of patent term reduction on quality-adjusted patents, while Figure 5c provides similar estimates using the Poisson specification (7). Like in Figure 3b, there are no significant pre-trends. Positive estimates after the news confirm that fields with *higher* patent

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

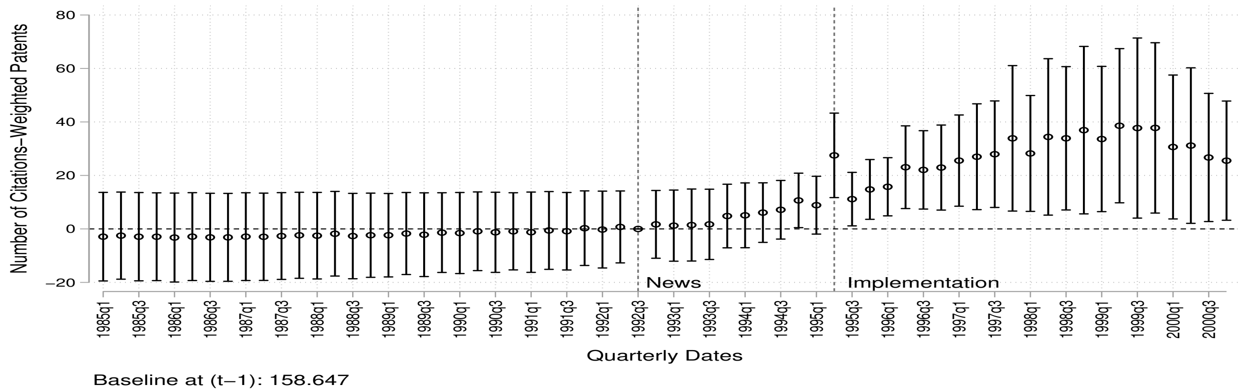

Notes: The figure depicts the relative effect of a one-percentage-point *higher* patent term reduction probability on the level of 5-year citations-weighted (granted) patents by application quarter, based on estimates of DiD specification (6) with $X_j = PL_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

term reduction probability see a *larger* increase in quality-adjusted patenting during the anticipation phase. Section 6.5 provides further evidence on additional *average* patent quality measures and links the results 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 6.2 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-percentage-point increase in the probability of a patent term reduction results in 1.44 additional quarterly patents (+2.7 percent in Poisson estimates), and a one-month shorter patent term extension corresponds to 3.3 more patents (+5.9 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 (Hegde, Herkenhoff and Zhu, 2023). Appendix C.2 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

Figure 5: Poisson DiD estimates by field

(a) Granted Patents and Patent Term Reduction Probability

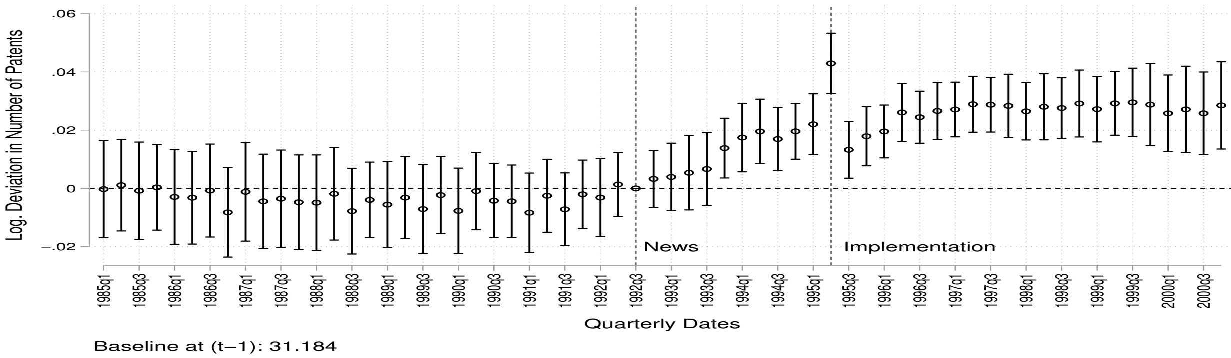

(b) Granted Patents and Average Patent Term Change

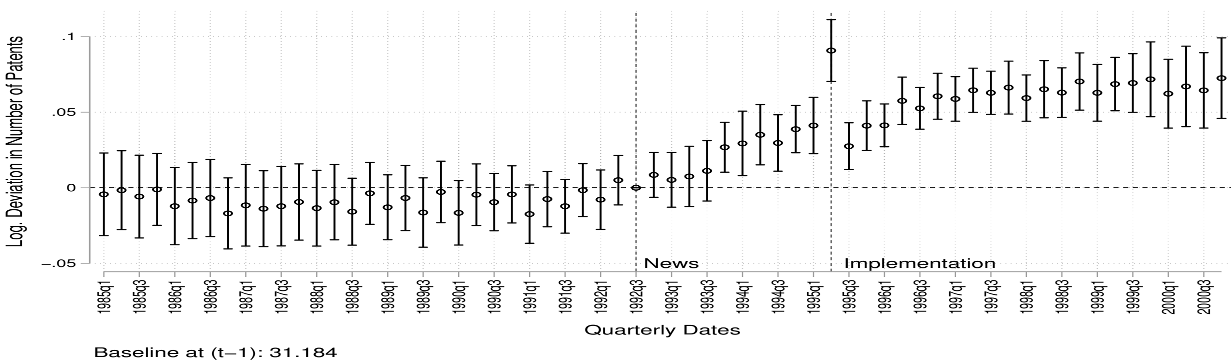

(c) 5-Year Forward Citations-Weighted Patents and Patent Term Reduction Probability

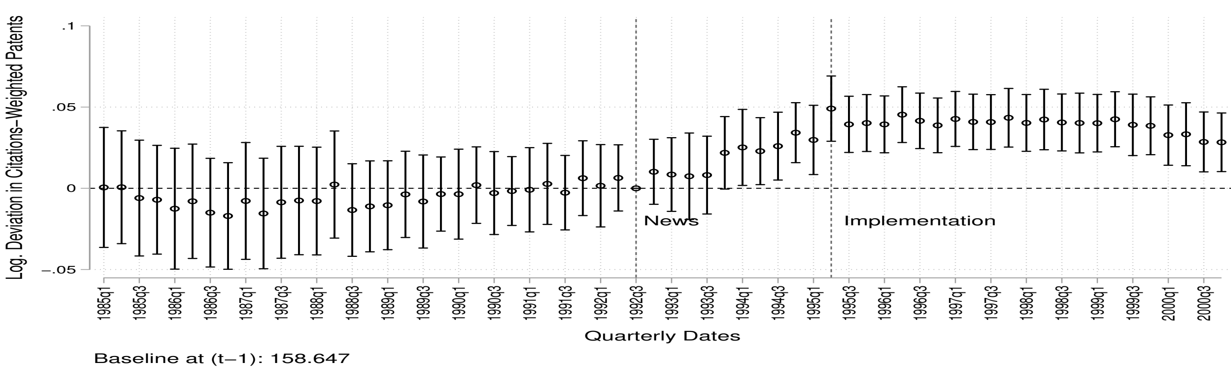

Notes: Panels (a) and (c) of the figure represent the *relative* effects of a one percentage point *higher* patent term reduction probability on log-deviations in (granted) patent applications (panel a) and citations-weighted patents (panel c) by application quarter, based on pseudo-Maximum Likelihood estimates of Poisson DiD specification (7). Panel (b) represents the *relative* effects of a one-month *shorter* extension in average patent term on log-deviations in (granted) patent applications by application quarter, based on estimates of Poisson DiD specification (7). 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 a shorter average patent term extension.

As previously noted, the persistence of a positive relationship between a shorter patent term and innovation following implementation is surprising. In the framework developed in Section 5, this persistence reflects the interaction between news-driven acceleration and technology disclosure externalities, for which I provide supporting evidence in Section 7. Importantly, Section 6.3 shows that the *direct* post-implementation effect of a higher probability of patent term reduction on innovation is in fact *negative*, consistent with prior literature.

Heterogeneity. The sensitivity of patenting to changes in patent term may vary systematically with observable field characteristics. To explore this heterogeneity, I conduct complementary triple-difference analyses. Appendix C.3.1 shows that the policy impact is larger in fields where a higher share of patents are renewed to the maximum term, suggesting greater salience of this policy tool. Appendix C.3.2 finds stronger DiD estimates in fields with higher patent litigation rates, indicative of greater reliance on patents to appropriate returns from innovation. Appendix C.3.3 documents larger effects in fields with lower uncertainty around the average patent term change. Finally, Section 7 presents additional evidence of heterogeneity based on field-level differences in technological dependence and technical competition intensity.

6.2 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.[21]

I use a direct measure of R&D expenditures from a yearly panel (1985-2000) of 2,410 listed U.S. firms, based on the NBER-Compustat dataset compiled by Hall, Jaffe and Trajtenberg (2001). For each firm i, the TRIPs-induced patent term reduction probability, PL_i, is calculated as a weighted average of field-specific probability, PL_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.[22] I collect firm-level R&D expenditures

[21] Section 6.5.1 further addresses issues related to innovation measures and secrecy. Appendix E explores how patent term changes affect sectoral Total Factor Productivity (TFP) and prices.

[22] To account for potential changes in firms' technological focus over time, I verify that all DiD results are robust to computing technological exposure based on a shorter time window, e.g., 1986-1991.

(`xrd`), sales, and other variables from COMPUSTAT (Standard&Poor's, 2022). The firm-level analysis uses the following Poisson DiD specification:

$$Y_{i,t} = \exp\left\{\alpha_i + \sum_{\substack{k=1987\\k\neq1991}}^{2000} \gamma_k \mathbf{1}_{(t=k)} + \sum_{\substack{k=1987\\k\neq1991}}^{2000} \beta_k \mathbf{1}_{(t=k)} PL_i + \theta' \mathbf{X}_{i,t} + \varepsilon_{i,t}\right\} \tag{9}$$

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 6.5. Equation (9) 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 PL_i on the log deviation (or approximate percentage change) of R&D expenditures or innovation. I estimate the regression using pseudo-Maximum Likelihood.

Figures 6a and 6b illustrate the impact of a one-percentage-point higher probability of a patent term reduction 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 6.1. During the news phase, a one-percentage-point higher probability of a patent term reduction is associated with 1.9% increase in yearly firm-level R&D (1993 estimate), equivalent to approximately \$1.7 million. Similarly, news of a one-percentage-point higher probability of a patent term reduction corresponds to 3.3% larger increase in citations-weighted patents at the firm level. The timing of these effects is also consistent with Section 6.1, as the change in citations-weighted patents is more gradual than the R&D response. These dynamics 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. 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 U.S. applicants, Appendix D.2 shows evidence of within-firm reallocation of patenting and inventors across technical fields that is consistent with field-level findings.

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

(a) R&D Expenditures

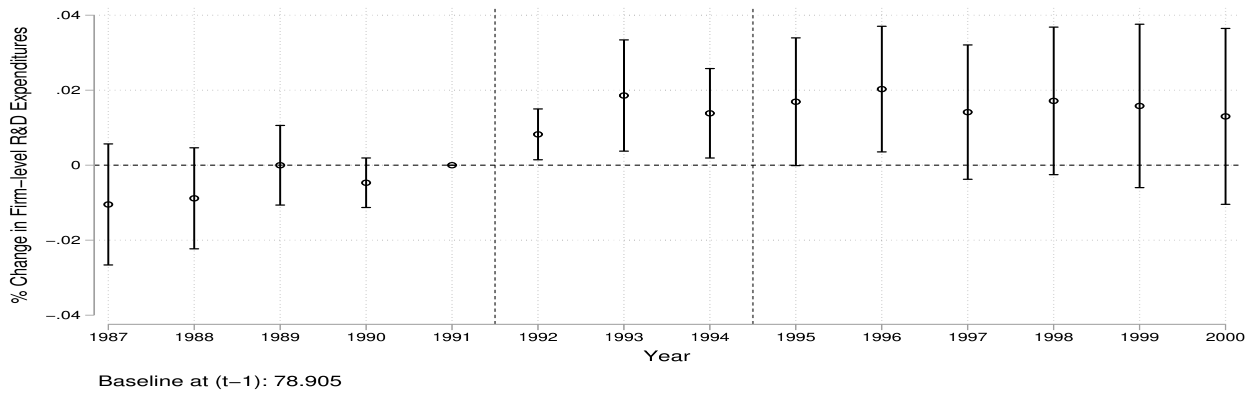

(b) Citations-Weighted Patents

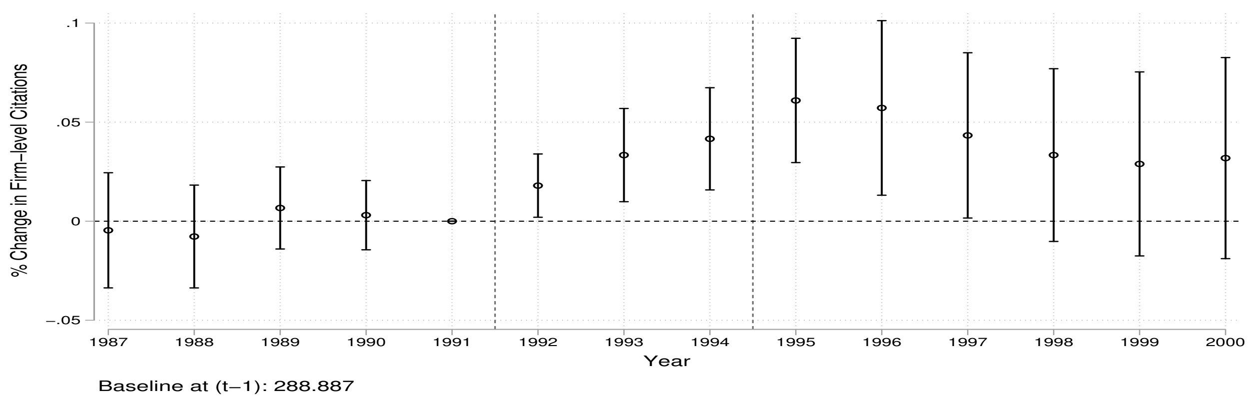

Notes: The figure depicts the effect of a one-percentage-point *higher* patent term reduction probability at the firm level—based on the Poisson DiD specification (9)—on yearly R&D expenditures (panel a) and citations-weighted (granted) patents by application year (panel b). Point estimates represent 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.

6.3 Direct Implementation Effect

I now present evidence on the *direct* effect of patent term changes on innovation. As shown in expression (5), the post-implementation DiD estimates from Sections 6.1 and 6.2 reflect both indirect effects from policy anticipation and direct effects from the actual implementation of the new patent term. The augmented DiD specification (8) 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-percentage-point higher probability of patent term reduction on the number of granted patents by field and application quarter, based on DiD estimates ($\widehat{\phi}_k$) from specification (8).

The figure highlights two main findings. First, the direct effect estimates remain positive *before* implementation, consistent with expression (4) stating that the news effect is mostly driven by the policy's direct impact. Thus, in fields with a higher probability of a patent term reduction—and thus a relatively shorter patent term extension—innovation accelerates more during the news phase (Fact 1).

Second, the direct effect estimates turn *negative* and stable *after* policy implementation. A higher fraction of patents obtaining a patent term *reduction* leads to a relative *decline* in innovation, once controlling for innovation dynamics during the news phase. On average, a one-percentage-point higher probability of patent term reduction results in 0.5 fewer patents per quarter, or about a 1.1% reduction in the average quarterly number of patents in the post-implementation period.

Using the field-specific change in average patent term as the main regressor yields similar results: a one-month *shorter* extension in average patent term results in 1.7% *lower* number of patents per quarter in the post-implementation period.

Appendix C.11.1 provides consistent evidence for quality-adjusted innovation (citations-weighted patents) by field. Furthermore, Appendix C.11.2 shows that controlling for innovation 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 higher probability of patent term reduction—i.e., a relatively *shorter* patent term—determines relatively *less* 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

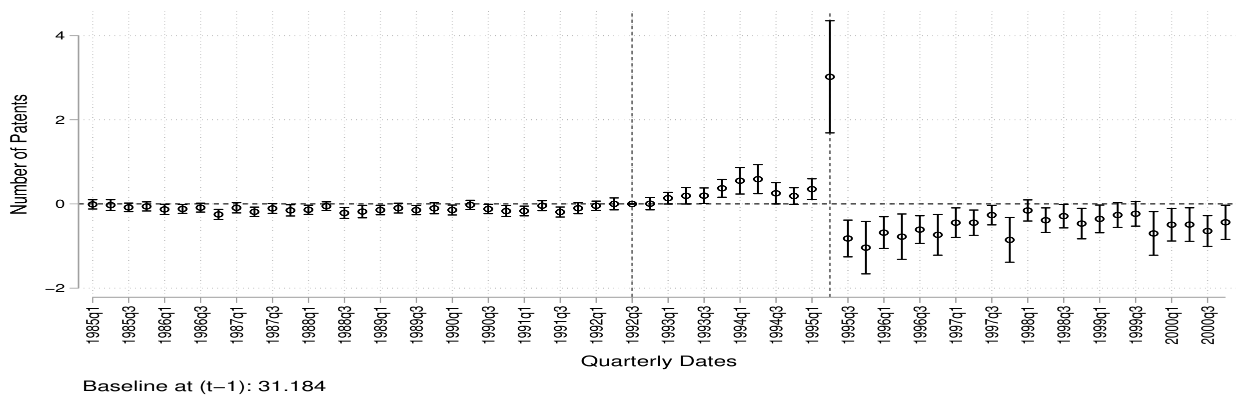

Notes: The figure depicts the relative *direct* effect of a one-percentage-point *higher* patent term reduction probability on the level of (granted) patents by application quarter, based on estimates of DiD specification (8) with $X_j = PL_j$. Bands represent 95% confidence intervals based on standard errors clustered by field and treatment phase. The first and second vertical lines mark news and implementation dates, respectively.

6.4 Elasticity and Comparison with Previous Literature

Estimates of the direct post-implementation effect indicate that a one-month longer patent term increases innovation by 1.7%, implying a semi-elasticity of 20.9% for a one-year extension. This magnitude 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, the model of Hémous et al. (2023) predicts that a one-month increase in U.S. patent terms would raise U.S. innovation by 1.2%, close to my estimate of 1.7%.

The elasticity of innovation to patent term may vary across technologies. While a full analysis of this heterogeneity is not feasible in the current setting, I assess robustness by re-estimating the semi-elasticity after excluding fields likely to be particularly responsive to patent protection. First, based on fields with *below*-median patent renewal rates to maximum term—thus less sensitive to the policy—the estimated semi-elasticity remains positive and significant at 1.1%. Second, excluding fields related to chemistry—known to be especially patent-sensitive—yields a semi-elasticity of 1.4%, also positive and significant.[23]

These results are also consistent with the elasticity of TFP to market size estimated by Beerli et al. (2020), who find 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 6.3 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 0.3%, close to the estimate of Beerli et al. (2020).

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 2.4% *increase* in patenting during the anticipation phase.[24] These findings on anticipation effects are novel to the literature and highlight the relevance of intertemporal trade-offs.

6.5 Identification and Measurement Concerns

I conclude this section by addressing key identification and measurement concerns, focusing on the main DiD specification (6). This specification provides unbiased estimates of the news and

[23] I exclude 21 4-digit fields included in 3-digit classes "Inorganic Chemistry" (C01), "Organic Chemistry" (C07), and "Organic Macromolecular Compounds" (C08).

[24] 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. The resulting figure 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.

implementation effects (4) and (5) 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.

6.5.1 Measurement of Innovation, Patent Quality, and Secrecy

Patent-based measures of innovation have known limitations: not all inventions are patented, and changes in patent term can affect both innovation incentives and the choice between patenting and secrecy. Section 6.2 addressed this concern by showing that changes in patent term also influenced firm-level R&D expenditures—an essential input to innovation—mirroring the observed effects on patenting. Appendix E further links these patent-based measures to changes in sectoral productivity and prices.

This subsection examines whether the policy affected average patent quality, which could shift due to changes in selection into patenting. The expected direction of quality changes is ambiguous and depends on the treatment phase. At news, innovators in fields expecting a patent term reduction might choose to reduce quality to accelerate filings, an effect that may be persistent. After implementation, a longer patent term could encourage patenting of higher-value inventions that are harder to imitate, increasing observed quality and novelty.[25] However, if only a few inventions hover around the patenting threshold, selection effects may be minor.

To test these hypotheses, I analyze the impact of the policy on three dimensions: (1) scientific and economic quality, proxied by average citations, generality, originality, and estimated economic value (Hall, Jaffe and Trajtenberg, 2001; Kogan et al., 2017); (2) novelty, based on new words or text similarity with prior and subsequent art (Arts, Hou and Gomez, 2021; Kelly et al., 2021); and (3) scope, measured by claim characteristics such as length and number of independent claims (Marco, Sarnoff and deGrazia, 2016). Results are presented in Appendix C.4 (quality and novelty) and C.5 (scope).

During the news phase, the policy has no significant impact on average quality, novelty, or scope. This finding suggests that firms did not file lower-quality patents when accelerating innovation in anticipation of term reductions.

After implementation, quality and novelty indicators move in different directions. Fields facing a higher probability of patent term reduction show weak increases in citations and economic value, but declines in novelty. This result suggests that innovators may have responded to a

[25] Harder-to-imitate inventions tend to be of higher quality and more novel, and face lower risk of imitation, making secrecy a more viable alternative. A longer patent term can shift this trade-off, leading innovators to patent more valuable inventions.

relatively shorter patent term by boosting quality while avoiding riskier, more novel projects. Patent scope remains unchanged, indicating no systematic shift in patenting strategy, such as narrowing claims to form patent thickets in response to weaker patent protection (Shapiro, 2001). Taken together, these results seem inconsistent with strong selection effects, which should similarly affect quality, novelty, and scope.

6.5.2 Technology Shocks and Definition of Fields

Field- and quarterly fixed effects cannot fully account for field-specific technological trends or technology demand shocks 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 (6) by adding quarter-by-3-digit IPC fixed effects, controlling for technology demand shocks or trends across over 100 3-digit technical areas. These controls ensure that the DiD estimates reflect only the variation in patent term and innovation outcomes across 4-digit fields within the same 3-digit category.[26] Appendix C.6 shows that the DiD estimates are unaffected by this adjustment. Third, I verify in Appendix C.6 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 (6).

6.5.3 Heterogeneous Effects and Selection on Doses

In DiD settings with continuous treatments (dose), such as PL_j or ΔT_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 (4) and (5). Appendix C.7 elaborates on this issue, using Callaway, Goodman-Bacon and Sant'Anna (2021) decomposition of TWFE to show that the DiD estimates in Figure 3 are constant across doses and generally stable. This result 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.

[26]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."

6.5.4 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—the DiD coefficients in (6) could be biased for the marginal responses (4) and (5).

Since ungranted U.S. patent applications were not published before 2001, assessing the impact of the patent term change on grant rates across fields is not feasible. However, Appendix C.8 complements the evidence of Section 2.4 showing that pre-news patent term reduction probability does not correlate with *systematic changes* and trends in 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.[27] Finally, I verify that DiD estimates are unaffected when using quarter-specific realized patent term reduction probability as treatment in (6) instead of its pre-news counterpart.

6.5.5 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 (6) 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.9, 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.

6.5.6 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 improving market access in LMICs, uncertainty surrounding adoption may have deterred firms from investing in innovations targeting those markets.[28] If the U.S. patent term change correlates with

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

[28] 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.

these factors across fields, specification (6) may be biased. To address this concern, Appendix C.10 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.

7 Interpretation and Transmission of News Shock

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

7.1 Interpretation of News Effects

Following policy news, innovation and R&D accelerate more in fields with higher probability of patent term reduction and shorter average extensions. In this subsection, I explore various mechanisms related to innovators attitudes toward risk that could account for these effects.

I begin by framing innovators decision during the anticipation phase as a choice between a certain outcome—the 17-year patent term guaranteed under the old regime—and a "lottery" under the new policy, where the term depends on the uncertain duration of the pending period. The risk associated with this lottery is captured by the variance of the field-specific distribution of pending periods. A stronger acceleration in innovation during this phase indicates a stronger preference for the pre-TRIPs patent term over the uncertain post-implementation outcome.

This behavior aligns with several theories of preferences under uncertainty. Under *risk aversion*, innovators prefer a certain patent term to a lottery with equal or even higher expected value. To secure the 17-year term under the old regime, they accelerate R&D and patenting before implementation, incurring two costs: higher R&D expenditures and, in most fields, an expected loss in patent term. The incentive to accelerate weakens as the expected gains from the new policy grow. Hence, innovation would accelerate more in fields with shorter expected extensions—consistent with the empirical findings.

An alternative explanation is *loss aversion*, whereby agents experience greater disutility from losses than utility from equivalent gains. Specifically, theoretical and empirical work on loss aversion suggests that agents exhibit a kink in their utility function around zero and *convex* utility over losses (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992). Under loss aversion, innovators may over-weight the risk of a patent term reduction and thus accelerate innovation following policy news to secure the longer term under the old regime. Descriptive evidence presented in Section 4.1 and historical records of the TRIPs policy debate support this interpretation.

A third mechanism is *ambiguity aversion*. Following the news shock, ambiguity may arise because the post-implementation distribution of pending periods is itself uncertain due to possible future events whose probability is hard to assess. Ambiguity-averse innovators are usually modeled as min-max decision makers, who only consider the worst-case scenario when making optimal decisions.

Crucially, these three mechanisms have distinct implications for how risk, i.e., the dispersion in pending periods, affects innovation dynamics during the news phase. Under risk aversion, greater variance in the pending period distribution (conditional on the mean) should strengthen the incentive to accelerate, as the certainty of the old regime becomes more appealing relative to a riskier lottery. In contrast, loss aversion with convex utility over losses implies that innovators may be risk-seeking in losses. In this case, higher variance might dampen the acceleration of innovation, since agents prefer more risky outcomes under a loss domain. Finally, under min-max behavior, higher downside risk should make the worst-case outcome worse, thus reinforcing the incentive to accelerate innovation.

To empirically assess these mechanisms, I enrich the main DiD specification (6) to estimate the period-specific effects of the variance of the pre-shock field-specific distribution of pending periods. Specifically, I interact the treatment variable—either patent term reduction probability or the mean of the pending periods distribution, which determines the average expected change in patent term—and the variance with dummies for each treatment period (pre-news, news, and post-implementation). This specification mirrors the main DiD specification (6) but focuses on estimating average quarterly effects within each period.

Table 4 presents the results, omitting post-implementation coefficients for brevity. The positive coefficient on patent term reduction probability during the news period ($d_{news} \times \overline{PL}_j$, column 1) confirms Fact 1: higher patent term reduction probability relates to a stronger acceleration of innovation upon news. Similarly, the positive effect of $\overline{PP}_j$ in the news period (column 2) implies that a longer average pending period—i.e., a shorter average patent term extension—induces a stronger acceleration of innovation at news.

Critically, in both columns, the *negative* effect of the variance ($d_{news} \times \sigma^2_{PP,j}$) implies that, conditional on the mean, greater dispersion in the pending period distribution—i.e., a riskier lottery—*dampens* the acceleration in innovation during the news period.

This finding provides suggestive evidence in support of loss aversion with convex utility over losses. In contrast, risk aversion would predict the opposite sign. Because the certain option becomes more attractive when the lottery is riskier, risk-averse innovators should accelerate *more* during the anticipation phase.

Table 4: DiD Effects of Different Moments of the Pending Period Distribution

	(1) Patents	(2) Patents	(3) Patents	(4) Patents
$d_{pre} \times PL_j$	-0.249		-0.259	
	(0.448)		(0.479)	
$d_{news} \times PL_j$	1.113***		1.205***	
	(0.323)		(0.347)	
$d_{pre} \times \overline{PP}_j$		-0.552		-0.544
		(0.871)		(0.914)
$d_{news} \times \overline{PP}_j$		2.010***		2.092***
		(0.614)		(0.651)
$d_{pre} \times \sigma^2_{PP,j}$	0.002	0.003		
	(0.004)	(0.005)		
$d_{news} \times \sigma^2_{PP,j}$	-0.006***	-0.007**		
	(0.002)	(0.003)		
$d_{pre} \times \sigma^2_{PP,j} \vert \Delta T_j < 0$			0.000	0.000
			(0.001)	(0.001)
$d_{news} \times \sigma^2_{PP,j} \vert \Delta T_j < 0$			-0.001**	-0.001*
			(0.001)	(0.001)
$d_{pre} \times \sigma^2_{PP,j} \vert \Delta T_j > 0$			-0.025	-0.007
			(0.074)	(0.082)
$d_{news} \times \sigma^2_{PP,j} \vert \Delta T_j > 0$			0.053	0.014
			(0.047)	(0.054)
Observations	39168	39168	35648	35648

Notes: The table reports the effect of key moments of the within-field pending period distribution on quarterly (granted) patent application by treatment period, based on DiD specification (6) estimated on a quarterly sample 1985Q1-2000Q4. Column (1) shows the effects of patent term reduction probability PL_j and the variance $\sigma^2_{PP,j}$. Column (2) reports the effects of the mean $\overline{PP}_j$ and the variance $\sigma^2_{PP,j}$. Columns (3) and (4) separately estimate the effect of the variance conditional on a positive and a negative patent term change ($\sigma^2_{PP,j}|\Delta T_j > 0$ and $\sigma^2_{PP,j}|\Delta T_j > 0$, respectively). Coefficients for the post-implementation period are not reported due to space constraints. The specification always includes field and quarter fixed effects. Standard errors are clustered by 4-digit technical field and treatment period. Statistical significance levels: * ($p < 0.1$), ** ($p < 0.05$), *** ($p < 0.01$).

To further explore the possible presence of loss aversion, I separately estimate the effects of downside and upside risk. Columns (3) and (4) of Table 4 replace the overall variance with two conditional variances: one for the part of the pending period distribution implying a patent term reduction ($\sigma^2_{PP,j}|\Delta T_j < 0$), the other for extensions ($\sigma^2_{PP,j}|\Delta T_j > 0$). The former proxies downside risk, the latter upside risk.

Only downside risk has a statistically significant negative effect, again consistent with loss aversion and convex utility over losses. In contrast, ambiguity aversion with min-max behavior lacks comparable support in the data, because it would predict that higher downside risk fosters the pre-implementation acceleration of innovation. Moreover, the effect of upside risk is not significant, aligning with the standard assumption of risk neutrality of firms over gains.

Lastly, the salience of patent term reduction probability and downside risk may have been reinforced by the potential anticipation of a clause in the URAA granting term extensions to patents still in force at TRIPs implementation, if the new rules implied a later expiry date. With anticipation of such provision, delaying patent applications after implementation would entail negative outcomes only, enhancing the incentives to accelerate innovation. Nonetheless, historical sources suggest that this detail was not anticipated.[29]

Beyond behavioral drivers, these findings on news effects contribute to a broader literature on the responsiveness of R&D and innovation to temporary shifts in innovation value. Theoretical work shows that R&D investment is procyclical, responding to fluctuations in the market value of innovation. Empirical studies similarly document procyclical R&D and patenting behavior. In this context, TRIPs policy news acted as a temporary shock, shifting the relative value of innovation across time, and inducing the observed short-run increase in R&D and patenting activity.

7.2 Transmission Channels

The empirical analysis also shows that news effects persist after implementation (Fact 2). Specifically, in fields with higher patent term reduction probability—where innovation accelerated the most following the news—patenting activity remains sustained for at least five years after implementation. This persistence reflects indirect effects that outweigh the direct impact of the policy, which would otherwise depress innovation in fields facing shorter patent protection (Fact 3). The result aligns with endogenous growth theory, where knowledge creation is cumulative: new innovations build on past advances through a "standing on the shoulders of giants" effect (Romer, 1990). Temporary shocks, like those from TRIPs news, can thus have lasting effects, consistent with medium-term business cycle models where short-run R&D shocks produce persistent productivity and output effects (e.g., Comin and Gertler, 2006).

In particular, *recent* discoveries are especially valuable for follow-on innovation. Hegde, Herkenhoff and Zhu (2023) show that a permanent reduction in patent publication time leads to stable increase in follow-on innovation, as faster knowledge diffusion improves the efficiency of subsequent R&D. A similar mechanism may temporarily affect post-implementation innovation in my setting, because the news shock accelerates diffusion during the anticipation phase.

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 insights of the innovation and endogenous growth literatures, I first focus on knowledge externalities. I then explore competition in the technology space as an alternative channel in Appendix F.2.

[29] Historical records from the U.S. policy debate—e.g., the hearings of the Senate Advisory Committee on Patent Law Reform—never mention that the new rules would apply retroactively. Moreover, other countries, such as Canada, did not adopt retroactive provisions, which triggered a WTO dispute with the U.S. in 1999.

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.[30] Since most knowledge externalities occur within the same technology (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.11.2). Thus, I proxy technological dependence by the average number of backward citations made by patents in field j to earlier patents in the same field, published within three years of the citing patent 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 probability of patent term reduction. To test this hypothesis, I estimate the triple-difference specification (10), which interacts within-field technological dependence $\overline{BB}_{jj}$—measured before the news shock—with patent term reduction probability PL_j.[31]

$$\begin{aligned} Y_{j,t} = \alpha_j &+ \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \gamma_k \mathbf{1}_{(t=k)} + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \eta_k \mathbf{1}_{(t=k)} \overline{B}_{jj} \\ &+ \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \beta_k \mathbf{1}_{(t=k)} PL_j + \sum_{\substack{k=1985Q1 \\ k\neq 1992Q3}}^{2000Q4} \theta_k \mathbf{1}_{(t=k)} PL_j \times \overline{B}_{jj} + \varepsilon_{j,t} \end{aligned} \tag{10}$$

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 PL_j with quarterly dummies capture the relative effect of one-percentage-point higher patent term reduction probability 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 PL_j varies with the level of technological dependence, $\overline{B}_{jj}$.

Figure 8a shows the triple-difference effects of a one-percentage-point *higher* patent term reduction probability, 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-

[30] 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 the sample period of this analysis, despite limitations highlighted for more recent years.

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

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 the news of a higher probability of a reduction in patent duration (Fact 1). This finding aligns with theoretical implications of endogenous growth models.

Furthermore, Figure 8b shows that, in the absence of technological dependence (i.e., conditional on $\overline{BB}_{jj} = 0$), a one-percentage-point *higher* patent term reduction probability results in *lower* post-implementation innovation. This result is consistent with the sign of the direct effect of patent term after policy implementation (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 (4).

In a second analysis, I examine how within-field technological dependence endogenously responds to the policy change. If post-implementation persistence of high innovation is driven by stronger knowledge diffusion, patents filed after implementation should increasingly cite recent inventions. Thus, proxies of technological dependence should rise more in fields with higher patent term reduction probability, where patents accelerated the most in the news phase.

To test this hypothesis, I estimate DiD specification (6) 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 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-percentage-point *higher* probability of patent term reduction on $\overline{B}_{jj,t}$ and $S_{jj,t}$, respectively. The treatment 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 higher patent term reduction probability.

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—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 estimates of a 1.5-year gestation period by Pakes and Schankerman (1984).

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 result supports the interpretation of these effects as the result of

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

(a) Triple Difference Effects

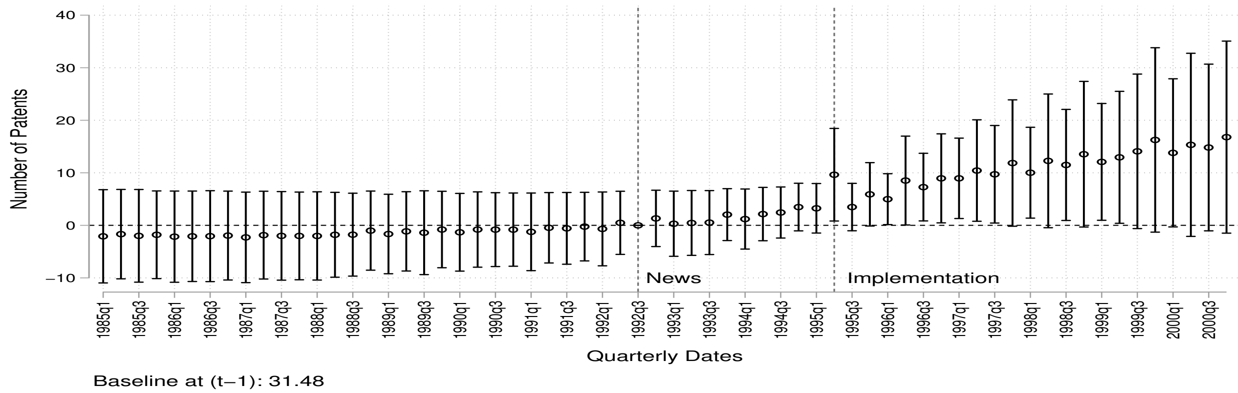

(b) Difference-in-Difference Effects conditional on $\overline{BB}_{jj} = 0$

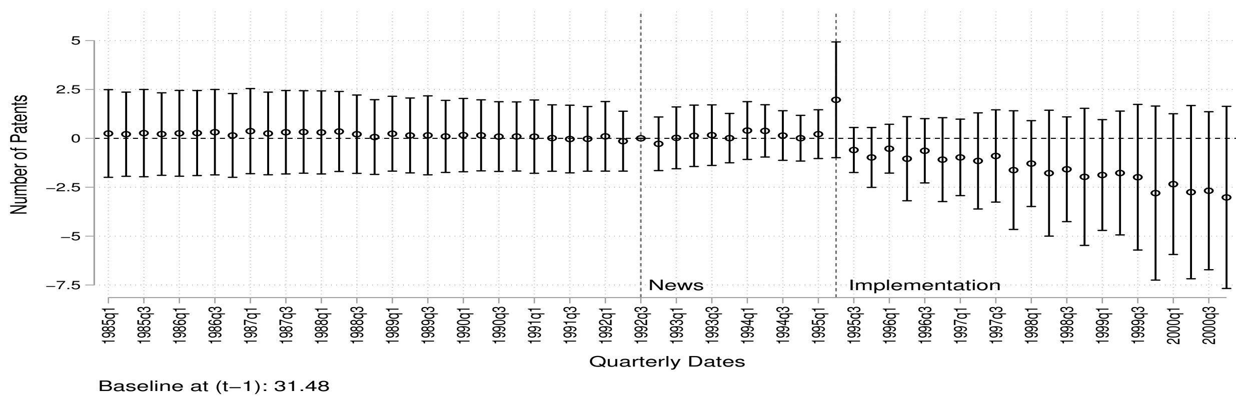

Notes: The figure depicts triple-difference (panel a) and difference-in-difference (panel b) effects of patent term reduction probability on the number (granted) patent applications by application quarter based on estimates of triple-difference specification (10). In panel (a), coefficients represents the relative change in the level of the outcome variable due to a one-percentage-point *higher* patent term reduction probability 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-percentage-point *higher* patent term reduction probability 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.

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 (8) 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

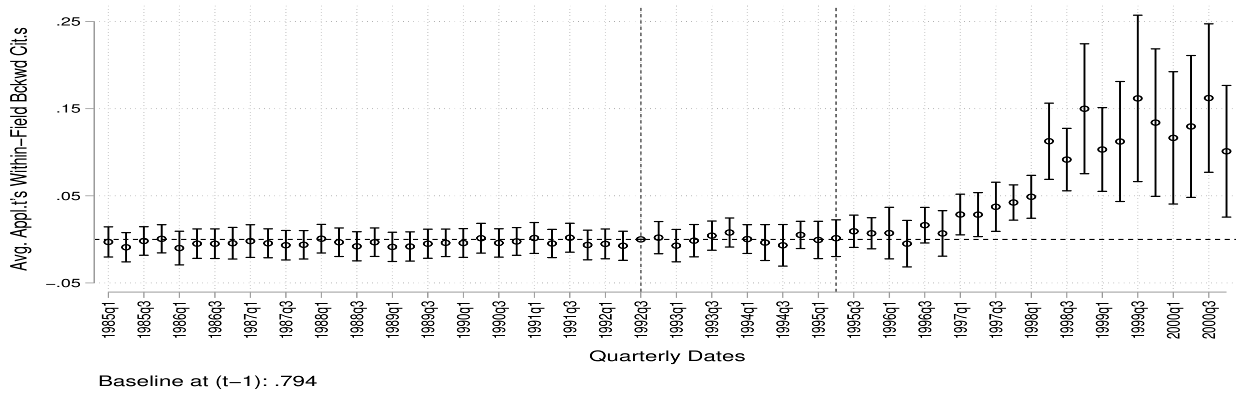

(b) Share of Patents with Applicant-made Within-field Backward Citation

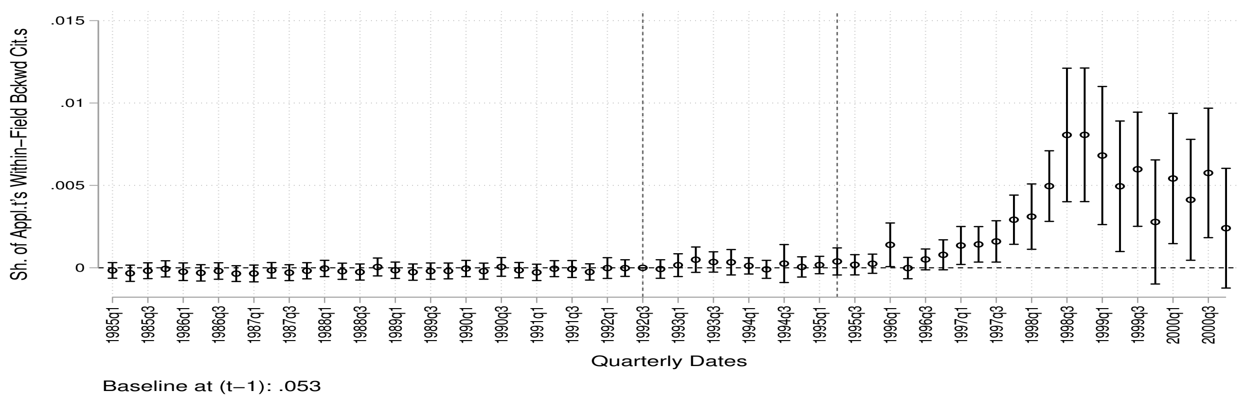

Notes: The figure illustrates the relative effect of a one-percentage-point *higher* patent term reduction probability on the level of the average number of applicant-made backward citations per patent classified in field j and applied for in quarter t (panel a) and on the level of the share of patents 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 (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.

7.3 Alternative Channels

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

Regarding competition in the technological space, the patent term change has a greater impact in fields where competition in patenting is higher, as measured by the concentration of patents across applicants or the entry rate of new applicants. When technological competitive pressure is higher, patent protection becomes more valuable and existing knowledge more accessible to new entrants, making innovators more responsive to changes in patent duration.

However, the patent term change had no direct effect on technological competition metrics, suggesting that shifts in the latter 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.

8 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 higher probability of patent term reduction—i.e., a relatively *shorter* average patent term—leads to relatively *less* innovation as a direct effect after policy implementation, with magnitudes consistent with prior literature on the impact of patent protection 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 higher future 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.

Anticipated innovation policy shocks can create short-term dynamics that go in the opposite direction of long-run effects, with larger magnitudes than other shocks due to their specific targeting of R&D and innovation. To accurately assess the positive and normative 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.

Future research can leverage the empirical estimates presented in this paper to discipline key parameters of the innovation production function in structural models of endogenous growth. These models can, in turn, inform normative analyses of optimal patent term in environments with forward-looking agents and policy anticipation.

Data Availability Statement

The data and code underlying this research are available on Zenodo at https://doi.org/10.5281/zenodo.17513718.

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, Douglas Hanley, and Nicolas Serrano-Velarde. 2020. "Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth." *The Review of Economic Studies*, 88(1): 1–43.

Akcigit, Ufuk, Murat Celik, and Jeremy Greenwood. 2016. "Buy, Keep, or Sell: Economic Growth and the Market for Ideas." *Econometrica*, 84: 943–984.

Akcigit, Ufuk, Santiago Caicedo, Ernest Miguelez, Stefanie Stantcheva, and Valerio Sterzi. 2018. "Dancing with the Stars: Innovation Through Interactions." National Bureau of Economic Research, Inc NBER Working Papers 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.

Hegde, Deepak, Kyle Herkenhoff, and Chenqi Zhu. 2023. "Patent Publication and Innovation." *Journal of Political Economy*, 131(7): 1845 – 1903.

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.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2): 263–292.

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.

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. 2021. "Patent Screening, Innovation, and Welfare." *The Review of Economic Studies*, 89(4): 2101–2148.

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."

Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, 5(4): 297–323.

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.