Patenting and Information Disclosure*

Xizhao Wang†

Abstract

Invention disclosure facilitates knowledge spillovers, supporting future progress but potentially limiting appropriability for the inventor. In this paper, I examine invention disclosure behavior by analyzing the readability of patent texts, using both traditional and novel AI-based readability scores. Using two difference-in-differences analyses, I find that following the 1980 Bayh-Dole Act and the staggered establishment of Technology Transfer Offices, two institutional changes that strengthened commercialization incentives in university innovation, university patents became less readable in their detailed descriptions. This decrease in readability does not extend to patent summary texts, suggesting that university inventors selectively obfuscate technical details on how to make and use the invention. The effect is intensified in technological fields with higher competitive pressure and is associated with a decline in external follow-on citations. These findings suggest that disclosure choices arise not only in the decision of whether to patent or rely on trade secrecy, but also in the degree of information disclosure within patents. Institutional shifts toward commercialization may weaken the patent system's disclosure function, with patterns consistent with reduced downstream knowledge diffusion.

*I am deeply grateful to my advisors, Benjamin F. Jones, Dimitris Papanikolaou, Dashun Wang, and Ryan Hill, for their invaluable guidance and support. I thank Jason Bell, El Hadi Caoui, Abhijit Sengupta, and Xuan Teng for their insightful discussions as conference discussants. I thank seminar participants at Columbia University, Georgia Institute of Technology, National University of Singapore, Northwestern University, the University of Chicago, the University of Hong Kong, and the University of Toronto, as well as participants at various conferences, for their valuable feedback. All errors are my own.

†Email: xizhao.wang@kellogg.northwestern.edu.

1 Introduction

The concept of "standing on the shoulders of giants" emphasizes that innovation builds upon the collective knowledge and achievements of previous generations. Extensive literature (Romer 1990, Scotchmer 1991, Kortum 1993, Jones 1995a, Weitzman 1998, Jones 1995b, Young 1998, Cockburn and Henderson 1996, Jones 2009, Jones 2010, Hall and Harhoff 2012, Akcigit et al. 2018) underscores the idea that knowledge gives rise to further knowledge. This principle aligns closely with the idea that knowledge disclosure – the open sharing of ideas and discoveries – generates broad societal benefits. Such disclosure is crucial in technological and scientific fields, where each breakthrough lays the groundwork for future advancements. Recognizing the social benefits of innovation disclosure, policymakers strive to promote the dissemination of knowledge, encouraging inventors and researchers to make their discoveries publicly accessible.[1]

From the standpoint of the inventor, however, there may be reluctance to disclose technological details due to concerns about appropriability – the ability to capture profits from their inventions. In many cases, inventors or firms may opt to keep innovation information private rather than disclosing it publicly. Previous studies have highlighted how inventors and firms utilize trade secrecy and lead time as appropriation mechanisms, allowing them to retain control over their ideas and potentially gain a competitive edge in the market (Mansfield 1986, Cohen et al. 2000, Mezzanotti and Simcoe 2023). Thus, there is a trade-off in invention disclosure between the societal perspective, which emphasizes social benefits, and the individual viewpoint of inventors or firms, who consider their private gains and appropriability conditions.

The patent system mixes both private incentives for encouraging invention while also trying to promote invention disclosure (Gallini 2002). Specifically, under the patent system, patent applicants disclose their inventions publicly in exchange for short-term monopoly rights. The U.S. Supreme Court aims to use the disclosure requirement to promote the dissemination of knowledge and further innovation.[2] To fulfill this objective, the U.S. Patent and Trademark Office (USPTO) and the U.S. Code (USC) require that a patent include a clear and comprehensive written description of the invention, along with instructions for how to make and use it. The patent's detailed description should be thorough enough to enable individuals with ordinary expertise in the relevant field, science, or domain to implement and apply the invention without the need for extensive trial and error.[3]

[1] An example is the implementation of the "Bermuda Principles" by the Human Genome Project (HGP), where both scientists and policymakers agreed that rapid release of the project's genomic data was essential for advancing scientific discovery and improving human health (see, e.g., Contreras 2011).

[2] The U.S. Supreme Court stated in *Kewanee Oil Co. v. Bicron Corp.*, 416 U.S. 470 (1974): "The information in patents adds to the general store of knowledge and is of such importance to the public weal that the Federal Government is willing to pay the high price of 17 years of exclusive use for its disclosure, which disclosure, it is assumed, will stimulate ideas and the eventual development of further significant advances in the art."

[3] Source: https://www.uspto.gov/patents/basics/apply/utility-patent. The explanation must be accessible to anyone skilled in the relevant field, allowing them to replicate the invention (as stated in 35 U.S.C. § 112(a)).

While patent applications are crafted within a regulatory environment designed to incentivize disclosure, in reality, legal and social science scholarship argues that patent documents often contain inadequate and unclear information (Roin 2005, Fromer 2008, Devlin 2009, Seymore 2009, Lemley 2011, Ouellette 2011, Sampat 2018, Kong et al. 2023). Certain legal scholars point out that, despite the requirements of 35 U.S.C. § 112(a), patent documents often fail to effectively convey the intricacies of advanced technologies (Devlin 2009). They suggest that patents can be difficult to comprehend and are often "unreadable" (Seymore 2009), with patent writers sometimes using vague or unclear language to safeguard their proprietary knowledge (Roin 2005, Fromer 2008). Additionally, it is argued that the Federal Circuit has permitted vague disclosures, so that patent writers are incentivized to draft them this way (Lemley 2011). These observations raise critical questions: Do inventors strategically disclose their inventions in patent documents, and if so, how? Can we measure this disclosure behavior? Under what circumstances are inventors more likely to disclose more or less knowledge information?

This study examines how institutional environments related to invention commercialization shape knowledge disclosure in patents. On one hand, clear patent disclosures can enhance customer acquisition, foster partnerships, support licensing opportunities, and improve access to financing by making inventions more comprehensible and appealing to investors and collaborators.[4] On the other hand, obfuscation can help inventors prevent their inventions from being easily imitated or surpassed by competitors, thereby preserving their competitive advantage.[5] This tension suggests that the extent of disclosure in patents may be a decision by inventors, shaped by their institutional environment.

To study this, the paper focuses on how shocks to the university innovation commercialization environment affect patent disclosure behavior. To the best of my knowledge, research on the degree of patent disclosure is limited. Addressing this question systematically and empirically presents two main challenges. First, measuring patent disclosure behavior is difficult. Second, we need causal identification to study the impact of the institutional environment on patent disclosure. To address the first challenge, I use textual analysis to construct readability scores to proxy the clarity and sufficiency of the patent document. To address the second, I leverage two natural experiments and employ difference-in-differences analyses to identify causal effects.

To measure the patent information disclosure, I conduct textual analysis to create readability scores for patent documents. Traditional readability measures have been widely used in finance and accounting studies to analyze financial documents (Li 2008, You and Zhang 2009, Miller

[4]Disclosure can signal the value and credibility of an invention, attracting potential investors and deterring competitors (Spence 1974, Bhattacharya and Ritter 1983, Long 2002, Scotchmer 2004, Graham and Hegde 2015, Hegde and Luo 2018).

[5]By revealing key details, disclosure exposes inventions to imitation and intensifies competition, which can diminish their commercialization potential (Bhattacharya and Ritter 1983, Horstmann et al. 1985, Mansfield 1986, Anton and Yao 1994, Cohen et al. 2000, Anton and Yao 2004, McGurk and Lu 2015, Mezzanotti and Simcoe 2023). Previous studies have argued that to maintain a competitive edge, inventors may choose not to patent at all (Horstmann et al. 1985, Mansfield 1986, Cohen et al. 2000, McGurk and Lu 2015, Mezzanotti and Simcoe 2023).

2010, Beyer et al. 2010, Loughran and McDonald 2014, Campbell et al. 2014, Loughran and McDonald 2016). Recent literature applies those measures outside financial documents. For instance, Hengel (2022) examines articles in the peer review process using traditional linguistic measures. Funk and Iriberri (2020) applies traditional readability measures to explore whether the complexity of abstracts varies by gender. Ashtor (2022), Kong et al. (2023), and Dyer et al. (2024) extend these linguistic features to construct readability measures for patent documents. For my analysis, I first follow the previous literature to construct classical readability measures typically used for technical documents.

I validate the traditional readability measures using data from patent examiners' decisions on rejections. In the patent application rejections dataset, each document-action is associated with specific action types and subtypes. The 112(a) criteria mandate that a patent's detailed description must include a comprehensive and clear account of the invention, along with instructions on how to construct and use it. I calculate traditional readability scores for the detailed descriptions of rejected applications and find that these scores can predict the likelihood of rejection based on the 112(a) criteria.

I then develop a novel AI-based readability measure using the transformer-based language model SciBERT (Beltagy et al. 2019) to analyze the detailed descriptions of rejected patents. By combining text embeddings generated by SciBERT with patent examiners' rejection decisions, I train supervised learning models to predict the probability of rejection under the 112(a) rule. These predicted probabilities serve as AI-based readability scores, offering a novel measure of text readability as interpreted by patent examiners.

To identify the causal impact of institutional environments on patent disclosure, I leverage two natural experiments. The enactment of the Bayh-Dole Act in 1980 marked a plausibly exogenous shift in the institutional environment of U.S. universities. It removed key restrictions and granted universities the right to pursue the commercialization of inventions funded by the federal government (Eisenberg 1996, Henderson et al. 1998a, Henderson et al. 1998b, Mowery and Sampat 2001, Rabitschek and Latker 2005). After the passage of the Bayh-Dole Act, there was an explosion in the growth of Technology Transfer Offices (TTOs) in U.S. universities and a substantial increase in the commercialization of university inventions and resulting revenue (Henderson et al. 1998a, Henderson et al. 1998b, Chukumba and Jensen 2005, Litan et al. 2007). TTOs introduced the professionalization of the patenting process and commercialization by providing patent attorneys and business managers to assist university inventors. The establishment of TTOs can be seen as an institutional shock for inventors at specific universities. Instead of the one-time policy shock of the Bayh-Dole Act, TTO establishments represent staggered shocks across universities.[6]

[6]These two shocks are institutional changes in university innovation and can be viewed as mutual robustness checks. The Bayh–Dole Act and the staggered rollout of university TTOs provide complementary quasi-experiments. Bayh–Dole is a one-time federal shock, though its effects may be confounded by contemporaneous events. In contrast, TTO formation varies across universities and over time, supplying cross-sectional and temporal variation. I analyze data at the inventor level to reduce organizational confounds and test parallel trends to

Taken together, using all granted patents' detailed descriptions from 1976 to 1990, I provide large-scale evidence on whether these institutional changes in the university commercialization environment positively, negatively, or insignificantly affect university patent disclosure. Specifically, I employ two difference-in-differences analyses to compare the readability change in university patents and corporate patents before and after these institutional shocks. Corporate patents, which are not directly affected by these policy changes, serve as a control group to isolate the impact of the shocks on university patents. The analyses are conducted both at the organizational level, comparing university and corporate assignees, and at the individual level, tracking changes in patent disclosure for given university and corporate inventors.

The results show a decline in the readability of university patents' detailed descriptions following the 1980 Bayh-Dole Act, relative to corporate patents. A similar reduction is observed after the establishment of Technology Transfer Offices (TTOs). Analyses at both the organizational and inventor levels show a decline in the clarity and sufficiency of detailed descriptions in university patents, while summary and background sections remain unchanged. This pattern suggests selective adjustments in disclosure following these institutional changes.

Further analyses suggest that changes in commercialization incentives may be one key mechanism driving this disclosure behavior. Evidence shows that shifts in patent characteristics, such as basic versus applied orientation, originality and generality, and overall quality, are unlikely to fully account for the decline in readability. The results are also difficult to reconcile with a uniform shift in legal drafting: while readability declines in detailed descriptions, I find little comparable change in the claims section, which is more directly tied to legal scope. By separately examining the timing of the Bayh-Dole Act and the establishment of Technology Transfer Offices (TTOs), and by disentangling the role of legal expertise from inventors, the results are consistent with selective adjustments in disclosure. Heterogeneous treatment effects across technological fields reveal that university inventors are more likely to reduce readability in areas with higher competitive pressure.[7] Together, these results highlight a broader response in innovation disclosure: when inventors anticipate commercialization potential, they may reduce the readability of patent descriptions to limit imitation, thereby diminishing openness in sharing how to make and use the invention.

I also analyze citation patterns to better understand knowledge spillovers. The results reveal a significant decline in external citations to university patents in the treatment group following the Bayh-Dole Act and the establishment of TTOs, relative to a control group. Using the Bayh-Dole Act as an instrument for readability, I interpret the 2SLS results as descriptive evidence rather than causal due to possible violations of the exclusion restriction. Nevertheless, the association between reduced patent readability and lower external citations suggests that diminished disclosure clarity may have contributed to limited knowledge diffusion, alongside other

assess the DID assumption, acknowledging residual endogeneity (e.g., influential inventors affecting TTO timing). Together, Bayh–Dole and TTO variation strengthen identification.

[7] This interpretation is consistent with broader evidence that weaker intellectual property protection can intensify competitive pressure and shift firms toward alternative appropriability strategies (e.g., Acikalin et al. 2023)

potential mechanisms. These findings suggest important welfare implications: the shift in the university innovation commercialization landscape brought about by the Bayh-Dole Act and the rise of TTOs has reduced the clarity and sufficiency of patent descriptions and diminished knowledge spillovers. Policies designed to encourage commercialization might thus generate unintended consequences for transparency and knowledge diffusion. This potential decline in how openly universities share their scientific discoveries should be carefully considered when shaping future innovation policy.

The rest of the paper is organized as follows. Section 2 reviews the related literature. Section 3 develops the conceptual framework. Section 4 describes the data, measures, and baseline sample. Section 5 outlines the empirical strategy and presents the main results, examining how institutional changes affect patent disclosure. Section 6 explores underlying mechanisms and discusses effects on follow-on innovation. Section 7 concludes.

2 Prior Literature

Previous research on invention disclosure has discussed the strategies of inventors and firms, particularly their decisions on whether to file a patent, the timing between patent filing and licensing or data sharing practices, the allocation of rights, and the associated costs and benefits of patenting (Horstmann et al. 1985, Mansfield 1986, Aghion and Tirole 1994, Green and Scotchmer 1995, Cohen et al. 2000, Long 2002, Landes et al. 2003, Scotchmer 2004, Grushcow 2004, Baker and Mezzetti 2005, Gans and Murray 2011, Graham and Hegde 2015, Gans et al. 2017, Hegde and Luo 2018, McGurk and Lu 2015, Mezzanotti and Simcoe 2023). This paper goes beyond the binary decision of whether to patent to examine the degree of patent clarity and sufficiency. The issue of inadequacy and lack of clarity in patent documents has been discussed in prior literature (Roin 2005, Fromer 2008, Devlin 2009, Seymore 2009, Lemley 2011, Ouellette 2011, Sampat 2018), though systematic empirical studies remain limited (Kong et al. 2023). This paper extends prior work by introducing novel AI-based readability scores and analyzing the causal impact of institutional changes on patent disclosure.

This paper further contributes to the literature on university commercialization policy. Prior studies have analyzed patenting rates, entrepreneurial formation, and invention commercialization following the Bayh–Dole Act and university TTOs (Henderson et al. 1998a, Henderson et al. 1998b, Jaffe 2000, Mowery and Ziedonis 2002, Thursby and Thursby 2002, Jensen et al. 2003, Mowery et al. 2001, Mowery and Sampat 2004, Rothaermel and Thursby 2005, Rothaermel et al. 2007).[8] Shane (2004) finds that the Bayh-Dole Act provided incentives for universities to shift patenting in lines of business where licensing is effective. Hvide and Jones

[8] Henderson et al. (1998a), Henderson et al. (1998b), Jaffe (2000) show that university patenting increased during this period while others (Mowery and Ziedonis 2002, Mowery et al. 2001, and Mowery and Sampat 2004) argue that the rise of university patenting cannot be credited entirely to the Act, as university patenting appears to have been rising before the passage of the Act.

(2018) analyzes the end of the "professor's privilege" in Norway and finds a decline in both entrepreneurship and patenting rates after this reform. Hsu and Kuhn (2023) examines that technology transfer offices (TTOs) allocate resources for patent applications and finds that inventors' licensing experience is a strong predictor of patent impact and commercialization success. Lerner et al. (2024) examine how the commercial impact of academic research changes when scholars move across universities. While previous studies have centered on patent volume, citation, and entrepreneurship, less attention has been paid to patent knowledge disclosure. This paper addresses this gap by incorporating patent disclosure as a critical factor in evaluating the impact of university commercialization policy.

This paper also relates to the literature on text readability and obfuscation. In financial reporting, firms often make annual reports harder to read to conceal poor performance (Li 2008, Bushee et al. 2018). In loan approvals, strategic verbosity is employed (Larrimore et al. 2011). Academia and scientific grants also see obfuscation, with complex language and selective reporting used to mask truth or exaggerate proposals (Nosek et al. 2012, Head et al. 2015, Markowitz 2019). In the patent domain, inventors may engage in partial disclosure (Anton and Yao 2002, Anton and Yao 2004). Additionally, obfuscation tactics include cluttered patent claims (Freilich 2017), irrelevant prior-art citations (Kuhn et al. 2020), and strategies to hinder examiners' assessments (Figueroa and Lemus 2024). This paper empirically examines obfuscation in patents. Patents serve as strategic tools for protecting, communicating, and at times concealing technological knowledge. Understanding what information firms and inventors choose to disclose or to obscure reveals how they manage the trade-off between appropriability and knowledge diffusion. My work sheds light on this process by analyzing disclosure behavior as a choice shaped by institutional and policy regimes.

This paper expands on readability measures derived from textual analysis in related literature. In addition to using traditional readability measures that have been utilized in previous literature (Li 2008, You and Zhang 2009, Miller 2010, Beyer et al. 2010, Loughran and McDonald 2014, Campbell et al. 2014, Loughran and McDonald 2016, Hengel 2022, Funk and Iriberri 2020, Ashtor 2022, Kong et al. 2023, Dyer et al. 2024, Bearson and Roche 2025), I also use rejection information from patent examiner decisions to construct AI-based readability scores. Following recent papers on transformer-based language models (Devlin et al. 2018, Beltagy et al. 2019) and supervised learning models to construct large-scale text-based measures from unstructured documents (Miric et al. 2023), I use the transformer-based language model SciBERT, combined with supervised learning models, to construct AI-based readability scores.

More broadly, this paper speaks to a large literature on knowledge spillovers and cumulative innovation, which shows that technological knowledge can spill across firms, that patent rights and disclosure rules can shape downstream research and the diffusion of such knowledge, and that access to existing knowledge can itself affect subsequent scientific production (Jaffe et al. 1993, Bloom et al. 2013, Williams 2013, Galasso and Schankerman 2015, Biasi

and Moser 2021, Hegde et al. 2023). Existing evidence, much of it from biomedical settings, is mixed (Sampat 2018): some studies find limited barriers from patents, while others show that patenting, patent invalidation, disclosure, and ownership rules can materially affect downstream research and cumulative innovation. Walsh et al. (2007) found that academic researchers rarely felt constrained by existing patents, though commercially oriented researchers reported more barriers. Murray and Stern (2007) and Huang and Murray (2009) observed that patenting and patent strategy can impede the open diffusion and long-run supply of public scientific knowledge. Sampat and Williams (2019) found that gene patents do not significantly impede follow-on innovation, whereas Williams (2013) earlier showed that private ownership of genomic data caused notable declines in related research. Other work shows that patent invalidation increases subsequent citations to focal patents, that access to information disclosed in patents facilitates cumulative innovation, and that stronger patent protection can in some settings stimulate follow-on corporate scientific research (Galasso and Schankerman 2015, Furman et al. 2021, Arora et al. 2024). Relatedly, work on mandatory disclosure in the patent system shows that forcing firms to reveal proprietary information can be privately costly but socially beneficial (Kim and Valentine 2021). This paper contributes to this body of work by providing evidence on how follow-on innovation responds to changes in the innovation commercialization environment.

3 Conceptual Framework

3.1 Institutional Settings

3.1.1 The 1980 Bayh-Dole Act

In this paper, I first use the 1980 Bayh-Dole Act as an exogenous shock to universities. Before 1980, the federal government held the right to claim royalties and income from patents when such innovations were supported by federal funding. While university inventions could be assigned as patents to the universities themselves, commercialization faced major institutional constraints. These included mandates for non-exclusive licensing, restrictions on granting exclusive licenses, and government oversight that allowed federal agencies to revoke commercialization rights if certain conditions were not met (Eisenberg 1996, Henderson et al. 1998a, Henderson et al. 1998b, Mowery and Sampat 2001, Rabitschek and Latker 2005).[9]

As a result, many university patents before 1980 were pursued primarily for prestige and academic recognition, rather than for commercial value (Henderson et al. 1998a, Henderson

[9]For instance, under the 1963 Kennedy Memorandum, contractors were typically permitted to retain title to federally funded inventions, but this was subject to a non-exclusive, royalty-free license granted to the government. If contractors held rights beyond this non-exclusive license, they were required to submit regular reports on their commercialization uses. Moreover, the government retained "march-in" rights, allowing it to revoke exclusivity (Eisenberg 1996).

et al. 1998b).[10]

The passage of the Bayh-Dole Act in 1980 fundamentally altered this landscape. The Act granted universities rights to retain title to inventions derived from federally funded research, effectively allowing them to commercialize them—including through exclusive agreements—and collect royalties and incomes with far fewer bureaucratic hurdles (Henderson et al. 1998a, Mowery and Sampat 2001, Eisenberg 1996). Moreover, Bayh-Dole required that universities share patent-related income with the inventors, stipulating in 35 U.S.C. § 202(c)(7) that inventors must receive a share of royalties. While the law does not mandate a fixed distribution formula, universities typically follow a standardized practice of splitting net revenues among the inventors, their departments, and the institution after deducting administrative costs (Ouellette and Tutt 2020, Contreras 2022).

The shift in the university commercialization landscape brought about by the Bayh-Dole Act provides a valuable opportunity to examine how institutional systems influence the disclosure behavior of institutions and inventors. Assessing the impact of this shift offers insights not only into the specific policy change but also into how institutional frameworks shape behavior in the broader context of technological advancement.

3.1.2 University Technology Transfer Offices

I next exploit the establishment of university Technology Transfer Offices (TTOs) as a second institutional shock affecting inventors within universities. Following the passage of the Bayh-Dole Act, which enabled universities to commercialize their federally funded inventions, U.S. universities increasingly established TTOs to manage and commercialize their intellectual property. These offices were viewed as vehicles for generating revenue and facilitating the transfer of innovation to the market (Orozco 2019, Contreras 2022).

TTOs are typically staffed by professionals such as attorneys and business managers with expertise in patents, licensing, and commercialization processes (Contreras 2022). Rather than pursuing decentralized or discipline-specific commercialization pathways, many universities adopted a centralized model, relying on TTOs as the primary channel for disseminating innovation (Litan et al. 2007). These offices play a key role in promoting university patent commercialization by transferring usage rights to businesses, entrepreneurs, and startups, and by providing legal and managerial support for licensing activities (Henderson et al. 1998a, Henderson et al. 1998b).[11] The establishment of TTOs thus marks a shift toward professionalized, institutionally driven patenting and commercialization processes within universities.

[10] This posed a significant barrier to widespread university patenting, particularly considering that approximately 70% of university research was federally funded. As a result, patents granted to universities during this time fell into one of three categories: patents representing the outcomes of research that did not rely on federal funding; patents pursued primarily for prestige and recognition rather than for commercial potential; and patents stemming from federal research that required a specific title rights waiver from the funding agency (Henderson et al. 1998a, Henderson et al. 1998b).

[11] Source: https://patents.princeton.edu/about-us/what-technology-transfer

3.2 Theoretical Hypotheses

University inventors face competing incentives that shape their disclosure decisions. Following prior literature (Merton 1973, Stern 2004, Sauermann and Roach 2012, Lam 2011, Gans et al. 2017, Ouellette and Tutt 2020, Burrage et al. 2024), I consider two primary returns for university inventors in patenting: (1) **financial returns** and (2) **broader returns**. Broader returns include reputational gains through recognition of novel ideas, which can influence career trajectories, and social impact by facilitating the adoption of their technology, advancing the public good, and validating their research. Financial returns, on the other hand, stem from income received through commercialization. While many studies suggest that university researchers place relatively large emphasis on broader returns (Merton 1973, Stern 2004, Evans and Leighton 1989, Hamilton 2000, Shane et al. 2003, Sauermann and Roach 2012, Gans et al. 2017, others argue that financial incentives do matter to university inventors (Lam 2011, Ouellette and Tutt 2020, Burrage et al. 2024).

It is natural to think that broader returns increase with disclosure, as greater openness fosters recognition and societal contributions (Merton 1973, Partha and David 1994, Shane et al. 2003, Stern 2004, Murray and Stern 2007, Evans and Leighton 1989, Hamilton 2000). The impact of disclosure on financial returns, however, remains unclear. Two contrasting arguments might emerge from the literature. First, financial returns may increase with disclosure. Disclosure can signal the value and credibility of an invention, securing financing on terms that reflect its innovation potential, attracting potential investors and deterring competitors (Spence 1974, Bhattacharya and Ritter 1983, Long 2002, Scotchmer 2004, Graham and Hegde 2015, Hegde and Luo 2018). Under this "signaling" view, greater transparency enhances financial returns, incentivizing inventors to disclose more. Second, financial returns may decrease with disclosure. By revealing key details, disclosure exposes inventions to imitation and intensifies competition, which can diminish their commercialization potential (Bhattacharya and Ritter 1983, Horstmann et al. 1985, Mansfield 1986, Cohen et al. 2000, Anton and Yao 1994, Anton and Yao 2004, McGurk and Lu 2015, Mezzanotti and Simcoe 2023). In this "imitation" view, higher disclosure reduces financial returns, prompting inventors to limit disclosure to safeguard their invention's competitive advantage.

To formalize the disclosure trade-off in a simple way, consider a university inventor who selects a disclosure level d, where $d \in \mathbb{R}$ and higher values represent greater levels of disclosure.

Let the inventor's objective function be:

$$U(d;\theta) = (1-\theta)B(d) + \theta\pi(d) - C(d),$$

where $B(d)$ captures the broader returns to disclosure, such as contributing to public knowledge and scientific progress. The broader returns increase when d increases, so $B'(d) > 0$. The term $\pi(d)$ represents the private financial returns to disclosure. $\pi'(d)$ depends on contextual factors: greater clarity and sufficiency can increase value through signaling or reduce value by

inviting imitation. Assume both $B(d)$ and $\pi(d)$ are concave ($B''(d) < 0$ and $\pi''(d) < 0$). The cost function $C(d)$ captures the non-financial cost of disclosure, and is assumed to be increasing and convex: $C'(d) > 0, C''(d) \geq 0$.

The parameter $\theta \in [0,1]$ denotes the institutional emphasis on commercialization. A low θ corresponds to environments with weak commercialization infrastructure, such as the pre-Bayh-Dole era, where university inventors derive most value from broader returns. A high θ, in contrast, reflects institutional settings with strong commercialization incentives, such as post-Bayh-Dole reforms or the establishment of university technology transfer offices. Thus, θ can be interpreted as the weight placed on financial versus broader returns.

The first-order condition for the optimum d^* is:

$$\frac{\partial U}{\partial d} = (1-\theta)B'(d^*) + \theta\pi'(d^*) - C'(d^*) = 0. \tag{1}$$

Differentiate (1) with respect to θ:

$$\frac{\partial d^*}{\partial \theta} = \frac{B'(d^*) - \pi'(d^*)}{(1-\theta)B''(d^*) + \theta\pi''(d^*) - C''(d^*)}.$$

Because the denominator is negative under standard concavity assumptions, the sign of $\frac{\partial d^*}{\partial \theta}$ depends on the difference in marginal returns. This gives rise to three distinct theoretical cases:

Case 1 — Strategic obfuscation. If $B'(d^*) - \pi'(d^*) > 0$, then $\frac{\partial d^*}{\partial \theta} < 0$. This implies that when institutions place greater emphasis on financial incentives, inventors reduce their level of disclosure, provided that the marginal broader returns to disclosure exceed the marginal financial returns. A plausible scenario is when the marginal financial benefit of disclosure is relatively weak or even negative – for instance, when imitation concerns outweigh the signaling value of clarity and sufficiency. As a result, inventors may strategically obscure technical details in their patents when they anticipate a higher likelihood of commercializing their invention.

Case 2 — Enhanced clarity. If $\pi'(d^*) - B'(d^*) > 0$, then $\frac{\partial d^*}{\partial \theta} > 0$. Here, an increase in commercialization support leads to greater disclosure. This inequality implies that, at the margin, the financial returns to disclosure exceed the marginal broader returns. This can occur when in the financial return the signaling value of clear patents is strong, and the risks of imitation are comparatively limited. In such environments, inventors benefit from making their patents more accessible and transparent, increasing the readability of descriptions to signal and communicate value in the market.

Case 3 — Neutral effect. If $B'(d^*) = \pi'(d^*)$, then $\frac{\partial d^*}{\partial \theta} = 0$. This case implies that commercialization incentives have no effect on the optimal level of disclosure when, at the margin, the financial returns from disclosure equal the broader returns.

Integrating these theoretical perspectives, the effects of changes in the university innovation commercialization environment on university-based innovation are ambiguous. Prior to the Bayh-Dole Act, university inventors were more likely to achieve reputational and social impacts but had limited opportunities to benefit financially from patents. The Bayh-Dole Act and TTO establishment, which represented a significant shift by enabling university inventors for commercialization, introduced stronger financial incentives into patenting activities. This shift in the university innovation commercialization landscape provides a valuable opportunity to examine how institutional systems influence the disclosure behavior of inventors. The actual effects remain an empirical question, which I address in the empirical sections.

4 Data and Measures

4.1 Patent Data

4.1.1 Basic Information from Patent's Application

I obtain patent information from PatentsView.org and the USPTO. As shown in Appendix Figure A1, the patent data contains much information, such as application date, inventor name, assignee name, cited patents, technology classification, and so on.

I use the assignee's name(s) in patent documents to determine whether a patent is associated with a university. If the assignee's name includes terms such as "university," "institution," "institute," "college," "Uni.," or similar variations, I classify the patent as university-affiliated. Similarly, if the assignee's name contains terms like "company," "LLC," "corporation," "Co.," or other corporate designations, I categorize the patent as a corporate patent.

4.1.2 Patent Citations

I use citation data from the NBER patent citations database (Hall et al. 2001)[12] which includes nearly 3 million U.S. patents issued between January 1963 and December 1999, along with over 16 million citations made to these patents from 1975 to 1999.

A key focus is on the extent to which patents cite prior inventions by the same assignee ("self-citations") versus those by different, unrelated assignees. The NBER data provide details on the proportion of self-citations each patent receives. I incorporate this proportion into my analysis, distinguishing between self-citations (internal knowledge transfers) and citations from other assignees (external citations, indicating potential knowledge spillovers).

4.1.3 Office Action Research Dataset for Patents

The dataset from the United States Patent and Trademark Office (USPTO) captures details stemming from Office actions administered by patent examiners to patent applications through-

[12]Source: https://www.nber.org/research/data/us-patents-1975-1999

out the patent evaluation procedure during the 2008 to mid-2017 period. An "Office action" is a written communication to the applicant regarding the examiner's determination concerning patent granting eligibility. Typically, it discloses the rationale behind any rejections, objections, or requirements and includes relevant information or references that the applicant might find beneficial in formulating a response to the examiner and determining whether to pursue the application further (Lu et al. 2017).

Within the patent application rejections data file, each document-action is paired with action types and subtype combinations, such as 101 rejection (pertaining to subject matter eligibility, statutory double patenting, utility, etc.), 102 rejection (related to lack of novelty), 103 rejection (centered on obviousness), 112 rejection (covering written description, indefinite claims, etc.), double patenting rejection, and so on. I use 112(a) rejections, which are issued when a patent application fails to provide a clear and enabling written description of the invention, to train the AI-based readability scores, as detailed in Section 4.3.2 below.

4.1.4 Patent Text Extraction

I then extract the text from each patent, encompassing the patent background and summary, and patent detailed description. In my main analysis, I use U.S. granted patent texts from 1976 for my main sample, as full-text access is available for all U.S. patents granted since 1976.[13] To validate traditional readability measures and construct the SciBERT-based readability score, I extract U.S. patent application texts from 2008 to 2017, aligning with the time span covered by the Office Action Research Dataset. For reference, an illustration of a patent document text is provided in Appendix Figure A2, with the red portion representing the patent summary, and the green section containing the detailed description. Each segment of the patent text serves distinct purposes, and further information can be accessed on the USPTO website[14]:

– Patent Background and Summary Text. The background outlines the relevant field of the invention, often incorporating paraphrased U.S. patent classification definitions or describing the subject matter known to the inventors. The summary provides a concise overview of the invention's core features and stated objectives. Together, these components introduce the invention and situate it within its broader technological context.

– Patent Detailed Description Text. The patent detailed description defines the exact invention for which a patent is sought, differentiating it from other inventions and prior art. It describes a specific embodiment of the process, machine, manufacture, composition of matter, or improvement, and explains the mode of operation or underlying principle when applicable. The description is aimed to be sufficient so that any person of ordinary skill in the pertinent art, science, or area could make and use the invention without extensive experimentation. The

[13] Source: https://bulkdata.uspto.gov/
[14] Source: https://www.uspto.gov/patents/basics/apply/utility-patent

best mode contemplated by the inventor of carrying out the invention must be set forth in the description. Ultimately, the patent detailed description teaches how to make and use the invention.

– Patent Claim Text. Additionally, I obtain the patent claim text for each patent. The claims within the specification precisely define the subject matter considered by the inventor(s) as the invention. These claims establish the extent of the patent's protection and significantly influence whether the patent is granted. In essence, the patent claim defines the scope of legal protection of the invention.

4.2 TTO Data

The Association of University Technology Managers (AUTM) provides data on U.S. universities' Technology Transfer Offices (TTOs), including their establishment dates. I obtained TTO data from AUTM and used university names as unique identifiers to match with patent assignees. Merging the datasets involved manually cross-referencing university names between the USPTO and AUTM datasets. To address cases where different assignee names refer to the same university, I perform manual verification and consolidate these variations into a single university entity.[15]

4.3 Readability Measures

4.3.1 Traditional Readability Scores

To measure the readability of a patent document, I first apply textual analysis to construct traditional readability scores – metrics originally developed by linguists to assess the clarity and sufficiency and complexity of written text based on the patent's language. As illustrated in Figure 1, traditional readability metrics primarily depend on two factors: the recognition level of semantic elements and the intricacy of syntactic elements. To simplify application, the traditional readability formulas are based on two key assumptions. Firstly, they approximate semantic and syntactic elements using easily calculable variables. For instance, the number of syllables in a word often serves as an indicator of its semantic complexity, while the length of a sentence in words is commonly used to gauge syntactic complexity. Secondly, these formulas disregard the sequential arrangement of words and sentences, averaging the semantic metrics across all words and syntactic metrics across all sentences without considering their order (Collins-Thompson 2014).[16]

[15]Some TTO data in AUTM are not specific to individual universities but operate at a system-wide level. For example, multiple assignees within the University of California System, such as UCLA, UCSD, UC Merced, UC Davis, UC Berkeley, UC Irvine, UC Riverside, UCSF, UC Santa Barbara, and UC Santa Cruz, fall under the umbrella of the University of California System in AUTM datasets.

[16]Using natural language processing and textual analysis method to construct the traditional readability scores are commonly used in accounting and finance studies on financial reports (see Li 2008, You and Zhang 2009, Miller 2010, Beyer et al. 2010, Loughran and McDonald 2014, Campbell et al. 2014, Loughran and McDonald

This paper applies four readability measures to assess patent readability. These four readability scores are developed either to measure the readability of technical manuals or assess a manuscript's complexity during the editing process.[17] Developed for the U.S. Air Force, the Automated Readability Index is tailored to accommodate technical documents.[18] The Flesch–Kincaid grade level was developed (Kincaid et al. 1975) to describe how many grades (years of formal education) are needed to understand the text.[19] The Coleman-Liau Index approximates the U.S. grade level needed to understand the text like the readability of textbooks (Coleman and Liau 1975). Flesch Reading Ease score represents how easy a text is to understand, using the same core measures as the Flesch–Kincaid grade level (word length and sentence length), but different weighting factors (Flesch 1948). Thus, if we see text A has a lower Flesch Reading Ease score or a larger Flesch-Kincaid Grade Level, Automated Readability Index, and the Coleman-Liau Index than text B, we could conclude that text A is harder to read than text B. In my empirial analysis, except for the Flesch Reading Ease, all readability scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score is already scaled such that lower values indicate more difficult text and is not transformed.

Since patent documents contain many technical terms, those readability scores for technical manuals and/or textbooks might be suitable in my research context. The details of those readability scores' formulas are shown in Figure 1.

4.3.2 AI-Based Readability Scores Derived from SciBERT

In this section, I use patent examiners' decisions to construct novel AI-based readability scores.

– Patent Examiner's Decisions. The Office Action Research Dataset from the United States Patent and Trademark Office (USPTO) captures office actions administered by patent examiners to patent applications throughout the patent evaluation procedure. Specifically, 112 rejection (a) is the focus of my analysis. The following is a quotation of 35 U.S.C. § 112(a): "IN GENERAL.—The specification shall contain a written description of the invention, and

2016). Kong et al. (2023) and Dyer et al. (2024) expand such readability measures to the patent documents.

[17] See, e.g., DuBay (2004), and https://readabilityformulas.com/how-to-decide-which-readability-formula-to-use/

[18] "The Air Force makes extensive use of written materials such as manuals, reports, staff studies, training documents, letters, etc. The readability of a document greatly influences the time required to extract needed information from the document. Likewise, it influences the probability that the information extracted will be correctly understood and used. The costly effects of inadequate communication are well-known. Deficiencies in written communications often could be precluded by the preparation of more readable Air Force documents. An Automated Readability Index, as described in this report, offers a fast and economical means of obtaining an index of readability of Air Force materials such as textbooks and technical manuals. Use of such an Automated Readability Index would contribute significantly to the efficiency of many Air Force operations." – Smith and Senter (1967).

[19] The formula was under contract to the U.S. Navy in 1975 by J. Peter Kincaid and his team and first used by the Army for assessing the difficulty of technical manuals in 1978 and soon after became a United States Military Standard (McClure 1987).

of the manner and process of making and using it, in such full, clear, concise, and exact terms as to enable any person skilled in the art to which it pertains, or with which it is most nearly connected, to make and use the same, and shall set forth the best mode contemplated by the inventor or joint inventor of carrying out the invention."

Therefore, when a patent application receives a 112 rejection (a) from the patent examiner, it indicates that the examiner perceives the written description of the invention as insufficient and ambiguous, making it difficult for others to comprehend how to create and utilize said invention. I view this form of rejection as an evaluation of the human-examiner readability regarding the detailed description of the patent.

I use the rejection data file for patent applications and link the detailed description text for each of them. I choose final versions of patent applications that were never granted.[20] I then create the predefined classes variable "rejected by 112(a)" representing 112(a) rejection. This is a binary variable, taking a value of 1 if there were any instances of 112(a) rejection(s), and 0 otherwise. It is worth mentioning that a single application action may contain multiple rejections from different types. If an application faced at least one rejection categorized as 112(a), I interpret it as reflecting examiner concerns with whether the invention is sufficiently disclosed and comprehensible by others and assign it a dummy value of 1.[21]

The training process used to generate AI-based readability scores is described below. Further details are discussed in Appendix C.

Step 1: Generate SciBERT Embeddings on Detailed Description Text. I use the text of the rejected patent application's detailed description as input and apply the pre-trained language model SciBERT on the input text to generate the SciBERT embedding – a representation in a 768-dimensional numerical vector space.[22] SciBERT is a specialized large-language model designed for scientific text. Figure 2 displays the overview of the SciBERT model.

Step 2: Train Supervised Learning Models on Readability. After generating SciBERT embeddings for the detailed descriptions of all rejected patent applications, I use the embeddings to predict patent rejection. Specifically, I undertake supervised learning, using positive and negative instances of 112(a) rejections as the outcome, and employing both Ordinary Least Squares (OLS) regression and XGBoost models to predict rejection decisions from the text em-

[20] I do so to avoid some cases, as some applications might initially face rejection and then be resubmitted, making it challenging to distinguish between the initially rejected documents and those that were resubmitted. Please note one limitation that the rejected patent application dataset is excluded from my main analysis sample. For the main analysis, I utilize all patents granted from 1976 to 1990. The patent rejection dataset spans from 2008 to 2017 from the USPTO (Lu et al. 2017).

[21] The dataset is highly imbalanced, with only about 13.8% of observations classified as rejected by 112(a).

[22] SciBERT retains BERT's architecture but undergoes pretraining on scientific publications (1.14M scientific articles). Patent detailed descriptions often exceed the maximum token length that SciBERT can handle, which is 512 tokens. To address this, I extract the first 128 tokens, the last 128 tokens, and the middle 256 tokens from each patent description, creating shorter descriptions that preserve essential information while fitting within the model's constraints.

beddings.[23] For training, 75% of the samples are randomly selected from each stratum of the rejected patent applications, while the remaining 25% serves as the test data.[24]

Step 3: Evaluate Model Performance. Despite the dataset's imbalance – only 13.8% of observations are positive – the models perform reasonably well (see Appendix C for details). Both the OLS and XGBoosting models achieve strong recall rates (83–84%) in identifying true positives. While the ROC AUC scores (0.58–0.60) indicate moderate discriminative power, they are sufficient for constructing a proxy measure of readability, given the goal is not high-stakes prediction, but rather to estimate the likelihood of a 112(a) rejection as a function of textual clarity and sufficiency. Additional evaluation using F_β scores further supports the models' ability to prioritize recall, suggesting they are effective in identifying patents likely to be rejected under 112(a). In this context, the predicted probability of rejection serves as a meaningful and appropriate proxy for measuring variation in patent readability.

Traditional readability scores yield comparable recall but exhibit slightly lower ROC AUC, indicating that both AI-based and traditional measures capture meaningful variation aligned with examiner decisions.

Step 4: Apply Supervised Learning Models to Generate Readability Scores. I apply the trained supervised learning models with SciBERT embedding to all patent detailed descriptions in my main sample to generate predicted probabilities of rejection by 112(a).

Figure 3 presents an overview of the steps involved in constructing the readability scores using SciBERT. I interpret the predicted probability as a AI-based readability score. Higher predicted probabilities indicate greater difficulty in comprehending the patent document.

4.3.3 Comparison of Different Readability Measures

This paper uses both traditional and AI-based readability measures, each with distinct strengths and limitations. Traditional readability scores are simple to compute, applicable across patent sections, and capture syntactic and semantic features of text. However, they rely on surface-level characteristics and overlook deeper cognitive aspects of comprehension. In contrast, AI-based readability measures provide a rich, high-dimensional assessment of text. These scores estimate the likelihood of 112(a) rejections due to unclear writing but are restricted to the description section, as 112(a) applies only to the patent specification and not to other sections. Additionally, the scores are subject to limitations imposed by token constraints.[25] They also

[23] OLS, a linear regression model, is advantageous for its simplicity and interpretability but may struggle with complex, non-linear relationships in the data. XGBoost, an advanced gradient boosting technique, excels in handling non-linear patterns and interactions but is more computationally intensive.

[24] This approach ensures that both datasets maintain the same proportion across the relevant strata, enabling robust model evaluation.

[25] Full detailed descriptions often exceed SciBERT's maximum token limit of 512. To address this, I extract the first 128 tokens, the last 128 tokens, and the middle 256 tokens from each patent description, creating shorter

function as black boxes and may reflect biases from their training data.[26]

Given that both traditional readability measures and AI-based readability measures have their pros and cons and they capture different features of the readability of patent texts, I use both in my analysis. This combined approach allows for a more comprehensive assessment of readability, leveraging the strengths of each method to better understand the clarity and sufficiency and complexity of patent documents. By incorporating both traditional and AI-based readability measures, my analysis aims to provide a well-rounded evaluation of patent text readability, ensuring that comprehensive aspects are considered.

In Appendix D.1, I present real examples of how traditional readability measures and AI-based readability scores are applied to patent texts.

4.4 Baseline Sample

Table 1 presents summary statistics for the baseline sample, which includes patents filed between 1976 and 1990. Approximately 2.44% of the observations have university assignees, while 97.56% are assigned to corporations.[27] Except for the Flesch Reading Ease, all readability scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score is already scaled such that lower values indicate more difficult text and is not transformed.

5 Empirical Analyses

5.1 Initial Descriptive Evidence

Figure 4 provides a first look at the data, comparing conditional mean readability scores of university and corporate patents. Prior to 1980, university patents were more readable, based on both traditional (Automated Readability Index) and AI-based (SVL-XGBoost Readability Score) metrics. After 1980, readability scores for university patents decreased – indicating reduced clarity and sufficiency – and the gap with corporate patents narrowed. Appendix Figure A3 presents corresponding results based on additional readability measures: Flesch-Kincaid Grade Level, Coleman-Liau Index, Flesch Reading Ease, and SVL-OLS Readability Score, all of which show similar trends to Figure 4. This suggests university patents' detailed descriptions have been getting harder to read after the year 1980.

inputs that preserve essential information while conforming to the model's constraints. However, this approach may not fully capture the entirety of the original text.

[26] The model is trained on 112(a) rejections from 2008-2017, while the main analysis focuses on patents from 1976-1990. Readability standards and drafting practices may vary across these periods. Also, the application text is from the published application (PGPub). According to USPTO consultations, amendments may occur after publication, meaning examiners may review a different version than what appears in PGPub.

[27] I exclude the patents containing both university assignee(s) and corporate assignee(s).

Appendix D.1 provides a real-world example of an individual inventor at a U.S. university whose patent disclosure became less readable following the Bayh-Dole Act and the establishment of a Technology Transfer Office. This example traces the same inventor over time, holding constant the technological field and assignee, and demonstrates a decline in the clarity and sufficiency of detailed descriptions after the institutional shocks.

5.2 Empirical Strategy

Both the Bayh–Dole Act and the staggered establishment of Technology Transfer Offices (TTOs) can be viewed as quasi-experiments that provide potentially exogenous shocks to the institutional environment, and they can serve as complements to each other in empirical analysis. The Bayh–Dole Act represents a one-time federal policy change in 1980, but one might be concerned that other contemporaneous events could confound its estimated effects. In contrast, TTO formation occurred on a staggered timeline across universities, which offers useful cross-sectional and temporal variation to help address such concerns. Since TTOs are established at the university level, analyzing data at the individual inventor level offers a more robust specification than purely organizational-level analysis. However, even with an inventor-level focus, endogeneity may persist; for example, a highly influential inventor might directly or indirectly influence their university's decision to establish a TTO. While this possibility cannot be fully ruled out, I conduct parallel trends tests to assess the plausibility of the DID assumption and to reduce this concern. In addition, the Bayh–Dole Act provides a complementary source of potential exogenous variation that further strengthens identification.

5.2.1 DID Setting I: The 1980 Bayh-Dole Act

The main analysis uses difference-in-differences regressions. Using the 1980 Bayh-Dole Act, which was implemented in December 1980 as the policy shock, I divide the period from 1976 to 1980 as the pre-trend period and the period from 1981 to 1990 as the post-trend period, and compare the change in patent readability scores in the control group (corporation patents) and the treatment group (university patents), as shown in Figure 5a.

I first run the baseline regression specification at the organization level:

$$Y_p = \alpha + \beta Treated_j \times Post_t + \mu_t + \omega_j + \nu_c + \varepsilon_p, \tag{2}$$

where Y_p is the readability score for patent p. $Treated_j = 1$ if patent's assignee j is a university, and 0 otherwise. $Post_t = 1$ if patent p's application year is after the 1980 Bayh-Dole Act, and 0 otherwise.

I run the specification (2) with multiple fixed effects, including application year fixed effects (μ_t), assignee fixed effects (ω_j), and technology classification fixed effects (ν_c). Year fixed effects absorb any systematic shocks or trends common to all applications within a given year.

Assignee and technology classification fixed effects control for time-invariant heterogeneity across assignees and technology fields, capturing persistent differences in organizational characteristics and technological domains.[28] The standard errors are clustered at the assignee level.

In addition to running the regression at the assignee level, I also examine the effects of the shock at the individual level. I construct DID specification at the individual level as:

$$Y_p = \alpha + \beta Treated_j \times Post_t + \lambda_i + \mu_t + \omega_j + \nu_c + \varepsilon_p, \tag{3}$$

Specification (3) extends the baseline model in (2) by incorporating inventor fixed effects (λ_i) to control for time-invariant inventor-level heterogeneity and by clustering standard errors at the inventor level.[29]

In this setting, I generally analyze the effects of innovation commercialization by using the DID with a one-time shock on university patents and corporate patents. One might be concerned that the 1980 Bayh-Dole Act might not fully consider variations in innovation commercialization to each U.S. university. Thus, in the next section, I consider TTO establishment within a university as an innovation commercialization shock for that university.

5.2.2 DID Setting II: University TTO Establishment

I next consider TTO establishment within each university as a shock to inventors within that university and run the difference-in-differences regressions, as shown in Figure 5b.

I first follow the organization-level specification in equation (2) and then apply the individual-level specification in equation (3), in both cases modifying $Post_t$ to equal 1 if the patent p's application year is after TTO establishment year for organization j, and 0 otherwise.

5.2.3 Regressions for Dynamics

To study dynamic effects, I include a full set of leads and lags around the Bayh-Dole Act. I construct a time variable $time$ = year - 1980 (the Bayh-Dole Act year). Given my sample period is from 1976 - 1990, I run the specification (4) by the dummy variable $Time_\tau^{time}$, where $\tau \in \{-4, -3, ..., 9, 10\}$ and $Time_\tau^{time} = 1$ if $\tau = time$ and 0 otherwise.

$$Y_p = \alpha + \beta_\tau \sum_{\tau=-4}^{10} Treated_j \times Time_\tau^{time} + \lambda_i + \mu_t + \omega_j + \nu_c + \varepsilon_p \tag{4}$$

[28] The patent classification file is structured in ascending order according to the patent application document numbers. Initially, it lists all Primary U.S. Classifications, followed by the Secondary U.S. Classifications. Given that a single patent can fall under multiple classifications, I utilize the primary technology classification (primary CPC code) for each patent in my analysis. See, e.g., USPTO Bulk Downloads: Patent Classification Information(https://www.google.com/googlebooks/uspto-patents-class.html)

[29] To address a concern that inventors might change their affiliations before and after year 1980, I also run specifications which only include inventors who stayed in either universities or corporations (without switching between the two groups) before and after year 1980.

In the DID setting II, the time variable $time_j$ = year - TTO establishment year for each assignee j. As TTO establishment year varies across universities, I run the specification (5) by the dummy variable $Time_{\tau_j}^{time_j}$, where $\tau_j \in \{-4,-3,...,3,4,5\}$ and $Time_{\tau_j}^{time_j} = 1$ if $\tau_j = time_j$ and 0 otherwise.[30]

$$Y_p = \alpha + \beta_\tau \sum_{\tau_j=-4}^{5} Treated_j \times Time_{\tau_j}^{time_j} + \lambda_i + \mu_t + \omega_j + \nu_c + \varepsilon_p \tag{5}$$

If the identification assumption holds, β_τ represents the causal effect of the Bayh-Dole Act or TTO establishment on the outcome of interest τ years (before) after the implementation of the Bayh-Dole Act or the establishment of the TTO. The parallel trend assumption is likely satisfied if β_τ fluctuates around zero and shows no significant deviations in the years leading up to the Bayh-Dole Act in 1980 or prior to the establishment of the TTO.

5.3 Empirical Results

5.3.1 The Patent Detailed Description Have Become Harder to Comprehend

– DID Setting I: The 1980 Bayh-Dole Act. In the DID setting I analysis, I first utilize the traditional readability score (Automated Readability Index) and the AI-based readability score derived from SciBERT(SVL-XGBoost Readability Score)[31], and subsequently conduct regressions with additional readability measures for robustness checks. Panel A of Table 2 presents evidence of a significant decrease of 0.454 in the Automated Readability Index at the inventor level. Relative to the control group mean of -13.124, this corresponds to a 3.5% decline in text readability, indicating that university patents' detailed descriptions became less readable following the 1980 Bayh-Dole Act. Columns (1) and (2) report results at the assignee level, while columns (3) and (4) report results at the inventor level. Panel B of Table 2 shows a 0.009 decrease in the SVL-XGBoost Readability Score (relatively, the control group mean is -0.143), representing a 6.3% decline in readability among university inventors following the 1980 Bayh-Dole Act, compared to corporate inventors. Across all specifications, whether at the assignee level or inventor level, and regardless of the fixed effects applied, the coefficients of the DID interaction term are positive and statistically significant, as shown in Table 2.

– DID Setting II: University TTO establishment. In the second DID setting, I again utilize the Automated Readability Index and the SciBERT-derived SVL-XGBoost Readability Score, followed by additional robustness checks with other readability measures in the appendix. Panel A of Table 3 shows a significant decrease of 0.356 in the Automated Readability Index at the inventor level, relative to a control group mean of -13.124. This suggests that

[30] The $time_j$ equals -4 for all $time_j$ that are 4 or more years before the shock, and $time_j$ equals 5 for all $time_j$ that are 5 or more years after to the shock.

[31] For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text.

the readability of university patent descriptions has declined following the establishment of Technology Transfer Offices (TTOs). Columns (1) and (2) present results at the assignee level, while columns (3) and (4) focus on the inventor level. Panel B of Table 3 indicates a 0.014 decrease in the SVL-XGBoost Readability Score, relative to a control group mean of -0.143, among university inventors compared to corporate inventors after the TTOs were established. As in the previous setting, the coefficients of the DID interaction term are consistently positive and statistically significant across all specifications.

– Dynamic Effects. Following the Act's implementation in December 1980, I perform a regression where the treatment effect is interacted with year-specific indicator variables relative to 1980. Figure 6 illustrates the estimated coefficients and their corresponding 95-percent confidence intervals. The figure reveals that the readability scores of both the treatment and control groups fluctuated around zero with no significant differences prior to the implementation of the Bayh-Dole Act in 1980. This pattern supports the identifying assumption of parallel trends for both groups. Additionally, the coefficients generally turn significantly negative after the Bayh-Dole Act, indicating that university patent descriptions have become harder-to-read after the policy change. The dynamic estimates show a further decline in readability around 1985. One possible explanation is the 1984 amendment to the Bayh-Dole Act, which relaxed restrictions on certain exclusive licenses and may have further strengthened universities' commercialization incentives.

Figure 7 plots the estimated coefficients and their 95-percent confidence intervals to show that the identifying assumption where the trends are parallel for the treatment and the control group is likely to hold in the DID setting II. The readability score of the two groups did not significantly differ before TTO establishment, and the effect became significantly negative after TTO establishment.

Overall, these dynamic analyses underscore the impact of the Bayh-Dole Act and TTO establishment on "the harder to read" of patent descriptions, highlighting a shift in the way university patents are drafted in the wake of the 1980 Bayh-Dole Act and TTO establishment.

5.3.2 Unchanged Readability in Patent Background and Summary Text

One possible concern is that university inventors may have altered their writing behavior or other aspects of patent applications more broadly following the institutional shocks, rather than selectively changing the detailed description section. To address this, I conduct a falsification test by constructing readability scores for the background and summary sections of each patent and running the same difference-in-differences regressions. It is important to note that Section 112(a) rejections primarily focus on whether the detailed description and disclosed embodiments adequately show possession of the invention and teach how to make and use it, rather than on the background or summary sections. Consequently, because the AI-based readability

scores use examiner 112(a) rejections linked to detailed descriptions, applying this model to background and summary sections would not be appropriate due to fundamental differences in content and structure. For this analysis, I therefore rely on traditional readability measures for background and summary sections.

The insignificant coefficients in Panel A and Panel B in Table 4 reveal that there is no decrease in the readability of patent background and summary text in two difference-in-differences settings. Figure 8 plots the estimated coefficients and their 95-percent confidence intervals over year, reinforcing the findings that there is no significant readability change in patent background and summary text from university inventors after the 1980 Bayh-Dole Act or TTO establishment.

These results suggest that university inventors tend to make the detailed descriptions of their patents more difficult to read, while keeping the background and summary sections unchanged. This distinction is important because the detailed description is meant to teach how to make and use the invention, whereas the background and summary provide context. These findings imply that after the institutional changes brought by the Bayh-Dole Act and TTO establishment, university inventors may be strategically limiting the disclosure of core techniques in their patents, particularly when it comes to explaining how to create and utilize their inventions.

5.3.3 Additional Checks

– Multiple Measures. As an additional check, I assess the robustness of the main findings using multiple traditional and AI-based readability metrics. Table A10 reports results for DID Setting I (Bayh-Dole Act), and Table A11 for DID Setting II (TTO establishment). Each column corresponds to a different readability measure: the Automated Readability Index, Flesch-Kincaid Grade Level, Coleman-Liau Index, Flesch Reading Ease, and two AI-based scores (SVL-XGBoost and SVL-OLS). Except for Flesch Reading Ease, all scores are transformed so that lower values consistently indicate more difficult-to-read text.

Across both settings, university patents exhibit significant declines in readability following the institutional shocks. As expected, there are no significant changes in the readability of background and summary sections (Tables A12 and A13).

– Alternative Staggered Difference-in-Differences Settings. To address concerns arising from the staggered timing of TTO establishments across universities, I implement the CSDID method developed by Callaway and Sant'Anna (2021) as a robustness check for DID Setting II. A key advantage of CSDID is its ability to appropriately aggregate treatment effects under staggered adoption, by constructing counterfactual outcomes for treated units based on trends among a carefully selected control group of untreated units (Callaway and Sant'Anna 2021, Roth 2022).[32]

[32] One limitation, however, is that CSDID requires panel data with a single observation per unit per time period. To meet this requirement, I aggregate the patent data to the inventor-year level by computing the mean readability

Appendix Table A14 presents additional robustness checks using CSDID to assess the treatment effect of TTO establishment. Panel A of Table A14 shows the traditional readability score and Panel B of A14 shows the AI-based readability score. Appendix Figure A4 illustrates the treatment effects over time using CSDID. These results are consistent with the findings in the main analysis.

– Intensive Margin of "Stayers". In the main analysis, the dynamic treatment assignment is captured by the time-varying treated indicator, ensuring that each observation is correctly classified at each point in time. This generalized DID framework compares the average treatment effect among those classified as treated at each time point. To address potential noise or confounding effects from inventors who transition between universities and corporations, I focus on a subgroup of inventors who never moved before and after the shocks. This approach helps isolate the treatment effect but may have drawbacks, such as reducing the sample size and potentially excluding informative cases, which could impact statistical power.

Table A15 examines the intensive margin of "stayers," defined as university inventors who patented with university assignees both before and after the Bayh-Dole Act (in DID setting I). The control group comprises corporate inventors who patented with corporate assignees both before and after the Bayh-Dole Act. Table A16 examines the intensive margin of "stayers," defined as university inventors who patented with the same university assignee before and after the establishment of their respective TTOs (in DID setting II). For the control group, I include inventors who consistently patented within the same corporate assignee throughout the sample period, as there is no analogous shock for them.

– Control-group Contamination. A potential concern is that the institutional shocks may also affect the control group through local spillovers, since prior work shows that such regime change can generate geographically localized effects on nearby economic activity and innovation (Hausman 2022, Fehder et al. 2025). I address this concern in two ways. First, the raw trends suggest that the divergence in readability is driven primarily by the treatment group: university patents exhibit a clear decline in readability after the shock, while the corporate control group remains relatively stable. This pattern is inconsistent with strong spillover effects on the control group.

Second, I implement a geographic exclusion test. I re-estimate the main specifications after excluding corporations located in the same cities as treated universities, since these firms are the most likely to be affected by local knowledge spillovers. The results remain robust. Tables A17 and A18 report estimates that are consistent with the main findings.

A related concern is that the Bayh-Dole Act may also affect small businesses. To address this issue, I rerun the baseline difference-in-differences specification after excluding observa-

score for each inventor annually. This aggregation may reduce precision, as it limits the ability to control for patent-level characteristics such as technological fields.

tions marked with `small_entity_indicator` in the PatEx dataset (Graham et al. 2015). Table A19 reports results that are again consistent with the main analysis.

The findings remain consistent with the main results, providing robust evidence that the Bayh-Dole Act and the establishment of TTOs affected the intensive margin: university inventors made their patent descriptions less readable after the shocks than before.

6 Discussion of Potential Mechanisms

6.1 Obfuscation in Response to Commercialization Incentives

The main analysis documents a significant decline in the readability of university patents' detailed descriptions after the Bayh-Dole Act and the establishment of Technology Transfer Offices (TTOs), while no such change is observed in the summary or background sections. This pattern is consistent with strategic disclosure behavior driven by commercialization incentives, as predicted by my theoretical framework. This section further explores the underlying mechanisms.

6.1.1 Technological Fields with Greater Competition

The main results of this paper show that university inventors reduced the readability of their patents' detailed descriptions following the shocks. Based on the theoretical hypothesis, this finding may reflect that concerns about imitation dominate inventors' disclosure decisions when they anticipate greater future commercialization potential. To further test this hypothesis, I examine heterogeneity in the effects across technological fields with varying degrees of competition, as measured by the Herfindahl-Hirschman Index (HHI). The HHI captures the concentration of patent ownership within each Cooperative Patent Classification (CPC) code. It is computed by aggregating patents at the CPC-year level, constructing a three-year backward-looking average based on lagged values of the Herfindahl-Hirschman Index (HHI). The HHI for each year is calculated by summing the squared patent shares of all assignees within a given CPC class. This rolling average smooths fluctuations.[33] A lower HHI indicates a less concentrated technological space and, therefore, greater competitive pressure and imitation risk.

Panel A and Panel B in Table 5 present the heterogeneous treatment effects by technological

[33] For a given CPC class c and year t, let $s_{i,c,t}$ denote the share of patents granted to assignee i. Then the yearly HHI is defined as:

$$\mathrm{HHI}_{c,t} = \sum_{i=1}^{N_{c,t}} s_{i,c,t}^2,$$

where $N_{c,t}$ is the number of assignees active in class c in year t. To construct a rolling measure of competition, I generate lagged HHI values for the previous one, two, and three years within each CPC class. I then compute a three-year backward-looking average of HHI, denoted $\overline{\mathrm{HHI}}_{c,t}^{(\mathrm{lag})} = \frac{1}{3}\left(\mathrm{HHI}_{c,t-1} + \mathrm{HHI}_{c,t-2} + \mathrm{HHI}_{c,t-3}\right)$.

field competitiveness.[34] "Competitive Fields" is a dummy variable equal to 1 if the technology field falls within the bottom 50% of the Herfindahl-Hirschman Index (HHI), indicating higher competition, and 0 otherwise. Across both policy shocks, university inventors tended to reduce the readability of their patents' detailed descriptions more in competitive fields. However, no significant effects were observed in the readability of the background and summary sections. In DID Setting I (Bayh-Dole Act), the triple interaction terms are statistically significant for both the traditional and AI-based readability scores in the detailed description section, indicating a robust effect. In DID Setting II (TTO establishment), the effects in the detailed description section remain strong and statistically significant for the AI-based readability score. The traditional readability measure exhibits similar directional patterns, although its coefficients are less statistically significant.

The findings suggest that increased competition amplifies the tendency to obscure technical details, as inventors seek to protect their innovations from potential imitation and preserve commercialization opportunities. These results are consistent with economic theory and empirical evidence in previous literature suggesting that heightened technological competition can lead inventors to reduce their innovation disclosure (Horstmann et al. 1985, Anton and Yao 2004, Glaeser and Landsman 2021, Fadeev 2023).[35] This pattern provides further support for the theoretical hypothesis of strategic obfuscation in response to commercialization incentives: when institutions place greater emphasis on commercialization incentives, inventors reduce their level of disclosure, particularly in contexts where imitation concerns are high.

6.1.2 Timing of Readability Declines

Universities' disclosure practices may have been shaped either by the Bayh-Dole Act's direct enhancement of commercialization incentives, by the subsequent establishment of TTOs which provided additional commercialization incentives and institutional support such as legal and technical expertise, or by the combined effects of both. The analysis below disentangles the timing of readability declines to assess whether they occurred after Bayh-Dole alone, after TTOs alone, or after both.

To further test whether enhanced commercialization potential influences disclosure behavior, I focus on periods prior to TTO establishment in treated universities. Since the Bayh-Dole Act primarily introduced commercialization rights, this approach allows a cleaner identification of its direct impact, separate from the additional institutional changes (such as professionalization and legal support) introduced by TTOs. As shown in Table 6, university inventors had already begun reducing the readability of their patent descriptions even before TTOs were established. Notably, there was no corresponding change in the readability of the summary or

[34] CPC classes are classified as "competitive" if their lagged HHI average falls below the median of the CPC-year HHI distribution.

[35] Glaeser and Landsman (2021) finds that a 1% increase in realized technological competition leads to a 0.04% increase in the delay before patent disclosure. This behavior arises from concerns that competitors may use disclosed information to circumvent or surpass the innovation.

background sections. This pattern further supports the hypothesis that the Bayh-Dole Act, by enhancing commercialization incentives, encouraged inventors to disclose less in order to safeguard the commercial value of their inventions.

As central actors in university commercialization, TTOs manage intellectual property rights, oversee licensing, and guide inventors through the patenting process (Hellmann 2007, Hvide and Jones 2018, Faccin et al. 2022). To isolate their effect, I restrict the sample to the post-1980 period, conditional on the Bayh-Dole Act being in place.[36] As shown in Table 7, inventors further reduced the readability of patent descriptions following the establishment of TTOs. Again, there was no corresponding change in the readability of the summary or background sections. This pattern suggests that TTOs, by reinforcing inventors' commercialization incentives, may have contributed to more strategic disclosure. However, this does not fully rule out the alternative possibility: that professional guidance provided by TTOs may have influenced disclosure practices. I address this possibility in the next section through additional robustness checks, including controls for professional legal support.

6.1.3 Impacts on Knowledge Diffusion

The obfuscation of patent text may limit knowledge spillovers, undermining a primary motive of the patent system. Here I consider a measure of knowledge spillovers, drawing on citation behavior (Hall et al. 2001, Henderson et al. 1998a, Mowery and Ziedonis 2002). I begin by estimating difference-in-differences regressions using total citations as the main outcome. As shown in column (1) in Table 8, patents from the treatment group receive significantly fewer citations following the Bayh-Dole Act and the subsequent establishment of Technology Transfer Offices (TTOs), relative to the control group. This result is consistent with prior research by Henderson et al. (1998a) and Mowery and Ziedonis (2002), who found that patents granted to institutions that began patenting after Bayh-Dole tended to receive fewer citations. Previous literature suggests that the lower citation rates of university patents issued after 1980 may reflect their being perceived as less "important" or "general."

While existing literature points to factors such as reduced importance or generality of post-1980 university patents, I propose an additional channel: the decline in citations may also stem from changes in disclosure behavior. This mechanism does not rule out other explanations but offers a complementary perspective on reduced knowledge diffusion. To further test this, I disaggregate citations into internal and external types. Internal citations are defined as those made by patents assigned to the same organization as the cited patent and do not reflect broader knowledge diffusion. In contrast, external citations are those made by patents from different organizations and more directly capture spillovers across institutional boundaries. Disaggregated results, presented in Table 8, reveal that the decline in total citations is driven primarily by a re-

[36] Restricting the analysis to the post-1980 period creates a uniform IP environment under Bayh-Dole but sharply limits pre-treatment observations. This loss of statistical power makes the key coefficient less precise, and its significance declines relative to the main specification.

duction in external citations, with internal citations remaining largely unchanged. This pattern is further supported by year-by-year estimates shown in Figures 9 and 10, which demonstrate that the decline in citations following the policy shocks is concentrated in external citations.

Further, I utilize the 1980 Bayh-Dole Act as an instrumental variable (IV) to estimate the readability scores of patent disclosures. These scores are then used in a two-stage least squares (2SLS) regression to determine their relationship with citation. The results, presented in Table 9, show that a decrease in the readability of patent detailed descriptions significantly predicts a reduction in both total and external citations. These 2SLS findings reinforce the hypothesis that reduced readability in patent disclosures negatively correlates with knowledge spillover.

It is important to note that the two-stage least squares regression does not provide a fully causal estimate in this context. Accordingly, I interpret the IV estimates as descriptive evidence: the observed reduction in follow-on external citations is consistent with a readability-based mechanism. These findings offer a complementary perspective by highlighting readability as a potential channel through which knowledge diffusion is reduced, while acknowledging that multiple other mechanisms may simultaneously be at work.

6.2 Alternative Explanations and Confounding Factors

To further test confounding factors that might explain the observed decline in patent readability, I examine alternative explanations. Could the effects be driven by a shift in the basic versus applied orientation of university patents? Are changes in originality or generality responsible for the patterns observed? Could changes in legal support and attorney involvement be influencing how patents are drafted? The following analyses address each of these questions in turn. Together, these tests help isolate the effect of commercialization incentives on disclosure behavior from other potential sources of variation.

First, to account for potential shifts in the basic-applied orientation of patents. The results remain robust after controlling for the orientation variables (Tables A20 and A21). Next, I examine the originality and generality of patents. Adding these controls does not alter the main results (Tables A22 and A23), reinforcing the view that these patent characteristics are not fully driving the observed effects. Patent quality is another potential confounder. To explore this, I use two proxies for patent quality. First, I use the journal associated with the paper paired to the patent from the Patent-Paper Pairs (PPP) dataset (Marx and Scharfmann, 2024) as a proxy for the quality of the invention. These journal indicators are included as fixed effects in the regressions, and the results remain consistent (Tables A24 and A25). Second, I use market value from the KPSS dataset (Kogan et al., 2017), which I examine separately in a correlational analysis. The KPSS results show only weak and mixed correlations between market value and readability (Figure A5), further supporting the conclusion that shifts in patent quality are unlikely to fully explain the main findings.

Finally, I consider the role of professional legal support in driving the observed changes. A

natural alternative explanation is that the decline in readability reflects increased attorney involvement. However, note that increased legal support, for example through university TTOs, may itself be one dimension of the institutional change associated with commercialization environment, rather than an entirely separate mechanism. While attorneys are typically responsible for drafting claims in precise legal language, inventors are primarily responsible for providing the technical content, as noted in prior research (Bryan et al. 2020, Andriosopoulos et al. 2023) and confirmed by my discussions with legal practitioners (Discussion of inventor and attorney's roles in Appendix D.3). While disentangling the effects of attorneys is beyond the primary scope of this paper and still an open question, I address potential bias from attorney-specific drafting styles by including attorney fixed effects in robustness checks. Including attorney fixed effects (Tables A26 and A27) slightly attenuates the coefficients but does not eliminate the main effect. This suggests that professionalization may contribute to the changes but is unlikely to be the sole driver.

An additional comparison comes from patent claims, the section most directly tied to legal scope and legal drafting. If the observed pattern primarily reflected a general shift in legal drafting, one might expect corresponding changes in claims readability. By contrast, I find little evidence of a comparable decline in claims readability after either institutional change (Appendix Table A29).[37] This pattern suggests that the estimated effect is concentrated in technical disclosure rather than reflecting a uniform change in patent drafting style.

Together, these robustness checks provide little support for these alternative perspectives. Full methodological details and robustness results are provided in Appendix D.2.

7 Conclusion

This paper examines patent disclosure behavior by constructing readability scores as a proxy for patent information disclosure, using both traditional and novel AI-based readability scores through natural language processing methods. I employ difference-in-differences approaches to identify the impact of institutional changes on patent disclosure. First, I use the 1980 Bayh-Dole Act, which granted universities patent property rights and the potential for economic benefits from their inventions, as an institutional shock. Second, I consider the establishment of technology transfer offices (TTOs) as shocks to university inventors, introducing support for the innovation commercialization processes. Across both settings, university patents become less readable in their detailed descriptions, while background and summary sections change little.

The results are consistent with the view that stronger commercialization incentives can reduce the clarity of patent disclosure, particularly in settings where imitation concerns are likely to be stronger. The larger effects in more competitive technological fields support this interpre-

[37] Because patents typically contain multiple claims, I construct a patent-level measure of claims readability by averaging readability scores across all claims within a patent.

tation. The contribution of the paper is to document a systematic and previously underexplored margin through which institutions may shape the informational content of patents.

I also show that these institutional changes are associated with declines in external citations to university patents. Although these citation patterns should be interpreted cautiously, they are consistent with the possibility that lower disclosure clarity may reduce downstream knowledge diffusion. More broadly, the findings suggest that policies designed to stimulate commercialization may affect not only whether inventions are patented, but also how clearly patented knowledge is disclosed.

These results point to a broader issue in innovation policy. The costs and benefits of commercialization policy are usually evaluated in terms of patenting, licensing, startup formation, or other observable outputs. This paper highlights an additional margin: the clarity of disclosure within patents themselves. If institutions that strengthen commercialization also weaken the disclosure function of patents, then the social effects of innovation policy may depend not only on how much invention is produced, but also on how effectively patented knowledge remains available for follow-on research.

References

Acikalin, U. U., Caskurlu, T., Hoberg, G., and Phillips, G. M. (2023). Intellectual property protection lost and competition: An examination using large language models. Working paper.

Aghion, P. and Tirole, J. (1994). The management of innovation. *The Quarterly Journal of Economics*, 109(4):1185–1209.

Akcigit, U., Caicedo Soler, S., Miguélez, E., Stantcheva, S., and Sterzi, V. (2018). Dancing with the stars: Innovation through interactions. Working paper, National Bureau of Economic Research.

Andriosopoulos, D., Czarnowski, P., and Marshall, A. (2023). Do corporate lawyers matter? evidence from patents. *Journal of Corporate Finance*, 83:102473.

Angori, G., Marzocchi, C., Ramaciotti, L., and Rizzo, U. (2024). A patent-based analysis of the evolution of basic, mission-oriented, and applied research in European universities. *The Journal of Technology Transfer*, 49(2):609–641.

Anton, J. J. and Yao, D. A. (1994). Expropriation and inventions: Appropriable rents in the absence of property rights. *The American Economic Review*, pages 190–209.

Anton, J. J. and Yao, D. A. (2002). The sale of ideas: Strategic disclosure, property rights, and contracting. *The Review of Economic Studies*, 69(3):513–531.

Anton, J. J. and Yao, D. A. (2004). Little patents and big secrets: managing intellectual property. *RAND Journal of Economics*, pages 1–22.

Arora, A., Belenzon, S., Marx, M., and Shvadron, D. (2024). When does patent protection spur cumulative research within firms? *The Journal of Law, Economics, and Organization*, 40(3):694–728.

Ashtor, J. H. (2022). Modeling patent clarity. *Research Policy*, 51(2):104415.

Baker, S. and Mezzetti, C. (2005). Disclosure as a strategy in the patent race. *The Journal of Law and Economics*, 48(1):173–194.

Bearson, D. F. and Roche, M. P. (2025). *Transferability MATTRs: Towards Understanding Antecedents of Strategic Licensing*. Harvard Business School.

Beltagy, I., Lo, K., and Cohan, A. (2019). SciBERT: A pretrained language model for scientific text. *arXiv preprint arXiv:1903.10676*.

Beyer, A., Cohen, D. A., Lys, T. Z., and Walther, B. R. (2010). The financial reporting environment: Review of the recent literature. *Journal of Accounting and Economics*, 50(2-3):296–343.

Bhattacharya, S. and Ritter, J. R. (1983). Innovation and communication: Signalling with partial disclosure. *The Review of Economic Studies*, 50(2):331–346.

Biasi, B. and Moser, P. (2021). Effects of copyrights on science: Evidence from the WWII book republication program. *American Economic Journal: Microeconomics*, 13(4):218–260.

Bloom, N., Schankerman, M., and Van Reenen, J. (2013). Identifying technology spillovers and product market rivalry. *Econometrica*, 81(4):1347–1393.

Bryan, K. A., Ozcan, Y., and Sampat, B. (2020). In-text patent citations: A user's guide. *Research Policy*, 49(4):103946.

Burrage, A., Dasgupta, N., and Ganguli, I. (2024). Gender diversity in academic entrepreneurship: Social impact motives and the NSF I-Corps program. Working Paper 32885, National Bureau of Economic Research.

Bushee, B. J., Gow, I. D., and Taylor, D. J. (2018). Linguistic complexity in firm disclosures: Obfuscation or information? *Journal of Accounting Research*, 56(1):85–121.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Campbell, J. L., Chen, H., Dhaliwal, D. S., Lu, H.-m., and Steele, L. B. (2014). The information content of mandatory risk factor disclosures in corporate filings. *Review of Accounting Studies*, 19(1):396–455.

Chukumba, C. and Jensen, R. (2005). University invention, entrepreneurship, and start-ups. Working Paper 11475, National Bureau of Economic Research.

Cockburn, I. and Henderson, R. (1996). Public–private interaction in pharmaceutical research. *Proceedings of the National Academy of Sciences*, 93(23):12725–12730.

Cohen, W. M., Nelson, R. R., and Walsh, J. P. (2000). Protecting their intellectual assets: Appropriability conditions and why US manufacturing firms patent (or not). Working Paper 7552, National Bureau of Economic Research.

Coleman, M. and Liau, T. L. (1975). A computer readability formula designed for machine scoring. *Journal of Applied Psychology*, 60(2):283.

Collins-Thompson, K. (2014). Computational assessment of text readability: A survey of current and future research. *ITL-International Journal of Applied Linguistics*, 165(2):97–135.

Contreras, J. L. (2011). Bermuda's legacy: policy, patents, and the design of the genome commons. *Minn. JL Sci. & Tech.*, 12:61.

Contreras, J. L. (2022). In the public interest: University technology transfer and the nine points document-an empirical assessment. *UC Irvine L. Rev.*, 13:435.

Devlin, A. (2009). The misunderstood function of disclosure in patent law. *Harv. JL & Tech.*, 23:401.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

DuBay, W. H. (2004). The principles of readability.

Dyer, T. A., Glaeser, S., Lang, M. H., and Sprecher, C. (2024). The effect of patent disclosure quality on innovation. *Journal of Accounting and Economics*, 77(2-3):101647.

Eisenberg, R. S. (1996). Public research and private development: Patents and technology transfer in government–sponsored research. *Va. L. Rev.*, 82:1663.

Evans, D. S. and Leighton, L. S. (1989). Some empirical aspects of entrepreneurship. *The American Economic Review*, pages 519–535.

Faccin, K., De Beer, C., Martins, B. V., Zanandrea, G., Kela, N., and Schutte, C. (2022). What really matters for TTOs efficiency? An analysis of TTOs in developed and developing economies. *The Journal of Technology Transfer*, 47(4):1135–1161.

Fadeev, E. (2023). Creative construction: Knowledge sharing and cooperation between firms. Working paper, Duke University Fuqua School of Business.

Fehder, D. C., Hausman, N., and Hochberg, Y. V. (2025). Innovation and capital. *Journal of Financial Economics*, 169:104029.

Figueroa, N. and Lemus, J. (2024). The benefit of the doubt: Patent examination under strategic obfuscation. *Management Science*, 70(9):5853–5870.

Flesch, R. (1948). A new readability yardstick. *Journal of Applied Psychology*, 32(3):221.

Freilich, J. (2017). Patent clutter. *Iowa L. Rev.*, 103:925.

Fromer, J. C. (2008). Patent disclosure. *Iowa L. Rev.*, 94:539.

Funk, P. and Iriberri, N. (2020). Are referees and editors in economics gender neutral? *The Quarterly Journal of Economics*, 135:269–327.

Furman, J. L., Nagler, M., and Watzinger, M. (2021). Disclosure and subsequent innovation: Evidence from the patent depository library program. *American Economic Journal: Economic Policy*, 13(4):239–270.

Galasso, A. and Schankerman, M. (2015). Patents and cumulative innovation: Causal evidence from the courts. *The Quarterly Journal of Economics*, 130(1):317–369.

Gallini, N. T. (2002). The economics of patents: Lessons from recent US patent reform. *Journal of Economic Perspectives*, 16(2):131–154.

Gans, J. S. and Murray, F. (2011). Funding scientific knowledge: Selection, disclosure and the public-private portfolio. In *The Rate and Direction of Inventive Activity Revisited*, pages 51–103. University of Chicago Press.

Gans, J. S., Murray, F. E., and Stern, S. (2017). Contracting over the disclosure of scientific knowledge: Intellectual property and academic publication. *Research Policy*, 46(4):820–835.

Glaeser, S. A. and Landsman, W. R. (2021). Deterrent disclosure. *The Accounting Review*, 96(5):291–315.

Graham, S. and Hegde, D. (2015). Disclosing patents' secrets. *Science*, 347(6219):236–237.

Graham, S. J., Marco, A. C., and Miller, R. (2015). The USPTO patent examination research dataset: A window on the process of patent examination.

Green, J. R. and Scotchmer, S. (1995). On the division of profit in sequential innovation. *The RAND Journal of Economics*, pages 20–33.

Grushcow, J. M. (2004). Measuring secrecy: A cost of the patent system revealed. *The Journal of Legal Studies*, 33(1):59–84.

Hall, B. H. and Harhoff, D. (2012). Recent research on the economics of patents. *Annu. Rev. Econ.*, 4(1):541–565.

Hall, B. H., Jaffe, A. B., and Trajtenberg, M. (2001). The NBER patent citation data file: Lessons, insights and methodological tools.

Hamilton, B. H. (2000). Does entrepreneurship pay? An empirical analysis of the returns to self-employment. *Journal of Political Economy*, 108(3):604–631.

Hausman, N. (2022). University innovation and local economic growth. *Review of Economics and Statistics*, 104(4):718–735.

Head, M. L., Holman, L., Lanfear, R., Kahn, A. T., and Jennions, M. D. (2015). The extent and consequences of p-hacking in science. *PLoS biology*, 13(3):e1002106.

Hegde, D., Herkenhoff, K., and Zhu, C. (2023). Patent publication and innovation. *Journal of Political Economy*, 131(7):1845–1903.

Hegde, D. and Luo, H. (2018). Patent publication and the market for ideas. *Management Science*, 64(2):652–672.

Hellmann, T. (2007). When do employees become entrepreneurs? *Management Science*, 53(6):919–933.

Henderson, R., Jaffe, A. B., and Trajtenberg, M. (1998a). Universities as a source of commercial technology: a detailed analysis of university patenting, 1965–1988. *Review of Economics and Statistics*, 80(1):119–127.

Henderson, R., Jaffe, A. B., and Trajtenberg, M. (1998b). University patenting amid changing incentives for commercialization. In *Creation and Transfer of Knowledge: Institutions and Incentives*, pages 87–114. Springer.

Hengel, E. (2022). Publishing while female: Are women held to higher standards? Evidence from peer review. *The Economic Journal*, 132(648):2951–2991.

Horstmann, I., MacDonald, G. M., and Slivinski, A. (1985). Patents as information transfer mechanisms: To patent or (maybe) not to patent. *Journal of Political Economy*, 93(5):837–858.

Hsu, D. H. and Kuhn, J. M. (2023). Academic stars and licensing experience in university technology commercialization. *Strategic Management Journal*, 44(3):887–905.

Huang, K. G. and Murray, F. E. (2009). Does patent strategy shape the long-run supply of public knowledge? Evidence from human genetics. *Academy of Management Journal*, 52(6):1193–1221.

Hvide, H. K. and Jones, B. F. (2018). University innovation and the professor's privilege. *American Economic Review*, 108(7):1860–98.

Jaffe, A. B. (2000). The US patent system in transition: Policy innovation and the innovation process. *Research Policy*, 29(4-5):531–557.

Jaffe, A. B., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *The Quarterly Journal of Economics*, 108(3):577–598.

Jensen, R. A., Thursby, J. G., and Thursby, M. C. (2003). Disclosure and licensing of university inventions:'the best we can do with the s** t we get to work with'. *International Journal of Industrial Organization*, 21(9):1271–1300.

Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317.

Jones, B. F. (2010). Age and great invention. *The Review of Economics and Statistics*, 92(1):1–14.

Jones, C. I. (1995a). R & D-based models of economic growth. *Journal of Political Economy*, 103(4):759–784.

Jones, C. I. (1995b). Time series tests of endogenous growth models. *The Quarterly Journal of Economics*, 110(2):495–525.

Kim, J. and Valentine, K. (2021). The innovation consequences of mandatory patent disclosures. *Journal of Accounting and Economics*, 71(2-3):101381.

Kincaid, J. P., Fishburne Jr, R. P., Rogers, R. L., and Chissom, B. S. (1975). Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2):665–712.

Kong, N., Dulleck, U., Jaffe, A. B., Sun, S., and Vajjala, S. (2023). Linguistic metrics for patent disclosure: Evidence from university versus corporate patents. *Research Policy*, 52(2):104670.

Kortum, S. (1993). Equilibrium R&D and the decline in the patent-R&D ratio: U.S. evidence. *American Economic Review: Papers and Proceedings*, 83(2):450–7.

Kuhn, J., Younge, K., and Marco, A. (2020). Patent citations reexamined. *The RAND Journal of Economics*, 51(1):109–132.

Lam, A. (2011). What motivates academic scientists to engage in research commercialization:'gold','ribbon'or 'puzzle'? *Research Policy*, 40(10):1354–1368.

Landes, W. M., Posner, R. A., et al. (2003). *The economic structure of intellectual property law*. Harvard university press.

Larrimore, L., Jiang, L., Larrimore, J., Markowitz, D., and Gorski, S. (2011). Peer to peer lending: The relationship between language features, trustworthiness, and persuasion success. *Journal of Applied Communication Research*, 39(1):19–37.

Lemley, M. A. (2011). The myth of the sole inventor. *Mich. L. Rev.*, 110:709.

Lerner, J., Manley, H., Stein, C., and Williams, H. (2024). The wandering scholars: Understanding the heterogeneity of university commercialization. Working paper, National Bureau of Economic Research.

Li, F. (2008). Annual report readability, current earnings, and earnings persistence. *Journal of Accounting and Economics*, 45(2-3):221–247.

Litan, R. E., Mitchell, L., and Reedy, E. (2007). Commercializing university innovations: Alternative approaches. *Innovation Policy and the Economy*, 8:31–57.

Long, C. (2002). Patent signals. *The University of Chicago Law Review*, pages 625–679.

Loughran, T. and McDonald, B. (2014). Measuring readability in financial disclosures. *Journal of Finance*, 69(4):1643–1671.

Loughran, T. and McDonald, B. (2016). Textual analysis in accounting and finance: A survey. *Journal of Accounting Research*, 54(4):1187–1230.

Lu, Q., Myers, A., and Beliveau, S. (2017). USPTO patent prosecution research data: Unlocking office action traits. *USPTO economic working paper*.

Mansfield, E. (1986). Patents and innovation: An empirical study. *Management Science*, 32(2):173–181.

Markowitz, D. M. (2019). What words are worth: National science foundation grant abstracts indicate award funding. *Journal of Language and Social Psychology*, 38(3):264–282.

Marx, M. and Fuegi, A. (2022). Reliance on science by inventors: Hybrid extraction of in-text patent-to-article citations. *Journal of Economics & Management Strategy*, 31(2):369–392.

Marx, M. and Scharfmann, E. (2024). Does patenting promote the progress of science? Working paper.

McClure, G. M. (1987). Readability formulas: Useful or useless? *IEEE Transactions on Professional Communication*, (1):12–15.

McGurk, M. R. and Lu, J. W. (2015). The intersection of patents and trade secrets. *Hastings Sci. & Tech. LJ*, 7:189.

Merton, R. K. (1973). The sociology of science: Theoretical and empirical investigations. *The University of Chicago*.

Mezzanotti, F. and Simcoe, T. (2023). Innovation and appropriability: Revisiting the role of intellectual property.

Miller, B. P. (2010). The effects of reporting complexity on small and large investor trading. *The Accounting Review*, 85(6):2107–2143.

Miric, M., Jia, N., and Huang, K. G. (2023). Using supervised machine learning for large-scale classification in management research: The case for identifying artificial intelligence patents. *Strategic Management Journal*, 44(2):491–519.

Mowery, D. C., Nelson, R. R., Sampat, B. N., and Ziedonis, A. A. (2001). The growth of patenting and licensing by US universities: An assessment of the effects of the Bayh–Dole act of 1980. *Research Policy*, 30(1):99–119.

Mowery, D. C. and Sampat, B. N. (2001). University patents and patent policy debates in the usa, 1925–1980. *Industrial and Corporate Change*, 10(3):781–814.

Mowery, D. C. and Sampat, B. N. (2004). The Bayh-Dole Act of 1980 and university-industry technology transfer: A model for other OECD governments? *The Journal of Technology Transfer*, 30(1):115–127.

Mowery, D. C. and Ziedonis, A. A. (2002). Academic patent quality and quantity before and after the Bayh–Dole act in the United States. *Research Policy*, 31(3):399–418.

Murray, F. and Stern, S. (2007). Do formal intellectual property rights hinder the free flow of scientific knowledge?: An empirical test of the anti-commons hypothesis. *Journal of Economic Behavior & Organization*, 63(4):648–687.

Nosek, B. A., Spies, J. R., and Motyl, M. (2012). Scientific utopia: II. restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 7(6):615–631.

Orozco, D. (2019). Assessing the efficacy of the Bayh-Dole act through the lens of university technology transfer offices (TTOs). *NCJL & Tech.*, 21:115.

Ouellette, L. L. (2011). Do patents disclose useful information. *Harv. JL & Tech.*, 25:545.

Ouellette, L. L. and Tutt, A. (2020). How do patent incentives affect university researchers? *International Review of Law and Economics*, 61:105883.

Partha, D. and David, P. A. (1994). Toward a new economics of science. *Research Policy*, 23(5):487–521.

Rabitschek, J. H. and Latker, N. J. (2005). Reasonable pricing-a new twist for march-in rights under the Bayh-Dole act. *Santa Clara Computer & High Tech. LJ*, 22:149.

Roin, B. N. (2005). The disclosure function of the patent system (or lack thereof). *Harvard Law Review*.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy*, 98(5, Part 2):S71–S102.

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, 4(3):305–322.

Rothaermel, F. T., Agung, S. D., and Jiang, L. (2007). University entrepreneurship: a taxonomy of the literature. *Industrial and corporate change*, 16(4):691–791.

Rothaermel, F. T. and Thursby, M. (2005). University–incubator firm knowledge flows: assessing their impact on incubator firm performance. *Research policy*, 34(3):305–320.

Sampat, B. and Williams, H. L. (2019). How do patents affect follow-on innovation? Evidence from the human genome. *American Economic Review*, 109(1):203–236.

Sampat, B. N. (2018). A survey of empirical evidence on patents and innovation.

Sauermann, H. and Roach, M. (2012). Taste for science, taste for commercialization, and hybrid scientists. In *34th DRUID Celebration Conference*.

Scotchmer, S. (1991). Standing on the shoulders of giants: Cumulative research and the patent law. *Journal of Economic Perspectives*, 5(1):29–41.

Scotchmer, S. (2004). *Innovation and incentives*. MIT press.

Seymore, S. B. (2009). The teaching function of patents. *Notre Dame L. Rev.*, 85:621.

Shane, S. (2004). Encouraging university entrepreneurship? The effect of the Bayh-Dole Act on university patenting in the United States. *Journal of Business Venturing*, 19(1):127–151.

Shane, S., Locke, E. A., and Collins, C. J. (2003). Entrepreneurial motivation. *Human Resource Management Review*, 13(2):257–279.

Smith, E. A. and Senter, R. (1967). *Automated readability index*, volume 66. Aerospace Medical Research Laboratories, Aerospace Medical Division.

Spence, M. (1974). Competitive and optimal responses to signals: An analysis of efficiency and distribution. *Journal of Economic Theory*, 7(3):296–332.

Stern, S. (2004). Do scientists pay to be scientists? *Management Science*, 50(6):835–853.

Thursby, J. G. and Thursby, M. C. (2002). Who is selling the ivory tower? Sources of growth in university licensing. *Management science*, 48(1):90–104.

Trajtenberg, M., Henderson, R., and Jaffe, A. (1997). University versus corporate patents: A window on the basicness of invention. *Economics of Innovation and New Technology*, 5(1):19–50.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.

Walsh, J. P., Cohen, W. M., and Cho, C. (2007). Where excludability matters: Material versus intellectual property in academic biomedical research. *Research Policy*, 36(8):1184–1203.

Weitzman, M. L. (1998). Recombinant growth. *The Quarterly Journal of Economics*, 113(2):331–360.

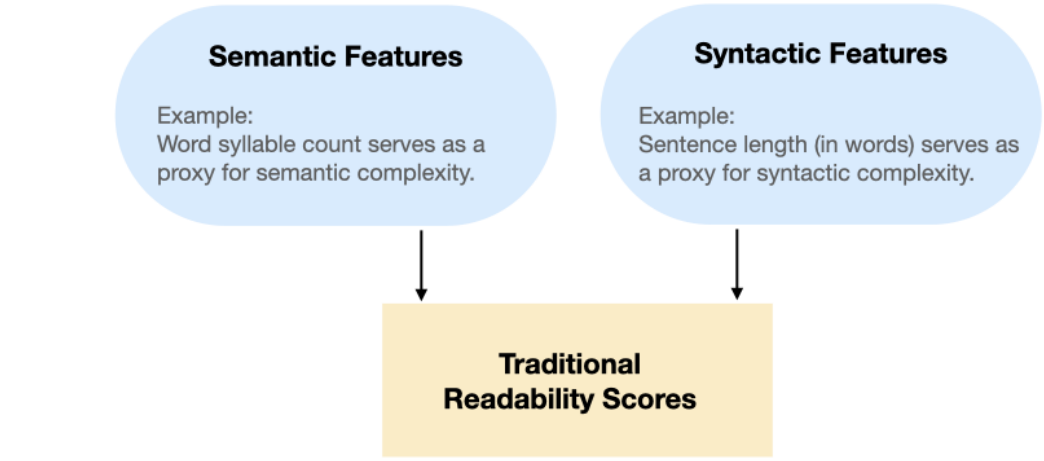
Williams, H. L. (2013). Intellectual property rights and innovation: Evidence from the human genome. *Journal of Political Economy*, 121(1):1–27.

You, H. and Zhang, X.-j. (2009). Financial reporting complexity and investor underreaction to 10-k information. *Review of Accounting Studies*, 14(4):559–586.

Young, A. (1998). Growth without scale effects. *Journal of Political Economy*, 106(1):41–63.

Figures

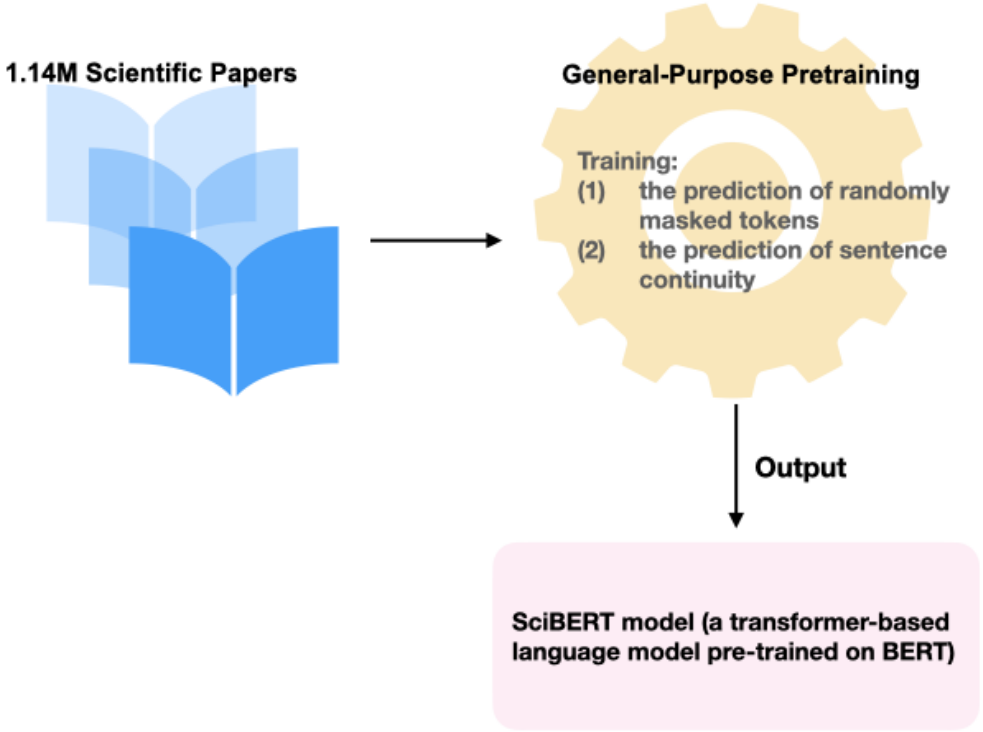
Figure 1: Traditional Readability Measures



Traditional Readability Measure	Formula	Interpretation
Automated Readability Index	4.71(charaters/words) + 0.5(words/ sentences) - 21.43	High score: hard-to-read
Flesch-Kincaid Grade Level	0.39(total words/total sentences) + 11.8(total syllables/total words) -15.59	High score: hard-to-read
The Coleman-Liau Index	0.0588L - 0.296S - 15.8; L : the average number of letters per 100 words and S : the average number of sentences per 100 words	High score: hard-to-read
Flesch Reading Ease	206.835 - 1.015(total words/total sentences) - 84.6(total syllables/total words)	Low score: hard-to-read

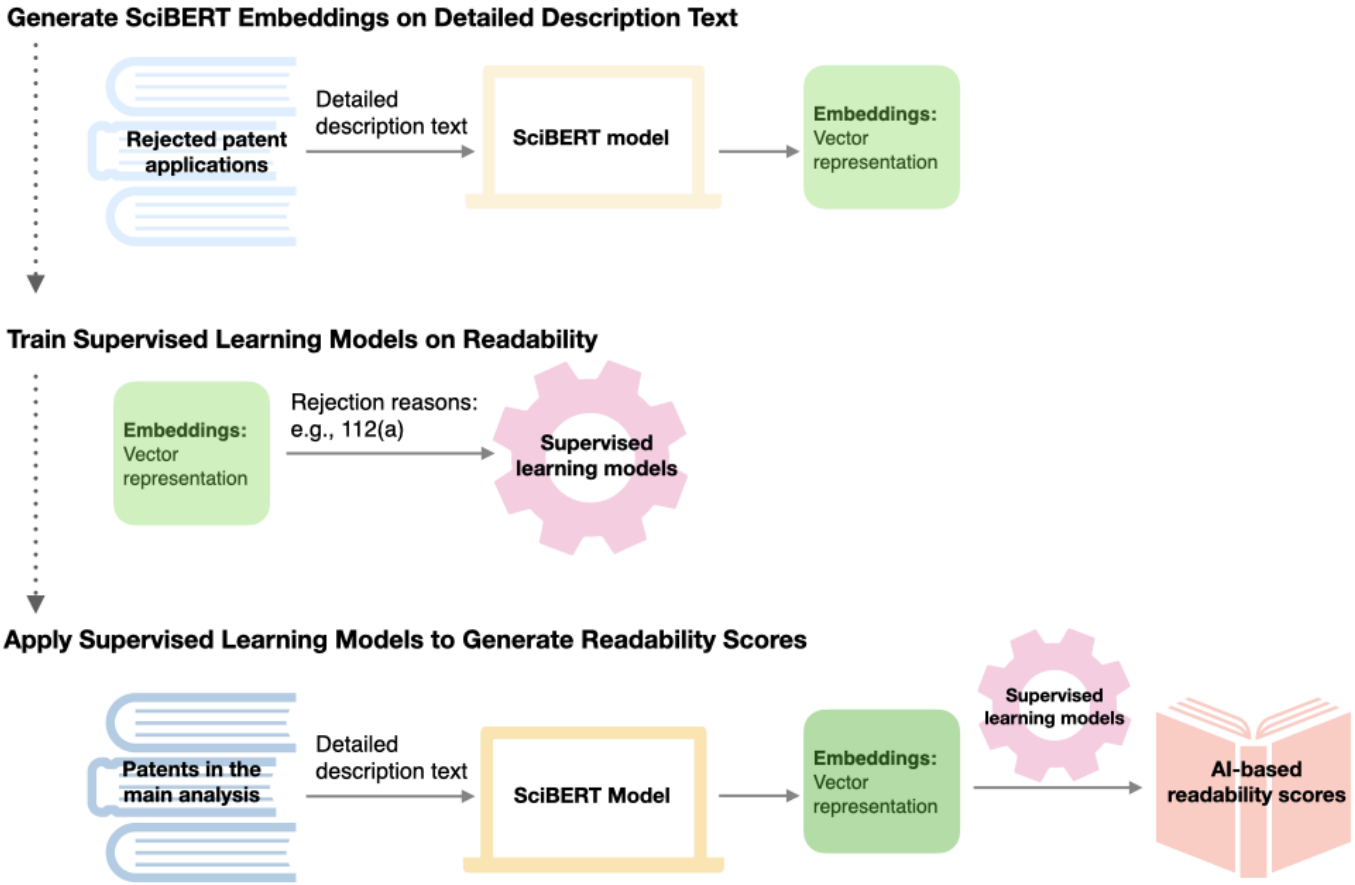
Notes: The figure shows the two factors considered in traditional readability scores and the original formulas for four traditional readability scores. Traditional readability metrics simplify the measurement of textual complexity through two main assumptions. First, they approximate semantic and syntactic difficulty using easily quantifiable variables. For instance, word syllable count serves as a proxy for semantic difficulty, while sentence length (in words) captures syntactic complexity. Second, these formulas ignore the sequential arrangement of words and sentences, instead averaging semantic and syntactic metrics across the entire text (Collins-Thompson, 2014). In the main analysis, all readability scores, except the Flesch Reading Ease score, are multiplied by negative one based on their original formulas so that lower values indicate harder-to-read text. The Flesch Reading Ease score is already scaled such that lower values correspond to more difficult text and is therefore not transformed.

Figure 2: The SciBERT Model



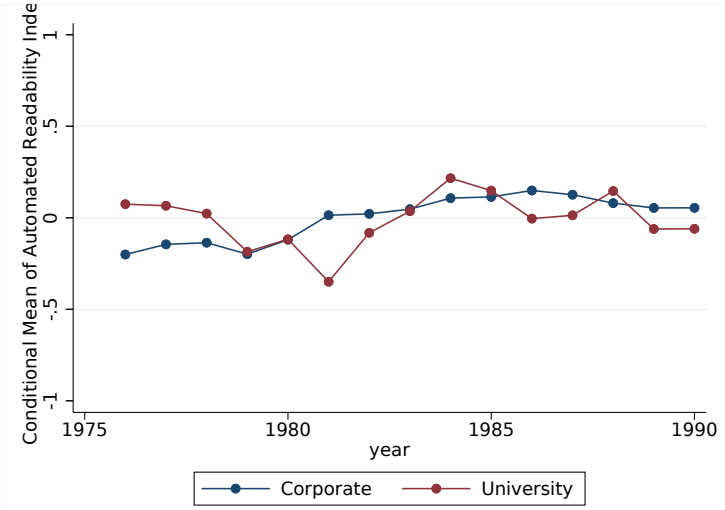
Notes: The figure displays the overview of the SciBERT model. SciBERT is a transformer-based language model built on BERT and pre-trained on 1.14 million scientific papers (Beltagy et al. 2019). It generates contextualized vector representations (known as embeddings) that capture meanings of text. The underlying BERT framework (Devlin et al. 2018) uses a deep neural network with a transformer architecture (Vaswani et al. 2017) to process text bidirectionally. It is trained on two tasks: predicting randomly masked tokens and determining sentence continuity. This architecture allows SciBERT to encode domain-specific language and understand textual dependencies, making it particularly suited for analyzing scientific and technical writing.

Figure 3: Construction of AI-based Readability Scores derived from SciBERT

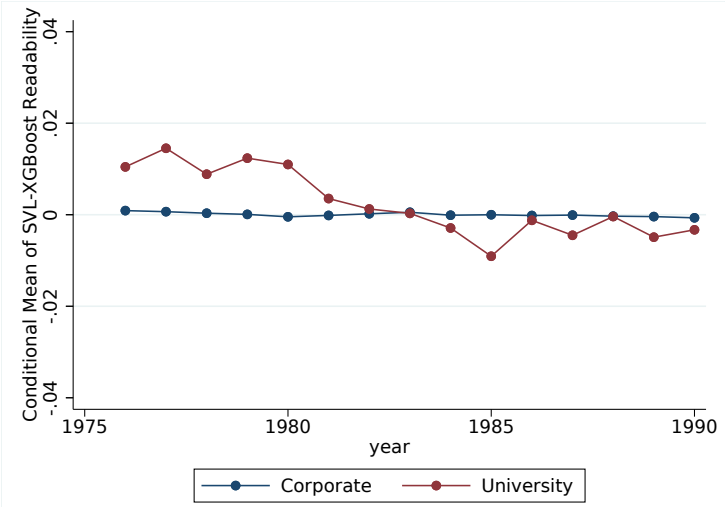


Notes: The figure displays the process of the construction of the AI-based readability scores derived from SciBERT. First, for each rejected patent application, the detailed description text is processed using the SciBERT language model to generate contextualized embeddings. Second, the resulting embeddings serve as inputs to supervised learning models –Ordinary Least Squares (OLS) and XGBoost– which are trained to predict 112(a) rejection outcomes. Third, trained models are then applied to all patents in the main sample to generate predicted rejection probabilities. These probabilities are interpreted as AI-based readability scores, with higher values indicating greater textual obscurity.

Figure 4: Conditional Mean of Readability Over Time



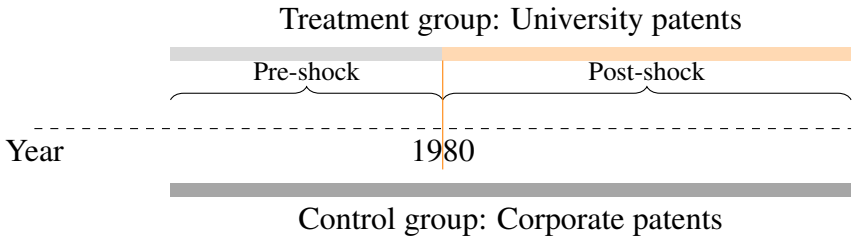
(a) Traditional Readability Score



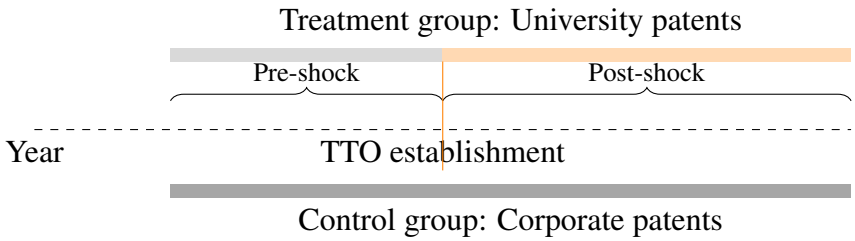
(b) AI-based Readability Score

Notes: The figure displays the (conditional) mean of readability of university patents and corporate patents over time, conditionally on technology classifications (CPC codes) and assignees. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text.

Figure 5: Illustration of Two DID Settings



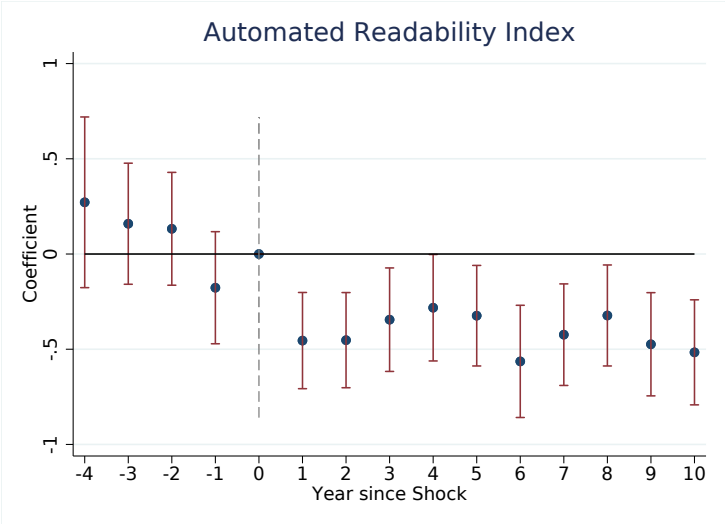
(a) Illustration of DID Setting I (The Bayh-Dole Act in 1980)



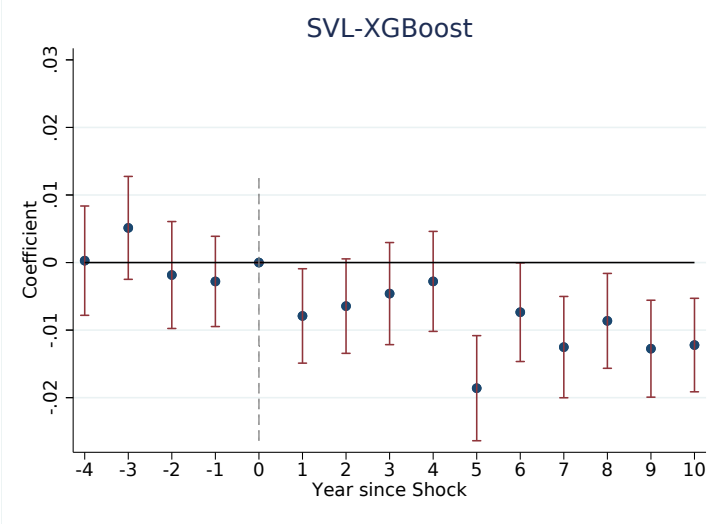
(b) Illustration of DID Setting II (TTO establishment)

Notes: This figure displays two difference-in-differences (DID) settings. Subfigure (a) illustrates the use of the 1980 Bayh-Dole Act as the policy shock. Subfigure (b) uses the staggered establishment of university Technology Transfer Offices (TTOs) as the shock. Curly braces indicate the pre-shock and post-shock periods.

Figure 6: DID Setting I: Readability of Detailed Description by 1980 Bayh-Dole Act Over Year



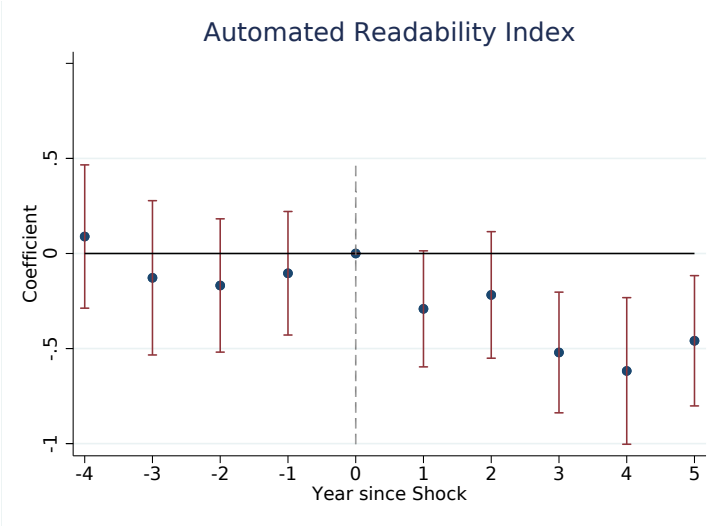
(a) Traditional Readability Score (Automated Readability Index)



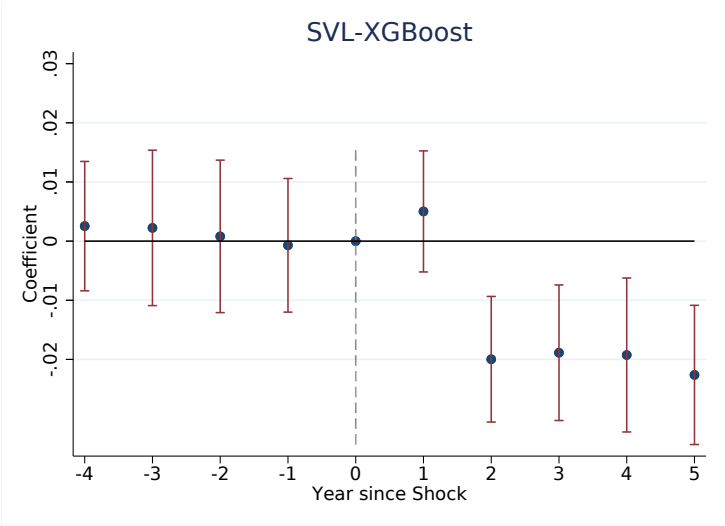
(b) AI-based Readability Scores Derived from SciBERT (SVL-XGBoost)

Notes: The figure displays changes in readability around 1980 Bayh-Dole Act relative to the control group. Subfigure (a) displays traditional readability score (Automated Readability Index), and subfigure (b) shows shows SciBERT readability score (SVL-XGBoost readability score). As the Act was implemented after December 1980, I conduct a regression where the treatment effect is interacted with indicator variables representing each year relative to 1980. The 95 percent confidence interval around the estimates is plotted with vertical red lines. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text.

Figure 7: DID Setting II: Readability of detailed description by TTO establishment Over Year



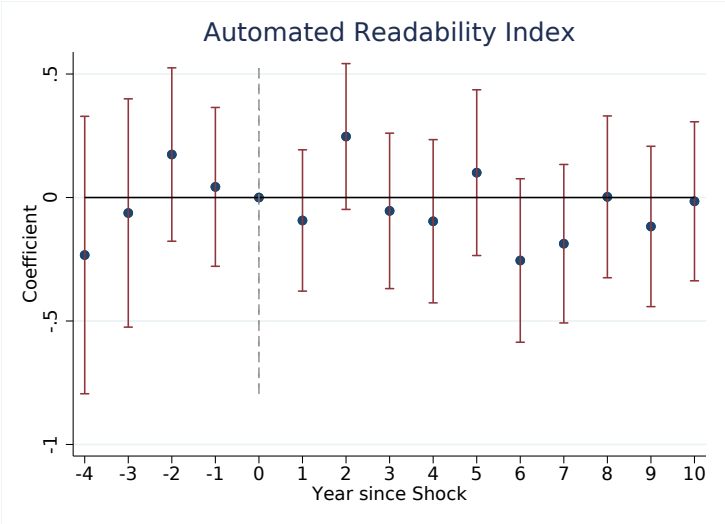
(a) Traditional Readability Score (Automated Readability Index)



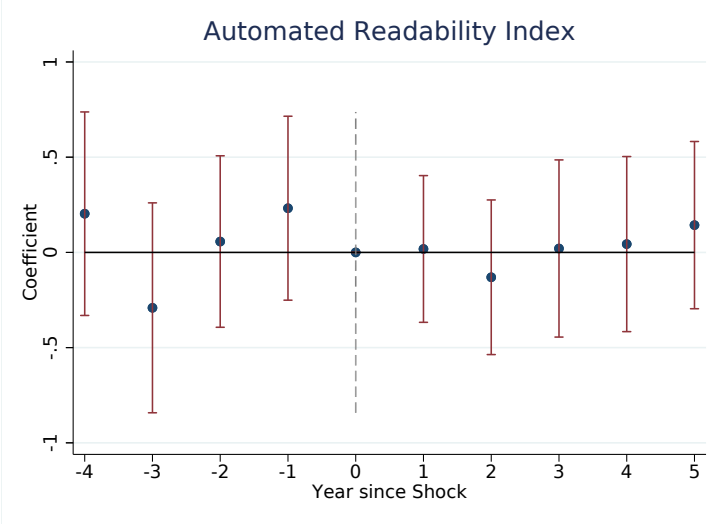
(b) AI-based Readability Scores Derived from SciBERT (SVL-XGBoost)

Notes: The figure displays changes in readability around TTO establishment relative to the control group. Subfigure (a) displays traditional readability score (Automated Readability Index), and subfigure (b) shows shows SciBERT readability score (SVL-XGBoost readability score). The 95 percent confidence interval around the estimates is plotted with vertical red lines. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text.

Figure 8: Readability Change of Patent Background and Summary Text Over Year



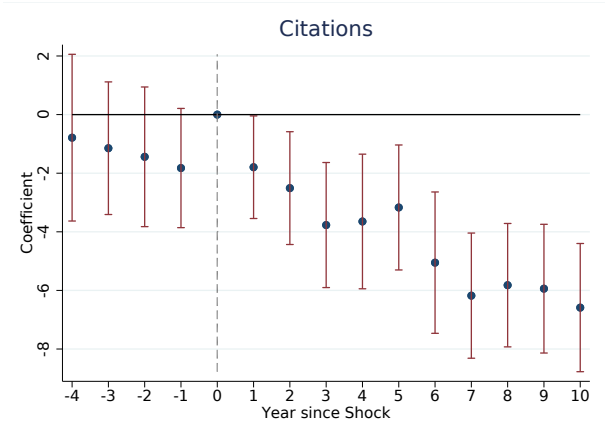
(a) DID setting I: Patent Background and Summary



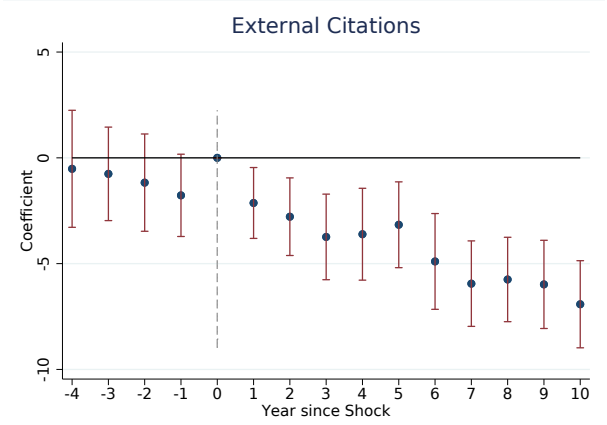
(b) DID setting II: Patent Background and Summary

Notes: The figure displays changes in readability change of patent background and summary text over years. Panel (a) shows the regression result in the DID I setting at the inventor level, and panel (b) shows the regression result in the DID II setting at the inventor level. The 95 percent confidence interval around the estimates is plotted with vertical red lines. The readability measure used is the Automated Readability Index. It is originally scaled such that higher values indicate lower readability (i.e., more difficult to understand), I multiply scores by -1, so that lower values correspond to harder-to-read text.

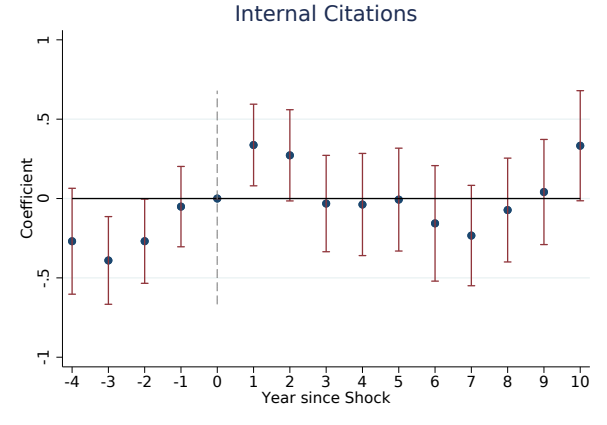
Figure 9: Citation Change by the 1980 Bayh-Dole Act Over Year



(a) Citations Change Over Year



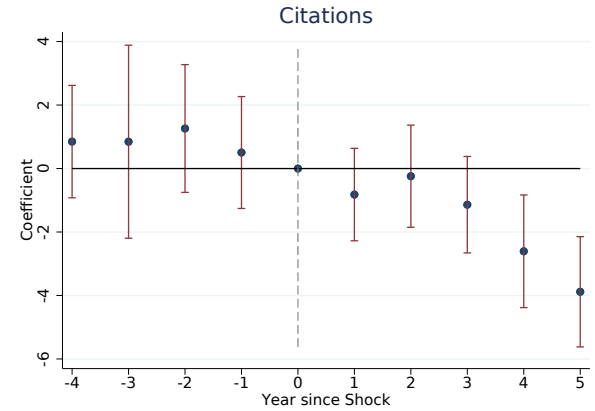
(b) External Citations Change Over Year



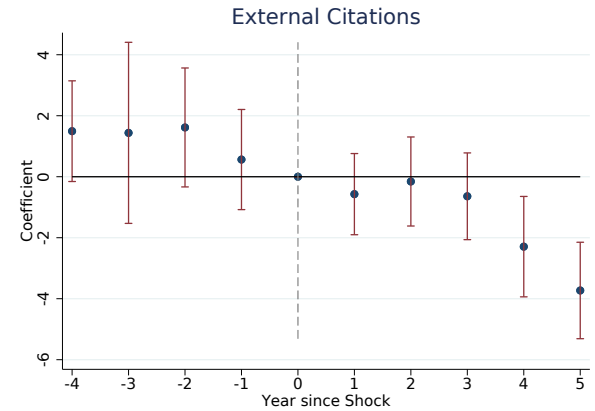
(c) Self-Citations Change Over Year

Notes: The figure displays changes in citations around 1980 Bayh-Dole Act relative to the control group. Subfigure (a) displays the overall citations received from patents, subfigure (b) shows external citations (citations excluding the self-citations), and subfigure (c) shows the self-citations. The 95 percent confidence interval around the estimates is plotted with vertical red lines. As the Act was implemented after December 1980, I conduct a regression where the treatment effect is interacted with indicator variables representing each year relative to 1980.

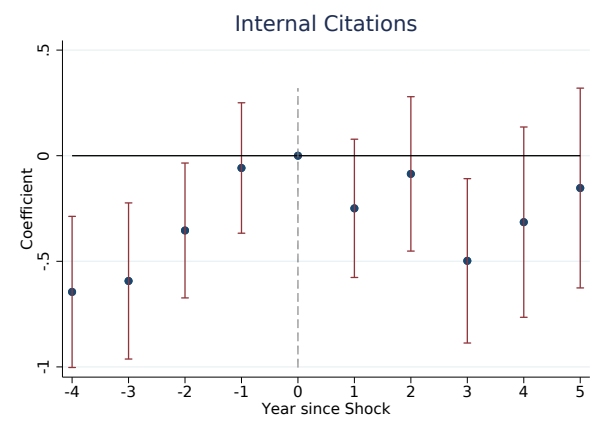
Figure 10: Citation Change by the TTO Establishment Over Year Over Year



(a) Citations Change Over Year



(b) External Citations Change Over Year



(c) Self-Citations Change Over Year

Notes: The figure displays changes in citations around the TTO establishment over year relative to the control group. Subfigure (a) displays the overall citations received from patents, subfigure (b) shows external citations (citations excluding the self-citations), and subfigure (c) shows the self-citations. The 95 percent confidence interval around the estimates is plotted with vertical red lines.

Tables

Table 1: Summary Statistics

Panel A: University Observations (N= 15011)

Variables	Mean	SD	Min	P25	Median	P75	Max
Automated Readability Index	-12.680	2.382	-42.700	-14.000	-12.500	-11.200	-4.200
Flesch-Kincaid Grade	-10.380	2.320	-25.500	-11.800	-10.400	-8.900	-3.000
Coleman-Liau Index	-11.918	2.074	-41.400	-13.090	-11.840	-10.680	0.340
Flesch Reading Ease	52.331	9.277	-11.940	46.510	52.260	58.080	88.130
SVL-XGBoost	-0.182	0.089	-0.709	-0.212	-0.153	-0.125	-0.008
SVL-OLS	-0.192	0.074	-0.556	-0.219	-0.174	-0.143	-0.007
Total Citations	10.565	14.130	0.000	3.000	6.000	13.000	292.000
External Citations	9.638	13.503	0.000	2.000	5.833	12.000	292.000
Internal Citations	0.928	2.100	0.000	0.000	0.000	1.000	36.094

Panel B: Corporate Observations (N= 599496)

Variables	Mean	SD	Min	P25	Median	P75	Max
Automated Readability Index	-13.124	3.316	-287.600	-14.600	-12.800	-11.200	1.900
Flesch-Kincaid Grade	-10.885	2.721	-131.200	-12.600	-10.900	-9.100	1.500
Coleman-Liau Index	-11.262	2.837	-210.970	-12.310	-10.910	-9.750	5.320
Flesch Reading Ease	53.601	10.266	-584.830	47.490	53.810	59.740	110.630
SVL-XGBoost	-0.143	0.052	-1.104	-0.158	-0.134	-0.114	0.059
SVL-OLS	-0.148	0.043	-0.577	-0.172	-0.145	-0.121	0.044
Total Citations	8.414	11.410	0.000	2.000	5.000	11.000	631.000
External Citations	6.987	9.891	0.000	1.333	4.000	9.000	619.453
Internal Citations	1.427	3.817	0.000	0.000	0.000	1.250	210.870

Notes: This table reports summary statistics for the baseline sample, including both traditional readability scores and AI-based readability scores for the detailed descriptions and citations of patent observations. Except for the Flesch Reading Ease, all readability scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score is already scaled such that lower values indicate more difficult text and is not transformed.

Table 2: DID Setting I: Readability Change of Patent Detailed Description

Panel A: Traditional Readability Score				
	Assignee Level		Inventor Level	
	(1)	(2)	(3)	(4)
Post × Treated	-0.246**	-0.259***	-0.453***	-0.454***
	(0.099)	(0.098)	(0.093)	(0.093)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE			X	X
ClassficationFE		X		X
R-squared	0.192	0.201	0.522	0.523
Observations	418271	418270	614507	614507
Panel B: AI-based Readability Score				
	Assignee Level		Inventor Level	
	(1)	(2)	(3)	(4)
Post × Treated	-0.016***	-0.016***	-0.009***	-0.009***
	(0.002)	(0.003)	(0.002)	(0.002)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE			X	X
ClassficationFE		X		X
R-squared	0.159	0.182	0.412	0.414
Observations	418271	418270	614507	614507

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act) for patent detailed descriptions. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. In Panel A, the dependent variable is the transformed traditional score ($-1\times$ Automated Readability Index); in Panel B, it is the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). Columns (1)-(2) are on the assignee level, and columns (3)-(4) are on the inventor level. Standard errors in columns (1)-(2) are clustered at assignee levels. Standard errors in columns (3)-(4) are clustered at inventor levels.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table 3: DID Setting II: Readability Change of Patent Detailed Description

Panel A: Traditional Readability Score				
	Assignee Level		Inventor Level	
	(1)	(2)	(3)	(4)
Post × Treated	-0.289**	-0.312**	-0.358**	-0.356**
	(0.128)	(0.131)	(0.160)	(0.161)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE			X	X
ClassficationFE		X		X
R-squared	0.192	0.200	0.522	0.523
Observations	417246	417245	612742	612742
Panel B: AI-based Readability Score				
	Assignee Level		Inventor Level	
	(1)	(2)	(3)	(4)
Post × Treated	-0.016***	-0.015***	-0.015***	-0.014***
	(0.004)	(0.005)	(0.004)	(0.004)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE			X	X
ClassficationFE		X		X
R-squared	0.158	0.182	0.412	0.414
Observations	417246	417245	612742	612742

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting II (TTO establishment) for patent detailed descriptions. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. In Panel A, the dependent variable is the transformed traditional score ($-1\times$ Automated Readability Index); in Panel B, it is the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). Columns (1)-(2) are on the assignee level, and columns (3)-(4) are on the inventor level. Standard errors in columns (1)-(2) are clustered at assignee levels. Standard errors in columns (3)-(4) are clustered at inventor levels.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table 4: Readability of Background and Summary in Two DID Settings

Panel A: DID Setting I (Bayh-Dole Act)				
	Assignee Level		Inventor Level	
	(1)	(2)	(3)	(4)
Post × Treated	-0.005	-0.025	-0.045	-0.039
	(0.134)	(0.137)	(0.110)	(0.109)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE			X	X
ClassficationFE		X		X
R-squared	0.233	0.244	0.564	0.565
Observations	418271	418270	614507	614507
Panel B: DID Setting II (TTO Establishment)				
	Assignee Level		Inventor Level	
	(1)	(2)	(3)	(4)
Post × Treated	-0.007	-0.022	-0.046	-0.055
	(0.140)	(0.139)	(0.155)	(0.155)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE			X	X
ClassficationFE		X		X
R-squared	0.232	0.244	0.564	0.565
Observations	417246	417245	612742	612742

Notes: This table reports the estimated coefficients and standard errors (in parentheses) from two difference-in-differences specifications examining the impact of policy shocks on the readability of patent background and summary sections. Panel A uses DID Setting I based on the 1980 Bayh-Dole Act. Panel B uses DID Setting II based on university TTO establishment. The dependent variable in all columns is the transformed traditional score ($-1\times$ Automated Readability Index). Columns (1)-(2) are on the assignee level, and columns (3)-(4) are on the inventor level. Standard errors in columns (1)-(2) are clustered at assignee levels. Standard errors in columns (3)-(4) are clustered at inventor levels.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table 5: Heterogeneous Effects across Technology Fields in Two DID Settings

Panel A: DID Setting I (Bayh-Dole Act)						
	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Competitive Fields	0.055***	-0.042**	-0.001**	0.001	0.057***	0.088***
	(0.013)	(0.021)	(0.000)	(0.000)	(0.019)	(0.033)
Post × Treated	-0.204*	-0.206*	-0.002	-0.003	-0.032	-0.004
	(0.116)	(0.115)	(0.003)	(0.003)	(0.158)	(0.159)
Post × Treated × Competitive Fields	-0.234**	-0.231**	-0.007***	-0.006***	-0.094	-0.118
	(0.092)	(0.092)	(0.002)	(0.002)	(0.131)	(0.132)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.527	0.527	0.420	0.421	0.570	0.571
Observations	575171	575171	575171	575171	575171	575171
Panel B: DID Setting II (TTO Establishment)						
	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Competitive Fields	0.054***	-0.042**	-0.001**	0.001	0.057***	0.088***
	(0.013)	(0.021)	(0.000)	(0.000)	(0.019)	(0.033)
Post × Treated	-0.157	-0.162	-0.005	-0.005	-0.055	-0.038
	(0.193)	(0.194)	(0.005)	(0.005)	(0.202)	(0.203)
Post × Treated × Competitive Fields	-0.201*	-0.188	-0.009***	-0.008***	-0.026	-0.045
	(0.117)	(0.119)	(0.003)	(0.003)	(0.150)	(0.153)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.526	0.527	0.420	0.421	0.569	0.570
Observations	573410	573410	573410	573410	573410	573410

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in regressions on heterogeneous effects across technology competition under two DID settings. Panel A uses DID Setting I based on the 1980 Bayh-Dole Act. Panel B uses DID Setting II based on university TTO establishment. Competitive Fields is a dummy variable equal to 1 if the technology field falls within the bottom 50% of the Herfindahl-Hirschman Index (HHI), indicating higher competition, and 0 otherwise. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table 6: Impact of the Bayh-Dole Act on University Inventors' Disclosure Behavior (Before TTO establishment)

	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.389***	-0.384***	-0.004*	-0.004*	-0.180	-0.172
	(0.115)	(0.114)	(0.002)	(0.002)	(0.144)	(0.143)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.523	0.524	0.396	0.397	0.564	0.565
Observations	605831	605831	605831	605831	605831	605831

Notes: This table presents the estimated coefficients and standard errors (in parentheses) from a DID analysis examining the impact of the Bayh-Dole Act on university inventors' disclosure behavior (pre-TTO establishment). Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table 7: Impact of TTO establishment on University Inventors' Disclosure Behavior (After Bayh-Dole Act)

	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.326**	-0.315*	-0.002	-0.001	-0.098	-0.086
	(0.160)	(0.161)	(0.005)	(0.005)	(0.157)	(0.157)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.578	0.579	0.459	0.460	0.594	0.595
Observations	417334	417334	417334	417334	417334	417334

Notes: This table reports the estimated coefficients and standard errors (in parentheses) from a DID analysis examining the impact of TTO establishment on university inventors' disclosure behavior, focusing on sample period after 1980. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table 8: Change of Follow-On Innovation in Two DID Settings

Panel A: DID Setting I (Bayh-Dole Act)			
	(1) Total Citations	(2) External Citations	(3) Internal Citations
Post × Treated	-3.216***	-3.437***	0.221**
	(0.669)	(0.636)	(0.092)
YearFE	X	X	X
AssigneeFE	X	X	X
InventorFE	X	X	X
ClassficationFE	X	X	X
R-squared	0.512	0.529	0.437
Observations	614408	614408	614408
Panel B: DID Setting II (TTO Establishment)			
	(1) Total Citations	(2) External Citations	(3) Internal Citations
Post × Treated	-2.313***	-2.325***	0.012
	(0.699)	(0.641)	(0.160)
YearFE	X	X	X
AssigneeFE	X	X	X
InventorFE	X	X	X
ClassficationFE	X	X	X
R-squared	0.512	0.529	0.437
Observations	612644	612644	612644

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in regressions on citations. Panel A uses DID Setting I based on the 1980 Bayh-Dole Act. Panel B uses DID Setting II based on university TTO establishment. Column (1) uses the total citations, column (2) uses external citations, and column (3) uses the internal citations. Standard errors are clustered at the inventor level.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table 9: Impact on Follow-On Innovation

(a) 2SLS Regression on Total Citations

	(1)	(2)
Traditional $\widehat{\text{Readability Score}}$	7.091***	
	(2.124)	
AI-based $\widehat{\text{Readability Score}}$		368.741***
		(117.483)
YearFE	X	X
AssigneeFE	X	X
InventorFE	X	X
ClassficationFE	X	X
Observations	614408	614408

(b) 2SLS Regression on External Citations

	(1)	(2)
Traditional $\widehat{\text{Readability Score}}$	7.579***	
	(2.140)	
AI-based $\widehat{\text{Readability Score}}$		394.094***
		(119.085)
YearFE	X	X
AssigneeFE	X	X
InventorFE	X	X
ClassficationFE	X	X
Observations	614408	614408

(c) 2SLS Regression on Internal Citations

	(1)	(2)
Traditional $\widehat{\text{Readability Score}}$	-0.488**	
	(0.220)	
AI-based $\widehat{\text{Readability Score}}$		-25.353**
		(11.683)
YearFE	X	X
AssigneeFE	X	X
InventorFE	X	X
ClassficationFE	X	X
Observations	614408	614408

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in 2SLS regressions on all citations, external citations, and internal citations. Table (a) is on total citations, Table (b) is on external citations, and Table (c) is on internal citations. Every column (1) uses the predicted Automated Readability Index, and column (2) uses the predicted SciBERT-XGBoost readability score. Standard errors are clustered at the inventor level. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Appendices

A Tables

Table A10: DID Setting I: Patent Detailed Description

	Traditional Readability Scores				AI-based Readability Scores	
	(1) Automated Readability Index	(2) Flesch-Kincaid Grade	(3) Coleman-Liau Index	(4) Flesch Reading Ease	(5) SVL XGBoost	(6) SVL OLS
Post × Treated	-0.454***	-0.277***	-0.234***	-1.164***	-0.009***	-0.004**
	(0.093)	(0.070)	(0.072)	(0.320)	(0.002)	(0.002)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X
R-squared	0.523	0.674	0.592	0.539	0.414	0.577
Observations	614507	614507	614507	614507	614507	614507

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act). Column (1) uses the Automated Readability Index score, Column (2) uses the Flesch-Kincaid Grade score, Column (3) uses the Coleman-Liau Index score, Column (4) uses the Flesch Reading Ease score, Column (5) uses the SVL-XGBoost readability score, and Column (6) uses the SVL-OLS readability score. Except for Column (4), all scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score (Column 4) is already scaled such that lower values indicate more difficult text and is not transformed. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A11: DID Setting II: Patent Detailed Description

	Traditional Readability Scores				AI-based Readability Scores	
	(1) Automated Readability Index	(2) Flesch-Kincaid Grade	(3) Coleman-Liau Index	(4) Flesch Reading Ease	(5) SVL XGBoost	(6) SVL OLS
Post × Treated	-0.356**	-0.144	-0.305**	-1.111**	-0.014***	-0.010***
	(0.161)	(0.096)	(0.134)	(0.533)	(0.004)	(0.003)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X
R-squared	0.523	0.674	0.592	0.539	0.414	0.578
Observations	612742	612742	612742	612742	612742	612742

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting II (TTO establishment). Column (1) uses the automated readability index score, column (2) uses Flesch-Kincaid grade score, column (3) uses the Coleman-Liau index score, column (4) uses the Flesch reading ease score, column (5) uses the SVL-XGBoost readability score, and column (6) uses the SVL-OLS readability score. Except for Column (4), all scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score (Column 4) is already scaled such that lower values indicate more difficult text and is not transformed. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A12: DID Setting I: Patent Background and Summary

	Multiple Readability Scores			
	(1) Automated Readability Index	(2) Flesch-Kincaid Grade	(3) Coleman-Liau Index	(4) Flesch Reading Ease
Post × Treated	-0.039	0.042	-0.006	0.295
	(0.109)	(0.090)	(0.054)	(0.348)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE	X	X	X	X
ClassficationFE	X	X	X	X
R-squared	0.565	0.614	0.599	0.552
Observations	614507	614507	614507	614507

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act) for patent background and summary text. Column (1) uses the automated readability index score, column (2) uses Flesch-Kincaid grade score, column (3) uses the Coleman-Liau index score, and column (4) uses the Flesch reading ease score. Except for Column (4), all scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score (Column 4) is already scaled such that lower values indicate more difficult text and is not transformed. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A13: DID Setting II: Patent Background and Summary

	Multiple Readability Scores			
	(1) Automated Readability Index	(2) Flesch-Kincaid Grade	(3) Coleman-Liau Index	(4) Flesch Reading Ease
Post × Treated	-0.055	0.088	-0.044	0.170
	(0.155)	(0.131)	(0.080)	(0.519)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE	X	X	X	X
ClassficationFE	X	X	X	X
R-squared	0.565	0.613	0.598	0.552
Observations	612742	612742	612742	612742

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting II (TTO establishment) for patent background and summary text. Columns (1) uses the automated readability index score, column (2) uses Flesch-Kincaid grade score, columns (3) uses the Coleman-Liau index score, and columns (4) uses the Flesch reading ease score. Except for Column (4), all scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score (Column 4) is already scaled such that lower values indicate more difficult text and is not transformed. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A14: Average Treatment Effect of Individual Inventor in DID setting II (CSDID)

Panel A: Traditional Readability Score	Coefficient	SE	Z	P-value	LL	UL
Automated Readability Index	-0.545	0.229	-2.379	0.017	-0.993	-0.096
Panel B: AI-based Readability Score	Coefficient	SE	Z	P-value	LL	UL
SVL-XGBoost	-0.022	0.007	-3.300	0.001	-0.035	-0.009

Notes: This table reports the average treatment effects using CSDID specification for DID setting II. In Panel A, the dependent variable is the transformed traditional score ($-1\times$ Automated Readability Index); in Panel B, it is the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). CSDID requires panel data with a single observation per unit per time period. To meet this requirement, I aggregate the patent data to the inventor-year level by computing the mean readability score for each inventor annually. This aggregation may reduce precision, as it limits the ability to control for patent-level characteristics such as technological fields. See Callaway and Sant'Anna (2021) for details.

Table A15: DID Setting I: Readability Change for Stayers

	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.515***	-0.512***	-0.009***	-0.009***	-0.184	-0.181
	(0.112)	(0.112)	(0.002)	(0.002)	(0.130)	(0.130)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.463	0.464	0.284	0.285	0.488	0.490
Observations	303937	303936	303937	303936	303937	303936

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act) for stayers. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A16: DID Setting II: Readability Change for Stayers

	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.483*	-0.484*	-0.012**	-0.011**	-0.263	-0.262
	(0.248)	(0.248)	(0.005)	(0.005)	(0.182)	(0.182)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.538	0.538	0.396	0.397	0.580	0.581
Observations	381688	381688	381688	381688	381688	381688

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting II (TTO establishment) for stayers. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A17: DID Setting I: Same-City Exclusion

	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.432***	-0.427***	-0.009***	-0.009***	-0.138	-0.129
	(0.097)	(0.096)	(0.002)	(0.002)	(0.117)	(0.117)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.557	0.558	0.443	0.444	0.578	0.579
Observations	331285	331285	331285	331285	331285	331285

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act) excluding control group located within a certain city of any treated university. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A18: DID Setting II: Same-City Exclusion

	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.313*	-0.310*	-0.014***	-0.013***	-0.090	-0.099
	(0.172)	(0.172)	(0.004)	(0.004)	(0.162)	(0.162)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.508	0.509	0.435	0.436	0.575	0.576
Observations	369604	369604	369604	369604	369604	369604

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting II (TTO establishment) excluding control group located within a certain city of any treated university. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A19: DID Setting I: Small Entity Exclusion

	Detailed Description				Background and Summary	
	Traditional Readability		AI-based Readability		Traditional Readability	
	(1)	(2)	(3)	(4)	(5)	(6)
Post × Treated	-0.442***	-0.439***	-0.007**	-0.007**	-0.205	-0.193
	(0.155)	(0.155)	(0.003)	(0.003)	(0.134)	(0.134)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE		X		X		X
R-squared	0.518	0.519	0.415	0.416	0.563	0.564
Observations	462459	462459	462459	462459	462459	462459

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act) excluding small entities. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. Columns (1), (2), (5), and (6) use the transformed traditional score ($-1\times$ Automated Readability Index). Columns (3)-(4) use the transformed AI-based score ($-1\times$ SciBERT-XGBoost Readability Score). The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A20: DID Setting I: Readability Change with Patent's Orientation as Control

	Detailed Description						Background and Summary		
	Traditional Readability			AI-based Readability			Traditional Readability		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post × Treated	-0.455***	-0.463***	-0.464***	-0.009***	-0.009***	-0.009***	-0.040	-0.042	-0.042
	(0.093)	(0.093)	(0.093)	(0.002)	(0.002)	(0.002)	(0.109)	(0.109)	(0.109)
Basicness (Share)	0.477***		0.454***	0.000		0.000	0.283***		0.277***
	(0.023)		(0.023)	(0.000)		(0.000)	(0.022)		(0.022)
Basicness (LDA)		0.733***	0.646***		0.003***	0.003***		0.239***	0.186***
		(0.043)	(0.043)		(0.001)	(0.001)		(0.046)	(0.046)
YearFE	X	X	X	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X	X	X	X
R-squared	0.524	0.524	0.525	0.414	0.414	0.414	0.565	0.565	0.565
Observations	614507	614507	614507	614507	614507	614507	614507	614507	614507

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in a baseline DID setting I (the 1980 Bayh-Dole Act) with patent's orientation features as controls (details in Appendix D.2). Details of the two control variables are provided in Appendix C.3. Columns (1)-(3) use the traditional readability score (Automated Readability Index) for detailed descriptions, while columns (4)-(6) use the SciBERT-XGBoost score. Columns (7)-(9) report the results for the background and summary sections using the traditional readability score. Columns (1)-(4) are on the patent detailed descriptions, and columns (5)-(6) are on the patent background and summary. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A21: DID Setting II: Readability Change with Patent's Orientation as Control

	Detailed Description						Background and Summary		
	Traditional Readability			AI-based Readability			Traditional Readability		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post × Treated	-0.354**	-0.353**	-0.352**	-0.014***	-0.014***	-0.014***	-0.054	-0.054	-0.053
	(0.164)	(0.163)	(0.166)	(0.004)	(0.004)	(0.004)	(0.155)	(0.155)	(0.155)
Basicness (Share)	0.474***		0.451***	0.000		-0.000	0.280***		0.274***
	(0.023)		(0.023)	(0.000)		(0.000)	(0.022)		(0.022)
Basicness (LDA)		0.732***	0.645***		0.003***	0.003***		0.243***	0.190***
		(0.043)	(0.043)		(0.001)	(0.001)		(0.046)	(0.046)
YearFE	X	X	X	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X	X	X	X
R-squared	0.524	0.524	0.524	0.414	0.414	0.414	0.565	0.565	0.565
Observations	612742	612742	612742	612742	612742	612742	612742	612742	612742

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in a baseline DID setting II (TTO establishment) with patent's orientation features as controls (details in Appendix D.2). Details of the two control variables are provided in Appendix C.3. Columns (1)-(3) use the traditional readability score (Automated Readability Index) for detailed descriptions, while columns (4)-(6) use the SciBERT-XGBoost readability score. Columns (7)-(9) report the results for the background and summary sections using the traditional readability score. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A22: DID Setting I: Readability Change with Patent's Originality and Generality as Control

	Detailed Description						Background and Summary		
	Traditional Readability			AI-based Readability			Traditional Readability		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post × Treated	-0.455***	-0.454***	-0.455***	-0.009***	-0.009***	-0.009***	-0.020	-0.038	-0.020
	(0.095)	(0.093)	(0.095)	(0.002)	(0.002)	(0.002)	(0.110)	(0.109)	(0.110)
Originality	-0.011		-0.013	0.000		0.000	0.003		0.009
	(0.024)		(0.024)	(0.000)		(0.000)	(0.019)		(0.019)
Generality		0.012	0.015		0.001***	0.001**		-0.053***	-0.048**
		(0.019)	(0.021)		(0.000)	(0.000)		(0.019)	(0.020)
YearFE	X	X	X	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X	X	X	X
R-squared	0.523	0.523	0.523	0.416	0.414	0.416	0.566	0.565	0.566
Observations	603732	614400	603724	603732	614400	603724	603732	614400	603724

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in a baseline DID setting I (the 1980 Bayh-Dole Act) with patent's originality and generality features as controls (details in Appendix D.2). Columns (1)-(3) use the traditional readability score (Automated Readability Index) for detailed descriptions, while columns (4)-(6) use the SciBERT-XGBoost readability score. Columns (7)-(9) report the results for the background and summary sections using the traditional readability score. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A23: DID Setting II: Readability Change with Patent's Originality and Generality as Control

	Detailed Description						Background and Summary		
	Traditional Readability			AI-based Readability			Traditional Readability		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Post × Treated	-0.358**	-0.354**	-0.358**	-0.015***	-0.015***	-0.015***	-0.011	-0.046	-0.011
	(0.164)	(0.161)	(0.164)	(0.004)	(0.004)	(0.004)	(0.156)	(0.155)	(0.156)
Originality	-0.009		-0.011	0.000		0.000	-0.001		0.005
	(0.024)		(0.024)	(0.000)		(0.000)	(0.019)		(0.019)
Generality		0.009	0.012		0.001***	0.001**		-0.053***	-0.048**
		(0.019)	(0.021)		(0.000)	(0.000)		(0.019)	(0.020)
YearFE	X	X	X	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X	X	X	X
R-squared	0.523	0.523	0.523	0.415	0.414	0.415	0.566	0.565	0.566
Observations	601996	612636	601988	601996	612636	601988	601996	612636	601988

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in a baseline DID setting II (TTO establishment) with patent's originality and generality features as controls (details in Appendix D.2). Columns (1)-(3) use the traditional readability score (Automated Readability Index) for detailed descriptions, while columns (4)-(6) use the SciBERT-XGBoost readability score. Columns (7)-(9) report the results for the background and summary sections using the traditional readability score. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A24: DID Setting I: Readability Change Adding Journal Fixed Effects

	Detailed Description		Background and Summary
	Traditional Readability	AI-based Readability	Traditional Readability
	(1)	(2)	(3)
Post × Treated	-0.413***	-0.010***	-0.032
	(0.090)	(0.002)	(0.112)
YearFE	X	X	X
ClassficationFE	X	X	X
AssigneeFE	X	X	X
InventorFE	X	X	X
JournalFE	X	X	X
R-squared	0.526	0.419	0.566
Observations	614103	614103	614103

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act) adding journal fixed effects (details in Appendix D.2). Column (1) uses the traditional readability score (Automated Readability Index) for detailed descriptions, while column (2) uses the SciBERT-XGBoost readability score. Column (3) reports the results for the background and summary sections using the traditional readability score. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A25: DID Setting II: Readability Change Adding Journal Fixed Effects

	Detailed Description		Background and Summary
	Traditional Readability	AI-based Readability	Traditional Readability
	(1)	(2)	(3)
Post × Treated	-0.379**	-0.015***	-0.087
	(0.176)	(0.004)	(0.160)
YearFE	X	X	X
ClassficationFE	X	X	X
AssigneeFE	X	X	X
InventorFE	X	X	X
JournalFE	X	X	X
R-squared	0.525	0.419	0.566
Observations	612338	612338	612338

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline under DID setting II (TTO establishment) adding journal fixed effects (details in Appendix D.2). Column (1) uses the traditional readability score (Automated Readability Index) for detailed descriptions, while column (2) uses the SciBERT-XGBoost readability score. Column (3) reports the results for the background and summary sections using the traditional readability score. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level.
*** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A26: DID Setting I: Readability Change Adding Attorney Fixed Effects

	Detailed Description		Background and Summary
	Traditional Readability	AI-based Readability	Traditional Readability
	(1)	(2)	(3)
Post × Treated	-0.423***	-0.007***	0.048
	(0.094)	(0.002)	(0.121)
YearFE	X	X	X
AssigneeFE	X	X	X
InventorFE	X	X	X
ClassficationFE	X	X	X
AttorneyFE	X	X	X
R-squared	0.561	0.442	0.608
Observations	610737	610737	610737

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting I (the 1980 Bayh-Dole Act) adding attorney fixed effects (details in Appendix D.2). Column (1) uses the traditional readability score (Automated Readability Index) for detailed descriptions, while column (2) uses the SciBERT-XGBoost readability score. Column (3) reports the results for the background and summary sections using the traditional readability score. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A27: DID Setting II: Readability Change Adding Attorney Fixed Effects

	Detailed Description		Background and Summary
	Traditional Readability	SciBERT-Readability	Traditional Readability
	(1)	(2)	(3)
Post × Treated	-0.351**	-0.011***	0.010
	(0.152)	(0.004)	(0.156)
YearFE	X	X	X
AssigneeFE	X	X	X
InventorFE	X	X	X
ClassficationFE	X	X	X
AttorneyFE	X	X	X
R-squared	0.561	0.441	0.608
Observations	608962	608962	608962

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in baseline DID setting II (TTO establishment) adding attorney fixed effects (details in Appendix D.2). Column (1) uses the traditional readability score (Automated Readability Index) for detailed descriptions, while column (2) uses the SciBERT-XGBoost readability score. Column (3) reports the results for the background and summary sections using the traditional readability score. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. Standard errors are clustered at the inventor level. *** denotes statistical significance at the 1% level, ** at the 5% level, and * at the 10% level.

Table A28: Impact on Follow-On Innovation: Multiple Readability Measures

(a) 2SLS Regression on All Citations

	(1)	(2)	(3)	(4)	(5)	(6)
Automated $\widehat{\text{Readability Index}}$	7.091***					
	(2.124)					
Flesch-$\widehat{\text{Kincaid Grade}}$		11.599**				
		(3.825)				
$\widehat{\text{Coleman-Liau Index}}$			13.780**			
			(5.360)			
Flesch $\widehat{\text{Reading Ease}}$				2.765**		
				(0.981)		
$\widehat{\text{SVL-OLS}}$					857.895**	
					(406.703)	
$\widehat{\text{SVL-XGBoost}}$						368.741***
						(117.483)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X
Observations	614408	614408	614408	614408	614408	614408

(b) 2SLS Regression on External Citations

	(1)	(2)	(3)	(4)	(5)	(6)
Automated $\widehat{\text{Readability Index}}$	7.579***					
	(2.140)					
Flesch-$\widehat{\text{Kincaid Grade}}$		12.396***				
		(3.909)				
$\widehat{\text{Coleman-Liau Index}}$			14.728***			
			(5.496)			
Flesch $\widehat{\text{Reading Ease}}$				2.955***		
				(1.002)		
$\widehat{\text{SVL-OLS}}$					916.879**	
					(422.227)	
$\widehat{\text{SVL-XGBoost}}$						394.094***
						(119.085)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X
Observations	614408	614408	614408	614408	614408	614408

(c) 2SLS Regression on Internal Citations

	(1)	(2)	(3)	(4)	(5)	(6)
Automated $\widehat{\text{Readability Index}}$	-0.488**					
	(0.220)					
Flesch-$\widehat{\text{Kincaid Grade}}$		-0.797**				
		(0.386)				
$\widehat{\text{Coleman-Liau Index}}$			-0.947**			
			(0.458)			
Flesch $\widehat{\text{Reading Ease}}$				-0.190**		
				(0.091)		
$\widehat{\text{SVL-OLS}}$					-58.984*	
					(32.149)	
$\widehat{\text{SVL-XGBoost}}$						-25.353**
						(11.683)
YearFE	X	X	X	X	X	X
AssigneeFE	X	X	X	X	X	X
InventorFE	X	X	X	X	X	X
ClassficationFE	X	X	X	X	X	X
Observations	614408	614408	614408	614408	614408	614408

Notes: This table reports the estimated coefficients and standard errors (in parentheses) in 2SLS regressions on citations. Table (a) is on total citations, Table (b) is on external citations, and Table (c) is on internal citations. Column (1) uses the predicted Automated Readability Index, column (2) uses the predicted Flesch-Kincaid Grade Level, column (3) uses the predicted Coleman-Liau Index, column (4) uses predicted Flesch Reading Ease, column (5) uses predicted SVL-OLS readability score, and column (6) uses predicted SVL-XGBoost readability score. Except for Column (4), all scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score (Column 4) is already scaled such that lower values indicate more difficult text and is not transformed. Standard errors are clustered at the inventor level.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

Table A29: Readability Change of Patent Claims

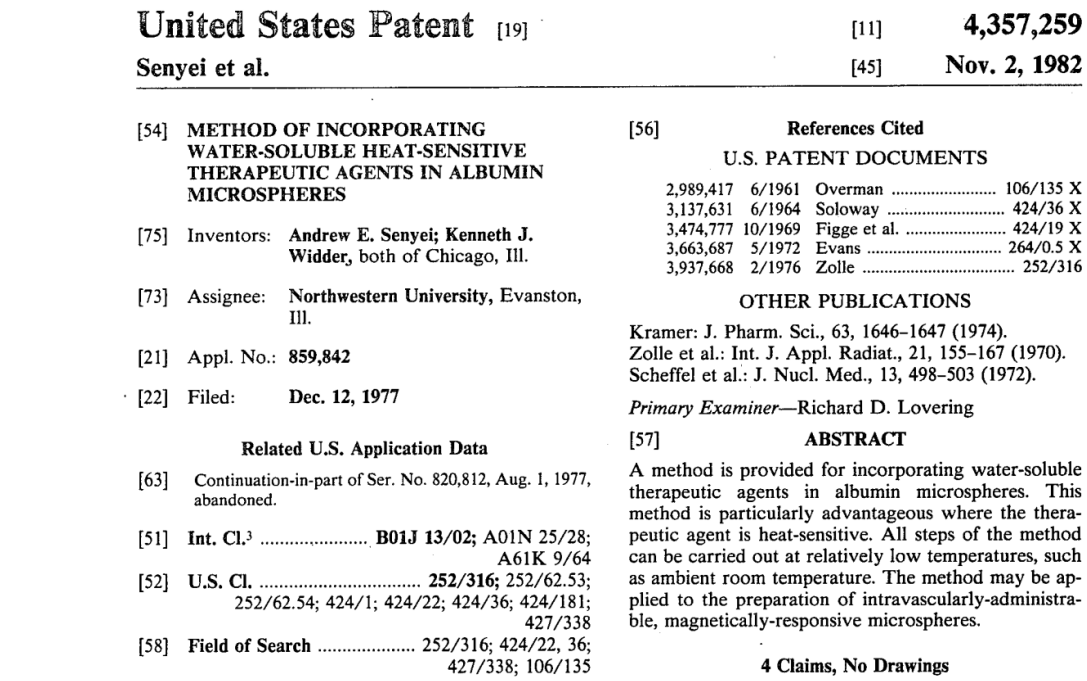
	Bayh-Dole Act		TTO Establishment	
	(1)	(2)	(3)	(4)
Post × Treated	-0.969	-0.961	0.820	0.801
	(0.600)	(0.599)	(0.826)	(0.824)
YearFE	X	X	X	X
AssigneeFE	X	X	X	X
InventorFE	X	X	X	X
ClassficationFE		X		X
R-squared	0.589	0.591	0.589	0.591
Observations	614490	614490	612725	612725

Notes: This table reports estimated coefficients for regressions using the claims section. Columns (1)–(2) use the Bayh-Dole Act as the shock, while columns (3)–(4) use TTO establishment as the shock. Because patents typically contain multiple claims, I construct a patent-level measure of claims readability by averaging readability scores across all claims within a patent. The dependent variable in all columns is the transformed traditional readability score ($-1\times$ Automated Readability Index). Standard errors are clustered at the inventor level.

*** denotes statistical significance at 1% level, ** 5% level, * 10% level.

B Figures

Figure A1: An Example of Patent Document (Basic Information)



United States Patent [19]

Senyei et al.

[11] **4,357,259**

[45] **Nov. 2, 1982**

[54] **METHOD OF INCORPORATING WATER-SOLUBLE HEAT-SENSITIVE THERAPEUTIC AGENTS IN ALBUMIN MICROSPHERES**

[75] Inventors: **Andrew E. Senyei; Kenneth J. Widder,** both of Chicago, Ill.

[73] Assignee: **Northwestern University,** Evanston, Ill.

[21] Appl. No.: **859,842**

[22] Filed: **Dec. 12, 1977**

Related U.S. Application Data

[63] Continuation-in-part of Ser. No. 820,812, Aug. 1, 1977, abandoned.

[51] **Int. Cl.**[3] **B01J 13/02;** A01N 25/28; A61K 9/64

[52] **U.S. Cl.** **252/316;** 252/62.53; 252/62.54; 424/1; 424/22; 424/36; 424/181; 427/338

[58] **Field of Search** 252/316; 424/22, 36; 427/338; 106/135

[56] **References Cited**

U.S. PATENT DOCUMENTS

2,989,417	6/1961	Overman	106/135 X
3,137,631	6/1964	Soloway	424/36 X
3,474,777	10/1969	Figge et al.	424/19 X
3,663,687	5/1972	Evans	264/0.5 X
3,937,668	2/1976	Zolle	252/316

OTHER PUBLICATIONS

Kramer: J. Pharm. Sci., 63, 1646–1647 (1974).
Zolle et al.: Int. J. Appl. Radiat., 21, 155–167 (1970).
Scheffel et al.: J. Nucl. Med., 13, 498–503 (1972).

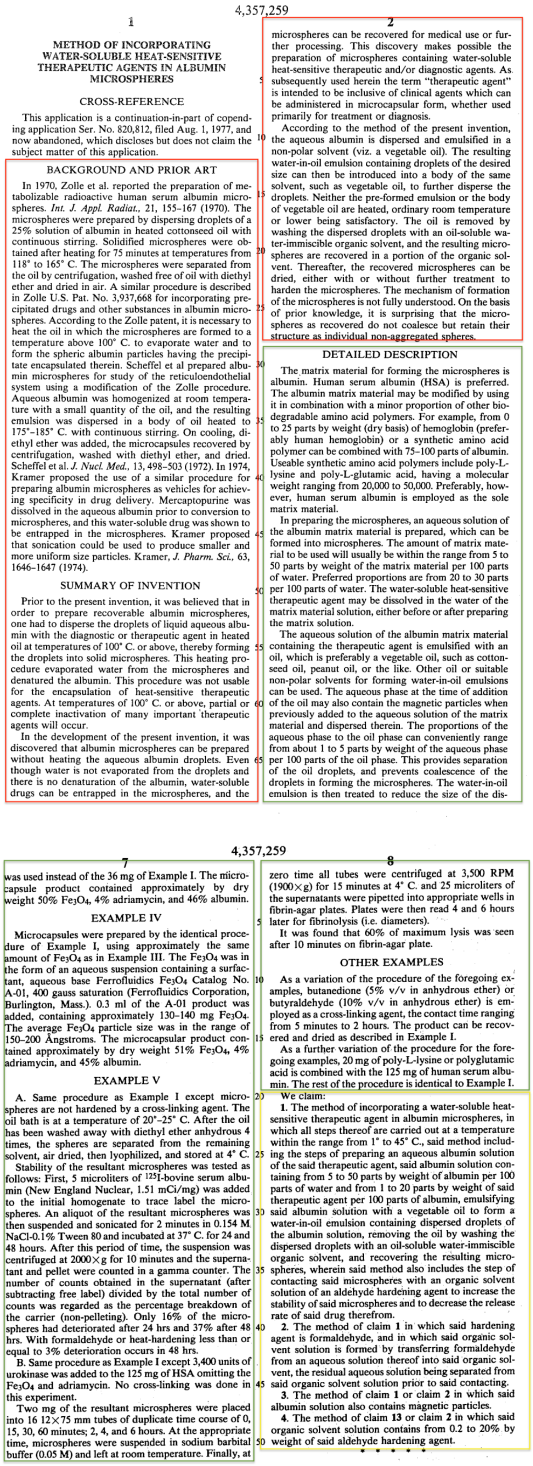
Primary Examiner—Richard D. Lovering

[57] **ABSTRACT**

A method is provided for incorporating water-soluble therapeutic agents in albumin microspheres. This method is particularly advantageous where the therapeutic agent is heat-sensitive. All steps of the method can be carried out at relatively low temperatures, such as ambient room temperature. The method may be applied to the preparation of intravascularly-administrable, magnetically-responsive microspheres.

4 Claims, No Drawings

Figure A2: An Example of Patent Document (Text)



4,357,259

1

METHOD OF INCORPORATING WATER-SOLUBLE HEAT-SENSITIVE THERAPEUTIC AGENTS IN ALBUMIN MICROSPHERES

CROSS-REFERENCE

This application is a continuation-in-part of copending application Ser. No. 820,812, filed Aug. 1, 1977, and now abandoned, which discloses but does not claim the subject matter of this application.

BACKGROUND AND PRIOR ART

In 1970, Zolle et al. reported the preparation of metabolizable radioactive human serum albumin microspheres. *Int. J. Appl. Radiat.*, 21, 155–167 (1970). The microspheres were prepared by dispersing droplets of a 25% solution of albumin in heated cottonseed oil with continuous stirring. Solidified microspheres were obtained after heating for 75 minutes at temperatures from 118° to 165° C. The microspheres were separated from the oil by centrifugation, washed free of oil with diethyl ether and dried in air. A similar procedure is described in Zolle U.S. Pat. No. 3,937,668 for incorporating precipitated drugs and other substances in albumin microspheres. According to the Zolle patent, it is necessary to heat the oil in which the microspheres are formed to a temperature above 100° C. to evaporate water and to form the spheric albumin particles having the precipitate encapsulated therein. Scheffel et al prepared albumin microspheres for study of the reticuloendothelial system using a modification of the Zolle procedure. Aqueous albumin was homogenized at room temperature with a small quantity of the oil, and the resulting emulsion was dispersed in a body of oil heated to 175°–185° C. with continuous stirring. On cooling, diethyl ether was added, the microcapsules recovered by centrifugation, washed with diethyl ether, and dried. Scheffel et al. *J. Nucl. Med.*, 13, 498–503 (1972). In 1974, Kramer proposed the use of a similar procedure for preparing albumin microspheres as vehicles for achieving specificity in drug delivery. Mercaptopurine was dissolved in the aqueous albumin prior to conversion to microspheres, and this water-soluble drug was shown to be entrapped in the microspheres. Kramer proposed that sonication could be used to produce smaller and more uniform size particles. Kramer, *J. Pharm. Sci.*, 63, 1646–1647 (1974).

SUMMARY OF INVENTION

Prior to the present invention, it was believed that in order to prepare recoverable albumin microspheres, one had to disperse the droplets of liquid aqueous albumin with the diagnostic or therapeutic agent in heated oil at temperatures of 100° C. or above, thereby forming the droplets into solid microspheres. This heating procedure evaporated water from the microspheres and denatured the albumin. This procedure was not usable for the encapsulation of heat-sensitive therapeutic agents. At temperatures of 100° C. or above, partial or complete inactivation of many important therapeutic agents will occur.

In the development of the present invention, it was discovered that albumin microspheres can be prepared without heating the aqueous albumin droplets. Even though water is not evaporated from the droplets and there is no denaturation of the albumin, water-soluble drugs can be entrapped in the microspheres, and the

2

microspheres can be recovered for medical use or further processing. This discovery makes possible the preparation of microspheres containing water-soluble heat-sensitive therapeutic and/or diagnostic agents. As subsequently used herein the term "therapeutic agent" is intended to be inclusive of clinical agents which can be administered in microcapsular form, whether used primarily for treatment or diagnosis.

According to the method of the present invention, the aqueous albumin is dispersed and emulsified in a non-polar solvent (viz. a vegetable oil). The resulting water-in-oil emulsion containing droplets of the desired size can then be introduced into a body of the same solvent, such as vegetable oil, to further disperse the droplets. Neither the pre-formed emulsion or the body of vegetable oil are heated, ordinary room temperature or lower being satisfactory. The oil is removed by washing the dispersed droplets with an oil-soluble water-immiscible organic solvent, and the resulting microspheres are recovered in a portion of the organic solvent. Thereafter, the recovered microspheres can be dried, either with or without further treatment to harden the microspheres. The mechanism of formation of the microspheres is not fully understood. On the basis of prior knowledge, it is surprising that the microspheres as recovered do not coalesce but retain their structure as individual non-aggregated spheres.

DETAILED DESCRIPTION

The matrix material for forming the microspheres is albumin. Human serum albumin (HSA) is preferred. The albumin matrix material may be modified by using it in combination with a minor proportion of other biodegradable amino acid polymers. For example, from 0 to 25 parts by weight (dry basis) of hemoglobin (preferably human hemoglobin) or a synthetic amino acid polymer can be combined with 75–100 parts of albumin. Useable synthetic amino acid polymers include poly-L-lysine and poly-L-glutamic acid, having a molecular weight ranging from 20,000 to 50,000. Preferably, however, human serum albumin is employed as the sole matrix material.

In preparing the microspheres, an aqueous solution of the albumin matrix material is prepared, which can be formed into microspheres. The amount of matrix material to be used will usually be within the range from 5 to 50 parts by weight of the matrix material per 100 parts of water. Preferred proportions are from 20 to 30 parts per 100 parts of water. The water-soluble heat-sensitive therapeutic agent may be dissolved in the water of the matrix material solution, either before or after preparing the matrix solution.

The aqueous solution of the albumin matrix material containing the therapeutic agent is emulsified with an oil, which is preferably a vegetable oil, such as cottonseed oil, peanut oil, or the like. Other oil or suitable non-polar solvents for forming water-in-oil emulsions can be used. The aqueous phase at the time of addition of the oil may also contain the magnetic particles when previously added to the aqueous solution of the matrix material and dispersed therein. The proportions of the aqueous phase to the oil phase can conveniently range from about 1 to 5 parts by weight of the aqueous phase per 100 parts of the oil phase. This provides separation of the oil droplets, and prevents coalescence of the droplets in forming the microspheres. The water-in-oil emulsion is then treated to reduce the size of the dis-

4,357,259

7

was used instead of the 36 mg of Example I. The microcapsule product contained approximately by dry weight 50% Fe_3O_4, 4% adriamycin, and 46% albumin.

EXAMPLE IV

Microcapsules were prepared by the identical procedure of Example I, using approximately the same amount of Fe_3O_4 as in Example III. The Fe_3O_4 was in the form of an aqueous suspension containing a surfactant, aqueous base Ferrofluidics Fe_3O_4 Catalog No. A-01, 400 gauss saturation (Ferrofluidics Corporation, Burlington, Mass.). 0.3 ml of the A-01 product was added, containing approximately 130–140 mg Fe_3O_4. The average Fe_3O_4 particle size was in the range of 150–200 Angstroms. The microcapsular product contained approximately by dry weight 51% Fe_3O_4, 4% adriamycin, and 45% albumin.

EXAMPLE V

A. Same procedure as Example I except microspheres are not hardened by a cross-linking agent. The oil bath is at a temperature of 20°–25° C. After the oil has been washed away with diethyl ether anhydrous 4 times, the spheres are separated from the remaining solvent, air dried, then lyophilized, and stored at 4° C.

Stability of the resultant microspheres was tested as follows: First, 5 microliters of ^{125}I-bovine serum albumin (New England Nuclear, 1.51 mCi/mg) was added to the initial homogenate to trace label the microspheres. An aliquot of the resultant microspheres was then suspended and sonicated for 2 minutes in 0.154 M NaCl-0.1% Tween 80 and incubated at 37° C. for 24 and 48 hours. After this period of time, the suspension was centrifuged at 2000×g for 10 minutes and the supernatant and pellet were counted in a gamma counter. The number of counts obtained in the supernatant (after subtracting free label) divided by the total number of counts was regarded as the percentage breakdown of the carrier (non-pelleting). Only 16% of the microspheres had deteriorated after 24 hrs and 37% after 48 hrs. With formaldehyde or heat-hardening less than or equal to 3% deterioration occurs in 48 hrs.

B. Same procedure as Example I except 3,400 units of urokinase was added to the 125 mg of HSA omitting the Fe_3O_4 and adriamycin. No cross-linking was done in this experiment.

Two mg of the resultant microspheres were placed into 16 12×75 mm tubes of duplicate time course of 0, 15, 30, 60 minutes; 2, 4, and 6 hours. At the appropriate time, microspheres were suspended in sodium barbital buffer (0.05 M) and left at room temperature. Finally, at

8

zero time all tubes were centrifuged at 3,500 RPM (1900×g) for 15 minutes at 4° C. and 25 microliters of the supernatants were pipetted into appropriate wells in fibrin-agar plates. Plates were then read 4 and 6 hours later for fibrinolysis (i.e. diameters).

It was found that 60% of maximum lysis was seen after 10 minutes on fibrin-agar plate.

OTHER EXAMPLES

As a variation of the procedure of the foregoing examples, butanedione (5% v/v in anhydrous ether) or butyraldehyde (10% v/v in anhydrous ether) is employed as a cross-linking agent, the contact time ranging from 5 minutes to 2 hours. The product can be recovered and dried as described in Example I.

As a further variation of the procedure for the foregoing examples, 20 mg of poly-L-lysine or polyglutamic acid is combined with the 125 mg of human serum albumin. The rest of the procedure is identical to Example I.

We claim:

1. The method of incorporating a water-soluble heat-sensitive therapeutic agent in albumin microspheres, in which all steps thereof are carried out at a temperature within the range from 1° to 45° C., said method including the steps of preparing an aqueous albumin solution of the said therapeutic agent, said albumin solution containing from 5 to 50 parts by weight of albumin per 100 parts of water and from 1 to 20 parts by weight of said therapeutic agent per 100 parts of albumin, emulsifying said albumin solution with a vegetable oil to form a water-in-oil emulsion containing dispersed droplets of the albumin solution, removing the oil by washing the dispersed droplets with an oil-soluble water-immiscible organic solvent, and recovering the resulting microspheres, wherein said method also includes the step of contacting said microspheres with an organic solvent solution of an aldehyde hardening agent to increase the stability of said microspheres and to decrease the release rate of said drug therefrom.

2. The method of claim **1** in which said hardening agent is formaldehyde, and in which said organic solvent solution is formed by transferring formaldehyde from an aqueous solution thereof into said organic solvent, the residual aqueous solution being separated from said organic solvent solution prior to said contacting.

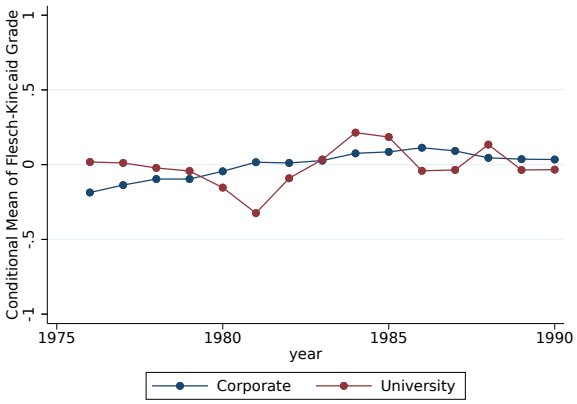
3. The method of claim **1** or claim **2** in which said albumin solution also contains magnetic particles.

4. The method of claim **13** or claim **2** in which said organic solvent solution contains from 0.2 to 20% by weight of said aldehyde hardening agent.

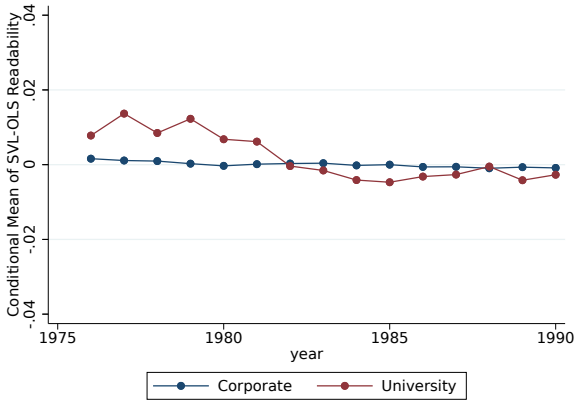
* * * * *

Notes: This is an example of a patent document text, with the red portion representing the patent summary, the green section containing the detailed description, and the yellow area containing the patent claims.

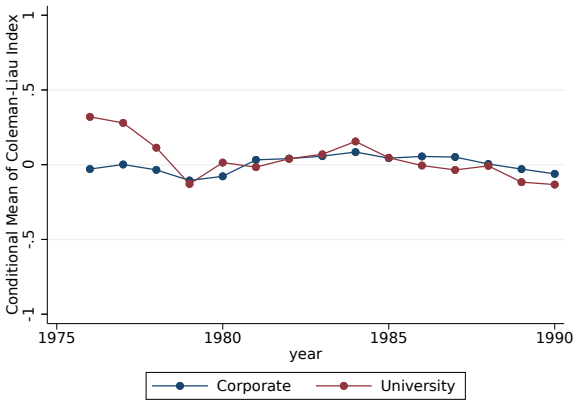
Figure A3: Conditional Mean of Multiple Readability Scores Over Time



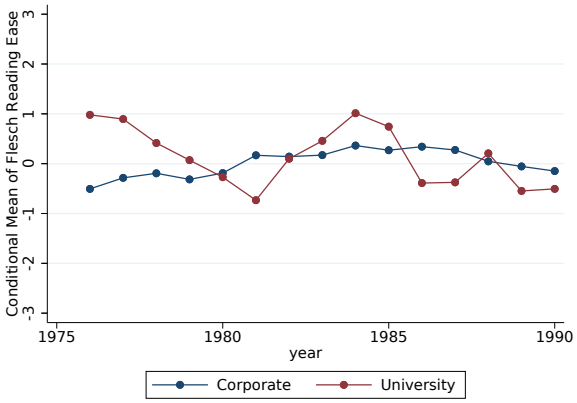
(a) Flesch-Kincaid Grade



(d) SVL-OLS Readability



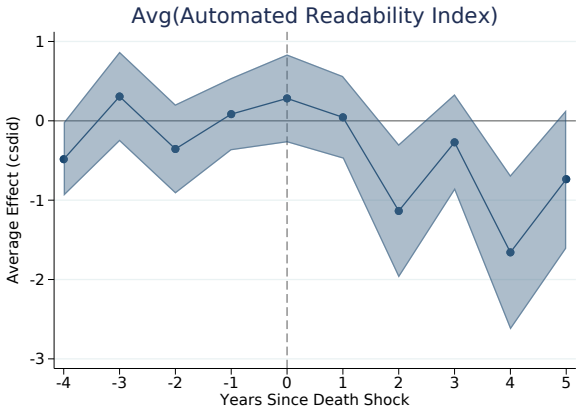
(b) Coleman-Liau Index



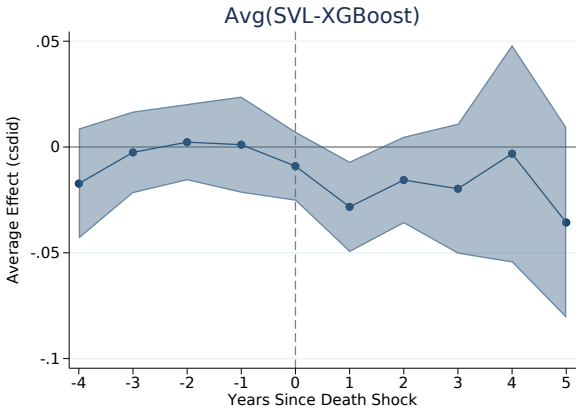
(c) Flesch Reading Ease

Notes: The figure displays the (conditional) mean of (different) readability scores of university patents and corporate patents over time, conditionally on technology classifications (CPC codes) and assignees. Except for the Flesch Reading Ease, all scores are multiplied by -1 so that lower values correspond to harder-to-read text. The Flesch Reading Ease score is already scaled such that lower values indicate more difficult text and is not transformed.

Figure A4: DID Setting II: Readability Change of Patent Detailed Descriptions by TTO establishment Over Year (CSDID)



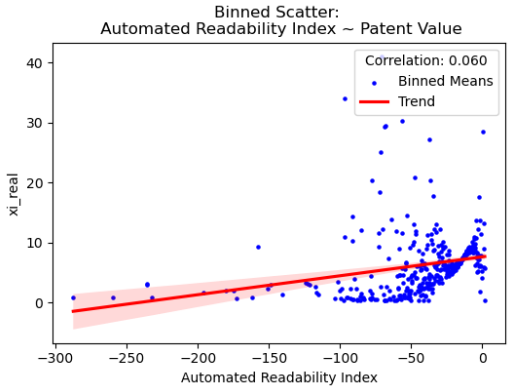
(a) Traditional Readability Score



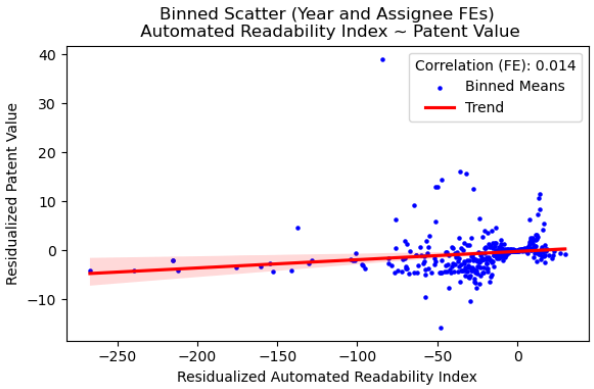
(b) AI-based Readability Score

Notes: The figure displays changes in aggregate (mean) readability around TTO establishment relative to the control group using CSDID. CSDID requires panel data with a single observation per unit per time period. To meet this requirement, I aggregate the patent data to the inventor-year level by computing the mean readability score for each inventor annually. This aggregation may reduce precision, as it limits the ability to control for patent-level characteristics such as technological fields.

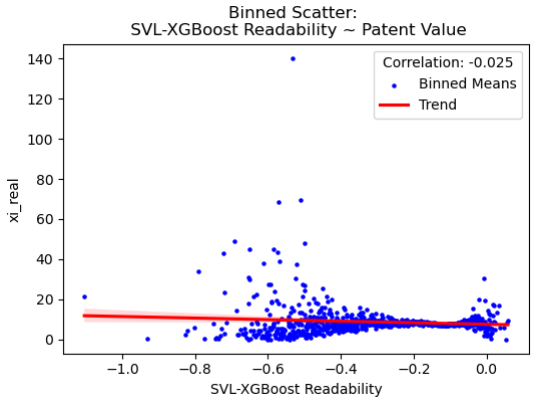
Figure A5: Correlation between Readability Scores and Patent Market Value



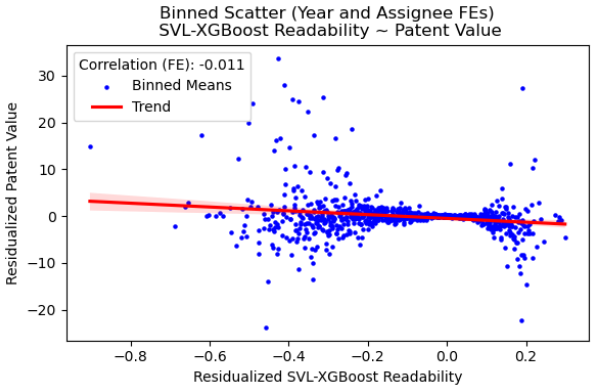
(a) Traditional Readability Score



(b) Traditional Readability Score (FEs)



(c) AI-based Readability Score



(d) AI-based Readability Score (FEs)

Notes: This figure shows the correlation between patent market value ("xi_real") and readability scores. Panels (a) and (c) present raw correlations. Panels (b) and (d) control for assignee and application year fixed effects. The two readability measures used are the Automated Readability Index and the SciBERT-XGBoost readability score. Both are originally scaled such that higher values indicate lower readability (i.e., more difficult to understand) or a higher likelihood of rejection. For interpretability, I multiply both scores by -1, so that lower values correspond to harder-to-read text. "xi_real" is defined in Kogan et al. (2017) as the value of innovation deflated to 1982 (million) dollars using the CPI.

C Technical Details

C.1 Use SciBERT to Process the Text

SciBERT is a transformer-based language model built on BERT (Beltagy et al. 2019) and pre-trained specifically on scientific literature. In the BERT framework (Devlin et al. 2018), embeddings are vector representations of text that capture contextual nuances, effectively condensing complex language into a manageable vector space. BERT, or Bidirectional Encoder Representations from Transformers, captures word context bi-directionally using a deep neural network with a transformer architecture (Vaswani et al. 2017). This architecture processes all text tokens through 12 layers of fully connected attention mechanisms, enabling the model to understand complex and long-range dependencies in text. BERT is trained on two tasks: predicting randomly masked tokens and determining whether one sentence logically follows another.

Unlike BERT, which is pre-trained on general domain corpora like news articles and Wikipedia, SciBERT is optimized for scientific texts, improving performance on scientific natural language processing (NLP) tasks. The SciBERT model generates 768-dimensional embeddings in its final layer (see Figure 2 for an overview). However, patent detailed descriptions often exceed the maximum token length that SciBERT can handle, which is 512 tokens. To address this, I extract the first 128 tokens, the last 128 tokens, and the middle 256 tokens from each patent description, creating shorter descriptions that preserve essential information while fitting within the model's constraints.

The [CLS] token, added at the beginning of the input sequence, is processed through a dense layer followed by a tanh activation function. Its embedding is designed to capture the overall meaning of the entire sequence (Devlin et al. 2018). For my analysis, I use this 768-dimensional embedding vector as input in model training.

C.2 Evaluate Model Performance

To evaluate the performance of models, I first utilize key metrics: recall,[38] and the Receiver Operating Characteristic Area Under the Curve (ROC AUC).[39]

The OLS model achieves a recall of 0.83. Similarly, the XGBoosting model shows a recall of 0.84. The high recall values indicate that these models are effective at identifying the majority of applications that should be flagged for potential rejection under 112(a), capturing approximately 83-84% of all true positive cases. The ROC AUC values, which measure the models' ability to distinguish between positive and negative instances, are 0.60 and 0.58 for the OLS and XGBoosting models, respectively. This suggests that the models have a roughly 60%

[38] Recall, also known as sensitivity or true positive rate, measures the model's ability to identify all relevant positive cases. It is defined as $\text{Recall} = \frac{\text{True Positives (TP)}}{\text{True Positives (TP)} + \text{False Negatives (FN)}}$.

[39] ROC AUC refers to the area under the Receiver Operating Characteristic curve, which quantifies the model's ability to distinguish between positive and negative cases. Values range from 0.5 (no discrimination) to 1 (perfect discrimination).

and 58% probability of correctly ranking a randomly chosen positive instance higher than a randomly chosen negative instance across various threshold levels.While modest, this reflects the challenge of achieving high precision in a highly imbalanced setting. Combined with the high recall scores, the model still provides a useful proxy for identifying unreadable patents under §112(a). These values reflect moderate discriminative ability but highlight the challenges posed by the highly imbalanced nature of the dataset (only around 13.8% observations are positive).[40]

To further assess model performance in this imbalanced setting, I report F_β-scores, which allow for differential weighting of recall and precision to better reflect the specific goals of the analysis.[41]

While the standard F_1-score ($\beta = 1$) assumes equal costs for both types of classification errors, my analysis favors sensitivity to identify potential "obfuscated" patents. As β increases, the metric becomes more lenient toward false positives in exchange for higher coverage of true positive cases. In this context, the F_2, F_3, and F_4 scores are 0.43, 0.57, and 0.66, respectively. These values underscore the model's effectiveness in retrieving the majority of relevant rejections despite the extreme sparsity of the positive class. Although the corresponding low precision reflects the inherent difficulty of predictive accuracy in imbalanced data, the high recall ensures that few "low-readability" cases are missed, which is critical for maintaining the integrity of the subsequent empirical analysis.

It is important to note that this model is not used for high-stakes classification tasks (such as automated rejection decisions), but rather to generate a continuous score that reflects the likelihood a patent would be deemed "unreadable" under §112(a). While the AUC is modest and may not support strong predictive classification, it is sufficient for constructing a proxy measure with weak-to-moderate signal for readability under the context.

Turning to the traditional readability scores, I present the performance metrics of the same models using these scores as predictors. The OLS model achieved a high recall of 0.96 but showed relatively low discriminative power, with a ROC AUC of 0.54. The XGBoosting model exhibited a lower recall of 0.78 but reached the same ROC AUC. These results suggest that while traditional readability metrics may have limitations in ranking performance, they nonetheless capture meaningful variation in patent text clarity and sufficiency—as evidenced by their alignment with examiner decisions. The traditional readability scores show comparable good recall to the machine learning models but slightly lower ROC AUC and PR AUC values. This indicates that traditional readability measures are somewhat less effective at ranking the readability of patent descriptions compared to the machine learning models, yet they still provide an assessment of patent readability.

[40] 40835 observations with 112(a) rejections, and 254968 observations with non-112 (a) rejections.

[41] The F_β-score is defined as:

$$F_\beta = (1+\beta^2) \cdot \frac{\text{Precision} \cdot \text{Recall}}{(\beta^2 \cdot \text{Precision}) + \text{Recall}}$$

C.3 Detailed Methodology for Classifying Patent Orientation

C.3.1 The Share-Based Method

The share-based method classifies patents by analyzing the proportion of predefined keywords associated with basic and applied research. Basic research is represented by keywords such as “theory,” “concept,” “theoretical,” “basic,” and “research-based,” while applied research is characterized by terms like “applied,” “application,” “process,” “industrial,” “practical,” and “commercial.”

To ensure consistency, patent descriptions are preprocessed by removing special characters, converting to lowercase, and eliminating stop words. After tokenization, the occurrences of basic and applied keywords are counted. The share of basicness is calculated as the proportion of basic keywords relative to the total number of basic and applied keywords in the patent, while appliedness is calculated analogously. For instance, a patent with 8 basic keywords and 2 applied keywords would have a basic share of 0.8 and an applied share of 0.2. If no keywords from either set appear, neutral scores of 0.5 are assigned to both metrics.

This method is computationally efficient and offers an intuitive way to classify patents based on their content, serving as a straightforward benchmark for thematic alignment.

C.3.2 The LDA-Based Method

The LDA-based method uses Latent Dirichlet Allocation (LDA) to identify latent thematic structures in patent descriptions. LDA assumes that each patent is a mixture of topics, and each topic is a distribution over words. This probabilistic approach enables a nuanced classification of patents.

First, patent descriptions are preprocessed and converted into a document-term matrix, where rows represent patents, and columns represent words. The LDA model then identifies a predefined number of topics (e.g., five). Each topic is characterized by a distribution over words, and the most representative keywords for each topic are extracted. Topics are classified as basic or applied based on the proportion of basic and applied keywords within them. For example, a topic dominated by terms such as "basic" and "theory" would have a high basicness score, while a topic emphasizing "application" or "industrial" would score higher in appliedness. The basicness and appliedness scores for each patent are calculated as weighted averages of the topic-specific scores, with weights determined by the patent’s distribution over topics.

This method captures latent patterns in patent descriptions, allowing for a probabilistic and flexible classification framework. By combining the LDA-based method with the share-based approach, I leverage the strengths of both methodologies. The share-based method provides a direct and interpretable measure, while the LDA-based method uncovers deeper thematic structures, creating a robust framework for classifying patents as basic or applied.

D Evidences

D.1 A Real-World Example

To illustrate both what the readability measures capture in patent documents and the mechanism proposed in the paper, I examine a concrete example from the dataset in which the readability of a university inventor's patent disclosures declines following the Bayh-Dole Act and the establishment of the university's Technology Transfer Office (TTO). For a fair comparison, I focus on patents filed by the same inventor at the same university, within the same technological classification (CPC class G09: *Education; Cryptography; Display; Advertising; Seals*), but across the pre- and post-shock periods. While the inventions themselves are different, the institutional setting, assignee, and technological field are held constant.

Specifically, I compare two patents filed by the same inventor at a U.S. university. The first was filed in 1976, prior to both the enactment of the Bayh-Dole Act and the establishment of the university's TTO. The second was filed in 1983, after both institutional changes had taken place. This comparison serves two purposes. First, it provides a concrete illustration of the textual features captured by the readability measures. Second, it offers a qualitative example consistent with the paper's mechanism: following an institutional shift that increased the value of commercialization, patent disclosures may become less transparent and more difficult for outside readers to parse.

Both traditional readability measures and the AI-based readability measure indicate that the later patent is substantially more difficult to read than the earlier one. The patent filed in 1976 is relatively straightforward and technical, using standard engineering terminology and schematic references. In contrast, the patent filed in 1983 is densely worded, repetitive, and filled with vague references to system behavior, making it considerably more difficult for an outsider to interpret. The latter text introduces layers of complexity that obscure rather than clarify the underlying invention.

Since the detailed descriptions are lengthy, I extract only the first paragraph from each as an illustrative example. I put these excerpts in ChatGPT, which highlights whether the language is clear or obscure. In the annotated examples, segments identified as vague or difficult to interpret are enclosed in parentheses, with brief explanations from ChatGPT clarifying why those parts are considered unclear (e.g., overuse of abstract terms, unclear referents, excessive length, or embedded technical complexity).

Pre-Shock Patent (filed in 1976):

> *Referring now to FIG. 1, there is schematically illustrated a gas discharge panel or plasma panel 10 of the type described in the aforementioned U.S. Pat. No. 3,559,190, having a plurality of crossing electrodes in respective arrays on each side of the panel separated by insulating material containing a gaseous medium.*

For purposes of describing the present invention, the plasma panel 10 shown in FIG. 1 is illustrated as including four electrodes in each array, and it is to be understood that the normal plasma panel may contain, for instance, 512 electrodes in each array – although panels with as many as 1024 electrodes in each array have been constructed (too many numerical details may distract from core idea). The electrodes in the Y matrix array are each connected to respective conductive lines 12, 14, 16 and 18; and the electrodes in the X matrix array are respectively connected to lines 20, 22, 24 and 26 (no explanation of why this numbering is important). In the illustrated plasma panel, there are therefore 16 gas cells or locations defined on the panel by an intersection of respective electrodes in the X and Y matrix array. For instance, gas cell or location 0, 0 on the panel is defined between lines 12 and 20; and gas cell 1, 1 is defined between lines 22 and 14 (overly technical and assumes domain knowledge).

Post-Shock Patent (filed in 1983):

FIG. 1 is a block diagram illustrating a cluster subsystem 10 for storing a plurality of instructional programs for execution by associated apparatus (vague: what apparatus? unclear reference), and for paging program segments to that other apparatus as required (ambiguous phrasing; difficult to know what "that other apparatus" refers to), a high speed communication subsystem interconnecting said cluster subsystem and a plurality of processor stations 20. Typically, the processor stations each include a microprocessor and associated electronic memory. Typically, the electronic memory at a processor station does not have the capacity to store even a single instructional program (longwinded and could be simplified; lacks clarity on why this is important). Rather the instructional programs are segmented into segments (redundant wording) (for example 2–3 kilobytes) and transmitted to a requesting processor station. The dimensions or parameters characterizing the various elements (overly general and abstract; lacks specificity) are arranged so that for up to some limiting number of processor stations (vague: "some limiting number" is undefined), the delay occasioned by the necessity for extracting a program segment from the cluster subsystem, shipping it to a processor station, storing it in the processor station memory, and executing the stored program segment, presents no more than a barely perceptible delay to the user (overly complex sentence; unclear metric). For those user keyboard inputs which can be accommodated by the program segment located at the processor station (wordy and indirect), no communication is necessary between the processor station and the cluster subsystem. There are in fact only two different occasions when the processor station must communicate with the cluster subsystem, and those are when the processor station requires a different program segment than the one located in

the electronic memory of the processor station (highly repetitive and unclear), and those occasions when information is transmitted to the cluster subsystem in order to build a user proficiency indication (vague phrase; "user proficiency indication" is undefined).

This example illustrates how university inventors may have responded to enhanced commercialization opportunities and institutional changes by reducing the readability of their patent's detailed descriptions.

D.2 Details about Alternative Explanations and Confounding Factors

D.2.1 Changes in Patent Characteristics

– Patent Basic vs. Applied Orientation. One potential concern is that the readability changes may be driven by shifts in the basic-applied orientation of the inventions themselves—whether patents are more basic or more applied in their orientation—rather than changes in disclosure strategy. Basic patents are defined as those typically grounded in fundamental research, emphasizing theoretical advancements, new principles, or foundational innovations. They may not have immediate commercial applications but can serve as a basis for further research and development (Trajtenberg et al. 1997, Henderson et al. 1998a, Angori et al. 2024). For instance, patents with a more basic orientation might naturally exhibit different readability patterns compared to more applied patents. If the shock leads to a systematic shift in the basic-applied orientation of patents filed, such a change could confound the interpretation of the results.

To address this, I implement two complementary measures of patent orientation: a share-based keyword method and an LDA-based topic model (Appendix C.3). As shown in Tables A20 and A21, the results remain consistent after controlling for patent orientation.

– Patent Originality and Generality. Another potential concern is that the observed changes in patent readability may be influenced by shifts in the originality and generality of patents. These measures reflect key aspects of a patent's role in the innovation ecosystem (Trajtenberg et al. 1997, Hall et al. 2001) and may be correlated with both disclosure strategies and readability patterns. Both measures are derived from the NBER patent database (Hall et al., 2001), which provides systematic methodologies for calculating originality and generality based on citation patterns.[42]

[42]Patent originality reflects the diversity of prior art citations, indicating whether an invention draws upon a wide range of technological fields. Highly original patents, which integrate diverse knowledge sources, may exhibit different readability characteristics compared to less original patents. For instance, patents with greater originality may involve more complex or interdisciplinary language, which could influence their measured readability.

Patent generality, on the other hand, captures the breadth of influence a patent has across subsequent technological fields. Patents with higher generality are cited by inventions spanning a wide array of fields, suggesting their broad applicability. The language used in such patents might be designed to appeal to a broader audience, potentially affecting readability.

As shown in Tables A22 and A23, the results remain robust even after accounting for these variables. Including these controls ensures that the observed changes in readability are not fully influenced by shifts in the originality and generality of patents.

– Patent Quality. An additional concern is that the observed changes in patent readability may reflect shifts in the quality of patents rather than changes in strategic disclosure. Patent quality, often associated with the economic and technological value of an invention, is an important factor that could influence both readability and disclosure decisions. High-quality patents may be written differently to maximize their impact, which could confound the analysis.

Measuring patent quality is inherently challenging, and previous literature often relied on citation counts as a proxy. However, this approach is imperfect, particularly in my context, as citations are likely to correlate with readability. For example, patents with clearer language may receive more citations simply because they are easier to interpret, introducing potential bias into the analysis.

Patent Quality: Journals from Patent-Paper Pairs as Proxies. First, I leverage the Patent-Paper Pairs (PPP) dataset from Marx and Scharfmann (2024), which matches patents from PatentsView to academic papers in OpenAlex authored by the same individual. Marx and Scharfmann (2024) constructed these PPPs using multiple criteria, including machine learning models that analyze textual similarities between titles and abstracts, as well as patent-paper citations identified by Marx and Fuegi (2022). These matched pairs ensure that the patent and the paper describe the same invention, providing a reliable link between the two. Importantly, many – but not all – PPP pairs include direct citations from the patent to the paper, further strengthening their connection.

I use the journal associated with the paper as a proxy for the quality of the invention, incorporating it as fixed effects in my main regressions. This approach accounts for potential variation in quality by controlling for differences in the perceived prestige or rigor of the journals in which the related research was published.[43]

The results from specifications adding the journal fixed effects are shown in Tables A24 and A25 and remain consistent with main results, indicating that the observed changes in patent readability are unlikely to be fully driven by shifts in invention quality. This robustness check strengthens the interpretation that the observed readability changes are more likely attributable to strategic disclosure adjustments following institutional changes, rather than just being a byproduct of variations in the quality of inventions.

[43]For patents without an associated paper, I assign a category labeled "NA" to ensure that these observations remain in the sample, thereby avoiding a reduction in sample size and preserving statistical power. For patents with more than one associated paper linked to different journals, I randomly select one journal as the representative proxy.

Patent Quality: Market Values as Proxies. I second use data from the KPSS dataset (Kogan et al., 2017), which estimates patent market value—a proxy for patent quality in monetary terms—based on patents held by public firms. Unlike citation-based measures, the KPSS dataset offers an independent assessment of patent quality by linking patents to their economic value, reducing potential endogeneity issues with readability. Note that the KPSS dataset covers patents for public firms and does not include university patents. Consequently, this data cannot be directly applied to my sample. Instead, I present correlation results as suggestive evidence.

As illustrated in Figure A5, there is a small negative correlation between patent market value and the SVL-XGBoost readability score (raw: -0.025; conditional on assignee and year fixed effects: -0.011), and a slight positive correlation with the automated readability score (raw: 0.060; conditional on assignee and year fixed effects: 0.014). The fact that these correlations are both modest and in opposite directions, yet both readability scores decrease in my main analysis, suggests that the observed decline in readability is unlikely to be driven by changes in patent quality. This finding supports the interpretation that the decrease in readability reflects a genuine shift in disclosure practices rather than changes in patent quality following the shock.

D.2.2 TTOs: Commercialization Incentives and Professionalization

TTOs likely influenced disclosure behavior through two mechanisms: (1) reinforcing commercialization incentives, and (2) introducing professional legal and technical expertise. TTOs play a central role in university commercialization by managing IP rights, overseeing licensing, and guiding inventors through the patenting process (Hellmann 2007, Hvide and Jones 2018, Faccin et al. 2022). To isolate their effect, I restrict the sample to post-1980, conditional on the Bayh-Dole Act being in place. As shown in Table 7, inventors further reduced patent readability after TTOs were established. This suggests that the presence of TTOs—either by reinforcing inventors' commercialization incentives or through the guidance of professionals—may have contributed to more strategic disclosure.[44]

Qualitative evidence (see Appendix D.3.1) suggests that patent attorneys often leverage their legal and technical expertise to craft claims that strategically protect intellectual property. Accordingly, the rise of Technology Transfer Offices (TTOs) likely influenced the patent drafting process by introducing attorneys into the preparation of applications, thereby shaping the clarity and sufficiency of patent language. However, prior research (Bryan et al. 2020, Andriosopoulos et al. 2023) and my discussions with legal practitioners suggest a division of labor: attorneys are primarily responsible for phrasing claims in precise legal language, while inventors remain principally responsible for providing the technical information disclosed in patents.

[44]Restricting the analysis to the post-1980 period creates a uniform IP environment under Bayh-Dole but sharply limits pre-treatment observations. The resulting loss of statistical power makes the key coefficient less precise and, in this subsample, its significance falls relative to the main specification.

To further disentangle the role of legal expertise from inventors' own incentives, I link patents to attorney identifiers from USPTO PatentsView and include attorney fixed effects in the analysis.[45] As shown in Tables A26 and A27, the main effects remain robust. Although the inclusion of attorney fixed effects modestly attenuates the coefficients, the results indicate that professionalization might contribute to—but did not solely drive—the observed decline in readability. The results suggest that attorneys serve as instruments in the technical disclosure process, executing strategies that align with inventors' incentives to manage how their inventions are revealed. Rather than independently shaping disclosure practices, attorneys act as executors of inventor-driven goals.

Overall, the findings support the view that inventors strategically adapted their disclosure practices in response to evolving commercialization incentives, amplified by institutional support through TTOs.

D.3 Patent Drafting and Motivation

Patent drafting in the United States may be undertaken by inventors themselves or, more commonly in institutional and commercial contexts, through collaboration between inventors and patent professionals (attorneys or agents), with each section of a patent document typically authored or heavily influenced by one or the other. The background and summary, detailed description, and claims serve distinct purposes and are written with different strategic and legal considerations in mind. Although inventors and attorneys may work together on the overall application, their contributions differ systematically across sections, reflecting their respective expertise and incentives.

The **Background and Summary** frame the invention in relation to prior art and emphasize its novelty. In U.S. practice, anything stated in the background can be treated as "applicant admitted" prior art—meaning the applicant is on record admitting certain information was known before the invention. This creates an incentive to keep the background generic and noncommittal to avoid limiting the scope of protection.[46]

The **Detailed Description of the Invention** is the heart of the technical disclosure and the section where inventor input is greatest. Inventors provide the substantive content—technical explanations, schematics, and examples—drawing on their "unmatched technical expertise" and unique understanding of the invention's novel features.[47] In practice, inventors often supply an initial write-up or invention disclosure describing how the invention works, sometimes including drawings or specific embodiments. Patent attorneys then refine this material to meet

[45] For patents without an attorney, I assign a category labeled "NA" to ensure that these observations remain in the sample, thereby avoiding a reduction in sample size and preserving statistical power. For patents listing multiple attorneys (around 30% in the sample), I use the first-listed attorney as the primary identifier. In practice, when multiple attorneys are of record, it is likely that official correspondence from the Office is directed to the first-listed attorney (https://blueironip.com/ufaqs/how-does-the-uspto-handle-correspondence-from-multiple-attorneys-in-a-patent-application/.

[46] Source:https://outlierpatentattorneys.com/patent-drafting-strategy-background-and-summary-sections

[47] Source:https://patentpc.com/blog/the-role-of-inventor-input-in-patent-drafting

legal standards such as enablement and clarity and sufficiency under 35 U.S.C. § 112(a).[48] While the attorney shapes the final language, the content largely reflects the inventor's perspective, making this section more sensitive to inventor-driven changes in disclosure practices.

The **Claims** define the legal boundaries of the invention and are almost always drafted by the attorney, with inventor input limited to clarifying scope and identifying essential features. Claims require specialized legal language and must satisfy doctrines such as definiteness, novelty, and non-obviousness. One guide notes that "one of the primary objectives of patent drafting is to establish claims that are both broad enough to prevent close variations of the invention by competitors and defensible in scope."[49] Attorneys frequently begin the application by drafting the claims first, then build the rest of the document to support them.

Understanding this division of labor is important for interpreting variation in patent readability. For instance, changes in the clarity and sufficiency of the detailed description are more likely to reflect inventor-side responses to institutional or policy shocks. While disentangling these effects is beyond the primary scope of this paper and still an open question, I address potential bias from attorney-specific drafting styles by including attorney fixed effects in robustness checks.

D.3.1 Examples of Patent Attorneys' Professional Services

To illustrate the role of patent attorneys in shaping disclosure practices, I reviewed publicly available legal commentary and guidance from both practicing attorneys and legal scholars. The following examples include perspectives from a full-service patent law firm and from a law professor. These sources highlight the tension between meeting disclosure requirements and protecting client interests, often encouraging strategic language and selective detail.

- **Example 1:**

  "A good patent lawyer will draft the patent application in such a way as to include a sufficient amount of detail to articulate the invention and thereby have it admitted while not limiting it too much so that it cannot be used as an enforcement device against infringers...

  Use broad and strategic language: Drafting approach: Draft the patent application using broad and strategic language that captures the essence of the invention without revealing every specific implementation detail."[50]

  – Cohn Legal, PLLC (a full-service patent law firm)

- **Example 2:**

[48] Source: https://www.upcounsel.com/patent-drafting
[49] Source: https://www.upcounsel.com/patent-drafting
[50] Source: https://www.cohnlg.com/practices/patents/

"Some amount of obfuscation will be familiar to patent attorneys today across various fields. Although not illegal, most patents avoid detailing the core invention... Likewise, some patent attorneys will have walked the tight-line between the patentability disclosure requirement and client-interest in keeping certain trade secrets... Perhaps all of this goes along with traditional core roles of attorneys — advocating and keeping secrets."[51]

– "Obfuscation and Patenting of @#$%" by Dennis Crouch, Professor, University of Missouri School of Law

[51] Source: https://patentlyo.com/patent/2021/05/obfuscation-and-patenting.html