Discussion Paper Series – CRC TR 224

Discussion Paper No. 119
Project B 02

Visibility of Technology and Cumulative Innovation: Evidence From Trade Secrets Laws

Bernhard Ganglmair*
Imke Reimers**

September 2024
(First version: August 2019)
(Second version: May 2021)

*University of Mannheim and ZEW Mannheim, E-mail: b.ganglmair@gmail.com
**Cornell University, E-mail: imke.reimers@cornell.edu

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Visibility of Technology and Cumulative Innovation: Evidence from Trade Secrets Laws*

Bernhard Ganglmair[†] Imke Reimers[‡]

September 23, 2024

Abstract

Innovation depends on the incentives to create new ideas as well as access to existing ones. Access, in turn, depends on an initial invention's inherent visibility as well as the inventor's decision to disclose it, for example through patenting. Compared to product innovations, process innovations tend to be less visible, making secrecy as an intellectual property strategy relatively more attractive; and stronger trade secrets protection laws could hamper their disclosure even further. Using exogenous variation in the level of trade secrets protection from the Uniform Trade Secrets Act, we show that stronger protection has a disproportionately negative effect on patenting of processes. We show in simulations that these changes in disclosure can have large implications for follow-on innovation and welfare.

Keywords: cumulative innovation; disclosure; intellectual property; Uniform Trade Secrets Act; visibility.

JEL Codes: D80; O31; O34

---

*We have benefited from conversations with participants at various conferences and seminars, as well as discussions with Kevin Bryan, James Dana, Francesc Dilmé, Fabian Gaessler, Daniel Gross, Deepak Hegde, Christian Helmers, Paul Hünermund, Jeffrey Kuhn, Matt Marx, Mike Meurer, Markus Nagler, Lisa Larrimore Ouellette, Ivan Png, Albert Roger, David Schwartz, Neel Sukhatme, and Mike Teodorescu. Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B02) is gratefully acknowledged. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]University of Mannheim and ZEW Mannheim. E-mail: b.ganglmair@gmail.com.

[‡]Cornell University. E-mail: imke.reimers@cornell.edu.

1 Introduction

When better protection of intellectual property improves the appropriability of R&D investment returns, firms have stronger incentives to invest and innovate. The fruits of such innovation serve as the proverbial shoulders on which future innovators can stand, thus fostering technological progress through more follow-on (or cumulative) innovation.[1] However, granting the inventor a temporary monopoly through a patent can have negative, "anticommons" effects on follow-on innovation when exclusivity renders the shoulders less accessible (Heller and Eisenberg, 1998). A negative effect on follow-on innovation also arises when inventors decide to disclose fewer of their inventions through patents and instead keep them secret. Nondisclosure is especially harmful to the diffusion of knowledge when the technologies are less visible per se, as might be the case with process innovations. In these cases, diffusion and follow-on innovation would be hampered if inventors kept more secrets, and even more so when legal trade secrets protection is strong. We study these effects of intellectual property policy on patenting of inventions featuring varying levels of visibility, and we highlight their implications for cumulative innovation.

Secrecy is an important tool in a firm's intellectual property management toolbox. There is ample survey-based evidence that secrecy is widely used and often more important as an appropriability mechanism than patents (e.g., Levin et al., 1987; Cohen et al., 2000; Arundel, 2001). Mansfield (1986) reports survey results suggesting that one out of three patentable inventions is kept secret when inventors have a choice between secrecy and patenting. Importantly, choosing secrecy does not mean that the invention is without any protection. The laws governing trade secrets offer protection against *misappropriation* of secrets—that is, the acquisition of a trade secret by *improper means* or the disclosure of a trade secret without consent.[2] For example, in a well-publicized legal case, Waymo LLC (a self-driving car startup

[1] In 1675, Sir Issac Newton wrote in a letter to Robert Hooke: "If I have seen further, it is by standing upon the shoulders of giants." See Scotchmer (1991) for the economics of giants' shoulders.

[2] Generally speaking, a trade secret is information (e.g., a customer list, a business plan, or a manufacturing process) that has some commercial value the secret holder wants to conceal from others (Friedman et al., 1991). We use the terms "secrecy" and "trade secrets" interchangeably.

under Google's Alphabet) accused Uber Technologies of violating both California state and federal trade secret laws, alleging that a former employee secretly downloaded data around a key piece of technology from its servers before resigning and launching a self-driving truck startup.[3] However, while trade secrets laws provide protection against such misappropriation, unlike patents they do not grant general exclusivity: A trade secret is not protected if it accidentally leaks or is uncovered through independent discovery or reverse engineering (Friedman et al., 1991); and others' ability to reverse engineer an invention depends on its visibility.

Stronger protection of trade secrets renders them more attractive relative to patents. In this paper, we ask how a change in the attractiveness of secrecy affects the diffusion of knowledge through (1) the decision to invest in different types of innovation, (2) the disclosure of these inventions, and (3) the ability to build on them. Our main empirical exercise studies the disclosure decision. We use exogenous variation in trade secrets protection across states and time brought about by the Uniform Trade Secrets Act (UTSA) of 1979/1985 to identify the effects of trade secrets protection on patenting of product vs. process inventions. Using an index of the strength of trade secrets protection (Png, 2017b) and data on the type of a patented invention (Ganglmair et al., 2022), we show that stronger trade secrets protection results in a disproportionate decrease of process patents (i.e., inventions that are less visible). Our estimated effects are largest among individual inventors and small firms and are driven by patents covering discrete rather than complex technologies. These results are robust to different modeling choices as well as the inclusion of controls for other changes in IP enforcement. Since patents provide insight into what is *not* kept secret, we interpret our findings as a relative increase in the propensity to keep process inventions secret.[4] This, by itself, limits opportunities for follow-on innovation.

[3] The startup was later acquired by Uber. See Waymo LLC v. Uber Technologies, Inc; Ottomotto LLC; Otto Trucking LLC. No. 3:17-cv-00939, N.D. Cal., San Francisco Division. The case settled in February 2018, five days into the trial.

[4] Our assumption of the choice between secrecy and patents, as opposed to joint use (Arora, 1997), comes without loss of generality as long as there is *some* degree of substitutability.

The welfare implications of such changes in intellectual property protection depend not only on the facilitation of follow-on innovation from existing inventions but also on the ex-ante incentives to innovate; and changes in initial innovation could in turn affect follow-on innovation. Because we do not observe inventions that are kept secret, we cannot estimate the effects of trade secrets protection on the total stock of innovations.[5] Instead, we simulate a simple framework of sequential innovation to make inferences about these broader effects. While stronger trade secrets lower disclosure of inventions, thus hampering opportunities for follow-on innovation, they can also raise incentives for initial innovation and thus provide additional opportunities. Both the negative and the positive effects are amplified for less visible technologies such as processes. If the incentives to innovate rise substantially, then this rise could outweigh the negative disclosure effect, especially in industries that are characterized by little sequential innovation. An ideal trade secrets policy therefore weighs the decrease in disclosure that we document against possible increases in initial innovation.

Beyond studies based on survey data, there is limited empirical work in economics on trade secrets, though a small literature presents indirect evidence on secrecy. Moser (2012) documents a shift toward patenting (and away from secrecy) in the chemical industry as the publication of the periodic table of elements has facilitated reverse engineering. Gross (2023) finds that a policy during World War II to keep certain patent applications secret resulted in slower dissemination of the patented technologies into product markets; de Rassenfosse et al. (2024) confirm these findings. Hegde and Luo (2018) show that a reduction of the duration of temporary secrecy of patent applications had a mitigating effect on licensing delays, and Hegde et al. (2023) find an acceleration in diffusion of knowledge and ideas.

A related strand of literature studies the effect of changes in legal trade secrets protection on innovation and patenting behavior. Png (2017a,b) finds that stronger trade secrets protection has a positive effect on firms' investment in R&D and renders patenting relatively less attractive. Related to this line of work, Contigiani et al. (2018) find that more employer-

[5] With the rare exception of Cunningham and Kapačinskaitė (2022), who provide empirical evidence of actual secrecy use, papers in the literature resort to indirect tests of the effects of trade secrets.

friendly trade secrets protection has a dampening effect on innovation, and Castellaneta et al. (2017) show a positive effect on firm value in industries with high mobility of skilled labor. Angenendt (2018) finds that patent applicants respond to stronger trade secrets protection by reducing the number of patent claims.

We add to these bodies of literature by examining the role of an invention's visibility in patenting and innovation decisions; and we explicitly study how these decisions affect the provision of opportunities for follow-on innovation, finding that optimal policies depend on the types of innovation that prevail in an industry. We further highlight that insights gained from the effects of patents on innovation do not necessarily apply to trade secrets.

The remainder of this paper is structured as follows: In Section 2, present a simple theoretical framework for the interaction of trade secrets protection, the visibility of an invention, and an inventor's patenting decision. We also provide institutional background on the Uniform Trade Secrets Act. In Section 3, we present our estimation sample that we use for our empirical analysis in Section 4. In Section 5, we use simulations of a model of sequential innovation to discuss the implications of trade secrets protection for follow-on innovation building on processes and products. We conclude with policy recommendations in Section 6.

2 Trade Secrets Protection and Innovation

Our empirical strategy leverages changes in the level of trade secret protection in the United States following the adoption of the Uniform Trade Secrets Act. In this section, we introduce a theoretical framework to lay out the effects of trade secrets protection on innovation, and we then provide some background information on the Act.

2.1 Theoretical Framework of Sequential Innovation

To guide our discussion of the effects of trade secrets protection on innovation, we consider a simple theoretical framework that captures three stages of an invention's life cycle: (1) the decision to invest in initial R&D, (2) the decision whether to patent the invention or to keep it secret, and (3) the resulting opportunities for follow-on innovation. All three stages depend both on the strength of trade secrets protection and the visibility of the invention.

In Stage 1, inventors have stronger incentives to invest in R&D if they believe their inventions will reap larger returns (including the value from the invention itself and licensing revenues from potential follow-on innovation) than their investment costs. The use of secrecy as an appropriability mechanism suggests that stronger trade secrets protection should weakly raise the expected returns to R&D. Png (2017b) provides evidence in line with this *appropriability effect*, showing that stronger trade secrets protection results in higher overall R&D investment. This effect should be driven by those inventions that benefit the most from secrecy and for which patenting is no profitable alternative. Because secrecy is a particularly viable option for production processes and business methods that are inherently difficult to observe, stronger trade secrets protection may raise incentives to invest in R&D for such process innovation disproportionately. This would lead to a rise in the underlying set of process inventions relative to product inventions.

In Stage 2, the inventor decides whether to patent and thus (partly) disclose the invention, or to keep the invention as a trade secret.[6] If the inventor patents the invention, they are granted monopoly rights for a limited period; and they can benefit from additional licensing fees.[7] On the other hand, trade secrets have no term limits but provide weaker opportunities for licensing and only limited protection against misappropriation. Still, stronger protection raises the attractiveness of trade secrets relative to patenting. While previous work shows

[6]We provide notation and formal proofs of this disclosure decision in the Appendix.

[7]While monopoly rights allow for larger profits to the inventor, monopoly market structures are also connected to inefficiencies in terms of total welfare. In Section 5, we account for this deadweight-loss effect stemming from limited access to the technology for other market participants

an overall decrease in patenting (e.g., Png, 2017a), these effects likely vary across invention types, for two reasons: First, for low-visibility inventions (such as new production processes), the information costs of disclosing the details in a patent are higher. Second, once disclosed, the enforcement of such a patent is more difficult because of higher monitoring costs (e.g., Goldstein, 2013; Hall et al., 2014). We therefore predict that stronger trade secrets protection disproportionately induces process innovators to choose secrecy over patenting, leading to a drop in the share of newly filed process patents. Testing this is the main objective of the empirical analysis below.

The investment and disclosure decisions affect follow-on innovation, which we describe as Stage 3. With stronger trade secrets, the appropriability effect likely increases the stock of initial inventions to build on (a positive *cumulation effect*, see Scotchmer, 1991), but the decision to keep these—and some other—inventions secret could render them invisible (a negative *disclosure effect*). Both effects on follow-on innovation are amplified for process innovations, which are inherently less visible than product innovations.

In what follows, we first examine the relative effects of trade secrets protection on patenting of processes, where we predict a disproportionately negative effect on the disclosure of processes. We return to the issue of how stronger trade secrets protection affects the opportunities of follow-on innovation in Section 5.

2.2 Uniform Trade Secrets Act (1979/1985)

The UTSA was published and recommended to the individual U.S. states for adoption in 1979 (revised in 1985) by the National Conference of Commissions on Uniform State Laws. Between 1979 and 2018, all U.S. states except New York and North Carolina have adopted the UTSA, with adoption dates ranging from 1981 (5 states) to 2018 (Massachusetts).[8]

The objective of the UTSA was to clarify and harmonize across U.S. states the legal protection of trade secrets. Most prominently, it attempted to standardize the definition

[8] In New York, State Assembly and Senate bills were introduced in 2024.

of a trade secret, the meaning of misappropriation, and remedies (including damages) for trade secret holders in case of a violation. For example, with the adoption of the UTSA, the Commonwealth of Virginia dropped the requirement of actual or intended use for something to be considered a trade secret and increased the punitive damages multiplier from 0.5 to 2.

Png (2017b) constructs an annual index that quantifies the strength of legal trade secrets protection at the state level on a scale from 0 to 1, for the years 1976 to 2008. For example, the changes in Virginia represent increases in two of the six inputs into the index.[9] On average, the UTSA implied a rise in the index of 42 percentage points across states (median = 46.7). In most states, the UTSA resulted in a strengthening of trade secrets protection, with the exception of Arkansas and Pennsylvania, where pre-UTSA trade secrets protection (under common law) was stronger, and Michigan, Nebraska, and Wyoming, where the protection index remained unchanged. There is no obvious pattern in the size of these changes over time and across states, and Png (2017b) cites anecdotal evidence that suggests that passing of the bills by state legislatures often happened for "whimsical" reasons.

2.3 Trade Secrets and the Disclosure Function of Patents

By using the UTSA to examine the effects of trade secrets protection on access to knowledge and follow-on innovation, we make two implicit assumptions. First, changes in the level of trade secrets protection affect firms' use and defense of trade secrets. Second, patents provide some disclosure of inventions.

The first premise, that the changes in trade secrets protection through the UTSA were sufficient to induce changes in behavior, is supported by empirical evidence (e.g., Png, 2017a,b; Castellaneta et al., 2017; Leblebicioglu and Savaser, forthcoming). Moreover, Almeling (2012) attributes part of the rise of trade secrets litigation over recent decades to the individual states' adoption of the UTSA, mostly because it raised awareness of the option

[9]In addition to these two factors, the index is higher (i) without a requirement that the trade secret holder exert reasonable effort to protect the secret, (ii) without a requirement that the information be used or disclosed, (iii) without a statute of limitation, and (iv) with unlimited length of injunction.

to keep trade secrets.

The second premise is that patents provide some disclosure of inventions. Some legal scholars have called the disclosure function of patents into question. For example, Ouellette (2012) argues that patents have lowered the level of openness in science. While acknowledging that "patents disclose useful, nonduplicative technical information" (p. 533), she notes they "could be even more informative." Others share these concerns (e.g., Roin, 2005; Fromer, 2009). In addition, Lemley (2008a) suggests that scientists and engineers might not pay attention to patents for strategic reasons, a phenomenon observed more often in complex than in discrete technologies (Bessen and Meurer, 2009).

Nevertheless, both legal and economics scholars seem to agree that patents provide *some* information. Exploiting variation across fields, Merges (1988) finds that many inventors rely on published patents for technical information. Recent work by economists also finds that innovators use existing patents for inspiration and information (e.g., Furman et al., 2021; Gross, 2023; Hegde et al., 2023). Importantly, our results hold as long as patents provide *some* disclosure.

3 Patent Data

In our empirical analysis, we examine how stronger trade secrets protection (measured by way of the state-level trade secrets protection index) affects patenting decisions. Our underlying data include patents with priority dates between 1976 and 2008—the years for which we have index data. We (a) match these patents to the relevant level of trade secrets protection by approximating the timing and location of the patenting decision, and we (b) determine each patent's type (process or product) based on the language used in its claims. Our analysis requires that we make decisions on both dimensions. Here we describe our choices for the main analysis, noting that our results are robust to these decisions. We supplement these data with additional patent characteristics.

3.1 Timing of the Disclosure Decision and Patent Location

To determine the timing of the disclosure (patenting) decision, we use the earliest priority date of the respective granted patent. This date reflects the application date of the first patent in a patent family—that is, the *parent application*, which applies to all its subsequent continuation and divisional applications.[10] The relevant disclosure decision was likely made at the time of the parent application, so that we use that application's priority date as the disclosure date for all related patents.

For the patent's location, we consider only patents for which all U.S. inventors and U.S. assignees are from the same state, and we use that state as the patent's location. Our approach ensures that the patent applicant's decision was driven by only that state's level of trade secrets protection, and not contaminated by laws in other states.[11] With our assumption of single-state patents, we limit our overall sample to 1,451,311 patents (out of 2,433,317 patents by U.S. applicants, and 4,370,594 total), granted between 1976 and 2014 and with priority dates between 1976 and 2008.[12]

3.2 Indicators for Process and Product Patents

We construct our indicators of process and product patents using information at the level of the patent's independent claims from Ganglmair et al. (2022).[13] A claim can be of one of three distinct types: (1) process (or method) claims describe the sequence of steps which

[10] For a continuation application, the applicant may not add new disclosures but may use claims from the parent application that have not been issued (or abandanced). Divisional applications involve splitting a parent application into two or more.

[11] An identifying assumption, which is supported by Paolino v. Channel Home Centers, 668 F.2d 721 724 n.2 (3d Cir. 1982), is that trade secrets protection is determined by the state where the secret was developed and not where it was misappropriated: "The law of the state of residence of the person who initially developed and protected the secret appears to be the obvious starting point for its protection."

[12] Our estimation sample slightly over-represents individual applicants and under-represents large firms. We document this selection in Section A.2.

[13] A patent claim describes what the applicant claims to be the invention for which the patent grants exclusive legal rights. Each patent can hold multiple claims of different types. An *independent* claim stands on its own whereas a *dependent* claim is in reference to an independent claim. Ganglmair et al. (2022) use information from both the preamble of the claim (that names what the invention is about) and the body of the claim (that lists steps of a process or the elements of a product) for their text-based categorization of patent claims.

Table 1: Summary Statistics

	N	Mean	Median	SD	Min	Max
Process patent	1,451,311	0.473	0	0.499	0	1
Number of process claims	1,451,311	0.871	0	1.407	0	60
Number of product claims	1,451,311	1.920	2	1.885	0	104
Product-by-process claims	1,451,311	0.042	0	0.288	0	30
Independent claims	1,451,311	2.883	2	2.286	1	116
Length of first claim (words)	1,451,311	169.194	148	106.034	1	7078
Length of description (chars.)	1,451,311	25992.144	15658	39439.832	4	3,608,036
Generality	1,096,154	0.638	0.719	0.244	0	1
Originality	1,276,719	0.626	0.694	0.244	0	1
4th year renewal	1,358,663	0.826	1	0.380	0	1
Observations	1,451,311					

Notes: This table provides summary statistics for all granted utility patents (between 1976 and 2014) with priority dates between 1976 and 2008 for which all U.S. inventors and assignees are from the same state.

together complete a task such as making an article; (2) product-by-process claims define a product through the process employed in the making of a product; and (3) product claims describe an invention in the form of a physical apparatus, a system, or a device.

We classify a patent as a *process patent* if at least one of its independent claims is either a process claim or a product-by-process claim, and as a *product patent* otherwise. We choose this rather aggressive definition because we are interested in whether any process-related aspects of an invention are disclosed at all, regardless of the disclosure of its product-related aspects.[14]

The top portion of Table 1 provides summary statistics for our patent-type indicators for all granted USPTO utility patents in our sample.[15] Almost half of all patents include a process claim, although that number increased steadily over the time period of our study, from just under 30% in the 1970s to almost 60% in the 2000s. This trend is universal across patent classes. In our empirical analysis, we examine whether the UTSA caused these trends to differ across states.

[14] We treat product-by-process claims as process claims because what they disclose is a process more than a product. Dropping patents with such claims leaves our results unchanged.

[15] For our final sample, we follow Strandburg (2004), who argues that business methods are highly visible "self-disclosing processes," and drop all business method patents (Lerner, 2006).

3.3 Additional Variables

We collect and construct additional patent characteristics to capture the complexity and value of the patented technology. The bottom of Table 1 summarizes these variables across all patents in our main sample. We proxy for a patent's breadth and complexity using the number of independent claims (see Lerner, 1994; Lanjouw and Schankerman, 2004) and the length (in words) of the first claim where shorter claims are likely broader (see Kuhn and Thompson, 2019). As an additional measure of a patent's complexity, we include the length of the patent's description text.

To capture the external value (or technological impact) of a patent, we construct measures of *patent generality* and *patent originality* as proposed by Trajtenberg et al. (1997). Patent generality captures the diversity of patents—measured by their respective patent classes—in which a given patent is (forward)-cited. A higher generality score implies more widespread impacts (Hall et al., 2001). Patent originality, on the other hand, captures the diversity of technologies from which a given patent (backward)-cites. A higher originality score means that the patented invention is combining ideas from different areas to create something new (or "original"). We construct these measures for each patent using the first USPC main class listed on the patent.[16] As a measure of a patent's internal or private value, we use information on whether the patent holder paid the patent maintenance fee during the 4^{th} year of the patent term (see, e.g., Pakes, 1986; Schankerman and Pakes, 1986).

4 Empirical Estimation and Results

Following the discussion of our theoretical framework in Section 2, we expect stronger trade secrets protection to lower inventors' incentives to patent, especially for process innovations. In this section, we empirically test two predictions. First, we test at aggregated levels whether rises in trade secrets protection lead to larger decreases in the patenting of processes than in

[16] There are about 450 main classes in the United States Patent Classification (USPC) system.

product innovations. Second, we examine at the patent level whether the relative decreases in patenting of process innovations are driven by certain types of applicants and technologies. In both sets of analyses, we take advantage of the staggered adoption of the UTSA across U.S. states to estimate the number of process and product patents.

4.1 The Number of Patents

We begin our empirical analysis with regressions at the state-year level. Our first set of regression equations are of the following form:

$$y_{st} = \beta_1 \, protection_{st} + \gamma_s + \mu_t + \epsilon_{st}.$$

Here, y_{st} denotes the number of patents of a type in state s in year t, where types can be all, process, or product patents. Our explanatory variable of interest, $protection_{st}$, is the trade secrets protection index ranging from zero to one described above, and γ_s and μ_t denote state and priority-year fixed effects, respectively. While these regressions capture overall patenting across states, they cannot capture changes in innovation and patenting across technologies. To account for this, our main focus is on regressions at the USPC class-state-year level. At this level, our regressions take on the form

$$y_{cst} = \beta_1 protection_{st} + \gamma_{cs} + \mu_{ct} + \epsilon_{cst}, \tag{1}$$

where the dependent variable y_{cst} describes functions of the number of patents of each type in USPC class c in state s and year t. The interacted fixed effects γ_{cs} for each USPC class and state, and μ_{ct} for each USPC main class and year, control for trends within a USPC class that are common to all states as well as for any USPC class-state-specific characteristics that do not vary over time.[17] Equation (1) thus ensures that the coefficient of interest β_1

[17] The fixed effects including the priority year control for nationwide policy changes such as the *Uruguay Round Agreements Act* of 1995 (extending the maximum validity of a patent to 20 years from filing) and the *American Inventors Protection Act* of 1999 (introducing pre-grant publication of patent applications).

captures within-technology class changes in patenting.

Columns 1 through 3 of Table 2 show coefficients from regressions at the state-year level. We estimate statistically and economically large decreases in patenting overall (column 1) as well as across both patent types (columns 2 and 3). The drop in process patenting outsizes the decrease in product patenting: a change in the trade secrets protection index from zero to one is associated with an almost 600-unit decrease in the number of process patents, compared to only 180 fewer product patents. At mean rises in the trade secrets protection index of 42 percent from UTSA adoption across states, this implies the UTSA led to about 250 fewer process patents and 75 fewer product patents for an average state. Note, however, that these results could be driven by corresponding shifts in invention technologies. Columns 4 through 6 report results from the USPC class-state-year level regressions from our preferred specification in Equation (1), which is designed to capture the effects of trade secrets protection on patenting within technology classes. Again, we find statistically significant decreases in patenting of both invention types, and the decrease is about three times as large for process patents as for product patents.[18]

4.1.1 Patenting within Technology Classes

The results above suggest large differences in the effects of trade secrets protection on patenting of processes vs. products, but they do not take into account initial patenting levels of each type; and these levels can vary across both states and technology classes. To facilitate interpretation, we estimate models with different functional forms in Table 3. The first three columns present results from log-OLS specifications, where the dependent variables are the zero-inflated natural logs of the numbers of patents in each USPC class, state and year. Column 1 shows that the total number of patents would decrease significantly, by 5.5

We cannot add state-year fixed effects because they are perfectly collinear with our explanatory variable of interest.

[18]Columns 1 and 4 confirm the results in Png (2017a). The point estimates are different because our analysis is at the state and technology class level rather than the firm level.

Table 2: Impact of Trade Secrets Protection on Patent Counts

	state-year level			USPC-state-year level		
	(1) total	(2) process	(3) product	(4) total	(5) process	(6) product
Trade secrets index	-778.186**	-597.642**	-180.544*	-2.463***	-1.822***	-0.641***
	(376.952)	(282.577)	(103.813)	(0.328)	(0.230)	(0.113)
$\overline{R^2}$	0.76	0.62	0.90	0.59	0.52	0.68
N	1,683	1,683	1,683	408,718	408,718	408,718

Notes: OLS regressions at the state-year level (columns 1 through 3) or USPC class-state-year level (columns 4 through 6). The dependent variable is the number of patents of the respective type (process, product, or all). The specifications in columns 1 through 3 include state and year fixed effects. Columns 4–6 include USPC class-state and USPC class-year fixed effects. Robust standard errors, clustered at the state level (Columns 1 through 3) and at the state-year level (Columns 4 through 6) in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

percent, with a full-point rise in the trade secrets protection index,[19] or about 2.3 percent at the mean UTSA-induced rise in trade secrets protection. Columns 2 and 3 show that this decrease is largely driven by process patents, which decrease by 11.9 percent, compared to only a 1.4 percent decrease among product patents.[20] Column 4 uses the share of all patents with a process claim as the dependent variable. It confirms the results from the first three columns, showing a decrease in the share of process patents of about one percentage point, or three percent of the pre-UTSA average process share of 32.7 percent.

Columns 5 through 7 of Table 3 report results from Poisson estimations to address potential issues stemming from the low patent counts in some technology classes and states. The results are in line with previous results: following strengthening of trade secrets protection, patent counts decrease across the board, and especially so for process patents.

4.1.2 Timing and Threats to Identification

The fixed effects regressions presented above could capture differences in trends leading up to a state's adoption of the UTSA. We address this concern by estimating the effects of UTSA

[19] This comes from $(e^{-0.057} - 1) \times 100$ percent $= 5.5$ percent.

[20] While we use the zero-inflated log patent counts to hold the number of observations fixed, we note that the results are qualitatively identical and quantitatively even more pronounced.

Table 3: Effects of Trade Secrets Protection within Technology Classes

	OLS				Poisson		
	(1) ln(total)	(2) ln(process)	(3) ln(product)	(4) Pr(process)	(5) ln(total)	(6) ln(process)	(7) ln(product)
Trade secrets index	-0.057***	-0.127***	-0.014	-0.010*	-0.200***	-0.265***	-0.078**
	(0.019)	(0.017)	(0.015)	(0.005)	(0.044)	(0.080)	(0.038)
$\overline{R^2}$	0.73	0.69	0.68	0.45			
N	408,718	408,718	408,718	246,745	408,875	374,159	397,191

Notes: Regressions at the USPC class-state-year level. Columns 1–4 present OLS regression results. In columns 1–3, the dependent variable is the zero-inflated natural log of the number of patents of the respective type (process, product, or all). In column 4, it is the share among all patents that includes a process innovation. Columns 5–7 show estimates from Poisson estimations of the the counts of process, product, and total patents, respectively. All estimations include USPC class-state fixed effects. The OLS regressions also include USPC class-year fixed effects; the Poisson regressions include year fixed effects. Robust standard errors in parentheses, in Columns 1 through 4, standard errors are clustered at the state-year level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

adoption in an event study setting. We allow for a flexible time structure by estimating annual changes in patenting in a USPC class-state pair relative to the year in which the state adopted the UTSA. Specifically, we estimate

$$y_{cst} = \sum_z \beta_z (UTSA)_s \times \mathbf{1}\{z\} + \gamma_{cs} + \mu_{ct} + \varepsilon_{cst}, \tag{2}$$

where the fixed effects γ_{cs} and μ_{ct} are defined as above; $(UTSA)_s$ equals one for states which adopted the UTSA and for which adoption implied an increase in trade secrets protection; and z represents the number of years since the state's UTSA adoption. To control for potential issues arising from heterogeneous treatment effects across cohorts (see, e.g., Goodman-Bacon, 2021), we estimate this regression using the estimator developed by Sun and Abraham (2021).

Figure 1 plots the coefficients of interest for two dependent variables, using the year before UTSA adoption as the omitted category. Panel (a) plots the overall effects of UTSA adoption on the zero-inflated log number of total patents, thus emulating column 1 of Table 3. It indicates, first, that patenting did not change systematically in states shortly before UTSA adoption, and second, that patenting decreased steadily and significantly after adoption,

Figure 1: Event Study Plots

(a) ln(total patents) **(b) Share of process patents**

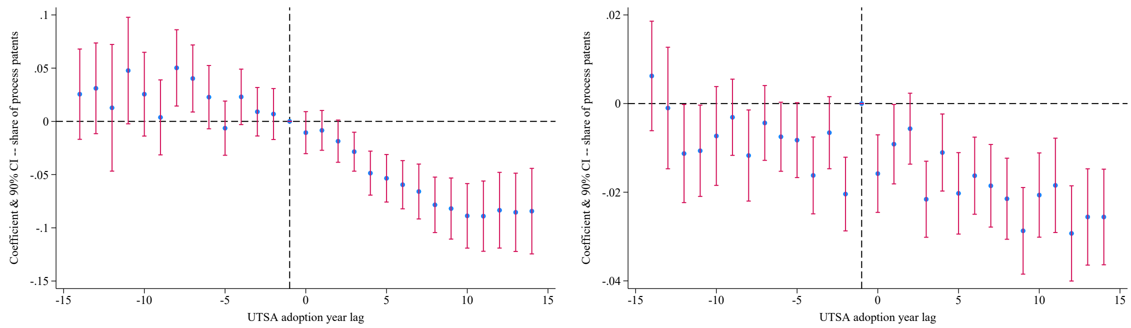

Notes: Point estimates and 90 percent confidence intervals from event study regressions for 15 years around a state's UTSA adoption, adjusting for the possibility of heterogeneous treatment effects (Sun and Abraham, 2021). The dependent variable is the natural log of patents (panel (a)) and the share of process patents (panel (b)). The omitted category is the year before a state's UTSA adoption, and the estimations are done at the USPC class-state-year level.

relative to patenting in other states. Panel (b) highlights the differences in effects between process and product patenting. Similar to column 4 of Table 3, it plots differences in the share of process patents in a state and USPC class before and after UTSA adoption. While the coefficients vary more from year to year, the overall trends support the main findings as well: with the exception of low process shares two years before UTSA adoption, there is no significant difference in the process patent share leading up to UTSA adoption, and the share decreases slightly after adoption.

4.2 Patent-level Analysis

To examine whether the relative decreases in process patents are driven by certain firms or technologies, we move to analyses at the patent level. We ask whether patents become more or less likely to cover a process invention when trade secrets protection changes. A benefit of this analysis is that we can easily examine the roles of additional patent characteristics such as their complexity and value, as well as technology and applicant characteristics. Formally, the regressions in this section are of the form:

$$process_{j,cst} = \beta_1\, protection_{st} + \beta_2 X_j + \gamma_{cs} + \mu_{ct} + \epsilon_j, \tag{3}$$

where the dependent variable is an indicator that is 1 if patent j belonging to USPC main class c and filed in state s and year t is a process patent; $protection_{st}$ and the interacted fixed effects γ_{cs} and μ_{ct} are defined as above. Our specification at the patent level is equivalent to an analysis at the state level where the states are weighted by the number of patents, but it also allows us to directly control for patent-specific measures of complexity and value, X_j. We estimate Equation (3) as a linear probability model, noting that logit estimations provide qualitatively identical results.

4.2.1 Overall Effects

Table 4 shows the coefficients from this specification, including different sets of control variables. Consistent with the aggregated analyses above, all specifications show a statistically significant, negative effect of a UTSA-related strengthening of trade secrets protection on the probability that a patent is a process patent. We are most interested in the specifications that include control variables on both patent complexity and value measures (Columns 4 and 5). Column 4, which includes separate fixed effects for USPC main class, state, and year, suggests that patents are 2.6 percentage points less likely to include a process innovation if the trade secrets protection index rises by a full point. Column 5 interacts the USPC main class dummies with the state and year dummies and, therefore, controls for state- and time-specific variation across technologies. It reports a 1.8 percentage point decrease. At a baseline process patent share across patents of 42.3 percent before UTSA adoption, and with a mean increase in the trade secrets protection index across patents of 0.36 points across all patents, our results correspond to respective mean decreases of 2.2 percent and 1.5 percent in the probability that a patent is a process patent when a state adopts the UTSA.

4.2.2 Effects Across Applicant and Invention Types

We examine the roles of firm sizes and technology complexities in Table 5. In the first two columns, we repeat the estimation from Columns 4 and 5 of Table 4, interacting the

Table 4: Patent-Level Impact of Trade Secrets Protection

	(1)	(2)	(3)	(4)	(5)
Trade secrets protection	-0.018***	-0.021***	-0.026***	-0.026***	-0.018***
	(0.006)	(0.006)	(0.007)	(0.007)	(0.006)
Log(indep. claims)		0.233***		0.231***	0.228***
		(0.002)		(0.002)	(0.002)
Log(length of first claim)		-0.044***		-0.052***	-0.045***
		(0.001)		(0.001)	(0.001)
Log(length of description)		-0.002**		0.001	0.004***
		(0.001)		(0.001)	(0.001)
Originality			0.027***	0.011***	0.013***
			(0.003)	(0.003)	(0.003)
Generality			0.062***	0.039***	0.031***
			(0.003)	(0.003)	(0.003)
4th year renewal			0.045***	0.025***	0.022***
			(0.002)	(0.001)	(0.001)
$\overline{R^2}$	0.297	0.342	0.288	0.335	0.357
N	1,451,307	1,451,307	894,956	894,956	892,296

Notes: Linear probability models at the patent level with 1[process patent] as the dependent variable, and the index of trade secrets protection as the independent variable of interest. Additional controls in columns (1)–(4) include indicator variables for the patent's location state, priority year, and USPC main class. Column (5) interacts USPC main class dummies with both state and year indicators. Robust standard errors, clustered by state and year, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

trade secrets index with indicators for small applicants (individuals and small firms) as well as large applicants (large firms). The estimated decrease in the probability that a patent is a process patent is largest for small applicants. In our preferred specification with interacted fixed effects (Column 2), the estimated coefficient (-0.022, se=0.006) corresponds to an average decrease in the probability of a process patent of 2.4 percent, which exceeds the estimated average effect of 1.5 percent. The negative coefficient is smaller and less statistically significant for large applicants.

Our findings confirm our expectations for three reasons. First, trade secrets are more important as a means to protect intellectual property for small firms than large firms (Hall et al., 2014). Second, individual states are only a small part of a large firm's overall market, and the adoption of the UTSA in just one of these states may not have a strong impact

Table 5: Impact of Trade Secrets Protection by Applicant Size and Technology Type

	Applicant size		Technology type	
	(1)	(2)	(3)	(4)
Trade secrets protection				
… × Small applicant	-0.036***	-0.022***		
	(0.007)	(0.006)		
… × Large applicant	-0.014*	-0.013*		
	(0.008)	(0.007)		
… × Discrete			-0.064***	-0.038***
			(0.010)	(0.008)
… × Complex			-0.008	-0.007
			(0.008)	(0.006)
$\overline{R}^2$	0.336	0.358	0.334	0.356
N	892,620	889,933	855,654	852,923

Notes: Linear probability models at the patent level with 1[process patent] as the dependent variable, and interactions of the trade secrets protection index with applicant size (columns 1 and 2) and with technology type (columns 3 and 4) as the independent variables of interest. All specifications include our sets of complexity and value controls. Columns 1 and 3 include state, year and USPC class fixed effects. Columns 2 and 4 include fixed effects for interactions of USPC class with states and with years, respectively. Robust standard errors, clustered by state and year, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

on patenting decisions. Third, findings by Crass et al. (2019) suggest a stronger degree of substitutability between secrecy and patents for small applicants.

In the next two columns of Table 5, we allow the effects to vary between "complex" and "discrete" technologies to explore the issue of substitutability between patenting and trade secrets more directly. Complex technologies (such as in electrical engineering, telecommunications, semiconductors, or machine tools) are more likely protected by a combination of patents and trade secrets, whereas discrete technologies (such as in chemicals, pharmaceuticals, or materials) are more likely to rely on just one IP strategy. Thus, the effect of stronger trade secrets protection should be most pronounced among discrete technologies. To test this, we assign a complexity indicator to each patent based on Graevenitz et al. (2013).[21] Interacting this indicator with the trade secrets index in our main specification (Column 4 of Table 5), we find that the probability of a process patent decreases by 2.7% at the baseline

[21] In our data, 73% of patents represent complex technologies.

(coef=0.038, se=0.008) among discrete technologies, whereas the effect is very small and statistically insignificant among complex technologies (coef=-0.007, se=0.006).

4.2.3 Robustness Analyses

Our data construction and empirical approach are based on a number of assumptions. We examine the robustness of our empirical results to these assumptions in additional regressions, replicating the specification from Column 5 of Table 4. Table 6 summarizes the results. In short, all specifications show a robust negative impact of trade secrets protection on the share of process patents.

In Panel (a), we vary our choices of the decision date as well as the invention location. The first two columns examine the timing of the decision. In Column 1, we use each individual patent's application date rather than the priority date. Column 2 circumvents the timing issue altogether by considering only the patent family head—the first patent within its family. The next three columns explore different choices for identifying the location of the decision maker. In a less conservative approach, in Column 3, we use *all patents* and assign the first assignee's state as the location of the disclosure decision, or the location of the first inventor if no U.S. assignee is listed. In a more conservative approach, in Column 4, we drop all patents with non-U.S. contributors. Finally, Column 5 uses only patents with a single decision maker. All approaches provide almost identical results to the main specification.

In Panel (b) of Table 6, we show that our results are also robust to process patent definitions and additional control variables. Columns 1 and 2 consider two alternative definitions of process patents, based on the first claim and on the weak majority of claims, respectively.[22] In Column 3, we drop software patents because these tend to be categorized as processes even though they do not inherently include process innovation.[23] Our results across these specifications are similar in magnitude to our preferred specification, although Column 1

[22] Kuhn and Thompson (2019) argue that under U.S. law the broadest claim is listed first.

[23] We follow Graham and Vishnubhakat (2013) in identifying patents as software patents. In our data, 66% of all software patents include a process claim, as opposed to 40% of non-software patents.

Table 6: Robustness Checks

Panel (a): Disclosure Date and Invention Location

	(1) Appl. Date	(2) Family Head	(3) Assignee Loc	(4) U.S. Only	(5) 1 Assignee
Trade secrets protection	-0.020***	-0.020***	-0.020***	-0.018**	-0.017***
	(0.006)	(0.005)	(0.005)	(0.007)	(0.006)
Observations	878512	796373	1435763	616992	849881
$\overline{R^2}$	0.357	0.364	0.356	0.344	0.358

Panel (b): Process Patent Definition and Control Variables

	(1) Process: 1st	(2) Process: Most	(3) No Software	(4) IDD	(5) Pre-Trends
Trade secrets protection	-0.008*	-0.015***	-0.017***	-0.017***	-0.029**
	(0.005)	(0.005)	(0.006)	(0.005)	(0.013)
Observations	886436	892296	652023	892296	892296
$\overline{R^2}$	0.331	0.279	0.335	0.357	0.357

Notes: Linear probability model with 1[process patent] as the dependent variable. In Panel (a): Column 1 sets the date of the disclosure decision as the patent's application date; Column 2 uses only the first patent in a patent family (the family head); Column 3 uses the location of the first assignee (or the first inventor if no assignee is listed); Column 4 is limited to patents for which all contributors are American and from the same state; and Column 5 drops patents with more than one assignee. In Panel (b), Columns 1–3 examine the definition of process patents. Column 1 uses the status of the patent's first claim; Column 2 considers a patent a process patent if at least half of its claims describe a process; Column 3 drops all software patents; Column 4 adds state-year specific control variables for changes (strengthening and weakening) in the court enforcement of the Inevitable Disclosure Doctrine (IDD); Column 5 adds state-specific linear pre-trends. Robust standard errors, clustered at the state and year, in parentheses. All specifications include the same control variables as the full specification in the main text.

shows a smaller and less statistically significant effect.

In the last two columns, we add additional state-level controls. First, we control for changes in the enforcement of the inevitable disclosure doctrine (IDD). The IDD is a common law doctrine that allows employers to seek protection against the disclosure of trade secrets by a former employee when working for a competitor.[24] Column 4 adds IDD enforcement changes from the case coding in Castellaneta et al. (2016:Table B1) as a control variable. Column 5 instead includes state-specific linear time trends before UTSA adoption to account

[24]For instance, in 1998 the Utah District Court found (in favor of the IDD) it "inevitable that defendants will traffic upon Novell's trade secrets and confidential technical information unless they are restrained from being in the same business Novell is in." Novell, Inc. v. Timpanogos Research Group, Inc., 46 U.S.P.Q.2d 1197 (Utah Dist. Ct. 1998).

for possible time-varying differences across states. The results from these specifications are unchanged or even larger.

4.3 Discussion of Identification

Our identification strategies rely on two assumptions. First, stronger trade secrets protection did not change the relative pool of process and product inventions that were at risk of being patented. Second, the adoption of the UTSA is not affected by an expectation that certain types of inventions will be more prevalent in the future. The first assumption is sensible: if stronger trade secrets protection incentivizes innovation, then this innovation is likely kept as a trade secret.[25] The second assumption requires additional analysis.

Although Png (2017b) provides evidence of the exogeneity of the UTSA with regard to firms' decisions to invest in R&D, one might be concerned that each state's decision to adopt the UTSA was motivated by changes in innovation and patenting behavior, rather than the other way around. With the caveat that patents are the results of investments made in the past, this would imply a change in the likelihood that a patent covers a process invention *before* a state adopts the UTSA.[26] We provide evidence in Figure 1 suggesting little pre-trends; and our results at the patent level are robust to the addition of linear time state-level trends as well. Still, we further examine the possibility of endogenous UTSA adoption in a set of placebo tests. Instead of the *true* UTSA adoption date for each state, we set an earlier date, dropping all patents with priority dates after the true UTSA adoption to avoid confounding our placebo effects with true ones.[27]

We then estimate the effect of placebo UTSA adoption—one, two, three and four years before the true adoption—on the probability that a patent is a process patent, in regressions that mirror Column 5 of Table 4. For all four placebo adoption dates, the coefficient of

[25]The pool could also change if firms and inventors move to states with stronger trade secrets protection. As shown by Png (2012), however, the adoption of the UTSA had no significant effect on inventors' mobility.

[26]Png (2017a) suggests an instrumental variables approach, which we adopt in an additional check; and we find even stronger results in that specification.

[27]We also drop all patents whose application date lies more than ten years before the state's true UTSA adoption to improve balance between treated and control groups.

interest is small and statistically insignificant, ranging from -0.002 (se=0.004) for placebo adoption two years earlier to +0.003 (se=0.004), four years earlier. These results suggest that states adopted the UTSA exogenously with respect to changes in the distribution of product and process patents.

5 Effects on Cumulative Innovation

While previous research has shown that stronger trade secrets protection raises overall R&D investment (Png, 2017b) but lowers technology spillovers (Wang, 2023) and overall patenting (Png, 2017a), the separate effects on product and process innovations are less well understood. In Section 4, we show that stronger trade secrets lead to a disproportionate decrease in process patents. As we highlight in Section 2, this could have profound implications on follow-on innovation. Here, we examine these implications more closely.

5.1 A Framework of Cumulative Innovation

To examine the trade-off between initial R&D and follow-on innovation and to illustrate how the tradeoff depends on the visibility of the initial invention, we use our framework of sequential innovation introduced in Section 2. We simulate the consequences of changes in trade secrets protection with simple assumptions on the inventions' visibilities, their access, and R&D costs.

In Stage 1 (initial R&D), we let an inventor observe a *potential* invention of a given type (process or product), its visibility, value, and the costs of development. The visibility follows an invention type-specific distribution with a higher average visibility for products than processes. The inventor undertakes the R&D project if the expected payoffs from the invention (as a function of the level of trade secrets protection) outweigh its cost. At Stage 2 (disclosure), the inventor decides whether to patent and thus disclose the realized invention (subject to patent protection) or keep it a secret (subject to trade secrets protection). At

Stage 3, follow-on innovation is a representative invention with random commercial value and R&D costs.[28]

For follow-on innovation to materialize, the value of the follow-on invention must exceed its costs (the invention must be *profitable*), and the initial invention must be accessible. The probability that the initial invention is accessible—and hence the probability that follow-on innovation is successful—depends on its inherent visibility and the disclosure decision. If the initial invention is *not patented*, then the access probability is determined solely by its visibility (from distributions that are bounded between 0 and 1). If it *is* patented, we mechanically set the access probability to 2/3, reflecting a combination of disclosure (making the invention highly visible) and potential anticommons effects (restricting access).

As discussed in Section 2, the level of trade secrets protection has three distinct effects on innovation in this model. First, trade secrets protection positively affects initial R&D by increasing the expected value of realized R&D projects (positive *appropriability effect*). Second, because more initial inventions are kept secret at Stage 2 when trade secrets protection is stronger, effective disclosure decreases, reducing the success probability of follow-on innovation (negative *disclosure effect*). Third, the increase in ex-ante R&D activity (from the appropriability effect) implies there is more initial innovation that can generate follow-on innovation, even if it is kept secret (positive *cumulation effect*).

5.2 Harm to Lower-Visibility Follow-On Innovation

Our goal is to simulate both the effects of stronger trade secrets protection on innovation of products and processes as well as their welfare implications. To do so, we need to make several assumptions. We describe both the details of the simulations and the underlying assumptions in Appendix Section A.1. The main takeaway from the simulation supports the results in Section 4: more realized initial inventions through stronger trade secrets protection

[28] The value of this representative invention can be interpreted as capturing the present discounted value of a stream of follow-on innovation (akin to quality-ladder models in Grossman and Helpman (1991), O'Donoghue et al. (1998), or Hopenhayn and Mitchell (2001), or the outcome of applied research building on basic research (Scotchmer, 2004:135ff) triggered by the initial invention.

(i.e., the appropriability effect) may come at the cost of harming follow-on innovation (via the disclosure effect), and this tradeoff is amplified for process innovations.

Figure 2 presents the results of our simulations. In panel (a), we plot the shares of innovations that are ultimately patented (Stage 2 decisions) for products (orange) and processes (blue). We distinguish between the share of all inventions that are patented (dash-dotted lines), and the patenting share of those inventions that are realized without trade secrets protection (baseline as solid lines). Because the newly incentivized inventions are always kept secret, the patenting propensity of the combined set of inventions is lower. The larger difference between the solid and dash-dotted lines for processes is a consequence of a stronger initial R&D effect for processes than products resulting from the relative attractiveness of trade secrets for low-visibility inventions. Finally, because stronger trade secrets protection disproportionately dampens the filing of process patents, but does not change the overall number of inventions that are at risk of being patented, the share of patents that include process innovations decreases when trade secrets protection increases, as the green graph in panel (a) shows, confirming our results in Section 4.

Our empirical results in Section 4 highlight the potential of the negative *disclosure* effect on follow-on innovation, but they cannot capture the positive *cumulation* effect from increased (undisclosed but nonetheless visible) initial R&D. In panel (b) of Figure 2, we disentangle the two effects using our simulated data. We plot the percentage change in follow-on innovation for products (orange) and processes (blue) in response to an increase in trade secrets protection. Consistent with panel (a), the solid lines depict the change in follow-on innovation for those initial inventions that are realized regardless of trade secrets protection changes (baseline). The decrease in follow-on innovation as protection increases is the negative disclosure effect. The dash-dotted lines depict follow-on innovation that builds on newly incentivized inventions, capturing the cumulation effect. Like in panel (a), both effects are stronger for processes.

In these simulations, trade secrets protection is more detrimental to follow-on innova-

Figure 2: Initial R&D and Follow-On Innovation

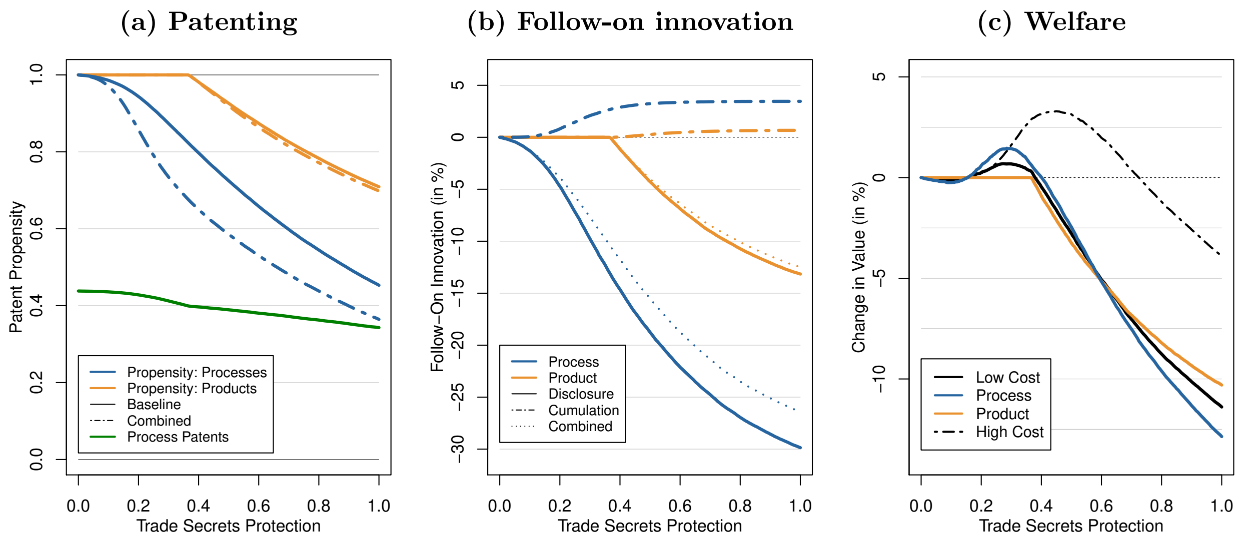

Notes: This figure illustrates the results from our simulations of the three-stage cumulative innovation model (with equal proportions of processes and products for the unrealized ideas). Panel (a) plots the patent propensities (i.e., the share of patented Stage-1 inventions) for products (orange) and processes (blue) as well as the share of process patents (green); panel (b) plots percentage changes in follow-on innovation (relative to $\tau = 0$); panel (c) plots the percentage change in value generated from initial and follow-on innovation (in % of the value when $\tau = 0$). In panels (a) and (b), we assume low average R&D costs (30% of the expected R&D project value); in panel (c), the solid lines depict values for low R&D costs, whereas the dash-dotted line depicts values for higher R&D costs (45% of the expected R&D project value). Visibilities are uniformly distributed with support $[0, 0.75]$ for processes and $[0.25, 1]$ for products.

tion from processes than from products (see the dotted lines in panel (b)). Trade secrets protection can encourage development of processes (appropriability effect), but too strong a protection can retard follow-on innovation if it leads to excessive nondisclosure unless the rise in initial innovation far outweighs the disclosure effect. Our empirical estimates in Section 4 suggest UTSA-related decreases in process patenting of about 5.0 percent and in product patenting of about 0.6 percent. Whether rises in initial innovation exceed these levels remains an open question.

5.3 Total Welfare Effects of Trade Secrets Protection

Finally, we examine how the three different channels add up and determine overall welfare through the value of initial R&D and follow-on innovation. We calculate "welfare" as the expected social surplus derived from the use of the invention, accounting for the inventions' values (and costs) and potential welfare losses from barriers to access to new technologies.

For each invention at Stage 1, the *potential surplus* is the value of the gains from trade without any barriers to access; and the *realized surplus* is the potential surplus net of the disclosure-state specific deadweight loss. For patented inventions, barriers to access (and thus the deadweight loss) increase in visibility as patents (and the ensuing exclusivity) are easier to enforce. For inventions kept as trade secrets, the barriers to access decrease in visibility (implying more exclusivity) and increase in trade secrets protection. For follow-on innovation, we assume free access and zero deadweight loss in both disclosure states.[29]

We begin with a baseline feature of our framework (not depicted): For no R&D costs, stronger trade secrets protection has an unambiguously negative effect on total welfare. Because of zero R&D costs, all initial inventions are realized for any value of trade secrets protection. Stronger trade secrets protection, however, introduces additional barriers to access to a technology, which increases the deadweight loss from monopoly power. The latter results in a decrease in the value associated with initial innovation as trade secrets protection increases. In addition, stronger trade secrets protection reinforces the disclosure effect, thus lowering follow-on innovation.

Panel (c) of Figure 2 illustrates the effects of trade secrets protection for non-zero R&D costs. The positive appropriability effect increases the social value associated with the initial R&D (albeit partially offset by the deadweight-loss effect). The negative disclosure effect lowers social value associated with follow-on innovation but is (partially) offset by the positive cumulation effect when more initial R&D is available to build on. The solid graph depicts the overall welfare effects (for low R&D costs): the percentage change in the value associated with both initial R&D and follow-on innovation. We observe positive effects for lower levels of trade secrets protection and negative welfare effects for higher levels.

We also plot the values separately for processes (blue) and products (orange) and show that processes benefit more from trade secrets protection than products when protection levels are low; but trade secrets protection becomes more detrimental to processes when it

[29]We provide more details on our welfare measure in Appendix Section A.1.

is too strong.

The dash-dotted graph in panel (c) depicts overall welfare for high R&D costs. Observe from the locations of the maxima of the graphs for low and high R&D costs that the optimal level of trade secrets protection increases in costs. This finding rationalizes existing law and practice, providing stronger protection for higher-cost projects. For example, in the State of New York (which has not adopted the UTSA but follows common law principles), one factor in determining whether something is a trade secret explicitly lists the costs of developing the information.[30] Moreover, under the UTSA, trade secrets holders must also show significant costs of duplication of the secret information to establish the validity of their case, for example, by referring to their own costs of R&D (Sandeen and Rowe, 2013:34).

6 Conclusion

The effects of intellectual property rights (IPR) on incentives to innovate are relatively well-studied (Williams, 2017). Still, we know less about the differences between the effects of IPR on initial and follow-on innovation or about the role of an invention's visibility. We add to this ongoing debate by arguing that the effects of IPR on follow-on innovation depend on both the details of the enforcement of IPR and the visibility of the original idea. For highly visible inventions, patents limit the ability of others to build on and develop further; for inventions whose technology is less visible, however, trade secrets limit access entirely. For such low-visibility inventions, patents incentivize inventors to disclose information, boosting the potential for follow-on innovation. Therefore, an intellectual property policy that particularly encourages the patenting of less visible inventions to lift the veil of secrecy could increase innovative activity as a whole.

We show that while stronger trade secrets protection is expected to increase incentives to invest in (initial) R&D (Png, 2017b), it lowers the incentives to patent (and disclose) the

[30] Restatement (First) of Torts, §757 cmt. b (1939). Despite the adoption of the UTSA and the publication of the Restatement (Third) of Unfair Competition (also governing aspects of trade secrets protection), courts and commentators in many states continue to cite this Restatement of Torts.

outcome of that R&D process and thus hampers the opportunities for follow-on innovation, especially for low-visibility inventions. The welfare implications of a policy of stronger protection are, therefore, far from clear. Because we do not observe inventions that are kept secret, we cannot estimate the effects of trade secrets protection on the total stock of innovations. Instead, we simulate a simple framework of sequential innovation to make inferences about these broader effects. Both the negative and the positive effects are amplified for less visible technologies such as processes. If the incentives to invest in R&D rise substantially from an increase in trade secrets protection, then this rise could outweigh the negative disclosure effect, especially in industries that are characterized by little sequential innovation. An ideal trade secrets policy weighs the decrease in disclosure that we document against possible increases in initial innovation.

Our results support arguments in favor of an optimal policy that accounts for different types of inventions and industries. We show that IP policy reforms (such as stronger trade secrets), intended to boost firms' incentives to invest in R&D, can dampen the diffusion of process knowledge that is not easily accessible to third parties without its disclosure in patents. Moreover, our simulations indicate that industries with high R&D costs (e.g., pharmaceuticals and chemicals) are most likely to benefit from increased trade secrets protection. By contrast, industries with relatively low R&D costs are likely to experience a welfare loss from stronger trade secrets protection.

Note that we specifically study the choice of secrecy for patentable inventions. A different, and broader, approach to trade secrets relates to contractual solutions (such as in non-disclosure agreements) and the design of the employment relationship (e.g., in the form of covenants-not-to-compete or noncompete agreements/NCA). Recent efforts by the Federal Trade Commission to ban covenants-not-to-compete and the ongoing debate on the differential effect of such a ban on innovation (Conti, 2014; Johnson et al., 2023) highlight the need for additional evidence. Given the mechanisms in our paper, we view our results as complementary to that literature.

References

Almeling, David S. (2012) "Seven Reasons Why Trade Secrets are Increasingly Important," *Berkeley Technology Law Journal*, 27, 1091–1117.

Angenendt, David T. (2018) "Easy to Keep, But Hard to Find: How Patentable Inventions are Being Kept Secret," Unpublished manuscript, University of Bologna.

Arora, Ashish (1997) "Patents, Licensing, and Market Structure in Chemicals," *Research Policy*, 26 (4-5), 391–403.

Arundel, Anthony (2001) "The Relative Effectiveness of Patents and Secrecy for Appropriation," *Research Policy*, 30 (4), 611–624.

Bessen, James and Michael J. Meurer (2009) *Patent Failure*, Princeton, N.J.: Princeton University Press.

Castellaneta, Francesco, Raffaele Conti, and Aleksandra Kacperczyk (2017) "Money Secrets: How Does Trade Secret Legal Protection Affect Firm Market Value? Evidence from the Uniform Trade Secret Act," *Strategic Management Journal*, 38 (4), 834–853.

Castellaneta, Francesco, Raffaele Conti, Francisco M. Veloso, and Carlos A. Kemeny (2016) "The Effect of Trade Secret Legal Protection on Venture Capital Investments: Evidence from the Inevitable Disclosure Doctrine," *Journal of Business Venturing*, 31 (5), 524–541.

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)," NBER Working Paper 7552, National Bureau of Economic Research, Cambridge, Mass.

Conti, Raffaele (2014) "Do Non-Competition Agreements Lead Firms to Pursue Risky R&D Projects?," *Strategic Management Journal*, 35 (8), 1230–1248.

Contigiani, Andrea, David H. Hsu, and Iwan Barankay (2018) "Trade Secrets and Innovation: Evidence from the 'Inevitable Disclosure' Doctrine," *Strategic Management Journal*, 39 (11), 2921–2942.

Crass, Dirk, Francisco Garcia-Valero, Francesco Pitton, and Christian Rammer (2019) "Protecting Innovation Through Patents and Trade Secrets: Evidence for Firms with a Single Innovation," *International Journal of the Economics of Business*, 26 (1), 117–156.

Cunningham, Colleen M. and Aldona Kapačinskaitė (2022) "Keeping Invention Confidential," unpublished manuscript, University of Utah.

de Rassenfosse, Gaétan, Gabriele Pellegrino, and Emilio Raiteri (2024) "Do Patents Enable Disclosure? Evidence from the Invention Secrecy Act," *International Journal of Industrial Organization*, 92, 103044.

Friedman, David D., William M. Landes, and Richard A. Posner (1991) "Some Economics of Trade Secret Laws," *Journal of Economic Perspectives*, 5 (1), 61–72.

Fromer, Jeanne C. (2009) "Patent Disclosure," *Iowa Law Review*, 94, 539–606.

Furman, Jeffrey L., Markus Nagler, and Martin Watzinger (2021) "Disclosure and Subsequent Innovation: Evidence from the Patent Depository Library Program," *American Economic Journal: Economic Policy*, 13, 239–270.

Ganglmair, Bernhard, W. Keith Robinson, and Michael Seeligson (2022) "The Rise of Process Claims: Evidence from a Century of U.S. Patents," unpublished manuscript, available at `https://papers.ssrn.com/abstract=4069994`.

Goldstein, Larry M. (2013) *True Patent Value: Defining Quality in Patents and Patent Portfolios*, Chicago, Ill.: True Value Press.

Goodman-Bacon, Andrew (2021) "Difference-in-Differences with Variation in Treatment Timing," *Journal of Econometrics*, 225, 254–277.

von Graevenitz, Georg, Stefan Wagner, and Dietmar Harhoff (2013) "Incidence and Growth of Patent Thickets: The Impact of Technological Opportunities and Complexity," *Journal of Industrial Economics*, 61 (3), 521–563.

Graham, Stuart J. H., Robert P. Merges, Pam Samuelson, and Ted Sichelman (2009) "High Technology Entrepreneurs and the Patent System: Results of the 2008 Berkeley Patent Survey," *Berkeley Technology Law Journal*, 24 (4), 1255–1328.

Graham, Stuart J. H. and Saurabh Vishnubhakat (2013) "Of Smart Phone Wars and Software Patents," *Journal of Economic Perspectives*, 27 (1), 67–86.

Gross, Daniel P. (2023) "The Hidden Costs of Securing Innovation: The Manifold Impacts of Compulsory Invention Secrecy," *Management Science*, 69 (4).

Grossman, Gene M. and Elhanan Helpman (1991) "Quality Ladders in the Theory of Growth," *Review of Economic Studies*, 58 (1), 43–61.

Hall, Bronwyn H., Christian Helmers, Mark Rogers, and Vania Sena (2013) "The Importance (or Not) of Patents to UK Firms," *Oxford Economic Papers*, 65 (3), 603–629.

——— (2014) "The Choice between Formal and Informal Intellectual Property: A Review," *Journal of Economic Literature*, 52 (2), 375–423.

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, Cambridge, Mass.

Hall, Bronwyn H. and Rosemarie H. Ziedonis (2001) "The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995," *RAND Journal of Economics*, 32 (1), 101–128.

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

Hegde, Deepak and Hong Luo (2018) "Patent Publication and the Market for Ideas," *Management Science*, 64 (2), 652–672.

Heller, Michael A. and Rebecca S. Eisenberg (1998) "Can Patents Deter Innovation? The Anticommons in Biomedical Research," *Science*, 280, 698–701.

Hopenhayn, Hugo A. and Matthew F. Mitchell (2001) "Innovation Variety and Patent Breadth," *RAND Journal of Economics*, 32 (1), 152–166.

Hsu, David H. and Rosemarie H. Ziedonis (2013) "Resources as Dual Sources of Advantage: Implications for Valuing Entrepreneurial-Firm Patents," *Strategic Management Journal*, 34 (7), 761–781.

Johnson, Matthew S., Michael Lipsitz, and Alison Pei (2023) "Innovation and the Enforceability of Noncompete Agreements," NBER Working Paper 31487, National Bureau of Economic Research, Cambridge, Mass.

Krishna, Vijay (2010) *Auction Theory*, Burlington, Mass.: Academic Press, 2nd edition.

Kuhn, Jeffrey M. and Neil C. Thompson (2019) "How to Measure and Draw Causal Inferences with Patent Scope," *International Journal of the Economics of Business*, 26 (1), 5–38.

Lanjouw, Jean O. and Mark Schankerman (2004) "Patent Quality and Research Productivity: Measuring Innovation With Multiple Indicators," *Economic Journal*, 114 (495), 441–465.

Leblebicioglu, Asli and Tanseli Savaser (forthcoming) "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *American Economic Journal: Economic Policy*.

Lemley, Mark A. (2008a) "Ignoring Patents," *Michigan State Law Review*, 2008, 19–34.

——— (2008b) "The Surprising Virtues of Treating Trade Secrets as IP Rights," *Stanford Law Review*, 61, 311–353.

Lerner, Josh (1994) "The Importance of Patent Scope: An Empirical Analysis," *RAND Journal of Economics*, 25 (2), 319–333.

——— (2006) "The New Financial Thing: The Origins of Financial Innovations," *Journal of Financial Economics*, 79 (2), 223–255.

Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter (1987) "Appropriating the Returns from Industrial Research and Development," *Brookings Papers on Economic Activity*, 3, 783–831.

Mansfield, Edwin (1986) "Patents and Innovation: An Empirical Study," *Management Science*, 32 (2), 173–181.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995) *Microeconomic Theory*, New York, N.Y.: Oxford University Press.

Merges, Robert P. (1988) "Commercial Success and Patent Standards: Economic Perspectives on Innovation," *California Law Review*, 76 (4), 803–876.

Moser, Petra (2012) "Innovation without Patents: Evidence from World's Fairs," *Journal of Law and Economics*, 55 (1), 43–74.

O'Donoghue, Ted, Suzanne Scotchmer, and Jacques-François Thisse (1998) "Patent Breadth, Patent Life, and the Pace of Technological Progress," *Journal of Economics and Management Strategy*, 7 (1), 1–32.

Ouellette, Lisa Larrimore (2012) "Do Patents Disclose Useful Information?," *Harvard Journal of Law and Technology*, 25 (2), 545–607.

Pakes, Ariel (1986) "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54 (4), 755–784.

Png, Ivan P.L. (2012) "Trade Secrets, Non-Competes, and Inventor Mobility: Empirical Evidence," unpublished manuscript, NUS Business School, National University of Singapore.

——— (2017a) "Secrecy and Patents: Theory and Evidence from the Uniform Trade Secrets

Act," *Strategy Science*, 2 (3), 176–193.

——— (2017b) "Law and Innovation: Evidence from State Trade Secrets Laws," *Review of Economics and Statistics*, 99 (1), 167–179.

Roin, Benjamin N. (2005) "The Disclosure Function of the Patent System (or Lack Thereof)," *Harvard Law Review*, 118 (6), 2007–2028.

Sampat, Bhaven and Heidi L. Williams (2018) "How do Patents Affect Follow-On Innovation? Evidence from the Human Genome," *American Economic Review*, 109 (1), 203–236.

Sandeen, Sharon K. and Elizabeth A. Rowe (2013) *Trade Secret Law in a Nutshell*, St. Paul, Minn.: West Academic Publishing.

Schankerman, Mark (1998) "How Valuable is Patent Protection? Estimates by Technology Field," *RAND Journal of Economics*, 29, 77–107.

Schankerman, Mark and Ariel Pakes (1986) "Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period," *Economic Journal*, 96 (384), 1052–1076.

Scotchmer, Suzanne (1991) "Standing on the Shoulders of Giants: Cumulative Research and the Patent Law," *Journal of Economic Perspectives*, 5 (1), 29–41.

——— (2004) *Innovation and Incentives*, Cambridge, Mass.: MIT Press.

Strandburg, Katherine J. (2004) "What Does the Public Get? Experimental Use and the Patent Bargain," *Wisconsin Law Review*, 2004, 81–174.

Sun, Liyang and Sarah Abraham (2021) "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 225 (2), 175–199.

Trajtenberg, Manuel, Rebecca Henderson, and Adam B. Jaffe (1997) "University Versus Corporate Patents: A Window on the Basicness of Invention," *Economics of Innovation and New Technology*, 5 (1), 19–50.

Wang, Yanzhi (2023) "Trade Secrets Laws and Technology Spillovers," *Research Policy*, 52 (7), 104794.

Webster, Elizabeth and Paul H. Jensen (2011) "Do Patents Matter for Commercialization?," *Journal of Law and Economics*, 54 (2), 431–453.

Williams, Heidi (2017) "How Do Patents Affect Research Investments?" *Annual Review of Economics*, 9 (1), 441–469.

Williams, Heidi L. (2013) "Intellectual Property Rights and Innovation: Evidence from the Human Genome," *Journal of Political Economy*, 121 (1), 1–27.

A Appendix

A.1 Theory Framework

A.1.1 A Model of Trade Secrets and Disclosure

We present a simple model of an inventor's decision whether to disclose a (patentable) invention through a patent or to keep the invention a secret. For this theoretical framework, we abstract away from factors that would ordinarily affect the patenting-secrecy tradeoff and instead examine the role of the invention's visibility.

An Inventor's Decision to Disclose: An invention i can be described by a tuple (ϕ, Θ, v). It is characterized by its visibility $\phi \in [0, 1]$, its type Θ, and its private commercial value $v \geq 0$ (from exclusive use). An inventor is given the choice to disclose an invention through a patent ($\tilde{d} = D$) or keep the invention secret ($\tilde{d} = S$).[31] We set the inventor's private returns $V_{\tilde{d}}$ equal to the *exclusivity-weighted* commercial value v, where we interpret v as the rents the inventor is able to appropriate from *exclusive* use of the invention.

In both disclosure states $\tilde{d} = D, S$, the probability of exclusive use depends on the *visibility* of the invention. Visibility is a two-way street. We refer to *disclosure-visibility*, denoted by ϕ_D, as the ease with which an inventor A can observe a firm B using A's (disclosed) invention. We refer to *secrecy-visibility*, denoted by ϕ_S, as the ease with which a firm B can observe inventor A using A's own (secret) invention. We will assume that, for a given invention, disclosure-visibility is higher than secrecy-visibility, $\phi_D \geq \phi_S$.[32] We set $\phi_D = \phi$ and $\phi_S = \xi\phi$ with $\xi \in (0, 1]$.

A patent for a more disclosure-visible invention is easier to enforce, and exclusivity prevails.[33] We can write the expected commercial value the inventor is able to materialize as

[31] This assumption of mutually exclusive states $\tilde{d}$ is for convenience and does not pose significant restrictions. Instead of a singleton invention, we can think of an invention that comprises both product and process elements, and for which the decision to patent or keep secret is made for each individual component.

[32] A simple argument for this assumption is that the inventor herself knows what to look for, whereas firm B has little prior guidance.

[33] Active monitoring of infringement is said to be a major source of the costs of patent enforcement (Hall et al., 2014).

$\phi_D v = \phi v$. In addition, the inventor receives a patent premium λ.[34] It captures the benefits from patenting over trade secrets and may even include licensing revenues from follow-on innovation. Collecting terms, we can summarize the inventor's private value of disclosing the invention (i.e., the value from patenting) as $V_D(\phi) = \phi\,(1+\lambda)\,v$.

The value from *secrecy* is determined by secrecy-visibility, $\phi_S = \xi\phi$; and the value of secrecy increases with the level of trade secrets protection. We denote the exogenous probability that a trade secret is protected by $\tau \in [0,1]$. Even with perfect trade secrets protection ($\tau = 1$), keeping the invention secret is of little value to the inventor if it is secrecy-visible. Conversely, weaker trade secrets protection reduces deterrence and results in more (unsanctioned) misappropriation of trade secrets (Friedman et al., 1991). We therefore assume that, without any trade secrets protection, the value of trade secrecy is zero even for non-visible inventions.[35] For tractability, we set the probability of detection of trade secret misappropriation equal to one, so that the only variation in the enforcement of trade secrets is through τ.[36] Collecting terms, we define the private value from secrecy as $V_S(\phi,\tau) = \tau\,(1-\xi\phi)\,v$.

The inventor chooses disclosure if, and only if, $V_D(\phi) \geq V_S(\phi,\tau)$, or

$$\phi \geq \frac{\tau}{1+\lambda+\xi\tau} =: \bar{\phi}(\tau). \tag{A.1}$$

For a given ϕ, we can write the decision to disclose as $\tilde{d} = D$ if $\phi \geq \bar{\phi}(\tau)$ and $\tilde{d} = S$ if otherwise. From the disclosure decision $\tilde{d}$ and the expression for $\bar{\phi}(\tau)$, we can conclude that an inventor is more likely to file for a patent as the degree of visibility ϕ increases, and she is less likely to patent as the degree of trade secrets protection τ increases.

[34] Patents are of additional value because, for instance, they signal the quality of the invention (Hsu and Ziedonis, 2013), convey reputation (Graham et al., 2009), or improve an inventor's bargaining position in license negotiations (Hall and Ziedonis, 2001). Webster and Jensen (2011) further provide evidence for a premium from commercialization.

[35] While the lack of legal sanctions is likely to encourage misappropriation, firms are expected to erect safeguards when trade secrets protection is weak (Friedman et al., 1991; Lemley, 2008b). These safeguards are often inefficient and their costs increase in v and decrease in τ. Without trade secrets protection, the effective commercial value may thus in fact fully dissipate.

[36] An alternative interpretation of τ is the product of the detectability of misappropriation and the strength of legal trade secrets protection.

Value of Trade Secrecy by Invention Type: We assume that an invention's visibility ϕ is unobservable but distributed on the unit support with cdf G_Θ. What is observable is an invention's *type* Θ that is correlated with its visibility. An invention is either a process (or method), $\Theta = M$, or a product, $\Theta = P$. The probability that the realized invention is a process is $\theta = \Pr(\Theta = M)$.

We assume that processes are on average less visible than products.[37] The (expected) values of secrecy $EV_{S|\Theta}(\tau)$ and disclosure $EV_{D|\Theta}(\tau)$ of an invention of type Θ are

$$EV_{S|\Theta}(\tau) = \int_0^1 \tau\,(1-\xi\phi)\,v dG_\Theta \qquad \text{and} \qquad EV_{D|\Theta}(\tau) = \int_0^1 \phi\,(1+\lambda)\,v dG_\Theta. \tag{A.2}$$

Proposition 1. *For a given level of trade secrets protection τ, the value of secrecy is higher for processes than for products. Conversely, the value of disclosure is lower for processes than for products.*

Proof. For the proof of this claim, we utilize the stochastic dominance property of our visibility distributions. As stated in the text, our assumption of hazard rate dominance implies first-order stochastic dominance (Krishna, 2010:276). It will be useful to first state the definition and general property of first-order stochastic dominance. We follow the treatment in Mas-Colell et al. (1995:195). Let $u(x)$ be a non-decreasing function in $x \in [0,1]$. Then

$$\int u(x)dG_P(x) \geq \int u(x)dG_M(x) \iff G_P(x) \overset{FOSD}{\succ} G_M(x). \tag{A.3}$$

Integrating by parts, we obtain

$$\int u(x)dG_\Theta(x) = [u(x)G_\Theta(x)]_0^1 - \int u'(x)G_\Theta(x)dx$$

[37] We formally capture this by assuming *hazard-rate dominance*. The distribution G_P hazard-rate dominates G_M if $\frac{g_p(\phi)}{1-G_P(\phi)} \leq \frac{g_M(\phi)}{1-G_M(\phi)}$ for all ϕ. Moreover, this implies that G_P first-order stochastically dominates G_M so that $G_P \leq G_M$ for all ϕ (Krishna, 2010:276).

Because $G_\Theta(0) = 0$ and $G_\Theta(1) = 1$ for $\Theta = M, P$, we can rewrite the condition in the claim as

$$\int u(x)dG_P(x) - \int u(x)dG_M(x) = \int u'(x)\left[G_M(x) - G_P(x)\right]dx \geq 0.$$

Because $G_P(x) \leq G_M(x)$ by first-order stochastic dominance, the condition holds for any non-decreasing function so that $u'(x) \geq 0$. Note that if $u(x)$ is strictly increasing and $G_P(x) < G_M(x)$ for some x, then the inequality is strict.

For the first claim in the proposition, $EV_{S|M}(\tau) > EV_{S|P}(\tau)$, $\tau\,(1 - \xi\phi)\,v$ is a strictly decreasing function in ϕ (because $\xi > 0$). We can simply rewrite the inequality as $-EV_{S|P}(\tau) > -EV_{S|M}(\tau)$ or

$$\int_0^1 \underbrace{-\tau\,(1-\xi\phi)\,v}_{u(\phi)}\,dG_P(\phi) \quad > \quad \int_0^1 \underbrace{-\tau\,(1-\xi\phi)\,v}_{u(\phi)}\,dG_M(\phi) \tag{A.4}$$

with $u(\phi)$ increasing in ϕ so that the general property above applies. We obtain a strict inequality by the implicit assumption that $G_M(\phi)$ and $G_P(\phi)$ are not identical so that $G_P(\phi) < G_M(\phi)$ for some ϕ. For the second claim, $EV_{D|M}(\tau) < EV_{D|P}(\tau)$, note that $\phi\,(1+\lambda)\,v$ is strictly increasing in ϕ, and the above general property applies. □

Evidence from survey data, finding that the propensity to patent is higher for products than processes and thus suggesting a higher value of secrecy for processes, comports with this theoretical finding (e.g., Levin et al., 1987; Cohen et al., 2000; Hall et al., 2013).

Probability of Disclosure for Invention Types: We derive the probability ρ that a given patent covers a process invention. We first establish two auxiliary results (for visualizations, see panel (a) in Figure 2). In Lemma 1, we show that the probability that a process is patented is weakly smaller than the probability that a product is patented. For this, let $d(\phi, \tau) = 1$ if $\tilde{d} = D$ and $d(\phi, \tau) = 0$ if $\tilde{d} = S$. The probability that an invention of type Θ

is patented is

$$d_\Theta(\tau) = \int_0^1 d(\phi,\tau)dG_\Theta(\phi) = \int_{\bar{\phi}(\tau)}^1 1 \cdot dG_\Theta(\phi) = 1 - G_\Theta(\bar{\phi}(\tau)). \tag{A.5}$$

Lemma 1. *For a given level of trade secrets protection τ, $d_M(\tau) \leq d_P(\tau)$.*

Proof. For any given τ, $d_M(\tau) \leq d_P(\tau)$ if, and only if, $G_P(\bar{\phi}(\tau)) \leq G_M(\bar{\phi}(\tau))$. The latter holds by first-order stochastic dominance of G_P over G_M. □

In Lemma 2, we establish the relationship between patenting probabilities and trade secrets protection.

Lemma 2. *The patenting probabilities for products $d_P(\tau)$ and processes $d_M(\tau)$ decrease in trade secrets protection τ.*

Proof. Patenting probabilities (weakly) decrease in τ if $d_\Theta(\tau)$ is (weakly) decreasing in τ. We have $\frac{\partial\bar{\phi}(\tau)}{\partial\tau} = \frac{1+\lambda}{(1+\lambda+\xi\tau)^2} > 0$ so that $G_\Theta(\bar{\phi}(\tau))$ increases in τ and $d_\Theta(\tau) = 1 - G_\Theta(\bar{\phi}(\tau))$ decreases in τ. □

Given the underlying distribution of invention types with $\theta = \Pr(\Theta = M)$, Bayes' rule gives us the probability that a given patent covers a process:

$$\rho(\tau) = \frac{\theta d_M(\tau)}{\theta d_M(\tau) + (1-\theta)\, d_P(\tau)}. \tag{A.6}$$

Proposition 2. *Given the pool of inventions, the probability that a given patent covers a process is decreasing as trade secrets protection increases.*

Proof. Using $d_M(\tau) = 1 - G_M(\bar{\phi}(\tau))$ and $d_P(\tau) = 1 - G_P(\bar{\phi}(\tau))$, the first derivative of $\rho(\tau)$ with respect to trade secrets protection τ is

$$\frac{\partial\rho(\tau)}{\partial\tau} = \frac{-(1-\theta)\,\theta\,[(1-G_P)\,g_M - (1-G_M)\,g_P]\,\bar{\phi}'}{(\theta\,(1-G_M) + (1-\theta)\,(1-G_P))^2}$$

where $\bar{\phi}' > 0$ is the partial derivative of $\bar{\phi}(\tau)$ with respect to τ and G_Θ and g_Θ are evaluated at $\bar{\phi}(\tau)$. The probability $\rho(\tau)$ decreases in τ if the term in brackets in the numerator is non-negative so that $(1 - G_P)\, g_M \geq (1 - G_M)\, g_P$ or $\frac{g_M}{1-G_M} \geq \frac{g_P}{1-G_P}$. The latter inequality holds by the assumption of G_P hazard-rate dominating G_M. □

The expression in Equation (A.6) can also be interpreted as the share of process patents in a sample of patents. Proposition 2 predicts that the probability that a given patent is a process patent decreases in response to an (exogenous) increase in trade secrets protection.

A.1.2 Simulation

We make the following assumptions for the simulations of the three-stage cumulative innovation model in Section 5 with 100,000 initial ideas i.

Stage 1 (Initial R&D): An inventor observes a *potential* invention (idea) i with characteristics (ϕ, Θ). The invention's visibility ϕ is drawn from an invention-type specific distribution with cdf F_Θ: the visibilities are uniformly distributed between with support $[0, 0.75]$ for processes and $[0.25, 1]$ for products. We assume that disclosure-visibility and secrecy-visibility are the same (so that $\xi = 1$ from Section A.1). Invention types Θ (product or process) are binary, and the probability that a potential invention is a process is $\theta^F = \Pr(\Theta = M)$. Before any investment is made, the inventor observes R&D costs C_i drawn from a logistic distribution with variance 0.5. We choose two location parameters for the cost distributions to describe industries with varying R&D costs: low average costs (30% of the expected value of the R&D investment) and high average costs (60%). The inventor always observes the invention's commercial value v_i that is exponentially distributed with rate parameter 1/10 with an expected value of 10. She undertakes the R&D project if the expected payoffs from the invention (including the value and licensing revenues from both the invention, i.e., the patent premium λ equaling 0.1—in line with Schankerman, 1998, and potential follow-on innovation) outweigh its cost.

Stage 2 (Patent or Trade Secret): The second stage of our augmented model is the disclosure model in Section A.1. Conditional on a positive R&D decision, the disclosure decision depends on the strength of trade secrets protection τ and the invention's realized visibility ϕ_i.

Stage 3 (Follow-on Innovation): For any realized initial invention i, we model follow-on innovation as one representative invention i_F with random value v_{i_F} (exponentially distributed with rate parameter 1/35 with an expected value of 35 that is 3.5 times the expected value of the initial invention) and cost C_{i_F} (logistically distributed with variance 0.5). Follow-on innovation can only happen if it is profitable (i.e., $v_{i_F} \geq C_{i_F}$). If it is, then the realization depends on how much of the initial invention i is visible after the inventor's disclosure decision. We refer to this measure as *effective visibility* of initial invention i and denote it by $\tilde{\phi}_i$. It is equal to

$$\tilde{\phi}_i = \begin{cases} \phi_i & \text{if R\&D in Stage 1 and } \textit{trade secret} \text{ in Stage 2;} \\ 1 & \text{if R\&D in Stage 1 and } \textit{patent} \text{ in Stage 2.} \end{cases} \tag{A.7}$$

Effective visibility is equal to the invention's visibility ϕ_i if the invention is realized but kept as a trade secret. We assume, without loss of generality, that the invention is fully disclosed through patenting so that effective visibility of a patented invention is equal to 1.[38]

In addition to the effective visibility, the probability that follow-on innovation is successful also depends on barriers to access to the initial invention. We capture how much patents—and their potential anticommons effect—lower the success probability of follow-on innovation by a scale parameter $\psi_D < 1$. For secrets, we normalize this parameter to $\psi_S = 1$. The success probability of follow-on innovation is then $\tilde{\psi}_{i_F,\tilde{d}} = \psi_{\tilde{d}}\tilde{\phi}_i$ following a realized initial invention with disclosure state $\tilde{d} \in \{S, D\}$. The baseline success probabilities for follow-on innovation are $\phi_S = 1$ for secret Stage 1 inventions (so that $\tilde{\psi}_S = \phi$) and $\phi_D = 2/3$ for

[38] The assumption of perfect disclosure through patenting is to simplify the analysis. Our results hold as long as patents provide more disclosure than secrecy.

patented inventions (so that $\tilde{\psi}_D = 2/3$).

This model for follow-on innovation at Stage 3 is simple but nonetheless consistent with stylized facts and other models proposed in the literature. We make three main assumptions. First, follow-on innovation as captured by v_{i_F} is by other firms rather than the inventor of the initial innovation. For instance, Sampat and Williams (2018) document that, for their sample of genome patents, most of follow-on research is done by firms other than the patent assignee. Follow-on innovation by the initial inventor does not explicitly enter our model but could be captured by v_i and is not dependent on the effective visibility of any part of the initial invention. Second, disclosure has a positive effect on follow-on innovation. Williams (2013) documents that a *restriction* of access to human genome data leads to a 20–40% *reduction* in follow-on research. Similarly, Gross (2023) and de Rassenfosse et al. (2024) find that compulsory secrecy of patent applications resulted in fewer citations.

Third, conditional on the effective visibility, the baseline probability of follow-on innovation to a trade secret is higher than that following a patent. This assumption reflects the anticommons effect (e.g., Heller and Eisenberg, 1998) where technologies are underused because patents on early ideas raise the costs of creating future ideas by introducing friction in the bargaining process over licenses.

Welfare Measure: We use the *expected total value added* of a given idea, denoted by $W(\tau)$, as our welfare measure. It is calculated as the weighted sum of the aggregate surplus from the realized initial invention, W_i (which depends on its disclosure state, $\tilde{d}_i$), and the aggregate surplus from realized follow-on innovation, W_{i_F}. The expected total value added of a potential idea i is equal to

$$\overline{W}(\tau) = E_{(\Theta_i, \phi_i, \tilde{d}_i, v_i, v_{i_F})} \left[\mathbf{R}_i(\tau) \left(W_i + \tilde{\psi}_{i_F, \tilde{d}_i} \mathbf{R}_{i_F} W_{i_F} \right) \right], \tag{A.8}$$

where expectations $E_{(\cdot)}$ are over the invention type Θ, visibility ϕ, disclosure state $\tilde{d}$, and commercial values v_i for initial and v_{i_F} for follow-on innovation. Further, $\mathbf{R}_i$ ($\mathbf{R}_{i_F}$) is an indicator that is equal to 1 if the initial (follow-on) R&D project is undertaken, and W_i and W_{i_F} are measures of aggregate surplus from initial and follow-on innovation.

We determine $\mathbf{R}_i$ and $\mathbf{R}_{i_F}$ as follows. Denote by EV_i the expected gross value of the invention to the inventor: the maximum between the expected value of secrecy ($EV_{S|\Theta}(\tau)$) and disclosure through patenting ($EV_{D|\Theta}(\tau)$). The inventor decides to undertake the initial R&D project ($\mathbf{R}_i = 1$) if $EV_i \geq C_i$. Similarly, the follow-on invention is realized ($\mathbf{R}_{i_F} = 1$) if it is profitable and successful. It is profitable if the commercial value covers the costs, $v_{i_F} \geq C_{i_F}$ and successful with probability $\tilde{\psi}_{i_F,\tilde{d}}$.

For the measures of aggregate surplus W_i, we assume that $2v_i$ is the *potential* aggregate surplus that materializes when there are no barriers to access to the invention. Because the barriers to access depend on the inventor's disclosure decision, the realized aggregate surplus is the potential aggregate surplus net of the disclosure-state specific deadweight loss. For instance, in the textbook case of linear demand with unit market size (and zero marginal cost), non-price discriminating monopoly profits ($=v_i$) are one-half of the aggregate surplus ($=2v_i$), and consumer surplus and deadweight loss are one quarter each ($=v_i/2$). This value represents the maximum deadweight loss (from a scenario with full barriers to access). We provide a concrete example below.

For patented inventions, barriers to access increase in visibility ϕ, and the aggregate surplus, W_D, as a function of visibility is equal to

$$W_D(\phi) = 2v_i - \frac{\phi v_i}{2} - C_i, \tag{A.9}$$

where C_i is the cost of R&D of the potential idea. For inventions kept as trade secrets, barriers to access decrease in ϕ and increase in trade secrets protection τ. As discussed in Section A.1, the probability that the inventor has exclusive access, implying full monopolistic deadweight loss, is equal to $\tau(1-\phi)$. Aggregate surplus, W_S for an invention that is kept secret is therefore equal to

$$W_S(\phi, \tau) = 2v_i - \frac{\tau\,(1-\phi)\,v_i}{2} - C_i. \tag{A.10}$$

To summarize, using the disclosure condition in Equation (A.1), the aggregate surplus of the initial invention is $W_i = W_D(\phi)$ if $\phi \geq \bar{\phi}(\tau)$ and $W_i = W_S(\phi, \tau)$ otherwise. For the aggregate surplus of any realized follow-on innovation, we assume free access, so that $W_{i_F} = 2v_{i_F} - C_{i_F}$.

A.2 Representativeness of the Sample

Because our main regression sample is limited to patents whose U.S. assignees and inventors are all from the same state, we introduce the possibility of sample selection. We examine this possibility by comparing our variables of interest across three samples: (1) *all* utility patents with priority dates between 1976 and 2008 and granted between 1976 and 2014 for which we observe the relevant information (4,287,180 patents); (2) the subset of patents with any U.S. assignee or inventor (2,391,486 patents); and (3) the subset of patents for which all U.S. assignees and inventors are located in the same state (our main estimation sample, 1,451,311 patents). Table A.1 shows summary statistics for our process patent indicator as well as the control variables. The regression sample (rightmost column) has a slightly higher share of process patents than the total population of patents, but smaller than the population of U.S. patents. They also seem to have slightly higher degrees of originality and generality. We control for these variables in the main estimation.

Figure A.1 illustrates the distributions of the sizes of the applicants. It shows that our regression sample over-represents individual applicants and under-represents large firms. Because small applicants (individuals and small firms) see the largest effect (see Section 4.2.2), our *average* treatment effects may be slightly over-estimated.

Table A.1: Summary Statistics for Different Subsamples

	All		All US		Single-State	
	Mean	SD	Mean	SD	Mean	SD
Process patent	0.459	0.498	0.507	0.500	0.473	0.499
Number of process claims	0.799	1.294	0.919	1.400	0.871	1.407
Number of product claims	1.781	1.798	1.875	1.872	1.920	1.885
Log(indep. claims)	1.185	0.450	1.246	0.452	1.242	0.453
Log(length of first claim)	4.989	0.582	4.953	0.594	4.976	0.584
Log(length of description)	9.716	0.965	9.759	0.959	9.699	0.951
Originality	0.602	0.253	0.632	0.240	0.626	0.244
Generality	0.606	0.263	0.634	0.249	0.638	0.244
4th year renewal	0.838	0.368	0.840	0.367	0.826	0.380
Observations	4287180		2391486		1451311	

Notes: This table provides summary statistics for all granted utility patents (between 1976 and 2014) with priority dates between 1976 and 2008. Column (1) shows statistics for all patents; Column (2) shows statistics for patents with at least one U.S. assignee or inventor; Column (3) uses single-state patents.

Figure A.1: Applicant Size Distributions for Different Subsamples

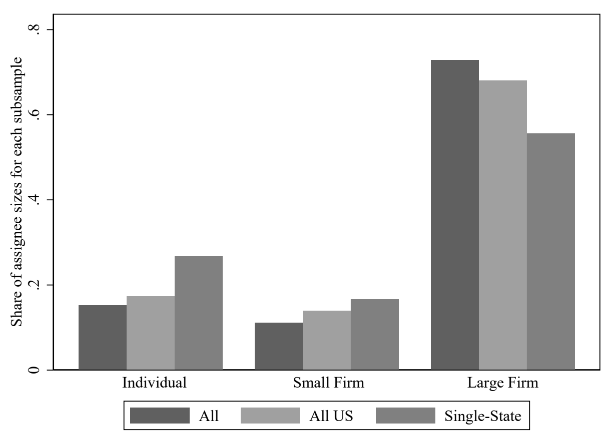

Notes: This figure presents shares of applicant sizes of different subsamples of all granted utility patents (between 1976 and 2014) with priority dates between 1976 and 2008. The darkest (leftmost) column shows statistics for all patents; the lightest (middle) column shows statistics for patents with at least one U.S. assignee or inventor; the rightmost column uses single-state patents.

ONLINE APPENDIX – NOT FOR PUBLICATION

Visibility of Technology and Cumulative Innovation: Evidence from Trade Secrets Laws

Data Appendix*

Bernhard Ganglmair† Imke Reimers‡

September 23, 2024

*Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B2) is gratefully acknowledged. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†University of Mannheim, ZEW Mannheim, and MaCCI. E-mail: b.ganglmair@gmail.com.

‡Cornell University.E-mail: imke.reimers@cornell.edu.

C Data Appendix

We construct our data sample using a number of sources. We obtain basic bibliographic information from PatentsView at `https://www.patentsview.org/download` for bulk download and `http://www.patentsview.org/api/doc.html` for API queries. We also use data from Ganglmair et al. (2022) for process patent indicators, the USPTO's Patent Maintenance Fee Events database at `https://bulkdata.uspto.gov/data/patent/maintenancefee` to calculate proxies for patent value as well as applicant size, the USPTO's Patent and Patent Application Claims Research Dataset at `https://bulkdata.uspto.gov/data/patent/claims/economics/2014/` for proxies of patent scope and complexity, and the Google Patents Research Data (`https://console.cloud.google.com/marketplace/partners/patents-public-data`) to construct data on the timing of disclosure. In Table C.1, we provide an overview of the steps of our sample construction. For further details, see the descriptions that follow.

Table C.1: Sample Construction and Sample Size

Sample/Variable	Source
Patents, granted January 1976 – December 2014	PatentsView
Priority dates: January 1976 – December 2008	Google Patents
U.S. only location	*constructed*
Exclude business method patents	PatentsView
Process patent indicator	Ganglmair et al. (2022)
	Main Estimation Sample: 1,451,311
Number of independent claims	USPTO Claims
Length of first claim	USPTO Claims
Length of detailed patent description	PatentsView (API)
Originality	*constructed*
Generality	*constructed*
4th year maintenance	USPTO Maintenance
USPC main classes	PatentsView
Applicant size	*constructed*
Complex vs. discrete technology	von Graevenitz et al. (2013)
NBER technology categories	PatentsView

Notes: Data sources are PatentsView (bulk data download page and API), Google Patents (Google Patents Research Data), USPTO Claims (USPTO's Patent and Patent Application Claims Research Dataset), USPTO Maintenance (USPTO's Patent Maintenance Fee Events database), and Ganglmair et al. (2022). *Constructed* means that variables are constructed/calculated by authors. For more details, see the descriptions below.

C.1 Main Sample

For our data sample, we start with the census of U.S. utility patents granted between 1976 and 2014. In order to obtain a clean assignment of the level of trade secrets protection to which the patent applicant was exposed at the time of the disclosure decision, we limit our sample to patents with disclosure *timing* between 1976 and 2008 and a *location* within the United States.

Timing: Priority Dates To identify the timing of the disclosure decision, we use the priority date of the head of a simple patent family (i.e., all patents that share the same priority claims). We implement this by using the earliest priority date for all patents from a given simple patent family. Information on simple patent family assignment and priority dates we obtain from the Google Patents Research Data.

Location To identify the location (i.e., U.S. state) of the disclosure decision, we use information on the locations of patent assignees and inventors. PatentsView provides data on disambiguated location, assignee, and inventor names. For each patent, we consider only assignees and inventors within the United States. Out of this subsample of names, we further consider only those patents for which all U.S. assignees and all U.S. inventors are located in the same state. We use this state as the respective state of the disclosure decision (and, by assumption, the relevant U.S. state for the UTSA adoption and trade secrets protection).

For a set of robustness results in the Appendix, we use a more aggressive location definition. First, we define the location of a patent by the location of the first assignee listed on the granted patent. If no assignee is listed, we use the location of the first inventor listed on the granted patent. Second, we limit our attention to U.S.-only patents, dropping all patents with an international contributor.

C.2 Patent Classification

For basic information on patent classification, we use the current United States Patent Classification (USPC) main classes (applied to all patents retrospectively) obtained from PatentsView. Where multiple main classes are listed on a patent, we use the first (by sequence).

For our main estimation sample, we exclude all business methods patents. We follow Lerner (2006) and define such patents as those with USPC main class 705 (i.e., the first main class listed on the patent). For a set of robustness results, we also rerun our analysis for a subsample that excludes software patents. For the classification of software patents, we follow Graham and Vishnubhakat (2013:fn 7).

Note that for our structural estimates, we use an extended sample that includes all granted patent through 2016. We discuss the reasons for this extension below.

C.3 Construction of Additional Variables

We further collect and construct three sets of variables to proxy a patent's "patent scope and complexity," its "external impact," and its "internal value." For our heterogeneity results, we also collect and construct a variable capturing the size of the patent applicant and an indicator variable of complex vs. discrete technologies.

C.3.1 Patent Scope and Complexity

We follow Lerner (1994) and Lanjouw and Schankerman (2004) and measure patent breadth and scope using the number of independent claims in a patent. Kuhn and Thompson (2019), however, argue that a simple count of (independent) claims may be a poor measure for patent scope.[1] They propose the length of the first patent claim as an alternative measure for patent scope, where shorter claims are broader. They use the first claim for their measure because under U.S. law the broadest claim should be listed first. We adopt their measure (length of the first claim in number of words) alongside the number of independent claims.

We collect the number of independent claims of a paper and the length of the first claim from the USPTO's Patent and Patent Application Claims Research Dataset at `https://bulkdata.uspto.gov/data/patent/claims/economics/2014`. This research dataset provides information on claims from patents granted between January 1976 and December 2014. For more details on the data, see Marco et al. (2016).

We further collect the length (in characters) of the detailed description of each patent from PatentsView through API queries (the data are not available for bulk data download at `http://www.patentsview.org/download`).

C.3.2 External Impact

We construct measures of patent generality and patent originality as proposed by Trajtenberg et al. (1997). See also Hall et al. (2001).

[1] Because each claim beyond 20 claims comes at an additional cost, patents with many claims may cover more valuable technologies, but need not be broader than patents with fewer claims.

Patent Originality: Patent originality of a patent j is defined as

$$1 - \sum_{k=1}^{n} \left(\frac{\text{backward citations}_{jk}}{\sum_{m=1}^{n} \text{backward citations}_{jm}} \right)^2 \tag{C.1}$$

where $s_{jk} = \frac{\text{backward citations}_{jk}}{\sum_{m=1}^{n} \text{backward citations}_{jm}}$ is the share of backward citations that patent j makes to patents in patent class $k = 1, \ldots, n$ over all backward citations made by patent j. A higher originality score means patent j draws on prior knowledge from a greater variety of fields. We construct this measure using the first listed USPC main class on a patent j. We have classification information for patents granted in and after 1976. This means that for patents granted early in our sample period that cite patents granted before 1976, we have little information about the classes of their cited patents. Because of this truncation issue, the originality measure is therefore noisier and coarser for earlier patents than for patents granted later in our sample period.

Patent Generality: Patent generality of a patent j is defined as

$$1 - \sum_{k=1}^{n} \left(\frac{\text{forward citations}_{jk}}{\sum_{m=1}^{n} \text{forward citations}_{jm}} \right)^2 \tag{C.2}$$

where $s_{jk} = \frac{\text{forward citations}_{jk}}{\sum_{m=1}^{n} \text{forward citations}_{jm}}$ is the share of forward citations that patent j receives from patents in patent class $k = 1, \ldots, n$ over all forward citations received by patent j. A higher generality score implies a higher widespread impact, influencing subsequent innovation in a broader variety of fields. A large number of patents never receive a patent citation, and our patent generality score is not defined for any patents without forward citations.

C.3.3 Internal Value

We use information on the applicant's renewal behavior as a measure of internal (or private) value of a patent (Pakes, 1986; Schankerman and Pakes, 1986). To this end, we construct a dummy variable equal to 1 if the applicant has paid the 4th-year maintenance fees (to be paid in the fourth year after patent grant).

We use information from the USPTO's Patent Maintenance Fee Events database at `https://bulkdata.uspto.gov/data/patent/maintenancefee` (Jan. 28, 2019). The database contains all recorded events related to the payment of maintenance fees for patents granted from September 1, 1981 and forward. A patent is said to have been maintained if one of the codes listed in Table C.2 is recorded.

Table C.2: Codes for Maintenance Fee Events

Code	Description
F170	Payment of Maintenance Fee, 4th Year
F173	Payment of Maintenance Fee, 4th Year, Undiscounted Entity
F273	Payment of Maintenance Fee, 4th Year, Small Entity
M1551	Payment of Maintenance Fee, 4th Year, Large Entity
M170	Payment of Maintenance Fee, 4th Year, PL 96-517
M173	Payment of Maintenance Fee, 4th Year, PL 97-247
M183	Payment of Maintenance Fee, 4th Year, Large Entity
M2551	Payment of Maintenance Fee, 4th Yr, Small Entity
M273	Payment of Maintenance Fee, 4th Yr, Small Entity, PL 97-247
M283	Payment of Maintenance Fee, 4th Yr, Small Entity
M3551	Payment of Maintenance Fee, 4th Year, Micro Entity

Source: Documentation file for Patent Maintenance Fee Events database at `https://bulkdata.uspto.gov/data/patent/maintenancefee`.

Because we have information on maintenance events through the end of 2018, covering the full four years after our main sample ends, we do not face any truncation issues for an applicant's 4th year maintenance decision. Note, however, that because maintenance information is available only for patents granted on or after September 1, 1981, we have 94,323 missing observations for patents granted between January 1976 and August 1981. Further note that we are not restricted by this truncation issue for our structural estimations and therefore use an extended sample with patents granted through December 2016.

C.3.4 Applicant Size

For our variable of applicant size (or entity size), we combine information from the USPTO's Patent Maintenance Fee Events database and bibliographic information on patents from PatentsView. Applicant size takes two values. It is equal to 1 if the applicant is an individual or a small firm (i.e., small entity but not an individual), and equal to 2 if the applicant is a large firm (i.e., large entity but not an individual).

The USPTO's Patent Maintenance Fee Events database provides information on the size of the entity for any recorded maintenance fee event. Entities are either micro or small ("small") or "large." This means, if an applicant's maintenance event for a patent j is recorded in the database, then we know the size of that patent j's applicant. Using assignee information (from PatentsView), we construct an applicant's size history (by year), based on recorded maintenance events. We hold the size of an applicant constant at the value of t until the next recorded event at $t' > t$ where it may or may not change. In addition, we use the size of the first entry for all previous years. With this size history, we can now assign an applicant size for all patents j of an assignee for which no maintenance event is recorded. This gives us size information for all patents by assignees that have at least one recorded maintenance event; patents by assignees without any maintenance events are without applicant size.

An applicant of a given patent j is an individual ($= 1$) if the first assignee listed on the patent is of type "individual" or if no assignee is listed on the patent. If the applicant is not an individual, then its size is equal to 2 if it is a small entity and equal to 3 if it is a large entity (as defined above). For the distribution of applicant size (for different definitions of patent location) see Figure B.2 in the Online Appendix.

C.3.5 Complex vs. Discrete Technologies

We classify technologies as either complex or discrete following the classification in von Graevenitz et al. (2013:Table VIII). Their classification builds on the 30 technology areas according to OST-INPI/FhG-ISI (see Schmoch, 2008:Table 1).

C.4 Process Patent Indicator

C.4.1 Summary of Indicator Construction

Ganglmair et al. (2022) employ text-analytical methods to identify the invention type of all independent claims in a given patent. We aggregate their claim-level data to obtain data at the patent level. In the sequel, we summarize their approach. Some of the material is also borrowed from Rosenberg (2012). An additional useful source of further background information is WIPO (2007).

The unit of analysis in Ganglmair et al. (2022) is an independent patent claim. A patent claim defines the scope of legal protection provided by a patent. It describes what the applicant claims to be its invention for which the patent grants exclusive rights. Each patent can hold multiple claims of different types. An *independent* claim stands on its own whereas a *dependent* claim is in reference to an independent claim, further limiting its scope.

Claims typically consist of two parts: a *preamble* and *body*. The preamble is an introductory phrase or paragraph that identifies the category of the invention of the claim. For example, an invention may be an apparatus or device (as in an *apparatus or device claim*, here referred to as *product claim*) or a method or process (as in a *method claim* or *process claim*). The *body* of a patent claim recites the elements of the claim. In many cases, these elements are *steps* (as in the steps of a process) or *items* (as in the items that define a product).

The approach in Ganglmair et al. (2022) uses information from both the preamble and the body. Both parts of the claim are classified as describing a process or a product. For the preamble, this classification is conducted via a simple keyword search (e.g., "process" or "method" for process-claim preambles; "apparatus" or "device" for product-claim preambles). For the body, the authors take a syntax-analysis approach, analyzing the linguistic structure of each line (or "bullet point") in the body. The steps of a process are listed using the gerund form of a verb, whereas the items of a product (an apparatus, a device) are listed as components. The authors' algorithm accounts for these drafting conventions when classifying a body as process-claim body or product-claim body. In the end, combining the classifications of the preamble and the body, a classification for the entire claim is obtained:

Process claim or method claim: A *process claim* (also called a *method claim*) describes the sequence of *steps* which together complete a task such as making an article of some sort. The preamble of a method claim often uses the terms "process" or "method." The body of a method claim typically consists of a listing of the "steps" of the process.

Product claim: A *product claim* (also called a "device claim" or "apparatus claim") describes an invention in the form of a physical apparatus, system, or device. For instance, a claim that covers a tripod for a camera or a window crank is an apparatus claim. In the preamble of a product claim, the patent applicant often recites what the product is and what it does. Then, in the body of the claim, the applicant lists the essential elements (i.e., "items") of the invention.

In addition to process claims and product claims, the special case of product-by-process claim is classified.

Product-by-process claim: A *product-by-process* claim is a claim that defines a product by the process of making it. The product-by-process claim defines a product by several process steps. Though, ultimately, the scope of the claim's coverage is directed toward a physical article (i.e., the "product") rather than the method, the claim includes elements of both product claiming (i.e., elements in the body that describe the items that comprise an article or product) and the sort of steps found in a process claim.

The authors' algorithm deals at great length with a number of issues: badly formatted claims, claims not following the usual drafting conventions, and two-part claims (also called improvement claims or Jepson claims). They have also compiled a dataset of close to 10,000 manually classified claims to test their algorithm and verify the results.

In Figure C.1, we plot the fraction of missing observations for each of our patent-level indicator. For both our main indicator and the process patent indicator with a majority of

Figure C.1: Share of Missing Observations

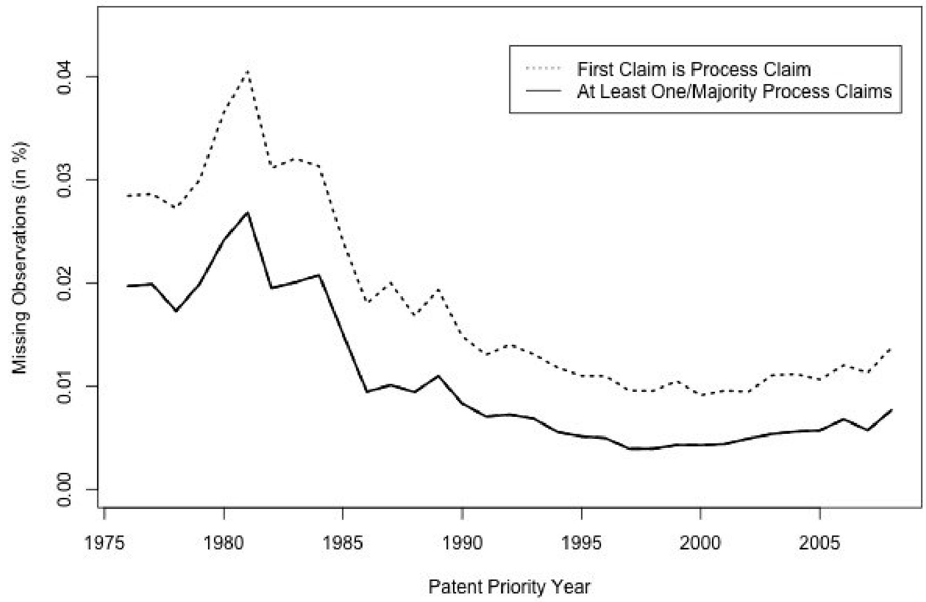

process claims, at least one patent claim must be classified - the graphs in the figure are therefore the same. The requirement for the indicator of the first process claim is stricter, and the number of missing observations is higher throughout. Notice, however, that the reliability of the approach increases over time as the percentage of missing observations (over all patents in our main sample) drops below 1% around 1985 (with higher numbers for patents with earlier priority dates).

C.4.2 Descriptive Figures

In Figure C.2, we plot the share of process patents by priority year. We show graphs for each of our three process patent indicators. The solid line depicts the share of process patents for our main indicator (at least one patent claim is a process claim, 'Any'). The dotted graph depicts the share of patents with the first patent claim a process claim ('First'); the dashed graph depicts the share of patents with a majority of process claims ('Most'). As we we have discussed in the main text, our main indicator is the most aggressive in terms of identifying patents as process patents. The overall time trends, however, are very similar. We also plot the average share of process claims in a patent (dash-dotted line). The graph follows similar trends.

In Figure C.3, we depict the share of process patents by applicant size. The share of process patents is higher in larger firms than in smaller firms, and lowest for individuals.

Figure C.2: Share of Process Patents (Multiple Indicators)

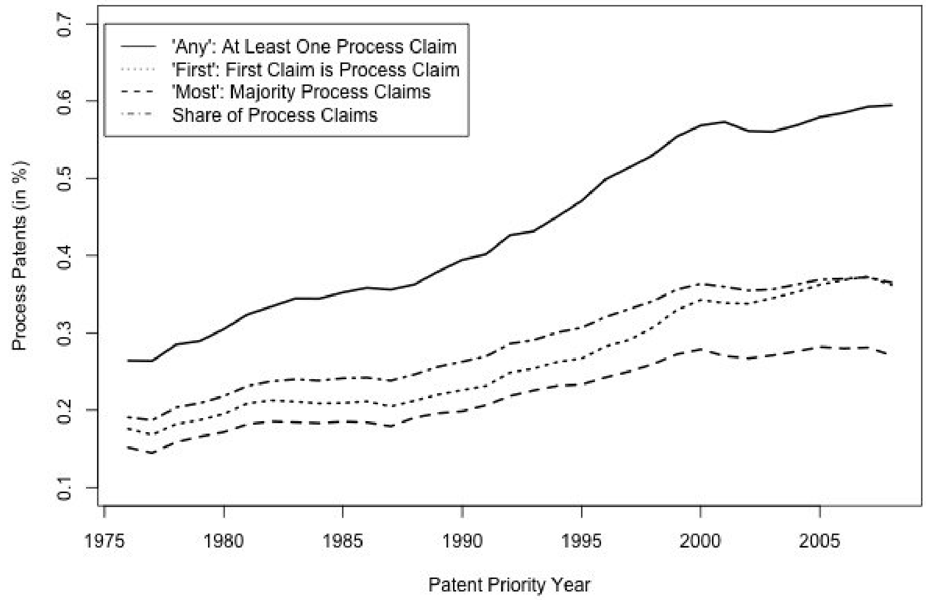

Figure C.3: Share of Process Patents by Applicant Size

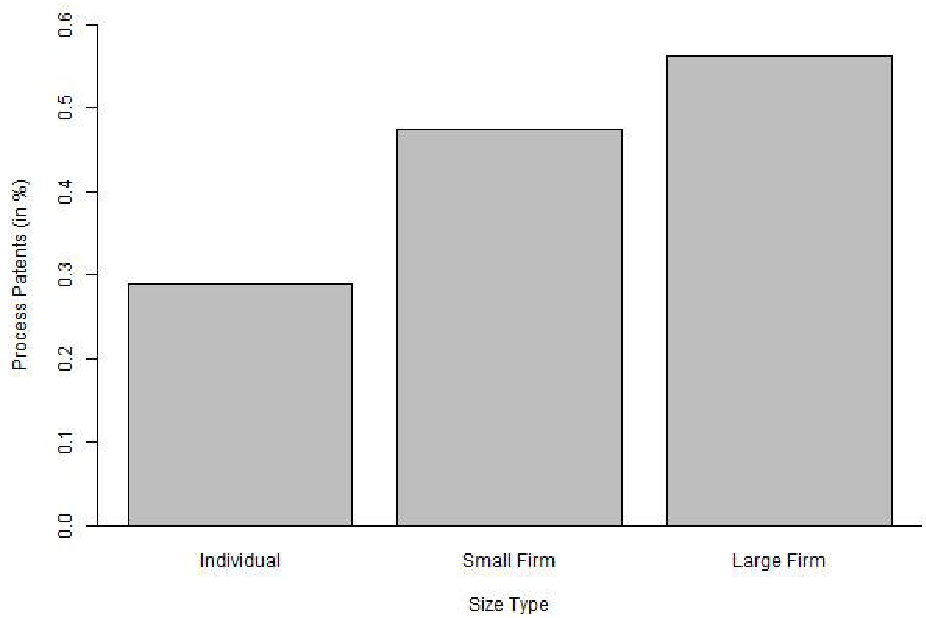

References

Ganglmair, B., W. K. Robinson, and M. Seeligson (2022): "The Rise of Process Claims: Evidence from a Century of U.S. Patents," Unpublished manuscript, available at https://papers.ssrn.com/abstract=4069994.

Graham, S. J. H. and S. Vishnubhakat (2013): "Of Smart Phone Wars and Software Patents," *Journal of Economic Perspectives*, 27, 67–86.

Hall, B. H., A. B. Jaffe, and M. Trajtenberg (2001): "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," NBER Working Paper 8498, National Bureau of Economic Research, Cambridge, Mass.

Kuhn, J. M. and N. C. Thompson (2019): "How to Measure and Draw Causal Inferences with Patent Scope," *International Journal of the Economics of Business*, 26, 5–38.

Lanjouw, J. O. and M. Schankerman (2004): "Patent Quality and Research Productivity: Measuring Innovation With Multiple Indicators," *Economic Journal*, 114, 441–465.

Lerner, J. (1994): "The Importance of Patent Scope: An Empirical Analysis," *RAND Journal of Economics*, 25, 319–333.

——— (2006): "The New Financial Thing: The Origins of Financial Innovations," *Journal of Financial Economics*, 79, 223–255.

Marco, A. C., J. D. Sarnoff, and C. deGrazia (2016): "Patent Claims and Patent Scope," USPTO Economic Working Paper 2016-04, United States Patent and Trademark Office, Office of the Chief Economist, Arlington, Va.

Pakes, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54, 755–784.

Rosenberg, M. D. (2012): *Patent Application Drafting: A Practical Guide*, New York, N.Y.: Oxford University Press.

Schankerman, M. and A. Pakes (1986): "Estimates of the Value of Patent Rights in European Countries During the Post-1950 Period," *Economic Journal*, 96, 1052–1076.

Schmoch, U. (2008): "Concept of a Technology Classification for Country Comparisons," Final report, World Intellectual Property Organization (WIPO), Geneva, Switzerland.

Trajtenberg, M., R. Henderson, and A. B. Jaffe (1997): "University Versus Corporate Patents: A Window on the Basicness of Invention," *Economics of Innovation and New Technology*, 5, 19–50.

von Graevenitz, G., S. Wagner, and D. Harhoff (2013): "Incidence and Growth of Patent Thickets: The Impact of Technological Opportunities and Complexity," *Journal of Industrial Economics*, 61, 521–563.

WIPO (2007): *WIPO Patent Drafting Manual*, IP Assets Management Series, Geneva, Switzerland: World Intellectual Property Organization.