// NO.19-035 | 02/2022

DISCUSSION PAPER

// BERNHARD GANGLMAIR AND
IMKE REIMERS

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

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

Bernhard Ganglmair† Imke Reimers‡

First version: August 7, 2019
This version: February 8, 2022

Abstract

Innovation depends on the incentives to create new ideas as well as the visibility of and access to existing ones. Using exogenous variation from the Uniform Trade Secrets Act, we show that stronger trade secrets protection has a disproportionately negative effect on patenting of less visible inventions (processes). We develop a framework of initial and follow-on innovation to determine the welfare effects of such shifts in disclosure. Stronger trade secrets may have negative effects on overall welfare by reducing opportunities for follow-on innovation, and optimal trade secrets policy depends on visibility, costs of R&D, and the value of cumulative innovation.

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, Paul Hünermund, Jeffrey Kuhn, Matt Marx, Mike Meurer, Markus Nagler, Lisa Ouellette, Ivan Png, David Schwartz, Neel Sukhatme, Mike Teodorescu, and others. Funding 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. Address: ZEW – Leibniz Centre for European Economic Research, L7, 1, 68161 Mannheim, Germany; E-mail: b.ganglmair@gmail.com.

‡Northeastern University, Department of Economics, 304 Lake Hall, 360 Huntington Ave, Boston, MA 02115, USA; E-mail: i.reimers@northeastern.edu.

"[S]ociety is giving something for nothing ... [when] concealable inventions remain concealed and only unconcealable inventions are patented."

Machlup and Penrose (1950)

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, especially in industries with technologies that are per se less visible or "self-disclosing" (Strandburg, 2004). In those industries, the diffusion of knowledge relies on the disclosure function of patents. Diffusion 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 and visibility of technology on patenting and 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

[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.

the disclosure of a trade secret without consent.[2] For example, in a well-publicized legal case, Waymo LLC (a self-driving car startup 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).

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 the decision to invest in different types of innovation, the disclosure of these inventions, and the ability to build on them. We use exogenous variation across states and time from the staggered adoption of the Uniform Trade Secrets Act (UTSA) of 1979/1985, which changed the strength of trade secrets protection in individual states, to identify the effects of trade secrets protection. Using the index of the strength of trade secrets protection introduced by Png (2017a) and new data on the type of a patented invention – product or process – to capture how visible or self-disclosing an invention is (Ganglmair et al., 2021), we show that stronger trade secrets protection results in a disproportionate decrease of process patents. Since patents provide insight into what is *not* kept secret, we interpret this change as a relative increase in the propensity to keep process inventions secret.[4] This, in turn, limits opportunities for follow-on innovation.

The welfare implications of such changes in intellectual property protection depend not only on the facilitation of follow-on innovation but also on the ex-ante incentives to innovate. To make inferences about these incentives, we develop a simple framework of sequential innovation. We illustrate insights from the model based on the causal estimates

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

[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, only 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.

from our descriptive analyses. We find that total welfare may in fact decline as trade secrets protection grows stronger. This is particularly the case when the costs of R&D are relatively small and stronger trade secrets protection does little to incentivize innovation. By contrast, stronger trade secrets protection can increase welfare when R&D is more costly as protection can lead to increased investment in initial R&D. Further, the effects are amplified when inventions are relatively invisible, because trade secrets are most attractive for them.

We provide more institutional background, including details about the UTSA and a discussion of the disclosure function of patents, in Section 2. In Section 3, we develop a simple model of an inventor's decision to disclose a new invention through a patent. Among other factors, the value of the invention from a patent increases with the underlying invention's visibility: Visibility allows for easier enforcement of the patent, thus guaranteeing exclusive access to the technology. By contrast, the value of an invention that is kept secret decreases in its visibility, because secrecy (and therefore exclusive access) is more difficult to maintain. Our model predicts that, for a given baseline share of (less visible) process *inventions*, the share of process *patents* decreases as trade secrets protection strengthens. This theoretical prediction serves as the basis for the empirical analysis in the rest of the paper.

In Section 5, we use the patent-level data introduced in Section 4 and the staggered adoption of the UTSA across states – which provides exogenous variation in the strength of trade secrets protection – to estimate the effect of stronger trade secrets protection on the likelihood that a patent covers a process innovation. Consistent with results from our theoretical model, we find that stronger legal protection of trade secrets leads to a disproportionate decrease in the patenting of processes. 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, and they are indicative of potentially large implications for follow-on innovation and welfare.

We examine these implications in Section 6. We add R&D decisions and follow-on innovation in a sequential innovation model to our simple disclosure framework from Section 3. Counterfactual simulations show that the optimal level of trade secrets protection is increasing in the costs of R&D. When costs are low, stronger legal protection of trade secrets has little effect on initial R&D but carries the unintended consequence of impeding follow-on innovation. On the other hand, for R&D projects that are relatively more costly, stronger legal protection improves welfare by encouraging initial R&D. Applying the model to our causal results, we further show that both positive and negative effects of trade secrets protection are more pronounced for processes than for products, and that the optimal level of trade secrets protection is lower when follow-on innovation carries more value.

Beyond studies based on survey data, there is limited empirical work 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 (2019) 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. 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. (2020) 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-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 accounting for the role of an invention's visibility in patenting and innovation decisions and in providing opportunities for follow-on innovation. We further highlight that insights gained from the effects of patents on innovation do not necessarily apply to trade secrets. This is particularly important in light of the U.S. Defend Trade Secrets Act and the EU Trade Secrets Directive 2016/943, for which impact evaluations are just now starting to become available. Results from the UTSA can thus inform an ongoing policy debate in both the U.S. and in Europe.

2 Institutional Background

2.1 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).

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 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 (2017a) constructs an annual index that measures the strength of legal trade secrets protection at the state level for the years 1976 to 2008. For example, the changes in Virginia represent increases in two of the six inputs into the index.[5] On average, the UTSA implied a rise in the index of 42 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 a handful of states with no change. There is no obvious

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

pattern in the size of these changes over time and across states, and Png (2017a) cites anecdotal evidence that suggests that passing of the bills often happened for "whimsical" reasons.

2.2 Trade Secrets and the Disclosure Function of Patents

By using the UTSA to examine the effects of trade secrets protection on follow-on innovation and welfare, 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). 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 to keep trade secrets.[6]

The second premise is that patents provide some disclosure of inventions. 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 (Roin, 2005; Fromer, 2009; Devlin, 2009; Seymore, 2009). In addition, Lemley (2008a) suggests that researchers 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, law 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 (Gross, 2019; Hegde et al., 2020; Furman et al., 2021). Importantly, while in our model we assume perfect disclosure for tractability reasons, our results hold as long as patents provdide

[6] For a comprehensive survey of trade secrets litigation in federal and state courts, see Almeling et al. (2010a,b).

some disclosure.

3 A Model of Trade Secrets and Disclosure

In this section, we consider an inventor's decision whether to disclose a (patentable) invention through a patent or to keep the invention a secret.[7] This decision is embedded (as Stage 2) in a three-stage sequential model, where Stage 1 describes the inventor's decision to invest in R&D and realize the initial invention, Stage 2 describes the disclosure decision, and Stage 3 captures the market's engagement in follow-on innovation. For the patenting decision at Stage 2, we focus on the roles of trade secrets and an invention's visibility. The model derives predictions for the analyses in Section 5, where we examine the empirical relationships between trade secrets protection, visibility, and patenting. We return to the full three-stage model in Section 6.

3.1 An Inventor's Decision to Disclose

An invention i at Stage 2 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). We discuss each of the invention's characteristics below.

An inventor is given the choice to disclose an invention through a patent ($\tilde{d} = D$) or keep the invention secret ($\tilde{d} = S$).[8] 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

[7] Given that we use patent data for our empirical analysis, we restrict our model interpretation to inventions that are patentable. In the U.S., this means they must exhibit patentable subject matter (35 U.S.C. §101), be useful (35 U.S.C. §101), novel (35 U.S.C. §102), and non-obvious (35 U.S.C. §103).

[8] 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.

(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, $\phi_D \geq \phi_S$. A simple argument for this is that the inventor herself knows what to look for, whereas firm B has little prior guidance. 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.[9] We can write the expected commercial value the inventor is able to materialize as $\phi_D v = \phi v$. In addition, the inventor receives a patent premium λ.[10] 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 $V_D(\phi) = \phi\,(1 + \lambda)\,v$.

Whereas disclosure-visibility is important to determine the value of a *patent*, 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 $\tau \in [0, 1]$. Recall that trade secrets laws provide protection against misappropriation of trade secrets but not against simple copying. This means that, 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.[11] Collecting terms, we define the private value from secrecy as $V_S(\phi, \tau) = \tau\,(1 - \xi\phi)\,v$.

An implicit assumption in V_S is that the secret holder can detect misappropriation,

[9] Active monitoring of infringement is said to be a major source of the costs of patent enforcement (Hall et al., 2014). Goldstein (2013:64) writes: "A patent claim whose infringement is very hard to discover is a claim with low or no value."

[10] 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; Sichelman and Graham, 2010), 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.

[11] 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.

and this ability is independent of the underlying visibility of the technology. A positive probability of detection is consistent with empirical evidence: Many instances of trade secrets litigation involve former employees or business partners stealing the secret holder's information (Almeling et al., 2010a,b). The Waymo case described in the introduction provides one prominent such example. For tractability, we set the probability of detection equal to unity, so that the only variation in the enforcement of trade secrets is through τ.[12]

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

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

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. Observe that in our model, the inventor's decision to patent an invention is not a function of the potential commercial value of the invention but rather of the *effective* value (given the invention's visibility).[13] 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 (Moser, 2012), and she is less likely to patent as the degree of trade secrets protection τ increases (Png, 2017b).

3.2 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.[14] The (expected)

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

[13] While the theoretical literature is divided (e.g., Anton and Yao, 2004; Jansen, 2011), most empirical studies find a positive relationship between the value of an invention and the propensity to patent (e.g., Moser, 2012; Sampat and Williams, 2018).

[14] 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).

values of secrecy $EV_{S|\Theta}(\tau)$ and disclosure $EV_{D|\Theta}(\tau)$ of an invention of type Θ are

$$EV_{S|\Theta}(\tau) = \int_{0}^{1} \tau\,(1-\xi\phi)\,v dG_{\Theta} \qquad \text{and} \qquad EV_{D|\Theta}(\tau) = \int_{0}^{1} \phi\,(1+\lambda)\,v dG_{\Theta}. \tag{2}$$

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.*

The proofs of this and all other results are relegated to Appendix Section A.1. 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; Arundel, 2001; Hall et al., 2013).

3.3 Probability of Disclosure for Invention Types

For our main theoretical result and prediction, we derive the probability ρ that a given patent covers a process invention. We first establish two auxiliary results. 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

$$d_{\Theta}(\tau) = \int_{0}^{1} d(\phi,\tau) dG_{\Theta}(\phi) = \int_{\bar{\phi}(\tau)}^{1} 1 \cdot dG_{\Theta}(\phi) = 1 - G_{\Theta}(\bar{\phi}(\tau)). \tag{3}$$

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

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.*

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:

$$\rho(\tau) = \frac{\theta d_M(\tau)}{\theta d_M(\tau) + (1-\theta)\, d_P(\tau)}. \tag{4}$$

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

The expression in Equation (4) 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. In the next two sections, we examine this prediction empirically.

4 Patent Data

In our empirical analyses, we estimate the probability that a patent includes a process innovation as a function of the trade secrets protection index described in Section 2, for patents with priority dates between 1976 and 2008 – the years for which we have trade secrets protection data. To do this, we (a) match a set of patents to the relevant level of trade secrets protection by identifying 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. We supplement these data with additional patent characteristics.

4.1 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.[15] 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.

[15] For continuations, the applicant may not add new disclosures but may delete claims. Divisions involve separating an earlier patent application into two or more.

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.[16] 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.[17]

4.2 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. (2021).[18] A claim can be of one of three distinct types: (1) process (or method) claims describe the sequence of steps which 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.[19]

[16] 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."

[17] Our estimation sample slightly over-represents individual applicants and under-represents large firms. We document this selection in Section A.2, and we show below that this selection does not drive our results.

[18] 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. (2021) 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.

[19] 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.

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.

The top portion of Table 1 provides summary statistics for our patent-type indicators for all granted USPTO utility patents in our sample.[20] 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.

4.3 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 (see Kuhn and Thompson, 2019), where shorter claims are likely broader. 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.

[20]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).

(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.[21] 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$^{\text{th}}$ year of the patent term (see, e.g., Pakes, 1986; Schankerman and Pakes, 1986).

5 Empirical Estimation and Results

5.1 The Impact of Trade Secrets Protection

We take advantage of the staggered adoption of the UTSA across U.S. states to estimate the likelihood that a patent includes a process (Proposition 2). In our main specification, we estimate the probability that a patent covers a process invention as a function of the patent's characteristics and the state's trade secrets protection index. Formally, we estimate

$$process_{j,cst} = \beta_1\, protection_{st} + \beta_2 X_j + \nu_{cs} + \mu_{ct} + \epsilon_j, \tag{5}$$

where the dependent variable is an indicator that is 1 if patent j, which belongs to USPC main class c and is filed by an entity in state s in year t, is a process patent; $protection_{st}$ is the trade secrets protection index in state s and year t. To control for trends within a USPC class that are common to all states and for any USPC class-state-specific characteristics that do not vary over time, we include interacted fixed effects for the patent's USPC class and priority year (μ_{ct}), as well as for its USPC class and

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

location state (ν_{cs}), respectively.[22] 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 cluster standard errors at the state-year level, and we estimate Equation (5) as a linear probability model, noting that logit estimations provide qualitatively identical results.

5.2 Baseline Results

Table 2 shows the coefficients from the baseline specifications, including different sets of control variables. 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 including 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 of 42.3% before UTSA adoption, and with a mean increase in the trade secrets protection index of 0.36 points across all patents, our results correspond to respective mean decreases of 2.2% and 1.5% in the probability that a patent is a process patent when a state adopts the UTSA. As we explain below, we can interpret these results as lower bounds of the effects on the disclosure decisions.

Our results foreshadow large potential effects of changes in trade secrets protection on follow-on innovation. We find that stronger protection leads to a disproportionate decrease in the disclosure of innovations that are inherently less visible and therefore rely on patents to facilitate follow-on innovation. In what follows, we first provide more detail

[22] 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). We cannot add state-year fixed effects because they are perfectly collinear with our variable of interest.

Table 2: Baseline Results – 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)
State FE	Yes	Yes	Yes	Yes	No
Year FE	Yes	Yes	Yes	Yes	No
USPC Mainclass FE	Yes	Yes	Yes	Yes	No
State/Year × USPC Mainclass FE	No	No	No	No	Yes
$\overline{R^2}$	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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

and test the robustness of this main finding, and we then develop a model that includes all stages of innovation – initial investment, the disclosure decision, and follow-on innovation – to provide intuition as well as estimates for the welfare effects of these law changes.

5.3 Timing of the Effect

To test whether the negative effect on process patenting we found above sets in immediately or gradually, we estimate the annual changes in the probability that a patent includes a process innovation relative to the year before a state's UTSA adoption. Specifically, we estimate

$$process_{j,cst} = \sum_z \beta_z 1(z)_{st} + \nu_{cs} + \mu_{ct} + \epsilon_j, \tag{6}$$

where ν_{cs} and μ_{ct} represent USPC class-state fixed effects and USPC class-year fixed effects, respectively, and $1(z)_{st}$ represents the "lag," or the number of years since patent j's filing state s adopted the UTSA. We choose the year before UTSA adoption as the omitted category. This analysis deviates from our main analysis in two ways. First, rather than an index of trade secrets protection, our explanatory variable of interest is a binary variable. Second, the analysis drops patents from states that did not adopt the UTSA because the lag variable is not clearly identified. As above, we estimate the regression as a linear probability model for ease of interpretation, and we cluster standard errors at the state-year level.

Figure 1 illustrates the results from this specification. Two takeaways are clear from the analysis. First, although almost all 95% confidence intervals include zero, there might be a slight decrease in the relative probability of a process patent in the years before UTSA adoption. Inspired by these patterns, we examine the validity of our research design in Section 5.5. Second, although the yearly coefficients are not statistically significant, they suggest an immediate and lasting negative effect of the UTSA on the probability of a process patent.[23] Given the average increase in the trade secrets protection index when the UTSA was adopted, the sizes of the coefficients are in line with those in Table 2.

5.4 Heterogeneous Effects

We examine the roles of firm sizes and technology complexities in Table 3. In the first two columns, we repeat the estimation from Columns (4) and (5) of Table 2, interacting the trade secrets index with indicators for small (individuals and small firms) and large

[23] Our setting differs from traditional difference-in-differences analyses because we only observe each patent once and our variable of interest is an index rather than a binary treatment variable. Still, the result in Figure 1, along with unreported regressions that show the estimated effects do not vary significantly across treatment cohorts, alleviates concerns about potential bias from weighting treatment effects across the staggered cohorts (Goodman-Bacon, 2021).

Figure 1: Yearly Effect of UTSA Adoption

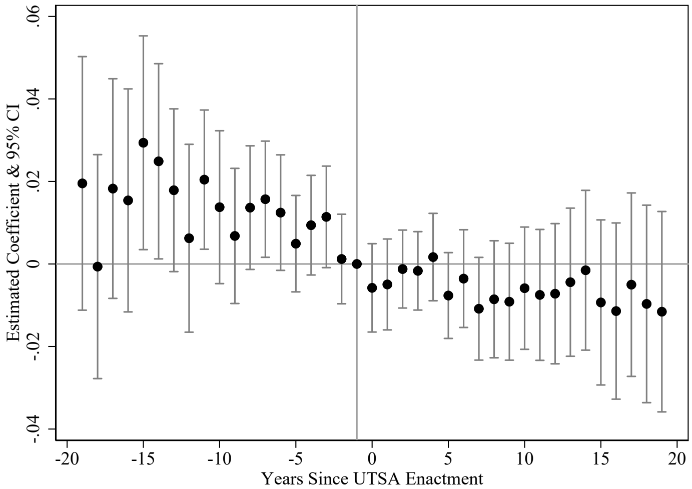

Notes: This figure presents coefficients on indicator variables that equal 1 in each year before and after UTSA adoption, as in Equation (6), for 20 years before and after adoption. The vertical line shows the year before UTSA adoption.

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% (compared to an estimated average effect of 1.5%). The (negative) coefficient is smaller, and less statistically significant, for large applicants.

Our findings confirm our expectations and thus provide support for our empirical design, 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 on patenting. Third, findings by Crass et al. (2019) suggest a stronger degree of substitutability between secrecy and patents for small applicants, which should in turn yield a stronger effect of trade secrets protection.

In the next two columns of Table 3, We allow the effects to vary between "complex" and "discrete" technologies to explore this 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

Table 3: 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)
State FE	Yes	No	Yes	No
Year FE	Yes	No	Yes	No
USPC Mainclass FE	Yes	No	Yes	No
State/Year × USPC Mainclass FE	No	Yes	No	Yes
$\overline{R}^2$	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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

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 von Graevenitz et al. (2013).[24] Interacting this indicator with the trade secrets index in our main specification (Column (4) of Table 3), we find that the probability of a process patent decreases by 2.7% at the baseline (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).

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

5.5 Discussion of Identification

Our identification strategy relies on two assumptions. First, the relative number of process and product *inventions* (rather than patents) does not vary systematically in response to the implementation of the UTSA. Second, the adoption of the UTSA is not affected by an expectation that certain types of inventions will be more prevalent in the future.

5.5.1 Innovation of Products and Processes

Png (2017a) shows that investment in R&D increases when trade secrets protection becomes stronger, which could change the pool of realized inventions.[25] While we are unaware of empirical evidence, it is likely that investment in process inventions is affected disproportionately, because less visible inventions benefit the most from secrecy. Thus, if a strengthening of trade secrets protection affected the *creation* of different types of inventions differently, then stronger trade secrets protection would likely lead to a relative *increase* in process patents absent changes in patenting behavior of existing inventions. That is, because we observe a relative *decrease*, our results can be interpreted as a lower bound of the effect of trade secrets protection on the patenting decision.

5.5.2 Placebo Tests

Although Png (2017a) 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 examine this possibility in a set of placebo tests. Instead of the *true* UTSA adoption date for each state, we set an earlier

[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 (2017b) suggests an instrumental variables approach, which we adopt in an additional check. Like in his paper, we find even stronger results in that specification. We continue without instruments to provide more conservative estimates, noting that all qualitative results hold if we include the instruments.

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 2. For all four placebo adoption dates, the coefficient of 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.

5.5.3 Randomizing Treatment

The negative coefficients in our main analyses could be the result merely of chance. We examine this possibility in a perturbation test similar to DeAngelo et al. (2017) and Alsan et al. (2019). We randomize the timing of UTSA adoption across states based on the true distribution of adoption dates. Then, we randomly assign these dates to the U.S. states and estimate the impact of this pseudo-adoption on the probability that a patent includes a process claim. We record the coefficient of interest – on an indicator variable that is 1 after pseudo-adoption and 0 before, rather than the numerical index – for 1000 such permutations. Using the true adoption dates, the coefficient of interest is -0.0109. Of the 1000 coefficients from the permutations, only 3 are more negative, suggesting that the UTSA indeed affected patenting of products and processes differently.

5.6 Robustness Analyses

Our data construction and empirical approach are based on a number of assumptions, which we examine in Appendix Section A.3. In short, we find that our results are robust. First, we vary the timing of the disclosure decision. Instead of assigning a patent's priority date, we use each patent's application date. We also limit our sample to the first patents in a patent family. Second, we examine our sample restriction to single-state patents.

[27] We also drop all patents that were applied for more than ten years before the state's true UTSA adoption to create a closer comparison group.

We consider both a broader definition of patent location (based on the first U.S.-based assignee) and two narrower definitions (limiting the analysis to U.S.-only patents and to single-assignee patents). Third, we examine our definition of a process patent. We consider two less stringent definitions based on individual claims, and we drop software patents. Fourth, we include control variables for state-level changes in the enforcement of the Inevitable Disclosure Doctrine, which may have affected patenting decisions beyond the effects of the UTSA. Fifth, we include state-specific linear time trends before UTSA adoption to account for possible time-varying differences across states. Finally, we also repeat our analysis after separately dropping each U.S. state to examine whether the effects are driven by changes in individual states. We do not find any evidence of this.

5.7 The Number of Patents

Lemma 2 states that a strengthening of trade secrets protection should lower the probability that any invention is patented, regardless of its type. We examine this claim here, by estimating the effects of the UTSA on the log-number of process and product patents. To control for changes in patterns across technologies, we create a balanced panel at the USPC class-state-year level, and we estimate equations such as

$$\ln(patents + 1)_{cst} = \beta_1 protection_{st} + \nu_{cs} + \mu_{ct} + \epsilon_{cst}. \tag{7}$$

Here, $patents_{cst}$ is the number of (process, product or total) patents in USPC main class c in state s and year t, and all other variables are as in our main analyses. We weight all observations according to the total number of patents in the state and USPC class in 1979 – one year before any state adopted the UTSA – and we impute zeros and unit weights for those state-year-classes without any patents.

Table 4 shows the results from these regressions. The first three columns include separate fixed effects for USPC classes, states, and years; and the last three columns include the interacted fixed effects – our preferred specification. Across the specifications, we find large and significant decreases in the number of process patents, between 11 and

Table 4: Impact of Trade Secrets Protection on Total Patenting

	Separate FEs			Interacted FEs		
	(1) Process	(2) Product	(3) All	(4) Process	(5) Product	(6) All
Trade secrets protection	-0.175***	-0.021	-0.096**	-0.113***	-0.003	-0.040
	(0.038)	(0.037)	(0.039)	(0.034)	(0.035)	(0.036)
$\overline{R}^2$	0.619	0.606	0.618	0.815	0.812	0.846
N	241,028	241,028	241,028	238,693	238,693	238,693

Notes: Log-OLS regressions at the USPC class-state-year level. The dependent variable is $\ln(patents + 1)$, where *patents* describes either the number of process patents, the number of product patents, or the total number of patents. The regressions are weighted based on the total number of patents in 1979 (the last year before the UTSA) in the respective USPC class, state and year. We infer zeros for class-state-years without any patents and assign these a frequency weight of 1. Columns (1) through (3) include state, year and USPC class fixed effects. Columns (4) through (6) interact USPC class fixed effects with state and year dummies, respectively. Robust standard errors, clustered by state and year, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

16%. By contrast, we find no significant effect on the number of product patents. On average, a one-unit increase in the trade secrets protection index decreases the number of patents by up to 10%.

6 Effects on Cumulative Innovation

In addition to affecting an inventor's patenting decision, strengthening trade secrets protection can incentivize investment in initial R&D, but it may also retard knowledge diffusion by reducing disclosure of less visible inventions. In what follows, we first introduce a three-stage model of sequential innovation that endogenizes an inventor's initial R&D decision (Stage 1) and accounts for the effect of the inventor's disclosure decision (Stage 2) on the intensity of follow-on innovation (Stage 3).[28] We then illustrate the implications of changes in trade secrecy protection, and we calibrate the model using the results from Section 5 and values of follow-on innovation from the literature to examine the role of an invention's visibility.

[28] Our model of follow-on innovation is simple but nonetheless consistent with stylized facts and other models proposed in the literature. See our discussion in Section A.4.1.

6.1 An Augmented Model of Cumulative Innovation

6.1.1 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. We assume that disclosure-visibility and secrecy-visibility are the same (so that $\xi = 1$ from Section 3). 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 and forms expectations of the invention's commercial value v_i based on a known distribution. She undertakes the R&D project if the expected payoffs from the invention (including the value and licensing revenues from both the invention and potential follow-on innovation) outweigh its cost. We refer to F_Θ and the distribution of invention types as *unconditional* distributions, that means, *before* the R&D decision is taken.

6.1.2 Stage 2 (Patent or Trade Secret)

The second stage of our augmented model is the disclosure model in Section 3. 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, where ϕ_i is drawn from the invention type specific *conditional* distribution of realized inventions with cdf G_Θ (*after* the R&D decision).

6.1.3 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} and cost C_{i_F}.[29] 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

[29]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)) triggered by invention i.

measure as *effective visibility* of initial invention i and denote it by $\tilde{\phi}_i$. It is equal to

$$\tilde{\phi}_i = \begin{cases} \phi_i & \text{if R\&D in Stage 1 and } \textit{trade secret} \text{ in Stage 2;} \\ 1 & \text{if R\&D in Stage 1 and } \textit{patent} \text{ in Stage 2.} \end{cases} \tag{8}$$

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.[30]

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 $\tilde{d} \in \{S, D\}$.

6.2 Surplus Implications

To provide some intuition from this model, we assume (1) that the unconditional visibilities (of potential inventions) are uniformly distributed between 0 and 1, (2) that the value of any (initial or follow-on) innovation is exponentially distributed with rate parameter 0.1, (3) that the patent premium λ equals 0.1 (in line with Schankerman, 1998), and (4) that 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 patented inventions (so that $\tilde{\psi}_D = 2/3$). We simulate the R&D and disclosure decisions, along with the follow-on innovation outcomes, given R&D costs drawn from a logistic distribution with variance 0.5. We choose three separate location parameters for the cost distributions to describe industries with varying R&D costs: no costs, low average costs (40% of the expected value of the R&D investment), and high average costs (80%).

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

Given simulated investment and disclosure decisions for 1,000,000 potential inventions, we calculate "welfare" $\overline{W}(\tau)$, a function of trade secrets protection, as the expected social surplus derived from the use of the invention.[31] 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 ϕ, whereas for inventions kept as trade secrets, the barriers to access decrease in visibility and increase in trade secrets protection, τ. For follow-on innovation, we assume free access and zero deadweight loss in both disclosure states. To calculate the expected surplus, we account for the inventor's decision to develop any potential inventions at Stage 1 and the probability of follow-on innovation at Stage 3.

Figure 2 illustrates how welfare from innovation varies with the trade secrets protection level τ, for each R&D cost level. The graphs on the left depict welfare from all innovation. For no R&D costs (Panel (a)), stronger trade secrets protection has an unambiguously negative effect on total welfare. But as R&D costs increase, stronger trade secrets protection can increase welfare. The right-hand side of Figure 2 separately depicts the surplus associated with initial R&D and with follow-on innovation to illustrate the channels that affect welfare.

Deadweight Loss from Monopoly Power: Stronger trade secrets protection increases barriers to access to a technology, which increases the deadweight loss from monopoly power. We show this effect in the solid line (for initial innovation) in Figure (*ii*) of Panel (a), which isolates this deadweight loss because, without R&D costs, all R&D projects are realized.

Decision to Innovate: When costs are nonzero, trade secrets protection has a positive effect on initial R&D by increasing the expected value of realized R&D projects. This in turn has a positive effect on welfare from innovation. We observe this effect in the solid-line graphs in Figures (*ii*) of Panels (b) and (c).

[31] We provide more details on our welfare measure in Section A.5.

Figure 2: Effect of Trade Secrets Protection on Welfare

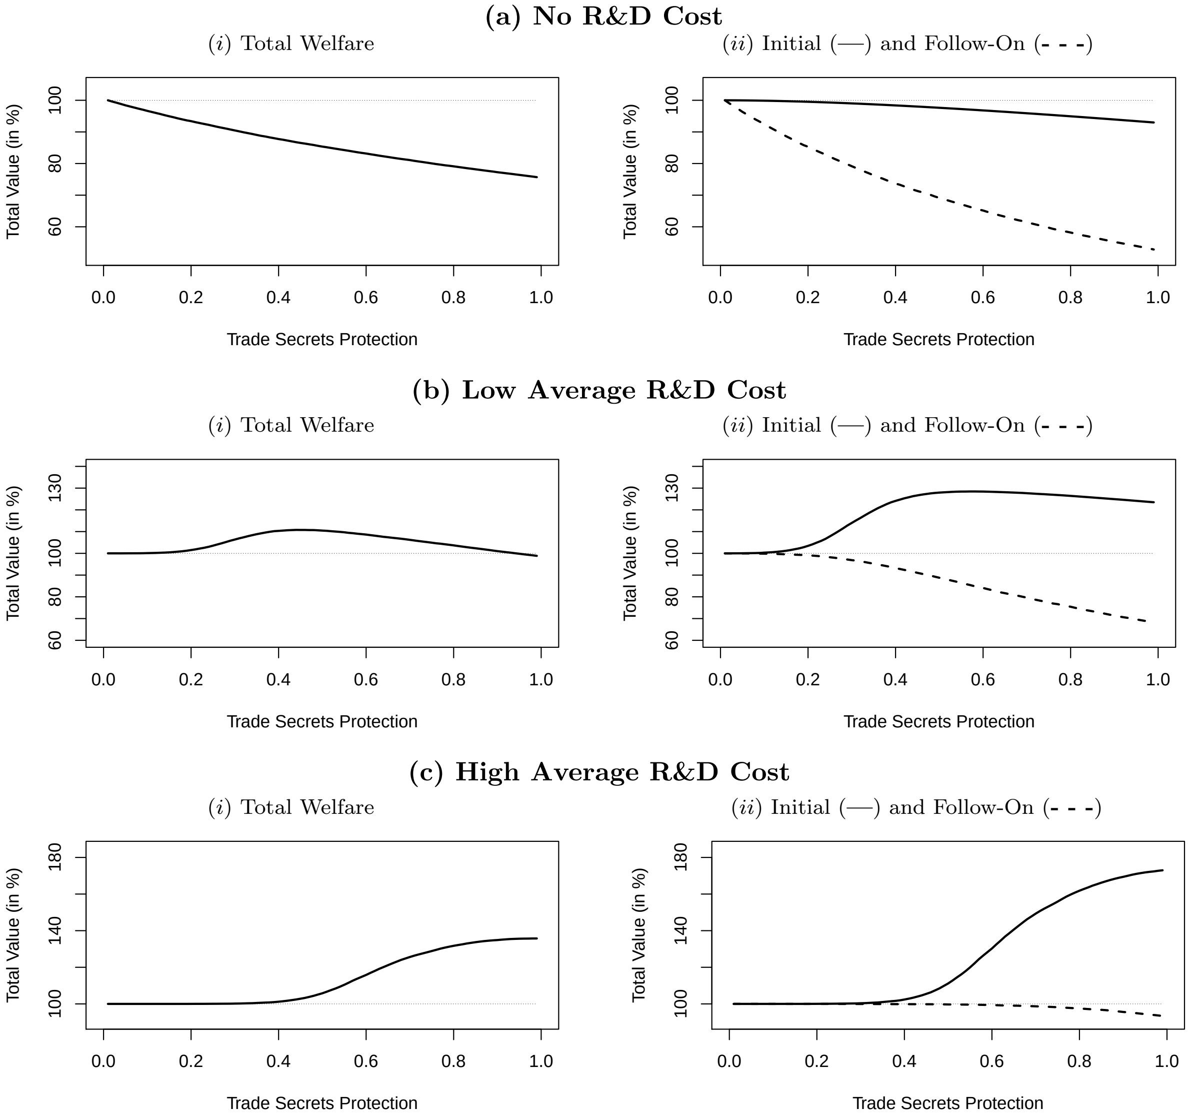

Notes: This figure plots welfare $\overline{W}(\tau)$ (in % of the value when $\tau = 0$) for values of $\tau \in [0, 1]$, when there are no R&D costs (Panel (a)), with low average R&D costs (40% of the expected R&D project value, Panel (b)), and with high average R&D costs (80%, Panel (c)). In each scenario, we simulate a sample of N = 1,000,000 inventions, assuming a uniform unconditional visibility distribution and baseline success probabilities of $\psi_S = 1$ and $\psi_D = 2/3$. In the left Panel (i), we show the aggregate value of both initial and follow-on innovation ($\overline{W}(\tau)$). In the figures on the right (ii), we separately plot the social value of initial R&D (solid) and follow-on innovation (dashed).

Follow-on Innovation (1): Stronger trade secrets protection affects welfare by lowering the share of inventions that are disclosed. This has a negative effect on overall welfare through $\tilde{\psi}_{i_F,\tilde{d}}$: effective visibility decreases, which in return reduces the success probability of follow-on innovation. We observe this negative effect of trade secrets protection in the dashed graphs on the right-hand side figures.

Follow-on Innovation (2): Stronger trade secrets protection has a secondary effect on follow-on innovation. The increased ex-ante R&D activity implies there is more initial

R&D to build on. This counteracts the negative effect of trade secrets on follow-on innovation from reduced disclosure, especially when R&D costs are high. To observe this, compare the dashed graphs in Figures (*ii*) for the value of follow-on innovation for Panels (b) and (c). For higher costs (Panel (c)), trade secrets protection has a stronger incentivizing effect on initial R&D. As a consequence, the decrease in the value of follow-on innovation is smaller here than for low costs.

Finally, observe from the locations of the maxima in the left graphs of Figure 2 that the optimal level of trade secrets protection increases in R&D costs. This rationalizes existing law and practice, which tends to provide stronger protection for higher-cost projects. In the State of New York (that has not adopted the UTSA but follows common law principles) one factor to determine whether something is a trade secret explicitly lists the costs of developing the information.[32] 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).

6.3 The Role of Visibility

The above results suggest important welfare effects of changes in trade secrets protection across all inventions, and these effects may be particularly strong among less visible inventions. Here, we study the role of visibility, using the invention's type as a proxy. To meaningfully distinguish the effects on process and product innovations, we need to obtain the unconditional visibility distributions F_Θ of both process and product innovations. In what follows, we briefly explain our estimation and simulation strategy and the set of assumptions we make for identification.[33] We then use the estimated values to illustrate the role of an invention's visibility in informing the effects of trade secrets protection.

[32]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 (Sandeen and Rowe, 2013:19).

[33]We provide a formal description of the estimation details, including a table summarizing the ingredients of each step, in Appendix Section A.4.

6.3.1 Estimation and Simulation Strategy

We proceed in four steps. In an initial step (Step 0), we predict the type distribution of disclosed inventions (process and product patents) based on our estimates in Section 5. Step 1 describes Stage 2 of the augmented model, in which we use maximum likelihood estimation to estimate the shares and distributions of the conditional (realized) inventions that explain our predicted patent shares. In Step 2, we simulate – given initial R&D costs – the distributions of unconditional inventions (ideas) that give rise to these estimated conditional distributions. We do this by matching simulated moments of the distributions of visibilities and invention types with those estimated in Step 1. Finally, in Step 3 we use the estimated conditional distributions to simulate follow-on innovation as we did in Section 6.2.

Step 1: To calculate the log-likelihood of the observed visibility distributions G_Θ and the distribution of invention types with θ, we use the predicted share of process patents for $\tau \in [0, 1]$ based on the results for discrete technologies from Column (4) of Table 3, and we make two main assumptions. First, as before, we set the patent premium $\lambda = 0.1$.[34] Second, the type-specific visibilities ϕ follow triangular distributions. We hold the mode for the distribution for products constant at 0.5 and estimate the distribution for processes without imposing hazard rate dominance.

The estimated parameters comport with our theoretical predictions. The estimated mode of the triangular distribution for processes is lower than the one we fixed for products (0.376 vs. 0.5).[35] Consequently, patenting probabilities for processes are lower than for products (Lemma 1) and decreasing in τ (Lemma 2), and the share of process patents decreases as trade secrets protection increases (Proposition 2).

Step 2: We estimate the unconditional distributions of invention types and visibilities through simulated method of moments. Specifically, we find the unconditional visibility and type distributions for which the moments of the simulated conditional distributions

[34] Our results are consistent for $\lambda \geq 0$.

[35] The mode is estimated very precisely, based on 1000 bootstrap replications. The triangular distribution for visibilities of products likelihood-dominates, implying that it also hazard-rate dominates the triangular distribution for visibilities of processes, as is our distributional assumption in Section 3.

match the moments of the conditional distributions we estimated in Step 1.[36] These simulated conditional distributions describe all potential inventions that the inventor decides to develop at Stage 1. In addition to the potential invention's visibility, this R&D decision is also driven by its value v_i and the costs of R&D C_i (both of which are simulated), as well as the strength of trade secrets protection τ, which informs the disclosure decision at Stage 2. Like we did for the conditional distributions in Step 1, we assume that visibilities follow triangular distributions, but unlike in Step 1, we estimate these distributions for both invention types.

Step 3: Given our results from Steps 1 and 2, we simulate follow-on innovation as in the previous section. We assume a baseline success probability of follow-on innovation of $\psi_D = 2/3$ for patented Stage-1 inventions and $\psi_S = 1$ for secret Stage-1 inventions. The invention values v_i and v_{i_F} are independent draws from the same distribution, as are R&D costs C_i and C_{i_F}.

6.3.2 Estimation Results

We calibrate our model for no R&D costs, low average costs, and high average costs for $N = 1{,}000{,}000$ simulated potential inventions, where the invention values v_i and the R&D costs are drawn from the same distributions as above. For all three cost levels, the results continue to satisfy our distributional assumption of hazard-rate dominance. Moreover, for both invention types, we observe a selection of higher-visibility inventions into development in Stage 1. The distributions of costs and values further imply relatively large R&D investment probabilities – ranging from 0.70, or 70% of all possible inventions, for high R&D costs to 1, or 100%, without any costs – in Stage 1. In Stage 2, 67% (high costs) to 77% (no costs) of realized inventions are patented, with patenting probabilities for processes considerably lower than for products (from 64% and 70%, respectively, for high costs to 70% and 80% for no costs). These results for patenting probabilities are in line with survey evidence reported by Mansfield (1986) who finds that between 66% and 84% of patentable inventions are indeed patented. Finally, in Stage 3, up to one half

[36] We use the means and variances of the visibility distributions and the means of the invention-type distributions.

Figure 3: The Role of Visibility

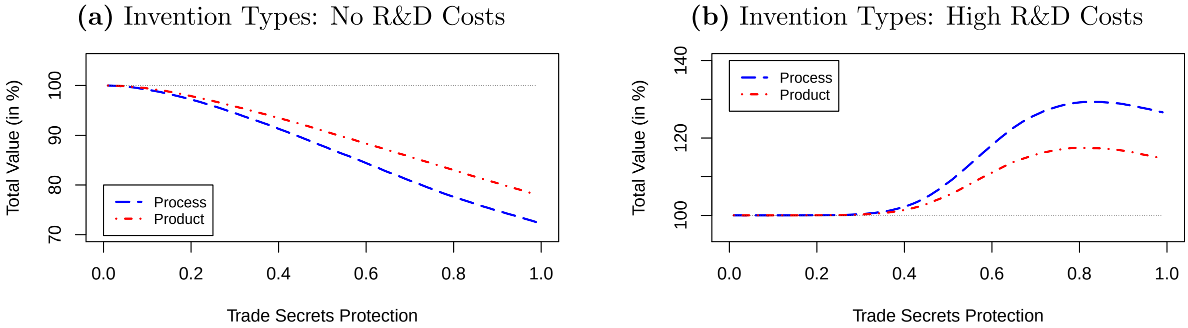

Notes: This figure plots the welfare function $\overline{W}(\tau)$ (in % of $W(0)$) for $\tau \in [0, 1]$. We simulate a sample of $N = 1{,}000{,}000$ inventions, using the estimates for unconditional distributions from Step 2 and assuming baseline success probabilities of $\psi_S = 1$ and $\psi_D = 2/3$. We plot the welfare functions separately for processes (dashed line) and for products (dash-dotted line). In Panel (a), we use the estimates for no R&D costs; in Panel (b), we use the estimates for high average R&D costs (such that costs are 80% of the expected R&D project value).

of all realized initial inventions lead to follow-on innovation (with the share decreasing in R&D costs). We see slightly lower probabilities of follow-on innovation building on processes (28% to 49%) than products (30% to 53%).

6.3.3 Visibility and Surplus

Both positive and negative effects of trade secrets protection are amplified for less visible inventions (processes). The positive incentive effect is stronger for processes, because they are less likely patented and stronger trade secrets protection increases appropriability. As we formalize in Section 3 and show in Section 5, the negative disclosure effect is also stronger for processes. The probability that processes are disclosed in patents decreases relative to products, which jeopardizes their follow-on innovation disproportionately.

Figure 3 illustrates these findings. We isolate the negative disclosure effect in the no-cost environment in Panel (a), in which ex-ante incentive effects do not play a role. As R&D costs increase, the positive effects on ex-ante incentives become more important. However, in the high-cost environment in Panel (b), these ex-ante incentive effects more than offset the negative disclosure effect.

6.4 The Role of Follow-on Innovation

In Figure 4, we compare, in a low cost environment, the results from the baseline model (where initial and follow-on value follow the same distributions) with a scenario in which follow-on innovation is on average twice as valuable. Larger relative value weights of initial and follow-on innovation can be explained, for instance, by more frequent or more numerous improvements of a product in a quality ladder when follow-on innovation is the present discounted value of all steps on that ladder. We find that the positive effect of trade secrets protection on the total value of innovation is weaker when follow-on innovation is more valuable, implying that the negative disclosure effect on follow-on innovation outweighs the positive effect from the increased potential for follow-on innovation due to more initial R&D. That is, even in an environment where the positive effects on follow-on innovation are more pronounced, the net effects of trade secrets on follow-on innovation remain negative.

Consequently, the optimal value of trade secrets protection for the initial invention is lower in industries that are characterized by a relatively larger value of follow-on innovation. We can see this by the shift to the left of the peak of the welfare functions when the value of follow-on innovation is doubled (panel (a)). In Panel (b), we plot the optimal value of trade secrets protection against the relative weight of follow-on innovation and observe a monotone negative relationship. In industries characterized by a lot of cumulative innovation, the impacts on follow-on innovation are amplified. For reasonable cost assumptions, the net effects are negative, making trade secrets protection less valuable as cumulative innovation gains importance. Moreover, because the effects on disclosure and follow-on innovation are even stronger for processes (Figure 3), the optimal trade secrets protection level is likely even lower for those inventions.

7 Conclusion

The effects of intellectual property rights on incentives to innovate are relatively well-understood, but we know less about the differences between the effects on initial and

Figure 4: The Role of Follow-on Innovation

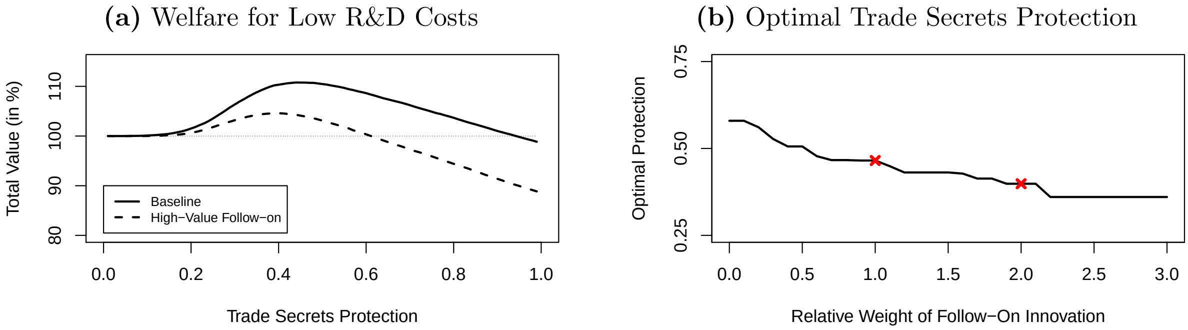

Notes: In Panel (a), we plot the welfare function $\overline{W}(\tau)$ (in % of $W(0)$) for $\tau \in [0,1]$ in the low-cost environment. We simulate a sample of $N = 1{,}000{,}000$ inventions, assuming a uniform unconditional visibility distribution and baseline success probabilities of $\psi_S = 1$ and $\psi_D = 2/3$. We plot the welfare function for two different values of follow-on innovation. The solid line describes the baseline from Figure 2 (where $v_{i_F} \sim \exp(1/10)$), the dashed line describes the scenario in which the value of follow-on innovation is drawn from $\exp(1/20)$ (twice the average value of the baseline). In Panel (b), we plot the optimal level of trade secrets protection for different values of follow-on innovation relative to initial ideas in a low-cost environment. We mark the baseline (=1) and the high-value scenario from Panel (a) (=2).

follow-on innovation. We add to recent discussions by arguing that the effects on follow-on innovation depend on both the intellectual property institutions and the visibility of the original idea. For highly visible inventions, patents limit the ability of others to build on said innovation. For inventions whose technology is less visible, however, trade secrets limit access entirely. For these inventions, patents can disclose information, which boosts the potential for follow-on innovation. Therefore, an intellectual property policy that particularly encourages patenting of less visible inventions could increase innovative activity as a whole.

The trade-off between the incentives to innovate and the ability of others to build on existing inventions also depends on the profitability of R&D investment. When R&D is relatively profitable (with low R&D costs), strengthening protection of a trade secret does little to incentivize additional investment in initial innovation, but it might discourage the disclosure of existing inventions. This hurts follow-on work, especially when the invention is not otherwise visible. On the other hand, when R&D is costly enough to prevent some innovation, a stronger trade secrets law could lead to more investment in initial R&D. If the increases in initial innovation are large, they could offset the losses from nondisclosure.

Our results support a body of literature that argues that an optimal policy distinguishes between different types of inventions and industries. 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 likely experience a welfare loss from stronger trade secrets protection.

Note that we specifically study secrecy of patentable inventions. A different, and broader, approach to trade secrets relates to the design of the employment relationship (e.g., in the form of covenants not to compete) or broader organizational concerns (such as in non-disclosure agreements). Given the mechanisms in our paper, we view our results as complementary to that literature.

References

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

Almeling, D. S., D. W. Snyder, M. Sapoznikow, W. E. McCollum, and J. Weader (2010a): "A Statistical Analysis of Trade Secret Litigation in Federal Courts," *Gonzaga Law Review*, 45, 291–334.

——— (2010b): "A Statistical Analysis of Trade Secret Litigation in State Courts," *Gonzaga Law Review*, 46, 58–101.

Alsan, M., O. Garrick, and G. Graziani (2019): "Does Diversity Matter for Health? Experimental Evidence from Oakland," *American Economic Review*, 109, 4071–4111.

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

Anton, J. J. and D. A. Yao (2004): "Little Patents and Big Secrets: Managing Intellectual Property," *RAND Journal of Economics*, 35, 1–22.

Arora, A. (1997): "Patents, Licensing, and Market Structure in Chemicals," *Research Policy*, 26, 391–403.

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

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

Castellaneta, F., R. Conti, and A. 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, 834–853.

Castellaneta, F., R. Conti, F. M. Veloso, and C. 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, 524–541.

Cohen, W. M., R. R. Nelson, and J. 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.

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

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

DeAngelo, G., B. R. Humphreys, and I. Reimers (2017): "Are Public and Private Enforcement Complements or Substitutes? Evidence from High Frequency Data," *Journal of Economic Behavior & Organization*, 141, 151–163.

Devlin, A. (2009): "The Misunderstood Function of Disclosure in Patent Law," *Harvard Journal of Law and Technology*, 23, 401–446.

Friedman, D. D., W. M. Landes, and R. A. Posner (1991): "Some Economics of Trade Secret Laws," *Journal of Economic Perspectives*, 5, 61–72.

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

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

Galasso, A. and M. Schankerman (2010): "Patent Thickets, Courts, and the Market for Innovation," *RAND Journal of Economics*, 41, 472–503.

Ganglmair, B., W. K. Robinson, and M. Seeligson (2021): "The Rise of Process Claims: Evidence from a Century of U.S. Patents," Unpublished manuscript.

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

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

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

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

Gross, D. P. (2019): "The Consequences of Invention Secrecy: Evidence from the USPTO Patent Secrecy Program in World War II," Unpublished manuscript, Harvard Business School.

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

Hall, B., C. Helmers, M. Rogers, and V. Sena (2014): "The Choice between Formal and Informal Intellectual Property: A Review," *Journal of Economic Literature*, 52, 375–423.

Hall, B. H., C. Helmers, M. Rogers, and V. Sena (2013): "The Importance (or Not) of Patents to UK Firms," *Oxford Economic Papers*, 65, 603–629.

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.

Hall, B. H. and R. 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, 101–128.

Hegde, D., K. Herkenhoff, and C. Zhu (2020): "Patent Disclosure and Innovation," Available at SSRN `https://ssrn.com/abstract=3158031`.

Hegde, D. and H. Luo (2018): "Patent Publication and the Market for Ideas," *Management Science*, 64, 652–672.

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

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

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

Jansen, J. (2011): "On Competition and the Strategic Management of Intellectual Property in Oligopoly," *Journal of Economics and Management Strategy*, 20, 1043–1072.

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

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.

Lemley, M. 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–.

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.

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

Machlup, F. and E. Penrose (1950): "The Patent Controversy in the Nineteenth Century," *Journal of Economic History*, 10, 1–29.

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

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

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

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

O'Donoghue, T., S. Scotchmer, and J. Thisse (1998): "Patent Breadth, Patent Life, and the Pace of Technological Progress," *Journal of Economics & Management Strategy*, 7, 1–32.

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

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

Png, I. P. (2012): "Trade Secrets, Non-Competes, and Inventor Mobility: Empirical Evidence," Unpublished manuscript, presented at DRUID 2012, available at `https://tinyurl.com/y58xd5z8`.

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

——— (2017b): "Secrecy and Patents: Theory and Evidence from the Uniform Trade Secrets Act," *Strategy Science*, 2, 176–193.

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

Sampat, B. and H. L. Williams (2018): "How do Patents Affect Follow-On Innovation? Evidence from the Human Genome," *American Economics Review*, 109, 203–236.

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

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

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.

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

Seymore, S. B. (2009): "The Teaching Function of Patents," *Notre Dame Law Review*, 85, 621–670.

Sichelman, T. and S. J. Graham (2010): "Patenting by Entrepreneurs: An Empirical Study," *Michigan Telecommunications and Technology Law Review*, 17, 111–180.

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

Trajtenberg, M., A. Jaffe, and R. Henderson (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.

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

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

A Appendix

A.1 Formal Proofs of Theoretical Results

Proof of Proposition 1: 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

$$\int u(x)dG_P(x) \geq \int u(x)dG_M(x) \iff G_P(x) \overset{FOSD}{\succ} G_M(x). \tag{A.1}$$

Integrating by parts, we obtain

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

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

$$\int_0^1 \underbrace{-\tau\,(1 - \xi\phi)\,v}_{u(\phi)}\,dG_P(\phi) \quad > \quad \int_0^1 \underbrace{-\tau\,(1 - \xi\phi)\,v}_{u(\phi)}\,dG_M(\phi) \tag{A.2}$$

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.

Proof of Lemma 1: 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.

Proof of Lemma 2: 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.

Proof of Proposition 2: 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

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

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.

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 5.4), our *average* treatment effects may be slightly overestimated.

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.

A.3 Robustness of the Empirical Results

The main analysis requires that we make several choices about variable definitions and the resulting sample selections. Here, we examine the robustness of our empirical results to these assumptions in additional regressions, replicating the specification from Column (5) of Table 2. All specifications show a robust negative impact of trade secrets protection on the share of process patents. We summarize all results in Table A.2.

Disclosure Date: We first assign the application date of each individual patent as the date of the disclosure decision (Column (1) of Panel (a)). The coefficient of interest remains strongly significant and almost identical to that in the main specification (-0.020 (se=0.006) instead of -0.018). Next, we circumvent the disclosure date issue altogether by considering only the patent family head – the first patent within its family. Again, the results are almost unchanged (Column (2) in Panel (a)).

Invention Location: We test the robustness of our results to the sample selection of single-state patents. In a less conservative approach, 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, we drop all patents with non-U.S. contributors, thus guaranteeing that the decision is made in the identified state. Again, both approaches provide almost identical results to the main specification (Columns (3) and (4) of Panel (a), respectively).

Decision Maker: Our focus on single-state patents also helps alleviate concerns about who makes the disclosure decision: if all assignees and inventors are located

Figure A.1: Applicant Size Distributions for Different Subsamples

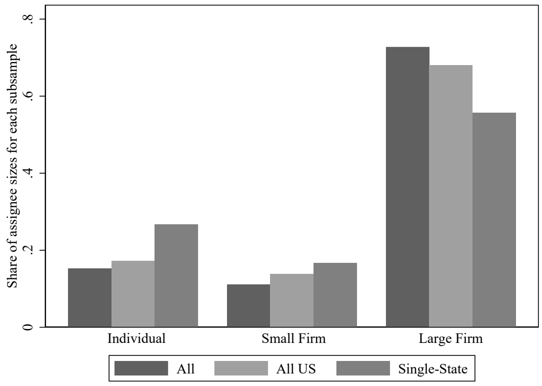

Notes: This figure presents shares of applicant sizes of different subsamples of all granted utility patents (between 1976 and 2014) with priority dates between 1976 and 2008. The darkest (leftmost) column shows statistics for all patents; the lightest (middle) column shows statistics for patents with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

in the same state, we know where the decision maker is located even if we do not know their identity. In another approach, we focus on patents with only one decision maker: those with just one assignee, or with just one inventor if no assignee is listed. The main result again remains almost unchanged (Column (5) in Panel (a)).

Definition of Process Patents: The main analysis defines all patents with at least one independent process claim as a process patent because we are interested in disclosure of any process component of the invention. Here, we use two alternative measures of a process patent: (1) a patent is a process patent if the *first* claim is a process claim,[37] and (2) a patent is a process patent if at least 50% of its independent claims are process claims. Our results are of similar magnitude to the main regression (Columns (1) and (2) of Panel (b)). Further, we drop all software patents, because software patents are often filed as process patents even though they do not inherently include process innovation.[38] The resulting coefficient on the trade secrets protection is almost identical as well (-0.017, se=0.006, Column (3)).

Other Changes in Relative Trade Secrets Strength: It is possible that the changes in protection due to the UTSA are correlated with other changes in the relative strength of trade secrets and patenting. Changes in the enforcement of the inevitable disclosure doctrine (IDD) pose one possible such confounding variable. The IDD is a common law doctrine that allows employers to seek protection against

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

[38] 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 A.2: 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.

the disclosure of trade secrets by a former employee when working for a competitor. Courts in the U.S. states have at various points in time either strengthened or weakened the judicial enforcement of the IDD.[39] Using the case coding from Castellaneta et al. (2016:Table B1), we control for changes in the judicial enforcement of the IDD (implying either a strengthening or weakening of employer-friendly trade secrets protection). Our results, reported in Column (4) of Panel (b), are again unchanged.

[39] For instance, in 1998 the Utah District Court found (in favor of the IDD) that "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). In 1999, the Fairfax County Circuit Court (in Virginia) found (against the IDD) that the "mere knowledge of a trade secret is insufficient to support an injunction order." Government Technology Services, Inc. v. Intellisys Technology Corp., 51 Va. Cir. 55 (Va. Cir. Ct., Oct. 20, 1999). Both cases cited in Castellaneta et al. (2016).

Accounting for Pre-Trends: Finally, the placebo tests in the main text suggest the share of process patents did not change in the years leading up to a state's UTSA adoption. Nevertheless, we add state-specific pre-trends to our difference-in-differences regression to account for the possibility that the shares of process patents were changing before UTSA adoption. The negative coefficient on the trade secret protection index is even stronger in this specification (coefficient=-0.029, se=0.013, Column (4) of Panel (b)).

A.4 Structural Estimation: Details

A.4.1 Modeling Follow-On Innovation: Discussion

Our 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.

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 where technologies are underused because patents on early ideas raise the costs of creating future ideas by introducing frictions in the bargaining process over licenses (Scotchmer, 1991; Galasso and Schankerman, 2010).

A.4.2 Estimation Steps

We summarize our estimation strategy in Table A.3, and we explain the details of our strategy for Steps 1 and 2 below.

Stage-2 Disclosure Decision (Step 1): We estimate the conditional distributions G_Θ and values for θ by maximizing the log-likelihood LL of the observed time-variant patent-type distribution. We observe two types of patents and use $\mathbf{M}_j \equiv \mathbf{M}_j(\Theta = M|\text{patent}) = 1$ to denote if a given patent j is a process patent, and $\mathbf{M}_j = 0$ if it is a product patent. Moreover, for each patent j, we observe the level

Table A.3: Estimation Strategy

Step 0	**Reduced form estimation of the share of process patents**	
Estimate	share of process patents	
Data	trade secrets protection index $\tau \in [0, 1]$	
Step 1	**ML estimation of conditional distributions (given realized R&D)**	
Estimate	conditional visibility distribution G_M for processes ($\phi \sim$ triangular(γ_M)) conditional invention-type distribution θ	
From Step 0	estimated share of process patents for each τ	
Calibration	patent premium λ visibility for products (G_P)	fixed ($= 0.1$) $\sim$ triangular(0.5)
Step 2	**SMM estimation of unconditional distributions (given R&D costs)**	
Estimate	unconditional visibility distributions F_Θ with $\Theta = M, P$ ($\phi \sim$ triangular($\cdot$)) unconditional invention-type distributions θ^F	
Moments	mean & variance of estimated and simulated conditional visibility distributions mean of estimated and simulated conditional invention-type distributions	
From Step 1	estimated conditional distributions G_Θ and θ_t	
Calibration	patent premium λ value v_i costs C_i	fixed ($= 0.1$) $\sim$ exp(0.1) $\sim$ logistic(C, 0.5)
Step 3	**Simulation of realized follow-on innovation**	
Simulate	$N =$ 1,000,000 potential inventions of full 3-stage sequential innovation model	
From Step 2	unconditional distributions F_Θ and θ^F	
Calibration	patent premium λ value v_i and v_{i_F} costs C_i and C_{i_F} baseline success probabilities ψ_D and ψ_S	fixed ($= 0.1$) independently $\sim$ exp(0.1) independently $\sim$ logistic(C, 0.5) fixed ($\psi_D = 2/3$ and $\psi_S = 1$)

of trade secrets protection τ_j at the time the decision to disclose the invention was made. Let $\rho(\tau_j)$ be the probability that a patent is a process patent as derived in Equation (4). Then, the log-likelihood of the data is given by

$$LL(G_M, G_P, \theta, \lambda) = \sum_j \mathbf{M}_j \log \rho(\tau_j) + (1 - \mathbf{M}_j) \log(1 - \rho(\tau_j)). \tag{A.3}$$

It is a function of the (conditional) distributions of visibilities G_Θ and the invention type θ, as well as the patent premium $\lambda = 0.1$. We estimate the model on the sample of single-state patents with priority dates between 1976 to 2008. For states that have adopted the UTSA, we exclude all patents with priority dates in the year of adoption.

Estimation of Unconditional Stage-1 Distributions (Step 2): In the second step, we estimate the *unconditional* distributions F_Θ of visibilities and θ^F of invention types, using as inputs the conditional distributions G_Θ and θ estimated in Step 1. For this second step, we follow a simulated-method-of-moments approach. First, for given unconditional distributions $\left(F_M, F_P, \theta^F\right)$ and some R&D cost C_i, we simulate a dataset of potential inventions and solve Stage 1 of our augmented model to obtain the *simulated* conditional distributions, $\delta \in \left\{\hat{G}_M, \hat{G}_P, \hat{\theta}\right\}$. Second, we calculate the simulated conditional moments $\hat{\mu}_m(\delta|F_M, F_P, \theta^F)$ for the simulated data and the estimated moments $\mu_m(\delta)$ based on the estimated conditional distributions G_Θ and θ from Step 1. Third, we define the quadratic score function

$$S(F_M, F_P, \theta^F) = \sum_{\delta} \sum_{m \in \mathcal{M}} \left(\hat{\mu}_m(\delta|F_M, F_P, \theta^F) - \mu_m(\delta)\right)^2 \tag{A.4}$$

where $\mathcal{M}$ is the set of moments (mean and variance for the visibility distributions and mean for the invention-type distribution). We minimize this score function over $\left(F_M, F_P, \theta^F\right)$ (specifically, the modes of the triangular visibility distributions and the shares of process inventions for the invention-type distributions) to obtain the unconditional distributions.

A.4.3 Estimation Results

In Table A.4, we report the parameters of both the conditional distributions (Step 1) and unconditional distributions (Step 2). For the conditional distributions, we obtain the distribution for the visibility of processes relative to the distribution for the visibility of products. A constant value of $\gamma_P = 0.5$ provides for a flexible specification without imposing our theoretical distributional assumptions. We find first-order stochastic dominance satisfied.

We report the parameters of unconditional distributions from Step 2 for no R&D costs ($C = 0$), low costs ($C = 2.295$ such that R&D costs are 40% of the expected R&D project value), and high costs ($C = 4.375$ such that R&D costs are 80% of the expected R&D project value). Note that, unlike in Step 1, where we fix G_P, in Step 2 we explicitly estimate F_P (i.e., the mode γ_P). First-order stochastic dominance (verified for the conditional distributions) continues to hold. The bottom panel of Table A.4 shows decisions at all three stages that are implied by the estimated parameters.

Table A.4: Estimates for Conditional and Unconditional Distributions

		(1)	(2)	(3)	(4)
			Stage 1: F_Θ, θ^F		
		Stage 2: G_Θ, θ	no cost	low cost	high cost
Mode for processes	γ_M	0.376 [0.3755, 0.3796]	0.373	0.328	0.196
Mode for products	γ_P	0.5	0.501	0.466	0.31
Share of processes	θ	0.454 [0.4536, 0.4550]	0.456	0.460	0.463
R&D intensity (Stage 1)			1.000	0.970	0.702
... of processes			1.000	0.964	0.683
... of products			1.000	0.975	0.718
Patents (Stage 2)			0.768	0.750	0.674
... for processes			0.734	0.709	0.637
... for products			0.797	0.784	0.704
R&D intensity (Stage 3)			0.512	0.399	0.291
... from processes			0.490	0.377	0.275
... from products			0.531	0.417	0.305

Notes: We report the parameter estimates for the conditional distributions from Stage 2 and Stage 1 of the augmented model. For Stage 2 (Step 1) in Column (1), we estimate the mode γ_M (of the triangular distribution over support $[0, 1]$) for processes and hold constant the mode $\gamma_P = 1/2$ for products. Invention types are Bernoulli distributed with parameter θ. We report in brackets the 99% confidence interval from 1000 bootstrap replications. The reported point estimates are from one single model using the full sample. For the simulated-method-of-moments approach for Stage 1 estimates (Step 2), we use the first two moments (mean and variance) for G_M and G_P and the first moment (mean) for θ. For the costs of the initial invention as well as the follow-on invention, we assume that $C_i = C + \varepsilon_i$ and $C_{i_F} = C + \varepsilon_{i_F}$ where ε_i and ε_{i_F} are (independently) logistically distributed with zero mean and scale $1/2$. We set $C = 0 = C_i$ (no cost) in Column (2), $C = 2.295$ (low cost, such that costs are 40% of the expected R&D project value) in Column (3), and $C = 4.375$ (high cost, such that costs are 80% of the expected R&D project value) in Column (4). We further assume that the value of the initial invention and follow-on innovation are (independently) drawn from the same distribution, $v_i, v_{i_F} \sim \text{Exp}(1/10)$. At the bottom of the table, we report R&D intensities at Stage 1 (share of inventions i that are developed) and Stage 3 (share of inventions i_F that are developed, conditional on Stage-1 R&D) and the share of patented inventions i (conditional on Stage-1 R&D) at Stage 2.

A.5 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

$$\overline{W}(\tau) = E_{(\Theta_i, \phi_i, \tilde{d}_i, v_i, v_{i_F})} \left[\mathbf{R}_i(\tau) \left(W_i + \tilde{\psi}_{i_F, \tilde{d}_i} \mathbf{R}_{i_F} W_{i_F} \right) \right], \tag{A.5}$$

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 inital (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

$$W_D(\phi) = 2v_i - \frac{\phi v_i}{2} - C_i, \tag{A.6}$$

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 3, 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

$$W_S(\phi,\tau) = 2v_i - \frac{\tau\,(1-\phi)\,v_i}{2} - C_i. \tag{A.7}$$

To summarize, using the disclosure condition in Equation (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 $W_{i_F} = 2v_{i_F} - C_{i_F}$.

For a concrete example, consider a market with linear demand $D(p) = 1 - p$. A firm with a new technology produces a homogeneous good at marginal production costs of c_L. This firm has many potential competitors that all produce at marginal costs $c_H > c_L$. Competition is in prices. We assume the invention is radical in the sense that the monopoly price (under low costs c_L) does not exceed the higher of the marginal costs, $p_L^m \leq c_H$. Moreover, for simplicity let $c_L = 0$. The monopoly profits in this case are $\pi_L^m = \frac{1}{4}$.

When the firm chooses to patent the technology, so that all potential competitors have (restricted) access to the technology, it is able to detect infringement of its patent (and enforce it) with probability ϕ. This means, with probability $1 - \phi$, there is at least one competitor who can freely use the low-cost technology. With at least one competitor producing at zero marginal cost, the equilibrium price (and deadweight loss) is equal to zero. The expected social surplus is $\phi \frac{3}{2\pi_L^m} + (1 - \phi) \cdot 0 = 2\pi_L^m - \frac{\phi \pi_L^m}{2}$.

Instead of a patent, let the firm keep the technology a secret. As discussed in the main text, the firm has exclusive access to the technology with probability $\tau(1 - \phi)$. This means that with probability $1 - \tau(1 - \phi)$ there is at least one competitor who can freely use the low-cost technology – and the equilibrium price and deadweight loss are equal to zero. The expected social surplus is

$$\tau(1 - \phi) \frac{3}{2\pi_L^m} + [1 - \tau(1 - \phi)] \cdot 2\pi_L^m = 2\pi_L^m - \frac{\tau(1 - \phi)\pi_L^m}{2}.$$

Let v denote the commercial value of the invention if the firm has exclusive access. In other words, let $v = \pi_L^m$, then the expressions for expected aggregate surplus are equal to the expression in Equations (A.6) and (A.7).

// IMPRINT

ZEW – Leibniz-Zentrum für Europäische Wirtschaftsforschung GmbH Mannheim

ZEW – Leibniz Centre for European Economic Research

L 7,1 · 68161 Mannheim · Germany
Phone +49 621 1235-01
info@zew.de · zew.de