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

Bernhard Ganglmair† Imke Reimers‡

March 19, 2019

Abstract

Patents grant an inventor temporary monopoly rights in exchange for the disclosure of the patented invention. However, if only those inventions that are otherwise already visible are patented (and others kept secret), then the bargain fails. We use exogenous variation in the strength of trade secrets protection from the Uniform Trade Secrets Act to show that a relative weakening of patents (compared to trade secrets) adversely affects patenting of processes more than that of products. Arguing that processes are on average less visible (or self-disclosing) than products, stronger trade secrets have thus a disproportionately negative effect on the disclosure of inventions that are not otherwise visible to society. We develop a structural model of initial and follow-on innovation to determine the effects of such a shift in disclosure on overall welfare in industries characterized by cumulative innovation. In counterfactual analyses, we find that while stronger trade secrets encourage more investment in R&D, they may have negative effects on overall welfare – the result of a significant decline in follow-on innovation. This is especially the case in industries with relatively profitable R&D.

Keywords: cumulative innovation; disclosure; process innovation; trade secrets; Uniform Trade Secrets Act.

JEL Codes: C15; D83; O31; O33; O34

*We thank Laura Grigolon, Dan Gross, Jeffrey Kuhn, Brian Love, Lisa Ouellette, Ivan Png, Jason Rantanen, David Schwartz, Neel Sukhatme, and seminar and conference participants at the Max Planck Institute for Competition and Innovation in Munich, Northeastern University, University of Connecticut, ZEW Mannheim, the IP Statistics for Decision Makers Conference (2017), the International Industrial Organization Conference (2018), the NOeG Winter Workshop (2018), and the Roundtable on Empirical Methods in IP at Northwestern University School of Law (2018) for comments and suggestions. Bernhard Ganglmair acknowledges financial support from the Collaborative Research Center (CRC) Transregio 224 EPoS (Bonn-Mannheim). The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†ZEW, MaCCI, and University of Mannheim. Address: ZEW – Leibniz Centre for European Economic Research, L7, 1, 68161 Mannheim, Germany; E-mail: b.ganglmair@gmail.com; Twitter: @ganglmair.

‡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 are expected to invest and innovate more. 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, instead keeping them secret. With stronger protection of such trade secrets (or with weaker patents), fewer of the proverbial shoulders become visible and therefore available for others to stand on. This effect is particularly prevalent 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. In this paper, we ask how a change in the attractiveness of secrecy relative to patents affects the diffusion of knowledge through the decision to invest in innovation, the disclosure of inventions, and the ability to build on these inventions.

Secrecy is an important tool in a firm's intellectual property management toolbox. Generally speaking, a trade secret is information ("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., 1991:61). There is ample survey-based evidence that (trade) secrets are 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. Trade secrets laws offer protection against *misappropriation* of secrets – that is, the acquisition of a trade secret by *improper means* (for instance, theft, bribery, misrepresentation, breach of contract, or espionage) or the disclosure of a trade secret without consent. However, unlike patents, trade secrets laws generally do not grant exclusivity.[2] This means, a trade secret is not protected if it accidentally leaks or

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

[2] Codified trade secrets laws in the U.S. go back to the Restatement (First) of Torts of 1939. The Uniform Trade Secrets Act of 1979/1985 was recommended for state-level adoption to clarify and

is uncovered through independent discovery or reverse engineering (Friedman et al., 1991).[3]

Stronger trade secrets protection renders trade secrets more attractive relative to patents (Png, 2017b). In this paper, we use exogenous variation of trade secrets protection across states and time from the staggered adoption of the Uniform Trade Secrets Act (UTSA) of 1979/1985 to study the trade-off between secrecy and disclosure through patenting for different technology types.[4] Using 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., 2019), we show that stronger trade secrets protection results in a disproportionate decrease of process patents. We then estimate the parameters of a structural model of sequential innovation that takes both ex-ante and follow-on innovation incentives into account. We find that total welfare may in fact decline as trade secrets protection grows stronger. This negative welfare effect is due to the reduced patenting of less visible inventions – processes – for which disclosure is essential to allow for follow-on innovation.

The paper proceeds in four steps. In Section 2, we develop a simple model of an inventor's decision to disclose a new technology through a patent. The value of the invention from a patent increases with the underlying invention's visibility: visibility allows for easier enforcement of the patent – guaranteeing exclusive access to the technology. The value of the invention that is kept secret, however, decreases in visibility, because secrecy (and therefore exclusive access) is more difficult to maintain. We assume that *processes* are on average less visible than *products*. The assumption implies that, on average, inventors of processes value secrecy more than those of products – consistent with survey evidence (Levin et al., 1987; Cohen et al., 2000; Arundel, 2001; Hall et al., 2013). For a given baseline share of process *inventions*, our model predicts that, as trade secrets protection improves, the share of process *patents* decreases. This theoretical prediction serves as the basis for the empirical analysis in the rest of the paper.

In Section 3, we discuss our two main datasets that we merge with basic bibliographic patent information. First, we use an index constructed by Png (2017a) that

harmonize trade secrets protection at the state level. With the passing of the Economic Espionage Act of 1996 (criminal) and the Defend Trade Secrets Act of 2016 (civil), the U.S. now has two federal law bodies governing trade secrets. In Europe, before the adoption of the EU Trade Secrets Directive (2016/943), trade secrets laws were fragmented, and relevant provisions found in labor law (France), law of unfair competition (Germany), or considered breach of confidence (UK).

[3]The Uniform Trade Secrets Act of 1979/1985, for instance, lists as such *proper means*: "discovery by independent invention; discovery by reverse engineering [...]; discovery under a license from the owern of the trade secret; observation of the item in public use or on public display; obtaining the trade secret from published literature."

[4]We do not consider the joint use of patents and secrecy (Arora, 1997; Crass et al., 2019) or disclosure without patenting (for instance, through academic publishing (Thursby et al., 2018) or corporate technical journals, such as the IBM Technical Disclosure Bulletin or the Xerox Disclosure Journal). Our assumption of the choice between secrecy and patents comes without loss of generality as long as there is *some* degree of substitutability between these two options.

measures the strength of legal protection of trade secrets before and after a state's adoption of the UTSA. It reflects the trade secrets protection to which an inventor in a given state was exposed at the time of her disclosure decision. Second, we use data compiled by Ganglmair et al. (2019) to construct process and product patent indicators.

In Section 4, we use these data to test the model implications. We use exogenous variation across locations and time in the level of trade secrets protection due to the staggered adoption of the UTSA by various U.S. states to estimate the effect of stronger trade secrets protection on the likelihood that a patent covers a process in a differences-in-differences estimation. Consistent with results from our theoretical model, we find that better legal protection of trade secrets leads to a disproportionate decrease of patenting of processes. Our estimated effects are largest among individual inventors (compared to large firms). We confirm our baseline results and the identifying assumptions in a number of robustness checks, including an instrumental variables strategy that uses state-specific adoption of other, unrelated policies to estimate a state's UTSA adoption (Png, 2017b).

In Section 5, we estimate the parameters of a dynamic model of sequential innovation and use these estimates to establish our welfare results. In addition to the disclosure decision from Section 2, we now take both an inventor's ex ante R&D decision and the effect of disclosure on follow-on innovation into account. We model follow-on innovation consistent with stylized facts: more disclosure of technical information boosts follow-on innovation (Williams, 2013; Gross, 2019), patents on early ideas raise the costs of creating future ideas (Scotchmer, 1991; Heller and Eisenberg, 1998; Galasso and Schankerman, 2015), and the information disclosed in patents is of sufficient quality and useful (Furman et al., 2018).

Our structural model provides estimates for the ex ante distributions of each invention type as well as their visibilities. These allow us to calculate the R&D intensity and the share and characteristics of trade secrets (over all realized inventions). The counterfactual analyses show that stronger trade secrets protection has a negative overall welfare effect in industries with relatively profitable R&D. When the benefits of trade secrets protection are inframarginal to an inventor, stronger legal protection of trade secrets has the unintended consequence of lowering total welfare by impeding follow-on innovation. This pattern is reversed for R&D projects that are relatively less profitable. In this case, the benefits are marginal and stronger legal protection improves welfare by encouraging initial R&D. We further show that the negative effects are more pronounced as the differences of the invention types in terms of their average visibilities increase.

This study contributes to several streams of literature. Beyond a number of studies based on survey data (e.g., Levin et al., 1987; Cohen et al., 2000; Arundel, 2001), there is limited empirical work on trade secrets – for the obvious data limitations. A small literature presents indirect evidence on secrecy. Moser (2012) documents a shift toward patenting (and away from secrecy) in the chemical industry as reverse engineering became easier with the publication of the periodic table of elements. Gross (2019) finds that a policy during World War II to keep certain patent applications secret resulted in fewer citations and 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 (implying more rapid knowledge diffusion) had a mitigating effect on licensing delays.

A related strand of literature studies the effect of changes in legal trade secrets protection on innovation and patenting behavior. Png (2017a) finds that stronger trade secrets protection has a positive effect on firms' investment in R&D, at least in the high-tech industry and for large companies. Similarly, Png (2017b) finds that strengthening trade secrets protection 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. Angenendt (2018) finds that patent applicants respond to stronger trade secrets protection through the UTSA by reducing the number patent claims and decreasing the amount of information disclosed. We add to this body of literature by analyzing the role of an invention's visibility in measuring the effect of an increase in trade secrets protection on the patenting and innovation decisions.

We explicitly model and estimate an inventor's behavioral response to stronger trade secrets (or weaker patents) and the subsequent decline in disclosure of inventions. Such a general equilibrium approach is critical to assessing the full welfare consequences of recent U.S. Supreme Court rulings that have narrowed the scope of what is and what is not patentable (see Sampat and Williams, 2018). Our welfare results provide new insights for the evaluation of these rulings. Moreover, to our knowledge, this is the first paper presenting welfare results explicitly for changes in trade secrets laws. This is particularly interesting in light of the EU Trade Secrets Directive 2016/943 adopted in June 2016 for which impact evaluations are not yet available. Results from the U.S. can thus inform an ongoing policy debate in Europe.

2 A Model of Trade Secrets and Disclosure

For our welfare analyses in a later section, we consider a 3-stage model of sequential innovation. Stage 1 describes the inventor's decision to invest in R&D and realize the initial invention. At Stage 2 (this section), the inventor decides to disclose the (patentable) invention or keep it a secret[5]. Stage 3 eventually captures the market's decision to engage in follow-on innovation that builds on the initial invention from

[5]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 it 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). Patentability of the invention in our context implies that the inventor is given a true choice between disclosure (through a patent) and trade secrecy.

Stage 1. In this section, we analyze the disclosure decision at Stage 2. We return to the full 3-stage augmented model in Section 5 when we present our welfare results.

2.1 An Inventor's Decision to Disclose

An invention i at Stage 2 can be described by a tuple (ϕ, Θ, v) and is characterized by its visibility ϕ, its type Θ, and its private commercial value v (from exclusive use). Visibility is the parties' ability to observe an invention or its use. We discuss each of the invention's characteristics below.

An inventor is given the choice to disclose an invention in a patent ($\tilde{\pi} = D$) or keep the invention secret ($\tilde{\pi} = S$). We set the inventor's private returns $V_{\tilde{\pi}}$ 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. A lower degree of exclusivity thus means the inventor reaps a smaller fraction of these rents. In both disclosure states $\tilde{\pi}$, the probability of exclusive use depends on the *visibility* of the invention.[6]

Once the inventor has disclosed the invention in a patent, she can accumulate profits only if that patent is enforceable and other firms can be excluded from its use. In order to enforce a patent, the patent holder must be able to detect the use of an invention by a potential infringer. A more visible invention with higher observability of its use is therefore easier to enforce (and exclusivity prevails). Patents for non-visible inventions, on the other hand, are not enforceable and of zero value because rents dissipate once the invention is freely available. The expected commercial value the inventor is able to materialize is therefore $\phi \cdot v$. In addition, the inventor receives a patent premium λ.[7] It captures the benefits from patenting over trade secrets.[8] We define the inventor's private value of disclosing the invention as

$$V_D(\phi) = \phi\,(1 + \lambda)\,v. \tag{1}$$

While a patent needs visibility to be of value, a trade secret's value *decreases* in the visibility of the underlying invention. Moreover, the value of trade secrecy increases with the level of trade secrets protection. We denote the exogenous proba-

[6] In certain applications, higher visibility can also be interpreted as a higher probability that the invention can be reverse-engineered. Scotchmer and Green (1990) show that an inventor of a patentable technology might not want to patent and keep the technology off the market to avoid reverse engineering. For a general treatment of reverse engineering, see Samuelson and Scotchmer (2002).

[7] 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 simply improve an inventor's bargaining position in license negotiations. Webster and Jensen (2011) further provide evidence for premium from commercialization, showing that being refused a patent has a significant negative effect on a firm's decision to launch and mass produce a product.

[8] For simplification, the patent premium λ captures these benefits in excess of what the inventor, if anything, could earn, for instance, from licensing the invention as a trade secret.

bility that a trade secret is protected by τ. Even with perfect trade secrets protection ($\tau = 1$), an invention is of little value to the inventor if it is visible and can easily be copied by others without the inventor's ability to police such use.[9] Moreover, weaker trade secrets protection reduces deterrence and results in more unsanctioned misappropriation (or unprotected use) of trade secrets (e.g., Friedman et al., 1991:68). We therefore assume that, without any trade secrets protection, the value of trade secrecy is zero even for non-visible inventions. This is because of disclosure through misappropriation.[10] We define the private value of a trade secret as

$$V_S(\phi, \tau) = \tau\,(1 - \phi)\,v. \tag{2}$$

That is, inventor of (ϕ, Θ, v) chooses disclosure if, and only if, $V_D(\phi) \geq V_S(\phi, \tau)$. This condition can be rearranged to read

$$\phi \geq \frac{\tau}{1 + \lambda + \tau}.$$

The inventor chooses disclosure through patenting if, and only if, visibility of the invention is sufficiently high (or trade secrets protection and the patent premium are sufficiently low). For a given ϕ, we can summarize the decision to disclose and patent, $\tilde{\pi} \in \{D, S\}$, as

$$\tilde{\pi} = \begin{cases} D & \text{if } \phi \geq \dfrac{\tau}{1 + \lambda + \tau} \\ S & \text{if otherwise.} \end{cases} \tag{3}$$

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).[11] The following lemma summarizes basic comparative statics of the inventor's decision to disclose. The proofs of this and all other results are relegated to the Appendix.

Lemma 1. *An inventor is more likely to disclose her invention by filing for a patent as the degree of visibility ϕ and the patent premium λ increase; she is less likely to patent as the degree of trade secrets protection τ increases.*

[9] Note that in our model we do not allow for independent discovery (that is independent of visibility ϕ). We also assume that if a competitor has rightfully acquired the invention, she cannot take out a patent.

[10] This is not as strong an assumption as it appears to be. Generally, the threat of legal sanctions will deter (at least some) misappropriation, and the lack of such a threat will encourage it. Friedman et al. (1991) and also Lemley (2008) have argued that if trade secrets protection is weak, firms erect often inefficient safeguards. The costs of these is expected to increase in v and decrease in τ. Without trade secrets protection, the effective commercial value may in fact fully dissipate.

[11] 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).

External sources provide corroborating evidence for our prediction. Moser (2012) provides empirical evidence for more patenting as visibility increases (captured by the ease of reverse engineering an invention), and Png (2017b) shows that patenting decreases as trade secrets protection increases.

2.2 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. More specifically, an invention is either a process (or method), $\Theta = M$, or a product, $\Theta = P$. Invention types are Bernoulli distributed where $\theta = \Pr(\Theta = M)$ is the probability that the realized invention is a process. We denote this distribution by $\mathcal{G}$. Note that these distributions (G_Θ for $\Theta = M, P$ and $\mathcal{G}$) are *conditional* distributions given the inventor's positive R&D decision at Stage 1 of the augmented model.

We assume that processes are on average less visible than products (e.g., Strandburg, 2004).[12] We formally capture this by assuming first-order stochastic dominance: G_P first-order stochastically dominates G_M so that $G_M \geq G_P$ for all ϕ. One implication of this assumption is a higher value of secrecy for processes than for products, given v. Conversely, the value of disclosure is lower for processes than for products. The (expected) value of secrecy of an invention of type Θ is

$$EV_{S|\Theta}(\tau) = \int_0^1 \tau\,(1-\phi)\,v dG_\Theta(\phi); \tag{4}$$

the expected value of disclosure is

$$EV_{D|\Theta}(\tau) = \int_0^1 \phi\,(1+\lambda)\,v dG_\Theta(\phi). \tag{5}$$

We show the claim in

Proposition 1. *Let $G_P(\phi) \leq G_M(\phi)$ for all ϕ. For a given level of trade secrets protection τ, the value of secrecy is higher for processes than for products, $EV_{S|M}(\tau) > EV_{S|P}(\tau)$. Conversely, the value of disclosure is lower for processes than for products, $EV_{D|M}(\tau) < EV_{D|P}(\tau)$.*

Empirical evidence comports with this theoretical finding. Using survey data, Levin et al. (1987), Cohen et al. (2000), Arundel (2001), or Hall et al. (2013) find that the propensity to patent is higher for products than processes, suggesting a higher value of secrecy for processes. In the Appendix, we present empirical evidence for the same. We exploit a change of the publication policy of pending U.S. utility

[12] We find support for this assumption both among legal practitioners (Goldstein, 2013:65–66) and managers (Federal Trade Commission, 2003:ch. 3, p. 32).

patent applications through the *American Inventors Protection Act of 1999*. Eligible patent applicants were given the option to delay the disclosure of their inventions (i.e., publication of their applications) and thus extend the period of temporary secrecy. While the baseline probability of opting out of disclosure is somewhat low (Graham and Hegde, 2015), we find strong evidence that applicants of process patents are more eager to extend the temporary secrecy of their inventions.

2.3 Probability to Disclose for Different 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 three simple auxiliary results. In Lemma 2, we show that the probability that a process is patented is weakly smaller than the probability that a product is patented. For this, let $\pi(\phi, \tau) = 1$ if $\tilde{\pi} = D$ and $\pi(\phi, \tau) = 0$ if $\tilde{\pi} = S$. The probability that an invention of type Θ is patented and disclosed is

$$\pi_\Theta(\tau) = \int_0^1 \pi(\phi, \tau) dG_\Theta(\phi). \tag{6}$$

Lemma 2. *For a given level of trade secrets protection τ, $\pi_M(\tau) \leq \pi_P(\tau)$.*

In Lemmas 3 and 4, we show that patenting probabilities are decreasing in trade secrets protection for both invention types, and that the patenting probability for products is decreasing at a lower rate than that for processes.

Lemma 3. *The patenting probabilities for products $\pi_P(\tau)$ and processes $\pi_M(\tau)$ are decreasing in τ.*

Lemma 4. *The difference between the patenting probabilities for products $\pi_P(\tau)$ and processes $\pi_M(\tau)$ is increasing in trade secrets protection τ.*

The patenting probability $\pi_\Theta(\tau)$ captures the probability that an invention of type Θ is disclosed through patenting. We do not observe, however, the characteristics of the underlying invention. Instead, we assume distributions G_Θ. Given the distribution $\mathcal{G}$ of invention types with $\theta = \Pr(\Theta = M)$, the probability that a given patent covers a process is

$$\rho(\tau) = \frac{\theta \pi_M(\tau)}{\theta \pi_M(\tau) + (1 - \theta)\, \pi_P(\tau)}. \tag{7}$$

The expression in (7) can be interpreted as the share of process patents in a sample of patents (where patents are either process or product patents). It is decreasing as trade secrets protection increases. We show this in

Proposition 2. *The share of process patents (patents covering a process or method invention) is decreasing as trade secrets protection increases.*

In other words, Proposition 2 predicts that, in response to an (exogenous) increase in trade secrets protection, the probability that a given patent is a process patent decreases. In Section 4, we take this prediction to the data, using the staggered introduction of the *Uniform Trade Secrets Act* to identify the effect of trade secrets protection on patenting behavior.

3 Institutional Background and Data

We examine the effects of a state-specific change in the strength of trade secrets protection to determine the role of an invention's visibility for the decision to disclose. Below, we provide institutional background and discuss our measure of trade secrets protection (Png, 2017a). We then introduce a dataset to identify process and product patents (Ganglmair et al., 2019) and discuss further data sources and definitions of additional control variables.[13]

3.1 Uniform Trade Secrets Act (1979/1985)

As our source of identification, we use the introduction of the Uniform Trade Secrets Act (UTSA). The UTSA is a body of laws relating to the protection of trade secrets. It was published and recommended to the individual U.S. states for adoption in 1979 (with a revision in 1985) by the National Conference of Commissions on Uniform State Laws. Since 1979, 47 states, the District of Columbia, Puerto Rico, and the U.S. Virgin Islands have adopted the UTSA, with adoption dates ranging from 1981 (5 states) to 2013 (Texas).[14] Using information on the level of trade secrets protection before and after a state's adoption of the UTSA, Png (2017a) constructs an index that measures the change in legal protection of trade secrets. We observe a strengthening of trade secrets protection if, for instance, the UTSA introduces a broader definition of what is a trade secret or a wider list of circumstances under which trade secrets law has been violated.[15]

Figure 1 illustrates the *change* in this index in individual states as they adopted the UTSA in a given year, with higher values implying larger increases in protection. In most states, the UTSA resulted in a strengthening of trade secrets protection, with the exception of Michigan, Nebraska, and Wyoming, where the UTSA had no effect, and Arkansas and Pennsylvania, where pre-UTSA trade secrets protection (under common law) was stronger. There is no obvious pattern in the size of these changes

[13]For more detailed information, see the Online Appendix.

[14]The list of adopting states includes all states except New York, Massachusetts, and North Carolina (Sandeen and Rowe, 2013).

[15]The index summarizes the inclusion of six different factors. On substantive law: continuous use requirement, requirement to take reasonable effort to protect trade secrets, and mere acquisition as misappropriation. On civil procedure: limitations on when a trade secret owner can take legal action. On remedies: limitations of injunctions, and availability of a punitive damages multiplier.

Figure 1: Change in Legal Protection of Trade Secrets (Png, 2017a)

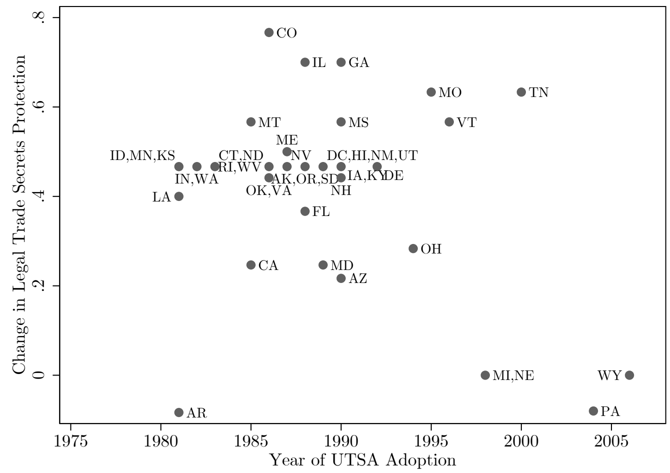

Notes: This figure presents data from Table 1 in Png (2017a). For the states that adopted the UTSA between 1981 and 2006, it depicts the change in legal protection of trade secrets across states as a result of the UTSA.

over time and across states, and anecdotal evidence suggests that passing of the bills often happened for "whimsical" reasons.[16]

We use annual data of Png's trade secrets protection index for all 50 states (plus the District of Columbia) for the years 1976 through 2008. Previous literature has shown that a change in trade secrets protection can have meaningful effects on firm innovation and patenting behavior (Png, 2017a,b; Contigiani et al., 2018; Angenendt, 2018).[17]

3.2 Timing of the Disclosure Decision and Patent Location

For a clean assignment of the level of trade secrets protection to which the patent applicant was exposed at the timeof the disclosure (and patenting) decision, we must identify the timing of the patenting decision as well as the invention's location (i.e., state of origin). We use the earliest priority date of the respective granted patent to determine the timing of the disclosure decision. The earliest priority date reflects the application date of the first patent application (i.e., the *parent application*) from which a patent's ultimate application draws. The priority date therefore accounts for

[16]See Pooley and Westman (1997) for more detail on trade secrets and the adoption of the UTSA.

[17]For a change in trade secrets protection in a given state to have an effect on innovation and patenting in that state, trade secrets protection must be determined by the state where the secret was developed and not where it was misappropriated. In Paolino v. Channel Home Centers, 668 F.2d 721 724 n.2 (3d Cir. 1982), the court argues in support of this identifying assumption, finding that "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."

an earlier disclosure decision for continuation or divisional applications.[18] We believe that the relevant disclosure decision was made at the time of the parent application, and we use that application's priority date as date for all related patents.[19]

For the location of the patent, we consider only patents for which both all U.S. inventors and U.S. assignees are from the same state and use that state as the patent's location.[20] While a significant number of patents list multiple inventors and assignees, oftentimes located in different states, our approach allows us to unambiguously identify a patent's location. It also ensures that the patent applicant's decision was driven by only that state's level of trade secrets protection, and not contaminated by changes in other states.

With our assumption of single-state patents, we limit our overall sample to 1,487,400 patents (out of 4,402,480 total), granted between 1976 and 2014 and with priority dates between 1976 and 2008.[21]

3.3 Indicators for Process and Product Patents

To categorize process and product patents, we use data constructed by Ganglmair et al. (2019), who employ text-analytical methods to identify the invention type of an independent claim in a given patent.[22] Claims are 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 aggregate the claim-level information in Ganglmair et al. (2019) to obtain an indicator for the invention type at the *patent* level. More specifically, 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; a patent is a *product patent* otherwise.[23,24]

[18] These are applications by the inventor that claim the priority date of a "parent" patent. For continuations, the applicant may not add new disclosures but may delete claims. Divisions involve separating a patent application into two or more.

[19] Our results are robust to using the more commonly used definition of the patent's application date.

[20] We disregard foreign inventors and assignees for this patent-state identification.

[21] For alternative specifications, we use as patent location the location of the first assignee or the location of the first inventor listed on the patent. As reported in the Appendix, results are very similar.

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

[23] For the purpose of our study, we treat product-by-process claims as process claims. Note that this assumption does not drive our results. If instead we treat product-by-process claims like product claims or drop them from the sample, our results follow through (available upon request).

[24] Our process patent indicator is rather aggressive. In the Appendix, we present robust results using two alternative and less aggressive definitions for process patents.

Table 1: Process Patent Indicators

	Before UTSA		After UTSA		
	Mean	Std. Dev.	Mean	Std. Dev.	T-stat
Process Patent	0.4258	0.4945	0.5173	0.4997	110
Number of Process Claims	0.7562	1.305	0.9637	1.477	89.8
Number of Product Claims	1.817	1.818	1.964	1.933	47.4
Observations	674,186		813,214		

Notes: This table provides summary statistics for the process patent indicator and the number of process and product claims per patent, for all granted utility patents (between 1976 and 2014) with a priority date between 1976 and 2008, and for which the location can be unambiguously identified. The process patent indicator variable is defined for all patents for which we can identify the claim type for at least one independent claim (666,131 patents before UTSA adoption and 808,980 patents after UTSA adoption).

For our final sample, we drop all business method patents.[25] Table 1 provides summary statistics for our patent indicators for all granted USPTO utility patents in our sample (all single-state patents granted between 1976 and 2014, with priority dates between 1976 and 2008), distinguishing between patents with priority dates before and after the UTSA adoption in the respective applicant's state. Generally, patents disclosed after UTSA adoption are much more likely to cover processes. In addition, both the number of process and product claims within a patent increases significantly.

Figure 2 plots the share of process patents in each priority year. The figure suggests that the increase in process patenting after the UTSA in Table 1 may simply be the result of a general trend toward more process innovation. An analysis of the effect of the UTSA therefore must take other factors into account.

3.4 Additional Variables

We collect and construct additional patent characteristics to capture the complexity, value, and importance of the patented technology. Table 2 summarizes these variables across all patents in our dataset of single-state patents. 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

[25]We loosely follow Lerner (2006) who identifies business methods patents as patents with a United States Patent Classification (USPC) main class 705. Our results are robust to this sample restriction (robustness results upon request). Strandburg (2004) argues that business methods are "self-disclosing processes" and thus highly visible. Note that this argument does not necessarily apply to all business-method patents. Boldrin and Levine (2008a:168–9) discuss why Amazon's one-click patent (U.S. patent number 5,960,411) may in fact not be as self-disclosing as one may suspect.

Figure 2: Process Patents Over Time

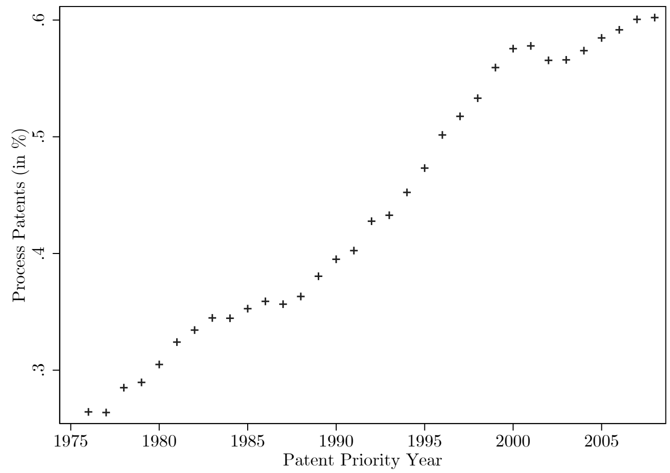

Notes: This figure presents annual data for the percentage of process patents. These numbers represent the annual share of process patents by the year of the respective patent's priority date, for our sample period of 1976 through 2008.

Kuhn and Thompson, 2017), where shorter claims are likely broader. As additional measures of a patent's complexity, we include the length (in characters) of the patent's detailed description text.

To capture the external value (or technological impact) of a patent, we construct measures of *patent generality* and *patent originality* as proposed by Trajtenberg et al. (1997). Patent generality captures the diversity of patents (measured by their respective patent classes) in which a given patent is (forward)-cited. A higher generality score implies a higher widespread impact, influencing subsequent innovation in more fields (Hall et al., 2001). Patent originality, on the other hand, captures the diversity of technologies from which a given patent draws – measured by the diversity of patent classes the patent (backward)-cites. A higher originality score means that the patented invention is combing ideas from different areas to create some new (or "original"). We construct these measures for each patent using the first USPC main class listed on the patent.[26] As a measure of internal or private value of a patent, we use information on whether the patent holder paid the patent maintenance fees during the 4th year of the patent term (see, e.g., Pakes, 1986; Schankerman and Pakes, 1986).[27]

[26]There are about 450 main classes and about 150,000 subclasses in the United States Patent Classification (USPC) system. For more information, see `http://www.uspto.gov/patents/resources/classification/overview.pdf`.

[27]For more information on patent maintenance, including the fee schedule, see `https://www.uspto.gov/patents-maintaining-patent/maintain-your-patent`.

Table 2: Additional Patent Characteristics

	N	Mean	Std. Dev.	Min	Max
Independent claims	1,477,567	2.876	2.287	0	138
Length of first claim (words)	1,477,563	169.580	107.045	0	7,078
Length of description (characters)	1,487,790	26,260.81	40,089.22	4	3,608,036
Generality	1,120,032	0.638	0.244	0	0.967
Originality	1,308,621	0.626	0.244	0	0.983
4th year renewal	1,393,420	0.825	0.380	0	1

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.

4 Empirical Estimation and Results

In this section, we examine whether stronger trade secrets protection has a disproportionate negative effect on process patenting. To do this, we take advantage of the staggered adoption of the UTSA across U.S. states over the course of more than 20 years. Providing evidence that the state-specific timing of the adoption was random for the purposes of this study, we estimate the likelihood that a patent includes a process (Proposition 2) in a difference-in-differences setting.

4.1 The Impact of Trade Secrets Policy

In our main specification, we estimate the probability that a patent covers a process invention as a function of the patent's characteristics as well as the state's trade secrets protection index. Formally, we estimate:

$$process_{jst} = \beta_1\, protection_{st} + \beta_2 X_{jt} + \nu_s + \mu_t + \eta_j + \epsilon_{jst}, \tag{8}$$

where the dependent variable is an indicator that is 1 if patent i filed in year t by an entity in state s is a process patent. $protection_{st}$ is the *value* of the trade secrets protection index relative to the state's base year. To control for any events that occur in all states simultaneously (such as the AIPA) and for an state- and USPC class-specific characteristics that do not vary over time, we include location-state (ν_s) and priority-year (μ_t) fixed effects, as well as dummy variables for patent j's first USPC main class (η_j). Further, X_{jst} includes patent-specific measures of complexity and value, as described in Section 3.[28] Thus, our coefficient of interest β_1 captures the effect of the *change* of protection. Finally, we cluster standard errors by the first USPC main class listed on the patent and the patent's state to allow for common

[28] While some of these variables are likely endogenous, we control for them regardless because we are interested in the impact of $protection_{st}$ on the probability of a patent including a process claim, and these covariates are likely correlated with this probability.

Table 3: Baseline Results – Impact of Trade Secrets Protection

	(1)	(2)	(3)	(4)
Trade secrets protection	-0.018*	-0.021**	-0.026***	-0.026***
	(0.009)	(0.009)	(0.009)	(0.008)
Log(indep. claims)		0.233***		0.231***
		(0.003)		(0.003)
Log(length of first claim)		-0.044***		-0.051***
		(0.004)		(0.003)
Log(length of description)		-0.002		0.001
		(0.002)		(0.002)
Originality			0.025***	0.010**
			(0.005)	(0.005)
Generality			0.061***	0.038***
			(0.004)	(0.004)
4th year renewal			0.044***	0.025***
			(0.002)	(0.002)
Observations	1475058	1465095	907867	899932
$\overline{R}^2$	0.300	0.345	0.289	0.337

Notes: Linear probability model with 1[process patent] as the dependent variable, and the index of trade secrets protection (Png, 2017a) as the independent variable of interest. Robust standard errors, clustered by USPC main class and state, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls include indicator variables for the patent's first listed USPC main class, the location state, and the priority year.

trends within these classes and states.

Table 3 shows the coefficients from the baseline specification for all granted single-state patents (between 1976 and 2014) with priority dates between 1976 and 2008, including different sets of control variables.[29] All specifications estimate a negative impact of a UTSA-related strengthening of trade secrets protection on the probability that a patent is a process patent. The specification including control variables on measures of patent complexity *and* measures of the patent's value (column (4)) finds that a patent is 2.5 percentage points less likely to be a process patent if the trade secrets protection index rises by a full point. At a baseline of 42.8% of process patents before UTSA adoption (see Table 1), and with a mean increase in trade secrets protection of 0.36 points across all patents, this corresponds to a mean decrease of 2.2% in the probability that a patent is a process patent when a state adopts the UTSA. This impact corresponds to economically significant changes in patenting decisions and is statistically significant and robust to various specifications.

[29] We report results of a linear probability model for ease of interpretation.

4.2 Identification and Instrumental Variables

Our identification strategy relies on two assumptions. First, the adoption of the UTSA is not affected by an expectation that certain types of innovation will be more prevalent in the future. To that end, Png (2017a) provides evidence of the exogeneity of the UTSA with regard to firms' decisions to invest in R&D. Second, the relative number of process and product *inventions* (rather than patents) does not vary systematically in response to the implementation of the UTSA. Below, we first explain that our results are inconsistent with changes in innovation behavior due to the strengthening of trade secrets protection. We then implement an instrumental variables estimation similar to Png (2017b) to address concerns about the causal relationship between trade secrets protection and patenting.[30]

4.2.1 Innovation of Products and Processes

It is possible that overall innovative activity increases as trade secrets protection increases at the margin. More specifically, it is possible that more creators of process inventions are affected – those that benefit the most from trade secrecy. This may be the direct result of an increase in the expected returns to an invention, or an indirect result of firms and inventors moving to locations with stronger trade secrets protection.[31] If that were indeed the case, then we would see a relative increase in the number of *process* inventions. Thus, if a strengthening of trade secrets protection affected the creation of different types of innovation differently, then stronger trade secrets protection would likely lead to a relative *increase* in process patents absent changes in patenting behavior of existing inventions.[32] However, we observe a relative *decrease*. Our results can therefore be interpreted as a lower bound of the effect of trade secrets protection.

4.2.2 Instrumental Variables

Our analysis utilizes exogenous variation in trade secrets protection dues to the UTSA. Despite anecdotal evidence that the UTSA was introduced in individual states for "whimsical" reasons, one might still be concerned that states chose to adopt the UTSA when firms were particularly interested in process innovation, compared to other states and years. To address this concern, we follow Png (2017b) and instrument for a state's adoption decision using four other state-level uniform laws as

[30] In the Appendix, we provide further evidence of exogeneity using Propensity Score Matching as well as a placebo test.

[31] Estimating the effects of changes in IP laws on firm location and market concentration is beyond the scope of this paper.

[32] Formally, consider the expression for the share (or probability) of process patents in equation (7). Assume for a moment that π_M and π_P do not change with τ; but let $\theta = \theta(\tau)$ be a function of τ. Then $\rho'(\tau) = \frac{\pi_M \pi_P \theta'(\tau)}{(\pi_P + (\pi_M - \pi_P)\theta(\tau))^2}$. If $\theta'(\tau) > 0$, then the share of process patents increases.

instruments. In particular, the Uniform Determination of Death Act (UDDA), the Uniform Federal Lien Registration Act (UFLRA), the Uniform Durable Power of Attorney Act (UDPAA), and the Uniform Fraudulent Transfer Act (UFTA) were introduced in 1978, 1978, 1979, and 1984, respectively. These four acts are not related to innovation or patenting behavior, but they are related to the UTSA as all were introduced by the Uniform Law Commission to harmonize state regulation around the same time. The identifying assumption is that states which adopted one uniform law early may have also been more likely to adopt other uniform laws early. Png (2017b) provides evidence that this assumption holds.

We therefore create four sets of instruments for a state's level of trade secrets protection. For each uniform law, we introduce a dummy variable that is 1 in state s if the state has implemented the law at the time of a patent's priority date. The first-stage results are strong: the coefficients on all four acts are highly statistically significant, for a first-stage F-statistic of 591.6.[33]

The second-stage results in this instrumental variables regression are shown in Table 4. They support our findings from the baseline estimation, as the coefficients on the trade secrets protection are negative and statistically significant in all four specifications. We continue in the following analyses without instruments to provide more precise estimates, noting that all qualitative results hold if we include the instruments.

4.3 Heterogeneity Effects

Trade secrets have been found to be more important as a means to protect intellectual property for small firms than large firms. A similar degree of heterogeneity is found with respect to technology. We find analogous patterns for the effect of trade secrets protection.[34] As for applicant size, the estimated decrease in the probability that a patent is a process patent is largest for individuals. At the means of the change in trade secrets protection and the initial share of process patents for individuals, the effect corresponds to an average decrease in the probability of a process patent of 6.0% (compared to an average effect of 2.1%). The (negative) impact is smaller for small firms, and no longer statistically significant for large firms. Moreover, much of the effect reported in Table 3 seems to be driven by innovation in the "chemical", "electrical and electronic", "mechanical", and "other" technology categories – that is, NBER categories (Hall et al., 2001). At odds with our theoretical predictions, however, we find a positive effect of trade secrets protection on the probability that a patent is a process patent in the "computers and communications" technology category.

[33] We present the first-stage results in the Appendix.

[34] Detailed estimation results are presented in the Appendix.

Table 4: Impact of Trade Secrets Protection – Instrumental Variables Regressions

	(1)	(2)	(3)	(4)
Trade secrets protection	-0.101***	-0.077*	-0.118***	-0.110**
	(0.037)	(0.042)	(0.042)	(0.054)
Log(indep. claims)		0.232***		0.231***
		(0.004)		(0.004)
Log(length of first claim)		-0.042***		-0.055***
		(0.005)		(0.005)
Log(length of description)		-0.002		0.003
		(0.002)		(0.002)
Originality			0.026***	0.010**
			(0.005)	(0.005)
Generality			0.068***	0.038***
			(0.008)	(0.007)
4th year renewal			0.050***	0.029***
			(0.006)	(0.004)
Observations	1475058	1465095	907867	899932

Notes: Linear probability model with 1[process patent] as the dependent variable, and instrumenting for trade secrets protection with indicators for UDDA, UDPAA, UFTA, and UFLRA adoption. Robust standard errors, clustered by USPC main class and state, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls include indicator variables for the patent's first listed USPC main class, the location state, and the priority year.

4.4 The Number of Patents

Our above results indicate a relative decrease in process patents compared to product patents. But the impact on the absolute number of patents of each type is still unclear. We create a panel at the state-year level to estimate the effect of trade secrets protection on the number of process and product patents. Formally, we estimate

$$patents_{st} = \beta_1 protection_{st} + \gamma_s + \mu_t + \epsilon_{st}, \tag{9}$$

where $patents_{st}$ is the number of (process or product) patents in state s in year t, $protection_{st}$ is the trade secrets protection index, and γ_s and μ_t denote state and priority-year fixed effects, respectively.

Table 5 displays the results of this specification, for process patents (column (1)), product patents (column (2)), and all patents (column (3)). We find that an increase in trade secrets protection decreases the number of both process and product patents. We see an (imprecisely estimated) UTSA-related decrease of 613 process patents per state and year per point increase in the trade secrets protection index. At an average

Table 5: Effect of Trade Secrets Protection on the Number of Patents

	(1) Process	(2) Product	(3) All
Trade secrets protection	-613.337**	-179.248*	-787.421*
	(299.346)	(106.004)	(395.774)
Observations	1683	1683	1683
$\overline{R}^2$	0.109	0.119	0.107

Notes: Fixed effects models with the number of patents as the dependent variables, and the trade secrets protection index as the independent variable of interest. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Fixed effects for the location state and priority year included.

of 358 process patents per state and year before UTSA adoption, and with an average trade secrets protection index change of 0.44 points across states, the point estimate suggests a decrease in patenting of process inventions by 75% on average. The number of product patents decreases with a strenghtening of trade secrets protection as well, albeit less dramatically. At the mean pre-UTSA number of product patents (483), the mean change in trade secrets protection implies a decrease in patenting of product patents of 16%.

5 Welfare Implications

In the previous section, we showed a negative effect (as claimed in Proposition 2) of trade secrets protection on the patenting of processes. Because of a reduction of disclosure of less visible invention, strengthening trade secrets protection can retard knowledge diffusion. In what follows, we evaluate the total welfare effects of stronger trade secrets. We first introduce an augmented three-stage model 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). We then estimate and calibrate the parameters of this three-stage model before simulating the effect of trade secrets protection on total welfare.

5.1 An Augmented Model of Cumulative Innovation

5.1.1 Stage 1 (Initial R&D)

An inventor observes a *potential* invention or R&D project i with characteristics (ϕ, Θ), where ϕ denotes the invention's visibility and Θ its type. Visibility ϕ is drawn from an invention-type specific distribution with cdf F_Θ. The invention type Θ is drawn from a (Bernoulli) distribution $\mathcal{F}$ where $\theta^{\mathcal{F}} = \Pr(\Theta = M)$. The commercial value v_i of the invention is not yet observed. We assume the inventor forms expectations based on the known distribution of v_i.[35] The inventor further observes costs C_i and undertakes the R&D project if the expected payoffs from the invention outweigh its cost. We refer to both F_Θ and $\mathcal{F}$ as *unconditional* distributions, that means, *before* the R&D decision is taken.

5.1.2 Stage 2 (Disclosure or Trade Secret)

The second stage of our augmented model is the model in Section 2. Given the realized invention (upon a positive R&D decision in Stage 1) and observing commercial value v_i, the inventor takes her disclosure decision. The invention is either patented (disclosed) or kept as a trade secret. The expected value of the potential invention in Stage 1 is a function of this disclosure decision in Stage 2. We refer to the type-specific distribution of visibilities and types, G_Θ and $\mathcal{G}$, as *conditional* distributions.

5.1.3 Stage 3 (Follow-on Innovation)

For any potential initial invention i, there is a potential follow-on invention i_F with random value v_{i_F} and cost C_{i_F}, to be realized by another inventor. The realization depends on how much of the initial invention i is visible after the inventor's disclosure decision. We denote the *effective visibility* of initial invention i by $\tilde{\phi}_i$. It is equal to

$$\tilde{\phi}_i = \begin{cases} 0 & \text{if no R\&D in Stage 1;} \\ \phi_i & \text{if R\&D in Stage 1 and trade secret in Stage 2;} \\ 1 & \text{if R\&D in Stage 1 and patent in Stage 2.} \end{cases} \tag{10}$$

Effective visibility is equal to zero if the invention has not been realized and equal to the invention's visibility ϕ_i if the invention is realized but kept as a trade secret. We assume the disclosure function of patents is working, that means, the invention is fully disclosed through patenting. This implies that if the inventor decides to patent her invention in Stage 1, then effective visibility is equal to 1.

We assume that follow-on innovation is probabilistic and depends on the disclosure state of the realized initial invention. Given the effective visibility, the success probability of follow-on innovation is $\tilde{\beta}_{i_F,\tilde{\pi}} = \beta_{\tilde{\pi}}\tilde{\phi}_i$ where $\beta_{\tilde{\pi}}$ is the baseline success probability of follow-on innovation following a realized initial invention with disclosure state $\tilde{\pi}$. For the remainder of our analysis, we assume $\beta_S = 1$ and $\beta_D < \beta_S$.[36]

5.1.4 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 four main

[35] We do not estimate this distribution and therefore, for brevity, refrain from introducing more notation.

[36] We discuss the motivation for this assumption below.

assumptions: First, follow-on innovation is by other firms rather than the inventor of the initial innovation. Consistent with this assumption, 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. Second, disclosure has a positive effect on follow-on innovation. In line with this, Williams (2013) documents that a *restriction* of access to human genome data leads to a 20–40% *reduction* in follow-on research. Similarly, Gross (2019) finds that a policy during World War II to keep certain patent applications secret resulted in fewer citations and slower dissemination of the patented technologies into product markets.

Third, given the effective visibility, the baseline probability of follow-on innovation following a trade secret is higher than that following a patent. This assumption reflects the "anticommons" effect (Heller and Eisenberg, 1998) where technologies are underused because patents on early ideas raise the costs of creating future ideas by creating frictions in the bargaining process over licenses (Scotchmer, 1991; Boldrin and Levine, 2004, 2008b; Green and Scotchmer, 1995; Bessen and Maskin, 2009; Galasso and Schankerman, 2010).[37] For our welfare analysis, we set $\beta_D = {}^2\!/\!_3$, a number consistent with empirical findings in Galasso and Schankerman (2015).[38]

Fourth, we assume that disclosure through patenting is perfect. By law, patent applicants are required to provide a written description of the invention in sufficient detail to allow any person of skill in the field to make and use the invention (35 U.S.C. §112(a)). While the quality of such disclosures has been called into question (Roin, 2005; Fromer, 2009), Furman et al. (2018), for instance, document that the opening of patent libraries (during the pre-internet era) had a positive effect on patenting by local firms.[39]

5.2 Welfare Measure

We use the expected total value added of a given invention, denoted by $W(\tau)$, as our welfare measure. It is calculated as the sum of the aggregate surplus from the initial invention, W_i, and the aggregate surplus from realized follow-on innovation, W_{i_F}.

To construct this expected total value added, let an indicator variable $\mathbf{R}_i(\tau) = 1$ if the inventor makes the initial invention, and $\mathbf{R}_i(\tau) = 0$ if otherwise. The inventor decides to undertake the initial R&D project if $EV_i \geq C_i$. We denote by EV_i the expected gross value of the invention to the inventor; that is, the expected value

[37] Of course, if the need for licensing is not given, the issue does not arise. See Sampat and Williams (2018:227–228) for a discussion of this point. (Robust) results for a model without the anticommons effects so that $\beta_S = \beta_D = 1$ are available upon request.

[38] Using data for U.S. patents, Galasso and Schankerman (2015) find an average increase in forward citations of 50% in response to the invalidation of the cited patent. Gaessler et al. (2018) find an increase of 20% using data for European patents. Related results on the effect of patents rights on follow-on innovation from historical episodes of compulsory licensing can be found in Moser and Voena (2012) and Watzinger et al. (2019).

[39] See Ouellette (2012, 2017) for more evidence on the value of disclosure in patents.

of secrecy, $EV_{S|\Theta}(\tau)$, or of disclosure through patenting, $EV_{D|\Theta}(\tau)$. Moreover, let an indicator variable $\mathbf{R}_{i_F} = 1$ if (conditional on being observed with probability $\tilde{\beta}_{i_F,\tilde{\pi}}$) the follow-on invention is realized (when the commercial value covers the costs, $v_{i_F} \geq C_{i_F}$), and $\mathbf{R}_{i_F} = 0$ otherwise. The expected total value added of an invention i is thus equal to

$$W(\tau) = E_{(\Theta_i,\phi_i,\tilde{\pi}_i,v_i,v_{i_F})}\left[\mathbf{R}_i(\tau)\left(W_i + \tilde{\beta}_{i_F,\tilde{\pi}_i}\mathbf{R}_{i_F}W_{i_F}\right)\right] \tag{11}$$

where expectations $E[\cdot]$ are over the invention type Θ, visibility ϕ, disclosure state $\tilde{\pi}$, and commercial values v for initial and follow-on innovation.

For the measures of aggregate surplus, we assume that $2v$ 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 (at Stage 1) is the potential surplus aggregate surplus net of the disclosure-state specific deadweight loss. The maximum deadweight loss from a scenario with full barriers to access (and certain exclusive use) is $v/2$.[40]

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

$$W_D(\phi) = 2v - \frac{\phi v}{2} - C, \tag{12}$$

where C_i is the cost of research and development 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 2, the probability that the inventor has exclusive access, implying full monopolistic deadweight loss, is equal to $\tau(1-\phi)$. Aggregate surplus, W_S, as a function of visibility and trade secrets protection is equal to

$$W_S(\phi,\tau) = 2v - \frac{\tau(1-\phi)v}{2} - C. \tag{13}$$

In a world with perfect legal protection of trade secrets, $\tau = 1$, the aggregate surplus from trade secrets is a mirror image of the aggregate surplus from a patent, with visibility playing the reverse role. In a world without trade secrets protection, $\tau = 0$, access to the technology is always free (implying a zero deadweight loss) because of disclosure through misappropriation.

[40]For instance, in the textbook case of linear demand with unit market size (and zero marginal cost), non-price discriminating monopoly profits ($=v$) are one half of the aggregate surplus ($=2v$), and consumer surplus and deadweight loss are each one quarter ($=v/2$). In the Online Appendix, we provide a simple competition model to derive the reduced-form aggregate surplus from invention i.

For the aggregate surplus of the initial invention we use

$$W_i = \begin{cases} W_D(\phi) & \text{if } \phi \geq \dfrac{\tau}{1+\lambda+\tau}, \\ W_S(\phi,\tau) & \text{if otherwise.} \end{cases} \tag{14}$$

For the aggregate surplus of follow-on innovation, conditional on initial invention i being realized, we assume free access, so that $W_{i_F} = 2v_{i_F} - C_{i_F}$.

We will for the remainder of the paper assume that the patent premium $\lambda \leq 1/2$, so that the private returns do not exceed the social returns from R&D (both when kept as trade secret and when patented), $W_{\tilde{\pi}_i} \geq V_{\tilde{\pi}_i} - C_i$. With this assumption, the implications from our model are in line with results shown by Bloom et al. (2013).[41]

5.3 Estimation

We use the state- and year-specific trade secrets protection index along with information on the type of U.S. patents to estimate the parameters of our model. We proceed in two steps:

Step 1: We begin by examining the second stage of our augmented model. We use our sample of U.S. utility patents and the trade secrets index in Png (2017a) to estimate the *conditional* distributions $\mathcal{G}$ of invention type Θ, and G_Θ of their type-specific visibilities ϕ_i. These distributions are conditional on the inventor having (successfully) invested in R&D at Stage 1.

Step 2: We then look at the inventor's decision at the first stage. Applying a simulated-method-of-moments approach, we estimate the *unconditional* distributions F_Θ (of visibility ϕ_i) and $\mathcal{F}$ (of invention type Θ), of potential inventions (ϕ, Θ) given R&D costs C_i.

5.3.1 Estimation of Stage-2 Disclosure Decision (Step 1)

We estimate the conditional distributions G_Θ and $\mathcal{G}$ by maximizing a log-likelihood function. We observe a patent's type 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 the patent is a product patent. Moreover, for each patent j, we observe the level of trade secrets protection τ_j

[41] Higher social returns to R&D are typically linked to knowledge spillovers and the public goods aspect of research (Nelson, 1959; Arrow, 1962). The inventor's disclosure decision is socially optimal (with aggregate surplus $W_{\tilde{\pi}}$ as benchmark) only for intermediate values of visibility. The inventor discloses for $\phi \geq \frac{\tau}{1+\lambda+\tau}$. Disclosure is socially optimal and $W_D(\phi) \geq W_S(\phi,\tau)$ if $\phi \leq \frac{\tau}{1+\tau}$. For intermediate values of ϕ, the inventor's decision to disclose is socially optimal. For high values of ϕ, the inventor discloses when it is optimal to keep the invention a secret; for low values of ϕ, the inventor keeps the invention a secret when it is optimal to disclose.

to which the applicant was subject at the time the decision to disclose the invention was made. The log-likelihood of the data is given by

$$LL(G_M, G_P, \mathcal{G}, \lambda) = \sum_j \mathbf{M}_j \log \rho(\tau_j) + (1 - \mathbf{M}_j) \log(1 - \rho(\tau_j)) \tag{15}$$

with $\rho(\tau_j)$ the probability that a patent is a process patent, subject to trade secrets protection τ_j, as derived in equation (7).

The log-likelihood is a function of the (conditional) distributions of visibilities G_Θ, the invention type $\mathcal{G}$, and the patent premium λ. Given data limitations, we estimate our model parameters making a number of assumptions:

1. Visibility ϕ follows a triangular distribution with support $[0, 1]$ and mode γ_Θ. We hold the mode for products constant at $\gamma_P = {}^1\!/_2$ and estimate the mode γ_M for processes. Note that G_P first-order stochastically dominates G_M (as is our working assumption) if $\gamma_M \leq {}^1\!/_2 = \gamma_P$.

2. The patent premium λ is a fixed parameter in our model, and we use values between 0 and ${}^1\!/_2$, based on values estimated in previous literature.[42]

3. We assume a time-variant distribution of invention types. The probability that invention i is a process is denoted by θ_t, $t = 1, \ldots T$. We assume $T = 3$ with θ_1 for all inventions with disclosure decisions from 1976 through 1989, θ_2 for 1990 through 1999, and θ_3 for 2000 through 2008.[43]

We estimate the model on the sample of single-state patents with priority dates between 1976 to 2008.[44] The value for τ_k is the value of the trade secrets protection index in the patent's state in the year of its priority date. We report the results for the conditional distributions from Step 1 in Table 6.

In all models in Table 6, we keep the mode for products constant at $\gamma_P = 0.5$ and estimate the mode for processes γ_M. We therefore obtain the distribution for the visibility of processes relative to the distribution for the visibility of products. This value for γ_P provides for a flexible specification without imposing our theoretical assumption of first-order stochastic dominance. For $\lambda = 0$ in model (1), we find our assumption of first-order stochastic dominance violated. The assumption is satisfied, however, for $\lambda = 0.1$, our preferred value based on existing literature (Schankerman and Schuett, 2017), and $\lambda = 0.5$, the highest value for which the social benefits from R&D outweigh the private benefits (Bloom et al., 2013).

[42] Schankerman and Schuett (2017) estimate a patent premium of approximately 0.1. Similar values are obtained by Lanjouw (1998) and Schankerman (1998) who estimate the patent premium to be an R&D subsidy of 10% and 15–25%, respectively. Arora et al. (2008) further document that for firms with a positive premium, the average patent premium is 50%.

[43] We present results with alternative assumptions about T in the Online Appendix.

[44] For states that have adopted the UTSA, we exclude all patents with priority dates in the year of adoption.

Table 6: Estimates for Conditional Distributions at Stage 2 (Step 1)

		(1)	(2)	(3)
License revenues [fixed]	λ	0.0	0.1	0.5
Mode for processes (G_M)	γ_M	0.572	0.374	0.249
		$[0.539, 0.616]$	$[0.374, 0.391]$	$[0.222, 0.312]$
Mode for products (G_P) [fixed]	γ_P	0.5	0.5	0.5
Share of process inventions (1976–1989)	θ_1	0.327	0.331	0.331
		$[0.325, 0.329]$	$[0.328, 0.333]$	$[0.329, 0.336]$
Share of process inventions (1990–1999)	θ_2	0.475	0.490	0.489
		$[0.473, 0.478]$	$[0.487, 0.491]$	$[0.486, 0.507]$
Share of process inventions (2000–2008)	θ_3	0.575	0.591	0.590
		$[0.573, 0.577]$	$[0.588, 0.593]$	$[0.586, 0.612]$
Observations N (no. of patents)		1,465,351	1,465,351	1,465,351
Log-likelihood/N		-0.672	-0.672	-0.672

Notes: We report the parameter estimates for the conditional distribution from Stage 2 of the augmented model. We estimate our structural model on the sample of single-state patents filed between 1976 and 2008. For states that have adopted the UTSA, we exclude patents from the year the UTSA was adopted. We estimate the mode γ_M (of the triangular distribution over support $[0, 1]$) for processes and fix the mode γ_P for products. Invention types are Bernoulli distributed ($\mathcal{G}$) with parameter θ_t, where $t = 1$ for patents with priority dates in 1976–1989 [$N = 383{,}020$], $t = 2$ for 1990–1999 [$N = 523{,}704$], and $t = 3$ for 2000–2008 [$N = 558{,}627$]. We report in brackets the 99% confidence interval from 800 bootstrap replications. The reported point estimates are from one single model using the full sample.

We use model (2) as our preferred model, both finding support in external sources and comporting with theory. With the parameters in this model, patenting probabilities for processes are lower than products (Lemma 2), decreasing in τ (Lemma 3), and decreasing at different rates so that $\pi_P(\tau) - \pi_M(\tau)$ is increasing (Lemma 4). Stronger trade secrets protection has a relatively stronger (positive) effect on the value of trade secrets when the invention is a process. This implies a decreasing share of process patents as trade secrets protection increases – as predicted in Proposition 2 and shown empirically in Section 4.[45]

Together with the empirical distribution $\hat{\tau}$ of the trade secrets protection index, the estimates of the time-variant innovation type distributions with parameters θ_t such that $\theta_1 < \theta_2 < \theta_3$ imply that the share of process patents, $\rho_t \equiv \rho(\hat{\tau}|\theta_t)$, is increasing over time. The implied process shares are $\rho_1 = 0.327$, $\rho_2 = 0.469$, and $\rho_3 = 0.580$. Our estimates and implied shares closely trace the empirical time series.[46]

[45] Figure E.1 in the Online Appendix illustrates these patterns. In panel (a), we provide a graphical illustration of first-order stochastic dominance, plotting the density (top) and distribution functions (bottom) for G_M (dashed line) and G_P (solid line). We illustrate the effect of τ on $\pi_\Theta(\tau)$ and $\rho(\tau)$ in panels (b) and (c) of the figure.

[46] We show this in panel (d) of Figure E.1 in the Online Appendix. Note that the estimates for θ_t are robust to the models reported in Table 6.

5.3.2 Estimation of Unconditional Stage-1 Distributions (Step 2)

In the second step of our procedure, we estimate the *unconditional* distributions F_Θ of visibilities and $\mathcal{F}$ of invention types, using as input the conditional distributions G_Θ and $\mathcal{G}$ estimated in Step 1. We use the specification and results of our preferred model (2) from Table 6. For this second step, we follow a simulated-method-of-moments approach to find F_Θ and $\mathcal{F}$ that yield (in simulations of Stage 2 of the augmented model) the estimated distributions G_Θ and $\mathcal{G}$. We proceed as follows:

1. For given unconditional distributions $(F_M, F_P, \mathcal{F})$ and some R&D cost C, we simulate a dataset of N potential inventions and solve Stage 1 of our augmented model to obtain the *simulated* conditional distributions, $\delta \in \left\{\hat{G}_M, \hat{G}_P, \hat{\mathcal{G}}\right\}$.

2. We calculate the simulated conditional moments $\hat{\mu}_m(\delta|F_M, F_P, \mathcal{F})$ for the simulated data. We also calculate the estimated moments $\mu_m(\delta)$ based on the estimated conditional distributions G_Θ and $\mathcal{G}$ in Table 6.

3. We define the quadratic score function

$$S(F_M, F_P, \mathcal{F}) = \sum_{\delta} \sum_{m \in \mathcal{M}} \left(\hat{\mu}_m(\delta|F_M, F_P, \mathcal{F}) - \mu_m(\delta)\right)^2 \tag{16}$$

   where $\mathcal{M}$ is the set of moments. We minimize this score function over $(F_M, F_P, \mathcal{F})$ to obtain the optimal unconditional distributions.

In Table 7, we report the parameters of unconditional distributions for no R&D costs ($C = 0$), low costs ($C = 2$), and high costs ($C = 4$).[47] Note that, unlike in Step 1, where we hold G_P constant, in Step 2 we explicitly estimate F_P (i.e., the mode γ_P). Our assumption of first-order stochastic dominance (verified for the conditional distributions) continues to hold. Moreover, for both inventions types, we observe a selection of higher-visibility inventions into development at Stage 2. The average visibility is between 0.368 and 0.445 (compared to 0.458 in Stage 2) for processes, and between 0.397 and 0.486 (compared to 0.500) for products.

The bottom panel of Table 7 shows decisions at all three stages that are implied by the estimated parameters. Our estimates imply relatively large R&D intensities – ranging from 0.592 for high R&D costs to 0.998 without any costs – in Stage 1. In Stage 2, over 79% of realized inventions are indeed patented, and the fraction is larger for lower R&D costs. Finally, at Stage 3, for the no-cost scenario, a bit more

[47] We provide graphical illustrations of these results in the Online Appendix. In panel (a) of Figure E.2, we plot the distribution functions for the estimated conditional distributions G_M (top) and G_P (bottom) as solid lines and the distribution functions for the simulated conditional distributions $\hat{G}_M$ and $\hat{G}_P$ as dashed lines. The graphs illustrate that the simulated conditional distributions well match the estimated distributions. In panel (b), we plot the density functions (top) and probability functions (bottom) of the unconditional distributions F_M and F_P.

Table 7: Estimates for Unconditional Distributions at Stage 1 (Step 2)

			(1)	(2)	(3)
			Stage 1: F_Θ, $\mathcal{F}$		
		Stage 2: G_Θ, $\mathcal{G}$	no cost	low cost	high cost
Mode for processes	γ_M	0.374	0.370	0.335	0.103
Mode for products	γ_P	0.5	0.497	0.458	0.191
Share of processes (1976–1989)	θ_1	0.331	0.329	0.339	0.352
Share of processes (1990–1999)	θ_2	0.490	0.489	0.491	0.501
Share of processes (2000–2008)	θ_3	0.591	0.596	0.595	0.596
R&D intensity (Stage 1)			0.998	0.954	0.592
Patents (Stage 2)			0.858	0.850	0.796
R&D intensity (Stage 3)			0.553	0.465	0.357

Notes: We report the parameter estimates for the unconditional distribution from Stage 1 of the augmented model. For the simulated-method-of-moments approach, we use the first two moments (mean and variance) for G_M and G_P and the first moment (mean) for $\mathcal{G}_t$. 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 (1), $C = 2$ (low cost) in column (2), and $C = 4$ (high cost) in column (3). We further assume that the value of the initial invention and follow-on innovation are (independently) drawn from the same distribution, $v_i, v_{i_F} \sim \mathrm{Exp}(1/10)$. 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.

than one half of all realized initial inventions lead to follow-on innovation, and this intensity is decreasing in R&D costs.

Our estimated values are in line with survey evidence. Mansfield (1986), for instance, finds (for a small sample of innovators for 1981–1983) that in industries in which patenting is relatively important, 84% of patentable inventions are patented, whereas the share is only 66% in other industries.[48]

5.4 Welfare Results

In this section, we vary the level of trade secrets protection to simulate its impact on welfare. We use the estimates for unconditional distributions from Table 7 (with patent premium $\lambda = 0.1$) to simulate a sample of initial inventions i and respective follow-on inventions i_F for each set of parameters, and we calculate our welfare measure $W(\tau)$ in equation (11) based on these.

[48] Mansfield's results suggest that patenting is relatively more important in pharmaceuticals, chemicals, petroleum, machinery, and fabricated metal products, whereas it is of less importance in primary metals, electrical equipment, office equipment, instruments, motor vehicles, rubber, and textiles.

5.4.1 Effect of Trade Secrets Protection

Figure 3 illustrates the simulated welfare impacts under varying levels of trade secrets protection. Panel (a) plots the value of $W(\tau)$ in percent of the value of $W(0)$ for varying levels of R&D costs. For no or low R&D costs, stronger trade secrets protection has an unambiguously negative effect on total welfare. Only for higher R&D costs do we see a positive effect and a welfare improvement from stronger trade secrets protection. This effect comes through various channels. To illustrate these channels, panel (b) of Figure 3 separately depicts the surplus from initial R&D and from follow-on innovation for varying levels of τ, in percent of the value for $\tau = 0$.

First, conditional on the distributions of patents and trade secrets, stronger trade secrets protection results in an anticompetitive effect. This manifests itself in a larger deadweight loss from a trade secret, as captured by the aggregate surplus $W_S(\phi, \tau)$ in equation (13). We can see this effect in the solid-line graph in the top panel of panel (b). Without R&D costs, all R&D projects are realized regardless of τ. The negative effect of higher τ is thus the result of a decrease in $W_S(\phi, \tau)$.

Second, conditional on innovation taking place, higher trade secrets protection leads to fewer inventions being disclosed through a patent. This has a negative effect on overall welfare $W(\tau)$ in equation (11) through $\tilde{\beta}_{i_F,\tilde{\pi}}$. The overall effective visibility decreases, which in return reduces the probability of follow-on innovation. We observe this negative effect on follow-on innovation in the dashed graphs in panel (b). As τ increases, the value of follow-on innovation decreases unambiguously.

Third, τ also affects the decision to innovate. 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 $W(\tau)$ in equation (11) through $\mathbf{R}_i(\tau)$. We observe this effect in panel (b) for positive R&D costs. For high R&D costs in particular, the positive effect through higher investment incentives more than offsets the negative effect on W_S.

Finally, the increased R&D activity implies there is more initial R&D to build on. This should counteract the negative effect on follow-on innovation laid out above, especially when R&D costs are high. We can observe this when we compare the dashed graph in panel (b) for the value of follow-on innovation for high costs with that for low costs. For higher costs, 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 for high costs than for low costs.

5.4.2 Different Distributions of Visibilities

We next use counterfactual distributions for the visibilities of process and products to better understand the role that differences in distributions for invention type play. Setting $\theta_t = 0.5$ for all t for convenience, we illustrate the results of this exercise in Figure 4. We compare the results from three scenarios to the total value from the estimated distributions from Table 7. In scenario 1 (solid line), we assume equal

Figure 3: Effect of Trade Secrets Protection on Welfare

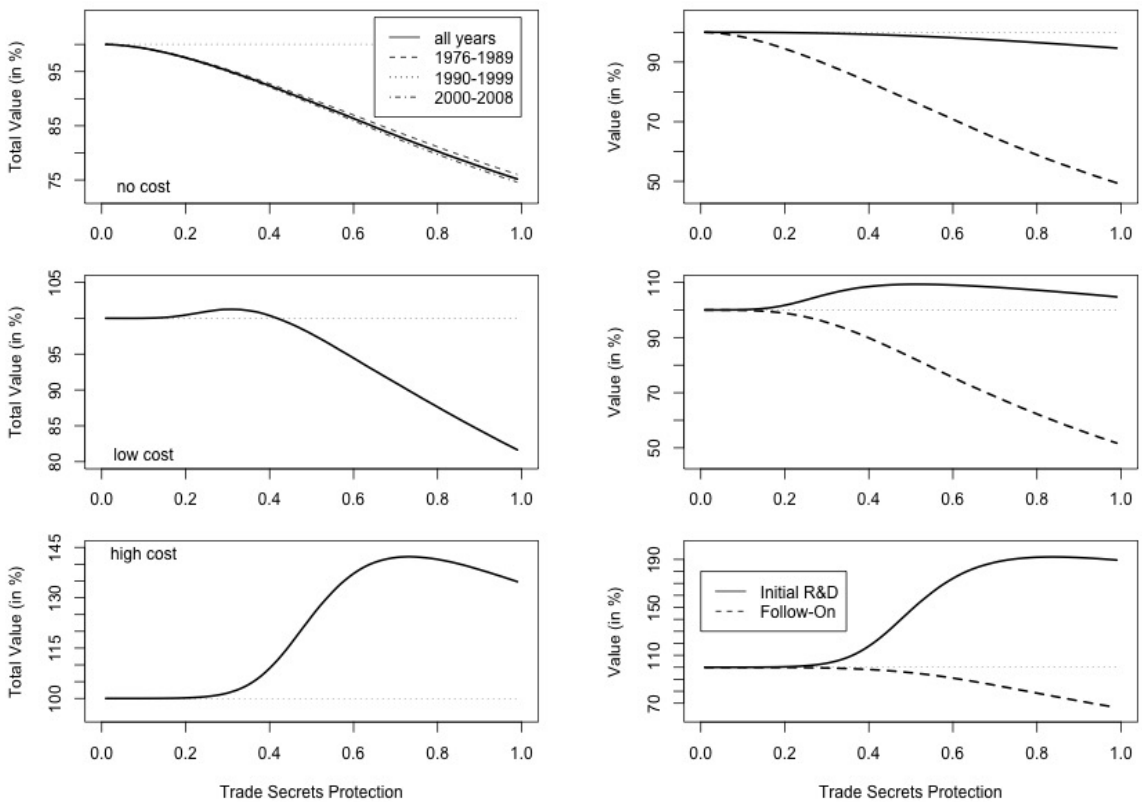

(a) Total Welfare

(b) Initial and Follow-On

Notes: This figure presents our welfare results. In panel (a), we plot the welfare function $W(\tau)$ (in % of $W(0)$). For values of $\tau \in [0, 1]$, 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 $\beta_S = 1$ and $\beta_D = 2/3$. We show the total value for our entire sample period (where a proportional number of simulated inventions have θ_t) as well as for the three subsample periods (for no cost). In panel (b), we plot the social value of initial R&D (solid) and follow-on innovation (dashed), again in % of the value for $\tau = 0$. For the top panels, we use the estimates for $C = 0$ (no cost); for the center panels, we use the estimates for $C = 2$ (low cost); and for the bottom panels, we use the estimates for $C = 4$ (high cost).

Figure 4: Visibility and the Effect of Trade Secrets Protection

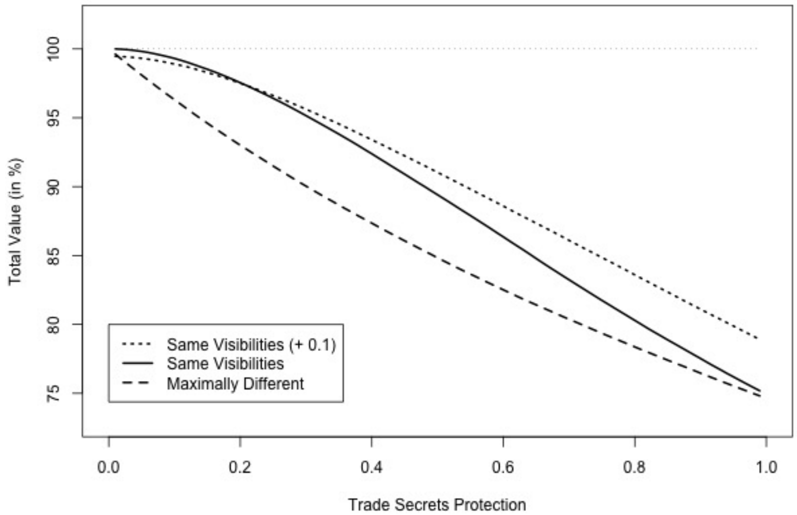

Notes: In this figure, we illustrate the effect of visibilities of different invention types on total welfare for the no-cost scenario ($C = 0$) from Figure 3. We plot total welfare for equal distributions for the two invention types (solid) line and maximally different distributions (dashed line) while keep the overall mean of visibility constant. More specifically, for *Same Visibilities*, we set $\theta_t = 0.5$ for all t and $\gamma_M = \gamma_P = \hat{\gamma}$ where $\hat{\gamma}$ is such that the mean of the triangular distribution with model $\hat{\gamma}$ is equal to the mean of estimated unconditional distribution. For *Maximally Different* we set $\gamma_M \geq 0$ as low as possible and $\gamma_P \leq 1$ as high as possible such that the overall mean is equal to the mean of estimated unconditional distribution. The estimated values are based on simulated data with N = 1,000,000.

distributions that imply the same mean visibilities as the estimated model (we calculate the mean value of visibilities from the estimated unconditional distribution in Table 7). In scenario 2 (dotted line), we assume equal distributions but increase the modes of the visibilities $\gamma_M = \gamma_P$ by 0.1. In scenario 3 (dashed line), we assume maximally different distributions, setting $\gamma_M = 0$ and $\gamma_P < 1$ so that, again, the mean is equal to the mean in the estimated model.

Comparing scenarios 1 and 2, we find that higher visibilities are associated with higher welfare. Higher visibilities enter the welfare function in three ways. Higher visibility implies more patenting (Lemma 1), and with higher patenting comes a higher deadweight loss (equation (12)). At the same time, higher patenting as well as higher visibilities increase effective visibility $\tilde{\phi}_i$ and thus increase follow-on innovation (equation (10)). Our results in Figure 4 show that the latter effect prevails.

By comparing scenarios 1 and 3, we can see what happens when the distributions of visibilites become more diverse – and products become on average more visible than processes, while overall average visibility remains constant. We find that stronger distributional differences have negative welfare effects. Welfare is consistently lower for the scenario with the maximally different distributions. This is evidence for a central role of visibilities in the welfare calculations.

Figure 5: Average Welfare Effect of the UTSA

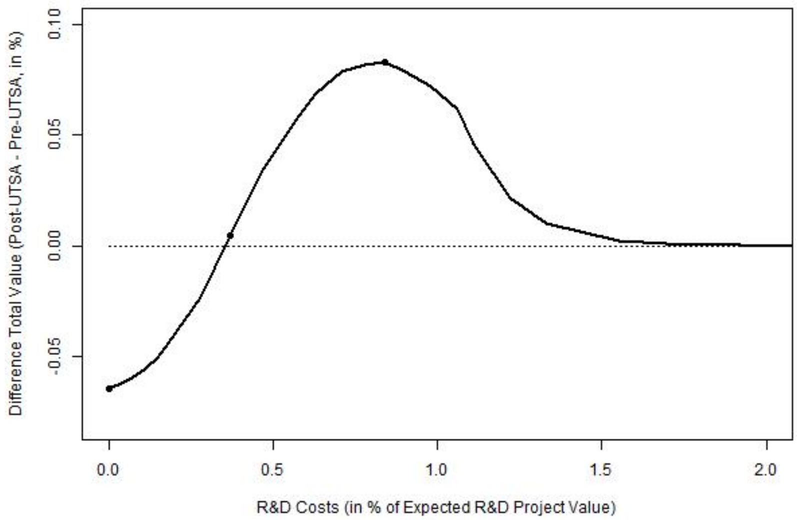

Notes: In this figure, we show the average welfare effect of the introduction of the Uniform Trade Secrets Act. We plot ΔW in equation (17), that is, the difference between total welfare (as fraction of pre-UTSA total welfare) evaluated at the average post-UTSA value of the trade secrecy index, $\tau^{\text{post}} = 0.394$, and the total welfare evaluated at the average pre-UTSA value, $\tau^{\text{pre}} = 0.071$. On the horizontal axis, we use R&D costs as fraction of the expected R&D project value (given expectations of invention type, visibility, commercial value, and the inventor's patenting decision). We mark the values of no cost, low costs, and high cost used in Figure 3.

5.4.3 Average Welfare Effect of the UTSA

We have shown the effect of hypothetical levels of trade secrets protection, $\tau \in [0, 1]$. Last, we calculate the average effect of the Uniform Trade Secrets Act as predicted by our model. We simulate data from our augmented model for the average value of trade secrets protection before the adoption of the UTSA, $\tau^{\text{pre}} = 0.071$, and after the adoption, $\tau^{\text{post}} = 0.394$. We then calculate the difference between the post-UTSA and pre-UTSA total welfare as fraction of pre-UTSA total welfare,

$$\Delta W = \frac{W(\tau^{\text{post}}) - W(\tau^{\text{pre}})}{W(\tau^{\text{pre}})}, \tag{17}$$

so that negative values of ΔW imply that the Uniform Trade Secrets Act had, on average, a negative effect on total welfare. We plot this average effect for varying values of R&D costs in Figure 5.

We find a negative effect of the UTSA for no R&D costs, a zero effect for low costs, and a positive for higher costs (as fraction of the expected R&D project value). More specifically, our model predicts a welfare loss of 7% for no R&D costs and a gain of 8% with high costs. These result suggest that in industries with relatively profitable R&D (that is, where R&D costs are very low), stronger trade secrets protection as implemented by the UTSA has negative effects on total welfare. In other words, when the benefits of trade secrets protection are inframarginal to an inventor who is deciding whether to invest in R&D, stronger legal protection of trade secrets has the unintended consequence of lowering total welfare by impeding follow-on innovation. This pattern is reversed for R&D projects that are relatively less profitable (when R&D costs are higher). In this case, that is, when the benefits are marginal, stronger legal protection improves welfare by encouraging initial R&D.[49]

The survey findings in Mansfield (1986) and our own estimates give some indication about which industries are more likely to have benefitted from increased trade secrets protection (pharmaceuticals and chemicals, and to a lower degree petroleum, machinery, and fabricated metal products) and in which industries the UTSA is likely to have resulted in welfare losses (electrical equipment, office equipment, motor vehicles, instruments, primary metals, rubber, and textiles).[50]

6 Conclusion

Intellectual property law tries to find a balance between encouraging initial innovation by granting monopoly rights to one's invention, and facilitating follow-on innovation by limiting these rights. While the effects of intellectual property rights on incentives to innovate in the first place are relatively well-understood, the incentives to use and build upon existing innovation have received less attention. The value of a patent in facilitating follow-on innovation depends largely on the original idea's visibility. For less visible innovations, a patent implies disclosure of an idea that may have otherwise not been accessible by others. On the other hand, patents for visible inventions limit the ability of others to use said innovation, which could otherwise be easily re-engineered.

This paper examines the role of intellectual property law in facilitating follow-on innovation by distinguishing between (less visible) process and (more visible) product inventions. We find that a strengthening of trade secrets protection leads to relatively fewer process invention, leading to a patenting culture that is less conducive to follow-on innovation. These results are promising, however: if relatively stronger trade secrets protection can lead to fewer process patents, then a relative weakening of trade secrets protection may lead to more process invention being patented and thus disclosed. Our results suggest that there is an optimal level of trade secrecy protection – depending on the costs of R&D – that encourages initial innovation *and* disclosure of the "right" inventions.

[49] Note that as R&D costs increase further, the average welfare effect converges to zero. This is because, as costs increase, few realized projects are kept as trade secrets, and stronger trade secrets protection becomes ineffective.

[50] Our estimates for the share of patented inventions that would not have been developed without any patent protection are increasing in R&D costs. Mansfield (1986) finds the highest share (30%) in pharmaceuticals and chemicals, a lower share (10–20%) in petroleum, machinery, and fabricated metal products, and the lowest share in the remaining industries.

References

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.

Arora, A., M. Ceccagnoli, and W. M. Cohen (2008): "R&D and the Patent Premium," *International Journal of Industrial Organization*, 26, 1153–1179.

Arrow, K. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, ed. by R. Nelson, Princeton University Press, 609–626.

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

Bessen, J. and E. Maskin (2009): "Sequential Innovation, Patents, and Imitation," *RAND Journal of Economics*, 40, 611–635.

Bloom, N., M. Schankerman, and J. van Reenen (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81, 1347–1393.

Boldrin, M. and D. K. Levine (2004): "Rent-Seeking and Innovation," *Journal of Monetary Economics*, 51, 127–160.

——— (2008a): *Against Intellectual Property*, New York, N.Y.: Cambridge University Press.

——— (2008b): "Perfectly Competitive Innovation," *Journal of Monetary Economics*, 55, 435–453.

Chung, S., K. Han, A. Animesh, and A. Pinsonneault (2015): "Competitive Impacts of IT Innovation: An Empirical Analysis of Software Patents in the IT Industry," Unpublished manuscript, available at `https://ssrn.com/abstract=2505119`.

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.

Dowd, M. J. and J. Crotty (2016): "The Limited Impact Of Provisional Patent Rights," Law360 (February 11, 2016), available at `https://www.law360.com/articles/752125/the-limited-impact-of-provisional-patent-rights`.

Federal Trade Commission (2003): "To Promote Innovation: The Proper Balance of Competition and Patent Law and Policy," Report, Federal Trade Commission, Washington, D.C.

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 (2018): "Disclosure and Subsequent Innovation: Evidence from the Patent Depository Library Program," Working Paper 24660, National Bureau of Economic Research, Cambridge, Mass.

Gaessler, F., D. Harhoff, and S. Sorg (2018): "Patents and Cumulative Innovation: Evidence from Post-Grant Patent Oppositions," Unpublished manuscript.

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

——— (2015): "Patents and Cumulative Innovation: Causal Evidence from the Courts," *Quarterly Journal of Economics*, 130, 317–369.

Ganglmair, B., W. K. Robinson, and M. Seeligson (2019): "Four Decades of Patent Claim Drafting: Theory, New Data, and Evidence," Manuscript in preparation.

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

Graham, S. and D. Hegde (2012): "Do Inventors Value Secrecy in Patenting? Evidence from the American Inventor's Protection Act of 1999," Unpublished manuscript.

——— (2015): "Disclosing Patents' Secrets," *Science*, 347, 236–237.

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.

Green, J. R. and S. Scotchmer (1995): "On the Division of Profit in Sequential Innovation," *RAND Journal of Economics*, 26, 20–33.

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.

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.

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.

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.

Kuhn, J. M. and N. C. Thompson (2017): "The Ways We've been Measuring Patent Scope are Wrong: How to Measure and Draw Causal Inferences with Patent Scope," Unpublished manuscript.

Lanjouw, J. O. (1998): "Patent Protection in the Shadow of Infringement: Simulation Estimations of Patent Value," *Review of Economic Studies*, 65, 671–710.

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. (2008): "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.

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

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

Moser, P. and A. Voena (2012): "Compulsory Licensing: Evidence from the Trading with the Enemy Act," *American Economic Review*, 102, 396–427.

Naqi, S. (2012): "Comment on Provisional Patent Rights," *Northwestern Journal of Technology and Intellectual Property*, 10, 595–606.

Nelson, R. R. (1959): "The Simple Economics of Basic Scientific Research," *Journal of Political Economy*, 67, 297–306.

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

——— (2017): "Who Reads Patents?" *Nature Biotechnology*, 35, 421–424.

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

Png, I. P. (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.

Pooley, J. and D. P. Westman (1997): *Trade Secrets*, New York, N.Y.: Law Journal Seminars-Press.

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

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

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.

Samuelson, P. and S. Scotchmer (2002): "The Law and Economics of Reverse Engineering," *Yale Law Journal*, 111, 1575–1663.

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.

Schankerman, M. and F. Schuett (2017): "Screening for Patent Quality: An Economic Framework for Policy Analysis," CEPR Discussion Paper 11688, Centre for European Policy Research, London, UK.

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

Scotchmer, S. and J. Green (1990): "Novelty and Disclosure in Patent Law," *RAND Journal of Economics*, 21, 131–146.

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.

Thursby, J. G., C. Haeussler, M. C. Thursby, and L. Jiang (2018): "[Prepublication Disclosure of Scientific Results: Norms, Competition, and Commercial Orientation]," *Science Advances*, 4.

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.

Watzinger, M., T. A. Fackler, M. Nagler, and M. Schnitzer (2019): "How Antitrust Enforcement Can Spur Innovation: Bell Labs and the 1956 Consent Decree," Unpublished manuscript, LMU Munich.

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.

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

Appendix

A Formal Proofs of Theoretical Results

Proof of Lemma 1

Proof. The proof follows from the disclosure decision in equation (3). □

Proof of Proposition 1

Proof. For the proof of this claim and later results, it will be useful to first state a more general property of first-order stochastic dominance. 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

$$\int u(x)dG_P(x) = [u(x)G_P(x)]_0^1 - \int u'(x)G_P(x)dx$$

and

$$\int u(x)dG_M(x) = [u(x)G_M(x)]_0^1 - \int u'(x)G_M(x)dx$$

Because $G_P(0) = G_M(0) = 0$ and $G_P(1) = G_M(1) = 1$, the two first RHS terms in these expression are equal. We can thus 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, $EV_{S|M}(\tau) > EV_{S|P}(\tau)$, note that $\tau(1-\phi)v$ is a strictly decreasing function in ϕ. The above general property applies because we can simply write $-EV_{S|M}(\tau) < -EV_{S|P}(\tau)$. 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. $\square$

Proof of Lemma 2

Proof. Because $\pi(\phi,\tau)$ is a non-decreasing function in ϕ, the general property in equation (A.1) in the proof of Proposition 1 applies. $\square$

Proof of Lemma 3

Proof. Because $\pi(\phi,\tau)$ is (weakly) decreasing in τ for all ϕ, the first derivative of $\pi_\Theta(\tau)$ with respect to τ,

$$\frac{d\pi_\Theta(\tau)}{d\tau} = \int_0^1 \frac{\partial \pi(\phi,\tau)}{\partial \tau} dG_\Theta(\phi), \tag{A.2}$$

is non-positive for $\Theta = M, P$. $\square$

Proof of Lemma 4

Proof. What is to be shown is

$$\frac{d\pi_P(\tau)}{d\tau} - \frac{d\pi_M(\tau)}{d\tau} = \int_0^1 \frac{\partial \pi(\phi,\tau)}{\partial \tau} dG_P(\phi) - \int_0^1 \frac{\partial \pi(\phi,\tau)}{\partial \tau} dG_M(\phi) \geq 0.$$

The cross-derivative of $\pi(\phi,\tau)$ is negative, $\frac{\partial^2 \pi(\phi,\tau)}{\partial\tau\partial\phi} < 0$. As ϕ increases, $\frac{\partial \pi(\phi,\tau)}{\partial \tau}$ is less negative and $\frac{\partial \pi(\phi,\tau)}{\partial \tau}$ increasing in ϕ. The general property in equation (A.1) in the proof of Proposition 1 applies. $\square$

Proof of Proposition 2

Proof. The proof follows from the result in Lemma 4 and the expression for the share of process patents in equation (7). $\square$

B Evidence from the American Inventors Protection Act (1999)

The analysis in the main text relies on the assumption that processes are less visible and patents covering processes are more difficult to enforce. Given this assumption, Proposition 1 implies that inventors of processes *should* be more likely to keep their inventions a secret. When they are given the choice, we expect process inventors to opt for secrecy more often – even if secrecy is only temporary.

We can test this implication of our working assumption by exploiting the enactment of the American Inventors Protection Act of 1999 (AIPA). The AIPA went into effect for all patent applications filed on or after November 29, 2000. It came with two important changes. First, all pending patent applications filed on or after the cutoff date are by default published 18 months after the filing date. This marks a significant change in policy as until the USPTO did not publish pending patent applications, but published only granted patents. Second, U.S.-only patents, for which applicants do not seek foreign protection, can opt out of automatic pre-grant publication.[51]

Because all patented inventions are (trade) secrets until the application is published, opting out of pre-grant publication represents a temporary extension of secrecy – until the granting of the patent. In 2001, the lag between filing a patent application and grant averaged about 38 months (Graham and Hegde, 2015), implying that opting out of pre-grant publication extended temporary secrecy by about 20 months. Graham and Hegde (2015) find that about 15% of all eligible patent applicants (filing after the effective date of the AIPA and asserting U.S.-only patent protection) opt out of pre-grant publication and choose temporary secrecy.[52]

We extend Graham and Hegde's data and analysis by adding our process patent

[51] There was a third, yet arguably ineffective, change. This last provision provision grants patent applicants provisional rights during the pendency of the patent (35 U.S.C. §154(d)). The provision adds (once the patent is granted) an applicant's right to collect reasonable royalties for infringement that occurred during the pendency of its application. This time window of *pre-grant infringement damages* begins with the pre-grant publication of the patent application. Pre-grant infringement damages, however, are limited. First, an applicant has provisional rights conditional on the patent being granted. This means, an applicant cannot sue for infringement damages while the application is still pending. Second, the provision requires actual notice of the alleged infringer, and the asserted claim must not substantially change in the examination process (between publication of the application and the granted patent) (Naqi, 2012; Dowd and Crotty, 2016).

[52] See Graham and Hegde (2012) for an extended version with additional results and details on the AIPA.

indicator and compare the applicants' choices of opting out of pre-grant publication across patent types. Our Proposition 1 implies that applicants of process patents will opt out of disclosure via pre-grant publication of their applications more often than those of product patents. Although opting out of publication grants the inventor only temporary secrecy, it is secrecy nonetheless. Results should therefore only give us a lower bound.

Our results comport with the prediction in Proposition 1 and provide support for our working assumption. Applicants of eligible patents with at least one process claim choose to keep their applications secret 16.1 percent of the time, whereas applicants of patents without any process claims choose secrecy only 13.5 percent of the time. The difference is highly statistically significant with a t-value of 25.8.[53] Figure B.1 plots the monthly shares of patent applications (of granted patents) that were opted out of pre-grant publication, distinguishing between applications with and without process claims. At any time in the five-year period after the AIPA (December 2000 through December 2005), a larger fraction of applicants of process patents (relative to product patents) decided to (temporarily) extend the secrecy of their patent applications.

More formal regression analyses, in which we estimate the probability that applicants of eligible patents opt out of pre-grant publication, controlling for patent and applicant characteristics, support these trends. In particular, we estimate the probability that a patent application (after passing of the AIPA) is kept secret until the patent's issuance. We estimate

$$secrecy_{jt} = \beta_1 process_{jt} + \beta_2 X_{jt} + \lambda_u + \mu_t + \epsilon_{jat}, \tag{B.1}$$

where the dependent variable is 1 if patent application j in year t is kept secret until the patent is granted. The independent variable of interest, $process_{jt}$, is 1 if the patent includes at least one process claim, X_{jt} includes patent-specific measures of complexity (the number of individual claims and the length of the first claim) and ex-post value (measures of the invention's originality and generality as well as indicators for whether the maintenance fee was paid after four, eight, and twelve years). We further include dummy variables for the patent's USPC class (λ_u) and the year of application (μ_t). Finally, we cluster standard errors by USPC main class to allow for

[53] The differences in means are similar when using our two alternative patent type indicator, that means, when comparing applications by predominantly product or process patents and and by first claim only.

Figure B.1: Probability of Extending Temporary Secrecy of Patent Applications

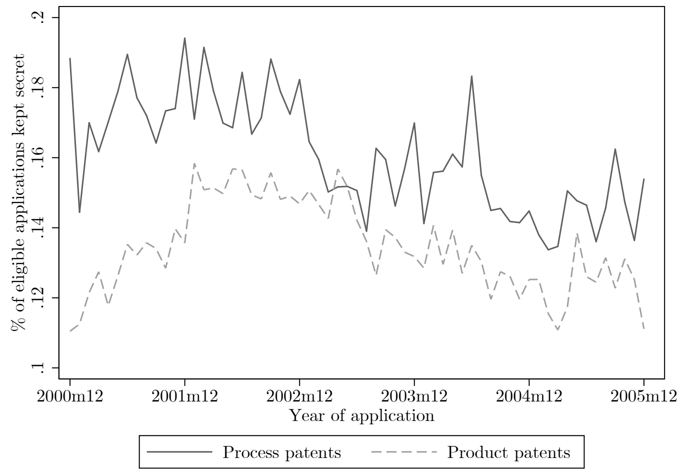

Notes: This figure plots monthly shares of applicants (of eligible patents) who opted out of pre-grant publication, by patent type (process or product), for granted patents whose applications were filed within the first five years after the AIPA went into effect. Note that for this picture (and analysis in this Appendix), we follow Graham and Hegde (2015) and use the application date (which is the relevant date for the option to opt out of publication).

common trends within these classes.

Table B.1 reports results of a linear probability model. It confirms what Figure B.1 suggests: even after controlling for patent specific characteristics, applicants of process patents are more likely to opt out of application disclosure when given the choice. The estimated decrease of 0.9 percentage points (column (5)) implies a decrease of 5.8% at the mean of 15.4% of patent applicants choosing secrecy. Overall, this evidence confirms the first proposition and thus also the identifying assumption that (on average) process inventions are less visible than product inventions.

C Heterogeneity Effects

The observation that the impact of stronger trade secrets protection depends on the invention's visibility informs an important policy debate. Trade secrets have been found to be more important as a means to protect intellectual property for small firms than large firms. A similar degree of heterogeneity is found with respect to technology. Hall et al. (2014) provide a comprehensive survey of the literature. We analyze the effect of stronger trade secrets protection on applicants of different size as well as on inventions in different technology classes.

Table B.1: Secrecy/Disclosure of Patent Applications After the AIPA

	(1)	(2)	(3)	(4)
Process patent (= 1)	0.015***	0.010***	0.015***	0.011***
	(0.003)	(0.003)	(0.004)	(0.003)
Log(indep. claims)		0.024***		0.021***
		(0.003)		(0.003)
Log(length of first claim)		0.009***		0.013***
		(0.003)		(0.003)
Log(length of description)		-0.023***		-0.026***
		(0.002)		(0.002)
Originality			-0.016**	-0.009
			(0.008)	(0.008)
Generality			0.051***	0.059***
			(0.005)	(0.005)
4th year renewal			-0.004	-0.000
			(0.007)	(0.006)
Observations	479379	479379	270839	270839
$\overline{R}^2$	0.055	0.058	0.058	0.062

Notes: Linear probability model with 1[application is kept secret] as the dependent variable, and 1[process patent] as the independent variable of interest. Robust standard errors, clustered by USPC main class, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls include indicator variables for the patent's first listed USPC main class and the year of application.

C.1 Firm Size

We consider three different sizes of patent applicants: individuals (=1), small firms (= 2), and large firms (=3).[54] We interact each of these with the trade secrets protection index, and we re-estimate equation (8) with these interactions. Table C.1 shows the coefficients of interest from these regressions without control variables (column (1)), with controls for the patent's complexity (column (2)), with controls for the patent's value (column (3)), and with all control variables (column (4)). Individuals are the omitted category. The estimated decrease in the probability that a patent is a process patent is largest for individuals, with an estimated decrease of 4.6 percentage points if the trade secrets protection index increases by 1 full point in column (4). At the means of the change in trade secrets protection and the initial share of process patents

[54]For more details on how we construct our size index, see the Online Appendix.

Table C.1: Effect of Trade Secrets Protection by Applicant Size

	(1)	(2)	(3)	(4)
Small firm	0.050***	0.027***	0.043***	0.021***
	(0.003)	(0.003)	(0.003)	(0.003)
Large firm	0.062***	0.038***	0.055***	0.035***
	(0.004)	(0.003)	(0.004)	(0.004)
Individual × Trade secrets protection	-0.037***	-0.041***	-0.051***	-0.046***
	(0.009)	(0.008)	(0.009)	(0.009)
Small firm × Trade secrets protection	-0.022**	-0.019**	-0.024**	-0.020**
	(0.009)	(0.009)	(0.010)	(0.009)
Large firm × Trade secrets protection	-0.002	-0.007	-0.008	-0.012
	(0.012)	(0.011)	(0.011)	(0.011)
Control variables: complexity	N	Y	N	Y
Control variables: value	N	N	Y	Y
Observations	1460358	1450434	905492	897578
$\overline{R}^2$	0.301	0.345	0.291	0.337

Notes: Linear probability model with 1[process patent] as the dependent variable. Firm size = 1 if applicant is an individual, = 2 if applicant is a small firm, and = 3 if applicant is a large firm. Omitted category: individuals before UTSA adoption. Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Control variables as in Table 3. Additional controls include indicator variables for the patent's first listed USPC main class, the location state, and the priority year.

for individuals, this corresponds to an average decrease in the probability of a process patent of 6.0% (compared to an average effect of 2.1%). The (negative) impact is smaller for small firms, and no longer statistically significant for large firms.

C.2 Technology Classes

We examine the variation of the effect of trade secrets protection across different technologies by estimating equation (8) with interactions of the trade secrets protection index with each NBER technology category.[55] Table C.2 shows the results from these regressions. Much of the impact reported in Table 3 seems to be driven by innovation in the "Chemical" (1), "Electrical and Electronic" (4), "Mechanical" (5), and "Other" (6) categories, while a stronger trade secrets protection *increases*

[55] Hall et al. (2001) construct six broad technology categories based on USPC main classes. These categories are Chemical (1), Computers & Communications (2), Drugs & Medical (3), Electrical & Electronic (4), Mechanical (5), and Others (6).

Table C.2: Effect of Trade Secrecy Protection by Technology Class

	(1)	(2)	(3)	(4)
Chemicals × TS protection	-0.064***	-0.060***	-0.059***	-0.053***
	(0.014)	(0.013)	(0.015)	(0.014)
Computers × TS protection	0.065***	0.061***	0.053***	0.046***
	(0.014)	(0.013)	(0.015)	(0.013)
Drugs × TS protection	-0.027	-0.020	-0.019	-0.017
	(0.021)	(0.020)	(0.020)	(0.019)
Electronics × TS protection	-0.010	-0.016	-0.033**	-0.036**
	(0.015)	(0.014)	(0.015)	(0.014)
Mechanics × TS protection	-0.031**	-0.036***	-0.040***	-0.038***
	(0.015)	(0.014)	(0.014)	(0.014)
Other × TS protection	-0.033***	-0.039***	-0.038***	-0.037***
	(0.010)	(0.010)	(0.010)	(0.010)
Control variables: complexity	N	Y	N	Y
Control variables: value	N	N	Y	Y
Observations	1475039	1465093	907858	899931
$\overline{R}^2$	0.300	0.345	0.289	0.337

Notes: Linear probability model with 1[process patent] as the dependent variable. We report interaction terms of the trade secrets protection index with NBER technology categories (Hall et al., 2001). These categories are Chemical (=1), Computers & Communications (=2), Drugs & Medical (=3), Electrical & Electronic (=4), Mechanical (=5), and Others (=6). Robust standard errors in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls include indicator variables for the patent's first listed USPC main class, the location state, and the priority year.

the probability that a patent in the "Computers and Communications" (2) category covers a process. Accordingly, a technology-specific patenting and trade secrets policy may facilitate follow-on innovation.

D Additional Robustness Checks

The main analysis makes several assumptions. First, we focus our analysis on single-state patents, that means, patents for which all assignees and inventors are from the same U.S. state. We take this conservative approach to avoid assigning patents to the "wrong" states. Second, we define a patent as a process patent if it includes at least one process claim, although the invention itself may still exhibit characteristics of a product. Third, we assume that states that adopted the UTSA early and states that adopted the UTSA later or not at all are similar and their patent portfolios would have evolved in the same way, were it not for the UTSA.

Here, we test the robustness of our results to these assumptions. First, we examine the strength of the instruments on the instrumental variables regressions by reporting the first stage results. Next, we use different definitions for a patent applicant's location, for the process patent variable, and for the level of trade secrecy protection. Finally, we use propensity score matching to choose the most appropriate control groups, and we test the exogeneity assumption in a Placebo exercise.

D.1 First Stage Results

Our instrument variables estimation relies on two assumptions. First, the instruments are unrelated to the dependent variable in the second stage. Second, they are strongly related with the endogenous variable. The former assumption is likely to hold because the laws we utilize as instruments do not concern innovation and patenting decisions. The latter is also likely to hold: bureaucratic red tape that slows down the state-specific implementation of one law may also affect the implementation of another state-specific law. Here, we provide further evidence that this assumption holds. Table D.1 shows the coefficients and partial F-statistic of the first stage. The coefficients on all instruments are strongly statistically significant, and the F-statistic is well beyond any critical value at 591.6.

D.2 Alternative Definitions of Process Patents and Location

The main analysis uses rather conservative definitions of both a process invention and the innovator's location. We define a patent as a process patent if at least one claim describes a process, and we consider only those patents for which all U.S. entities are located in the same state. Here, we relax these definitions. First, we utilize two alternatives measures of a process patent: (1) a patent is a process patent if the first claim is a process claim, and (2) a patent is a process patent if at least 50% of all independent claims are process claims. Second, we assign as the patent's location the location of the assignee who is listed first, and if no assignee is listed, we use the location of the first listed inventor.

Table D.3 shows the coefficients on the change in trade secrecy protection due to the UTSA, replicating the specification from column (4) of Table 3. The first three

Table D.1: First Stage Results of IV Regression (Trade Secrets Protection)

	(1) DV: UTSA Index
UDDA	0.0180***
	(0.0052)
UDPAA	-0.0975***
	(0.0035)
UFTA	0.0740***
	(0.0034)
UFLRA	0.0396***
	(0.0052)
Observations	1,487,739
$\overline{R}^2$	0.7911
F-stat for all instruments	591.58***

Notes: Dependent variable is the effective trade secrets protection index. Robust standard errors, clustered by USPC main class and state, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls include the complexity and value variables from the main analysis, as well as indicator variables for the patent's first listed USPC main class, the state, and the priority year.

columns use the same location definition as the main text, whereas the last three columns define the invention's location as the state of the first listed U.S. assignee, or the first U.S. inventor if no U.S. assignee is listed. For each location definition, we use three definitions for a process patent: (1) at least one claim is a process claim, (2) the first claim is a process claim, and (3) at least half of all claims are process claims. Column (1) of Table D.3 therefore exactly replicates column (4) of Table 3.

The results are robust to most definitions. While the location definition does not affect the qualitative results (and the quantitative results are similar as well), the estimated effect disappears when we define a patent as a process patent only if at least half of its claims describe a process – the most restrictive definition. Importantly, previous literature leads us to believe that the most appropriate definition of a process patent is either the first or the second. For instance, Kuhn and Thompson (2017) argue that under U.S. law the broadest claim should be listed first, suggesting the second definition should be used. Our main analysis uses the first definition because we are interested in disclosure of any process – regardless of the its role in the patent.

Table D.2: Alternative Process Patent Indicators

	Before UTSA		After UTSA		
	Mean	Std. Dev.	Mean	Std. Dev.	T-stat
Process Patent ['First']	0.2592	0.4382	0.3098	0.4624	67.5
Process Patent ['Most']	0.2208	0.4148	0.2505	0.4333	42.1
Observations	674,186		813,214		

Notes: This table provides summary statistics for the alternative patent indicators for all granted utility patents (between 1976 and 2014) with a priority date between 1976 and 2008, and for which the location can be unambiguously identified. The indicator variable 'Most' is defined for all patents for which we can identify the claim type for at least one independent claim (666,131 patents before UTSA adoption and 808,980 patents after UTSA adoption). The 'First' indicator is defined only for patents for which we can identify the claim type of the first claim: 660,720 patents before the UTSA and 804,217 patents after the UTSA.

Table D.3: Effect of the UTSA Using Different Definitions

	Single-State Patents			Broader U.S. Patents		
	(1) 'Any'	(2) 'First'	(3) 'Most'	(4) 'Any'	(5) 'First'	(6) 'Most'
Trade secrets protection	-0.026***	-0.022***	-0.008	-0.028***	-0.029***	-0.005
	(0.008)	(0.007)	(0.006)	(0.008)	(0.007)	(0.006)
Control vars: Complexity	Y	Y	Y	Y	Y	Y
Control vars: Value	Y	Y	Y	Y	Y	Y
Observations	899932	894030	899932	1446801	1436218	1446801
$\overline{R}^2$	0.337	0.307	0.278	0.336	0.303	0.268

Notes: Linear probability model with 1[process patent] as the dependent variable and the trade secrets protection index as the independent variable of interest. Column (1) is the same as column (4) of Table 3. Columns (1) through (3) look only at single-state patents. Columns (4) through (6) consider the location-state of the first assignee (if present) as the patent's location, and the first inventor if no assignee is listed. Columns (1) and (4) consider all patents with any process claims as process patents. Columns (2) and (5) define a patent's type by its first claim. Columns (3) and (6) consider a patent as a process patent if it has more process than product claims. Robust standard errors, clustered by USPC main class and state, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls are identical to the main analysis.

Table D.4: Effect of UTSA (Binary Indicator)

	(1)	(2)	(3)	(4)
After UTSA adoption	-0.008**	-0.011***	-0.012***	-0.013***
	(0.004)	(0.004)	(0.004)	(0.004)
Control vars: Complexity	N	Y	N	Y
Control vars: Value	N	N	Y	Y
Observations	1475058	1465095	907867	899932
$\overline{R}^2$	0.300	0.345	0.289	0.337

Notes: Linear probability model with 1[process patent] as the dependent variable and a binary variable that is 1 if the state has adopted the UTSA as the independent variable of interest. Robust standard errors, clustered by USPC main class and state, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls are identical to the main analysis.

D.3 UTSA as a Binary Variable

The main analysis uses the trade secrets protection index derived by Png (2017a) as the dependent variable of interest. One might instead be interested in the effect of the UTSA adoption itself, as it created well-defined rules and norms regarding trade secrecy protection, which may not have previously been present regardless of the initial strength of the law.

We examine the impact of the UTSA adoption here, repeating the analysis from Table 3 but using an indicator variable that is 1 if the state has adopted the UTSA (and zero otherwise), rather than the trade secrets protection index. Table D.4 reports the coefficients of interest from these regressions, controlling for the same variables as the main analysis. The estimated effect is even larger than that in the main estimation, with an average decrease in the probability that a patent includes a process claim of 1.05 percentage points, or 2.6% at the mean.

D.4 Propensity Score Matching

Our main analysis compares patents across all U.S. states, regardless of previous patenting trends and other demographic patterns. However, it is possible that some states are more closely related to each other than others. For instance, innovators in Nebraska may have more in common with innovators from Wyoming than with those from California. To account for this possibility, we match states that adopted the

UTSA before 1990 ("treated" states) with those adopting the UTSA later or not at all ("control" states) based on their propensity to adopt the UTSA early.

In particular, we estimate the probability of early adoption as a logistic function of several variables describing each state in 1980 – before any state adopted the UTSA: the number of patents per population, the share of process patents, the governor's party affiliation and the percent of the population voting for Ronald Reagan in the 1980 presidential election, and the state capital's latitude and longitude. In addition, we control for the year of the state's UDDA adoption. Our logit estimation provides a reasonably good fit, with a pseudo R-squared of 0.21. We match the states based on their estimated propensity scores, dropping those states that do not have a close neighbor.

Importantly, our matching approach allows us to illustrate the annual effect of the UTSA on the likelihood that a patent includes a process. Figure D.1 depicts the average share of process patents for treated and control states, in each year before and after the treated state's UTSA adoption.[56] The share of process patents increases for all states – and the trends for treated and control states are similar prior to the treated state's UTSA adoption. However, there is a small dip among treated states (solid line) just after UTSA adoption that is not seen among the control states.

This dip is also seen in a regression that mirrors the analysis from Table 3, using a binary variable for UTSA adoption instead of the trade secrets protection index used in the main analysis to allow for a more direct interpretation of the results. This regression, using both complexity and value controls (similar to column (4) of Table 3), estimates a decrease in the probability that a patent includes a process invention of 1.05 percentage points (significant at the 1 percent level) due to the UTSA.

D.5 Placebo Test – Earlier UTSA Adoption

One might still 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. In that case, we might see a significant change in the likelihood that a patent covers a process invention *before* a state adopts the UTSA. We examine this possibility in a Placebo test. Instead of the true UTSA adoption date, we set that

[56] For each matched pair, we assign the year of UTSA adoption by the "treated" state, and we count the years before and after that date accordingly for both states.

Figure D.1: Average Share of Process Patents Around UTSA Adoption

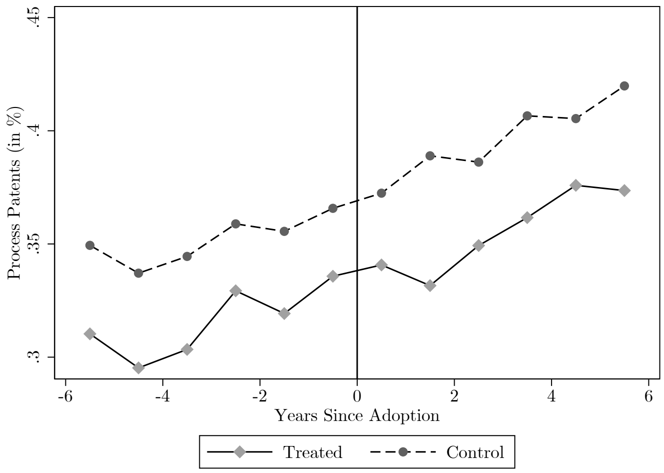

Notes: The figure shows the average share of process patents among all patents for the propensity score matched states. The solid line denotes the treated states (states that adopted the UTSA before 1980), and the dashed line shows the matched control states. For the treated states, year 0.5 denotes the year of the state's UTSA adoption, 1.5 is the year after that, etc. For the control states, the year of adoption is assigned as that of its matched treated counterpart. State pairs for which UTSA adoption is less than 5 years apart are dropped.

date two years earlier, dropping all patents that were applied for after the true UTSA adoption.[57] We then estimate the effect of the fake UTSA adoption on the likelihood that a patent is a process patent.

Table D.5 shows the coefficients of interest for specifications that mirror those in Table 3. While all specifications return a negative point estimate for the coefficient on the fake UTSA adoption, these estimates are smaller than those using the true UTSA adoption, and they are statistically insignificant throughout. These results provide evidence that the states adopted the UTSA exogenously with respect to the distribution of product and process patents.

[57]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.

Table D.5: Placebo Test: Effect of (Fake) UTSA

	(1)	(2)	(3)	(4)
After placebo UTSA adoption	-0.005	-0.004	-0.007	-0.007
	(0.004)	(0.004)	(0.004)	(0.004)
Control vars: Complexity	N	Y	N	Y
Control vars: Value	N	N	Y	Y
Observations	213024	209045	141063	137744
$\overline{R}^2$	0.282	0.325	0.273	0.318

Notes: Linear probability model with 1[process patent] as the dependent variable and a binary variable that is 1 in the two years before the state adopted the UTSA as the independent variable of interest. All observations after the state's actual adoption are dropped. Robust standard errors, clustered by USPC group and state, in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Additional controls are identical to the main analysis.

1204

Appendix for Online Publication

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

Bernhard Ganglmair[*]
ZEW, MaCCI, University of Mannheim

Imke Reimers[†]
Northeastern University

March 19, 2019

[*]ZEW, MaCCI, and University of Mannheim. Address: ZEW – Leibniz Centre for European Economic Research, L7, 1, 68161 Mannheim, Germany; E-mail: b.ganglmair@gmail.com; Twitter: @ganglmair

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

Online Appendix

E Additional Tables and Figures for Structural Results

E.1 Conditional and Unconditional Distributions

In Figure E.1, we show the estimated conditional density and probability functions for processes and products (panel (a)); the probability of patenting for processes and products from equation (6) (panel (b)); the share of process patents as function of τ from equation (7) (panel (c)); and the estimated share of process patents over time (panel (d)). Panel (b) depicts the results for Lemmata 2, 3, and 4.

In Figure E.2, we juxtapose the estimated conditional distriutions and the simulated conditional distributions from simulations using the unconditional distributions (panel (a)) and the estimatedun conditional density and probability functions for processes and products (panel (b)).

E.2 Estimated Distributions (by R&D Costs)

In Figure E.3, we plot the mode of the estimated unconditional distributions (Step 2) of visibilities for processes (dashed line) and products (dotted line). Analogous to the graph in Figure 5, we vary R&D costs and plot the outcome against R&D in % of Expected R&D Project Value.

As R&D costs increase and fewer initial ideas are realized, inventions become on average less visible. For no R&D costs, the conditional and unconditional distributions are the same as all initial inventions (unconditional) are realized (conditional). With higher R&D costs, we observe selection. In order for the conditional distributions to be realized (recall: the conditional distribution is constant, not dependent on the counterfactual value of C), the initial distributions must change with C. For sufficiently high costs, we hit the lower bound of $\gamma_\Theta = 0$.

In Figure E.4, we plot the implied patent propensity against R&D costs, for all realized inventions (solid line) and separately for processes (dashed line) and products (dotted line). First, observe that process inventions are less likely patented than product inventions. This is true across all values of R&D costs. Second, the implied

Figure E.1: Results from Structural Model (Conditional Distributions)

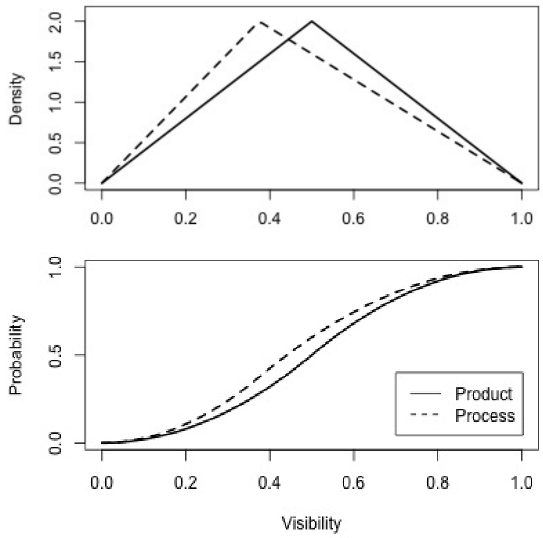

(a) PDF and CDF

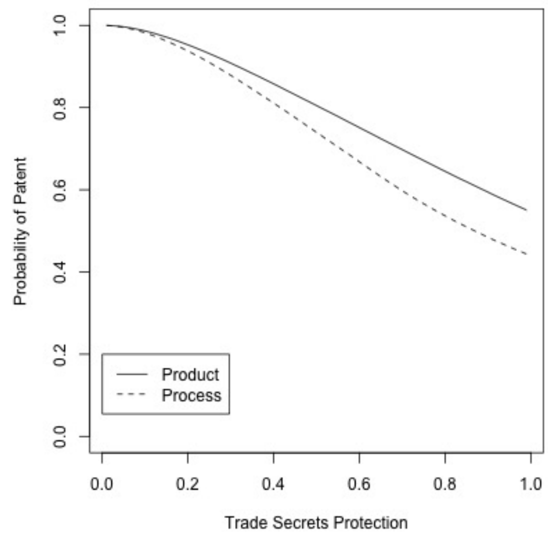

(b) Patenting Probability

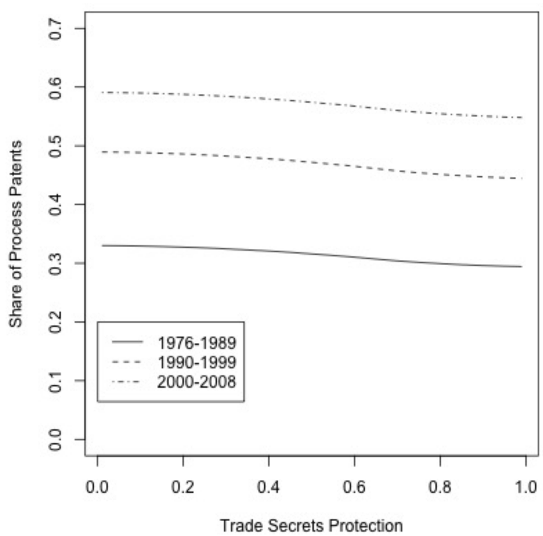

(c) Share of Process Patents (by τ)

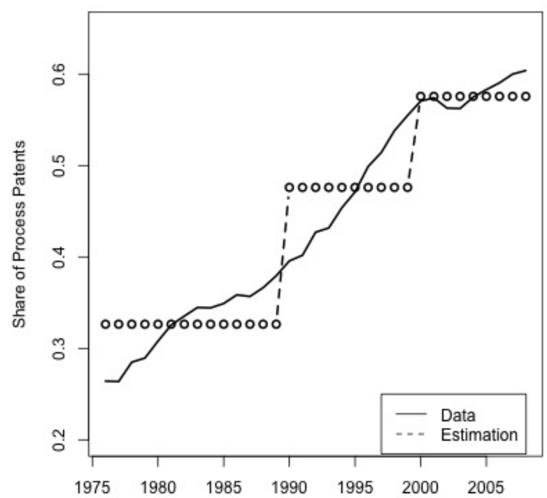

(d) Share of Process Patents (by year)

Notes: We depict the estimation results (Step 1) for model (2) in Table 6. In panel (a), we plot the density and probability functions of visibilities for products and processes. For panel (b), we plot the patenting probabilities $\pi_\Theta(\tau)$ (by invention type Θ) as function of trade secrets protection τ. For panel (c), we plot the share of process patents $\rho(\tau)$ as function of trade secrets protection (τ) for three different estimates of θ_t. For panel (d), we plot the share of process patents $\rho(\tau)$ over time (see Figure 2). The solid line depicts annual process shares from the data, the dash-dotted line depicts the estimated values given θ_t and the empirical distribution of τ for the respective t. Panels (b) through (d) are based on simulated data with $N = 1{,}000{,}000$.

Figure E.2: Results from Structural Model (Unconditional Distributions)

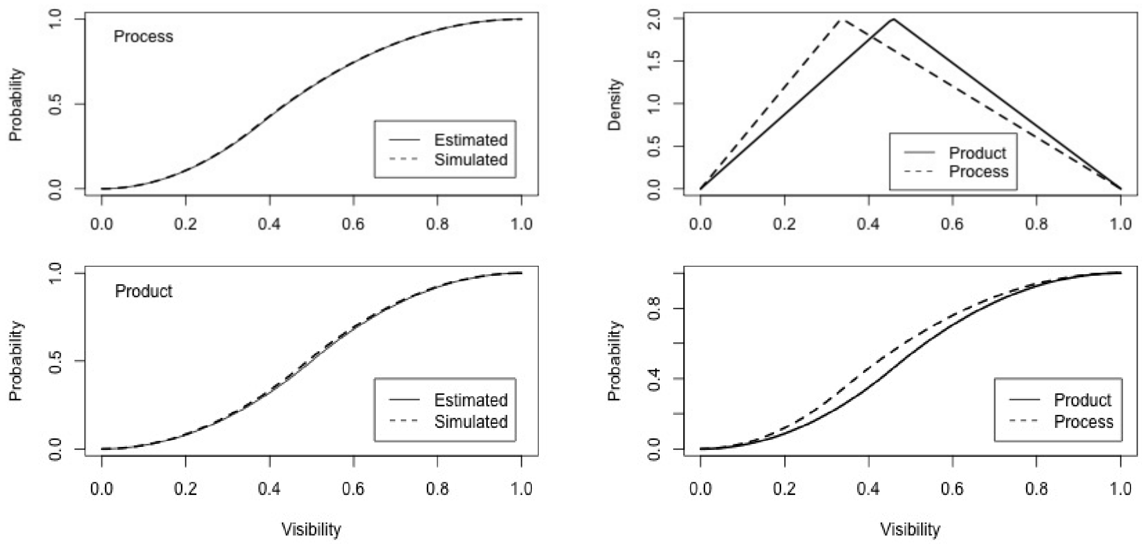

(a) Conditional CDFs　　**(b)** Unconditional PDF and CDF

Notes: We depict the estimation results from Step 2 of our estimation procedure for $C = 2$ (low cost). In panel (a), we plot the conditional distributions of visibilities for process (top panel) and products (bottom panel) from Step 1 ("estimated;" solid line) and the implied conditional distributions from simulations using the unconditional distributions from Step 2 ("simulated;" dashed line). In panel (b), we plot the density and probability functions of visibilities for products and processes for the unconditional distributions (Step 2).

patent propensities initially decrease as costs increase but increase (and converge to one) for higher R&D costs.

E.3 Time-Varying Distribution of Invention Types

For the specification of the structural model in the main text, we use a time-varying distribution of invention types with $T = 3$ different values for the share of process inventions, θ_t for $t = 1, 2, 3$. In Tables E.1 and E.2, we present estimation results for $T = 6$ and $T = 7$ with θ_t for $t = 1, \ldots, 6$ and $t = 1, \ldots, 7$. Our results are robust. First, our estimates of γ_M satisfy our assumption of first-order stochastic dominance (now also for $\lambda = 0$). Second, our estimates for the distribution of invention types imply an increasing share of realized process inventions. In Figures E.5 and E.6, we also plot the empirical and implied share of process patents. The solid line depicts annual process shares from the data, the dash-dotted line depicts the estimated values given θ_t and the empirical distribution of τ for the respective t.

Figure E.3: Unconditional Distributions (Modes of Triangular Distribution)

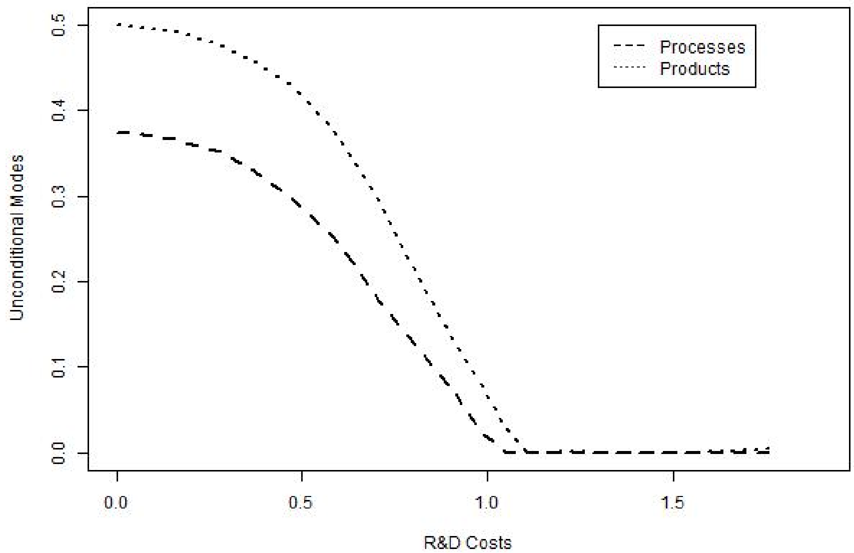

Notes: In this figure, we plot the estimated modes of the triangular distribution for visibilities of processes (dashed line) and products (dotted line). On the horizontal axis, we use R&D costs as fraction of the expected R&D project value (given expectations of invention type, visibility, commercial value, and the inventor's patenting decision).

Figure E.4: Implied Patent Propensities

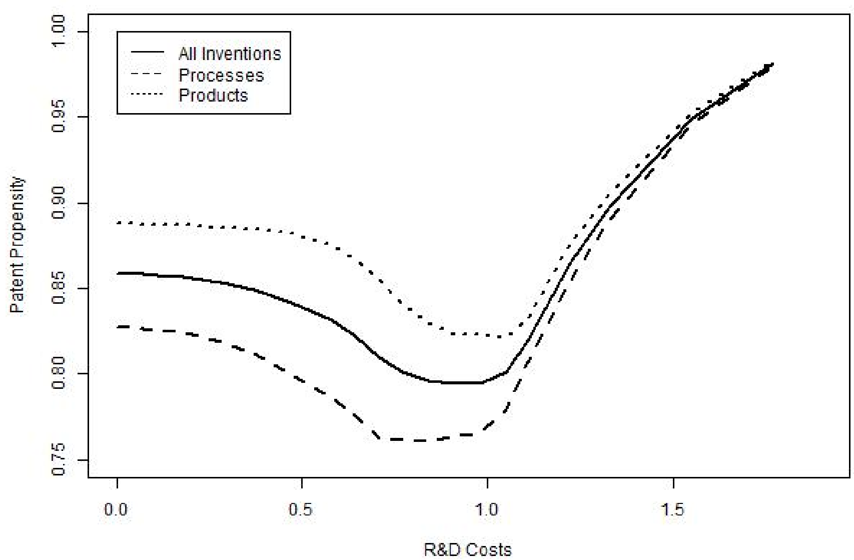

Notes: In this figure, we plot the implied patent propensities for all realized inventions (solid line), processes (dashed line), and products (solid line). On the horizontal axis, we use R&D costs as fraction of the expected R&D project value (given expectations of invention type, visibility, commercial value, and the inventor's patenting decision).

Figure E.5: Share of Process Patents ($T = 6$)

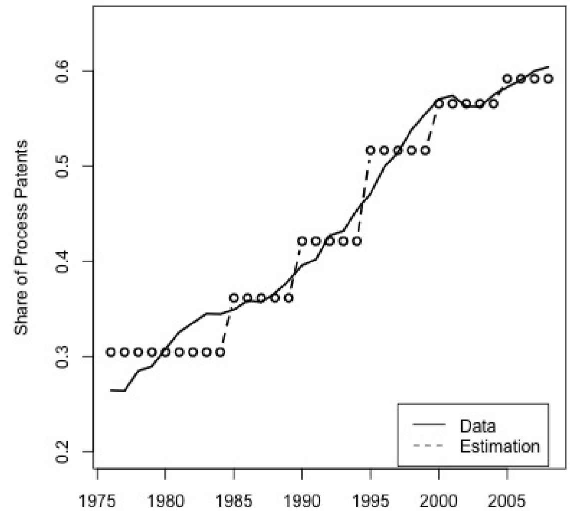

Notes: In this figure, we plot the share of process patents $\rho(\tau)$ over time (see Figure 2). The solid line depicts annual process shares from the data, the dash-dotted line depicts the estimated values given θ_t and the empirical distribution of τ for the respective t, where $t = 1, \ldots, 6$. The parameter estimates are reported in Table E.1, the estimated values are based on simulated data with $N = 1{,}000{,}000$.

Figure E.6: Share of Process Patents ($T = 7$)

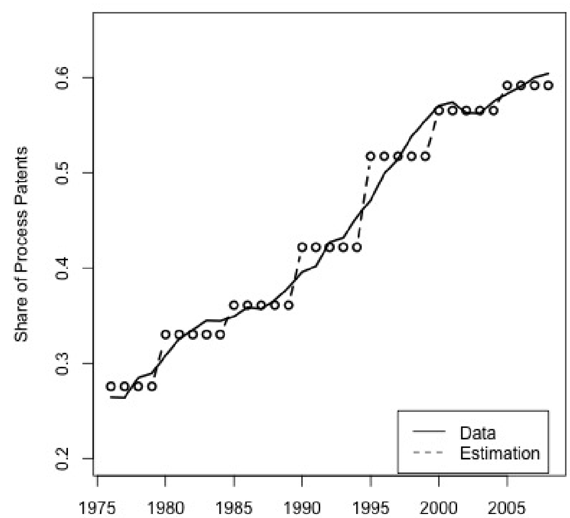

Notes: In this figure, we plot the share of process patents $\rho(\tau)$ over time (see Figure 2). The solid line depicts annual process shares from the data, the dash-dotted line depicts the estimated values given θ_t and the empirical distribution of τ for the respective t, where $t = 1, \ldots, 7$. The parameter estimates are reported in Table E.2, the estimated values are based on simulated data with $N = 1{,}000{,}000$.

Table E.1: Estimates for Conditional Distributions ($T = 6$)

		(1)	(2)	(3)
License revenues [fixed]	λ	0.0	0.1	0.5
Mode for processes (G_M)	γ_M	0.463	0.374	0.266
Mode for products (G_P) [fixed]	γ_P	0.5	0.5	0.5
Share of process inventions (1976–1984)	θ_1	0.305	0.306	0.308
Share of process inventions (1985–1989)	θ_2	0.365	0.369	0.368
Share of process inventions (1990–1994)	θ_3	0.425	0.433	0.435
Share of process inventions (1995–1999)	θ_4	0.520	0.530	0.530
Share of process inventions (2000–2004)	θ_5	0.571	0.581	0.581
Share of process inventions (2005–2008)	θ_6	0.596	0.606	0.608
Observations N (no. of patents)		1,465,351	1,465,351	1,465,351
Log-likelihood/N		-0.67	-0.67	-0.67

Notes: We report the parameter estimates for the conditional distribution from Stage 2 of the augmented model with six time periods. We estimate our structural model on the sample of single-state patents filed between 1976 and 2008. For states that have adopted the UTSA, we exclude patents from the year the UTSA was adopted. We estimate the mode γ_M (of the triangular distribution over support $[0, 1]$) for processes and fix the mode γ_P for products. Invention types are Bernoulli distributed ($\mathcal{G}$) with parameter θ_t, where $t = 1$ for patents with priority dates in 1976–1984 [N = 232,450], $t = 2$ for 1985–1989 [N = 127,825], $t = 3$ for 1990–1994 [N = 177,685], $t = 4$ for 1995–1999 [N = 253,815], $t = 5$ for 2000–2004 [N = 261,483], and $t = 6$ for 2005–2008 [N = 166,751]. The reported parameter estimates maximize the log-likelihood in equation (15).

E.4 Results from Subsample Excluding Software Patents

We rerun our analysis of Step 1 with a subsample that excludes software patents. Following Chung et al. (2015:Table 2), we identify software patents as patents with United States Patent Classification (USPC) main classes 341, 345, 370, 380, 382, 700–707, 710, 711, 713–715, 717, 726, and 902. For the sample used in the analysis in the text, we have already excluded business methods patents with USPC main class 705 (e.g. Lerner, 2006). We report the results from the sample without software patents in Table E.3. In Figure E.7, we also plot the empirical and implied share of process patents. The solid line depicts annual process shares from the data (without software patents – the dotted line depcits the annual shares including software patents), the dash-dotted line depicts the estimated values given θ_t and the empirical distribution of τ for the respective t.

Table E.2: Estimates for Conditional Distributions ($T = 7$)

		(1)	(2)	(3)
License revenues [fixed]	λ	0.0	0.1	0.5
Mode for processes (G_M)	γ_M	0.436	0.367	0.249
Mode for products (G_P) [fixed]	γ_P	0.5	0.5	0.5
Share of process inventions (1976–1979)	θ_1	0.277	0.276	0.274
Share of process inventions (1980–1984)	θ_2	0.331	0.333	0.333
Share of process inventions (1985–1989)	θ_3	0.368	0.369	0.366
Share of process inventions (1990–1994)	θ_4	0.429	0.434	0.434
Share of process inventions (1995–1999)	θ_5	0.523	0.531	0.530
Share of process inventions (2000–2004)	θ_6	0.574	0.582	0.580
Share of process inventions (2005–2008)	θ_7	0.599	0.607	0.607
Observations N (no. of patents)		1,465,351	1,465,351	1,465,351
Log-likelihood/N		-0.67	-0.669	-0.67

Notes: We report the parameter estimates for the conditional distribution from Stage 2 of the augmented model with seven time periods. We estimate our structural model on the sample of single-state patents filed between 1976 and 2008. For states that have adopted the UTSA, we exclude patents from the year the UTSA was adopted. We estimate the mode γ_M (of the triangular distribution over support $[0, 1]$) for processes and fix the mode γ_P for products. Invention types are Bernoulli distributed ($\mathcal{G}$) with parameter θ_t, where $t = 1$ for patents with priority dates in 1976–1979 [N = 109,264], $t = 3$ for 1980–1984 [N = 123,186], $t = 3$ for 1985–1989 [N = 127,825], $t = 4$ for 1990–1994 [N = 177,685], $t = 5$ for 1995–1999 [N = 253,815], $t = 6$ for 2000–2004 [N = 261,483], and $t = 7$ for 2005–2008 [N = 166,751]. The reported parameter estimates maximize the log-likelihood in equation (15).

F A Simple Competition Model

In this section, we derive the reduced-form social surplus functions in equations (12) and (13) from a simple competition model. We derive the expressions for process inventions; the case for product invention is analogous.

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, $p^m_L \leq c_H$. Moreover, for simplicity let $c_L = 0$. The monopoly profits in this case are $\pi^m_L = \frac{1}{4}$.

Now, suppose the firm has chosen to patent the technology. This means, all potential competitors have (restricted) access to the technology. The patent holder 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

Table E.3: Estimates for Conditional Distributions (Step 1) – Excluding Software Patents

		(1)	(2)	(3)
License revenues [fixed]	λ	0.0	0.1	0.5
Mode for processes (G_M)	γ_M	0.537	0.374	0.374
Mode for products (G_P) [fixed]	γ_P	0.5	0.5	0.5
Share of process inventions (1976–1989)	θ_1	0.322	0.325	0.323
Share of process inventions (1990–1999)	θ_2	0.448	0.461	0.453
Share of process inventions (2000–2008)	θ_3	0.523	0.538	0.529
Observations N (no. of patents)		1,328,068	1,328,068	1,328,068
Log-likelihood/N		-0.673	-0.673	-0.673

Notes: We report the parameter estimates for the conditional distribution from Stage 2 of the augmented model. We estimate our structural model on the sample of single-state patents filed between 1976 and 2008. For states that have adopted the UTSA, we exclude patents from the year the UTSA was adopted. We also exclude software patents, using the list of USPC main classes in Chung et al. (2015:Table 2) to identify software patents. We estimate the mode γ_M (of the triangular distribution over support $[0, 1]$) for processes and fix the mode γ_P for products. Invention types are Bernoulli distributed ($\mathcal{G}$) with parameter θ_t, where $t = 1$ for patents with priority dates in 1976–1989 [N = 372,333], $t = 2$ for 1990–1999 [N = 480,708], and $t = 3$ for 2000–2008 [N = 475,027].

Figure E.7: Share of Process Patents (Excluding Software Patents)

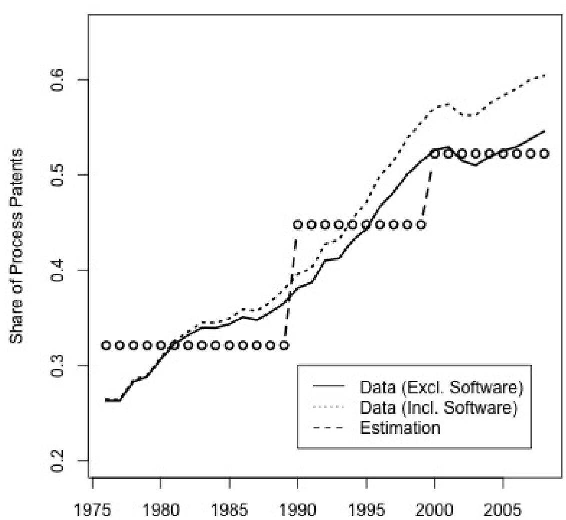

Notes: In this figure, we plot the share of process patents $\rho(\tau)$ over time (see Figure 2) for the model without software patents. The solid line depicts annual process shares from the data, the dash-dotted line depicts the estimated values given θ_t and the empirical distribution of τ for the respective t, where $t = 1, 2, 3$. The parameter estimates are reported in Table E.1, the estimated values are based on simulated data with N = 1,000,000.

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

$$\phi\frac{3}{2\pi_L^m} + (1-\phi)\cdot 0 = 2\pi_L^m - \frac{\phi\pi_L^m}{2}. \tag{F.1}$$

Instead of a patent, let the firm keep the technology a secret. As discussed in the Section 2, the firm has exclusive access to the techology 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. 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

$$\begin{aligned}\tau\,(1-\phi)\,\frac{3}{2\pi_L^m} + [1-\tau\,(1-\phi)]\cdot 2\pi_L^m = \\ 2\pi_L^m - \frac{\tau\,(1-\phi)\,\pi_L^m}{2}.\end{aligned} \tag{F.2}$$

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 (12) and (13).

G Data Appendix

We construct our data sample using a number of sources. We obtain basic bibliographic information from PatentsView at `https://www.patentsview.org/download` for bulk download and `http://www.patentsview.org/api/doc.html` for API queries. We also use data from Ganglmair et al. (2019) for process patent indicators, the USPTO's Patent Maintenance Fee Events database at `https://bulkdata.uspto.gov/data/patent/maintenancefee` to calculate our proxies for patent value as well as applicant size, the USPTO's Patent and Patent Application Claims Research Dataset at `https://bulkdata.uspto.gov/data/patent/claims/economics/2014/` for proxies of patent scope and complexity, and the Google Patents Research Data at `https://console.cloud.google.com/marketplace/partners/patents-public-data`

Table G.4: Sample Construction and Sample Size

Sample/Variable	Source	Obs.
Patents, granted January 1976 – December 2014	PatentsView	X
Priority dates: January 1976 – December 2008	Google Patents	X
U.S. only location	*constructed*	X
Exclude business method patents	PatentsView	X
	Main Estimation Sample:	X
Process patent indicator	Ganglmair et al. (2019)	X
Number of independent claims	USPTO Claims	X
Length of first claim	USPTO Claims	X
Length of detailed patent description	PatentsView (API)	X
Originality	*constructed*	X
Generality	*constructed*	X
4th year maintenance	USPTO Maintenance	X
USPC main classes	PatentsView	X
Applicant size	*constructed*	X
NBER technology categories	PatentsView	X

Notes: Data sources are PatentsView (bulk data download page and API), Google Patents (Google Patents Research Data), USPTO Claims (USPTO's Patent and Patent Application Claims Research Dataset), USPTO Maintenance (USPTO's Patent Maintenance Fee Events database), and Ganglmair et al. (2019). Constructed means that variables are constructed/calculated by authors. For more details, see the descriptions below.

to construct data on the timing of disclosure.[1] In Table G.4, we provide an overview of the steps of our sample construction. For further details, see the descriptions that follow.

G.1 Main Sample

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

Timing: Priority Dates To identify the timing of the disclosure decision, we use a patent's priority date. More specifically, we use the priority date of the head

[1]We thank Jeffrey Kuhn for his support with Google's Big Query.

of a simple patent family (i.e., all patents that share the same priority claims). We implement this by using the earliest priority date for all patents from a given simple patent family. Information on simple patent family assignment and priority dates we obtain from the Google Patents Research Data.

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

For a set of robustness results in the Appendix, we use a more aggressive location definition. There, we define the location of a patent by the location of the first assignee listed on the granted patent. If no assignee is listed, we use the location of the first inventor listed on the granted patent.

G.2 Patent Classification

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

For our main estimation sample, we exclude all business methods patents. We follow Lerner (2006) and define such patents as those with USPC main class 705 (i.e., the first main class listed on the patent). For a set of robustness results in this Online Appendix, we also rerun our analysis for a subsample that excludes software patents. Following Chung et al. (2015:Table 2), we identify software patents as those with USPC main classes 341, 345, 370, 380, 382, 700–707, 710, 711, 713–715, 717, 726, and 902. We list the descriptions of these main classes in Table G.5.

Our main sample comprises patents covering XXX distinct main classes. The five most frequent main classes are XXX.

Table G.5: USPC Main Classes for Software Patents

Class	Description
341	Coded data generation or conversion
345	Computer graphics processing and selective visual display systems
370	Multiplex communications
380	Cryptography
382	Image analysis
700	Data processing: generic control systems or specific applications
701	Data processing: vehicles, navigation, and relative location
702	Data processing: measuring, calibrating, or testing
703	Data processing: structural design, modeling, simulation, and emulation
704	Data processing: speech signal processing, linguistics, language translation, and audio compression/decompression
705	Data processing: financial, business practice, management, or cost/price determination
706	Data processing: artificial intelligence
707	Data processing: database and file management or data structures
710	Electrical computers and digital data processing systems: input/output
711	Electrical computers and digital processing systems: memory
713	Electrical computers and digital processing systems: support
714	Error detection/correction and fault detection/recovery
715	Data processing: presentation processing of document, operator interface processing, and screen saver display processing
717	Data processing: software development, installation, and management
726	Information security
902	Electronic funds transfer

Source: `https://www.uspto.gov/web/patents/classification/selectnumwithtitle.htm` and Chung et al. (2015:Table 2).

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

G.3 Construction of Additional Variables

We further collect and construct three sets of variables to proxy a patent's "patent scope and complexity," its "external impact," and its "internal value." For our heterogeneity results, we also collect and construct variables capturing the size of the patent applicant and the broader technology class of the patent.

G.3.1 Patent Scope and Complexity

We follow Lerner (1994) and Lanjouw and Schankerman (2004) and measure patent breadth and scope using the number of independent claims in a patent. Kuhn and Thompson (2017), however, argue that a simple count of (independent) claims may be a poor measure for patent scope.[2] They propose the length of the first patent claim as an alternative measure for patent scope, where shorter claims are broader. They use the first claim for their measure because under U.S. law the broadest claim should be listed first. We adopt their measure (length of the first claim in number of words) alongside the number of independent claims.

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

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

G.3.2 External Impact

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

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

$$1 - \sum_{k=1}^{n} \left(\frac{\text{backward citations}_{jk}}{\sum_{m=1}^{n} \text{backward citations}_{jm}} \right)^2 \tag{G.3}$$

where $s_{jk} = \frac{\text{backward citations}_{jk}}{\sum_{m=1}^{n} \text{backward citations}_{jm}}$ is the share of backward citations that patent j makes to patents in patent class $k = 1, \ldots, n$ over all backward citations made by patent j. A higher originality score means patent j draws on prior knowledge from a greater variety of fields. We construct this measure using the first listed USPC main class on a patent j. We have classification information for patents

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

granted in and after 1976. This means that for patents granted early in our sample period that cite patents granted before 1976, we have little information about the classes of their cited patents. Because of this truncation issue, the originality measure is therefore noisier and coarser for earlier patents than for patents granted later in our sample period.

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

$$1 - \sum_{k=1}^{n} \left(\frac{\text{forward citations}_{jk}}{\sum_{m=1}^{n} \text{forward citations}_{jm}} \right)^2 \tag{G.4}$$

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

G.3.3 Internal Value

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

We use information from the USPTO's Patent Maintenance Fee Events database at `https://bulkdata.uspto.gov/data/patent/maintenancefee` (January 28, 2019). The database contains all recorded events related to the payment of maintenance fees for patents granted from September 1, 1981 and forward. A patent is said to have been maintained if one of the codes listed in Table G.6 is recorded.

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

Table G.6: Codes for Maintenance Fee Events

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

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

through December 2016.

G.3.4 Applicant Size

For our variable of applicant size (or entity size), we combine information from the USPTO's Patent Maintenance Fee Events database and bibliographic information on patents from PatentsView. Applicant size takes three values. It is equal to 1 if the applicant is an individual, equal to 2 if the applicant is a small firm (i.e., small entity but not an individual), and equal to 3 if the applicant is a large firm (i.e., large entity but not an individual).

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

An applicant of a given patent j is an individual ($= 1$) if the first assignee listed on the patent is of type "individual" or if no assignee is listed on the patent. If the applicant is not an individual, then its size is equal to 2 if it is a small entity and equal to 3 if it is a large entity (as defined above). In our main estimation sample, XXX% of applicants are individuals, XXX are small firms, and XXX% are large firms.

G.3.5 Technology Class

We obtain NBER technology classifications from PatentsView. The NBER technology categories are constructed by Hall et al. (2001). Patents are assigned to six categories: Chemical (1), Computers & Communications (2), Drugs & Medical (3), Electrical & Electronic (4), Mechanical (5), and Others (6). We provide a list of the categories with their respective 36 sub-categories in Table G.7. Note that software patents (see above) predominantly fall into category Computers & Communications and sub-category Computer Hardware & Software.

Filling some gaps in the data, we assign USPC main class 532 to category 1 (Chemical) and sub-category 14 (Organic Compounds); and USPC main classes 901 (robots) and 902 (electronic funds transfers) to category 2 (Computers & Communications) and sub-category 22 (Computer Hardware & Software).

G.4 Process Patent Indicator

G.4.1 Summary of Indicator Construction

Ganglmair et al. (2019) employ text-analytical methods to identify the invention type of all independent claims in a given patent. We aggregate their claim-level data to obtain data at the patent level. In the sequel, we summarize their approach. Some of the material is also borrowed from Rosenberg (2012). An additional useful source of further background information is WIPO (2007).

The unit of analysis in Ganglmair et al. (2019) is an independent patent claim. A patent claim defines the scope of legal protection provided by a patent. It describes what the applicant claims to be its invention for which the patent grants exclusive rights. Each patent can hold multiple claims of different types. An *independent* claim

Table G.7: NBER Technology Categories and Sub-Categories

NBER Category	NBER Sub-Categories
Chemical (1)	Agriculture, Food, Textiles (11); Coating (12); Gas (13); Organic Compounds (14); Resins (15); Miscellaneous-chemical (19)
Computers & Communications (2)	Communications (21); Computer Hardware & Software (22); Computer Peripherals (23); Information Storage (24); Electronic Business Methods and Software (25)
Drugs & Medical (3)	Drugs (31); Surgery & Medical Instruments (32); Biotechnology (33); Miscellaneous-Drug&Medical (39)
Electrical & Electronic (4)	Electrical Devices (41); Electrical Lighting (42); Measuring & Testing (43); Nuclear & X-rays (44); Power Systems (45); Semiconductor Devices (46); Miscellaneous-Elec. (49)
Mechanical (5)	Materials Processing & Handling (51); Metal Working (52); Motors, Engines & Parts (53); Optics (54); Transportation (55); Miscellaneous-Mechanical (59)
Others (6)	Agriculture, Husbandry, Food (61); Amusement Devices (62); Apparel & Textile (63); Earth Working & Wells (64); Furniture, House Fixtures (65); Heating (66); Pipes & Joints (67); Receptacles (68); Miscellaneous-Others (69)

Source: Hall et al. (2001) and PatentsView. Appendix 1 in Hall et al. (2001) also lists the respective USPC main classes (version 1999) for each sub-category.

stands on its own whereas a *dependent* claim is in reference to an independent claim, further limiting its scope.

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

The approach in Ganglmair et al. (2019) uses information from both the preamble and the body. Both parts of the claim are classified as describing a process or a product. For the preamble, this classification is conducted via a simple keyword search (e.g., “process” or “method” for process-claim preambles; “apparatus” or “device” for product-claim preambles). For the body, the authors take a syntax-based approach, analyzing the linguistic structure of each line (or “bullet point”) in the body. The steps of a process are listed using the gerund form of a verb, whereas the items of a product (an apparatus, a device) are listed as components. The authors’ algorithm accounts for these drafting conventions when classifiying a body as process-claim body or product-claim body. In the end, combining the classifications of the preamble and the body, a classification for the entire claim is obtained:

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

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

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

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

The authors’ algorithm deals at great length with a number of issues: badly formatted claims, claims not following the usual drafting conventions, and two-part

Figure G.8: Share of Missing Observations (All Three Patent-Level Indicators)

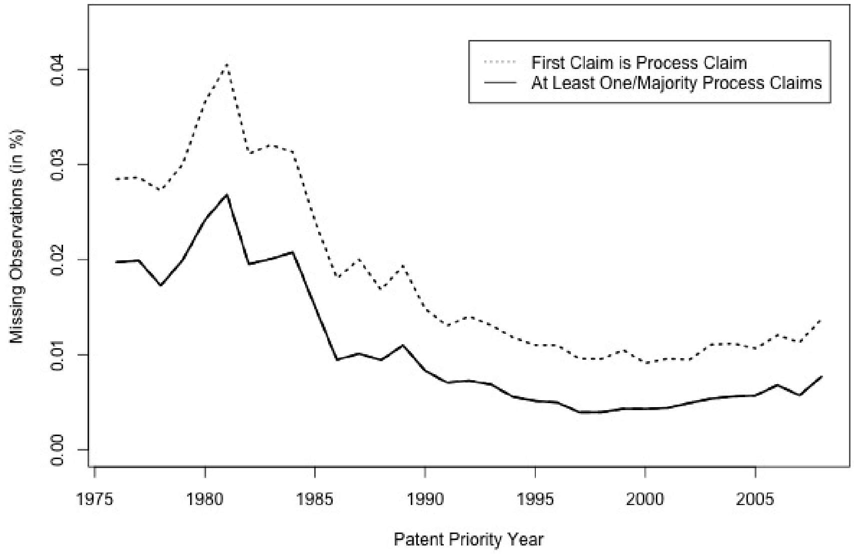

claims (also called improvement claims or Jepson claims). In Figure G.8, we plot the fraction of missing observations for each of our patent-level indicator. For both our main indicator and the process patent indicator with a majority of process claims, at least one patent claim must be classified - the graphs in the figure are therefore the same. The requirement for the indicator of the first process claim is stricter, and the number of missing observations is higher throughout. Notice, however, that the reliability of the approach increases over time as the percentage of missing observations (over all patents in our main sample) drops below 1% around 1985 (with higher numbers for patents with earlier priority dates).

G.4.2 Descriptive Figures

In Figure G.9, we plot the share of process patents by priority year. We show graphs for each of our three process patent indicators. The solid line depicts the share of process patents for our main indicator (Figure 2: at least one patent claim is a process claim, 'Any'). The dotted graph depicts the share of patents with the first patent claim a process claim ('First'); the dashed graph depicts the share of patents with a majority of process claims ('Most'). As we we have discussed in the main text, our main indicator is the most aggressive in terms of identifying patents as process patents. The overall time trends, however, are very similar. We also plot the average

Figure G.9: Share of Process Patents (Multiple Indicators)

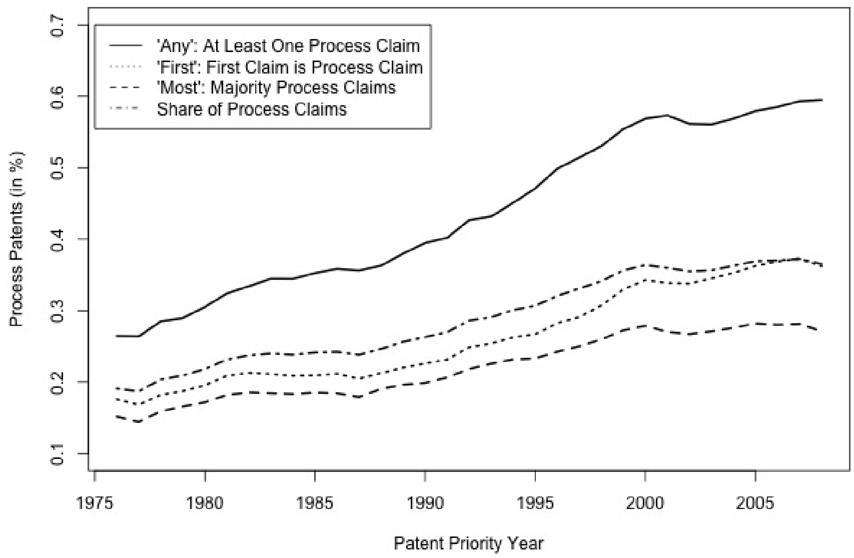

share of process claims in a patent (dash-dotted line). The graph follows similar trends.

In Figure G.10, we depict the share of process patents by applicant size (panel (a)) and NBER category (panel (b)) – the two dimensions we use for our analysis of heterogeneous treatment effects in the main text. The share of process patents is higher in larger firms than in smaller firms, and lowest for individuals. In panel (a) of Figure G.11 we can further observe this pattern in all NBER categories except "Drugs and Medicals" (Category 3) in which small firms exhibit the highest numbers for process patents, followed by large firms and individuals.

In panel (b) of Figure G.10, we see that the NBER Category "Computers and Communication" (Category 2) has the highest share of process patents. Within this category, "Computer Hardware & Software" (Sub-Category 22) and "Electronic Business Methods and Software" (Sub-Category 25) stand out. This implies that even without business methods (or: business method patents), category 2 is the a leading category for process patents. On the other end of the spectrum, the catch-all category "Others" (Category 6) exhibits the lowest share. Within this latter category, "Earth Working & Wells" (Sub-Category 64) has the highest share (with more than 50%), whereas "Furniture, House Fixtures" (Sub-Category 65) comes with the lowest share of process patents.

Figure G.10: Share of Process Patents

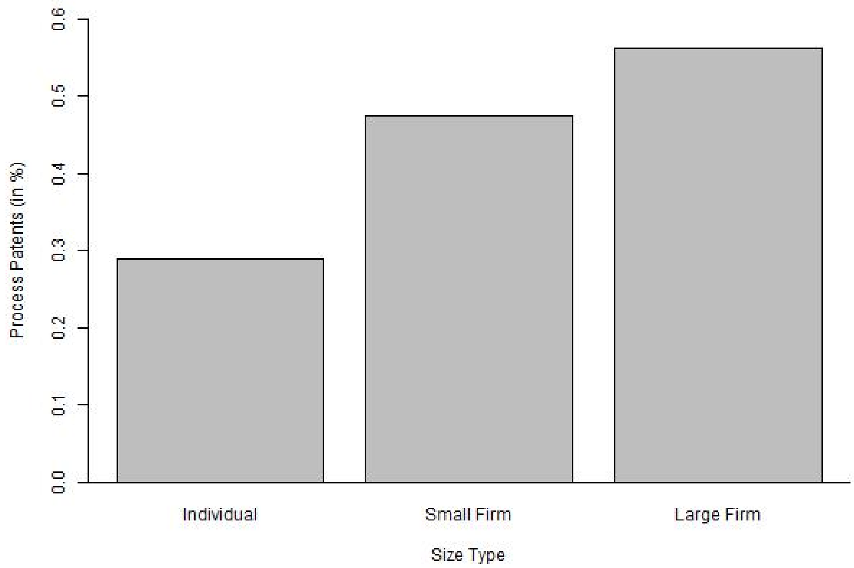

(a) Share of Process Patents by Applicant Size

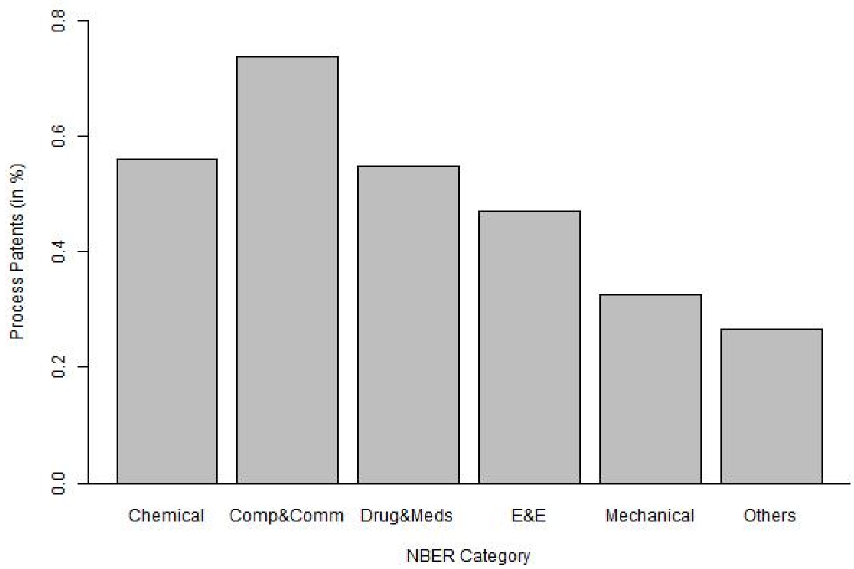

(b) Share of Process Patents by NBER Category

Figure G.11: Share of Process Patents (by NBER Category and Time Period)

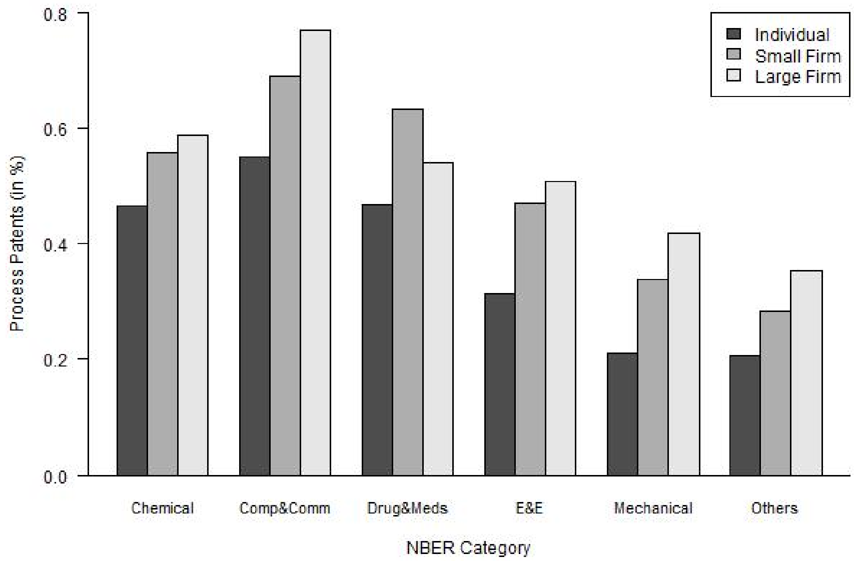

(a) By NBER Category and Applicant Size

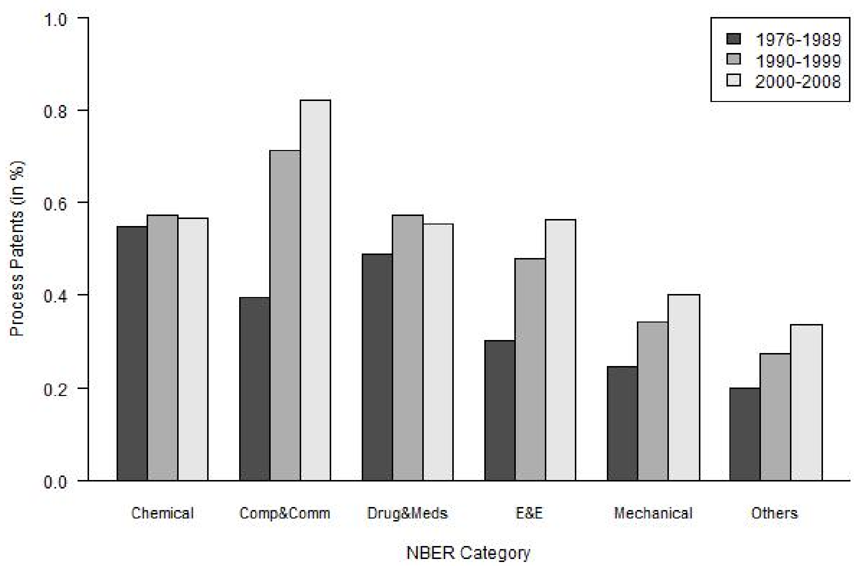

(b) By NBER Category and Time Period

Figure G.12: Share of Process Patents (by NBER Sub-Category)

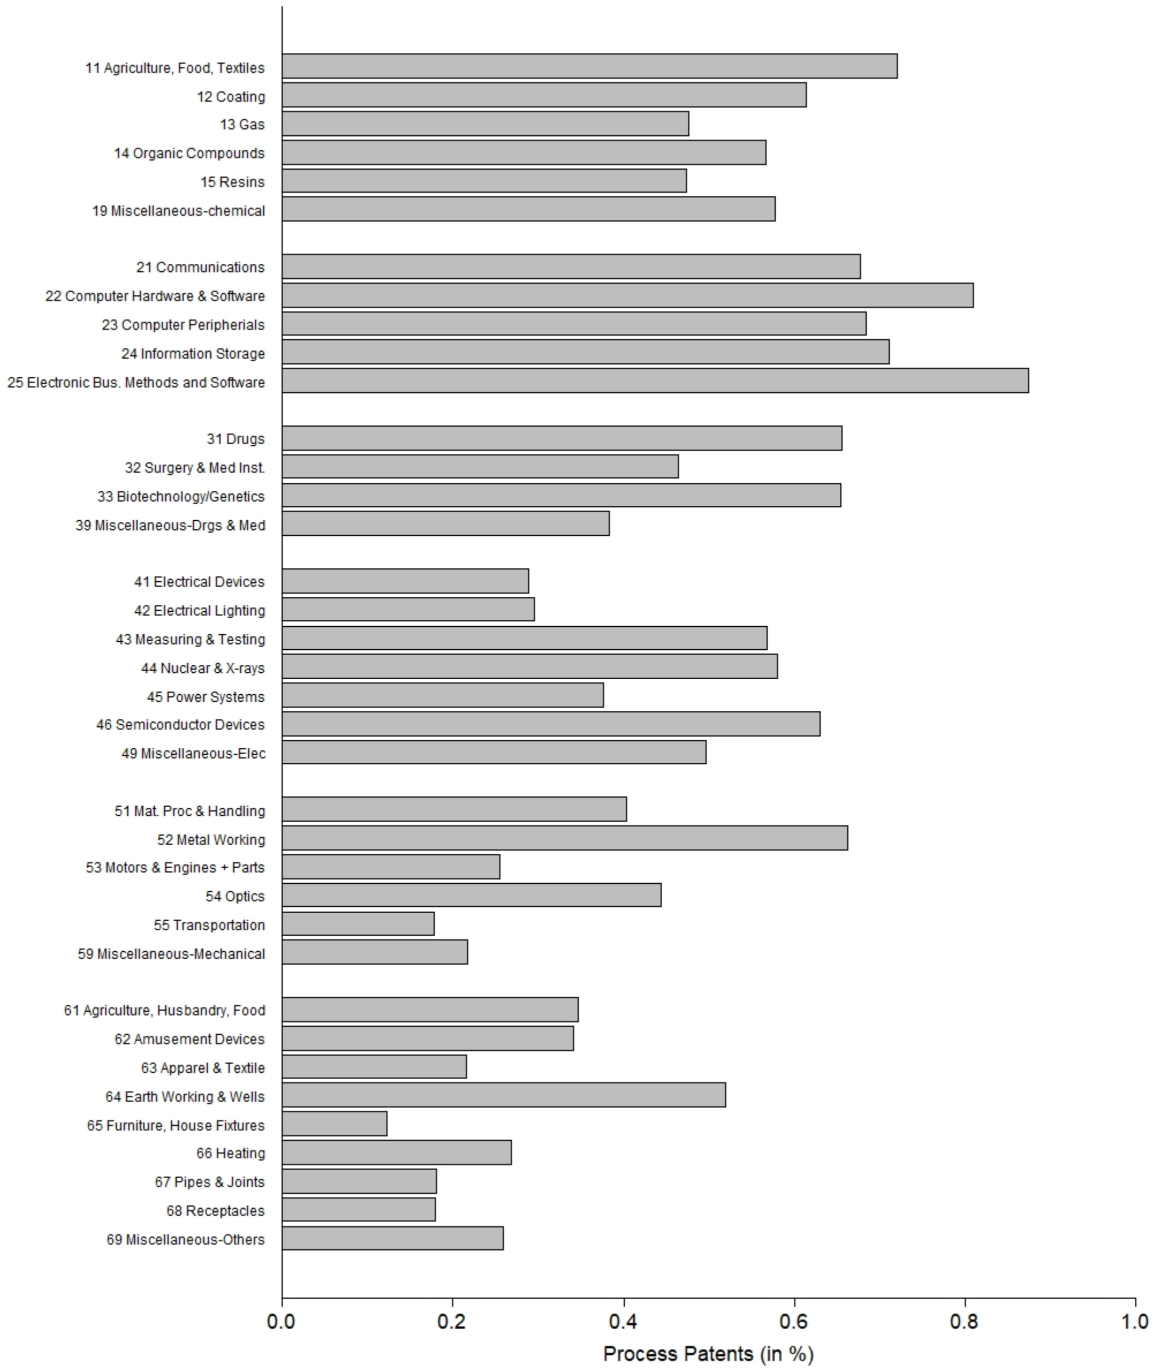

Last, in panel (b) of Figure G.11 we capture time trends in the share of process patents for different NBER categories. We see strong positive time trends for "Computers and Communication" (Category 2) and weaker trends for "Electrical and Electronic" (Category 4), "Mechanical" (Category 5), and the catch-all category "Others" (Category 6). We see little or no time trends for "Chemical" (Category 1) or "Drugs and Medical" (Category 3).

G.5 Data for AIPA Analysis (Graham and Hegde, 2015)

We use the supplementary data provided by the authors at `http://science.sciencemag.org/content/347/6219/236`.