Patent or Trade Secret? Generative AI, Innovation Protection, and Strategic Disclosure*

Terrence Tianshuo Shi
terrence.shi@anderson.ucla.edu
Anderson School of Management
University of California, Los Angeles

Charles C.Y. Wang
charles.cy.wang@hbs.edu
Harvard Business School
Harvard University

Danye Wang
danye-wang@uiowa.edu
Tippie College of Business
University of Iowa

Yuxiang Zheng
yuxiang.zheng@rutgers.edu
Rutgers School of Business Camden
Rutgers University

July 2026

Abstract

We examine how the emergence of generative artificial intelligence (GenAI) reshapes firms' choice between the two primary forms of protecting innovation: patents and trade secrets. The effect is ex ante ambiguous. GenAI may make secrecy harder to maintain by increasing the risk of inadvertent innovation leakage and lowering competitors' cost of reverse engineering. Yet it may also make patenting more costly, since it can help rivals more quickly imitate or invent around technologies after patent filings disclose the underlying technical details. Using the public launch of ChatGPT in November 2022 as a shock to the availability and salience of GenAI, we find that trade-secret-reliant firms shift toward patenting, consistent with rising secrecy costs increasing the relative appeal of legal exclusivity despite the disclosure required by patents. The shift is stronger when firms face greater information leakage risk, as captured by GenAI exposure and dispersed inventor networks, and when their industries and technologies are more susceptible to reverse engineering. Despite the overall increase in patenting, firms appear to limit the information they reveal. They delay publication, provide shorter and less specific disclosures, concentrate filings in areas with faster technological obsolescence, and seek broader claims. Taken together, our findings suggest that GenAI shifts the boundary between private knowledge and public disclosure, increasing legal exclusion without necessarily producing a proportional increase in useful technical disclosure.

Keywords: Generative AI, patents, trade secrets, innovation, disclosure

JEL Codes: K11, M41, O31, O34

*We appreciate helpful comments from Mary Billings, Henry Friedman, Cristi A. Gleason, Camilla Hrdy, Yiwei Li, Vincent Yongzhao Lin, Kristen Valentine, Jenny Li Zhang, and Yuan Zou. We thank attorneys Will Cheng (LinkedIn), Tim Dwyer (Holland & Hart), Ryan McCormick (M&B IP), and Bradford Newman (Eversheds Sutherland) for their invaluable insights. All errors are our own.

1 Introduction

We ask whether and how the emergence of generative artificial intelligence (hereafter, GenAI) reshapes firms' choice between the two principal mechanisms for protecting innovation: patents and trade secrets. These two mechanisms rest on fundamentally different arrangements over innovation disclosure and enforcement rights (Glaeser and Lang, 2024). A patent grants legal exclusivity, but only in exchange for public, enabling disclosure of the invention (Arrow, 1962; Hall and Harhoff, 2012). A trade secret, by contrast, requires no disclosure, but the holder must take "reasonable measures" to keep the information confidential; it offers no protection against lawful reverse engineering or independent discovery, and maintaining secrecy requires firms to guard against employee leakage and misappropriation (Friedman et al., 1991; Samuelson and Scotchmer, 2002; Lemley, 2008). By reshaping both how firms innovate and how proprietary information flows, GenAI may fundamentally alter the trade-off between patents and trade secrets. The direction of this effect, however, is ex ante ambiguous.

On the one hand, GenAI can weaken trade-secret protection by increasing information leakage and facilitating reverse engineering (Levine, 2023; Hrdy, 2025; Bombard and Tibbetts, 2025). Effectively deploying GenAI in operations and innovation often requires firms to provide AI systems with proprietary data and domain knowledge, creating what Satya Nadella terms the "reverse information paradox"[1]: to benefit from AI, firms must make proprietary knowledge accessible in ways that may expose it to deliberate extraction by rivals or inadvertent leakage through employees' use of AI tools.[2] A recent ruling (*Trinidad v. OpenAI*) reinforces this

[1] See Microsoft CEO Satya Nadella's recent article for further discussion. Nadella argues that firms must reveal institutional knowledge to make AI useful, yet doing so may expose the very knowledge they seek to protect; firms should therefore retain control over proprietary information within enterprise boundaries.

[2] Illustrating the possibility of deliberate attacks by rivals, OpenEvidence alleged that Doximity engineers used prompt-injection attacks to extract proprietary information from its AI product. Employee-related inadvertent leakage has also raised concerns across firms. See recent examples from Samsung, Amazon, and JPMorgan.

concern by holding that information shared with a GenAI tool without adequate safeguards does not qualify for trade secret protection (Toren, 2026).[3] Moreover, by helping rivals synthesize public information, product observations, and leaked fragments into usable technical knowledge, GenAI can substantially lower the cost of reverse engineering. Recent technical demonstrations show that GenAI can help translate product observations into editable designs, reconstruct material structures, and interpret obfuscated software (e.g., Johansen et al., 2026).[4] As secrecy becomes less reliable, patents become relatively more valuable because they remain enforceable against independent invention and reverse engineering, giving trade-secret-reliant firms stronger incentives to patent vulnerable know-how (Michael Best & Friedrich, 2026).[5]

On the other hand, patenting may also become more costly in the GenAI era. Patent applications require firms to disclose technical details publicly, potentially giving rivals a roadmap for learning from, imitating, or designing around an invention (Dyer et al., 2024; Zhang, 2024). GenAI may amplify this disclosure cost by accelerating innovation cycles and helping competitors more efficiently build on disclosed knowledge once patents are published (Besiroglu et al., 2024; McKinsey & Company, 2025; Calvino et al., 2025). Because patent prosecution often takes several years, competitors may use the disclosed technical details to

[3] Although firms can mitigate leakage risk through enterprise GenAI products, restrictions on public models, contractual and technical safeguards, and employee training, these measures remain costly and imperfect. For example, conventional data-loss-prevention systems often miss AI-chat leaks because such interactions resemble ordinary encrypted web traffic (Intuition Labs, 2024), while evolving extraction methods may outpace existing defenses. GenAI can therefore increase trade-secret leakage risk despite available safeguards.

[4] For example, Wonderful PCB describes AI systems that analyze images of a printed circuit board, identify its components and electrical connections, and automatically generate an editable circuit schematic. In software, Check Point researchers used GenAI to analyze heavily obfuscated malware, recover encrypted information, and explain its hidden logic—tasks that ordinarily require substantial manual expertise.

[5] In a recent public speech at the 2026 IPBC Global conference, Will Cheng, Senior Patent Counsel at LinkedIn, note that, in the age of AI, protecting trade secrets "has a surprisingly large cost" and that "[f]iling a patent on an iterative invention may just be preferable to attempting to preserve it as a trade secret." Our conversations with IP practitioners and attorneys are consistent with this view: they emphasized that AI can make trade secrets harder to preserve by facilitating reverse engineering and increasing leakage risks, which may make patents a more attractive form of protection for some firms.

develop alternatives or a newer generation of the technology before the patent is granted, leaving the applicant with rights over an invention that may have less commercial value by then. Recent industry remarks reflect similar concerns.[6] Moreover, GenAI may raise the bar for patentability by expanding what counts as prior art or what would be obvious to a person having ordinary skill in the art, an issue the USPTO has explicitly raised in its request for comments on AI, prior art, and the obviousness standard (U.S. Patent and Trademark Office, 2024; Gould, 2026).

To answer this question, we use the public launch of ChatGPT in November 2022 as a largely unanticipated shock to the availability and salience of GenAI, following recent accounting research (Bertomeu et al., 2025; Jia et al., 2025; Bradshaw et al., 2026). In our setting, the shock matters for innovation protection because it sharply increased the salience of GenAI-enabled leakage and reverse engineering for firms choosing between secrecy and patent protection. We examine a sample of innovating firms from 2019 to 2025 and compare changes in patent filing behavior across firms with different ex-ante reliance on trade secrecy using a continuous-treatment difference-in-differences design. Specifically, our treatment variable is the firm's pre-ChatGPT trade-secret intensity scaled by its R&D stock, capturing the extent to which the firm relies on secrecy to protect innovation relative to its innovation investment.

Our estimates suggest that, following the public launch of ChatGPT, innovators with greater ex-ante reliance on trade secrecy substantially increased their patent filings, consistent with a shift from secrecy toward patenting as the mechanism for protecting innovation. Specifically, firms with a one-standard-deviation higher level of pre-ChatGPT trade secret intensity exhibit

[6] In a February 9, 2026 post on X, Mark Cuban argued that LLMs can learn from patent disclosures and help competitors identify workarounds, potentially making patenting less attractive. This comment suggests that practitioners perceive GenAI as increasing the proprietary costs of patent disclosure, thereby pushing firms toward secrecy.

an approximately 14.8% larger increase in patent filings after ChatGPT's launch. The findings are supported by the absence of differential pre-trends between trade-secret reliant and less-reliant innovators, and are robust to specifications that absorb time-invariant firm characteristics, common time shocks, and observable time-varying firm controls.[7,8]

To explore the mechanisms underlying this increase in patent filings, we conduct a series of cross-sectional tests. First, if the patenting response reflects AI-related leakage risk, it should be stronger when GenAI is more likely to be used in firms' innovation and operations. Following Eisfeldt et al. (2026), we measure a firm's GenAI exposure based on the extent to which GenAI can help perform the tasks carried out by its workforce. Consistent with our prediction, the significant increase in patent filings is more than twice as large among high-GenAI-exposure, trade-secret-reliant firms as among their lower exposure counterparts. Further decomposing this measure, we find that the response is concentrated in exposure through core tasks tied to innovation, such as technical problem solving and research, with no significant differential response for supplemental tasks such as reporting and administrative support. Likewise, firms employing more AI-skilled workers are likely to use AI more intensively and thus face greater leakage risk. Consistent with this prediction, we obtain similar results using the number of AI-skilled employees as an alternative measure of firm-level AI adoption (Babina et al., 2024).

The second set of tests focuses on employees, who represent another primary conduit

[7] Because our treatment variable, *Trade Secret Intensity*, is continuous, we conduct additional analyses using the continuous-treatment DiD framework of Callaway et al. (2024). This framework does not impose a linear relation between treatment intensity and patent filings. Instead, it estimates the level effect at each positive treatment dose relative to the zero-dose group and then aggregates these effects over the distribution of positive doses. The results support our baseline inferences. First, a joint test fails to reject the null that the pre-event coefficients are jointly zero, providing no evidence of differential pre-trends across exposure groups. Second, the full-sample aggregate ATT is positive and directionally consistent with our baseline estimate. Third, the results are robust to replacing the continuous treatment measure with an above-median treatment indicator.

[8] Our results are also robust to alternative estimation approaches and outcome definitions, including Poisson estimation, an extensive-margin filing indicator, a truncation-adjusted filing measure, and a higher-frequency quarterly design. The results remain consistent after excluding software firms, suggesting that the effect is not driven solely by software companies or by post-ChatGPT filings of LLM-related AI patents.

through which trade secrets leak. Protecting a trade secret requires the firm to continuously restrict access to sensitive know-how and to monitor the people who carry it (Kang and Lee, 2022; Cunningham and Kapacinskaite, 2026), and innovation knowledge is harder to safeguard when it is dispersed across more inventors and more locations (Singh, 2005; Alcácer and Zhao, 2012). Consistent with this channel, the patenting response among trade-secret-reliant firms is substantially stronger for firms with larger inventor bases and for firms with more inventors located outside the headquarters state, where monitoring is more difficult. Additionally, we find that firms are more likely to shift toward patenting when they experience greater employee turnover, as departing inventors may carry proprietary know-how to competitors, and such knowledge becomes more valuable to its recipients when GenAI can help exploit it.

Third, we examine whether firms respond more strongly when their innovations face greater reverse engineering risk, a particularly important threat to trade secret protection in the age of AI. Following Kraft and Rammer (2025), we use survey evidence on firms' reverse engineering practices to construct an industry-level measure of reverse engineering prevalence. We find that the shift toward patenting is substantially stronger in industries where reverse engineering is more common, consistent with such risk being an important driver of firms' shift from trade secrets to patents. As complementary evidence, we find that the increase in patenting is concentrated in product rather than process patents. This pattern is consistent with product innovations being embodied in observable artifacts that rivals can more readily analyze and reverse engineer using AI (Cohen et al., 2000).

Having documented the shift toward patenting and examined the mechanisms underlying that shift, we next consider how GenAI affects knowledge diffusion. Although the shift from trade secrets to patents may increase knowledge disclosure overall, firms concerned about the

proprietary costs of patent disclosure can strategically choose what and when to disclose, potentially limiting the resulting gains in knowledge diffusion. Consistent with this view, we document four complementary patterns. First, firms delay the public release of their applications, lengthening the time between filing and publication (Glaeser and Landsman, 2021). Second, patent disclosures become shorter and contain fewer numerical details and figures, suggesting a decline in disclosure quality. Third, the increase in filings is concentrated in technology areas with faster obsolescence, measured using citation trends as in Ma (2025). In these areas, disclosed information is likely to lose its proprietary value more quickly, reducing the cost of patent disclosure. Fourth, firms file patents with broader claims, consistent with seeking wider exclusionary rights without providing commensurately richer technical disclosure. Taken together, these patterns suggest that GenAI-induced patenting may expand legal exclusion without a proportional increase in useful technical disclosure, potentially limiting the knowledge spillovers and subsequent innovation that patent disclosure is intended to promote (Galasso and Schankerman, 2015; Kim and Valentine, 2021).

Finally, we document several patterns consistent with the costs and benefits shaping firms' shift from secrecy to patenting in response to GenAI. First, the patenting response is stronger among firms with greater internal legal resources, consistent with legal expertise lowering the cost of assessing the secrecy risks created by GenAI and converting vulnerable proprietary knowledge into patent rights (Armstrong et al., 2026). Second, the increase in filings is concentrated in technology fields where firms that rely on trade secrets have less prior patent coverage. This pattern suggests that firms direct new applications toward areas where existing patent protection is weaker and the marginal value of converting trade secrets into patents is greater. Third, because a rejected patent application may disclose technical information without providing legal

protection, firms have stronger incentives to file when their applications are more likely to be granted. Consistent with this prediction, the increase in filings following ChatGPT's launch is concentrated in technology subclasses with more lenient examiners.

Our paper contributes to two related streams of literature. First, we contribute to the emerging accounting and finance literature on the consequences of GenAI. Recent studies use the public launch of ChatGPT to examine firms' information environments, corporate attention to GenAI, and AI use by capital market intermediaries (Chang et al., 2026; Blankespoor et al., 2026; Ecker et al., 2026; Kim et al., 2026). We shift the focus from information production and disclosure in capital markets to firms' real decisions about how to protect innovation and the resulting implications for knowledge spillovers. Our findings show that GenAI can redraw the boundary between private knowledge and public disclosure, shaping how innovation is protected, disclosed, and diffused.

Second, we contribute to the literature on patent disclosure, innovation protection, and proprietary costs. Prior research shows that firms manage the timing, content, and scope of patent disclosures and that disclosure quality shapes knowledge diffusion and follow-on innovation (Glaeser et al., 2020; Glaeser and Landsman, 2021; Shao et al., 2023). Related studies document how firms' ability to protect trade secrets and retain proprietary knowledge shapes their disclosure choices (Glaeser, 2018; Li et al., 2018; Aobdia, 2018; Boot and Vladimirov, 2025). We connect these literatures by showing that an external technological shock—the public launch of ChatGPT—changes firms' choice of innovation-protection mechanism and subsequent patent disclosures. Trade-secret-reliant firms increase patent filings, particularly when leakage and reverse-engineering risks are more pronounced. Importantly, firms manage the composition, timing, scope, and information content of their

filings, suggesting that increased patenting can reflect a shift from secrecy to legal exclusion rather than a proportional increase in inventive activity or useful technical disclosure.

2 Institutional Background and Hypothesis Development

2.1 Patents, Trade Secrets, and the Proprietary Costs of Disclosure

Firms can protect different aspects of an innovation through patents, trade secrecy, or a combination of the two. A patent grants its owner the right to exclude others from practicing the claimed invention, even when competitors independently develop or reverse engineer the technology. In return, however, the applicant must publicly disclose the invention in sufficient detail to enable a person skilled in the relevant field to practice it. This broader right to exclude therefore comes at a cost: Patent applicants face uncertainty over the scope of the claims ultimately granted (Armstrong et al., 2026), while disclosure provides competitors with technical information they may use to learn from the invention, develop superior alternatives, or invent around the patent (Glaeser et al., 2020; Glaeser and Lang, 2024). Moreover, obtaining and enforcing patent rights is costly and time-consuming, and litigation outcomes remain uncertain (Kim et al., 2025). Even when a patent holder prevails in court, the remedy may not include an injunction preventing the competitor from continuing to sell its product, allowing product-market competition to persist (Mezzanotti, 2021).

Trade secret protection operates differently. It requires neither registration nor public disclosure and can continue indefinitely as long as the information remains secret. To qualify for protection, the information must derive economic value from not being generally known, and its owner must take reasonable measures to preserve its secrecy. Trade secret law provides recourse against improper acquisition, use, or disclosure, but generally not against independent

discovery or lawful reverse engineering. Thus, when a competitor independently develops or lawfully reverse engineers the protected knowledge, the firm cannot rely on trade secret law to prevent that competitor from using it (Friedman et al., 1991; Lemley, 2008; Glaeser, 2018).

Survey evidence indicates that firms often regard secrecy as at least as important as patents, with the relative value of the two mechanisms depending on the characteristics of the innovation and the environment in which it is developed (Arundel and Kabla, 1998; Cohen et al., 2000; Arundel, 2001; Jankowski, 2012). Secrecy is more attractive when knowledge can be confined within the firm, when the underlying technology is difficult to infer from the final product, and when disclosure would provide competitors with valuable technical guidance. Patenting becomes more attractive when an innovation is readily observable, when reverse engineering is feasible, or when the expected value of enforceable exclusion exceeds the costs of disclosure and prosecution.

This tradeoff over how to protect innovation also connects directly to the literature on proprietary costs of disclosure. In standard disclosure models, firms optimally withhold information when disclosure would transfer valuable knowledge to competitors (Verrecchia, 1983; Einhorn, 2007; Einhorn and Ziv, 2008). Consistent with these models, managers identify the loss of competitive information as an important reason for limiting disclosure (Graham et al., 2005). Although disclosure is mandatory once a firm applies for a patent, the decision to enter the patent system remains discretionary. When keeping secrecy is feasible, a firm can avoid disclosure by retaining the innovation as a trade secret. The relevant proprietary cost of patenting is therefore the knowledge disclosed through an application that could otherwise have remained secret. When secrecy becomes more difficult to maintain, the marginal cost of disclosure falls because competitors are more likely to obtain the underlying knowledge even without patent disclosure. Consistent with this logic, prior work shows that firms' disclosure

choices respond to changes in their ability to preserve proprietary information (Aobdia, 2018; Glaeser, 2018; Li et al., 2018; Azinovic-Yang et al., 2025).

2.2 Hypothesis Development: GenAI and Innovation Protection

2.2.1 Challenges to Trade Secret Protection

Since the public launch of ChatGPT in late 2022, GenAI has rapidly entered firms' routine operations and innovation activities. While GenAI can enhance firms' operations and innovation processes, its diffusion has outpaced many firms' capacity to govern how proprietary information is supplied to AI systems and how those systems store, process, or reuse it (LayerX Security, 2025; Cyera and Cybersecurity Insiders, 2025). GenAI can therefore challenge trade secret protection through several distinct channels (Levine, 2023; Hrdy, 2025).

First, GenAI may create new channels for inadvertent trade secret leakage, weakening trade secrecy as a mechanism for protecting innovation. As these tools become embedded in everyday tasks, prompts and file uploads may transmit source code, technical documents, and internal data to external systems that firms neither monitor nor fully control. Industry research finds that 77% of employees paste data directly into generative AI tools, with 82% of that activity occurring through unmanaged personal accounts that fall entirely outside corporate visibility (LayerX Security, 2025). Another survey finds that sensitive corporate information has been found in 22% of uploaded files, spanning M&A planning, financial modeling, and proprietary code (Morrone, 2025). Several reported incidents, for example at Samsung and JPMorgan, illustrate how employees' use of GenAI can cause proprietary information to leak, prompting many firms (e.g., Amazon) to restrict these tools and caution employees against entering confidential material (centrexIT, 2025). Despite these risks, oversight remains limited: only 13% of surveyed enterprises report strong visibility into

internal AI use, while 57% cannot block risky activity in real time (Cyera and Cybersecurity Insiders, 2025). Effective governance takes time to develop and is costly to implement. Firms must map AI use and data flows, establish policies and vendor protections, train employees, and, in some cases, make substantial investments in computing capacity, security expertise, and ongoing maintenance to operate locally hosted AI systems.

What makes this risk legally consequential is that inadvertent disclosure of trade secrets through generative AI may strip those secrets of legal protection under the Defend Trade Secrets Act (DTSA). The DTSA requires trade secret owners to take "reasonable measures" to maintain secrecy, and a recent ruling holds that sharing proprietary information with a public AI platform fails that standard. In *Trinidad v. OpenAI* (2026), a federal court dismisses the plaintiff's DTSA claims because she has voluntarily disclosed her proprietary frameworks to OpenAI while using ChatGPT to develop them. The decision signals that sharing proprietary information with a public generative AI platform, absent adequate contractual protections, is essentially equivalent to releasing that information into the public domain (Toren, 2026).

Second, GenAI can create a distinct leakage channel through deliberate attempts by rivals to extract proprietary information directly from a firm's deployed AI interface. Unlike inadvertent leakage through routine employee use, this risk arises when firms deploy GenAI in products or operational workflows that embed proprietary instructions, data, workflows, and domain knowledge. Microsoft Chairman and CEO Satya Nadella describes this tension as the "reverse information paradox," adapting Arrow's information paradox to enterprise AI: firms must reveal institutional knowledge to obtain value from AI, yet doing so may expose the very knowledge they seek to protect. Nadella therefore argues that firms should retain control over proprietary information and keep it within enterprise boundaries. Even when firms do so,

however, a deployed AI system with an external interface can create a new attack surface through which outsiders may attempt to recover protected information. Prompt injection, for example, uses adversarial instructions to induce a model to reveal hidden system prompts or other proprietary content. A recently filed lawsuit illustrates this risk. In *OpenEvidence Inc. v. Doximity, Inc.* (2025), OpenEvidence alleged that Doximity personnel used unauthorized credentials and prompt injection queries in an attempt to extract its system prompt and other proprietary information from its clinical AI product.

Third, generative AI also raises the risk that trade secrets will be uncovered through reverse engineering, the one avenue against which secrecy offers no legal recourse. Reverse engineering has traditionally been a costly, time-consuming process that required a skilled reverse engineer to recover a product's underlying design, formula, or code (Samuelson and Scotchmer, 2002). Generative AI can help convert the information embodied in a distributed product into usable competitive knowledge, and evidence of AI-assisted reverse engineering now spans both physical products and software. For example, on the manufacturing side, commercial vendors now market AI services that reconstruct a product's design from external measurements alone: one firm advertises a tool that photographs and X-rays a finished circuit board and automatically generates a complete schematic, claiming a roughly 70% reduction in time and an explicitly lower skill barrier (Wonderful PCB, 2025), while comparable tools perform chemical deformulation, recovering the ingredient composition of a proprietary formulation from a finished sample (DKSH, 2024; National Polymer, 2024). On the software side, large language models can read the compiled, obfuscated code that ships inside a product and explain in plain language what it does. They can recover logic and structure that firms deliberately strip away to keep secret, a task that typically requires rare reverse-engineering

expertise (Bukhteyev, 2025). Taken together, these examples illustrate how GenAI can reduce the time and expertise required to infer how a product works, thereby weakening a practical barrier on which trade secrecy depends.

2.2.2 Challenges to Patent Protection

On the other hand, GenAI can increase the cost of protecting innovation through patents by widening the timing mismatch between patent disclosure and patent protection. Patent applications must describe inventions in substantial technical detail and often become public before examination is complete and the final scope of protection is known (Kim and Valentine, 2021; Dyer et al., 2024). Once an application is published, GenAI can help competitors rapidly interpret the disclosure, combine it with other technical knowledge, and develop improvements or designs that avoid the proposed claims. The published application can therefore serve as a technical roadmap for competitors while the applicant remains uncertain about whether, and to what extent, enforceable rights will ultimately be granted.

This timing problem is particularly important because GenAI may accelerate research and product development, while patent prosecution can still take several years (Besiroglu et al., 2024; McKinsey & Company, 2025; Calvino et al., 2025). By the time a patent is granted, competitors may already have used its disclosure to develop a newer generation of the technology that falls outside the final claims. The patented invention may therefore have lost much of its commercial relevance, even though its disclosure helped rivals develop improved alternatives. In this setting, the applicant bears the proprietary cost of disclosure but receives exclusionary rights over technology that may already be outdated. This risk is especially pronounced in fast-moving fields and is reflected in industry commentary warning that LLMs can learn from published patent disclosures and help competitors identify workarounds (Cuban, 2026).

Moreover, GenAI may effectively raise the bar for patentability. The proliferation of publicly available AI-generated technical material could expand the relevant prior art, while widespread access to AI tools may alter what a person having ordinary skill in the art could readily accomplish, causing some inventions to be viewed as obvious. Firms may therefore disclose valuable technical information in their applications while facing a greater risk of rejection or receiving narrower patent rights. The legal treatment of these issues remains unsettled, but the USPTO has expressly sought public comment on how AI may affect prior art, ordinary skill, obviousness, and patentability (U.S. Patent and Trademark Office, 2024; Gould, 2026). This uncertainty increases the expected cost of prosecution and reduces the likelihood that disclosure ultimately produces meaningful exclusionary rights.

2.2.3 Hypothesis Development

The preceding discussion yields competing predictions for innovative firms' choice between patents and trade secrecy. This tradeoff should be especially consequential for firms that relied heavily on trade secrecy before GenAI became widely available. These firms have more proprietary knowledge exposed to the new risks created by GenAI and a greater scope to substitute from secrecy toward patenting.

If the declining durability of secrecy outweighs the increased costs of patent protection, firms with greater prior reliance on trade secrecy should increase their patent filings more strongly following the public launch of ChatGPT. This possibility has also attracted attention from intellectual property practitioners. At IAM's IPBC Global conference, LinkedIn Senior Patent Counsel Will Cheng suggests that, as AI makes trade secret risks harder to manage, firms may need to become more comfortable using patents to protect innovation (IAM, 2026). However, if the increased costs and reduced expected benefits of patenting instead dominate, trade secret

reliant firms may continue relying on secrecy despite its greater vulnerability, resulting in no increase or even a decline in patent filings.

Existing theory does not determine which force dominates. We therefore state the hypothesis in null form:

H1: The emergence of GenAI does not lead firms that previously relied more heavily on trade secrecy to increase their use of patent protection.

3 Empirical Design and Summary Statistics

3.1 Data Sources

We obtain data on applications and grants of U.S. patent from PatentsView, a database maintained by the USPTO. PatentsView provides cleaned and organized patent data, including application and grant dates, assignee and inventor identities, inventor locations, technology classifications, and examiner information.

To match patent applications to U.S. public firms, we use the patent-to-PERMNO linkage developed by Kogan et al. (2017), hereafter KPSS, to construct a bridge between PatentsView's persistent disambiguated assignee identifiers and CRSP PERMNOs. We retain unambiguous assignee-to-PERMNO links directly. When an assignee identifier is associated with multiple PERMNOs, we resolve the ambiguity using exact standardized-name matches and, when necessary, token-based textual similarity between PatentsView assignee names and public-firm names, following the matching approach of Kim and Valentine (2023). We then apply the resulting assignee-to-PERMNO bridge to PatentsView's pre-grant publication records. The persistence of PatentsView's disambiguated assignee identifiers allows us to identify published applications filed during 2023–2025. We aggregate distinct applications to

the public firm-year level based on their filing dates and obtain the corresponding financial data from Compustat.

3.2 Key Variables

3.2.1 Patent Filings

To capture firms' decisions to seek patent protection, we define *Patent Filings* as the inverse hyperbolic sine (IHS) transformation of the number of distinct published patent applications filed by a firm in a given year and observed in PatentsView.[9] We assign each application to its filing year rather than its publication or grant year and merge it with the corresponding Compustat firm-year. We focus on applications rather than grants because an application records the firm's decision to seek patent protection when that decision is made.

3.2.2 Trade Secret Intensity

To measure firms' reliance on secrecy before ChatGPT became publicly available, we construct *Trade Secret Intensity*. We first aggregate each firm's references to trade secrets in its Form 10-K filings over 2019–2021, based on the data provided by Glaeser (2018).[10] We then scale this total by the firm's R&D stock as of 2019, which we construct from R&D expenditures over the preceding 20 years using a 15% annual depreciation rate. We define *Trade Secret Intensity* as the inverse hyperbolic sine (IHS) transformation of this scaled measure.Scaling by accumulated R&D measures reliance on secrecy relative to the firm's stock of innovation investment and reduces the influence of firm size. A higher value of *Trade Secret Intensity* therefore indicates that a firm discusses trade secrecy more extensively relative to its

[9] Because patent filings are highly right skewed and contain many zero observations, this transformation reduces skewness while retaining firm years with no observed filings.

[10] Glaeser (2018) validates this disclosure-based measure by showing that it is concentrated in innovation-intensive industries, aligns with survey evidence on trade-secret use, overlaps substantially with a litigation-based measure of trade-secret value, and is associated with predicted determinants of trade secrecy, including stronger legal protections, firm size, and R&D intensity.

accumulated R&D investment. Because the measure is constructed entirely from information available before ChatGPT's launch and remains fixed throughout the sample period, it cannot mechanically respond to the firm's subsequent patent filings.[11]

3.2.3 Other Control Variables

We control for firm characteristics that may be associated with changes in patent filings, including *Size*, *ROA*, *Loss*, *Leverage*, *BTM*, *Cash/Assets*, *Capex/Assets*, *PPE/Assets*, *Sales Growth*, *R&D/Assets*, *R&D Stock*, and *Patent Stock*. *BTM* is measured as total assets divided by the sum of market capitalization and total debt. *R&D Stock* and *Patent Stock* are based, respectively, on R&D expenditures and granted patents over the preceding 20 years; both are constructed using a 15% annual depreciation rate and transformed using the inverse hyperbolic sine. Appendix A provides complete variable definitions.

3.3 Sample Construction

We begin with the universe of U.S. public firms covered by Compustat from 2019 through 2025. Our research design exploits the widespread introduction of generative AI following the launch of ChatGPT in November 2022 as the shock and compares the three years before and after its introduction. Accordingly, we define 2019–2021 as the pre-period and 2023–2025 as the post-period, excluding 2022 because ChatGPT was released near the end of the year. To facilitate the difference-in-differences analysis, we require firms to be present in both the pre- and post-periods. We also require firms to have a positive R&D stock as of 2019 as R&D stock serves as the denominator in our trade-secret intensity measure. Finally, to focus on firms

[11] One caveat of this trade secret intensity measure is that it is necessarily indirect. Because secrecy is by nature unobservable, the measure does not capture a firm's complete stock of secret knowledge, but rather the intensity with which the firm discusses trade secrets in its mandatory filings. It may therefore reflect variation in both firms' underlying reliance on secrecy and their disclosure practices. Accordingly, we use it to rank firms by their prior reliance on trade secrecy rather than to estimate the number or value of their trade secrets.

with an established capacity for patentable innovation and to construct measures of firms' ex ante innovation characteristics for the cross-sectional analyses, we require firms to have filed at least one patent during the 10-year period preceding the sample period. These sample selection procedures yield 5,294 firm-year observations representing 929 patent-active public firms.

3.4 Research Design

We use a continuous-treatment difference-in-differences design to examine whether firms that rely more heavily on trade secrecy before the introduction of generative AI subsequently change their patenting behavior differently from firms that rely less heavily on trade secrecy. Our baseline specification is

$$Patent\ Filings_{it} = \beta\ Trade\ Secret\ Intensity_i \times Post_t + \gamma X_{it} + \alpha_i + \delta_t + \varepsilon_{it}. \tag{1}$$

where i and t index firms and years, respectively, and α_i and δ_t denote firm and year fixed effects. Firm fixed effects absorb all time-invariant differences across firms, including stable variation in technology focus, innovation strategies, and baseline propensities to patent. Year fixed effects absorb changes in patenting common to all firms, including the main effect of *Post*. We cluster standard errors by firm to allow for serial correlation within firms.

The identification assumption of the DID design is that, absent the introduction of widely accessible generative AI, firms with different levels of pre-event trade-secret intensity would have experienced similar changes in patent filings. To evaluate this assumption, we estimate the following event-study specification:

$$Patent\ Filings_{it} = \sum_{\tau \neq 2021} \beta_\tau \left(Trade\ Secret\ Intensity_i \times \mathbf{1}[t = \tau] \right) + \gamma X_{it} + \alpha_i + \delta_t + \varepsilon_{it}. \tag{2}$$

where 2021 is the omitted reference year. The specification estimates separate trade-secret-

intensity coefficients for 2019, 2020, and each post-event year from 2023 through 2025.

3.5 Summary Statistics

Table 1 reports the industry distribution and summary statistics for the sample. Panel A shows that the sample is concentrated in manufacturing, which accounts for 742 firms, or 79.9% of the sample.[12] Services is the second-largest sector, with 122 firms, or 13.1%. Other industries account for relatively small shares of the sample, consistent with our focus on patent-active public firms. Mean trade-secret intensity also varies across industries, with relatively high values in retail trade and mining.

Panel B presents summary statistics for the main variables. More than 40% of firm years file patents in our sample period. The IHS-transformed patent-filing variable has a mean of 1.316 and a median of zero, indicating that annual patenting activity remains right-skewed even among patent-active firms. The IHS-transformed trade-secret-intensity variable is also highly skewed, with a mean of 0.293 and a median of 0.040.

4 Empirical Results

4.1 GenAI, Trade Secret Reliance, and Patent Filings

4.1.1 Main Results

Table 2 reports the estimates from the main difference-in-differences analysis specified in Equation (1). Column (1) reports the baseline specification without fixed effects or controls. Column (2) adds time-varying firm-level controls, while Column (3) further includes firm and

[12] The concentration of manufacturing firms reflects our sample selection criteria, which require firms to have positive pre-period R&D stock and patenting activity. Public firms with active R&D and patenting are disproportionately represented in manufacturing industries, consistent with the sample selection approach in Arora et al. (2021). This pattern suggests that our findings are not limited to core AI developers, but also extend to firms that may integrate GenAI into their operations and innovation processes (Chen et al., 2024).

year fixed effects to absorb, respectively, time invariant differences across firms and common changes in patenting over time. Across these specifications, the coefficients on *Trade Secret Intensity* × *Post* are positive and statistically significant at the 1% level (coef.=0.221, s.e.=0.030 in Column (1); coef.=0.230, s.e.=0.031 in Column (2); coef.=0.226, s.e.=0.029 in Column (3)). Notably, although the R^2 increases from 0.07 in Column (1) to 0.89 in Column (3), largely because of the inclusion of firm fixed effects, the coefficient of interest changes only modestly. In the spirit of Oster (2019), the combination of a substantial increase in explanatory power and a stable coefficient provides reassurance that the documented differential increase in patent filings is unlikely to be driven solely by omitted firm characteristics.

To assess the economic magnitude, we rely on the fully specified model in Column (3), which yields a coefficient of 0.226. Based on the standard deviation of *Trade Secret Intensity* reported in Table 1, firms with a one-standard-deviation higher level of pre-ChatGPT trade secret intensity exhibit a 0.138-unit larger post-ChatGPT increase in IHS-transformed patent filings, corresponding to an approximate 14.8% differential increase in the underlying number of applications.[13] Overall, these results are consistent with the view that GenAI increases the relative cost of maintaining secrecy, leading trade-secret-reliant firms to rely more heavily on patent protection after the launch of ChatGPT.

4.1.2 Dynamic Evidence

The difference-in-differences design relies on the assumption that, absent ChatGPT's launch, firms with different levels of prior trade secret reliance would have followed similar patenting trends. We assess this assumption using the event time specification in Equation (2).

[13] Because the dependent variable is inverse hyperbolic sine (IHS)-transformed, we apply the standard log-like approximation, $\exp(\beta \times \sigma) - 1$, where β is the coefficient on *Trade Secret Intensity* × *Post* and σ is the standard deviation of *Trade Secret Intensity*. In Column (3), this calculation is $\exp(0.226 \times 0.609) - 1 = 0.148$, implying an approximately 14.8% increase in patent filings for a one-standard-deviation increase in trade-secret intensity.

Figure 1 plots the dynamic coefficients from interacting pre-event trade-secret intensity with year indicators, using 2021 as the reference year. The coefficients for the pre-event years, 2019 and 2020, are close to zero and statistically insignificant, suggesting that firms with greater and lower trade-secret reliance did not exhibit differential patenting trends before the launch of ChatGPT. In contrast, the coefficient becomes positive in 2023 and increases steadily thereafter, consistent with firms that previously relied more heavily on trade secrecy increasing their patent filings following the widespread availability of GenAI. This pattern supports the interpretation that the main results are driven by a post-ChatGPT change in patenting behavior rather than by pre-existing divergent trends.

The post-launch dynamics offer two additional insights. First, the response emerges within the first full year after ChatGPT's launch, a timing more consistent with firms rapidly changing how they protect existing proprietary knowledge than with firms increasing innovation activity itself, which would generally take longer to produce patentable inventions. Second, the progressively larger estimates are also consistent with secrecy risks intensifying as AI models become more capable and efficient, particularly as recent agentic systems improve their ability to synthesize technical information and facilitate reverse engineering (Kwa et al., 2025; Crawford et al., 2026). Taken together, these patterns support the interpretation that our findings reflect a post ChatGPT change in firms' innovation protection strategies rather than preexisting differences in patenting trends.

4.2 Mechanism Tests: The Information Leakage Channel

To explore the channels underlying the shift from secrecy toward patenting, we first examine whether the response is stronger when firms face greater information leakage risk. If GenAI makes secrecy more costly by increasing the risk that proprietary knowledge is exposed

or becomes harder to control, then trade-secret-reliant firms should respond more strongly when their operations, inventors, and employees create more opportunities for such leakage.

4.2.1 GenAI Exposure and AI Capability

If our documented patenting response reflects AI-related leakage risk, it should be stronger when GenAI is more likely to be used in firms' innovation and operations. We first use the GenAI exposure measures developed by Eisfeldt et al. (2026), which capture the extent to which GenAI can help perform the tasks carried out by a firm's workforce. Firms with greater task exposure are more likely to integrate GenAI into their innovation activities, often supplying AI systems with firm-specific data and domain knowledge to generate useful outputs. These interactions create additional opportunities for proprietary information to leave the firm, weakening the relative value of secrecy and strengthening firms' incentives to shift toward patent protection. The results reported in Panel A of Table 3 are consistent with our prediction. As shown in Columns (1) and (2), the coefficient on *Trade Secret Intensity* $\times$ *Post* is positive and significant in both the low-exposure subsample (coef.=0.143, s.e.=0.033 in Column (1)) and the high-exposure subsample (coef.=0.390, s.e.=0.061 in Column (2)), but is approximately 2.7 times as large among high-exposure firms, and the difference (0.247) is statistically significant at the 1% level.

We further follow Eisfeldt et al. (2026) and distinguish between core and supplemental exposure. Core exposure captures tasks closely connected to firms' operations and innovation, including technical problem solving, product development, and research, where information supplied to GenAI is more likely to contain proprietary knowledge. Supplemental exposure instead captures support activities such as reporting, documentation, and administrative work, which are less likely to involve the technical knowledge underlying firms' innovations. We

repeat the analysis using these two measures in Columns (3) through (6). The high-minus-low difference is positive and significant for core exposure (coef.= 0.227; s.e.= 0.068), but smaller and statistically insignificant for supplemental exposure (coef.= 0.078; s.e.= 0.064). This pattern suggests that the shift toward patenting is concentrated among firms whose core activities are more exposed to GenAI, rather than among firms whose exposure arises primarily from peripheral tasks.

Table 3 Panel B provides a complementary test based on firms' capacity to adopt AI. Firms employing more AI-skilled workers possess greater human capital for implementing AI and are therefore more likely to use it intensively, creating additional opportunities for proprietary information to enter AI systems. Following Babina et al. (2024), we measure this capacity using the average number of employees with AI-related skills during 2009–2018. The coefficient on *Trade Secret Intensity* × *Post* is 0.422 among firms with more AI-skilled employees, compared with 0.124 among firms with fewer such employees. The estimated response is approximately 3.4 times as large in the high-capacity group, and the difference of 0.298 is statistically significant at the 1% level. These results are consistent with firms that use AI more intensively facing greater proprietary-information leakage risk and therefore finding patent protection relatively more attractive.

4.2.2 Inventor Networks and Employee Turnover

Employees represent another important channel through which trade secrets can leave the firm. Protecting proprietary knowledge requires firms to restrict access to sensitive information and monitor the employees who possess it (Kang and Lee, 2022; Cunningham and Kapacinskaite, 2026). These tasks become more difficult when innovative knowledge is distributed across a larger inventor network or across geographically dispersed locations (Singh, 2005; Alcácer and Zhao,

2012). We therefore expect the patenting response to be stronger when more employees possess proprietary knowledge and when those employees are more difficult to monitor.

Table 4 Panel A examines these predictions. Columns (1) and (2) partition firms based on the number of unique inventors. The coefficient on *Trade Secret Intensity* × *Post* increases from 0.067 among firms with smaller inventor networks (s.e. = 0.026) to 0.572 among firms with larger inventor networks (s.e. = 0.134). The high-minus-low difference is 0.505 and statistically significant at the 1% level (s.e. = 0.137). Columns (3) and (4) produce a similar pattern based on geographic dispersion. The estimated patenting response is approximately 3.4 times as large among firms with more geographically dispersed inventors as among firms with fewer inventors outside the headquarters state, and the difference is statistically significant at the 1% level. These findings are consistent with greater monitoring costs and more potential leakage points making secrecy more difficult to preserve.

Table 4 Panel B focuses on employee turnover, a more direct channel through which proprietary knowledge may leave the firm. GenAI may increase employee turnover by displacing workers whose tasks become automated. And those departing employees can carry technical know-how to competitors or new employers, and GenAI may further make that knowledge easier for recipients to interpret and use. This concern is particularly relevant for scientists and engineers, who are more likely to possess technical know-how related to the firm's innovation activities. Consistent with these predictions, the patenting response is stronger among firms with larger increases in overall employee turnover (Diff. (High - Low) =0.197, s.e.=0.063) and among firms with larger increases in turnover among scientists and engineers (Diff. (High - Low) =0.115, s.e.=0.060). These findings suggest that trade-secret reliant firms shift more strongly toward patenting when proprietary knowledge is more likely to

leave the firm through employee mobility.

4.3 Mechanism Tests: Reverse Engineering as a Threat to Secrecy

We next examine whether firms respond more strongly when their innovations face greater risk of reverse engineering. This threat is particularly important for trade-secret protection because lawful reverse engineering allows competitors to use the recovered knowledge without liability (Samuelson and Scotchmer, 2002). By reducing the time and expertise required to infer how an observable product or technology works, GenAI can weaken a practical barrier on which secrecy depends. If this mechanism contributes to the shift toward patenting, the response should be stronger where reverse engineering is more prevalent and where innovations are more readily observable.

Table 5 Panel A examines ex ante reverse-engineering exposure, measured as the industry-level reverse-engineering prevalence, constructed using Eurostat CIS 2018 survey evidence on firms' use of reverse engineering as a source of information for innovation, following Kraft and Rammer (2025).[14] The coefficient on *Trade Secret Intensity* $\times$ *Post* is 0.201 in industries with low reverse-engineering prevalence (s.e. = 0.033), compared with 0.358 in industries with high prevalence (s.e. = 0.067). Thus, the estimated response is approximately 1.8 times as large in high-prevalence industries. The difference of 0.157 is statistically significant at the 5% level (s.e. = 0.075). These results are consistent with the view that firms shift more strongly toward patenting when secrecy is more vulnerable to reverse engineering.

[14] We construct the reverse-engineering exposure measure at the SIC3 industry level using Eurostat CIS 2018 and SIC-based Fama-French industry descriptions. First, we use the Eurostat CIS 2018 variable reporting the percentage of firms that identify reverse engineering as a source of information for innovation. For each NACE Rev. 2 industry, we take the median value across countries. We then expand the SIC code ranges in the Ken French/Fama-French 48 industry portfolio definitions to SIC3 industries and assign each SIC3 industry a SIC-based industry description. Each SIC3 industry is matched to the closest NACE Rev. 2 industry using text similarity between the SIC-based description and the NACE industry label. The SIC3 reverse-engineering exposure measure is the cross-country median reverse-engineering prevalence of the best-matched NACE industry. We then merge this SIC3-level measure to firms based on their pre-event SIC3 industry.

Table 5 Panel B provides complementary evidence by separating filings into process and product patents, following Bena et al. (2022). The increase in patenting is concentrated in product patents rather than process patents. Taken together, Table 5 suggests that reverse-engineering risk helps explain why trade-secret-reliant firms shift toward patenting after the launch of ChatGPT.

4.4 Strategic Patent Disclosure and Knowledge Diffusion

Having documented the shift toward patenting and examined its underlying mechanisms, we next consider the implications for knowledge diffusion. Although the shift from trade secrets to patents may increase public disclosure overall, firms concerned about the proprietary costs of patent disclosure can strategically choose what to disclose and when, potentially limiting the resulting increase in useful public knowledge. Table 6 examines whether firms manage the timing, content, technological composition, and claim scope of their applications after shifting toward patenting.

4.4.1 Disclosure Timing and Information Content

Table 6 Panel A examines patent disclosure timing (Glaeser and Landsman, 2021). If firms seek patent protection while limiting the immediate release of technical information, they may delay the public publication of their applications. We classify applications as fast-disclosure filings if they are published within 18 months of the filing date, and as slow-disclosure filings if they are published 18 months or more after the filing date. Consistent with this prediction, we document fewer fast-disclosure filings (coef. $= -0.064$, s.e. $= 0.009$ in Column (1)) but more slow-disclosure filings (coef. $= 0.232$, s.e. $= 0.027$ in Column (2)). Additionally, Column (3) shows that the average time between filing and publication increases significantly (coef. $= 0.773$; s.e. $= 0.120$). Together, these findings suggest that trade-secret-reliant firms seek

more patents after ChatGPT's launch while delaying when the underlying technical information becomes publicly available.

Table 6 Panel B examines the information content of patent disclosures, including disclosure length, numerical details, and the number of figures and drawing sheets. The coefficients on *Trade Secret Intensity* $\times$ *Post* are all negative and statistically significant. These patterns suggest that, conditional on filing patents, trade-secret-reliant firms provide shorter and less detailed technical disclosures. This evidence is consistent with firms seeking legal protection while limiting the amount of useful information revealed to competitors.

4.4.2 Technological Obsolescence and Claim Scope

Table 6 Panel C examines whether firms patent more when the benefits of obtaining patent rights are greater relative to the costs of disclosure. In rapidly obsolescing technology fields, disclosed technical knowledge loses its competitive value more quickly, reducing the lasting advantage that disclosure provides to rivals. A patent may nevertheless provide enforceable rights over the claimed invention, making patent protection relatively more attractive. Following Ma (2025), we measure technological obsolescence using citation trends and divide technology fields at the sample median. The coefficient on *Trade Secret Intensity* $\times$ *Post* is positive and significant in high-obsolescence fields (coef. $= 0.047$, s.e. $= 0.013$), but small and insignificant in low-obsolescence fields (coef. $= 0.006$, s.e. $= 0.013$). This pattern is consistent with firms being more willing to patent when the value of securing exclusionary rights is higher relative to the proprietary cost of disclosure.

Finally, we examine the scope of legal protection that firms seek in their patent applications. Table 6 Panel D uses the length of independent claims as a proxy for claim scope, because shorter claims generally contain fewer limitations and therefore seek broader protection. The coefficient

on *Trade Secret Intensity* × *Post* is negative and statistically significant (coef. = −0.154, s.e. = 0.048). This result suggests that trade-secret-reliant firms seek broader exclusionary rights as they shift toward patenting, thereby increasing the expected legal benefit of entering the patent system.

Taken together, the results in Table 6 are consistent with firms seeking broader legal protection while managing the proprietary costs of patenting through their choices of which technologies to patent, when their applications become public, and how much technical detail they disclose. These strategic choices have important implications for knowledge diffusion. Although shifting from secrecy to patenting can make more knowledge available to follow-on innovators, these benefits may be attenuated when applications become public later, contain less technical detail, and disclose knowledge that loses value quickly. The GenAI-induced shift toward patenting may therefore expand legal exclusion without producing a proportional increase in useful technical disclosure and knowledge spillovers.

4.5 Additional Analyses

4.5.1 Internal Legal Resources

Converting vulnerable know-how into patent rights requires firms not only to recognize the changing risks of secrecy but also to possess the legal capacity to respond. Internal legal expertise can help firms identify patentable knowledge, evaluate disclosure risks, draft applications and claims, and navigate patent prosecution. We therefore expect the patenting response to be stronger when firms have greater internal legal capacity and face lower costs of converting trade secrets into patent applications. Table 7 tests this prediction by partitioning firms based on their pre-ChatGPT employment of IP-experienced legal personnel.[15] The estimated patenting

[15] We follow Ahn et al. (2025) to identify legal employees with IP experiences.

response is approximately three times as large among firms with greater legal resources, and the high-minus-low difference is statistically significant at the 1% level. These results are consistent with internal legal capacity helping firms translate concerns about GenAI-related secrecy risks into patent filings.

4.5.2 Expected Returns to Seeking Patent Protection

Table 8 next examines whether firms direct new applications toward settings where the expected benefits of patenting are greater. Panel A separates filings into core and non-core technology fields. The response is concentrated in non-core fields, where firms have less prior patent coverage and the marginal benefit of obtaining additional protection is likely higher.

Panel B examines how expected examination outcomes shape firms' filing incentives. A rejected application may reveal technical information without conferring enforceable patent rights, making patenting less attractive when rejection is more likely. We classify technology subclasses as strict or lenient based on whether their historical rejection rates are above or below the sample median. The estimated patenting response is positive and statistically significant in lenient fields but small and insignificant in strict fields, consistent with firms filing more when disclosure is more likely to result in patent protection.

4.5.3 Robustness Tests

In Table 9, we conduct several robustness tests. Panel A examines whether the main result is driven by software firms, which may be especially affected by GenAI and may rely heavily on both trade secrets and patenting. The coefficient on *Trade Secret Intensity* $\times$ *Post* is positive but statistically insignificant for software firms, while the coefficient remains positive and statistically significant for non-software firms. The difference between the two groups is not statistically significant. These results indicate that the positive patenting response extends beyond the

software industry.

Panel B evaluates the robustness of the main result to alternative model specifications and outcome definitions. First, we use Poisson estimation to address the count nature and skewness of patent filings. Second, we use an extensive-margin indicator for whether a firm files any patent application in a given year, which tests whether the result reflects a broader shift into patenting rather than only changes in filing intensity among firms that already patent. Third, we use a truncation-adjusted patent-filing measure to address the fact that more recent patent applications are less likely to have been published and therefore may be undercounted in patent databases. Fourth, we use a shorter quarterly event window around the ChatGPT launch to test whether the result is robust to a higher-frequency design closer to the event. Finally, we replace our continuous treatment variable with an indicator for above-median trade secrecy intensity. Across all five tests, the coefficient on *Trade Secret Intensity* (*High TradeSecrecy Dummy*) $\times$ *Post* remains positive and statistically significant. These results suggest that the main finding is not sensitive to the functional form, treatment or outcome definition, publication-truncation adjustment, or event-window choice.

5 Conclusion

This paper provides some of the first evidence on how GenAI shapes firms' choices between patents and trade secrets for protecting innovation. Using the public launch of ChatGPT as a common shock, we find that firms with greater prior reliance on trade secrecy increase their patent filings more strongly following the launch. This response emerges after the launch and becomes more pronounced over time. It is also stronger when GenAI is more likely to enter firms' innovation activities, when proprietary knowledge is dispersed across more employees and

locations, and when reverse engineering is more prevalent. Together, these patterns are consistent with GenAI making secrecy more difficult to preserve and thereby increasing the relative value of patent protection.

However, firms do not simply disclose more knowledge as they turn toward patenting. The resulting applications are published later, contain less detailed technical disclosures, seek potentially broader claims, and are directed more toward rapidly obsolescing technologies, suggesting that firms continue to manage the proprietary costs of patent disclosure. GenAI may therefore encourage firms that previously relied on secrecy to seek patent rights more often without producing a proportional increase in useful technical disclosure. Our findings highlight an important tradeoff between innovation protection and knowledge diffusion and underscore the need to consider not only how GenAI changes patenting activity, but also what and when the resulting patents disclose.

References

Ahn, Jaehan, Rani Hoitash, Udi Hoitash, and Jaeyoon Yu, 2025, The costs and benefits of an in-house legal function, *Available at SSRN 5375605* .

Alcácer, Juan, and Minyuan Zhao, 2012, Local r&d strategies and multilocation firms: The role of internal linkages, *Management Science* 58, 734–753.

Aobdia, Daniel, 2018, Employee mobility, noncompete agreements, product-market competition, and company disclosure, *Review of Accounting Studies* 23, 296–346.

Armstrong, Chris, Jinhwan Kim, Jalal Sani, and Terrence Tianshuo Shi, 2026, Patent disclosures, examiners, and greenwashing, *Available at SSRN 6994139* .

Arora, Ashish, Sharon Belenzon, and Lia Sheer, 2021, Knowledge spillovers and corporate investment in scientific research, *American Economic Review* 111, 871–898.

Arrow, Kenneth J., 1962, Economic welfare and the allocation of resources for invention, in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–626 (Princeton University Press, Princeton, NJ).

Arundel, Anthony, 2001, The relative effectiveness of patents and secrecy for appropriation, *Research policy* 30, 611–624.

Arundel, Anthony, and Isabelle Kabla, 1998, What percentage of innovations are patented? empirical estimates for european firms, *Research policy* 27, 127–141.

Azinovic-Yang, Li, Yue Chen, Shunsuke Matsuno, and Sunho Yoo, 2025, Strategic patenting under financial disclosure mandates, *Available at SSRN 5767542* .

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson, 2024, Artificial intelligence, firm growth, and product innovation, *Journal of Financial Economics* 151, 103745.

Bena, Jan, Hernan Ortiz-Molina, and Elena Simintzi, 2022, Shielding firm value: Employment protection and process innovation, *Journal of Financial Economics* 146, 637–664.

Bertomeu, Jeremy, Yupeng Lin, Yibin Liu, and Zhenghui Ni, 2025, The impact of generative AI on information processing: Evidence from the ban of ChatGPT in italy, *Journal of Accounting and Economics* 80.

Besiroglu, Tamay, Nicholas Emery-Xu, and Neil Thompson, 2024, Economic impacts of AI-augmented R&D, *Research Policy* 53, 105037.

Blankespoor, Elizabeth, Ed Dehaan, and Qianqian Li, 2026, Generative ai in financial reporting, *Journal of Accounting Research* .

Bombard, Gregory S., and Andrew A.J. Tibbetts, 2025, Reverse engineering in the age of AI: Are your trade secrets still safe?, Greenberg Traurig LLP Insights, By Gregory S. Bombard and Andrew (A.J.) Tibbetts.

Boot, Arnoud, and Vladimir Vladimirov, 2025, Disclosure, patenting, and trade secrecy, *Journal of Accounting Research* 63, 5–56.

Bradshaw, Mark T., Chenyang Ma, Benjamin P. Yost, and Yuan Zou, 2026, Generative AI use by capital market information intermediaries: Evidence from Seeking Alpha, *Journal of Accounting Research* .

Bukhteyev, Alexey, 2025, Beating XLoader at speed: Generative ai as a force multiplier for reverse engineering, Check Point Research, Accessed July 14, 2026.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro HC Sant'Anna, 2024, Difference-in-differences with a continuous treatment, *National Bureau of Economic Research* .

Calvino, Flavio, Jelmer Reijerink, and Lea Samek, 2025, The effects of generative AI on productivity, innovation and entrepreneurship, OECD Artificial Intelligence Papers 39, OECD Publishing, Paris.

centrexIT, 2025, ChatGPT data leakage: How employees accidentally expose proprietary data, centrexIT Blog.

Chang, Anne Yanru, Xi Dong, Xiumin Martin, and Changyun Zhou, 2026, Ai democratization and trading inequality, *Journal of Accounting Research* .

Chen, Wilbur, Terrence Tianshuo Shi, and Suraj Srinivasan, 2024, The value of ai innovations, *Available at SSRN* .

Cohen, Wesley M., Richard R. Nelson, and John P. Walsh, 2000, Protecting their intellectual assets: Appropriability conditions and why U.S. manufacturing firms patent (or not), *National Bureau of Economic Research* .

Crawford, Brian, Justin Phillips, and Patrick McClure, 2026, Automatically attacking software reverse engineering ai agents.

Cuban, Mark, 2026, Post on AI, patent filing, and trade secrets, X post, Accessed July 14, 2026.

Cunningham, Colleen, and Aldona Kapacinskaite, 2026, Keeping invention confidential, *Management Science* 72, 1658–1679.

Cyera, and Cybersecurity Insiders, 2025, State of AI data security report, Technical report, Cyera.

DKSH, 2024, Using AI to formulate and deformulate, DKSH Insights (Podcast), Accessed July 14, 2026.

Dyer, Travis A., Stephen Glaeser, Mark H. Lang, and Caroline Sprecher, 2024, The effect of patent disclosure quality on innovation, *Journal of Accounting and Economics* 77, 101647.

Ecker, Frank, Xitong Li, Yilan Li, and Fan Wu, 2026, How stock market participants use generative artificial intelligence: Evidence from user-platform interaction data, *Journal of Accounting Research* .

Einhorn, Eti, 2007, Voluntary disclosure under uncertainty about the reporting objective, *Journal of Accounting and Economics* 43, 245–274.

Einhorn, Eti, and Amir Ziv, 2008, Intertemporal dynamics of corporate voluntary disclosures, *Journal of Accounting Research* 46, 567–589.

Eisfeldt, Andrea L, Gregor Schubert, Miao Ben Zhang, and Bledi Taska, 2026, Generative ai and firm values, *Journal of Finance forthcoming* .

Friedman, David D., William M. Landes, and Richard A. Posner, 1991, Some economics of trade secret law, *Journal of Economic Perspectives* 5, 61–72.

Galasso, Alberto, and Mark Schankerman, 2015, Patents and cumulative innovation: Causal evidence from the courts, *The Quarterly Journal of Economics* 130, 317–369.

Glaeser, Stephen, 2018, The effects of proprietary information on corporate disclosure and transparency: Evidence from trade secrets, *Journal of Accounting and Economics* 66, 163–193.

Glaeser, Stephen, and Mark Lang, 2024, Measuring innovation and navigating its unique information issues: A review of the accounting literature on innovation, *Journal of Accounting and Economics* 78, 101720.

Glaeser, Stephen, Jeremy Michels, and Robert E. Verrecchia, 2020, Discretionary disclosure and manager horizon: Evidence from patenting, *Review of Accounting Studies* 25, 597–635.

Glaeser, Stephen A, and Wayne R Landsman, 2021, Deterrent disclosure, *The Accounting Review* 96, 291–315.

Gould, Daniel, 2026, The patent time bomb: How AI might rewrite IP strategy, Castle Fund.

Graham, John R, Campbell R Harvey, and Shiva Rajgopal, 2005, The economic implications of corporate financial reporting, *Journal of Accounting and Economics* 40, 3–73.

Hall, Bronwyn H., and Dietmar Harhoff, 2012, Recent research on the economics of patents, *Annual Review of Economics* 4, 541–565.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg, 2001, The NBER patent citation data file: Lessons, insights and methodological tools, NBER Working Paper 8498, National Bureau of Economic Research.

Hrdy, Camilla A., 2025, Trade secrecy meets generative AI, *Chicago-Kent Law Review* 100, 317–357.

IAM, 2026, Linkedin senior patent counsel urges businesses to "get comfortable" with risk in the ai era, https://www.iam-media.com/trade-secrets/article/linkedin-senior-patent-counsel-urges-businesses-get-comfortable-risk-in-the-ai-era, Accessed July 14, 2026.

Intuition Labs, 2024, Prevent ChatGPT proprietary data leaks, Accessed: 2026-05-28.

Jankowski, John E, 2012, Business use of intellectual property protection documented in nsf survey, *NSF info brief, NSF* 12, 307.

Jia, Ning, Ningzhong Li, Guang Ma, and Da Xu, 2025, Corporate responses to generative AI: Early evidence from conference calls, *Review of Accounting Studies* .

Johansen, Frederik Lizak, Ulrik Friis-Jensen, Erik Bjørnager Dam, Kirsten Marie Ørnsbjerg Jensen, Rocío Mercado, and Raghavendra Selvan, 2026, deCIFer: Crystal structure prediction from powder diffraction data using autoregressive language models, *Transactions on Machine Learning Research* .

Kang, Hyo, and Wyatt Lee, 2022, How innovating firms manage knowledge leakage: A natural experiment on the threat of worker departure, *Strategic Management Journal* 43, 1961–1982.

Kim, Ji Ho, Botir Kobilov, Terrence Tianshuo Shi, and Yuan Zou, 2026, When ai distorts investor interpretation: Evidence from forced-balance summaries, *Available at SSRN 6871778* .

Kim, Jinhwan, Terrence Tianshuo Shi, and Rodrigo S. Verdi, 2025, The innovation consequences of judicial efficiency, *Journal of Accounting and Economics* 80, 101813.

Kim, Jinhwan, and Kristen Valentine, 2021, The innovation consequences of mandatory patent disclosures, *Journal of Accounting and Economics* 71, 101381.

Kim, Jinhwan, and Kristen Valentine, 2023, Public firm disclosures and the market for innovation, *Journal of Accounting and Economics* 76, 101577.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *The quarterly journal of economics* 132, 665–712.

Kraft, Kornelius, and Christian Rammer, 2025, Reverse engineering and innovation: empirical evidence from firm-level data, *The Journal of Technology Transfer* 1–54.

Kwa, Thomas, Ben West, Joel Becker, et al., 2025, Measuring ai ability to complete long software tasks, in *Advances in Neural Information Processing Systems*.

LayerX Security, 2025, Enterprise AI and SaaS data security report 2025, Technical report, LayerX Security, Available at layerxsecurity.com.

Lemley, Mark A., 2008, The surprising virtues of treating trade secrets as IP rights, *Stanford Law Review* 61, 311–354.

Levine, David S., 2023, Generative artificial intelligence and trade secrecy, *Journal of Free Speech Law* 3, 559–588.

Li, Yinghua, Yupeng Lin, and Liandong Zhang, 2018, Trade secrets law and corporate disclosure: Causal evidence on the proprietary cost hypothesis, *Journal of Accounting Research* 56, 265–308.

Ma, Song, 2025, Technological obsolescence, *The Review of Financial Studies* hhaf059.

McKinsey & Company, 2025, The next innovation revolution: Powered by AI, McKinsey & Company, By Alex Singla, Alexander Sukharevsky, Elia Berteletti, Lareina Yee, and Michael Chui.

Mezzanotti, Filippo, 2021, Roadblock to innovation: The role of patent litigation in corporate r&d, *Management Science* 67, 7362–7390.

Michael Best & Friedrich, 2026, Rethinking trade secrets vs. patents for software and AI in light of the Anthropic code leak, Michael Best & Friedrich LLP, By Derek C. Stettner.

Morrone, Megan, 2025, Workers are spilling secrets to chatbots, Citing research by Harmonic Security.

National Polymer, 2024, Deformulation chemistry, National Polymer Blog, Accessed July 14, 2026.

Oster, Emily, 2019, Unobservable selection and coefficient stability: Theory and evidence, *Journal of Business & Economic Statistics* 37, 187–204.

Samuelson, Pamela, and Suzanne Scotchmer, 2002, The law and economics of reverse engineering, *Yale Law Journal* 111, 1575–1663.

Shao, Xiang, Sang Wu, and Yue Zhang, 2023, Strategic disclosure in r&d races: Theory and evidence, *Available at SSRN 5895603* .

Singh, Jasjit, 2005, Collaborative networks as determinants of knowledge diffusion patterns, *Management science* 51, 756–770.

Toren, Peter J., 2026, Navigating recent developments in generative AI and trade secret protection, IPWatchdog.

U.S. Patent and Trademark Office, 2024, Request for comments regarding the impact of the proliferation of artificial intelligence on prior art, the knowledge of a person having ordinary skill in the art, and determinations of patentability made in view of the foregoing, Federal Register.

Verrecchia, Robert E, 1983, Discretionary disclosure, *Journal of accounting and economics* 5, 179–194.

Wonderful PCB, 2025, AI-powered PCB reverse engineering: Automated schematic generation, Wonderful PCB Blog, Accessed July 14, 2026.

Zhang, Yue, 2024, Corporate r&d investments following competitors' voluntary disclosures: Evidence from the drug development process, *Journal of Accounting Research* 62, 335–373.

Appendix A. Variable Definitions

Variable	Definition	Source
Dependent Variables		
Patent Filings	The inverse hyperbolic sine (IHS)-transformed number of patent applications filed by a firm in a given year.	PatentsView
Patent Filings_ Dummy	An indicator variable that equals one if a firm filed at least one patent application in a given year, and zero otherwise.	PatentsView
Patent Filings (Core Field)	The IHS-transformed number of patent applications filed by a firm in core technology fields, defined as the firm's top technology subclasses by patent count during 2009–2018, in a given year.	PatentsView
Patent Filings (Non-Core Field)	The IHS-transformed number of patent applications filed by a firm in non-core technology fields in a given year.	PatentsView
Product Patent Filings	The IHS-transformed number of product patent applications filed by a firm in a given year, classified based on product-related language in the first independent claim.	PatentsView
Process Patent Filings	The IHS-transformed number of process patent applications filed by a firm in a given year, classified based on process-related language in the first independent claim.	PatentsView
Patent Filings (Fast Disclosure)	The IHS-transformed number of patent applications filed by a firm in a given year that were published within 18 months after the filing date.	PatentsView
Patent Filings (Slow Disclosure)	The IHS-transformed number of patent applications filed by a firm in a given year that were published 18 months or more after the filing date.	PatentsView
Patent Filings (High Obsolescence)	The IHS-transformed number of patent applications filed by a firm in subclasses with above-median technological obsolescence rates in a given year.	PatentsView; Ma (2025)
Patent Filings (Low Obsolescence)	The IHS-transformed number of patent applications filed by a firm in subclasses with below-median technological obsolescence rates in a given year.	PatentsView; Ma (2025)
Patent Filings (Strict Fields)	The IHS-transformed number of patent applications filed by a firm in strict examiner fields, defined as technology subclasses with above-median patent-application rejection rates among examiners, in a given year.	PatentsView
Patent Filings (Lenient Fields)	The IHS-transformed number of patent applications filed by a firm in lenient examiner fields, defined as technology subclasses with below-median patent-application rejection rates among examiners, in a given year.	PatentsView
Filings-to-Publication Gap	The average number of months between the filing date and the publication date of patent applications filed by a firm in a given year.	PatentsView
Length	The IHS-transformed average number of words in the brief summary section of patent applications filed by a firm in a given year.	PatentsView
Number of Numbers	The IHS-transformed average count of numerical values in the brief summary section of patent applications filed by a firm in a given year.	PatentsView
Number of Figures & Sheets	The IHS-transformed average number of figures and drawing sheets in patent applications filed by a firm in a given year.	PatentsView

Length of Indep. Claims	The IHS-transformed average number of words in the independent claims of patent applications filed by a firm in a given year.	PatentsView

Independent Variables

Trade Secret Intensity	The IHS-transformed ratio of a firm's total number of trade secret references in 10-K from 2019 to 2021 to the firm's R&D stock as of 2019, capturing the firm's ex-ante reliance on secrecy relative to innovation investment.	Glaeser (2018); Compustat
Post	An indicator variable that equals one for the post-ChatGPT period (2023–2025) and zero for the pre-period (2019–2021).	
Size	The natural logarithm of total assets.	Compustat
ROA	Net income scaled by total assets.	Compustat
Loss	An indicator variable that equals one if a firm reports negative net income, and zero otherwise.	Compustat
Leverage	Total debt scaled by total assets.	Compustat
BTM	Book-to-market ratio, measured as total assets divided by the sum of market capitalization and total debt.	Compustat
Cash/Assets	Cash and short-term investments scaled by total assets.	Compustat
Capex/Assets	Capital expenditures scaled by total assets.	Compustat
PPE/Assets	Net property, plant, and equipment scaled by total assets.	Compustat
Sales Growth	Year-over-year growth in net sales.	Compustat
R&D/Assets	Research and development expenditures scaled by total assets.	Compustat
R&D Stock	The IHS-transformed R&D stock computed using a 15% annual depreciation rate applied to historical R&D expenditures.	Compustat
Patent Stock	The IHS-transformed cumulative number of patents granted to a firm over the prior 20 years, with a 15% annual depreciation rate.	PatentsView

Cross-Sectional Variables

Unique Inventors	The number of distinct inventors associated with a firm's patenting activity during 2009–2018. Firms are partitioned into low and high groups based on the sample median of this pre-period measure.	PatentsView
Inventors in Non-HQ States	The number of distinct inventors located outside the firm's headquarters state during 2009–2018. Firms are partitioned into low and high groups based on the sample median of this pre-period measure.	PatentsView; Compustat
Employee Turnover Change	The change in a firm's employee turnover rate from the pre-period (2019–2021) to the post-period (2023–2025). Firms are partitioned into low and high groups based on the sample median of this change measure.	Revelio
Reverse-Engineering Exposure	A SIC3 industry-level measure based on the cross-country median share of firms reporting reverse engineering as a source of innovation in the Eurostat CIS 2018 survey, mapped from NACE Rev. 2 industries to SIC3 industries using text similarity. Firms are partitioned into low and high groups based on the sample median of this pre-event industry-level measure.	Eurostat CIS 2018
GenAI Exposure (Total)	A firm-level total GenAI exposure score following Eisfeldt et al. (2026). Firms are partitioned into low and high groups based on the sample median of this measure.	Eisfeldt et al. (2026)

GenAI Exposure (Core)	A firm-level core-business GenAI exposure score constructed following Eisfeldt et al. (2026). Firms are partitioned into low and high groups based on the sample median of this measure.	Eisfeldt et al. (2026)
GenAI Exposure (Supplemental)	A firm-level non-core business (supplemental) GenAI exposure score constructed following Eisfeldt et al. (2026). Firms are partitioned into low and high groups based on the sample median of this measure.	Eisfeldt et al. (2026)
AI Adoption	The firm's average number of employees with AI-related skills during 2009–2018, measured using AI-worker data from Babina et al. (2024). Firms are partitioned into low and high groups based on the sample median of this measure.	Babina et al. (2024)
IP-Experienced Personnel	The average number of IP-experienced legal employees at the firm during 2009–2018. Firms are partitioned into low and high groups based on the sample median of this pre-period measure. We follow Ahn et al. (2025) to identify legal employees with IP experience.	Revelio
Software Industry	An indicator variable that equals one if a firm is in the software industry (SIC 737), and zero otherwise. This variable is used to partition firms into software and non-software groups.	Compustat

Figure 1. Parallel Trend

This figure plots the dynamic treatment effects and their corresponding 99% confidence intervals over the six-year period, three years before and after the ChatGPT launch year, excluding the launch year. We interact the trade secrecy intensity with a yearly indicator relative to the ChatGPT launch year, except for year 2021, which serves as the reference year. The coefficient plotted on each interaction represents the treatment effect relative to the reference year.

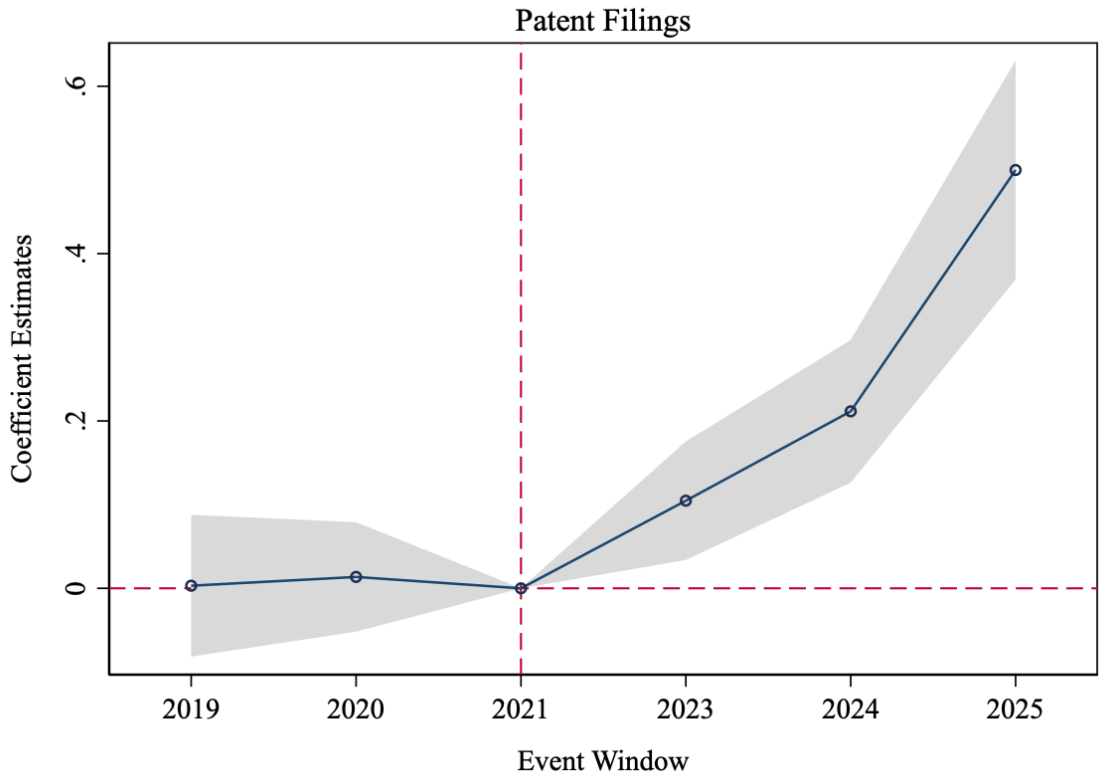


Table 1. Summary Statistics

This table presents the sample composition and summary statistics. Panel A reports the industry distribution of the sample firms. Panel B presents summary statistics for the key variables used in the analysis. Detailed variable definitions are provided in Appendix A.

Panel A. Industry Distribution

Industry	# Firms	% Firms	Mean Trade Secret Intensity
Mining (1000-1499)	15	1.61	0.507
Construction (1500-1799)	1	0.11	1.865
Manufacturing (2000-3999)	742	79.87	0.298
Transportation, Communications, and Utilities (4000-4999)	23	2.48	0.194
Wholesale Trade (5000-5199)	4	0.43	0.088
Retail Trade (5200-5999)	7	0.75	0.779
Finance, Insurance and Real Estate (6000-6799)	11	1.18	0.519
Services (7000-8999)	122	13.13	0.265
Nonclassifiable (9900-9999)	4	0.43	0.002
Total	929	100.00	

Table 1. Summary Statistics (Continued)

Panel B. Summary Statistics

	N	Mean	STD	P25	Median	P75
Patent Filings	5,294	1.316	1.876	0.000	0.000	2.312
Patent Filings_ Dummy	5,294	0.444	0.497	0.000	0.000	1.000
Patent Filings (Core Field)	5,294	0.791	1.412	0.000	0.000	0.881
Patent Filings (Non-Core Field)	5,294	1.087	1.732	0.000	0.000	1.818
Product Patent Filings	5,294	1.273	1.855	0.000	0.000	2.312
Process Patent Filings	5,294	0.347	0.810	0.000	0.000	0.000
Patent Filings (Strict Fields)	5,294	0.708	1.311	0.000	0.000	0.881
Patent Filings (Lenient Fields)	5,294	0.984	1.740	0.000	0.000	1.444
Patent Filings (Fast Disclosure)	5,294	1.167	1.728	0.000	0.000	2.095
Patent Filings (Slow Disclosure)	5,294	0.637	1.404	0.000	0.000	0.000
Patent Filings (High Obsolescence)	5,294	0.905	1.627	0.000	0.000	1.444
Patent Filings (Low Obsolescence)	5,294	0.951	1.561	0.000	0.000	1.444
Filings-to-Publication Gap	5,294	4.368	6.298	0.000	0.000	6.825
Length	5,294	2.863	3.800	0.000	0.000	7.509
Number of Numbers	5,294	1.407	2.267	0.000	0.000	2.998
Number of Figures & Sheets	5,294	1.118	1.551	0.000	0.000	2.873
Length of Indep. Claims	5,294	1.985	2.633	0.000	0.000	5.323
Trade Secret Intensity	5,294	0.293	0.609	0.002	0.040	0.240
Post	5,294	0.474	0.499	0.000	0.000	1.000
Loss	5,294	0.385	0.487	0.000	0.000	1.000
Size	5,294	7.381	2.553	5.533	7.472	9.213
ROA	5,294	-0.076	0.346	-0.087	0.029	0.080
Leverage	5,294	0.270	0.236	0.083	0.242	0.379
BTM	5,294	0.724	0.560	0.323	0.584	0.977
Cash/Assets	5,294	0.265	0.248	0.075	0.173	0.390
Capex/Assets	5,294	0.027	0.028	0.008	0.019	0.035
PPE/Assets	5,294	0.176	0.151	0.070	0.130	0.232
Sales Growth	5,294	0.127	0.711	-0.053	0.034	0.150
R&D/Assets	5,294	0.104	0.167	0.015	0.046	0.116
R&D Stock	5,294	6.247	2.267	4.810	6.296	7.628
Patent Stock	5,294	4.122	2.250	2.475	3.997	5.632

Table 2. Trade-Secret Reliance, GenAI, and Patent Filings

This table reports the DiD results of the effect of the GenAI launch on patent filings. The dependent variable (Patent Filings) is the inverse hyperbolic sine (IHS)-transformed number of patent applications filed by a firm in a given year. Trade Secret Intensity is the inverse hyperbolic sine (IHS)-transformed ratio of a firm's total number of trade secret references in 10-K from 2019 to 2021 to the firm's R&D stock as of 2019. Post is an indicator variable equal to one for the post-period (2023–2025) and zero for the pre-period (2019–2021). Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

	(1)	(2)	(3)
D.V.:		Patent Filings	
Trade Secret Intensity × *Post*	0.221***	0.230***	0.226***
	(0.030)	(0.031)	(0.029)
Trade Secret Intensity	-0.844***	0.178***	
	(0.072)	(0.049)	
Post	-0.458***	-0.461***	
	(0.035)	(0.037)	
Size		0.103***	0.092**
		(0.027)	(0.039)
ROA		-0.472***	-0.121*
		(0.107)	(0.063)
Loss		-0.111	-0.001
		(0.070)	(0.038)
Leverage		-0.231*	-0.023
		(0.132)	(0.097)
BTM		0.111*	-0.016
		(0.057)	(0.034)
Cash/Assets		0.422***	0.116
		(0.154)	(0.114)
Capex/Assets		2.541*	-0.038
		(1.482)	(0.701)
PPE/Assets		0.322	-0.167
		(0.281)	(0.253)
Sales Growth		-0.016	-0.015
		(0.018)	(0.013)
R&D/Assets		-0.378*	-0.024
		(0.221)	(0.129)
R&D Stock		0.092***	0.081
		(0.033)	(0.058)
Patent Stock		0.479***	-0.050
		(0.026)	(0.060)
Observations	5,294	5,294	5,294
R-squared	0.070	0.507	0.891
Firm FE	No	No	Yes
Year FE	No	No	Yes
Controls	No	Yes	Yes

Table 3. GenAI Exposure and Patent Filings

This table reports subsample analyses based on firm exposure to and implementation of artificial intelligence. In Panel A, the sample is partitioned based on the extent to which GenAI can help perform the tasks carried out by its workforce, measured by total, core-business, and non-core-business AI exposure following Eisfeldt et al. (2026). In Panel B, the sample is partitioned based on a firm's AI implementation using the pre-period number of employees with AI-related skills. Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Panel A. GenAI Task Exposure

	GenAI Exposure		Core Exposure		Supplemental Exposure	
D.V.: Patent Filings	(1) Low	(2) High	(3) Low	(4) High	(5) Low	(6) High
Trade Secret Intensity × *Post*	0.143***	0.390***	0.146***	0.373***	0.199***	0.277***
	(0.033)	(0.061)	(0.034)	(0.059)	(0.036)	(0.053)
Diff. (High - Low)	0.247***		0.227***		0.078	
	(0.069)		(0.068)		(0.064)	
Observations	2,644	2,650	2,645	2,649	2,644	2,650
R-squared	0.908	0.876	0.908	0.876	0.902	0.879
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes	Yes	Yes

Panel B. AI Workforce Adoption

	AI Adoption		
D.V.: Patent Filings	(1) Low		(2) High
Trade Secret Intensity × *Post*	0.124***		0.422***
	(0.030)		(0.100)
Diff. (High - Low)		0.298***	
		(0.104)	
Observations	2,683		2,611
R-squared	0.829		0.896
Firm FE	Yes		Yes
Year FE	Yes		Yes
Controls	Yes		Yes

Table 4. Information Leakage Risks and Patent Filings

This table presents subsample analyses based on labor-related information leakage risks. Panel A partitions the sample using pre-period inventor measures constructed from patenting activity during 2009–2018. In columns (1)–(2), the sample is partitioned based on the number of unique inventors. In columns (3)–(4), the sample is partitioned based on the number of inventors located outside the headquarters states. Panel B partitions the sample based on the change in employee turnover from the pre-period to the post-period. In columns (1)–(2), the partition is based on turnover among all employees, while in columns (3)–(4), it is based on turnover among scientists and engineers. Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Panel A. Unique Inventors

	Unique Inventors		Inventors in Non-HQ States	
D.V.: Patent Filings	(1) Low	(2) High	(3) Low	(4) High
Trade Secret Intensity × *Post*	0.067**	0.572***	0.122***	0.413***
	(0.026)	(0.134)	(0.033)	(0.077)
Diff. (High - Low)	0.505***		0.291***	
	(0.137)		(0.084)	
Observations	2,664	2,630	2,668	2,626
R-squared	0.739	0.886	0.888	0.886
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes

Panel B. Employee Turnover Change

	All Workers		S&E Workers	
D.V.: Patent Filings	(1) Low	(2) High	(3) Low	(4) High
Trade Secret Intensity × *Post*	0.154***	0.351***	0.170***	0.285***
	(0.035)	(0.053)	(0.038)	(0.047)
Diff. (High - Low)	0.197***		0.115*	
	(0.063)		(0.060)	
Observations	2,897	2,397	2,647	2,647
R-squared	0.897	0.885	0.899	0.882
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Controls	Yes	Yes	Yes	Yes

Table 5. Exposure to Reverse-Engineering Risk and Patent Filings

This table reports how the treatment effect varies with ex ante reverse-engineering risk and firms' ex post patenting choices. Panel A presents subsample analyses based on ex ante reverse-engineering risk. Reverse-engineering risk is measured using the reverse-engineering prevalence, constructed from the 2018 Eurostat Community Innovation Survey based on the proportion of firms reporting reverse engineering as an information source. Panel B reports the treatment effect on firms' choice between product and process patents following Bena et al. (2022). Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Panel A. Ex-ante Reverse-Engineering Exposure

	Reverse-Engineering Prevalence		
D.V.: Patent Filings	(1) Low		(2) High
Trade Secret Intensity × *Post*	0.201***		0.358***
	(0.033)		(0.067)
Diff. (High - Low)		0.157**	
		(0.075)	
Observations	3,231		2,063
R-squared	0.882		0.898
Firm FE	Yes		Yes
Year FE	Yes		Yes
Controls	Yes		Yes

Panel B. Ex-post Product vs Process Patent Choice

	(1)	(2)
D.V.:	Process Patent Filings	Product Patent Filings
Trade Secret Intensity × *Post*	-0.019	0.011**
	(0.012)	(0.006)
Observations	5,294	5,294
R-squared	0.831	0.995
Firm FE	Yes	Yes
Year FE	Yes	Yes
Controls	Yes	Yes

Table 6. Strategic Patent Disclosure and Knowledge Diffusion

This table reports results on firms' information controls in patent filings. Panel A reports results on disclosure timing using Patent Filings (Fast Disclosure), Patent Filings (Slow Disclosure), and Filings-to-Publication Gap. Panel B reports results on information content using Length, Number of Numbers, and Number of Figures & Sheets. Panel C reports results on patent filings in technology fields with below- and above-median levels of technological obsolescence, measured following Ma (2025). Panel D reports results on claim scope using the length of independent claims. Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Panel A. Disclosure Timing (Speed of Patent Publication)

	(1)	(2)	(3)
D.V.:	Patent Filings (Fast Disclosure)	Patent Filings (Slow Disclosure)	Filings-to-Publication Gap
Trade Secret Intensity × *Post*	-0.064***	0.232***	0.773***
	(0.009)	(0.027)	(0.120)
Observations	5,294	5,294	5,294
R-squared	0.979	0.828	0.722
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Panel B. Information Content Quality in Patent Filings

	(1)	(2)	(3)
D.V.:	Length	Number of Numbers	Number of Figures & Sheets
Trade Secret Intensity × *Post*	-0.227***	-0.102**	-0.079***
	(0.073)	(0.044)	(0.028)
Observations	5,294	5,294	5,294
R-squared	0.790	0.739	0.784
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Controls	Yes	Yes	Yes

Table 6. Strategic Patent Disclosure and Knowledge Diffusion (Continued)

Panel C. Technological Obsolescence of Disclosed Fields

	(1)	(2)
D.V.:	Patent Filings (Low Obsolescence)	Patent Filings (High Obsolescence)
Trade Secret Intensity × *Post*	0.006	0.047***
	(0.013)	(0.013)
Observations	5,294	5,294
R-squared	0.963	0.959
Firm FE	Yes	Yes
Year FE	Yes	Yes
Controls	Yes	Yes

Panel D. The Scope of Patent Rights

	(1)
D.V.:	Length of Indep. Claims
Trade Secret Intensity × *Post*	-0.154***
	(0.048)
Observations	5,294
R-squared	0.796
Firm FE	Yes
Year FE	Yes
Controls	Yes

Table 7. Internal Legal Resources and Patent Filings

This table presents subsample analyses based on a firm's internal legal resources. The sample is partitioned based on the availability of IP-experienced legal personnel. Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

	IP-Experienced Personnel		
D.V.: Patent Filings	(1) Low		(2) High
Trade Secret Intensity × *Post*	0.116***		0.354***
	(0.034)		(0.081)
Diff. (High - Low)		0.239***	
		(0.088)	
Observations	2,656		2,638
R-squared	0.861		0.894
Firm FE	Yes		Yes
Year FE	Yes		Yes
Controls	Yes		Yes

Table 8. Additional Analyses

This table reports how the treatment effect varies across different aspects of patent portfolio choice. Panel A reports results on patent filings in core and non-core technology fields. Core (non-core) technology fields are defined based on whether the technology subclass is among the firm's top patenting subclasses during 2009–2018. Panel B reports results on patent filings in strict and lenient examiner fields. Strict (lenient) examiner fields are defined based on whether the patent-application rejection rate in a technology subclass is above (below) the sample median. Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Panel A. Core vs. Non-Core Technology Fields

	(1)	(2)
D.V.:	Patent Filings (Core Field)	Patent Filings (Non-Core Field)
Trade Secret Intensity × *Post*	0.017	0.030**
	(0.013)	(0.012)
Observations	5,294	5,294
R-squared	0.939	0.976
Firm FE	Yes	Yes
Year FE	Yes	Yes
Controls	Yes	Yes

Panel B. Strict vs. Lenient Examiner Fields

	(1)	(2)
D.V.:	Patent Filings (Strict Fields)	Patent Filings (Lenient Fields)
Trade Secret Intensity × *Post*	0.006	0.037***
	(0.013)	(0.013)
Observations	5,294	5,294
R-squared	0.926	0.967
Firm FE	Yes	Yes
Year FE	Yes	Yes
Controls	Yes	Yes

Table 9. Robustness Tests

This table reports robustness tests. Panel A compares software-oriented and non-software-oriented firms using SIC-based partitions. Panel B reports results using alternative specifications, including Poisson estimation in column (1), extensive margin on patent filings in column (2), publication-rate adjustment for truncation following Hall et al. (2001) in column (3), a shorter quarterly event window using 6 quarters before and after the launch of ChatGPT in column (4), and an indicator variable for above-median trade secret intensity as the treatment variable in column (5). Standard errors, clustered at the firm level, are reported below the coefficient estimates in parentheses. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% levels, respectively (two-tailed). Detailed definitions of all variables are provided in Appendix A.

Panel A. Software Industry vs. Non-Software Industry

	Software Industry		
D.V.:	(1) Software		(2) Non-Software
Trade Secret Intensity × *Post*	0.126		0.220***
	(0.179)		(0.029)
Diff. (Software - Non-Software)		-0.094	
		(0.178)	
Observations	552		4,742
R-squared	0.858		0.898
Firm FE	Yes		Yes
Year FE	Yes		Yes
Controls	Yes		Yes

Panel B. Alternative Specifications

D.V.: Patent Filings	(1)	(2)	(3)	(4)	(5)
Model:	Poisson	Extensive Margin	Truncation-Adj.	Short Window	High TradeSecrecy Dummy
Trade Secret Intensity × *Post*	0.409**	0.039***	0.174***	0.083***	
	(0.185)	(0.012)	(0.029)	(0.015)	
High TradeSecrecy Dummy × *Post*					0.311***
					(0.058)
Observations	3,397	5,294	5,294	11,050	5,294
R-squared		0.719	0.895	0.907	0.891
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	No	Yes
Year-Quarter FE	No	No	No	Yes	No
Controls	Yes	Yes	Yes	Yes	Yes